Journal of Financial Stability 71 (2024) 101232

Contents lists available at ScienceDirect

Journal of Financial Stability

journal homepage: www.elsevier.com/locate/jfstabil

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

Lucia Alessi [a,*], Stefano Battiston [b,c], Virmantas Kvedaras [a]

[a] European Commission, Joint Research Centre, Ispra (VA), Italy[1]
[b] University of Zurich, Zurich, Switzerland
[c] University Ca' Foscari of Venice, Venice, Italy

ARTICLE INFO

Keywords:
Sustainable finance
Climate finance
Transition risk
Paris agreement
Securities holdings

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 participation in the equities of high-carbon firms in response to the agreement, and the trend reverted after the US's announcement of withdrawal from the agreement. 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.

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)[2]-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 focus on investors' *participation* in firms, defined as share of each firm's market capitalization they hold in their portfolio of equities. We find evidence that investors have reduced their participation in high-carbon 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.

This work is relevant for financial stability because the exposure of the financial sector to high-carbon firms has been recognized by

☆ The peer-review for the present article was conducted outside of Elsevier's submission system. A record of this review was kept, and has been reviewed by the Publisher to ensure conformity with Elsevier's and the journal's standards for quality and ethics. Dates of submission and acceptance reflect the dates on which the final versions were uploaded to our submission system. The timeline for peer-review is available upon request. The Journal has now changed its policy to ensure that all peer review is conducted using Elsevier's Editorial Manager submission system.

* Corresponding author.
E-mail address: lucia.alessi@ec.europa.eu (L. Alessi).

[1] The opinions expressed are those of the authors only and do not necessarily represent the European Commission's official position. The authors are grateful to two anonymous referees for their valuable comments.

[2] Including also the United Kingdom (UK), as we consider the period before the Brexit took place.

https://doi.org/10.1016/j.jfs.2024.101232
Received 15 February 2022; Received in revised form 1 August 2022; Accepted 17 November 2022
Available online 7 February 2024

financial supervisors as a source of *climate transition risk*, under a scenario of so-called *disorderly transition* (NGFS, 2019).

Our finding suggest that climate policies do have an effect on the level of holdings of financial institutions in high-carbon firms, but that lack of coherence in policies announcement lead financial actor to revert to their financial positions. Moreover, we also document a potential buildup of transition risk in less regulated industries of the financial sectors and in particular jurisdictions.

The PA itself was a long process starting from the adoption by the UNFCCC[3] 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, 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 van 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., 2017b,a). The formal notice of intention to withdraw was given on November 4, 2019, abiding to Article 28 of the PA.[4]

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 high-carbon 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 high-carbon firms and possibly, in turn, their profitability.

For these reasons, we expect, *a priori*, that investors may have actually reduced their stakes in high-carbon 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.

Several recent works look at the impact of climate policies on prices and/or risk premia, focusing either on stocks, or bonds, or loans (Monasterolo and de Angelis, 2020; Ramelli et al., 2020; Alessi et al., 2021a, 2023; Allevi et al., 2019; Fatica and Panzica, 2021; Beyene et al., 2021; Bolton and Kacperczyk, 2021). Our work complement those studies by looking instead at the impact on the allocation of holdings of securities. A few works have looked at the impact of climate policies, and the Paris Agreement (PA) in particular, on financial holdings: Boermans and Galema (2019) focus on stock holdings of Dutch pensions funds, while Reghezza et al. (2021) focus on European banks; Boermans and Galema (2020) analyze carbon home bias; herding behavior on decarbonization of global holdings of stocks has been investigated by Benz et al. (2020). Our work focuses on the impact of PA on holdings of European stocks by different types of investors. Since we focus on holdings of stocks, our results do not rule out that financial actors may have adjusted their portfolios of bonds and loans in a different direction from what we find for stocks. However, our paper offers some clear findings on the holdings of stocks. The result happens to be in line the results on loans by Reghezza et al. (2021).

To measure investors' stakes in high-carbon 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. The first set, i.e the 'treatment' group, consists of EU firms that are expected 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 B), and are dubbed hereafter as 'high-carbon' (HC). The second set of firms, i.e. the 'control' group, comprises firms with low levels of emissions, which are expected to be little affected by environmental policies, as they typically are active in sectors of the economy that have a lower 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 (2021) 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[5] 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 can 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 (2021) study the changes in the perceptions of the importance of climate change. In line 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. (2021a) 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. Along the same lines, by looking at equity holding indexes from the US, Europe and global financial markets, Monasterolo and de Angelis (2020) find that while low-carbon assets were perceived as riskier than the market before the Paris Agreement announcement, the level of systematic risk associated to them significantly decreased afterwards.

[3] Acronym for United Nations Framework Convention on Climate Change.

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

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

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.

Our explanation of the change in investors' behavior is also linked to the literature on uncertainty about economic/environmental policy and its implications for financial investments. In particular, we link the investors' reaction to the US withdrawal to the increased uncertainty about the viability and continuity of the policies agreed in the PA. As stressed by Pindyck (2007), for environmental problems uncertainties are greater and more crucial than for most other private and public policy decisions because the impact of environmental policies is highly nonlinear, involves irreversibilities and much longer time horizons. Higher uncertainty over policy costs and benefits of environmental policies underlines the importance of their credibility in order for political communication to move financial markets in the intended direction (see, e.g., Ferrara and Sattler, 2018; Battiston et al., 2021).

The literature on climate-related financial risks distinguishes to source of risk. Physical risk refers to losses on financial assets as a result of climate-related hazards. Transition risk refers to losses in financial investments arising from (partially unanticipated) changes in value of assets related to economic activities affected by climate policies, especially in the context of a disorderly transition. In terms of financial stability implications, our findings lends support to the relevance of climate stress-testing models (see e.g. Battiston et al., 2017; Roncoroni et al., 2021), aimed at investigating the ability of the financial system to withstand severe but plausible climate-related losses. With respect to transition risk, losses can be modeled as arising from changes in investors' preferences and expectations towards high-carbon assets (see e.g. Battiston et al., 2021; Dunz et al., 2021; Alessi et al., 2022). We contribute to this stream of literature, investigating the impact of global climate change policies on financial market participants, with specific attention to the heterogeneity of the reactions across different types of investors. In particular, we find that climate policies can indeed induce shifts in investors' holding strategies even at the aggregate level.

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 C, 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).[6] 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. In fact, SHS data is quite representative of the EU as a whole, as in addition, most non-euro area EU countries (namely Bulgaria, the Czech Republic, Denmark, Hungary, Poland and Romania) also collect this data.[7] We focus on stakes into companies that are located in the EU, in the period between 2015Q1 and 2020Q3.[8] 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. Complementary firms' NACE codes (4-digits) and GHG emission levels are obtained from Bloomberg at the consolidate account level of the firms. We limit the analysis to GHG Scope 1 emissions (i.e. direct emissions from production).[9] We are aware that for some companies, e.g. in the extractive sector, Scope 3 (i.e. indirect emissions, except Scope 2) can be much larger than their Scope 1. However, Scope 3 data suffers from lack of comparability across firms. We also did not consider Scope 2 (i.e. indirect emissions associated with the energy purchased by the firms for production) in this study because for large emitters it tends to be smaller than Scope 1. In particular, we use the most populated indicator, which is total greenhouse gases (GHG) emissions in carbon dioxide equivalent (CO2e), if available, or total carbon dioxide (CO2) otherwise in thousands of metric tons (Total GHG/CO2 Emissions).

Refinitiv 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.[10]

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.[11] 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. Furthermore, it does not depend directly on the variation in prices or quantities of other stocks, while this would become a problem if shares of (weights in) investment portfolio were under consideration instead.[12]

Formally, the (log) participation of holder sector h, in terms of ESA2010 classification (in short "holder" hereafter) into company j at time t is calculated as follows:

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

where $H_{h,j,t}$ and $M_{j,t}$ stand for the market value of holdings of holder sector h into company j and the total market value of company j, respectively, in period t.[13]

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

[7] See Appendix A discussing the dataset in more detail.

[8] 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.

[9] https://ghgprotocol.org/corporate-standard.

[10] We did not go beyond this set of matching variables because the achieved matching quality seems to be very good (see Appendix B), while additional matching dimensions would shrink further the number of matched cases. Moreover, firm-specific variables reflecting the firm's funding strategy, such as indebtedness, are much less populated.

[11] 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.

[12] For instance, portfolio shares are susceptible to underlying changes in the supply of securities; e.g., an increase in the number of technology-intensive companies would imply, ceteris paribus, a reduction in the share of any other (also high-carbon) firms in portfolios containing both types of stocks. Our participation metric is fully stock-specific and does not suffer from this shortcoming.

