Alessi, Lucia; Battiston, Stefano; Kvedaras, Virmantas

Working Paper

Over with carbon? Investors' reaction to the Paris Agreement and the US withdrawal

JRC Working Papers in Economics and Finance, No. 2021/12

Provided in Cooperation with:
Joint Research Centre (JRC), European Commission

Suggested Citation: Alessi, Lucia; Battiston, Stefano; Kvedaras, Virmantas (2021) : Over with carbon? Investors' reaction to the Paris Agreement and the US withdrawal, JRC Working Papers in Economics and Finance, No. 2021/12, European Commission, Ispra

This Version is available at:
https://hdl.handle.net/10419/250162

European Commission

JRC TECHNICAL REPORT

Over with carbon? Investors' reaction to the Paris Agreement and the US withdrawal

Alessi, L.

Battiston, S.

Kvedaras, V.

2021

JRC Working Papers in Economics and Finance, 2021/12

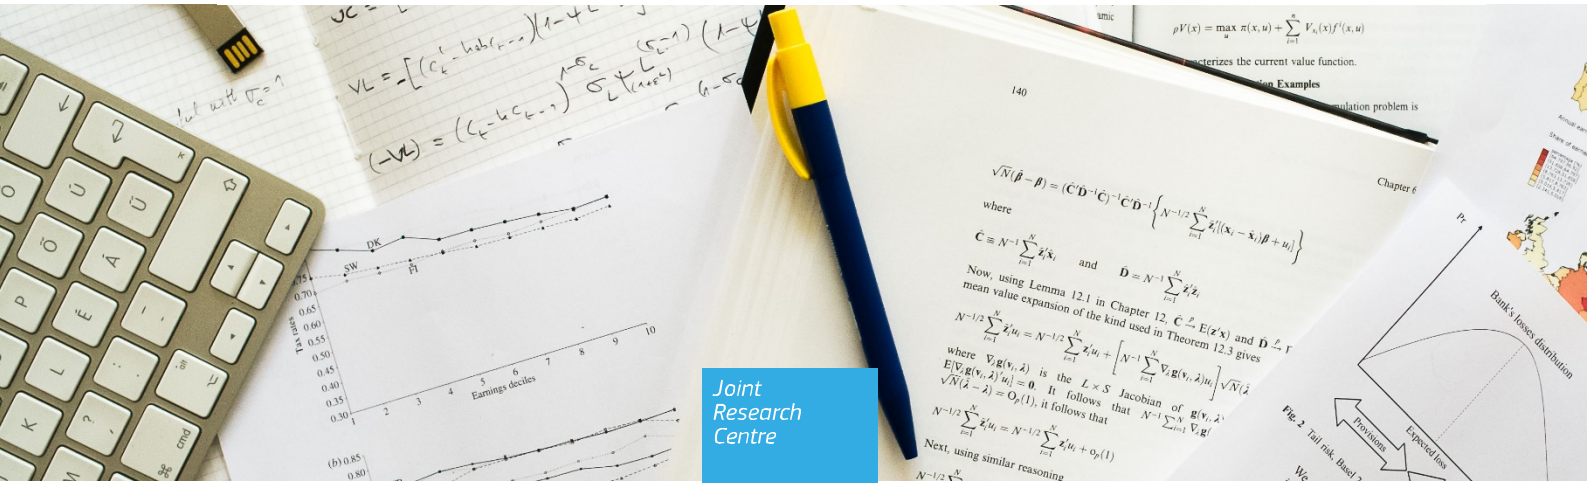

Contact information
Name: Virmantas Kvedaras
Address: Via E. Fermi, 2749, I-21027 Ispra (VA), ITALY
Email: virmantas.kvedaras@ec.europa.eu
Tel.: + 39 0332 783517

EU Science Hub
https://ec.europa.eu/jrc

JRC127845

Ispra: European Commission, 2021

How to cite this report: Alessi et al. (2021) *Over with carbon? Investors' reaction to the Paris Agreement and the US withdrawal*, JRC Working Papers in Economics and Finance, *2021/12,* European Commission, Ispra, 2021, JRC127845.

Executive summary

Climate change mitigation, i.e. the stabilization of global warming below 2 degrees Celsius compared to pre-industrial levels has become a central topic in finance and its implications for financial stability are today a key area of concern for central banks and financial supervisors. Furthermore, there is a consensus on the fact that climate change mitigation cannot be achieved without the engagement of the financial sector. In this context, the Paris Agreement (PA)—that was adopted by the UNFCCC on 12 December 2015 while becoming effective almost a year later—has marked a milestone as it is the first international agreement to state explicitly the role of finance.

In this paper, we study to what extent financial investors—who can play both an enabling or a hampering role depending on their perception of climate policies and their credibility—have adjusted their holdings of carbon-intensive securities in response to the PA and to the subsequent United States (US) withdrawal from the PA; which the US administration announced on June 1st of 2017. We focus on equities issued by European Union (EU)-resident firms, and we carry out a multi-period difference-in-difference analysis on matched high- and low-carbon firms that identifies the dynamics of the impact. To measure investors' stakes in carbon-intensive companies we focus on a participation metric, representing the share of stocks owned by a given holder in terms of the total market capitalization of a company. We construct it using data from a confidential database of securities holdings of the European Central Bank (ECB), namely the Securities Holding Statistics (SHS) database, where investors' holdings are aggregated at the level of the institutional sector and by country.

We find evidence that investors have significantly reduced their exposure to carbon-intensive assets in response to the PA and that the trend reverted after the US withdrawal announcement, in connection with the increased uncertainty about the viability and credibility of the agreement. However, the extent of the reaction and the intensity of the reversal vary across categories and geographies of the securities holders, their ownership size, and their institutional sector. First, a sharper and more consistent decrease of participation in high-carbon firms is observed for more regulated institutional investors and holders from high-income countries, while other financial institutions and holders from the BRIC (Brazil, Russia, India, and China) countries tended to increase

their participation in these firms. Second, the response of households is less steady over time with a clear change in the trend after the announcement of the US withdrawal from the PA, in contrast to regulated financial institutions for which the reduction is more persistent. Third, larger owners were less willing or able to reduce their participation in high-carbon companies, possibly because of the costs associated with selling large portions of stocks, or with a view to driving the low-carbon transition of these companies.

These results have implications in terms of transition risk transfer. On the one hand, the reduction in overall participation in high-carbon companies by the holders in our sample (i.e. covered in the SHS database) implies an increase in participation by investors who are not in the SHS sample, which are essentially non-EA financial investors. Indeed, based on the subset of holdings by non-EA investors we have in our dataset, we do see an increase in participation in European high-carbon companies by investors located in the BRIC region, in particular. Moreover, we document a transfer of transition risk from more regulated financial institutions towards other financial institutions within Europe. We also find that investors are less willing or able to reduce their participation in those high-carbon firms where they hold large stakes.

Our results have some relevant policy implications. First, global environmental policy has an impact on investors behavior in terms of portfolio allocation. Second, the successful redirection of global financial flows towards climate action (Article 2c of the PA) requires a clear and unanimous signal from the global community of policy makers. Third, as the low-carbon transition picks up speed, a close monitoring of the buildup of transition risk in particular sectors and jurisdictions is warranted.

Over with carbon? Investors' reaction to the Paris Agreement and the US withdrawal

Lucia Alessi[1,*], Stefano Battiston[3,4], and Virmantas Kvedaras[1,*]

[1]European Commission, Joint Research Centre, Ispra (VA), Italy

[*]The opinions expressed are those of the authors only and do not necessarily represent the European Commission's official position.

[3]University Ca' Foscari of Venice, Venice, Italy

[4]University of Zurich, Zurich, Switzerland

16th December 2021

Abstract

How financial investors may react to policy events related to sustainability and climate change mitigation in particular, is a key question with implications for sustainable finance and financial stability. We address this question by carrying out a multi-period difference-in-difference approach on a confidential database of securities holdings of the European Central Bank, and we provide evidence of several effects related to the Paris Agreement. In aggregate, investors reduced their exposure to carbon-intensive assets in response to the agreement, and the trend reverted after the US withdrawal announcement. However, the reaction varies across categories and geographies of the securities holders, their ownership size, and the emissions of owned firms. In particular, transition risk has been taken up by less regulated financial institutions and the BRIC countries. Our results highlight that the redirection of global financial flows towards climate action requires clear and unanimous signals from the global community of policy makers.

Keywords: high-carbon firms, finance, Paris Agreement, stock holdings, US withdrawal.

1 Introduction

Climate change mitigation has become a central topic for sustainable finance and its implications for financial stability are today a key area of concern for central banks and financial supervisors (NGFS, 2019). In this context, the Paris Agreement (PA) has marked a milestone as it is the first international agreement to state explicitly the role of finance. Furthermore, there is a consensus on the fact that climate change mitigation, i.e. the stabilization of global warming below 2 degrees Celsius compared to pre-industrial levels, cannot be achieved without the engagement of the financial sector. At the same time, financial investors can play both an enabling or a hampering role depending on their perception of climate policies and their credibility. Hence, it is crucial to understand how financial investors react to policy developments.

In this paper, we study to what extent financial investors have adjusted their holdings of carbon-intensive (high-carbon, hereafter) securities in response to the PA and to the subsequent United States (US) withdrawal from the PA. We focus on equities issued by European Union (EU)[1]-resident firms, and we carry out a multi-period difference-in-difference (DiD) analysis. We use data from a confidential database of securities holdings of the European Central Bank (ECB), namely the Securities Holding Statistics (SHS) database, where investors' holdings are aggregated at the level of the institutional sector and by country. We find evidence that investors have reduced their exposure to carbon-intensive assets in response to the PA and that the trend reverted after the US withdrawal announcement. However, the extent of the reaction varies across categories and geography of the securities holders, their ownership size, and the level of emissions of owned firms. Our results shed new light on the role of the financial sector in relation to the policy objectives of achieving sustainability goals. However, our results also point to the buildup of financial risks related to the low-carbon transition in less regulated segments of the financial sectors and in particular jurisdictions.

The PA itself was a long process starting from the adoption by the UNFCCC[2] on 12 December 2015 and becoming effective almost a year later, i.e., since 4 November 2016. It marked a shift in the global attitude towards climate change mitigation, adaptation,

[1]Including also the United Kingdom (UK), as we consider the period before the Brexit took place.
[2]Acronym for United Nations Framework Convention on Climate Change.

and finance. Indeed, in addition to providing a legal framework for an international commitment to country-specific emission targets via a variety of mechanisms, it has been a landmark for mobilizing financial investments in climate mitigation (see, e.g., Ellis and Moarif, 2017, Law and Zhang, 2019, Mehling, 2021, Reins and Calster, 2021). On June 1st, 2017, the US administration announced that the US would withdraw from the PA, raising global concerns about the viability of the PA objectives (see, e.g., Dai et al., 2017, Steinhauer, 2018, Zhang et al., 2017a,b). The formal notice of intention to withdraw was given on November 4, 2019, abiding to Article 28 of the PA.[3]

Since the EU has been playing a leading role in global climate action, we test whether investors' attitude towards high-carbon firms located in the EU has changed after the PA. A reduction of investments in carbon-intensive firms could be due to the expectation that EU relevant regulation would become stricter, e.g. via an extension of the EU Emission Trading System (ETS), as well as a removal of exemptions and reduced rates that currently encourage the use of fossil fuels. Furthermore, the heightened attention in the EU towards firms' environmental performance and the introduction of more detailed, mandatory sustainability-related disclosures could negatively impact the reputation of carbon-intensive firms and possibly, in turn, their profitability.

For these reasons, we a priori expect that investors may have actually reduced their stakes in carbon-intensive firms after the PA. Still, whether investors reacted to the PA at all is not obvious, as their reaction would depend on the expectations on scope (how broad and how severe), speed (how quickly), and likelihood of the policy impact. Looking at US withdrawal, what to expect as a reaction is less straightforward. On the one hand, increased uncertainty about the viability of the PA could have halted EU investors' progressive shift away from high-carbon firms. On the other hand, investors could have expected that the US decision would have not impacted the EU plans.

To measure investors' stakes in carbon-intensive companies we focus on a price-invariant stock participation metric, representing the share of stocks owned by a given holder in terms of the total market capitalization of a company. We test whether the PA had a significant impact on this participation metric considering two sets of firms.

[3]The formal withdrawal took place in November, 2020, whereas the US rejoined the PA in February 2021.

The first set, i.e the 'treatment' group, consists of EU firms that are likely to be affected in a negative way by environmental policy changes. These firms are identified based on their greenhouse gases (GHG)/carbon dioxide (CO2) emission levels and their sector of economic activity (see Appendix 6.1), and are dubbed hereafter as 'high-carbon' (HC). The second set of firms, i.e. the 'control' group, comprises firms that will be largely unaffected by environmental policies, as they are active in sectors of the economy that have a low impact on climate and the environment. Firms in the first set are matched to firms in the second set, so that the analysis ultimately only focusses on similar firms, based on size and other characteristics. In order to evaluate the impact on the participation of investors into these two sets of firms, we employ a multi-period DiD approach, which allows to detect gradual adjustments and is suitable to detect trend changes after subsequent events, such as the PA and the announcement of the US withdrawal from it. In particular, for our benchmark exercise we use the Callaway and Sant'Anna (2020) approach, building on the Sant'Anna and Zhao (2020) doubly-robust DiD estimator.

