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

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

March 2025

Abstract

We analyze how a negative shock to the profitability of oil-extracting firms may lead to a shift from dirty to clean R&D along the supply chain. First, we develop a theoretical framework, showing that adjustment costs in R&D give firms in the fossil energy supply chain an additional incentive to shift R&D activity towards clean innovation as a consequence of a negative shock. Next, we leverage the 2014 oil price drop to empirically investigate the impact of reduced profitability in the fossil energy supply chain on clean R&D. We propose a novel method to identify firms' exposure to the shock. In line with the predictions from the model, we find that more exposed firms increased their clean R&D more than other firms. Our findings imply that carbon pricing will induce clean innovation not only by increasing demand for clean technologies, but also by lowering profitability in 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, Eugenie Dugoua, Rikard Forslid, Bård Harstad, Linnea Lorentzen, Christoph Meinerding, Sjak Smulders, and Andre Trindade, along with seminar audiences at University of Cambridge, Geneva Graduate Institute, University of Surrey, 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, and CCA Workshop on International Economics for valuable comments and discussion. This project has received funding from the Norwegian Ministry of Trade, Industry and Fisheries (INNOPROD). Data from Statistics Norway have been crucial for this project.

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

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

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

1 Introduction

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

While there is an extensive literature investigating how clean R&D responds to changes in energy prices, the 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 energy prices may 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 profitability in the fossil energy supply chain triggers a reallocation of R&D resources towards clean innovation.[2]

We start by developing a theoretical framework building on the literature on directed technical change, in particular Acemoglu et al. (2012). In this literature, 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 activity, regardless of what type of production and innovation they currently undertake. We introduce within-firm dynamics into this framework by adding two key features, which give rise to additional responses to price changes.

The first feature we introduce is adjustment costs for R&D activity, as proposed by Bloom et al. (2013b). Adjustment costs in R&D can arise for many reasons – for example, in connection with research team dynamics or from search and hiring costs. The presence of adjustment costs implies that in the aftermath of a negative shock to expected returns

[1]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 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 2.2.

[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 2.4.

to dirty innovation, firms more exposed to the shock will face a stronger incentive to divert their R&D activities towards clean innovation. The second feature we introduce is spillovers from the more mature, i.e., dirty, R&D activity towards clean R&D. This extension works in the opposite direction: After a negative shock to the returns to dirty innovation, the presence of spillovers will incentivize firms exposed to the shock to scale down both dirty and clean R&D.

The theoretical model 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 spillovers pull in opposite directions, the net effect is theoretically ambiguous. We turn to the data to determine which of these forces dominates by studying the impact of a negative oil price shock on clean R&D activity in input suppliers.

In our empirical analysis, we use detailed micro data from Norway. 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.[3] We exploit the large negative shock to the contemporaneous and future price of oil in the latter half of 2014, when the crude oil price dropped by approximately 60 percent. This quasi-natural experiment allows us to study the dynamics of resource reallocation in the Norwegian manufacturing industries. We interpret the empirical results through the lens of the theoretical model.

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

The purpose of our empirical analysis is to estimate the causal effect of the oil price shock on input suppliers' clean R&D investments. Specifically, we investigate whether the manufacturing firms more exposed to the shock reacted differently from other firms. To this end, 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. We compute this measure based on the share of products in the firm's export basket that are used as inputs by oil-extracting firms. This firm-level exposure measure allows us to compare

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

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

firms in the same industry that otherwise have similar characteristics at baseline but are differentially exposed to the shock because of their 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 firms more exposed to the oil price shock increased their clean R&D efforts more than less exposed firms. This is true whether we look at the firm's likelihood of investing in clean R&D, the share of clean R&D in its R&D investments, or the value of its clean R&D expenditures. In light of our theoretical model, our empirical findings suggest that adjustment costs encouraged exposed firms to shift their R&D activity towards clean innovation after the shock.

Our theory suggests that firms' R&D decisions are driven by expected future profits, which are hard to measure in the data. However, we investigate whether more exposed firms experienced decreases in sales and profits after the shock. This is indeed the case, as both sales per employee and profits decreased. We also explore whether the exposed firms simply intensified their R&D investments across the board, finding that this was not the case. We find no effect on total R&D investments, number of R&D employees or share of R&D employees in total employment.

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. In light of the theoretical framework, we link these findings to within-firm adjustment costs in R&D activity. 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, a method that we believe can be applied to many other settings.

Our findings may at first glance appear counterintuitive. It is well established in the literature that if the fossil energy price falls, demand for dirty goods increases, and firms invest more in dirty innovation (Aghion et al., 2016). At the same time, Acemoglu et al. (2012) show how a decline in fossil energy prices can shift R&D investments towards clean innovation as a consequence of the change in relative profitability. We show that, in addition, a fall in profitability for input suppliers in the fossil energy supply chain leads them to shift their R&D activity *more* towards clean innovation than it does other firms, in line with the existence of adjustment costs in R&D activity.