[13] Security-by-security data are aggregated by issuer and holder before taking the logarithmic transformation. We use several different aggregations in terms of sectors (e.g. across countries and/or across industries). Note that as the original data are unconsolidated, there are cases in which the total sum of holdings reported in the SHS, $\sum_h H_{h,j,t}$, is larger than the firm's market capitalization, $M_{j,t}$. The discrepancy affects a small fraction of firms and the

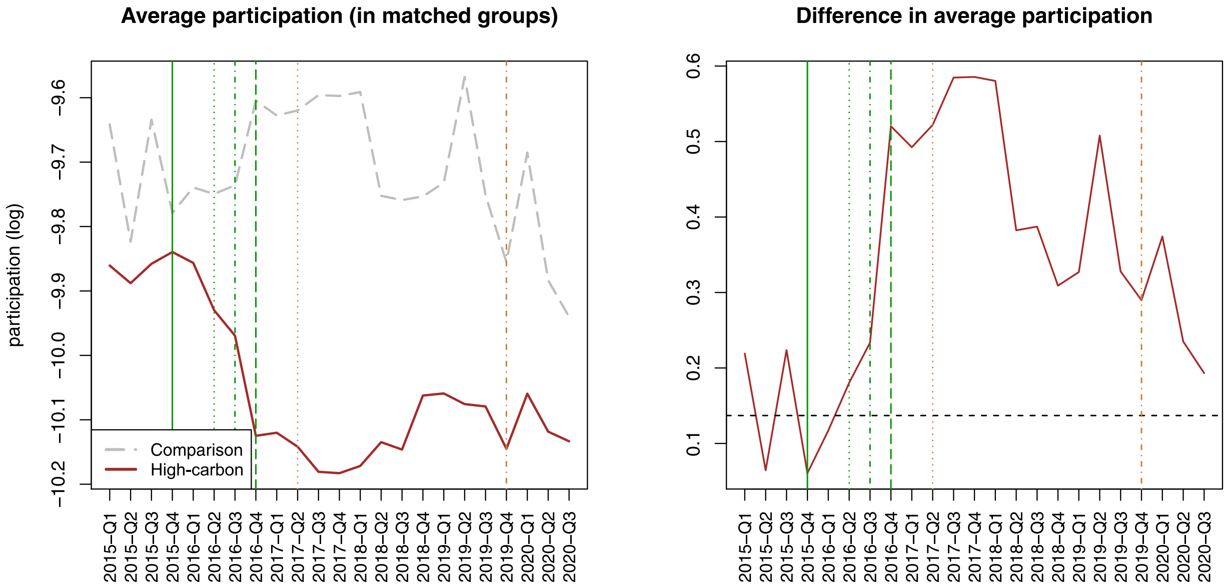

Fig. 1. Dynamics of average participation in the matched groups of treated (high-carbon) and control firms and difference between the two groups. The left panel plots the weighted average of (log-)participation, as defined in Eq. (1), in the matched set of companies from the high-carbon (treated) and low-carbon unaffected (comparison) companies, correspondingly. Vertical lines signify different stages of the Paris Agreement process. The right panel plots the respective difference between the two lines presented in the left panel, i.e., comparison minus high-carbon. A dashed horizontal line in the right panel denotes the average difference observed before the Paris Agreement.

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 aggregated along a particular dimension, e.g., type, country, or investment size.

Note that the indicator defined above captures the allocation of aggregate holdings of selected holders (e.g. the bank sector) in HC companies, relative to each company's market capitalization (and not in terms of total amount). It is both computed at a more aggregate level (i.e. for all EU holders) and by type and geography of holders, as recorded in the SHS database. However, since the total amount of shares in each firm is given, if the indicator of a specific holder decreases, this implies that, on average, investors in that holder sector have decreased their participation while investor in other holder sectors must have increased their holdings. When the decrease affects EU holders overall, then participation must have increase for holders outside the EU.

To motivate the estimations that will be presented later on, the left panel of Fig. 1 plots the dynamics of average participation indicator in the two matched sets of treated 'high-carbon' (HC) 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) high-carbon activities such as those in the fossil value chain, cement, air transportation, manufacturing of motor vehicles, etc. Firms are matched based on company size (market value), as well as on yield and volatility of their stock returns (see Appendix B for details).

ratio of the two quantities remains mostly below 1.15. We address this issue by excluding the cases with values of the ratio in the top 1% and by rescaling the holdings proportionally by using the correction factor $s_{h,j,t} = \frac{M_{j,t}}{\sum_h H_{h,j,t}}$. As there is a tendency for the number of such cases to increase with the size of market capitalization, their percentage in our matched sample from large companies is higher Instead of using the above correction, one can also exclude all the instances with a ratio larger than one, and work with a smaller number of matched firms. More details are reported in Section 4, item "Results without participation correction".

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 becoming 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, 2021) 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 Fig. 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 holders in the sample (which still leaves out many other holders, as discussed earlier) and EU-based 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 modeling 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 can be expected to be affected by the PA, while in the control group we include firms in the bottom tercile.[14] We further exclude from the control group firms whose main activity falls directly or indirectly in HC activities such as the fossil value chain, electricity, steel and cement, air transportation and motor vehicles manufacturing. We do so for two reasons. First, some firms active operating in these sectors might have low direct emission levels and yet belong to high-emissions value-chains, thus be negatively affected. Second, within each of the above sectors, comparatively lower-emission firms could be positively affected by the PA. This leaves in the control group firms operating in sectors such as health, education, etc., whose relevance for climate change mitigation is comparatively very limited. Still, we acknowledge that firms in these sectors might have indirectly positively benefited from the PA, as investors shifting away from HC industries might have diverted their investments to these industries. While this is a theoretical possibility and a potential limitation to our study, in the following section we show that the change and variation in the control group relative to the treated HC group was relatively small during and after the PA.

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 neighbor matching (see Appendix B 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 herafter, see Callaway and Sant'Anna, 2021; de Chaisemartin and d'Haultfoeuille, 2020b, 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 $D_{j,t} = \mathbb{1}\{j \in \mathbb{T}\} \cdot \mathbb{1}\{t \geq 2016Q2\}$ 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 possible outcomes with and without the treatment, with the actual outcome being $y_{j,t} = Y_{j,t}(D_{j,t})$, depending on the treatment state. The ATT is then defined as follows:

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

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.

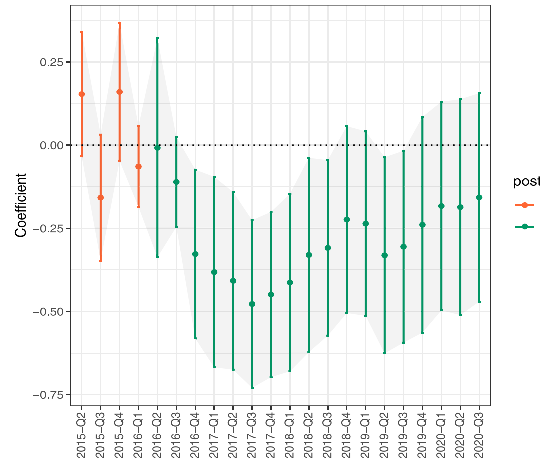

Fig. 2. The estimated period-specific ATTs. The figure plots the estimated average treatment effect on the treated in terms of the average (log-) participation, as defined in Eq. (2), with the ninety percent bootstrap confidence bands. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively). Notice that the normalization of point estimates is with respect to the first observation of the non-treatment period. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

For the main analysis we employ the doubly-robust ATT estimator proposed by Callaway and Sant'Anna (2021), applying it to the properly weighted matched set of treated and control firms (see Appendix B) and focusing on the case where all treated firms are treated at the same time. Fig. 2 plots the corresponding estimated ATTs with their 90% bootstrap-based confidence bands.[15]

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.[16] Finally, there is a clear increase in the

[14] Similar results appear also using the top quartile, but this shrinks substantially the number of matched firms. Analogous dynamics appears also using the half split, however the significance of the impact becomes weaker.

[15] The bootstrap-based inference is used with clusters/blocks at the issuer level and 1000 replications. The implementation relies on the `att_gt()` function of package `did` for R (see https://CRAN.R-project.org/package=did).

[16] It should be pointed out that a potential presence of interaction and spillover effects might bias the ATT estimation (see, e.g., Berg et al., 2021).

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 could 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 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 analyzed period.[17] However, the effect might have vanished also in connection with other reasons. First, the emergence of immediate risks related to the onset of the Covid-19 pandemic 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 implementation and achievements. Third, EU policies announced and implemented, with a particular reference to the EU green taxonomy, have clarified that even some high-carbon activities can be called green, if they have an enabling role or use the most efficient available technologies in terms of emissions (in several manufacturing sectors, green activities refer to the top 5% by emission performance).[18] 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 a more articulated definition of 'green activity' was not available —thus diminishing the relevance of the indicator used here.

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 investors' 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 Eq. (2) but calculated by considering only the holders belonging to a given type of holder sector or geography. Whereas it remains the same as in the aggregate analysis in the last Sections 3.2.3 and 3.2.4.

However, the dynamic path—i.e., the participation reduction after the PA and its reversal after the US withdrawal—would still remain present even if the estimated impact level was biased. Nevertheless, we intentionally cut off some potential spillover channels. for example, we exclude 'green energy' sector companies from the controls, as they could have otherwise absorbed part of the financing diverted from HC companies.

[17] 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.

[18] 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.

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 Fig. 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).

The following observations can be drawn. First, not all holder sectors reduce their relative participation in HC firms. In particular, the sector called "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 Fig. 3). This can be interpreted as a transfer of transition risk from more regulated financial institutions (banks, insurance firms, investment and pension funds) towards less regulated financial institutions. As a possible explanation, these investors, being typically active in the trading of derivatives and the intermediation to foreign acquisitions, and thus mostly driven either by speculative investments or by increased demand by non-EU investors, could be more willing to acquire stakes in HC companies (see next section).