Throughout the paper, based on this approach, we are able to document the following effects. First, the participation of investors in HC firms was significantly shrinking after the PA, compared with non-HC firms, with an overall reduction of HC holdings by about a quarter in relative terms. This trend reversed after the US withdrawal announcement, which increased uncertainty, and whose impact vanished by the end of 2020. Second, a sharper and more consistent decrease of participation in HC firms is observed for more regulated institutional investors and holders from high-income countries, while other financial institutions and holders from the BRIC[4] countries tended to increase their participation in these firms. Third, larger owners were less willing or able to reduce their participation in HC companies, possibly because of the costs associated with selling large portions of stocks, or with a view to driving the low-carbon transition of these companies.

Our research contributes to a better understanding of the implications of global climate policy actions on investors' behavior, with findings being consistent—but not overlapping—with a number of recent studies. The importance of a coordinated global policy is underlined in Bartram et al. (2021), who show that local climate policies are

[4]BRIC stands for Brazil, Russia, India, and China.

likely to fail due to the possibility of firm reallocation when environmental policies are only local. However, they do not explore the effects of the increased uncertainty about the viability of climate policies brought about by one of the key policy participants deciding to renege.

Baiardi and Morana (2020) study the changes in the perceptions of the importance of climate change. In full concordance with our findings in terms of the sign of the impact, they uncover significant changes in the concerns about the awareness of climate change in relation with the PA, the US withdrawal, as well as the Global Climate Strikes. However, they investigate only the impact on perceptions and not the actual financial outcomes. Ramelli et al. (2020) also find a reaction of the European stock market to the first Global Climate Strike. Still on European stock prices, and fully in line with our results on quantities, Alessi et al. (2021) find that the greenium, i.e. the risk premium asked by investors to hold greener stocks, decreased after the PA and the first Global Climate Strike, while it increased after the US withdrawal.

Finally, Reghezza et al. (2021) study the impact of the PA and the US withdrawal on bank lending. They find that, after the PA, European banks reallocated credit away from polluting firms, whereas in the aftermath of the US announcement, lending by European banks to polluting firms in the US further decreased. We find that banks' also reduced their investments in equities of European HC firms.

In this paper, by considering equity holdings of various types of investors in EU firms, we contribute to this stream of literature, investigating the impact of global climate change policies on financial market participants, also stressing the heterogeneity of the reactions across different types of investors.

The rest of the paper is structured as follows. Section 2 links the discussed global policy events with the dynamics appearing in HC and non-HC matched firms. Section 3 presents the econometric estimation results, applying the methodology characterized in Appendix 6.2, and covers estimations at the aggregate level (Section 3.1), several sources of potential heterogeneity (Section 3.2), and a number of robustness evaluations (Section 4). Section 5 concludes.

2 Data, metric and basic illustration

Our analysis is based on confidential security-by-security databases hosted by the European Central Bank. The main source of data is the Securities Holding Statistics Database - Sector module (SHS).[5]. SHS data include holdings by investors that are grouped into institutional sectors, classified according to the ESA2010 methodology (e.g. banks, government, etc.) and available at a quarterly frequency. The SHS database covers holdings of investors residing in the euro area and non-resident investors' holdings of euro area securities that are deposited with a euro area custodian. We focus on stakes into companies that are located in the EU, in the period between 2015Q1 and 2020Q3.[6] The holding information is complemented with information on the issuer side from the Eurosystem's Centralised Securities Database (CSDB), such as issuer name, issuer's sector of economic activity (NACE), and outstanding amounts.

Further information on the issuers is retrieved via commercial databases. Emission data is sourced from Bloomberg. In particular, we use the most populated indicator, which is total greenhouse gases (GHG) emissions in carbon dioxide equivalent if available, else total carbon dioxide (CO2) in thousands of metric tons (Total GHG/CO2 Emissions). Refinitv Eikon is the source for the covariates used for the matching procedure, i.e. the dividend yield, the historical stock return volatility, and the market value.

The key metric that we use in our analysis is investors' stock participation, defined as the (logarithm of the) share of stocks owned by holders in terms of the total market capitalization of a company, both expressed in market value.[7] This metric is invariant to stock price fluctuations, while the level or change in investments or shares in investors' portfolios would not enjoy this property.

Formally, the (log) participation of holder h into company j at time t is calculated as follows:

[5]See `https://www.ecb.europa.eu/stats/financial_markets_and_interest_rates/securities_holdings/html/index.en.html`

[6]SHS data started being collected in the fourth quarter of 2013; however, the quality of the first vintages is not optimal. Equity holdings are recorded as F-511 in the SHS database.

[7]The logarithm transform better satisfies the parallel trends assumption needed for identification of the effect. The difference-in-difference effect thus will establish the relative and not absolute decrease in the participation intensity.

$$y_{h,j,t} = \log\left(\frac{H_{h,j,t}}{M_{j,t}}\right), \tag{1}$$

where $H_{h,j,t}$ and $M_{j,t}$ stand for the market value of holdings of holder h into company j and the total market value of company j, respectively, in period t.[8]

In our main estimation exercises, the dependent variable is the average participation indicator, calculated as follows:

$$y_{j,t} = \frac{1}{N_H}\sum_h y_{h,j,t} \tag{2}$$

with N_H denoting the number of holders. In heterogeneity analyses we consider furthermore only some subsets of holders separated by a particular dimension, e.g., their sector, country, or investment size.

To motivate the estimations that will be presented later on, the left panel of Figure 1 plots the dynamics of average participation indicator in the two matched sets of treated 'high-carbon' firms, on the one hand, and of control (untreated, or 'comparison') firms, on the other. These latter are firms characterized by low emission levels and not belonging to (nor serving) the fossil, cement, and other directly or indirectly pollution-intensive industries. Firms are matched based on company size (market value), as well as the yield and volatility of their stock returns (see Appendix 6.1 for details).

[8]Security-by-security data are aggregated by issuer and holder before taking the logarithmic transformation. As data are unconsolidated, in some cases (about a quarter of all issuers) the total sum of holdings is greater than the market capitalization of a company. We exclude the 1% most extreme values, while for the others we shrink all holdings proportionally by using the correction factor $s_{h,j,t} = \frac{M_{j,t}}{\sum_h H_{h,j,t}}$ whenever the total sum of holdings reported in the SHS, i.e. $\sum_h H_{h,j,t}$, is larger than the market capitalization $M_{j,t}$.

Instead of using this correction, one could drop issuers, for which the total sum of holdings in the SHS is greater than the market capitalization. The outcome is qualitatively similar but the procedure results in a smaller number of matched firms.

Figure 1: Dynamics of average participation in the matched groups of treated (high-carbon) and control firms and difference between the two groups

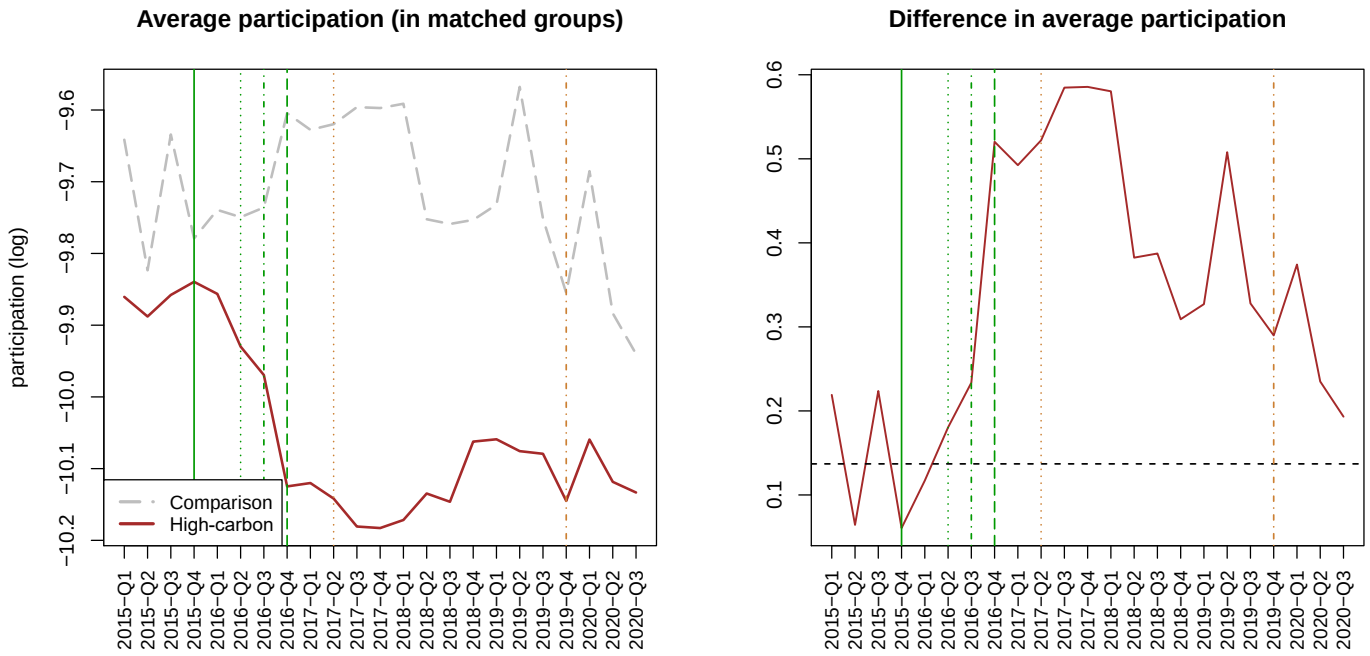

The first four vertical lines (in green) are connected with the process linked to the Paris Agreement: on the 12th of December 2015 (2015-Q4 in the figure) the text was adopted by consensus by the Parties of the UNFCCC; on the 22nd of April 2016 (2016-Q2) the Agreement was opened for signature; in October 2016 (2016-Q3) a large enough number of ratifying countries was reached for the Agreement to enter into force; and on the 4th of November 2016 (2016-Q4) it actually went into effect. The remaining two vertical lines (in light brown) mark the dates related to the US withdrawal, namely the 1st of June 2017 (2017-Q2), when the US announced the withdrawal, and the 4th of November 2019 (2019-Q4), when the formal notice of intention to withdraw was given. Looking at the right panel, the horizontal (black) dashed line indicates the initial difference between the average participation in the control group and the HC group observed before the PA, i.e. during the period from the first quarter of 2015 until the first quarter of 2016, while the solid line indicates how this difference in participation evolved over time.

This picture reveals that, after the PA, investors reduced their participation in HC firms relatively to the control group. After the announcement of the US withdrawal, this trend reversed, with the difference in participation between the two groups be-

coming progressively smaller. The difference in participation spikes up in the second quarter of 2019, possibly in connection with the first two Global Climate Strikes for Future that took place on the 15th of March and the 24th of May, which seemingly influenced the climate change awareness (see e.g. Baiardi and Morana, 2020) and financial markets (see, e.g., Ramelli et al., 2020). Although our analysis might be also capturing other processes that could have had an impact on equity holdings of high-carbon relative to other companies, the largest changes of magnitude and direction seem to be dominated by and well correlated with the dating of the Paris Agreement and the US withdrawal announcement.

In the next sections, we use several econometric approaches to evaluate whether the difference visible in Figure 1 is statistically significant and to check whether the established pattern still holds using a more refined analysis framework.

3 Empirical evaluation

In this section, we present the main empirical findings on the dynamic pattern of the impact of the PA on investments. In Section 3.1, we start by considering the impact at the aggregate level, i.e. looking at all investors and issuers. Then, in Section 3.2, we look at four possible sources of heterogeneity in the responses: i) investor institutional sector; ii) investor geographic location; iii) investor participation size; and iv) issuer GHG emissions. Finally, in Section 3.3 we discuss the statistical significance of the results presented in the previous sections.

The DiD framework involves two crucial modelling choices. One is establishing the timing of the treatment, the other is the definition of the treated and control groups. With respect to the former, the reaching of the PA was a long process marked by a number of events. Hereafter, we adopt the quarter of the opening for signature of the PA (2016-Q2) as the beginning of the treatment, since the negotiation of the text by the UNFCCC parties was not binding as yet in terms of any implications. Nevertheless, even this moment might be somewhat early, as we actually find that the largest adjustment took place when the PA was ratified and went into force.

With respect to the definition of the groups, the treatment group includes firms in the top tercile of the emission distribution (HC firms), as they are most likely to

be affected by the PA, while in the control group we include firms in the bottom tercile.[9] We further exclude from the control group firms whose main activity falls in the airlines, cement, electricity, fossil, and steel sectors. This we do for two reasons. First, some firms working in these sectors might have low emission levels but belong to high-emissions value-chains. Second, the control set should be unaffected by the treatment. However, low-emission firms active in the above sectors could in fact be positively affected by the PA, as funding could arguably move from more polluting to less polluting firms in the same sector.

Finally, instead of looking at all treated and control firms, we only consider similar firms across the two groups. As matching procedure, for the main analysis we use the Coarsened Exact Matching (CEM), while for the robustness checks we employ the genetic matching algorithm (GEN1) with generalized Mahalanobis distance, as well as a greedy nearest neighbour matching (see Section 6.1 for details).