Our findings indicate that the substantial R&D capabilities in the fossil energy supply chain can be leveraged in the clean transition. Moreover, the presence of adjustment costs for R&D gives rise to reallocation of resources *within a given firm* in response to changes in profitability. More broadly, 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 findings suggest 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.

Although the focus of our empirical analysis is input suppliers in the Norwegian fossil energy supply chain, we conjecture that our results extend to other settings under certain conditions. First, we expect the same mechanism to be salient for the oil-extracting firms themselves: The presence of adjustment costs in R&D activity should lead oil-extracting firms to redirect innovation activity towards clean innovation when they are faced with lower expected returns to dirty innovation. Our focus on input suppliers is primarily motivated by identification issues, as it allows us to compare otherwise-similar firms that are differentially exposed to the shock because of differences in their product mix. This enables us to control for other potential confounders occurring at the same time. Second, while it is clear that our main finding of increased clean innovation as a consequence of a relative drop in the oil price would not extend to firms globally in a naive way, our results may well also apply to firms in the fossil energy supply chain in other countries as long as the oil price drop reduces their expected future profitability.

The paper relates to several strands of literature. The theoretical framework builds directly on the well-established literature on directed technical change and climate, in particular Acemoglu et al. (2012). Dechezleprêtre and Hémous (2022) provide a recent review.[5] We add to this literature by including within-firm reallocation of R&D investments in a model allowing innovation to take place in both clean and dirty technologies at the same time. This feature of parallel innovation in both sectors 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

[5]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.

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.

On the empirical side, the paper contributes to the literature investigating the relationship between relative energy prices and clean innovation. The findings of Newell et al. (1999) indicate that higher energy prices will induce more rapid development of energy-saving technologies. Similarly, Popp (2002) finds that higher energy prices are associated with a significant increase in energy-saving innovations at the industry level. Aghion et al. (2016) show that firms in the auto industry tend to innovate more in clean and less in dirty technologies when they face higher tax-inclusive fuel prices. A similar pattern, with a positive relationship between energy prices and energy-saving innovation, is 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 related to clean innovation activity. Dugoua and Gerarden (2024), Dugoua and Dumas (2021) and Dechezleprêtre and Kruse (2022) provide recent contributions to the literature on energy prices, environmental policies and clean innovation.[6] The main focus of the empirical literature has been on the demand-side effect of changes in fossil energy prices on clean innovation. Moreover, there is little empirical evidence on the role of within-firm reallocation in the shift from fossil to clean technologies. We add to this literature by focusing on the supply-side effect of lower fossil energy prices, and specifically, we add to the understanding of how lower expected future profitability can induce a shift in R&D activity within firms.

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

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

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

[7]See Golosov et al. (2014) for a key contribution to the literature on optimal carbon pricing and Timilsina (2022) and Köppl and Schratzenstaller (2023) for reviews of this literature. 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).

on investments. We also propose a novel method for identifying firm-level exposure to shocks, relying on firm–product trade data that are increasingly available to researchers.

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

2 Data and Stylized Facts

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

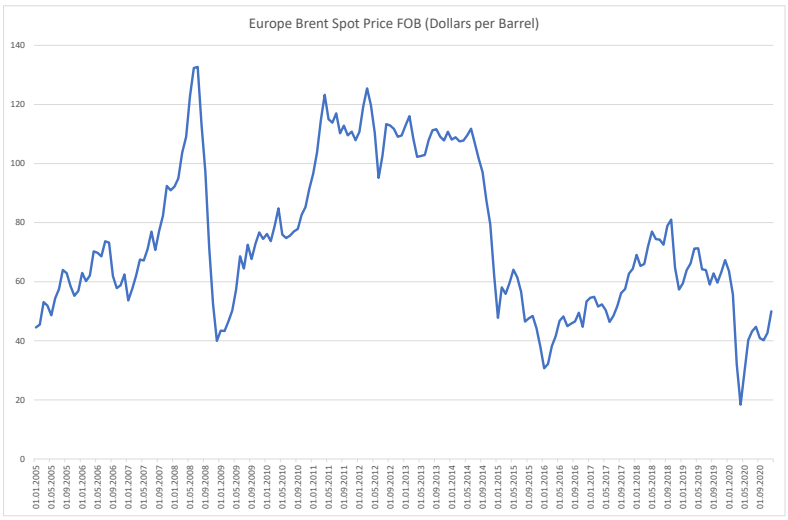
2.1 The oil shock

Our analysis relies on a quasi-natural experiment created by 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 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.[9]

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

[9]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.