Second, looking at regulated financial institutions, insurance corporations and investment funds[19] 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. This behavior of households could be 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.

Finally, it is interesting to notice that the results for the household sector are only marginally significant (see also Table 2), while those for NFCs are not significant.[20] This evidence seems to suggest that the 'sophistication' of investors could play a role in shaping their reaction. Indeed, professional investors in the financial sector seem to have reallocated holdings more (and more steadily) in anticipation of the impact of transition risk, whereas more 'naive' investors such as households and NFCs did so to a lesser extent or not at all.

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, our sample covers all the Euro Area (EA) countries and most of the EU but non-euro area (non-EA) countries (the latter though on a best-effort reporting basis), including a total of twenty three countries. We have information on non-EA investors only through EA custodians, which are mandated to

[19] Excluding Money Market Funds (MMF).

[20] As holdings by NFCs are the fourth largest in value, and much larger than those of households, we still report them separately in Fig. 10 in Appendix D despite insignificance.

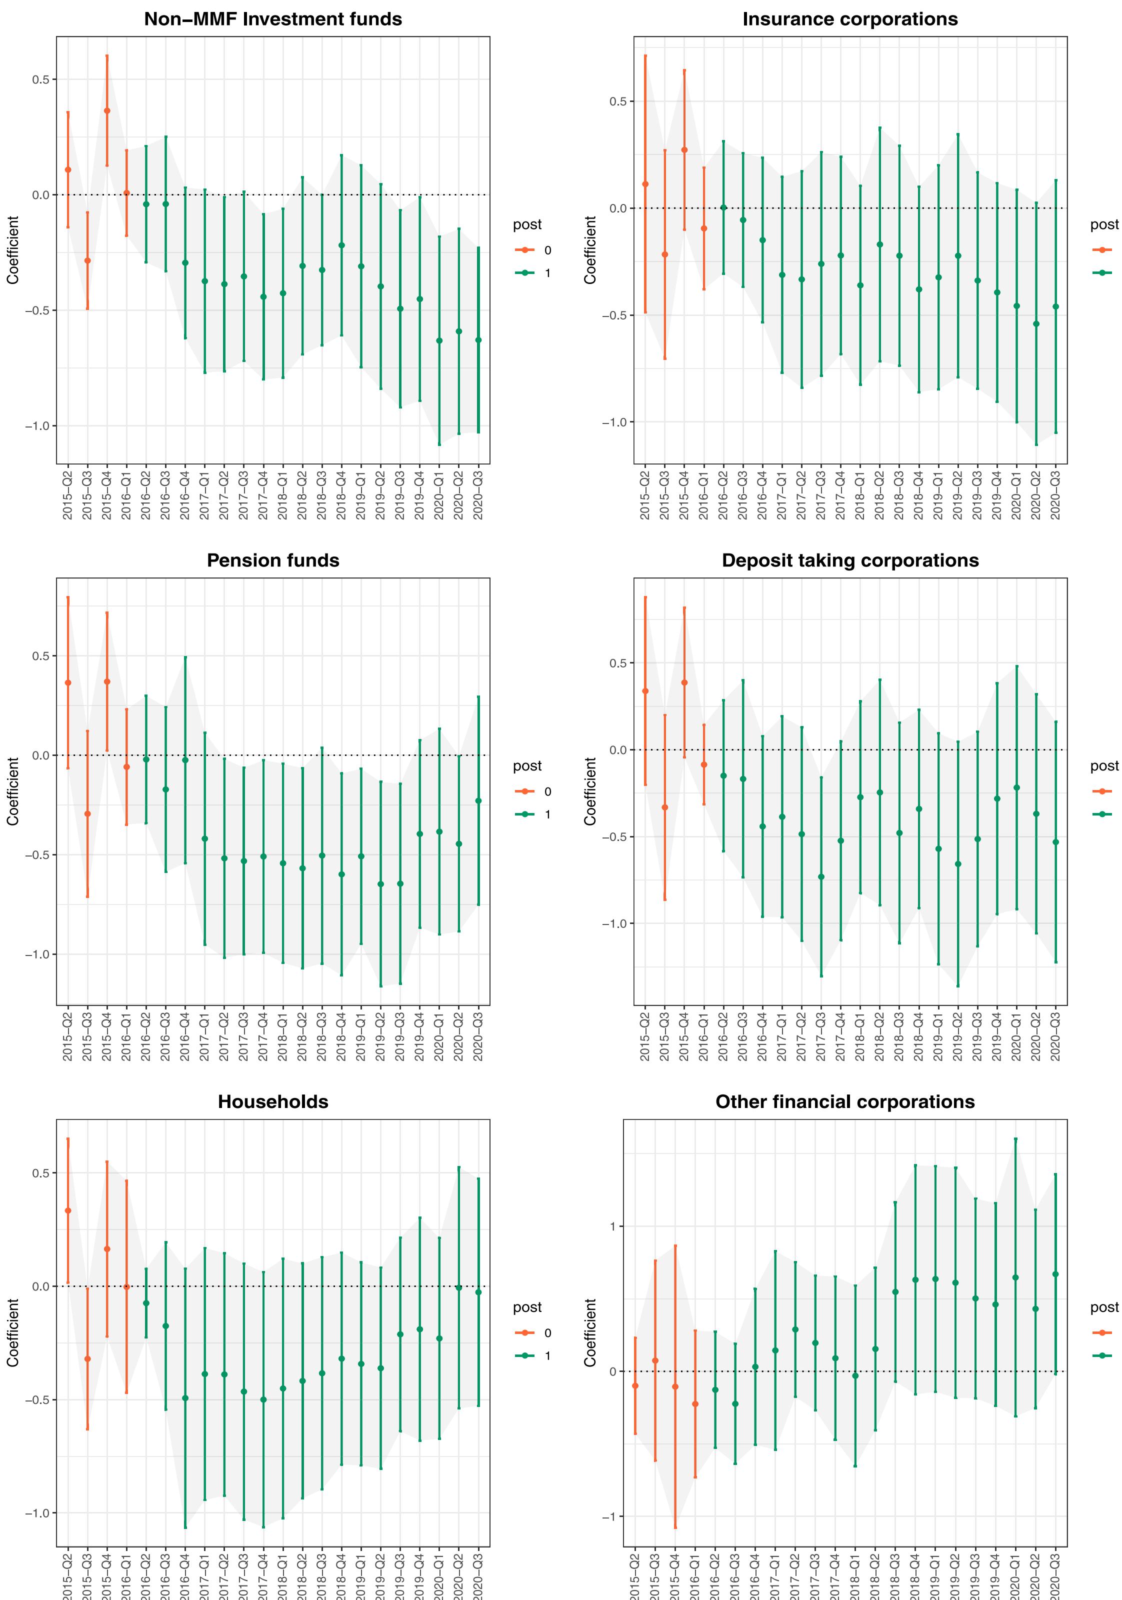

Fig. 3. Holder-sector and period-specific ATTs. The figures plot the average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. Separate panels correspond to different holder sectors. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

report on their holdings of EA securities by investors resident inside and outside the EA. 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, which is admittedly a strong assumption given the small size of non-EU holdings in our sample compared to global holdings (see also Appendix A). 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.[21] Finally, about a quarter of the records are investments from tax havens,[22] 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 Fig. 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. The two groups of large EU economies in Fig. 4 are based on the individual significance of the impact on the holder country. The impact for Germany, France, and Spain was insignificant individually, whereas the impact for Italy, the Netherlands, Sweden, and the UK was significant. An attempt to explain such differences could build on the evidence discussed in the previous section regarding more 'sophisticated' investors versus more 'naive' ones. In particular, in some countries, notably including Germany, the insignificant reaction might be explained by the large share of NFCs relatively to other types of holders, whereas the impact for NFCs was found to be insignificant (see Appendix D). By the same token, a significant reaction in the Netherlands and the UK can be partially connected with these countries being big financial investment centers.

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 Brazil, Russia, India, and China.[23] There are several possible explanations for these patterns. 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 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.[24] Fig. 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.[25]

Fig. 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 Fig. 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.

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 treated group is heterogeneous in terms of emission levels and, therefore, the reaction of investors may also differ across HC investee companies. To test the presence of such an effect, we consider the following panel data model of the (log) participation ($y_{j,t}$) with individual issuer and period effects (α_j and λ_t, correspondingly)

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

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 $z_{j,t}$ 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[26] 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.

Despite that the use 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 Fig. 9 in Appendix B, these two indicators also exhibit

[21] 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.

[22] 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.

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

[24] The Athey and Imbens (2010) 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.

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

[26] The variance–covariance matrix is clustered by issuers.