3.1 Evaluation at the aggregate level

Given the possibility of gradual realization of the impact and the regime changes expected in connection with the PA and the US withdrawal, our quantity of interest is the *period-specific* 'average treatment effect on the treated' (ATT, see Callaway and Sant'Anna, 2020, Chaisemartin and Haultfoeuille, 2020, and Xu, 2017). To fix ideas, consider periods indexed by t and firms indexed with $j \in \{\mathbb{T}, \mathbb{C}\}$, where $\mathbb{T}$ and $\mathbb{C}$ are the sets of indexes connected with treated and control firms. Let $y_{j,t}$ stand for the average holdings relative to the total market capitalization (in logarithmic terms) as defined in eq. (2). Next, let $D_{j,t} = \mathbb{1}\{j \in \mathbb{T}\} \cdot \mathbb{1}\{t \geq 2016\text{Q2}\}$ denote the treatment status which takes value one for treated firms starting from the second quarter of 2016 and zero otherwise. Furthermore, let $Y_{j,t}(1)$ and $Y_{j,t}(0)$ denote the potential outcome with and without the treatment, with the actual outcome $y_{j,t} = Y_{j,t}(D_{j,t})$ depending on the treatment state. The ATT is then defined as follows:

$$ATT_t = \mathbb{E}\big[Y_{j,t}(1) - Y_{j,t}(0) | D_{j,t} = 1\big], \tag{3}$$

[9] Similar results appear also using the top quartile, but this shrinks substantially the number of matched firms.

potentially, conditioning additionally on a vector of other explanatory variables. Notice that, although our aim is to evaluate the impact of two separate events—the PA and the US withdrawal—it would not be possible to evaluate them separately, for two reasons. First, the two events arguably affect the same set of firms but in opposite directions; hence, a non-dynamic DiD estimator taking the PA as treatment would actually yield the average effect of the two events on the treated firms, which may be overall insignificant. On the other hand, focussing only on the US withdrawal would not be appropriate either, as the PA already induced trend differences between the two groups of firms, which would violate the parallel trend assumption.

For the main analysis we employ the doubly-robust ATT estimator proposed by Callaway and Sant'Anna (2020), applying it to the properly weighted matched set of treated and control firms (see Appendix 6.1) and focusing on the case where all treated firms are treated at the same time. Figure 2 on the following page plots the corresponding estimated ATTs with their 90% bootstrap-based confidence bands.[10] The figure reveals a few important patterns. First, before the PA, there is no significant trend difference between the treated and comparison groups: the null hypothesis of parallel trends cannot be rejected at the usual significance levels neither taken individually nor if tested jointly (see also Table 2 in Section 3.3). Second, a sharp deviation appears from zero towards highly significant ATTs after the PA. The effect continues to increase (in absolute terms) approximately until the period of the US announcement about the intention to withdraw from the PA; namely, 2017-Q2. The maximum effect is reached just one quarter later than that of the announcement, which is not exceptional, as the announcement was made during the second half of 2017-Q2, i.e., in June. Third, after the US intention to withdraw became public, the ATTs started to decrease (in absolute terms) lagging from the announcement by a quarter.

Finally, there is a clear increase in the confidence bands of ATTs: first, after the PA and, even further, after the US withdrawal. Apart from a genuinely larger uncertainty after the US withdrawal, this increase is likely to be also driven by heterogeneous reactions of different groups of investors (to be explored in the next section). Moreover, the larger confidence bands around the estimated ATTs in the post-US withdrawal

[10] The bootstrap-based inference is used with clusters/blocks at the issuer level. The implementation relies on the `att_gt()` function of package `did` for `R` (see `https://CRAN.R-project.org/package=did`).

Figure 2: The estimated period-specific ATTs

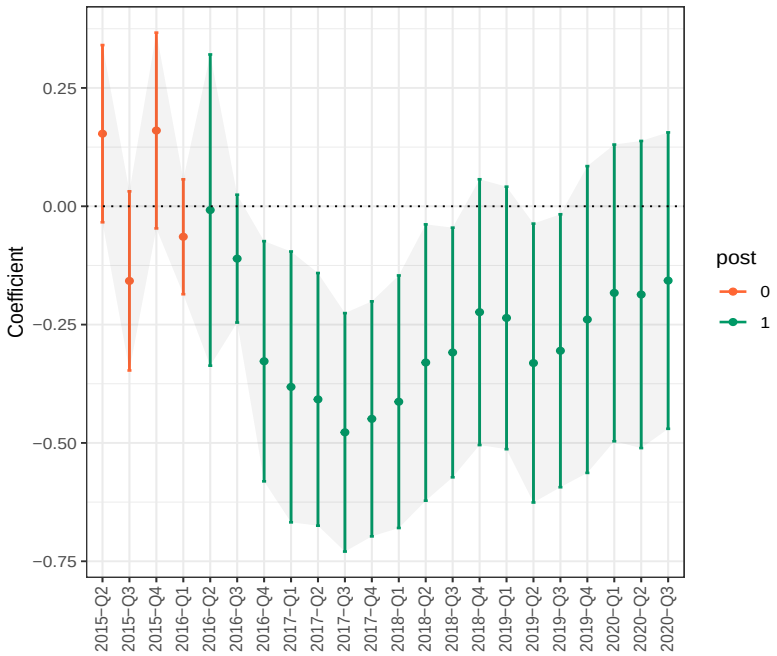

Note: The normalization is with respect to the first observation of the non-treatment period.

announcement might be further driven by two counteracting forces. First, the US withdrawal increased doubts about the viability and success of global climate policy, lacking the US commitment. This reaction would shrink the effect. Second, part of investors might have even decided to reallocate a part of their HC investments towards US HC firms, which would result in a reduced participation in European HC firms just like right after the PA.

The increase in uncertainty about the viability and credibility of the PA and the emergence of even large heterogeneity of reactions after the US announcement resulted in a no more significant difference between treated and untreated by the end of the analysed period.[11] However, the effect might have vanished also in connection with other reasons. First, the emergence of more tangible problems and risks due to the appearance of the Covid-19 could have changed the perception of priorities and the reaction of investors. Second, the initial investors' valuation and expectations with reference to EU policies could have been in contrast with the perceived actual imple-

[11]In January 2021, i.e. after the end-date of our sample, President Biden announced that the US would rejoin the PA and the US officially rejoined the following month.

mentation and achievements. Third, EU policies announced and implemented, with a particular reference to the EU green taxonomy, have clarified that even high-carbon activities can be green, if they nevertheless minimize emission levels by using state-of-the-art technologies.[12] Hence, investors might have progressively started looking more closely at HC companies and screen them based on e.g. the existence of a commitment to emission reduction and/or a broader transition strategy, the greenness of their capital expenditure, or in comparison to peers, and not just based on the current absolute level of emissions—which was a natural criterion for investors when there was no more sophisticated definition of 'green activity'—thus diminishing the relevance of this particular indicator.

3.2 Heterogeneous responses

In this section, we take a look at different potential sources of heterogeneity in the responses of investors. In particular, we cover four types of heterogeneity. We consider that investors belong to different institutional sectors and are located in different countries, and also that the size of their stakes in investee companies compared to companies' total market capitalization can be larger or smaller. Finally, we investigate whether investor responses could also vary based on the level of emissions of the issuers.

For the estimation of the ATT in the following two subsections, the dependent variable $y_{j,t}$ is defined as in Equation 2 but calculated by considering only the holders belonging to a given sector or geography. In Section 3.2.3, the dependent variable is without pre-averaging as given in Equation (1), whereas it remains the same as in the aggregate analysis in the last subsection 3.2.4.

3.2.1 Holder sector

Here, we look at various types of investors split into institutional sectors (e.g. banks, households, non-financial corporations, etc.). In Figure 3, we plot the estimated dynamic reactions only for those sectors that have a significant overall response (either negative or positive) at least at the 10% significance level (see Table 2).

[12] An early-feedback EU Taxonomy proposal was put forward in December 2018 by the Technical Expert Group on sustainable finance established by the European Commission, which published a draft report in June 2019 and a final report in March 2020. The Taxonomy Regulation entered into force in July 2020.

Figure 3: Holder-sector and period-specific ATTs

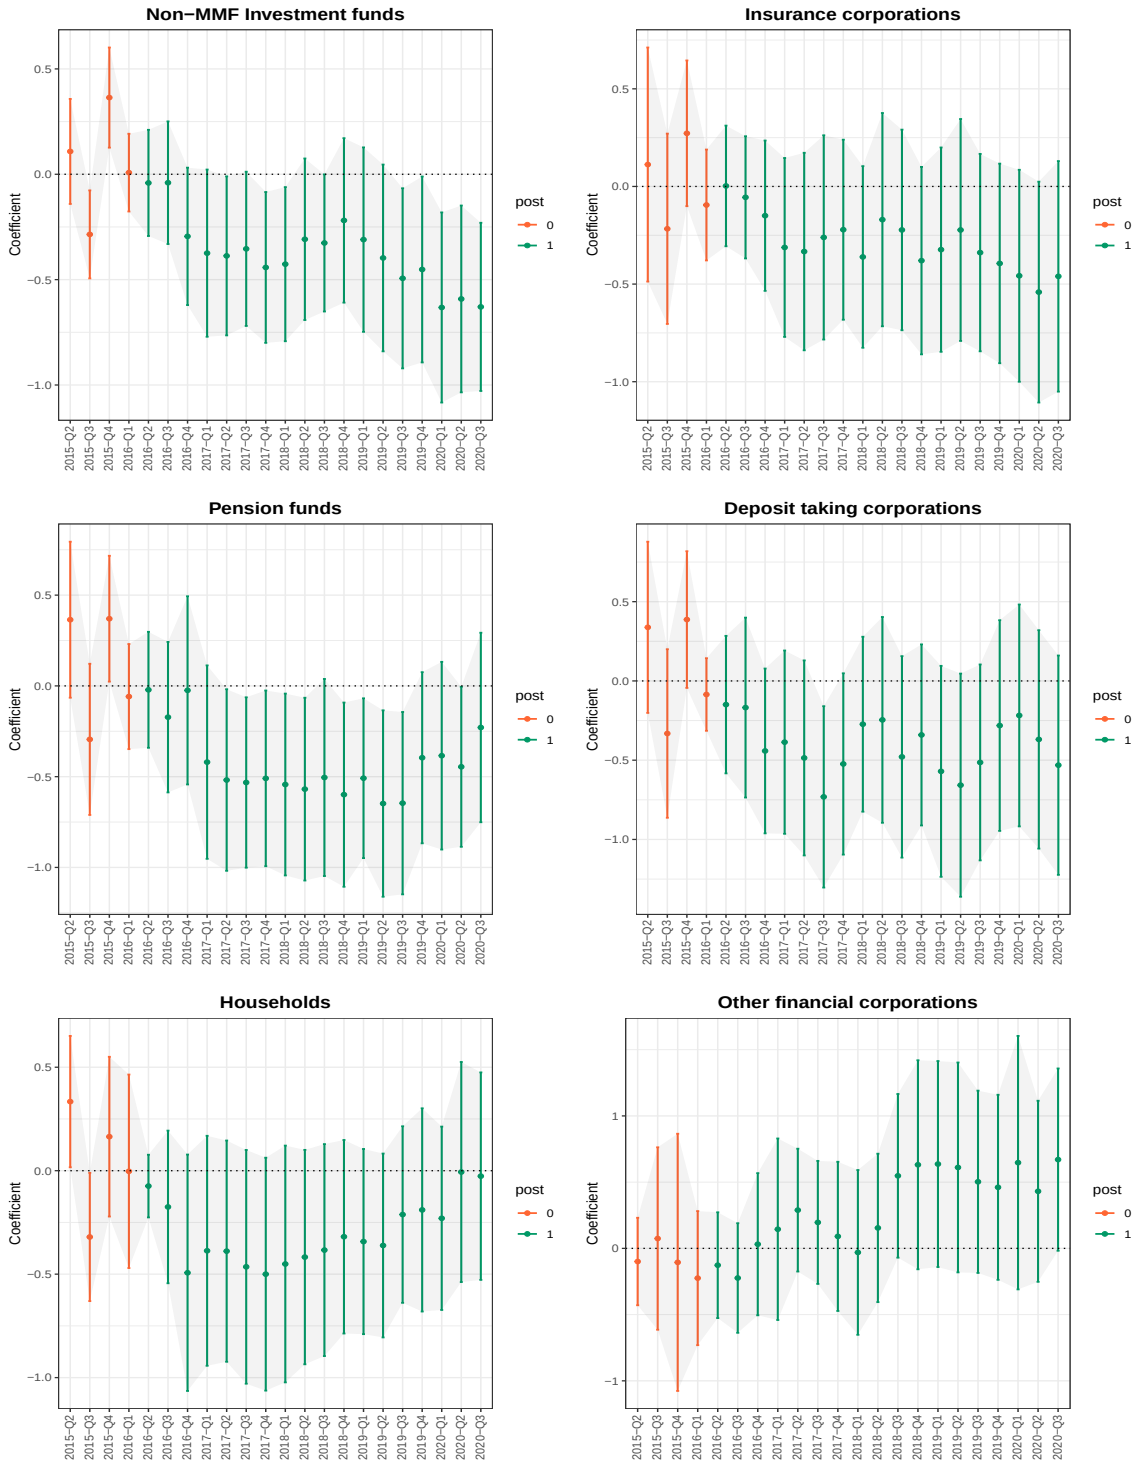

A couple of observations can be drawn. First, not all sectors reduce their relative participation in HC firms. In particular, financial corporations other than financial intermediaries, i.e. financial institutions which trade only little of either their assets

or their liabilities on open markets, even tend to increase their stake in HC firms (see the bottom-right panel in Figure 3). Hence, there is a shift of transition risk[13] connected with HC firms from more regulated financial institutions (banks, insurance firms, investment and pension funds) towards less regulated financial institutions. The increase in participation by the latter, typically active in the trading of derivatives and the intermediation to foreign acquisitions, is likely to be driven either by speculative investments or by increased demand by non-EU investors, willing to acquire stakes in HC companies (see next section).

Second, looking at regulated financial institutions, insurance corporations and investment funds[14] have reduced even further their relative participation in HC firms since around the middle of 2019. On the other side, pension funds seem to have slightly softened their initial response since about the same time, whereas the response of banks (deposit taking corporations), after an initial reduction in their participation in HC firms, did not show any clear trend. Overall, banks, investment funds, pension funds and insurers display a consistent trend in reduction of participation in HC firms even after the withdrawal announcement.