Figure 1: Oil price: Brent blend



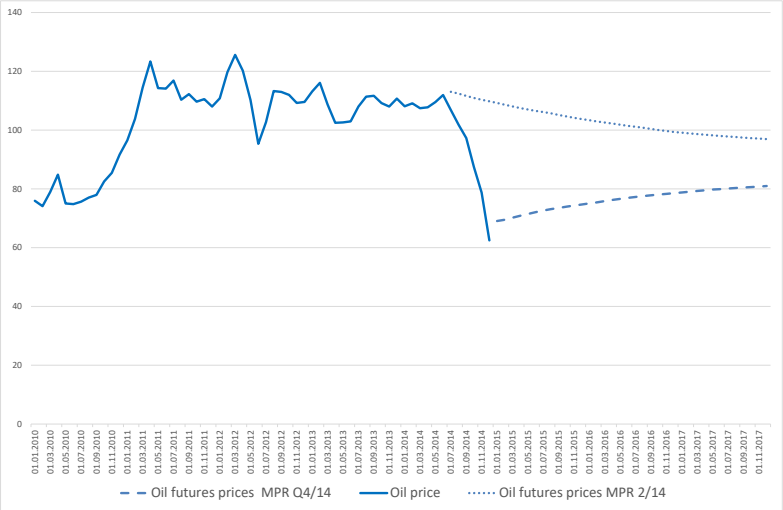
Note: The figure plots the European Brent spot price for oil per quarter for the period 2005–2020. Source: US Energy Information, https://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=PET&s=RBRTE&f=D.

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. Firms most likely did not expect this dramatic decline. 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 potentially 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.

This oil price shock has appealing methodological properties in the setting we study. First, the oil price is set in the world market and is completely exogenous from the perspective of firms in a small open economy.[10] Additionally, the technology, intermediates and capital goods used for offshore oil extraction differs substantially from the technology used for onshore 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, as the Norwegian fossil energy supply chain is based exclusively on offshore activity.

[10]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/.

Figure 2: Crude oil price and oil future prices



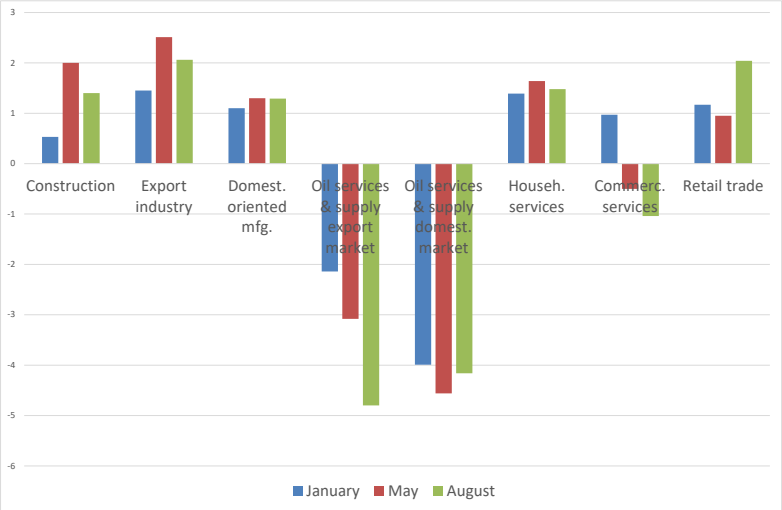
Note: The figure plots the crude oil price until end 2014 and oil future prices from 2014 to 2017. Source: The Norwegian Central Bank Monetary Policy Reports 2/2014 and 4/2014.

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 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.[11]

[11]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.

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.

2.2 Data

Our empirical analysis is based on five datasets. 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, allowing us to calculate the firm's age. The register also holds data on firms' number of employees.

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 data from annual accounting reports that, according to Norwegian law, must be filed with the public Register of Company Accounts.

The third dataset is the R&D survey.[12] The survey provides information about the value of total R&D expenditures and the number – and wage bill – of R&D personnel.

[12] 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.

Importantly, the survey also gives information about the share of clean R&D. 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.[13] We link all datasets with a unique firm identifier.

2.3 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).[14] This restriction gives us an unbalanced panel of approximately 1,300 firms per year. 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.

2.4 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.