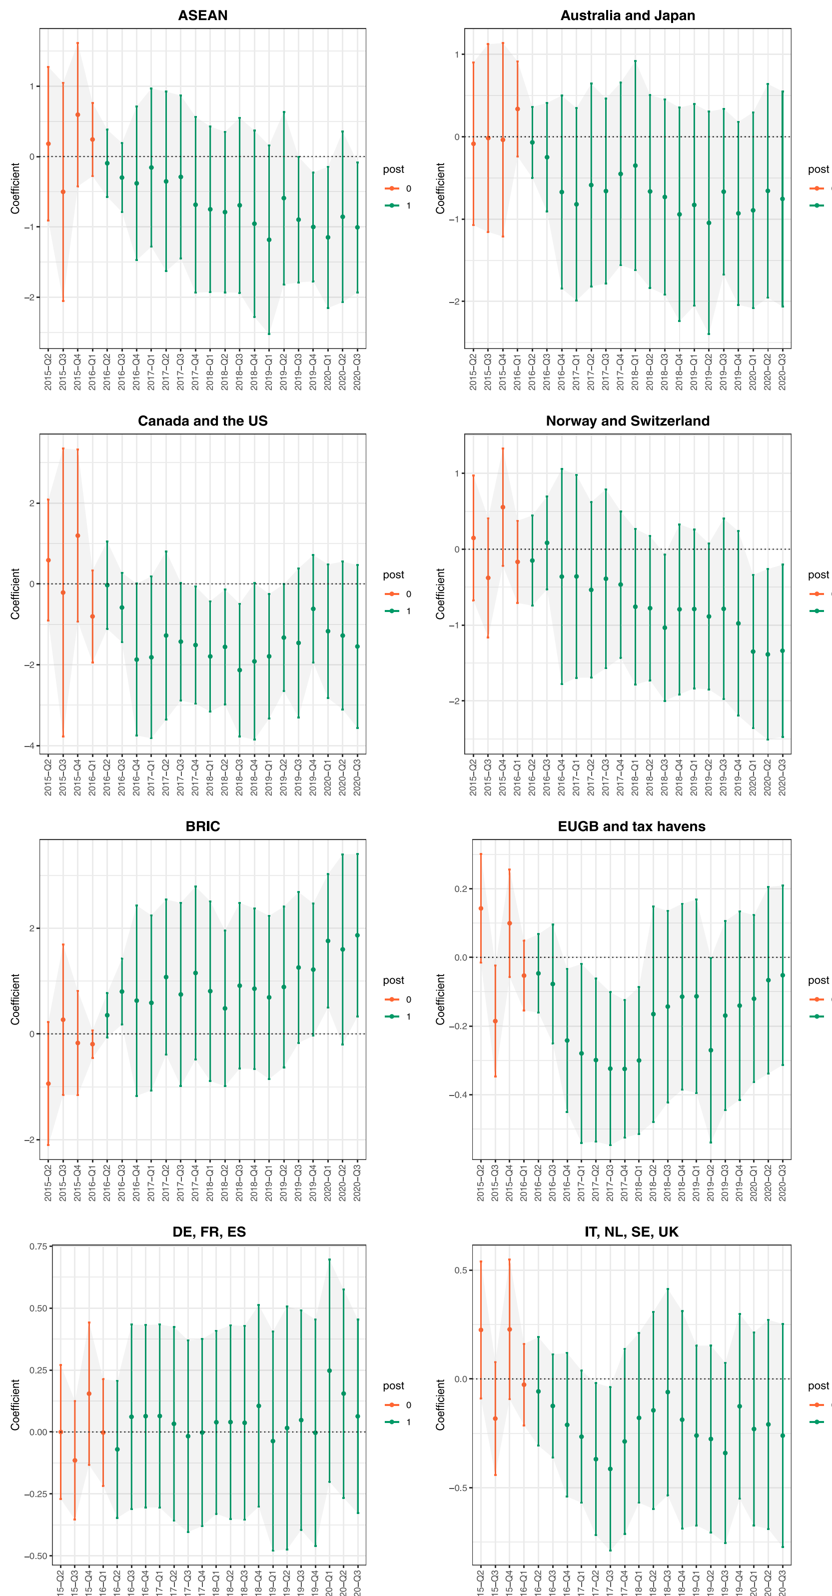

Fig. 4. Holder-area and period-specific ATTs. The figures plot the average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. Separate panels correspond to different holder areas reporting the ownership: ASEAN stands for the Association 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. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

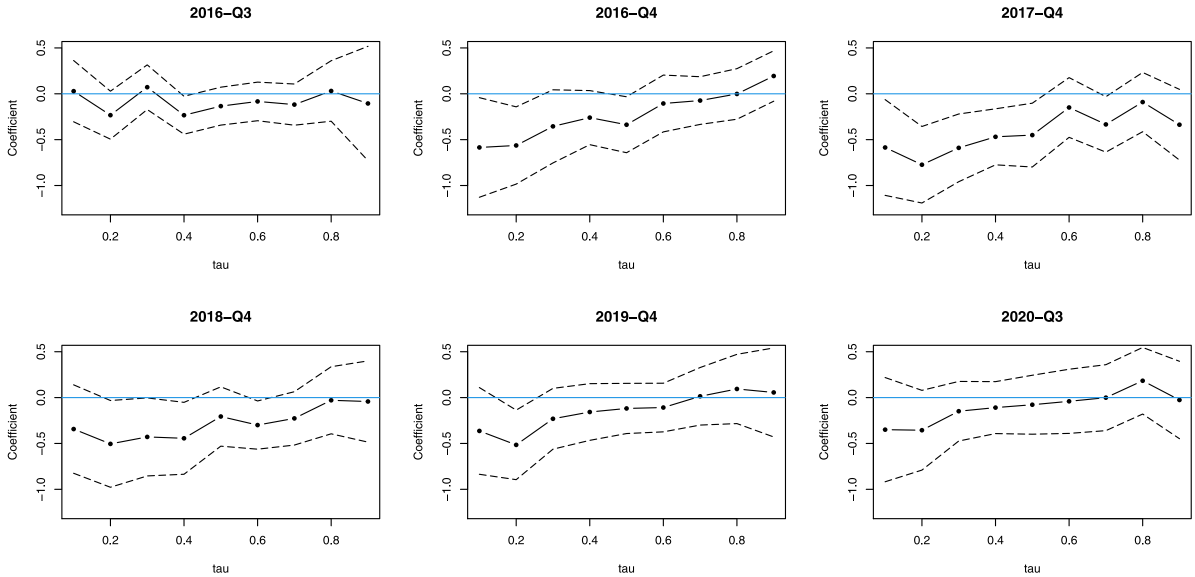

Fig. 5. Estimated QTTs. The figures plot the quantile treatment effects on the treated in terms of the (log-) participation with the ninety percent confidence bands. The panels correspond to different periods under consideration. The x-axis reports the various quantile levels (tau) of the distribution of the participation indicator.

Table 1
Heterogeneity in terms of emissions. The table summarizes the panel estimation results using the two-way fixed effects estimator, allowing for the interaction effects with emissions (in column (2)) and the intensities of emissions: relative to sales in column (3) and relative to assets in column (4). To avoid endogeneity, emissions and their intensities are fixed at their pre-Paris Agreement level of 2015.

	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	2772	2772	2160	2160
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.

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.[27] 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 (2021), 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 results indicate that, in the aftermath of the PA, financial investors significantly reduced their participation in European HC companies. In terms of magnitude of the effect, the European investors in our sample reduced their relative participation in HC companies by about a quarter.

4. Robustness checks

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

Additional estimators. First, we evaluate the robustness of the presented findings to different estimation methods. The main results underlying Fig. 2 were obtained using the Callaway and Sant'Anna (2021) approach. As robustness checks, we apply the DID_l estimator of de Chaisemartin and d'Haultfoeuille (2020a), which is unbiased under heterogeneous and dynamic effects,[28] and the generalized synthetic

[27] The intensities remain insignificant also if considered alone, i.e., dropping the *treatment ∗ emissions* term from the specifications (3) and (4), although the p-values of the *treatment ∗ emissions-to-sales* and *treatment ∗ emissions-to-assets* are smaller than that of *treatment* in such reduced specifications.

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

Table 2
Overall significance of the PA impact. For each robustness evaluation, the table presents the estimated ATTs (see column named Coeff.), their bootstrapped standard errors (S.E.), the 90% and 95% bootstrapped confidence bands, the *p*-value of the null hypothesis of parallel trends before the treatment (p-val.(Par.Tr.)), the number of matched treated and comparison units together with their sum (n.treat., n.comp., and n, correspondingly), and the number of periods (T). Acronyms: ASEAN—the Association of Southeast Asian Nations; BRIC—Brazil, Russia, India, and China; EUGB—the former EU27 with the UK.

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
Holder sector	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
	Pension funds	−0.426**	0.142	−0.660	−0.193	−0.687	−0.166	0.630	58	67	125	23
	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
Holder area	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
	Norway and Switzerland	−0.725**	0.375	−1.300	−0.150	−1.406	−0.044	0.949	40	39	79	23
	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

control estimator proposed by Xu (2017), which further allows for certain dynamics of the error term.[29] The former has similar identification assumptions to that of Callaway and Sant'Anna (2021), 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.[30] Fig. 11 in Appendix E 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, the impact is consistent across all the employed approaches.

Relevance of the aggregation level. The base results provided in Fig. 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). Fig. 12 plots similar results but using non-aggregated cross-sectional data at the issuer-holder level (as in 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.

Brexit. Next, 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. Fig. 13 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.