Third, the response of Households is less steady over time with a clear change in the trend after the announcement of the US withdrawal from the PA, in contrast to regulated financial institutions. For households, indeed, the participation in HC firms reverts to the pre-PA situation in the aftermath of the US withdrawal. Rather than to speculative motives, the behavior of households is likely to be more sentiment-driven and connected with the increased uncertainty in the continuity of the global anti-pollution policy after the US withdrawal. Not only the increased uncertainty after the US withdrawal could have affected households' opinion more sizeably, but they might also have perceived a growing disconnect between their initial valuation and expectations of the PA implications, on the one hand, and the actual situation and progress of climate policies, on the other—which they might also have a less clear picture about, compared to professional investors.

[13]The literature identifies essentially two types of climate-related financial risks. Transition risk relates to the risk that exposures toward economic activities that will be negatively impacted by the low-carbon transition will result in non-performing loans and investments owing to the underlying assets becoming stranded. Physical risk relates to the consequences of more severe and frequent climate-related natural disasters.

[14]Excluding Money Market Funds (MMF).

3.2.2 Holder area

The inference about the impact that geographic differences may have on the behavior of investors has more potential caveats than the split by holder sector considered previously. First, we have information on non-EA investors only through EA custodians, which are mandated to report on their holdings of EA securities by investors resident inside and outside the EA.[15] Hence, our conclusions about the geographic patterns for non-EA investors are valid only as much as the behavior observed in the SHS can be extrapolated and generalized for all investors from those regions. Second, we observe only the behavior of the end-investor (e.g. a financial subsidiary), which can be located in a different country from the ultimate investor.[16] Finally, about a quarter of the records are investments from tax havens,[17] for which we also do not have information on the location of the ultimate investor.

Keeping these limitations in mind, some interesting patterns emerge. In Figure 4 we again report only a selection of more interesting cases having a significant total response at least at the 10% significance level (apart from a single specific case). Note that we merge tax havens with EU countries and the UK, as we assume that most of the ultimate holders investing first in a tax haven and then in Europe via a EA custodian are in fact European, as are most of the holders investing directly via a EA custodian.

[15]The number of reporting countries in the SHS is twenty three, including all countries from the Euro Area (EA) and a few additional ones. Most non-euro area EU countries also collect SHS data, though on a best-effort basis.

[16]It is of interest to note that the number of records reported as holdings of Luxembourgish investors is of about the same size as for German or French investors.

[17]Including the non-cooperative and gray countries indicated in `https://ec.europa.eu/taxation_customs/sites/taxation/files/eu_list_update_18_02_2020_en.pdf`. From here, we use the earliest available list of December 5, 2017.

Figure 4: Holder-area and period-specific ATTs

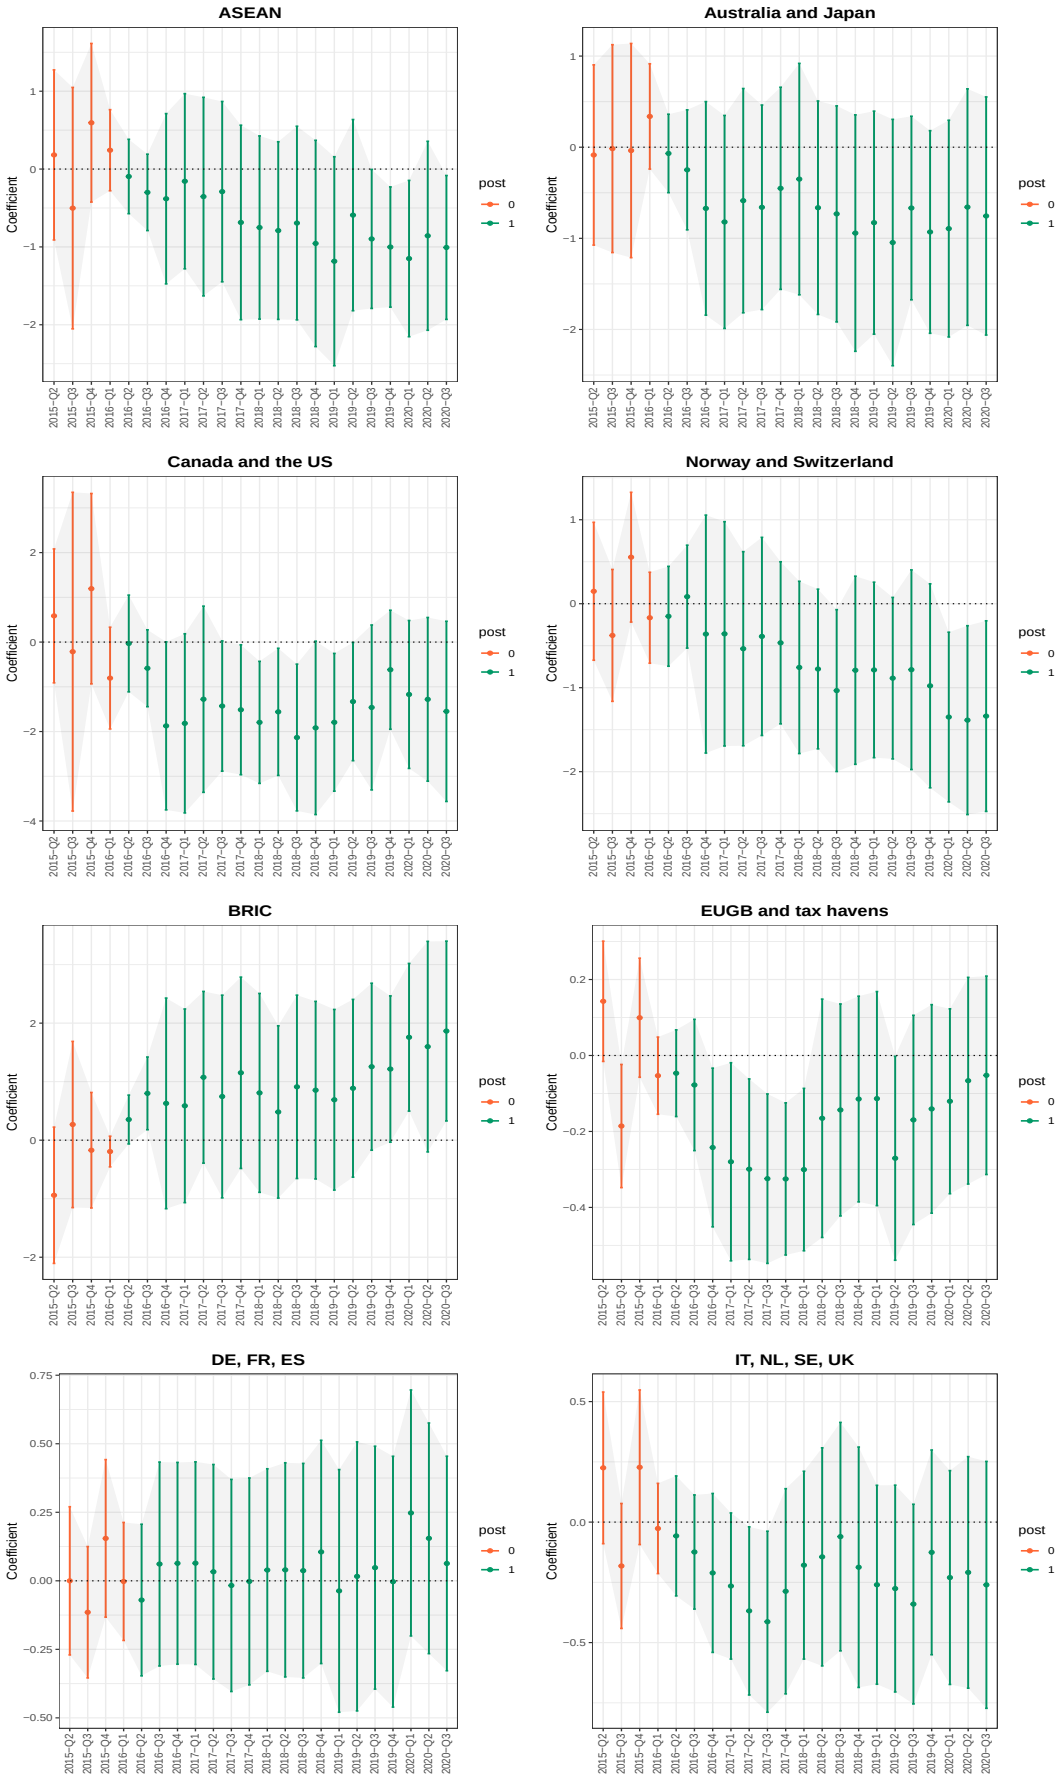

Note: ASEAN stands for the Association of Southeast Asian Nations; BRIC is a grouping acronym which refers to Brazil, Russia, India, and China; and EUGB signifies the former EU27 with the UK.

A typical reaction of holders from more developed countries is to reduce the participation in HC firms after the PA. The reduction is more sizable for holders from Canada and the US and from Norway and Switzerland. The decrease in the latter seems to become even more pronounced by the end of the investigated period.

On the contrary, the participation in HC firms tends to increase by investors from the BRIC region, covering Brasil, Russia, India, and China.[18] Different motivations are likely to be behind those acquisitions stemming from these countries. First, BRIC investors may be more willing to take up climate transition risk to earn higher returns. Second, there could be geopolitical interests underpinning such investments. In particular, being Russia the main EU supplier of crude oil, natural gas and solid fossil fuels, it has a direct interest in the European energy sector, whereas foreign direct investment is one of the key levers in China's approach to attain a dominant position in international markets.

Finally, the participation of holders from the EU countries (and the UK), taken as a whole together with those from tax havens, follows the previously established hump-shaped reaction pattern. However, investors from different countries may display a different reaction. In particular, there is practically no change after the PA in the relative holdings of HC firms by holders from Germany (DE), France (FR), and Spain (ES), whereas the participation in the HC sector tends to be significantly smaller after the PA for holders from Italy (IT), the Netherlands (NL), Sweden (SE), and the UK.

3.2.3 Participation size: Quantile treatment effects

In this section we test whether the size itself of the participation in HC companies at the time of the PA might have affected investors' reactions. In general, large shareholders might be less willing or able to reduce their participation in the firms where they hold large stakes compared to the total market capitalization, because of higher liquidation costs due to market impact, potential loss of influence in the decision-making process of the company, and, possibly, because of a better knowledge of the company's actual situation and plans. Some comparatively larger shareholders might have even tried to exploit the aftermath of the PA to reach control of some firms if, based on private

[18]This pattern is mostly driven by other countries than India, because there are very few records about Indian investors in the SHS.

information, they knew those firms would be not at risk as much as perceived by the market. Overall, we expect that smaller holders adjust their participation quicker and, in relative terms, more sizably than large investors.

To check the potential significance of the participation size, we look at the quantile treatment effects on the treated (QTT) by evaluating the changes at particular quantile levels of holdings in terms of the previously defined participation indicator.[19] Figure 5 plots the estimated effects against the various quantiles of the distribution of the participation indicator (tau) considering different periods: 2016-Q3 and 2016-Q4 as periods during which the initial adjustment takes place, and, additionally, the last available quarter of each consequent year; namely, 2017-Q4, 2018-Q4, 2019-Q4, and 2020-Q3.[20]

Figure 5: Estimated QTTs

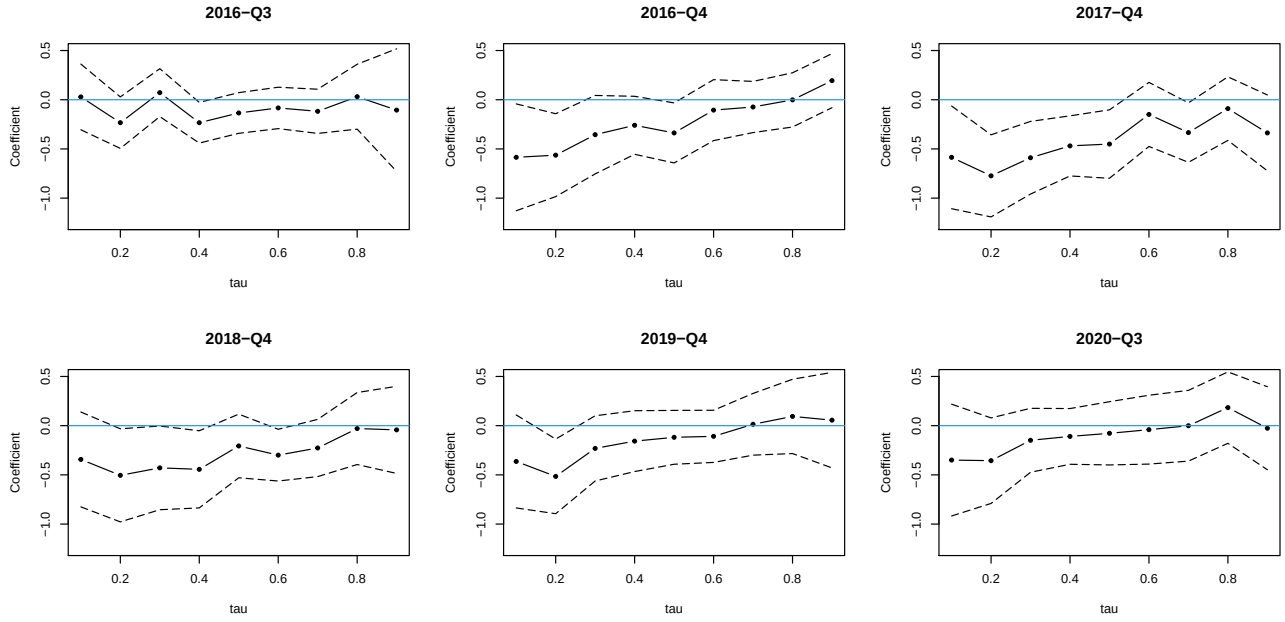

Note: The x-axis reports the various quantile levels of the distribution of the participation indicator.