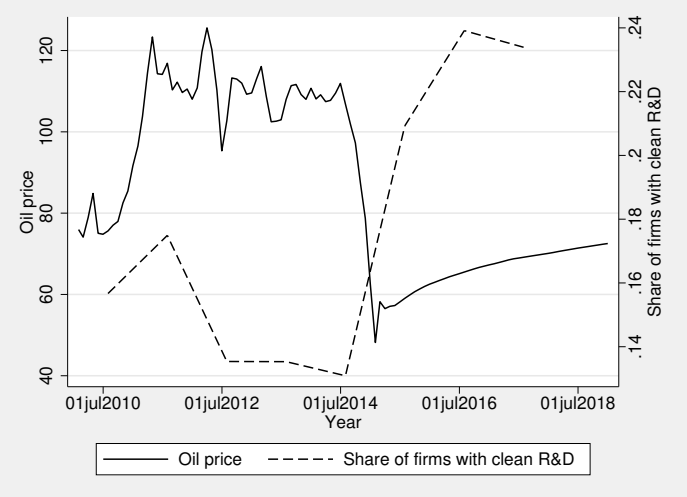
Fact 1: After the oil price shock, the share of firms investing in clean R&D increased in both absolute and relative terms. In 2013, the share of firms with positive R&D in our sample was 40 percent, and the share of firms investing in clean R&D was only 5 percent. By 2017, these numbers had increased 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

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

[14]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.

shows the increase in the share of firms with clean R&D expenditure relative to all firms engaged in R&D.

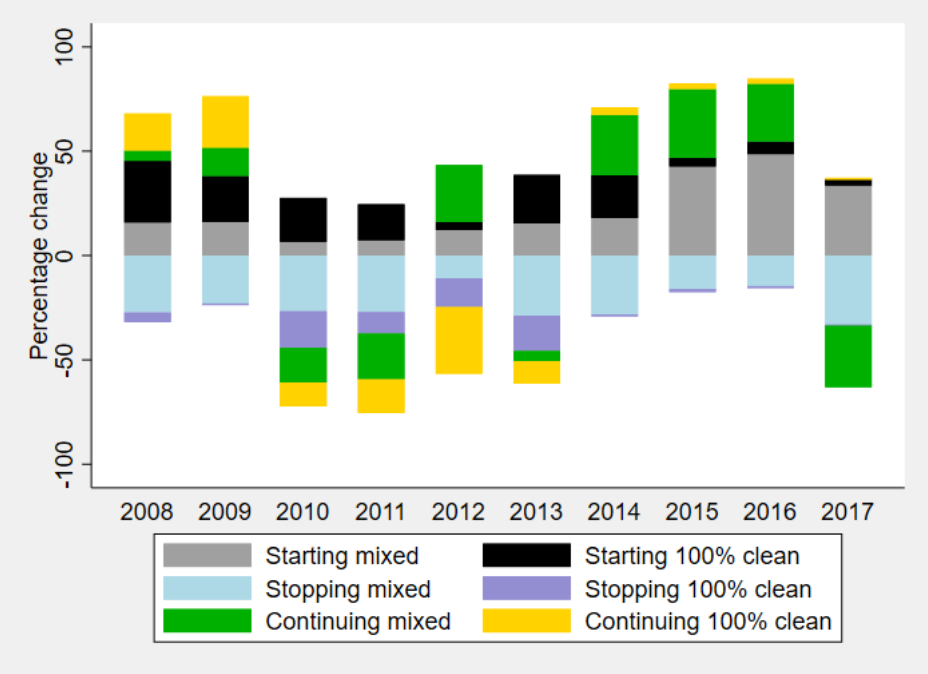
Figure 4: The oil price and the share of R&D firms with clean R&D



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.[15]

Fact 3: There is persistence in R&D activities 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 activities 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 activity is subject to adjustment costs.

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

3 A Theoretical Model of Directed Technical Change

In this section, we develop a stylized model of directed technical change to guide our empirical analysis of the impact of a permanent drop in the oil price on investments in clean R&D in input suppliers in the fossil energy supply chain. We build on the two-sector model of Acemoglu et al. (2012), where the dirty sector represents fossil fuel production. Based on this framework, we develop a model that allows us to investigate how a negative price shock to a dirty final good may propagate through the supply chain and induce a shift in R&D investment from dirty to clean R&D in input suppliers. The model allows us to illustrate how input suppliers may react to an oil price drop differently from other firms.