Alternative matching methods. The base results were obtained relying on the Coarsened Exact Matching (CEM) approach by Iacus et al. (2012).[31] Fig. 14 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 B for additional details).

Alternative periods for matching. For the main analysis we based our matching on the five-year (2011–2015) pre-treatment period averages of the matching variables. Fig. 15 plots in addition the results when the three-year (2013–2015) average and the value of 2015 alone are used instead.

Other minimum distances. As explained in Appendix B, 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. Fig. 16 plots the results obtained in the two cases, in the left and the right panel, respectively. The results 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. **Reduction of matching dimensions.** In order to achieve an increase in the number of matched cases, we might also consider reducing the number of the matching variables/dimensions. Given that the initial difference between the potential treated and control groups is smaller for historical volatility (see Fig. 7 in Appendix B), we retain the dividend yield and market value variables. In this situation, the number of matched cases increases to 98 for the treated and to 92 for the control group. The results remain as previously (see the left panel of Fig. 17 in Appendix E). However, if the matching dimensionality were further reduced performing the matching on a single market value variable, the outcome would become insignificant (see the right panel).[32]

Estimation without financial sector companies. The base matching procedure results in a substantial share of financial firms in the control group (see Table 7 in Appendix B), as they match well in terms of

[29] 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).

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

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

[32] In this case, 143 companies would be from the potentially treated group, whereas 138 firms would be from the control group.

size. When matching on companies other than those belonging to the Financial and insurance sector (NACE sector K), the key patterns remain similar (see Fig. 18).

Estimations with additional sectoral restrictions. In the baseline analysis, we excluded from the controls a number of potentially affected sectors (e.g. the fossil value chain, electricity, steel and cement, air transportation and motor vehicles manufacturing) which might still be too narrow.[33] One robustness check excludes from the control certain sectors at a broad classification level, i.e. the NACE "main sections": "B - Mining and quarrying", "C - Manufacturing", and "D - Electricity, gas, steam and air conditioning supply", "F - Construction", and "H - Transportation and storage". Results are shown in Fig. 19 in Appendix B (left panel).

Another robustness check excludes from the control a set of sectors according to the business classification available from Bureau van Dijk (BvD): "Construction", "Chemicals, Petroleum, Rubber & Plastic", "Miscellaneous Manufacturing", "Mining & Extraction", "Transport, Freight & Storage","Metals & Metal Products", "Transport Manufacturing", and "Utilities", Results are shown in Fig. 19 in Appendix B (right panel).

ESG scores. An relevant hypothesis is that investors used Environmental, Social, and Governance (ESG) scores as a proxy of emission levels for their decision making. We therefore present an additional robustness check using the company ESG score and not emission levels (see Fig. 20 in Appendix B). Although the dynamics is somewhat similar, the significance reduces substantially.

Results without participation correction. In the following robustness check, we remove from the analysis the cases where the value of the sum of the holdings in a given firm at a given time is greater than the firm's market capitalization. Fig. 21 plots the histogram of the ratio between the sum of the holdings in each firm in the matched sample and its market capitalization. The discrepancy affects a fraction of firms and the value of the ratio remains below 1.2. Our analysis looks at the changes in the participation of holders across all matched HC firms. Thus the correction of the holdings described in footnote 13 is expected to have limited or no effect. In this robustness check, we verify that this is indeed the case. Fig. 22 in Appendix B plots the respective estimates using the aggregate and holder-level data. In the former case, two instead of three matching variables are used to increase the number of the matched units—the number of matched treated and control companies increases from 26 and 36 to 43 and 58, correspondingly—as, otherwise, the dynamic path remains very similar, but the confidence bands are about 1.5 times wider. At the holder-level data (the right panel in the figure), no such changes are applied.

Additional controls. In the next set of robustness checks we add several additional control variables individually and also jointly (see Fig. 23 in Appendix B).[34] First, to account for potential exogenous changes in capital structure, notably new equity issues and share buybacks, we include as exogenous controls the percentage change in the number of issued stocks and the percentage of buybacks (relative to the number of issued stocks). The dynamic path remains analogous, whereas the confidence bands shrink substantially (see the top-left panel of Fig. 23).

Next, in the top-right panel of Fig. 23, we evaluate the robustness of our findings to the well-known home bias effect in investments (see, e.g., Boermans and Galema, 2020; Darvas and Schoenmaker, 2017, for the analysis of the European situation). Its presence implies a relative reduction of the potential set of buyers, because foreigners are less interested in acquiring stocks offered by the domestic sellers who were biased towards the domestic stocks themselves. Hence, the presence of home bias-implied segregation might also create larger costs when trying to get rid—mostly within a more limited domestic market again—of sizeable chunks of domestic stocks that became relatively unattractive. The top-right panel of Fig. 23 plots the estimated ATTs controlling for the home bias with an indicator function taking value one when an investor and a company are from the same country, and zero otherwise.[35]

The bottom-left panel of Fig. 23 presents further the outcome controlling for sector-specific trends (at the two-digit NACE sector level). The central dynamics is retained, but the significance decreases substantially. Nevertheless, whenever all these three additional control variables are included jointly, the outcome (see the bottom-right panel) is similar to the baseline result seen in Fig. 2 and the overall impact is also highly significant (see Table 3).

Intra-sectoral shifts. In the baseline analysis we studied the changes in the participation induced by the differences in emission levels across all relevant sectors. Although this accounts for the inter-sectoral shifts, it only partially captures the effect of potential intra-sectoral shifts, if present. Namely, a part of the within-sector change represented by the reduction of participation in the heaviest emitters might still be accounted for (for those emitters that also are among the top emitters across all sectors), but the participation shift towards *relatively* green companies with low emissions within a given sector, will not be identified using the previous framework. Such a change might be important, for instance, because of the presence of some (e.g., sector-knowledge and research-related) costs, linked with the inter-sectoral reallocation of funds, or policies encouraging a transformation within sectors.

Therefore, we next augment the previous analysis with two additional evaluations by changing the treatment group under consideration. Namely, from each HC sector we separate two types of companies based on their stance within a sector (in terms of emissions-to-sales[36]): (a) companies within high-carbon levels within a sector, and (b) companies with relatively low-carbon (LC) emissions within it. Note that the two groups are expected to be affected by the Paris Agreement in the opposite way due to the reallocation of funds: participation in group (a) is expected to shrink, whereas the participation in group (b) is expected to increase. In both cases, we use the same potential[37] control group of companies as in the baseline evaluation.[38]

A complication with the 'within-sector' approach is that it is not granted that all investors classify and use the same (general or specialized) classification or even the specific aggregation level for such decisions.[39] Therefore, to proxy for the potential underlying classification, we next present the results with two different classifications and aggregation levels. One is the Climate Policy Relevant Sectors (CPRS), developed in Battiston et al. (2017), and widely used by policy institutions, which regroups NACE 4 digit codes into classes of activities with different relevance for transition risk. A second classification we

[33] This additional check besides emissions aims mostly at filtering some cases where companies either are expected to benefit from reallocation of funds in HC sectors or have themselves low levels of direct emission but high levels of indirect emissions, e.g., mining of coal.

[34] Note that these robustness checks are performed using the de Chaisemartin et al. (2021) framework where the post-treatment controls can be added, although with the cost of losing the double-robustness property.

[35] The positive home bias effect is highly significant itself but its interaction with the level of emissions is (marginally) insignificant.

[36] We use the relative intensity of emissions here, because this ranking within a sector is performed on sector's companies before any matching.

[37] I.e., before the matching procedure.

[38] One could even suspect that these companies from the control group might have benefited even more than only relatively greener companies from within HC sectors after the Paris Agreement. However, in such a case, one would observe a reduction in the relative participation when comparing the groups of 'relatively green' companies (i.e., only within a sector) with 'globally green' companies, whereas the contrary holds empirically (see Fig. 24 in Appendix B).

[39] This is illustrated also by the complicated process of defining the taxonomy (see OECD, 2020).

use NACE[40] at 2-digits level, denoted as NACE2.[41] In each case, the panels of Fig. 24 plot the estimated dynamic impacts for the two within-sector effects: the LC and HC companies in the treatment group, correspondingly.

As expected, we find that the impact on the relatively LC companies (within a sector) is positive, whereas the relatively HC companies tend to experience a negative effect.[42] Another notable feature of the dynamics is that, in both cases, it starts to accelerate since 2019, especially, after the EU technical screening criteria for economic activities became more clearly defined.[43],[44] This also could explain, at least partially, the decreasing impact towards the end of the period observed in Fig. 2, where the estimated impact is driven more by the inter-sectoral reallocation of funds and less connected with the within-sectoral changes. Within the sectors, the impact on the relatively green companies seems to be somewhat more notable (at least, whenever considering the NACE2 case), which is in line with results in Monasterolo and de Angelis (2020). However, a long-term sustainable increase in the participation in sectoral LC companies need to have other sources than savings or borrowing, and the previously established inter-sectoral change might be an important counterpart of this process.

The result documented in this specific robustness check means that investors have reacted to the emission performance of firms within sectors and would support the interpretation that investors have taken a *best-in-class* approach. However, the result presented here co-exists with the result that the change in participation is explained by the differences in emission levels across sectors, as analyzed in the next point.