Figure 5 reveals that, first of all, the dynamics of the estimated QTTs over the considered periods are broadly consistent with those of the ATTs depicted in Figure 2. Namely, the largest adjustment takes place from 2016-Q3 to 2016-Q4 with a further mild reduction towards 2017-Q4. Afterwards, the impact on the treated generally decreases (in absolute terms), both over the years and the quantile levels.

[19]The Athey and Imbens (2006) estimator is employed here as implemented in the `qte` package for `R` (see `https://CRAN.R-project.org/package=qte`). Similar results hold using other estimators also available in the package.

[20]In each case, the change (log-difference) from 2015-Q4 is considered.

In addition, the presented QTTs reveal that, indeed, the adjustment by holders with large participation indicator (tau values close to one) is much smaller, if any. Whereas the reduction in participation by smaller holders is much larger and significant over most considered periods. However, this does not happen as much for the very smallest holdings, probably because of the much smaller potential loss. Hence, at least a part of the observed uncertainty around estimated ATTs seems to be driven also by this kind of heterogeneity.

3.2.4 Heterogeneity in terms of emissions

The uncertainty around estimated ATTs can be further due to the fact that the treated group is still quite heterogeneous in terms of emission levels and, therefore, the reaction of investors may also differ across HC investee companies. Just as a potential indication of such an effect, we consider further the following panel data model of the (log) participation ($y_{j,t}$) with individual issuer and period effects (α_j and λ_t, correspondingly)

$$y_{j,t} = \alpha_j + \lambda_t + \beta_0 D_{j,t} + \beta_1 D_{j,t} \cdot E_{j,t} + \boldsymbol{\theta}' \boldsymbol{z}_{j,t} + \xi_{j,t}, \tag{4}$$

which includes not only the treatment indicator ($D_{j,t}$) but also its interaction with emissions ($D_{j,t} \cdot E_{j,t}$) while vector $\boldsymbol{z}_{j,t}$ comprises other potentially relevant controls, and $\xi_{j,t}$ signifies the remaining error term.

The estimated parameters of interest β_0 and β_1 are reported in Table 1 together with their robust asymptotic standard errors[21] provided in regular parentheses below each coefficient. For comparability, Columns (2) and (1) report the results both with and without the interaction term with emissions, correspondingly.

[21] The variance-covariance matrix is clustered by issuers.

Table 1: Heterogeneity in terms of emissions

	Dependent variable: participation (in logs)			
	(1)	(2)	(3)	(4)
treatment (β_0)	-0.207^{***}	-0.135^{*}	-0.064	-0.050
	(0.074)	(0.073)	(0.089)	(0.046)
treatment $*$ emissions (β_1)		-0.010^{***}	-0.011^{*}	-0.009^{**}
		(0.002)	(0.006)	(0.004)
treatment $*$ emissions-to-sales			-0.003	
			(0.005)	
treatment $*$ emissions-to-assets				-0.012
				(0.007)
Observations	2,772	2,772	2,160	2,160
R^2	0.832	0.834	0.841	0.841
R^2 (within)	0.0134	0.0242	0.015	0.015
F Statistic (within)	41.98^{***}	35.96^{***}	12.14^{***}	12.14^{***}
Degrees of freedom (of F Stat.)	[1; 2621]	[2; 2620]	[3; 2034]	[3; 2034]
Issuer and period effects	+	+	+	+

Note: *p<0.1; **p<0.05; ***p<0.01

Despite that the usage of an aggregate parameter might not be proper and fully informative due to the already established heterogeneity and variation of the impact over time, the results are quite indicative about the importance of differing emission levels as a source of heterogeneity. Namely, the emission level for the treated is highly significant while the unconditional effect (β_0) not only shrinks in absolute terms by about 65% as we switch from Column (1) to (2) but also becomes less significant.

Columns (3) and (4) are further included to test the significance of emission levels as compared with indicators of emission intensity, defined as emissions over sales or emissions over assets, respectively. As shown in Figure 8 in Appendix 6.1, these two indicators also exhibit two very different distributions for treated and controls, so

could in principle work as reference indicators. However, we find that emission levels, interacted with the treatment indicator, remain significant even when interaction terms with emission intensity indicators are included in Columns (3) and (4). In fact, emission intensity indicators are not significant. This indicates that investors' decisions are mostly based on emission levels rather than on emission intensities. Furthermore, note that our DiD set-up, although based on emissions levels and not on intensities, takes care of the firm size dimension in the matching step, as size is one of the matching controls.

Finally, it should be pointed out that the aggregate impact reported in Column (1) of Table 1 and established here using the panel data modeling framework, is somewhat smaller than the one using the aggregate dynamic impact estimator suggested by Callaway and Sant'Anna (2020), which overall significance and size will be considered in the next section.

3.3 Overall significance of the impact

In previous sections the stress was on the dynamic pattern of the response, whereas in this section we summarize the results by presenting the overall significance of each previously considered case. Taking all post-PA periods into account, Table 2 reports the respective overall doubly-robust ATT coefficients. Their bootstrap-based standard errors are reported together with the simultaneous 90% and 95% bootstrap confidence bands. The ATTs that are significant at the 10% and 5% significance levels are correspondingly marked with * and **. Furthermore, the p-values are reported that are relevant for the pre-testing of parallel trends assumption (see the column named p-val.(Par.Tr.)).

Apart from one case, the overall PA impact is statistically significant at least at the 10% significance level. In all the considered cases, the parallel trends assumption cannot be rejected. Hence, our research shows that, in the aftermath of the PA, financial investors significantly reduced their participation in European HC companies.

Table 2: Overall significance of the PA impact

	Specification	Coeff.	S.E.	90% conf.bands		95% conf.bands		p-val.(Par.Tr.)	n.treat.	n.comp.	n	T
Aggregate	Base (overall)	-0.282**	0.088	-0.436	-0.128	-0.454	-0.110	0.802	59	69	128	23
	Non-MMF inv. funds	-0.373**	0.110	-0.553	-0.193	-0.626	-0.120	0.552	59	69	128	23
	Insurance corp.	-0.289*	0.151	-0.543	-0.035	-0.606	0.027	0.894	59	69	128	23
Holder	Pension funds	-0.426**	0.142	-0.660	-0.193	-0.687	-0.166	0.630	58	67	125	23
sector	Deposit taking institutions	-0.409**	0.173	-0.689	-0.129	-0.739	-0.080	0.695	59	69	128	23
	Households	-0.301*	0.163	-0.557	-0.046	-0.626	0.023	0.424	59	69	128	23
	Other financial corp.	0.315**	0.156	0.075	0.554	0.001	0.628	0.908	59	69	128	23
	ASEAN	-0.674**	0.336	-1.199	-0.149	-1.286	-0.063	0.847	36	30	66	23
	Australia and Japan	-0.665*	0.413	-1.323	-0.007	-1.475	0.145	0.990	36	26	62	23
	Canada and US	-1.396**	0.516	-2.256	-0.537	-2.351	-0.442	0.644	31	32	63	23
Holder	Norway and Switzerland	-0.725**	0.375	-1.300	-0.150	-1.406	-0.044	0.949	40	39	79	23
area	BRIC	0.981**	0.513	0.202	1.761	0.184	1.779	0.761	36	22	58	23
	EUGB and tax havens	-0.181**	0.074	-0.306	-0.056	-0.345	-0.016	0.525	59	69	128	23
	DE,FR, ES	0.047	0.129	-0.174	0.268	-0.217	0.311	0.969	59	69	128	23
	IT, NL, SE, UK	-0.222*	0.117	-0.412	-0.032	-0.463	0.018	0.759	59	69	128	23

4 Robustness checks

In this section we present a number of robustness checks by varying the estimation and matching methods, restricting issuer and holder countries, considering non-aggregated data, etc. The related figures are placed in Appendix 6.3.

First, we evaluate the robustness of the presented findings to different estimation methods. The main results underlying Figure 2 were obtained using the Callaway and Sant'Anna (2020) approach. As robustness checks, we apply the DID_l estimator of Chaisemartin and Haultfoeuille (2020), which is unbiased under heterogeneous and dynamic effects,[22] and the generalized synthetic control estimator proposed by Xu (2017), which further allows for certain dynamics of the error term.[23] The former has similar identification assumptions to that of Callaway and Sant'Anna (2020), including the parallel trends in the treated and comparison groups before the treatment. The generalized synthetic control approach of Xu (2017) embodies the idea of Abadie et al. (2010) about the synthetic matching and has a different set of assumptions for causal identification. Hence, the consistency of empirical results based on these different estimators would reveal robustness not only to different matching strategies but also to alternative identification assumptions. Finally, all these methods are robust to certain cross-sectional and temporal heterogeneity of the impact.[24] Figure 9 in Appendix 6.3 reports the findings based on these two additional methods in the left panel and the right panel, respectively. Despite some variation in the estimated level of the impact, the shape of it is consistent across all the employed approaches.

Second, we explore the relevance of the aggregation level. The base results provided in Figure 2 were obtained considering the average holders' participation aggregated by issuers, which is to say, by averaging over different holders of the same security. Thus, essentially, we considered a panel data structure over issuer and time, which, as a byproduct, also allowed the estimation of model (4). Figure 10 plots similar results but using non-aggregated cross-sectional data at the issuer-holder level (as in

[22] We used the Stata `did_multiplegt` command (see Chaisemartin et al., 2021).

[23] In particular, we allow for first order serial correlation. Furthermore, the optimal number of factors (for projections) is selected by the cross-validation procedure. We employed the `gsynth()` package for `R` (see `https://CRAN.R-project.org/package=gsynth`).

[24] It is important to note that, in all these cases, the bootstrap-based inference is used with clusters/blocks at the issuer level.

eq. (1)), where holder records vary by holder sector and country. In this case, the included cross-sectional fixed effects comprise any observed issuer-holder combination. The main dynamic pattern again remains similar to the one reported previously.

Third, we explore whether the results could be influenced by the Brexit process that also initiated in 2016-Q2 (period 0 in the figures under consideration), as the respective voting took place in June 23, 2016. Figure 11 plots the results when we drop UK issuers (left panel) and both the issuers and holders from the UK (right panel) from the dataset under consideration.We do this in order to eliminate potential interferences due to Brexit-related changes in the behavior of investors with respect to UK issuers, as well as in the behavior of UK investors. The qualitative picture remains similar after both adjustments.

Fourth, the base results were obtained relying on the Coarsened Exact Matching (CEM) approach by Iacus et al. (2012).[25] Figure 12 plots the results using alternative matching methods. Namely, the left panel relies on the genetic algorithm-based matching, whereas the right panel plots the results using the nearest-neighbor matching approach. All main patterns established previously using the CEM are also retained, although the estimated size of the impact seems to be more moderate. Partially, this can stem from the fact that a larger number of matched firms are selected by the two additional methods which also leads to some deterioration of the quality of the matching (see Appendix 6.1 for additional details).

Fifth, for the main analysis we based our matching on the five-year (2011-2015) pre-treatment period averages of the matching variables. Figure 13 plots in addition the results when the three-year (2013-2015) average and the value of 2015 alone are used instead.

Sixth, as explained in Appendix 6.1, after performing the matching we impose a minimum distance in terms of emission intensity between the lowest emitter in the group of treated firms and the highest emitter in the comparison group. While a minimum distance of 1.5 was imposed in the benchmark analysis, in this robustness check we first decrease it to zero, and then increase it to 3. Figure 14 plots the results obtained in the two cases, in the left and the right panel, respectively. The results

[25] The empirical implementation, characterized in more details in Appendix 6.1, uses the `matchit()` function from package `MatchIt` for `R` that automatically loads the `cem` package for R (Iacus et al., 2009).

remain very similar to the base case. Hence, we see again that the role of emission intensities is only marginal after the performed matching, as was also shown in Table 1.

Finally, to illustrate the adequacy of the performed evaluation under the null hypothesis of absent impact, we create a pseudo situation by using a random split of a joint pool of the previously treated HC (high emission) and comparison (low emission) firms.[26] The matching procedure now is applied to this pseudo split into treated and control firms.[27] Figure 15 presents a couple of typical realizations with different seeds of random number generator. They reveal that, indeed, there is no significant deviation between these artificially created 'treated' and 'comparison' groups.

Overall, despite all the alternative specifications, resulting also in a substantial variation of the number of matched firms, the general pattern remains quite consistent. Finally, Table 3 on the next page summarizes the overall significance of the results in all the robustness checks described above, including also information on the size of treated and control groups. In all cases, the overall PA impact is statistically significant at least at the 10% significance level.

5 Conclusion

The Paris Agreement and the US withdrawal affected significantly the participation of financial investors in European high-carbon companies. Holdings in such companies have decreased significantly relative to non-high-carbon firms since the PA went into force. This result is consistent with investors revising their expectations on HC firms in the direction of higher risk. However, the process reverted after the US announcement of withdrawal from the PA. This finding is consistent with the explanation that the announcement created uncertainty about the viability and credibility of the agreement.

These changes in participation have certain implications for the transfer of risks. On the one hand, the reduction in overall participation in HC companies by the holders

[26] Note that such a pseudo split of firms remains the same for all the periods under consideration. It is performed by generating random independent draws from the standard Gaussian distribution for each company. Firms with realized values below -0.25 are prescribed to 'controls', whereas those with above 0.25 are classified as 'treated'. Firms with values in between are dropped to get the number of the matched firms similar to that obtained in the base analysis.