Informed by the three stylized facts presented above, we extend the model of Acemoglu et al. (2012) by allowing input suppliers to differ in their innovation productivity with respect to dirty and clean R&D, which implies that firms will make heterogeneous choices regarding R&D investments. Moreover, we introduce two types of within-firm dynamics in the model. First, we allow for within-firm spillovers from the more mature (dirty) R&D activity to the less mature (clean) activity. Here, we aim to capture the possibility that dirty R&D activity primarily focused towards fossil energy production may also be valuable for developing technologies for renewable energy production or energy saving.[16]

[16]Note that our model can be extended in several directions to allow for more general spillovers. For example, it is possible to allow for spillovers in both directions, i.e., from clean to fossil R&D investments. Such an extension would not change our conclusions regarding the qualitative effects on innovation of a drop in the dirty energy price. Next, one could include between-firm spillovers or more general spillovers depending on aggregate knowledge stocks. 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,

Second, in line with Bloom et al. (2013b), we include adjustment costs in R&D activity. We think of these adjustment costs as capturing not only the direct costs of firing scientists but also other economic and behavioral responses that might arise within a team of highly specialized workers when downscaling is considered. For instance, managers might be reluctant to let go of hard-to-replace scientists with specialized skills. They may also prefer to avoid disrupting the dynamics within a team and instead to keep a well-composed team intact. This team cohesion effect occurs when a team's collective performance is greater than the sum of its individual members' performances and removing one or more of the scientists would lead to productivity losses for the remaining team members. Managers might also worry about a layoff contagion effect, whereby laying off one scientist makes colleagues worry about their own job security.[17]Managers might also care about their own careers and be reluctant to let their own teams shrink, even to the detriment of the firm's best interest. We model adjustment costs by including a convex cost of rescaling R&D activity within the firm, which reflects that many of the sources of such costs have characteristics likely to be associated with increasing marginal costs in rescaling within a given period.[18]

3.1 Model

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

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

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

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

also likely to be present within firms (see, e.g., Nix (2020)). Allowing positive spillovers from the most mature type of R&D activity to the other captures the possibility that dirty R&D activity in the energy sector can potentially also be valuable for developing technologies necessary for renewable energy production or energy saving, which is especially relevant for our understanding of how within-firm dynamics may lead input suppliers to react to the shock that we study differently from other firms. To keep the framework as simple and transparent as possible, we therefore include spillovers only from dirty to clean innovation within the firm.

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

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

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

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, we follow Acemoglu et al. (2012) and let $\phi = \alpha^2$.

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

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

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

To innovate, the 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 firm hiring a scientist chooses whether to put the scientist to work on dirty or clean R&D. A scientist hired for type-j R&D activity is randomly assigned to an input of that type. The input suppliers act competitively in the market for scientists. Let s_{kjt} be the number of scientists hired by firm k for type-j R&D activity in period t, and let $s_{kt} = s_{kdt} + s_{kct}$ be the total number of scientists hired by firm k in period t.

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

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

0, $g^{d\prime}(\cdot) > 0$, $g^{d\prime\prime} < 0$, $g^{d\prime\prime\prime} = 0$ and $g^d(1) \leq 1$.[19]

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

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

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

To reflect adjustment costs such as those discussed above, we include a convex cost of rescaling the R&D activity within the firm. For simplicity, the adjustment cost is separable from the innovation productivity and given by:

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

for firm k, with $c > 0$. Note that key to the results presented below is that the rescaling cost depends on the overall level of R&D activity in the firm, independently of whether the activity is related to clean or dirty inputs, as we aim to capture a relation between clean and dirty R&D activity.

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

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

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

with $\int_0^1 s_{kt} dk = 1$ in equilibrium. Allocation of scientists will therefore drive the technological state of the economy in the long run. In the following, we make one further simplifying assumption, namely, that all input producers are myopic, in the sense that they do not take future adjustment costs into account when they make their R&D investment decision.[20] The value of potential innovation as a result of R&D investment arises in the current period and is, of course, included in the firms' objectives.

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

[20] Fully forward-looking firms will take into account not only the current adjustment costs but also those expected to be incurred in the future. For a firm adjusting its R&D up or down, the expected future adjustment costs provide an incentive to rescale somewhat faster than what would otherwise be optimal for a myopic firm. However, as long as the future adjustment costs are discounted, the key

3.2 Market equilibrium

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

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

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

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

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

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

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

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

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

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

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

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

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

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

effect of including the adjustment costs in the firms' decisions is captured in the partly myopic problem presented in the following. Therefore, we keep the framework as simple as possible by assuming that firms are myopic with respect to future adjustment costs.

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

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

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

3.3 Comparative statics

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

We show in Appendix Section B that the direct effect of a fall in the price of the dirty final good on demand for scientists is negative for the input suppliers in group II and III, which supply the producers of the dirty final good with inputs: Their revenue falls, and returns to dirty R&D innovation fall. As a consequence, these firms will reduce their dirty R&D activity and thus their demand for scientists, which in turn leads to a decline in wages (see Equations (29), (32) and (37)). The decline in wages will induce all firms active in R&D, including the firms in group I, to expand their R&D activity (see Equations (33)–(36)). This result is the first-order effect of the change in relative prices, parallel to that described by Acemoglu et al. (2012).