Inter-sectoral shifts. The baseline results considered the emissions-based total change of participation in companies irrespective if driven by inter- or intra-sectoral shifts, while the previous robustness check evaluated the emissions-based shifts within a sector. One could further ask whether the disinvestment did not take place at the sectoral level irrespective of the carbon-intensity of individual companies. To investigate this point, this robustness check performs an analogous matching/ATT analysis for six HC industries (Fossil, Electricity, Airlines, Automotive, Cement, and Steel) considering all the companies belonging to them as treated, irrespective of their emission levels, while using the same control group as previously.[45] Recall that companies from these sectors are prevented from entering the control group *a priori*.

In none of the considered cases the results are significant.[46] Only the Fossil sector is almost significant at the 10% significance level, whereas its overall ATT would become significant even at the 5% significance level if the treatment were considered to have started since 2016-Q4 instead of 2016-Q2. Still, Fig. 25 reveals some interesting patterns. First, the participation in the Fossil sector tends to decrease quite steadily, whereas the participation in the Electricity sector tends to increase.[47] Airlines and Automotive industries feature a U-shape behavior over time, whereas the Cement and Steel sectors do not reveal any clearer pattern. With the discussed conditional exception for the Fossil sector, the presented insignificant sectoral outcome when emissions are not taken into account could indicate that firm-level carbon emissions (actual, estimated, or even perceived) were an important factor of decision making, at least during the first half of the analyzed period. It should however be acknowledged, that investors might have classified companies in different classes/sectors compared to the sectoral classification we use in this robustness check, which could be a reason why we do not find a significant impact.

Estimations with a random split of firms. 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.[48] The matching procedure now is applied to this pseudo split into treated and control firms.[49] Fig. 26 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.[50]

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 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 predominant number of cases, the overall PA impact is statistically significant at least at the 10% significance level.

5. Conclusion

In this paper, we find that the Paris Agreement and the US withdrawal from it 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. However, the process reverted after the US announcement of withdrawal from the PA. These finding are consistent with the explanation that investors revised their expectations on HC firms as becoming more risky after the PA and that the announcement created uncertainty about the viability and credibility of the agreement itself. One the one hand, the baseline analysis shows that the effect is driven by the absolute level of emissions across sectors. At the same time, one of the robustness checks on intra-sectoral shifts has highlighted the role of relative emission performance within sectors, suggesting some role for a best-in-class approach in investor's decisions.

These changes in participation can also be interpreted in terms of transfer of risks. On the one hand, the reduction in overall participation

[40] NACE is the standard classification of economic activities in Europe, maintained by Eurostat https://ec.europa.eu/eurostat/web/nace-rev2.

[41] The results with NACE codes at the most aggregate level (i.e. "main sections") are similar to the CPRS case.

[42] In each case, i.e., NACE2 and CPRS, there is either a HC or LC case with the significant overall impact of the PA (see Table 3). It should be pointed out that the results are even more significant whenever smaller shares (quartile, quintile, etc.) of LC companies within sectors are considered. Whereas, for HC companies within sectors, the significance increases with larger (and not smaller) shares (half and even more). This potentially points towards a certain asymmetry of the impact with a smaller subset of relatively green companies.

[43] Note that, in December 2018, the Technical Expert Group (TEG) on Sustainable Finance set up by the European Commission published a first draft proposal for the EU Taxonomy for Sustainable Activities and asked for public feedback. In, June 2019, the TEG released a technical report containing proposed technical screening criteria for substantial contribution to climate change mitigation across 67 economic activities (see TEG, 2020).

[44] The European Green Deal could have further contributed to this increase during the latest quarters.

[45] It should be pointed out that lifting the (zero or very low) emissions requirement on control group firms results in even greater insignificance than presented next.

[46] Because of their insignificance, these results are not included in Table 3.

[47] It is clear that investments in renewable energies could play a role here.

[48] 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.

[49] 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.

[50] Furthermore, we repeated such simulations 1000 times exploring the null hypothesis of overall ATT being zero and obtained a good correspondence between the empirical rejection frequencies (0.112 and 0.067) and the respective nominal sizes (0.1 and 0.05, correspondingly).

Table 3
Overall significance of the PA impact (robustness checks). For each robustness evaluation, the table presents the estimated ATTs (see column named Coeff.), their bootstrapped standard errors (S.E.), the 90% and 95% bootstrapped confidence bands, the *p*-value of the null hypothesis of parallel trends before the treatment (p-val.(Par.Tr.)), the number of matched treated and comparison units together with their sum (n.treat., n.comp., and n, correspondingly), and the number of periods (T). Acronyms: BvD—Bureau van Dijk; CPRS—Climate Policy Relevant Sectors; ESG—Environmental, Social, and Governance; HC—High-Carbon; LC—Low-Carbon; NACE—Nomenclature of Economic Activities.

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
2 matching dimensions	−0.201**	0.076	−0.327	−0.074	−0.351	−0.050	0.859	98	92	190	23
Without financials	−0.407**	0.132	−0.624	−0.189	−0.666	−0.147	0.742	26	27	53	23
Sectoral restrictions (NACE class.)	−0.174*	0.103	−0.341	−0.007	−0.391	0.043	0.988	47	44	91	23
Sectoral restrictions (BvD class.)	−0.214**	0.100	−0.369	−0.059	−0.420	−0.008	0.980	51	49	100	23
With ESG score	−0.160	0.138	−0.383	0.063	−0.435	0.115	0.595	17	31	48	23
Without cases requiring correction 1[a]	−0.198*	0.102	−0.367	−0.028	−0.411	0.015	0.741	43	58	101	23
Without cases requiring correction 2[b]	−0.383**	0.145	−0.598	−0.167	−0.676	−0.090	0.202	26	36	62	23
Additional controls 1[c]	−0.363**	0.098	−0.524	−0.201	−0.554	−0.172	0.406	59	69	128	23
Additional controls 2[d]	−0.340**	0.102	−0.438	−0.101	−0.469	−0.070	0.298	59	69	128	23
Additional controls 3[e]	−0.403	0.265	−0.842	0.035	−0.921	0.115	0.209	59	69	128	23
Additional controls (all)	−0.630**	0.245	−1.034	−0.226	−1.107	−0.153	0.287	59	69	128	23
LC intra-sector (CPRS class.)	0.199	0.151	−0.052	0.450	−0.074	0.472	0.908	32	28	60	23
LC intra-sector (NACE2 class.)	0.271*	0.149	0.032	0.511	−0.016	0.559	0.920	33	40	73	23
HC intra-sector (CPRS class.)	−0.201**	0.088	−0.348	−0.053	−0.386	−0.015	0.985	49	56	105	23
HC intra-sector (NACE2 class.)	−0.123	0.113	−0.320	0.075	−0.336	0.091	0.997	37	51	88	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

[a] With aggregate data and 2 matching variables.
[b] With holder-level data.
[c] Percentage change in stocks and the share of buybacks.
[d] Home bias indicator function taking value one when investor's and company's countries coincide and zero otherwise).
[e] Sector-specific trends (at the two-digit NACE level).

in HC companies by the holders 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. Furthermore, more financially 'sophisticated' investors and financially developed countries seem to reallocate holdings more (and more steadily) in anticipation of the impact of transition risk. We also find that investors reduce to a smaller extent their participation in those high-carbon firms where they hold larger stakes.

It should be noted that our findings on the reduction in participation of certain holder types in high carbon firms does not imply unconditionally a reduction of their exposure to transition risk. Under combinations of the following conditions, transition risk may actually not materialize: (i) if the low-carbon transition does not take place (e.g. in so-called hot-house-world scenario of NGFS (2019)); (ii) if high-carbon companies have low leverage ratios and high profit margins that allow them to deal financially with the adverse impact of climate policies; (iii) if high-carbon firms undertake sufficient green investment or change their business models such to adjust to a low-carbon economy environment.

However, under conditions in which transition risk can materialize, the reduction in participation of EU holders in high-carbon firms lowers their transition risks. On the one hand, the shift of HC-ownership towards holders located in the BRIC countries can be seen as a transfer of transition risk outside the EU, since holders from these countries tend to increase their participation in HC companies, which can also be expected to be less intertwined with the EU investors. On the other hand, the shift of HC-ownership towards less regulated financial sectors could imply a transfer of transition risk where this is less monitored.

As further research, it would be interesting to investigate the following aspects. First, our analysis focused only on equity holdings, while holdings of loans and bonds may, in principle, have been influenced differently. Related work on loans find results in similar direction, at least right after the PA (Reghezza et al., 2021). Therefore, future work could examine the combined impact on holdings of stock, bonds and loans. Second, the SHS aggregation at the level of the institutional sector does not allow to investigate whether different investors within the same sector and (e.g. different banks in the same country) 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. Third, 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. Fourth, this analysis could be extended beyond Scope 1 emissions.