[27] Other than in the original split by the level of emissions, the empirical distribution functions of the matched firms from these randomly formed groups are similar in terms of emissions.

Table 3: Overall significance of the PA impact (robustness checks)

Specification	Coeff.	S.E.	90% conf.bands		95% conf.bands		p-val.(Par.Tr.)	n.treat.	n.comp.	n	T
Holder-level estimation	-0.323**	0.121	-0.527	-0.120	-0.554	-0.093	1.000	59	69	128	23
Without UK issuers	-0.386**	0.154	-0.631	-0.141	-0.650	-0.122	0.998	19	17	36	23
Without UK holders and issuers	-0.439**	0.130	-0.647	-0.231	-0.694	-0.184	0.952	19	17	36	23
Genetic matching	-0.152*	0.086	-0.292	-0.011	-0.317	0.014	0.871	84	84	168	23
Nearest neighbor matching	-0.132*	0.074	-0.255	-0.009	-0.294	0.030	0.916	83	115	198	23
Matching on 2013-2015 averages	-0.325**	0.109	-0.502	-0.148	-0.521	-0.129	0.721	63	56	119	23
Matching on 2015 data	-0.213**	0.086	-0.352	-0.073	-0.395	-0.030	0.955	75	84	159	23
No constraint on relat. emissions	-0.274**	0.090	-0.422	-0.125	-0.436	-0.112	0.829	64	76	140	23
3 times higher relative emissions	-0.263**	0.101	-0.423	-0.103	-0.482	-0.045	0.609	53	62	115	23
Random draw 1	0.024	0.090	-0.128	0.175	-0.161	0.208	0.994	67	67	134	23
Random draw 2	-0.074	0.109	-0.251	0.104	-0.284	0.136	0.971	61	61	122	23

in our sample (i.e. covered in the SHS database) implies an increase in participation by the holders who are not in the sample, which are essentially non-EA financial investors. Indeed, based on the subset of holdings by non-EA investors we have in our dataset, we do see an increase in participation in European HC companies by investors located in the BRIC region, in particular. Moreover, we document a transfer of transition risk from more regulated financial institutions towards other financial institutions within Europe. We also find that investors are less willing or able to reduce their participation in those high-carbon firms where they hold large stakes.

As further research, it would be interesting to investigate the following aspects. First, our analysis focused only on equity holdings; however, loans and bonds may have been influenced differently. Second, the SHS aggregation at the level of the institutional sector does not allow to investigate whether different investors within the same sector have reacted differently, with their responses possibly averaging out at the aggregate level. The bank-level SHS module could be used to shed light on this particular aspect. Finally, extending the dataset to 2021 and beyond, i.e. covering the period with the US rejoining the PA and the recovery from the Covid-19, could help to discriminate between several possible explanations of why the aggregate impact of the PA becomes insignificant by 2020. A further monitoring of later global agreements, e.g. achieved in the Glasgow summit, is also worth pursuing in order to understand their perception by market participants.

Finally, our results have some relevant policy implications. First, global environmental policy has an impact on investors behavior in terms of portfolio allocation. Second, the successful redirection of global financial flows towards climate action (Article 2c of the PA) requires a clear and unanimous signal from the global community of policy makers. Third, as the low-carbon transition picks up speed, a close monitoring of the buildup of transition risk in particular sectors and jurisdictions is warranted.

References

Alessi, L., Ossola, E. and Panzica, R. (2021) When do investors go green? Evidence from a time-varying asset-pricing model. European Commission - Joint Research Centre, Working paper, forthcoming.

Athey, S., and Imbens, G. (2010) Identification and Inference in Nonlinear Difference-in-Differences Models. Econometrica 74(2), 431–497.

Austin, P.C. (2010) Statistical criteria for selecting the optimal number of untreated subjects matched to each treated subject when using many-to-one matching on the propensity score. American Journal of Epidemiology 172(9), 1092–1097.

Abadie, A., Alexis, D., and Jens, H. (2010) Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association 105(490), 493–505.

Baiardi, D., and Morana, C. (2021) Climate change awareness: Emprical evidence for the European Union. Energy Economics 96 (available online 11 February, 2021, at `https://doi.org/10.1016/j.eneco.2021.105163`).

Bartram, S.M., Hou, K., and Kim, S. (2021) Real Effects of Climate Policy: Financial Constraints and Spillovers. CEPR Discussion Papers 15986, C.E.P.R. Discussion Papers.

Callaway, B., and Sant'Anna, P.H.C. (2020) Difference-in-Differences with multiple time periods. Journal of Econometrics, forthcoming (available online 17 December, 2020, at `https://doi.org/10.1016/j.jeconom.2020.12.001`).

de Chaisemartin, C., and d'Haultfoeuille, X. (2020) Difference-in-Differences Estimators of Intertemporal Treatment Effects. Available at `http://dx.doi.org/10.2139/ssrn.3731856`.

de Chaisemartin, C., and d'Haultfoeuille, X. (2020) Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. American Economic Review 110(9), 2964–2996.

de Chaisemartin, C., d'Haultfoeuille, X., and Guyonvarch, Y. (2021) DID_MULTIPLEGT: Stata module to estimate sharp Difference-in-Difference designs with multiple groups and periods. Statistical Software Components S458643, Boston College Department of Economics, revised 13 Feb 2021. Available at `https://ideas.repec.org/c/boc/bocode/s458643.html`.

Dai, H.C., Zhang, H.B., and Wang, W.T. (2017) The impacts of U.S. withdrawal from the Paris Agreement on the carbon emission space and mitigation cost of China, EU, and Japan under the constraints of the global carbon emission space. Advances in Climate Change Research 8(4), 226–234.

Diamond, A., and Sekhon, J.S. (2013) Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. Review of Economics and Statistics 95(3), 932–945.

Ellis, J., and Moarif, S. (2017) Enhancing transparency of climate finance under the Paris Agreement: Lessons from experience. OECD/IEA Climate Change Expert Group Papers 2016/3, OECD Publishing (`http://www.oecd.org/environment/cc/Enhancing-transparency-climate%20finance-V2%20.pdf`).

Ho, D.E., Imai, K., King, G., and Stuart, E.A. (2011) `MatchIt`: Nonparametric Preprocessing for Parametric Causal Inference. Journal of Statistical Software 42(8).

Iacus, S.M., King, G., and Porro, G. (2009) `cem`: Software for Coarsened Exact Matching. Journal of Statistical Software 30(9).

Iacus, S.M., King, G., and Porro, G. (2011) `cem`: Multivariate Matching Methods That are Monotonic Imbalance Bounding. Journal of American Statistical Association 106(493),345–361.

Iacus, S.M., King, G., and Porro, G. (2012) Causal Inference Without Balance Checking: Coarsened Exact Matching. Political Analysis 20(1), 1–24.

Iacus, S.M., King, G., and Porro, G. (2019) `cem`: A Theory of Statistical Inference for Matching Methods in Causal Research. Political Analysis 27(1), 46–68.

Law, C., and Zhang, H. (2019) Implementing Provisions on Climate Finance Under the Paris Agreement. Climate Law 9(1-2), 21–39.

Mehling, M. (2021) "Article 9 (Finance)." In (eds. L.S. Reins and G. van Calster) A Commentary on the Paris Agreement on Climate Change. Cheltenham: Edward Elgar, forthcoming.

Network for Greening the Financial System (2019)

Ramelli, S., Ossola, E., and Rancan, M. (2020) Climate Sin Stocks: Stock Price Reactions to Global Climate Strikes. JRC Working Papers in Economics and Finance, 2020/03. Publications Office of the European Union, Luxembourg, 2020, doi:10.2760/422081, JRC120974.

Reghezza, A., Altunbas, Y., Rodriguez d'Acri, C., Marques-Ibanez, D, and Spaggiari , M. (2021) Do Banks Fuel Climate Change?. European Central Bank Working Paper 2550.

Reins, L., and van Calster, G. (2021) "Introduction – The Paris Agreement on Climate Change." In (eds. L.S. Reins and G. van Calster) A Commentary on the Paris Agreement on Climate Change. Cheltenham: Edward Elgar, forthcoming.

Sant'Anna, P.H.C., and Zhao, J. (2020) Doubly robust difference-in-differences estimators. Journal of Econometrics 219(1), 101–122.

Sekhon, J.S. (2011) Multivariate and Propensity Score Matching Software with Automated Balance Optimization: The Matching package for R. Journal of Statistical Software 42(1), 1–52.

Steinhauer, V. (2018) Leaving the Paris Agreement: The United States' Disengagement from the Global Climate Regime and its Impact on EU Climate Diplomacy. EU Diplomacy Paper 04/2018, available at `http://aei.pitt.edu/97353/`.

Stuart, E.A., and Rubin, D.B. (2008) Best Practices in Quasi–Experimental Designs: Matching Methods for Causal Inference. Ch. 11 in Best Practices in Quantitative Methods (ed. J. Osborne), pp. 155–176.

Xu, Y. (2017) Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models. Political Analysis 25(1), 57–76.

Zhang, H.B., Dai, H.C., Lai, H.X., and Wang, W.T. (2017) The withdrawal of the U.S. from the Paris Agreement and its impact on global climate change governance. Advances in Climate Change Research 8(4), 213–219.

Zhang, H.B., Dai, H.C., Lai, H.X., and Wang, W.T. (2017) U.S. withdrawal from the Paris Agreement: Reasons, impacts, and China's response. Advances in Climate Change Research 8(4), 220–225.

6 Appendices

6.1 Appendix A: Matching procedure

For the main analysis, we use the Coarsened Exact Matching (CEM) approach which leaves only subclasses containing treatment and control units that are exactly equal on the coarsened support of covariate values.[28] The CEM bounds the degree of model dependence and the treatment effect estimation error, eliminates the need for a separate procedure to restrict data to common empirical support, is robust to measurement error, etc. (see Iacus et al., 2011, 2012, 2019).

For the additional robustness checks we further employ the genetic matching algorithm—abbreviated as GEN1 in the sequel—with the generalized Mahalanobis distance which uses the genetic algorithm to determine the scaling factors for each covariate that minimize a criterion of covariate imbalance.[29]

Furthermore, we also included a greedy nearest neighbor matching (hereafter, abbreviated as NN2) with a propensity score estimated using logistic regression of the treatment on the covariates, allowing for up to two control units for a single treatment unit (see, e.g., Austin, 2010 for arguments to keep the ratio low and Stuart and Rubin, 2008, for a general discussion). A caliper of size 0.15 was applied both in the GEN1 and NN2 matching procedures with little changes when varying it between 0.1 and 0.2.

The matching is based on the pre-PA data on three covariates. As we look at financial investments, we first of all include the profitability (dividend yield) and riskiness (historical volatility) of stock returns. To further account for the size differences of firms, we also include the (logarithm of the) market value of firms among the matching covariates.[30] In the base analysis, the five-year average of pre-treatment data (2011-2015) of the covariates was employed. In the robustness checks, a three-year average (2013-2015) and a single pre-treatment year (2015) were also considered.

[28] We use the Sturge's rule for the coarsening (see, e.g., Iacus et al., 2009).

[29] Genetic matching was performed using the `MatchIt` package (Ho et al., 2011) in `R`, which calls functions from the `Matching` package (Diamond and Sekhon, 2013; Sekhon, 2011). In our case, the criterion is the p-value in covariate balance testing. We have also limited to a single control unit to be matched to a treated unit in this approach which yielded higher number of matched treated units.

[30] Quantitatively similar results remain including further a liquidity indicator (turnover by volume) with the implication of a shrinking number of matched firms. Given good quality of the matching, we do not include any additional firm characteristics, which would further reduce the size of matched samples.

For each covariate, Figures 6 and 7 illustrate the performance of the performed matching procedures in terms of the empirical Quantile-Quantile (eQQ) adequacy between the treated and comparison (control) units. Figure 6 plots additionally a simple covariate (im)balance evaluation in terms of the standardized mean difference in the treated and comparison groups (see the right panel). There is a sizable discrepancy between the distributions of treated and control units in the unmatched sample (All); it is especially large in terms of the company size indicator (log(MrktVal)).

In the CEM-matched case (see the right panel in Figure 6), the correspondence between the quantiles of empirical cumulative distribution functions in the treated and control groups is very good. In fact, it is seemingly better than that observed for the GE1-matched and NN2-matched cases (see the left and right panels in Figure 7, respectively). The CEM-based matching has not only a much smaller total multivariate imbalance but also a larger percentage of local common support (see Table 4). Therefore, despite somewhat smaller number of matched cases, we ground our base analysis on the CEM outcome. As part of the treated firms remain unmatched, the actual estimand under consideration is the feasible sample ATT.