Firms in group III are active in both dirty and clean R&D. The total impact of a price shock for dirty goods on their clean R&D activities depends on the wage effect, which is also experienced by firms in group I, and on within-firm allocation of R&D investments related to technological spillovers and adjustment costs: The presence of technological spillovers implies that the reduction in dirty R&D activity leads to lower marginal productivity of scientists working on clean innovation. This in turn leads to

reduced clean R&D activity. Adjustment costs, on the other hand, imply that these firms, relative to firms in group I, increase the resources they allocate towards clean innovation because the net marginal cost of their clean R&D activity decreases as they downscale their dirty R&D. In other words, the adjustment costs make firms shift some of their scientists from dirty to clean R&D instead of firing them. The effect of a change in the price of the dirty final good on the demand for scientists for clean R&D activity in group III relative to that among firms in group I depends on which effects dominate.

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

Proposition 1. *When there are no within-firm technological spillovers and no adjustment costs, a drop in the price of the dirty final good, p_{dt}, will induce the same change in clean R&D in all firms engaged in clean R&D, independently of whether the firms are also active in dirty R&D:*

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

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

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

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

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

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

Proof. See Appendix Section B, specifically Equation (41).

In the next section, we take these propositions to the data. Note that in the theoretical framework, we deliberately focus on the intensive margin to avoid corner solutions. We do this to highlight the key mechanisms that we propose, which would also hold 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 but also include firms conducting only dirty R&D before the shock. □

4 Empirical Strategy

The point of departure for our empirical analysis is the theoretical framework developed above. According to Proposition 2, the relative impact of an oil price shock on exposed firms (group *III*) relative to non-exposed firms (group *I*) is ambiguous.[21] 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.

[21] 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.

Therefore, we instead develop a novel firm-specific exposure measure leveraging our 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).[22] We denote these *oil-related* products ($j \in o$).[23]

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

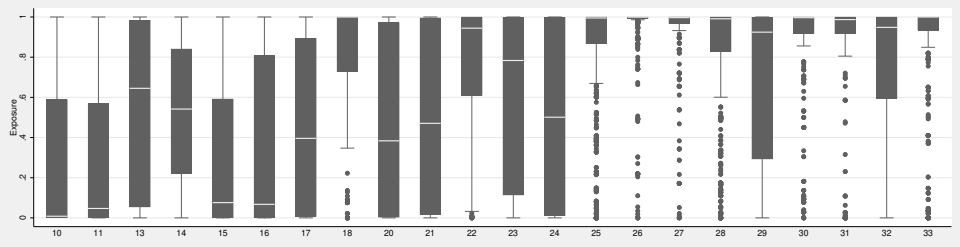
[22] We use the whole pre-shock period to account for the fact that firms' sourcing of capital goods will be subject to lumpiness over time.

[23] 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.

points outside this range are plotted individually as outliers, indicating data points that differ significantly from those in the rest of the distribution.

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.[24] 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:

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

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

where y_{kt} is an outcome variable related to the firm's R&D investments or performance. x_{ok} captures the firm's exposure to the shock based on its export portfolio (as described in Section 4.1).[25] 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. 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 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.[26]

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 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 public funding for R&D and innovation.

As all of these control variables themselves are possibly affected by the oil price shock,

[25]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.

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

we compute them based on the first year the firm is observed in the sample and interact them with year dummies. Standard errors are clustered at the firm level.

5 Empirical Results on Clean Innovation

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 activities. 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 (17) 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.[27] 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.

[27] 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 activities? 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 likelier 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

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:

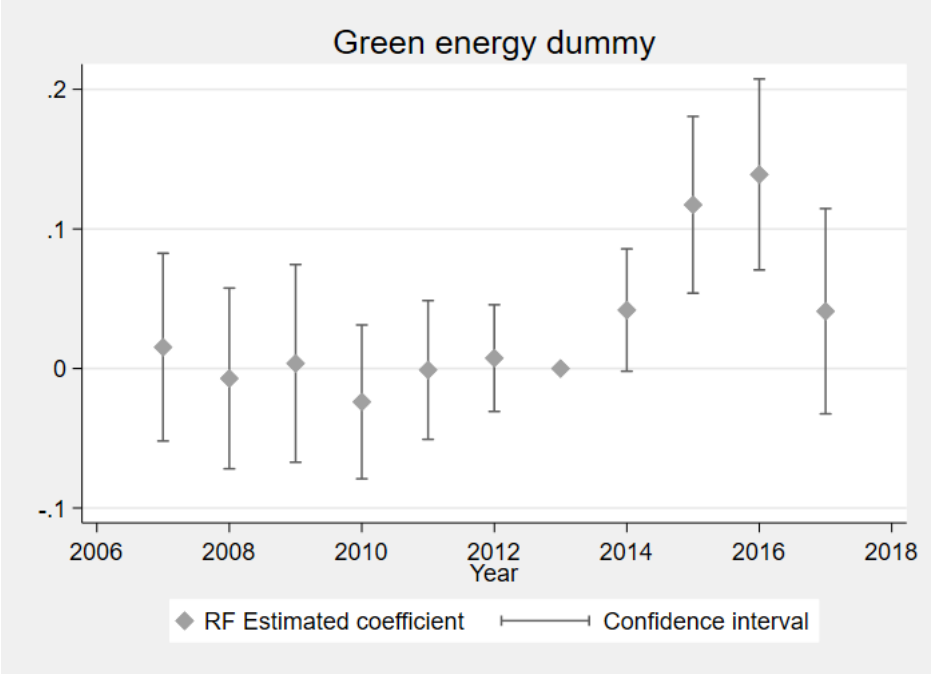
$$y_{kt} = \alpha_k + \beta x_{ok} \times \delta_t + \gamma Z_{kt} + \delta_{st} + \varepsilon_{kt} \tag{18}$$