Finally, our results have some relevant policy implications. We document a case in which a global environmental policy has an impact on investors behavior in terms of portfolio allocation and the lack of coherence at global level reverts the initial effects. This finding lends support to the idea that the successful redirection of global financial flows towards climate action (Article 2c of the PA) requires policy credibility (Battiston et al., 2021) and in particular a clear and unanimous signal from the global community of policy makers. In this regard, regional policy measures such as the Carbon Border Adjustment Mechanisms (already in place in California, and planned by the EU, Canada, and Japan) could help to achieve a more homogeneous response of the global investor community to limit carbon leakage. Overall, as the low-carbon transition picks up speed in some parts of the

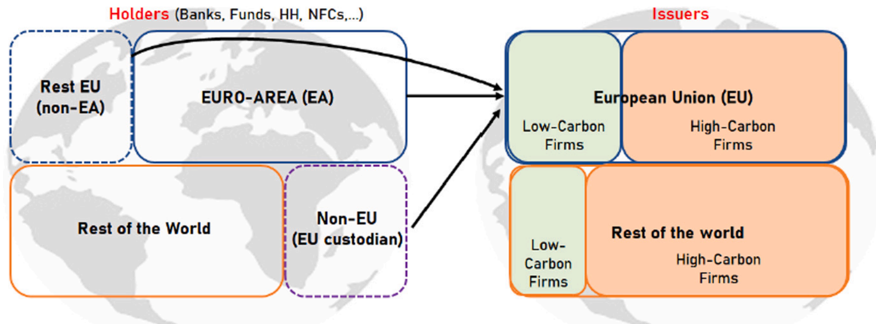

Fig. 6. Securities included in our sample. This figure shows the coverage of the employed ECB Securities Holdings Statistics data in terms of holders (holder sector and area) and issuers (area and carbon-intensity of companies). *Source:* Carola Müller (2021). Discussion: Over with carbon? Investors' reaction to the Paris Agreement and the US withdrawal, IV Conference on Financial Stability (23–25 November, 2021).

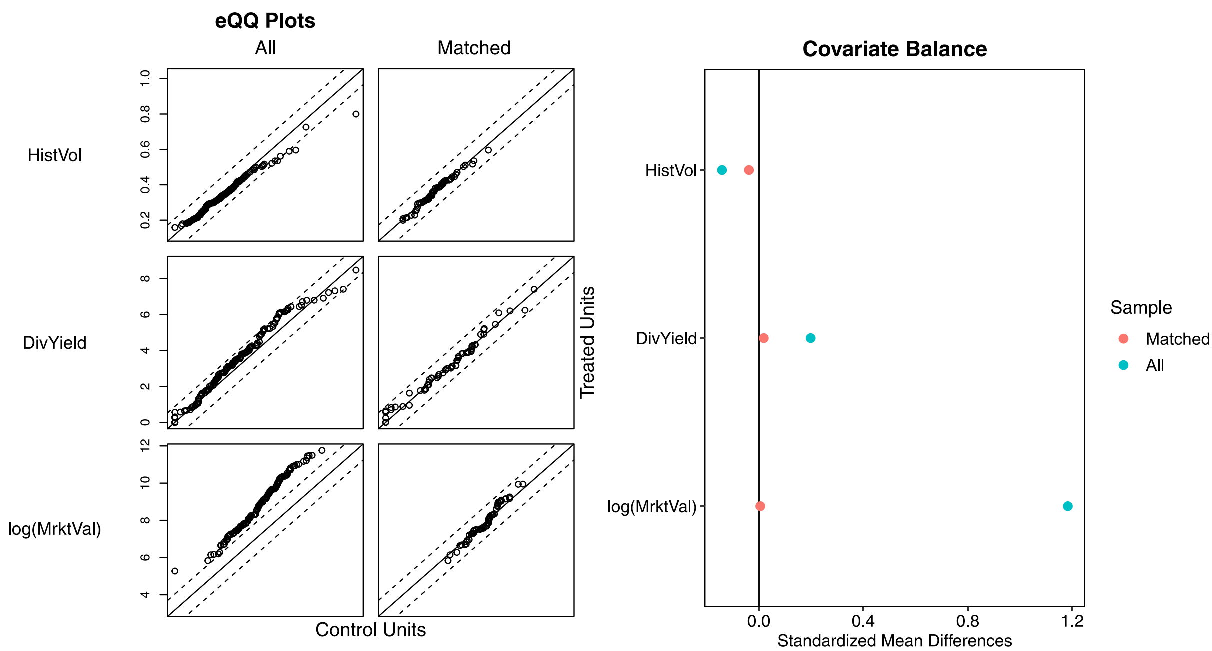

Fig. 7. Empirical quantile–quantile (eQQ) plots and covariate balance: Unmatched vs. CEM-matched. The left panel presents the empirical quantile–quantile plots of the treated and control units in terms of the three matching variables before and after matching. The right panel presents the corresponding standardized mean differences between the treatment and control units before and after matching. The matching here is performed applying the Coarsened Exact Matching (CEM). The underlying data are from the Refinitiv Eikon database.

world, a closer monitoring of the buildup of transition risk in particular sectors and jurisdictions is warranted.

CRediT authorship contribution statement

Lucia Alessi: Conceptualization, Formal analysis, Resources, Supervision, Writing – original draft, Writing – review & editing. **Stefano Battiston:** Conceptualization, Investigation, Resources, Supervision, Writing – original draft, Writing – review & editing. **Virmantas Kvedaras:** Data curation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing.

Appendix A. Data description

The SHS database covers holdings of investors residing in the euro area and several non-euro area EU countries, as well as non-resident investors' holdings of euro area securities that are deposited with a euro area custodian. The actual reporting population consists of resident monetary financial institution (MFIs), investment funds (IFs), and financial vehicle corporations (FVCs). MFIs and IFs report data on own holdings of securities and on securities they hold in custody on behalf of other investors. In this latter case, for each security holding they need to report the sector of the holder. In the case of resident investors, this can be insurance corporations, pension funds, FVCs, other financial intermediaries, non-financial corporations (NFCs), Government, households.[51] Looking at shares only, which are the focus of this paper, and investment funds shares/units in 2016, i.e. the beginning of our 'post' period, total holdings by euro area investors as covered by this dataset amounted to EUR 13 tn.[52] While the SHS includes securities issued by both euro-area and non-euro area issuers, we only focus on European issuers. Fig. 6 shows the coverage of our sample.

[51] The legal basis for the collection of SHS data is Regulation (EU) No 1011/2012 of the European Central Bank of 17 October 2012 concerning statistics on holdings of securities (ECB/2012/24). A detailed description of the SHS dataset is available in the ECB Economic Bulletin 2015 issue 2, Article 2 available at https://www.ecb.europa.eu/pub/pdf/other/eb201502_article02.en.pdf. Based on this article, the SHS database covers around 83% of the total outstanding amount of securities issued by euro area residents.

[52] See https://www.ecb.europa.eu/press/pr/date/2017/html/pr170202_1.en.html.

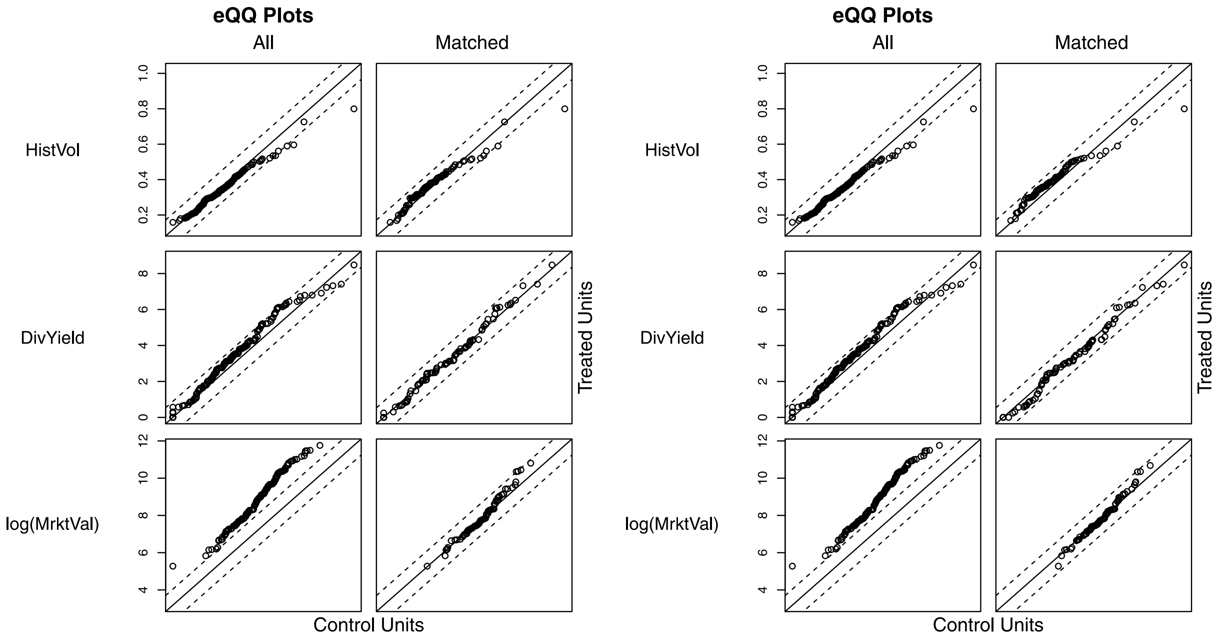

Fig. 8. Empirical quantile–quantile (eQQ) plots: Unmatched vs. genetic matching (left panel) and nearest neighbors (right panel). Both panels present the empirical quantile–quantile plots of the treated and control units in terms of the three matching variables before and after matching. The left panel relies on the genetic matching, whereas the panel on the right is based on the nearest neighbors-based matching. The underlying data are from the Refinitiv Eikon database.