Figure 6: Empirical quantile-quantile (eQQ) plots and covariate balance: Unmatched vs. CEM-matched

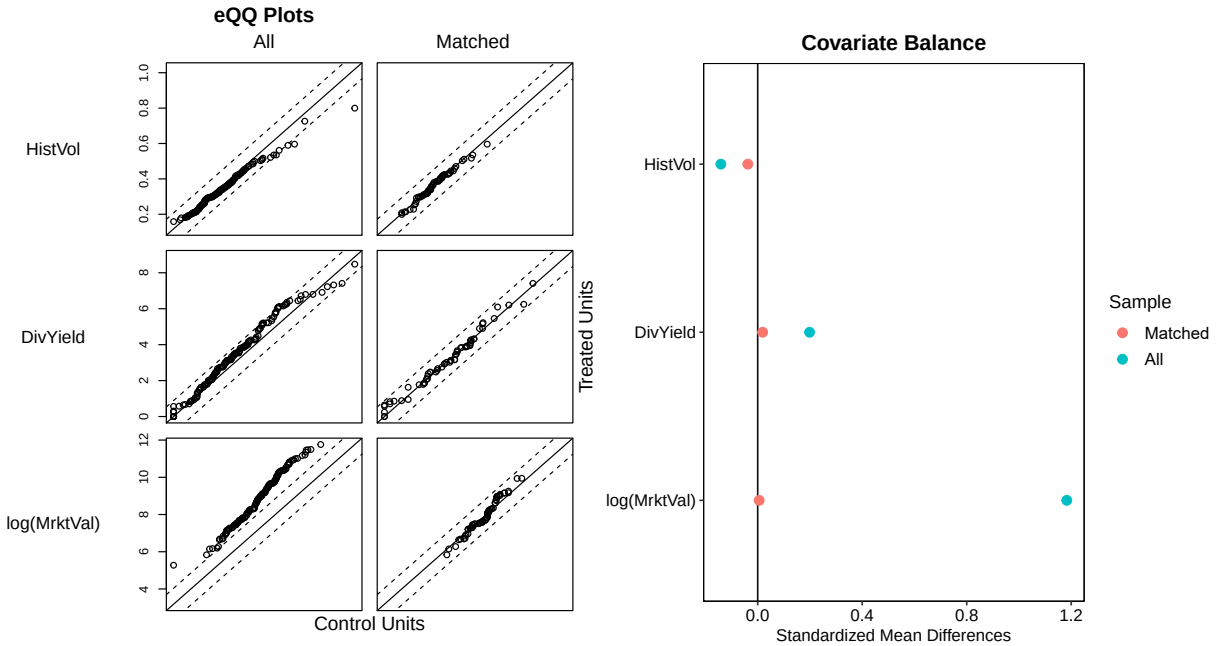

Figure 7: Empirical quantile-quantile (eQQ) plots: Unmatched vs. genetic matching (left panel) and nearest neighbors (right panel)

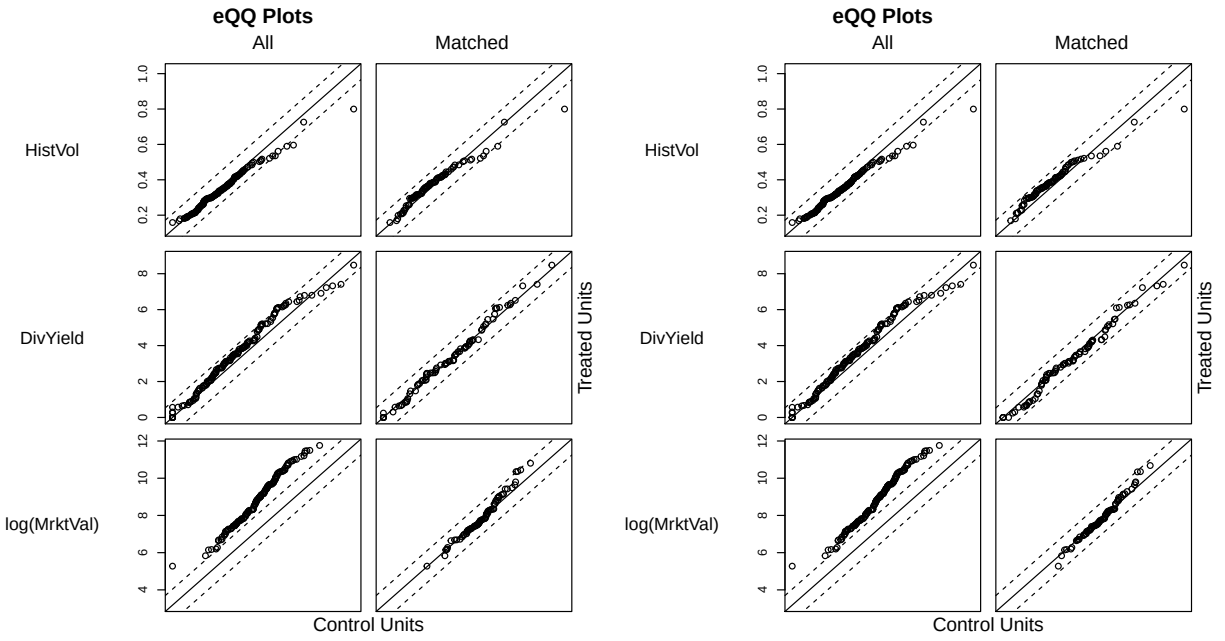

Table 4: Multivariate imbalance, local common support, and number of matched units

	All	CEM	GEN1	NN2
Multivariate Imbalance Measure:	0.56	0.38	0.48	0.49
Local common support (%):	27	51	43	38
Number of matched controls:		69	84	115
Number of matched treated:		59	84	83
Total number of controls:	152	152	152	152
Total number of treated:	164	164	164	164

The resulting distributions of treated and comparison firms by the broad NACE activity sectors are reported in Tables 5 and 6.

Finally, after performing the matching, we further drop firms having the overlapping or insufficiently distant relative emission levels—relative to sales and assets—in the comparison and the treated groups. In the base analysis, we require that the ratio between the minimum value observed in the treated group would be 1.5 times higher than the maximum observed in the comparison group. Further variations of this threshold are explored in the robustness checks considering the situations without any constraint and with the doubled requirement, i.e., 3 times separation. The matching is

Table 5: Activity sector of treated firms

NACE sector	Units
B - Mining and quarrying	8
C - Manufacturing	22
D - Electricity, gas, steam and air conditioning supply	9
E - Water supply; sewerage, waste management and remediation activities	2
F - Construction	7
G - Wholesale and retail trade; repair of motor vehicles and motorcycles	2
H - Transportation and storage	8
K - Financial and insurance activities	1

Table 6: Activity sector of comparison firms

NACE sector	Units
C - Manufacturing	14
E - Water supply; sewerage, waste management and remediation activities	1
F - Construction	5
I - Accommodation and food service activities	1
J - Information and communication	12
K - Financial and insurance activities	23
L - Real estate activities	3
M - Professional, scientific and technical activities	6
N - Administrative and support service activities	2
R - Arts, entertainment and recreation	1
S - Other service activities	1

repeated again, in order that such a removal of some units would not bias the weights. The resulting difference of the distribution between the relative emissions to sales and assets are illustrated in Figure 8 that plots the respective empirical cumulative distribution functions (top and bottom panel, respectively) in the groups of matched treated (high-carbon) and control (comparison) firms. For a better visibility of the difference between the minimum level in the treated group and the maximum level in the comparison group (marked by the vertical brown and grey lines, respectively), the support is cut at 500 in the figures on the right side that, otherwise, present the same information.

Figure 8: Relative emissions to assets and sales in the matched groups of treated (high-carbon) and comparison firms

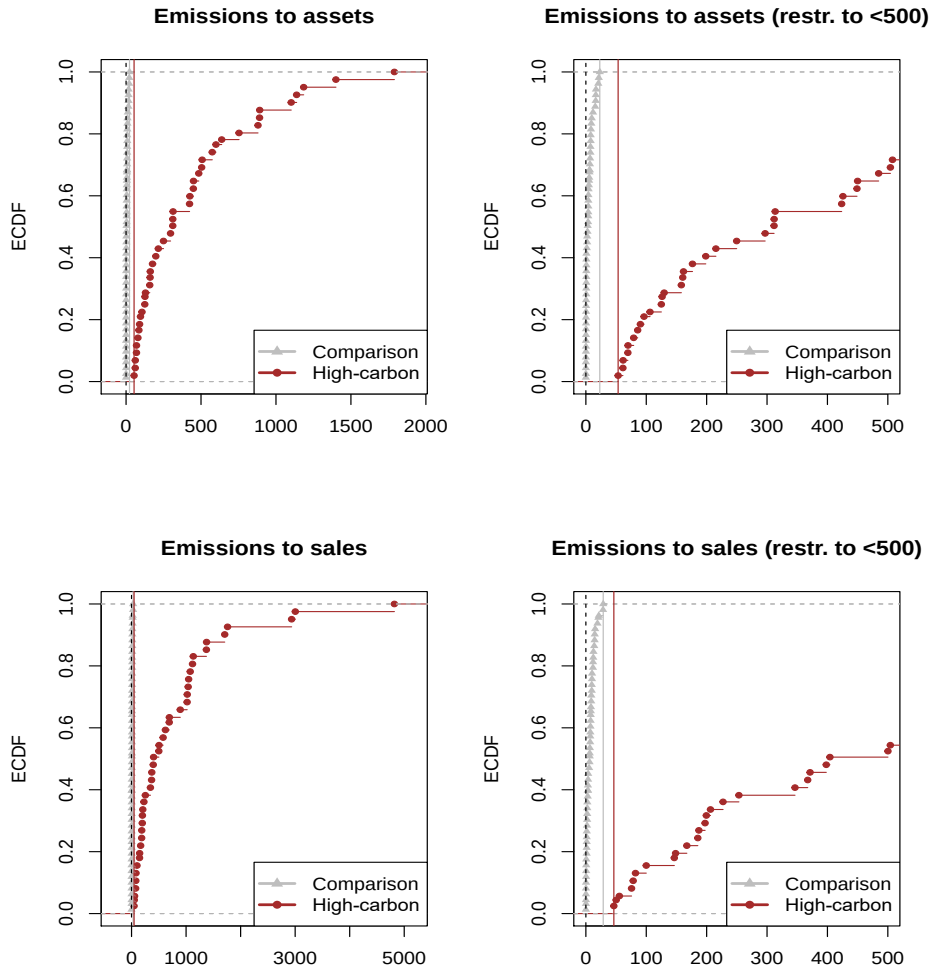

6.2 Appendix B: Estimators of the period-specific and overall ATTs

We separate between the two types of main results discussed in Section 3. First, there are dynamic effects established based on the estimates of the period-specific ATTs that vary over time, e.g., as presented in Figure 2. Second, there is an overall ATT estimate reported in Tables 2 and 3 that characterizes the effect during the whole post-treatment period. Next, we briefly present each of these estimators.

Our main results that provide the multi-period ATTs rely on the doubly-robust estimator of Callaway and Sant'Anna (2020) defined in their eq. (4.1) which identifies the period-specific ATTs from comparison with the never-treated group that, in our case, consists of not-HC firms. Furthermore, given that in our study there is a single treatment date (t_0) and no anticipation ($\delta = 0$), their estimator reduces, in our case, to

$$\widehat{ATT}(t) := \widehat{ATT}_{dr}^{nev}(t) = \mathbb{E}_n \left[\left(\frac{D}{\mathbb{E}_n[D]} - \frac{\frac{\hat{p}(X;\hat{\pi})(1-D)}{1-\hat{p}(X;\hat{\pi})}}{\mathbb{E}_n\left[\frac{\hat{p}(X;\hat{\pi})(1-D)}{1-\hat{p}(X;\hat{\pi})}\right]} \right) \left(y_t - y_{t_0-1} - \hat{m}_t(X;\hat{\beta}_t) \right) \right],$$

where $\mathbb{E}_n[Z] = \frac{1}{n}\sum_{i=1}^{n} Z_i$ for some generic Z, D is a binary variable that equals to one for treated units, whereas $\hat{p}(X;\hat{\pi})$ and $\hat{m}_t(X;\hat{\beta}_t)$ are parametric estimators of the propensity score $p(X;\pi)$, which defines the probability of being treated conditional on pre-treatment covariates X in a parametric (logistic) regression with its vector of parameters π, and the linear population outcome regression of the never-treated group conditional on pre-treatment covariates X with the respective parameter vector β_t.

Given these period-specific ATTs, we further apply the overall ATT estimator defined by Callaway and Sant'Anna (2020) in their eq. (3.11) that, in our case with a single group, coincides with their eq. (3.7) yielding a simple average of the previously described period-specific ATTs:

$$\hat{\theta}^{Overall} = \frac{1}{T - t_0 + 1} \sum_{t=t_0}^{T} ATT(t).$$

6.3 Appendix D: Robustness plots

Figure 9: Alternative estimators

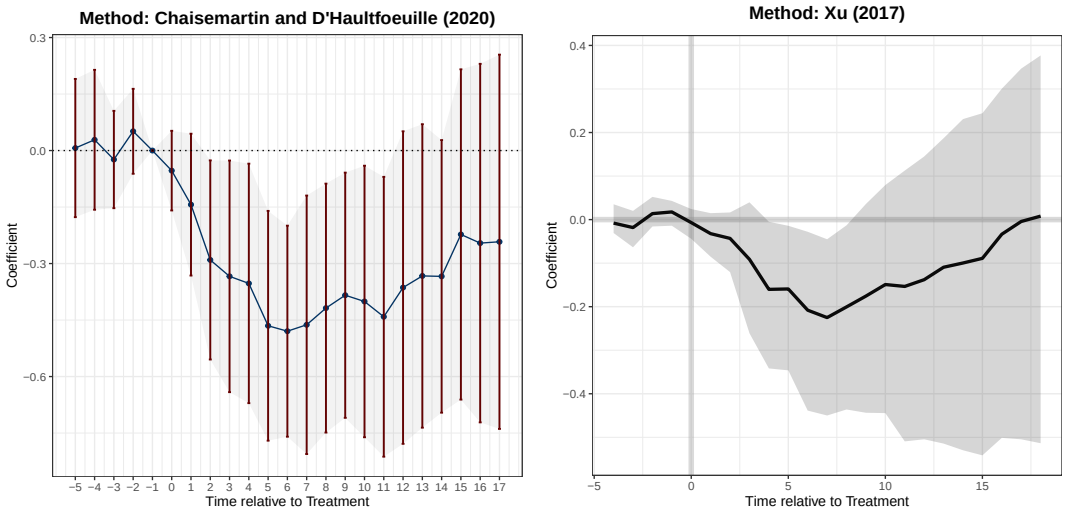

Note: To underscore the different methods, we keep the style of figures in correspondence with the respective packages: the `did_multiplegt` command for Stata and the `gsynth` package for `R`. Note that in the latter, the normalization is with respect to the first observation of the non-treatment period, whereas in former – with respect to the last observations of the non-treatment period (a period just before the treatment).