For ease of exposition, we show the graph for a binary treatment variable here.[28] 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

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

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 (18).

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 are very similar to our baseline results, suggesting that our coefficient of interest can be interpreted as the average level treatment effect.[29]

[29]The results using a binary dummy variable are available upon request.

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. $*p < 0.1, **p < 0.05, ***p < 0.01$.

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 zero values are an issue. PPML allows us to deal with both of these concerns.[30] 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.

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 (17) with these two new variables as dependent variables. Table 5 reports the results. The estimated coefficients are not significant.

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

[30]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.

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 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.

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 (17) with these two new variables as dependent variables. The results from these regressions are presented in Table 6. We find that sales per employee 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 activities? 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 do so by developing a theoretical framework that highlights the role of adjustment costs in R&D activity. 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 carbon pricing will induce clean innovation not only by increasing demand for clean technologies and shifting relative prices, but also by shifts in R&D induced by lower profitability in the fossil energy supply chain. We leave the task of quantifying the aggregate impact of these shifts to future research. Furthermore, our

results indicate that the substantial R&D capabilities in the fossil energy supply chain can be leveraged in the clean transition.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hémous, "The Environment and Directed Technical Change," *American Economic Review*, 2012, *102* (1), 131–66. 1, 5, 3, 3.1, 3.1, 3.3

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

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

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

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

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

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. 9

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

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393. 16

_ , 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, 2.4

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

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," Technical Report 2024. 5.2

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

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

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

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

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

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

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

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

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

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

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

Hamermesh, Daniel S., "Labor Adjustment and the Source of Adjustment Costs," *The Economic Journal*, 1995, *105*, 620–634. 18

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

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

Kilian, Lutz, "The impact of the shale oil revolution on US oil and gasoline prices," *Review of Environmental Economics and Policy*, 2016, *10* (2). 9

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

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

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

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

Marx, Matt and Bram Timmermans, "Hiring Molecules, Not Atoms: Comobility and Wages," *Organization Science*, 2017, *28* (6), 1115–1133. 17

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

Nix, Emily, "Learning spillovers in the firm," Technical Report, IFAU Working Paper No. 2020:14 2020. 16

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

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

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

Prest, Brian C, "Explanations for the 2014 oil price decline: Supply or demand?," *Energy Economics*, 2018, *74*, 63–75. 9

Prestmo, Joakim Blix, Birger Strøm, and Hilde Karoline Midsem, "Ringvirkninger av Petroleumsnæringen i Norsk Økonomi," Technical Report 8 2015. 3

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

Sun, Liyang and Jesse M. Shapiro, "A Linear Panel Model with Heterogeneous Coefficients and Variation in Exposure," *Journal of Economic Perspectives*, 2022, *36* (4), 193–204. 5.2

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

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 Organisation 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 activity 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 activities.

B Calculations for Analytical Model

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

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

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

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

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

For firms in group III, the two FOCs are given by:

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

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

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

The FOCs in Equations (19)–(22) give the demand for scientists from all three groups of firms.

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

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

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

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

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

with

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

The cross-derivative of the objective function of a firm in group III is positive if the spillover effect dominates the effect of the adjustment cost, while it is negative if the adjustment cost dominates. In either case, we assume that Equation (25) holds, i.e., that the effect of dirty-scientist hiring on the marginal value and cost of hiring scientists for clean innovation, independent of its direction, is not large enough to prevent an interior solution to the problem of any firm in group III.

First-order effects of the oil price drop.

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

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

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

Next, it follows directly from differentiating the FOCs and using the inequalities in Equations (23)–(26) that:

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

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

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

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

Effects of wage changes.

Similarly, it follows, again from differentiating the FOCs and using Equations (23)–(26), that:

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

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

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

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

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

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

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

Proof of Propositions 1 and 2.

Define

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

with

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

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

Insert from Equations (39), (40) and (37) and reorganize to obtain:

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

with

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

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

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

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

If the adjustment costs dominates, the full expression in the outer brackets will be negative, and X_2 will take the opposite value of $\partial H^{III}_{kct}/\partial s_{kdt}$. The same will be the case if the spillover effect dominates, as long as $\partial H^{III}_{kct}/\partial s_{kdt}$ is not too large. In the main body of the paper, we assume this is the case. However, it is worth noting that for the special case of very large spillover effects, X_2 may take the same sign as $\partial H^{III}_{kct}/\partial s_{kdt}$. As a consequence, if λ^{II} is sufficiently large, the same can become true for the full expression for Δ in Equation (38). The firms in group II will react to the price drop by decreasing their demand for scientists, which will push the wage down. If the share of firms belonging to this group is very large, the firms in group III may face a wage drop that is large relative to the drop in prices.

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

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

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

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

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.

First, we categorize all products imported to Norway as oil related or not oil related. In this first step, we rely on data on imports to oil-extracting firms in Norway. 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.

Building on Armington (1969), we assume that the production function of the oil-extracting firms encompasses different preferences for foreign and domestically produced

inputs:

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

with subscript D denoting domestically produced inputs and subscript M foreign inputs, while β_{di} is the value of the imported good in production relative to the domestically produced good. The size of β_{di} determines the home bias for input di. For higher values of β_{di}, input demand would be directed more towards domestically produced goods. If the home bias differs across industries and inputs, the share of imports of good x_{di} that is directed towards the oil-extracting firms will not be informative about the relative importance of good x_{di} 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

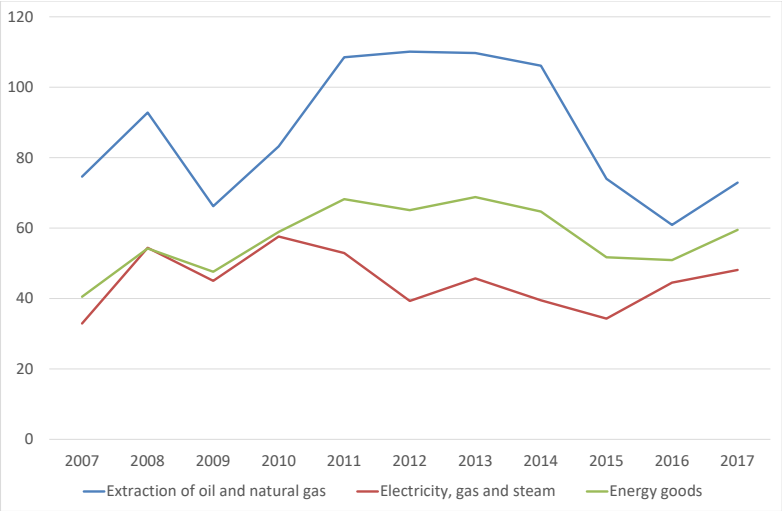
Table 8: Clean R&D: Dummy, share and value with PPML

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. $*p < 0.1, **p < 0.05, ***p < 0.01$.

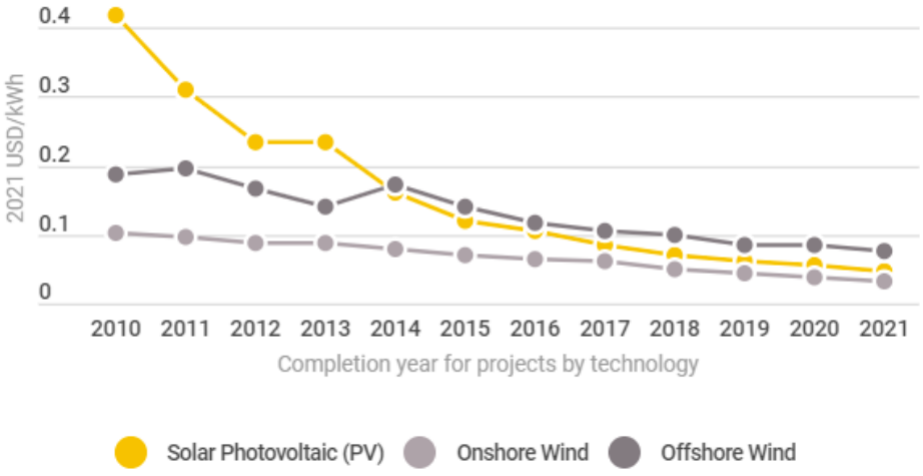
E Additional Figures

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).