Figure 10: Estimated ATTs with holder-level data

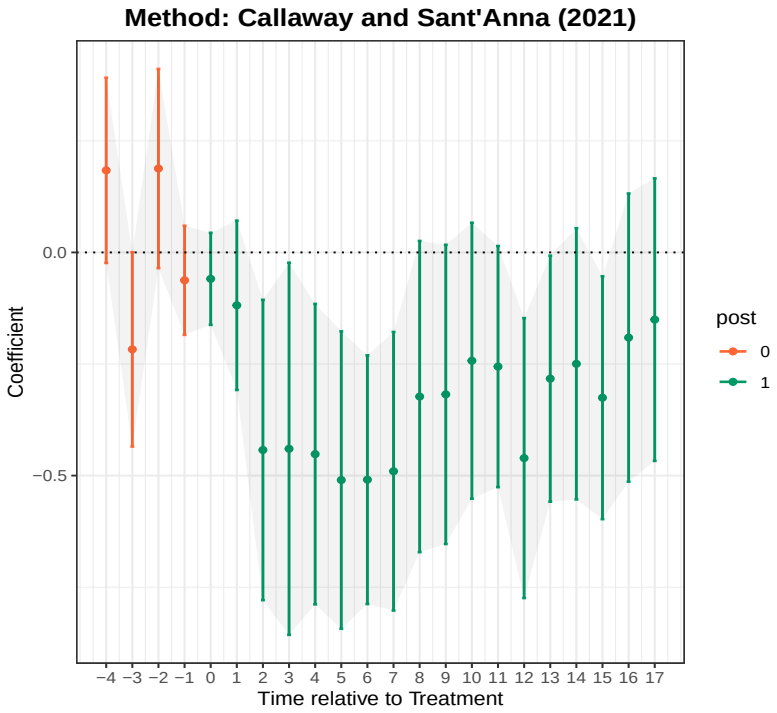

Figure 11: Estimated ATTs without the UK

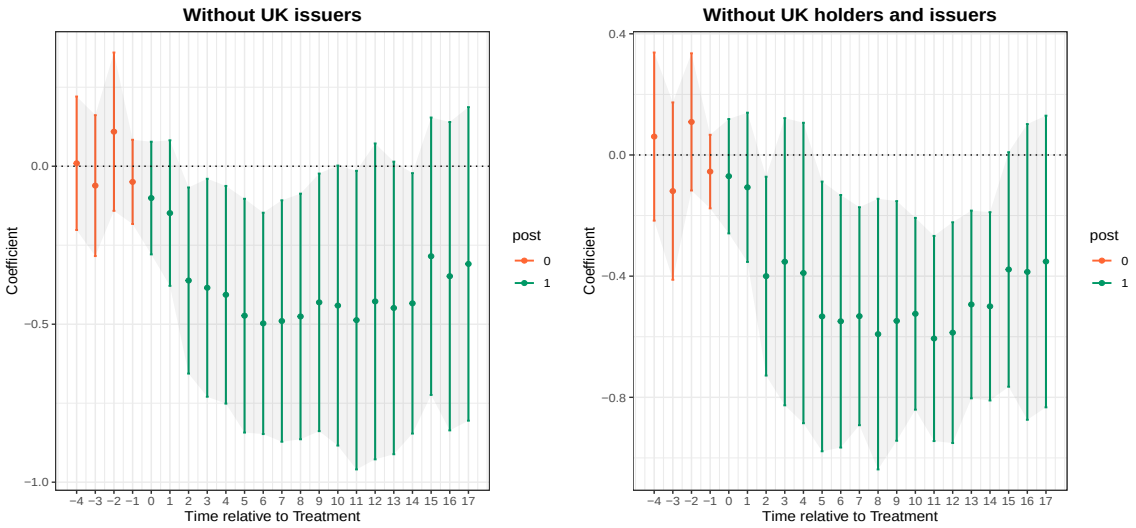

Figure 12: Estimated ATTs with alternative matching methods

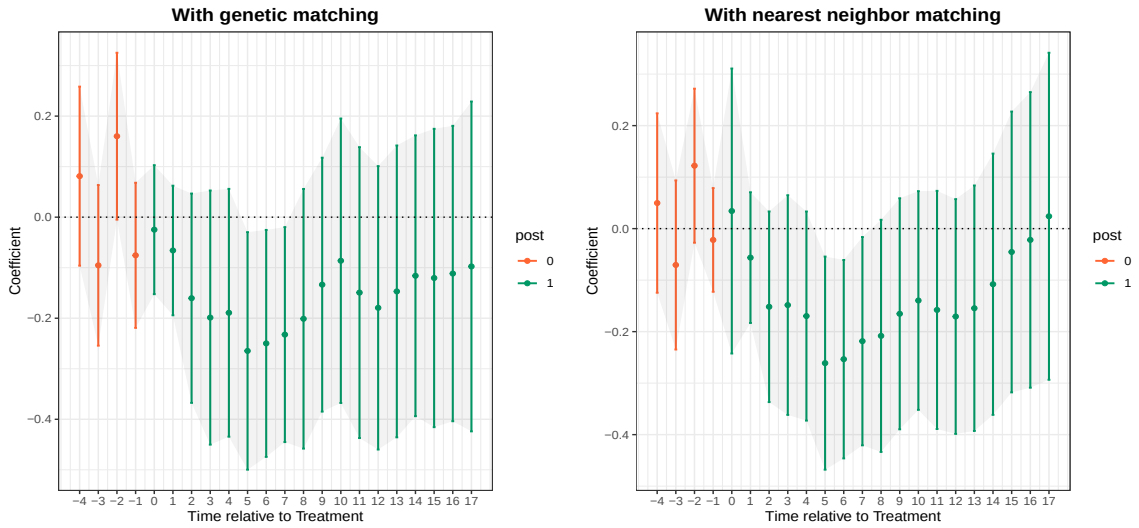

Figure 13: Estimated ATTs with alternative periods used for matching

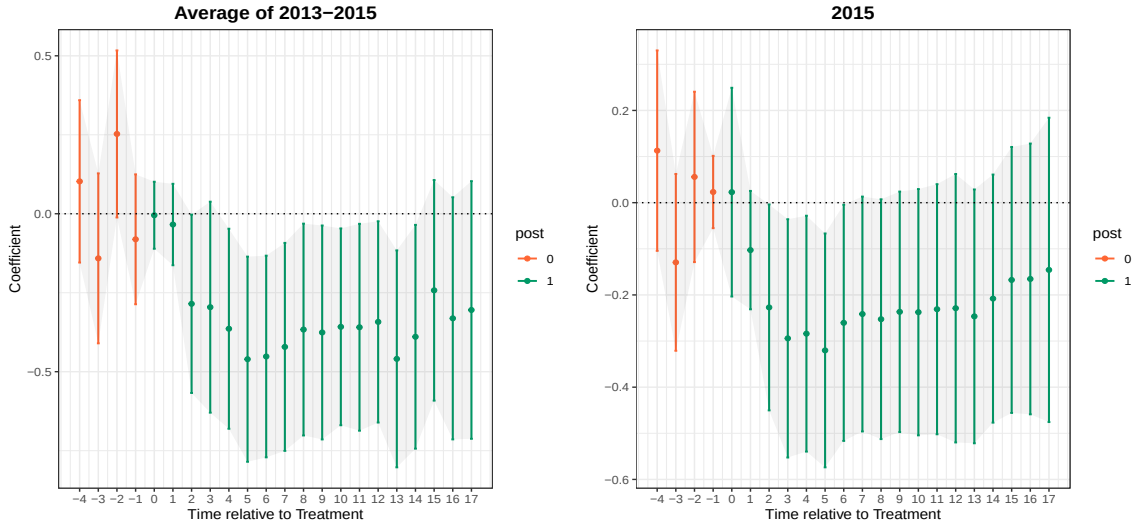

Figure 14: Estimated ATTs with alternative minimum distance of ratios of emissions to assets and sales in the matched treated and comparison groups

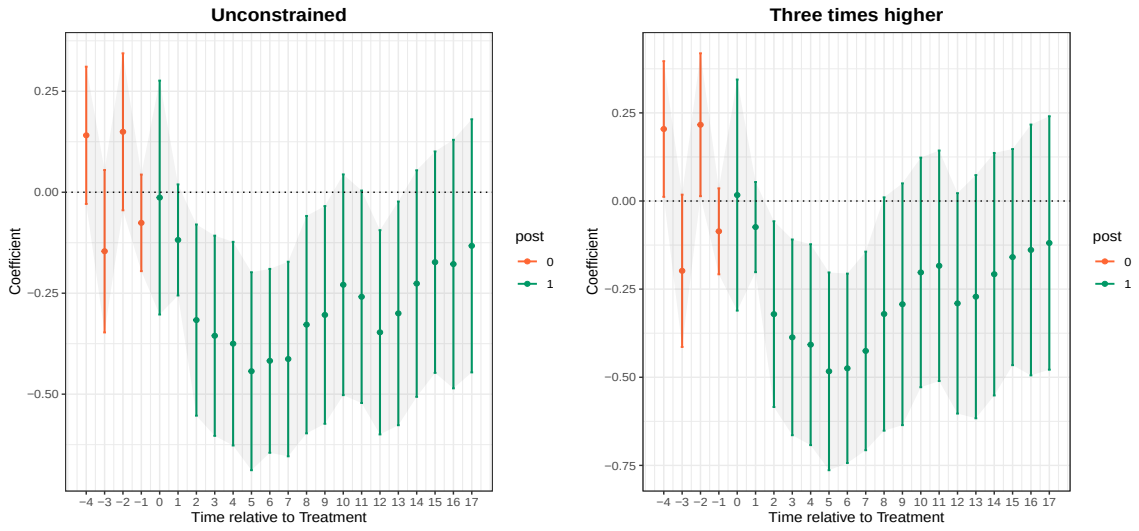

Figure 15: Estimated ATTs with a random split of firms

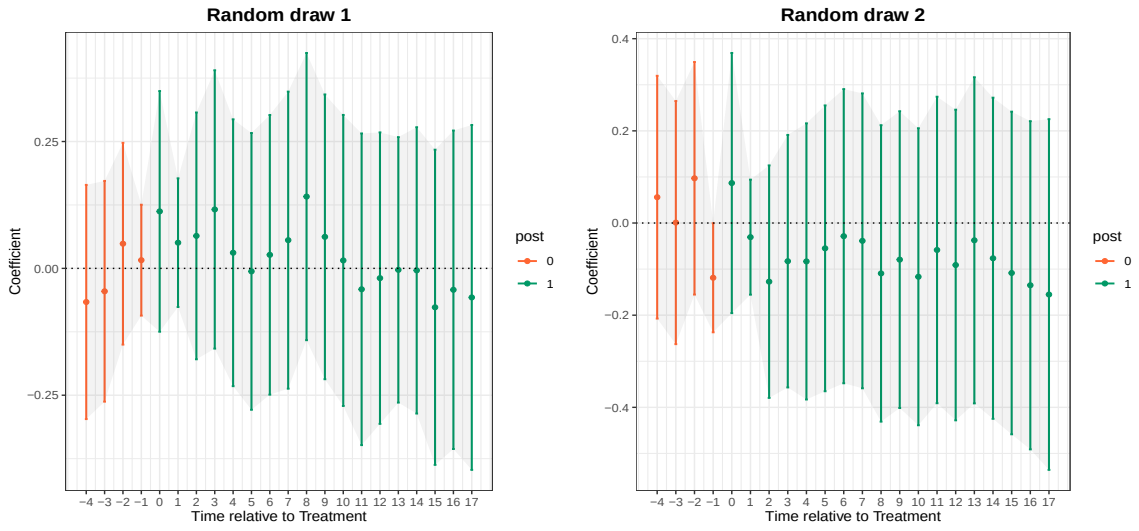

GETTING IN TOUCH WITH THE EU

In person

All over the European Union there are hundreds of Europe Direct information centres. You can find the address of the centre nearest you at: https://europa.eu/european-union/contact_en

On the phone or by email

Europe Direct is a service that answers your questions about the European Union. You can contact this service:

- by freephone: 00 800 6 7 8 9 10 11 (certain operators may charge for these calls),

- at the following standard number: +32 22999696, or

- by electronic mail via: https://europa.eu/european-union/contact_en

FINDING INFORMATION ABOUT THE EU

Online

Information about the European Union in all the official languages of the EU is available on the Europa website at: https://europa.eu/european-union/index_en

EU publications

You can download or order free and priced EU publications from EU Bookshop at: https://publications.europa.eu/en/publications. Multiple copies of free publications may be obtained by contacting Europe Direct or your local information centre (see https://europa.eu/european-union/contact_en).

The European Commission's science and knowledge service

Joint Research Centre

JRC Mission

As the science and knowledge service of the European Commission, the Joint Research Centre's mission is to support EU policies with independent evidence throughout the whole policy cycle.

EU Science Hub

ec.europa.eu/jrc

@EU_ScienceHub

EU Science Hub - Joint Research Centre

EU Science, Research and Innovation

EU Science Hub