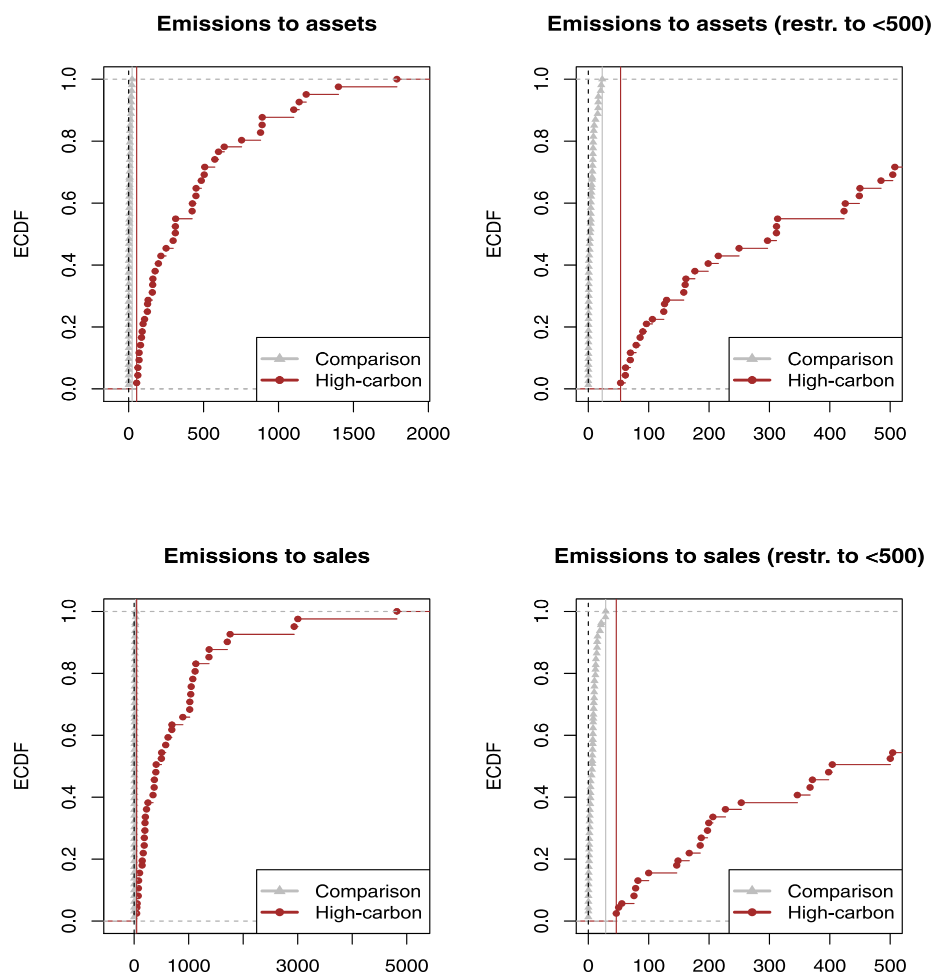

Fig. 9. Relative emissions to assets and sales in the matched groups of treated (high-carbon) and comparison firms. The figures plot the empirical cumulative distribution functions (ECDFs) of emissions to assets (upper panels) and emissions to sales (lower panels) for the high-carbon and comparison units. The panels on the left side plot the ECDFs for the whole support, whereas the support is shrunk to below 500 in both panels on the right side for a better visibility of the support separation. The underlying data stem from the Bloomberg and Refinitiv Eikon databases.

Table 4
Summary statistics of data used for baseline matching and estimation. The table provides the summary statistics of main variables (by companies) together with the respective data sources and transformations. The participation is characterized at the aggregate level.

Variable	1% quant.	Med.	99% quant.	Mean	St.Dev.	Obs.	Period	Source	Transf.
Participation (without log., ×100)	0.0004	0.007	0.090	0.014	0.034	14 408	2015Q1–2020Q3	ECB SHS	Log
GHG emissions (th.m.tons)	0.0003	0.121	59.0	3.51	13.9	611	2017	Bloomberg	None
GHG intensity (per assets)	0.011	21.7	1555	144.8	439.6	508	2017	Bloomberg	None
GHG intensity (per sales)	0.156	34.5	3231	247.0	755.5	507	2017	Bloomberg	None
Market value (EUR, mln.)	56.7	1787	85 332	6643	14 588	627	2011–2015 avg.	Refinitv	Log
Dividend yield (%)	0	2.75	8.65	3.01	2.24	627	2011–2015 avg.	Refinitv	None
Historical volatility	0.135	0.330	0.748	0.348	0.126	627	2011–2015 avg.	Refinitv	None

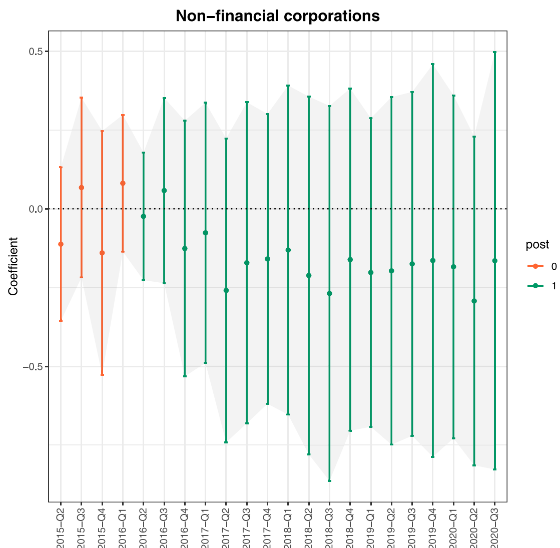

Fig. 10. Estimated ATTs for the holder sector of Non-Financial Corporations (NFC). The figure plots the average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

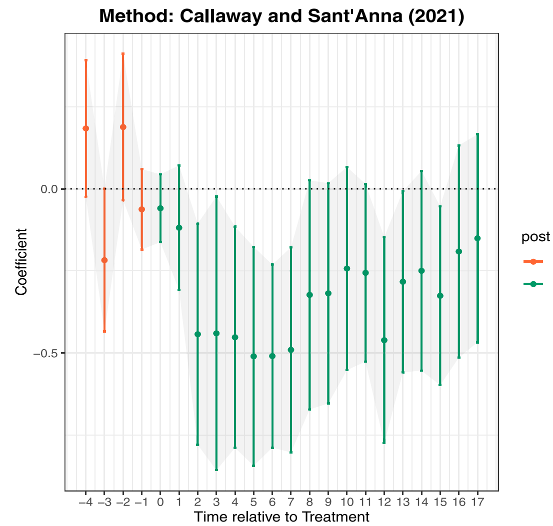

Fig. 12. Estimated ATTs with holder-level data. The figure plots the estimated average treatment effect on the treated in terms of the holder-level (log-) participation, as defined in Eq. (1), with the ninety percent bootstrap confidence bands. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

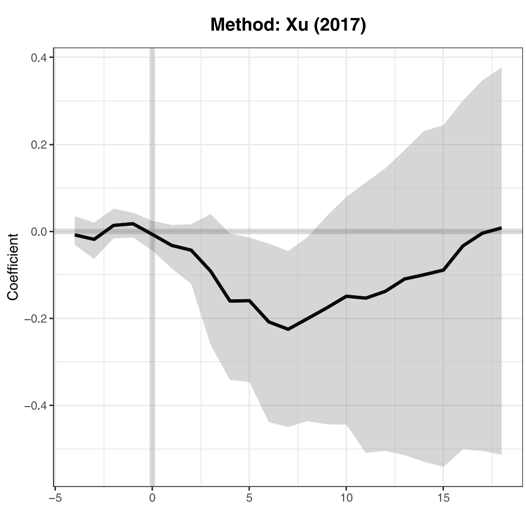

Fig. 11. Alternative estimators. The figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. To underscore the different methods, we keep the style of figures in correspondence with the respective packages: 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).

Among all European issuers, we select a subset based on data availability. Namely, looking at the 2000 largest EU companies (in terms of their market value), all covariates needed for the matching procedure—the dividend yield, the historical stock return volatility, and the market value—are available in 80% of the cases. However, the data on emissions is available only in about 45% of the cases under discussion. This leads to a substantial shrinkage of the number of issuers, which amounts to 627 after the 1% winsorization.[53] Table 4 reports the descriptive statistics of the respective sample.

Furthermore, we intentionally drop companies from the fossil value chain, electricity, steel and cement, air transportation and motor vehicles manufacturing from the control group even if they are not recorded as heavy GHG/CO2 emitters. The resulting numbers of issuers under comparison in the baseline scenario for the treated and control groups are reported in Table 5 in the next section.

Appendix B. 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.[54] The CEM bounds the degree of model dependence and the treatment effect estimation error, eliminates the need for a separate

[53] The winsorization is based on the ratio of the sum of holdings by all holders and the total market value.

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

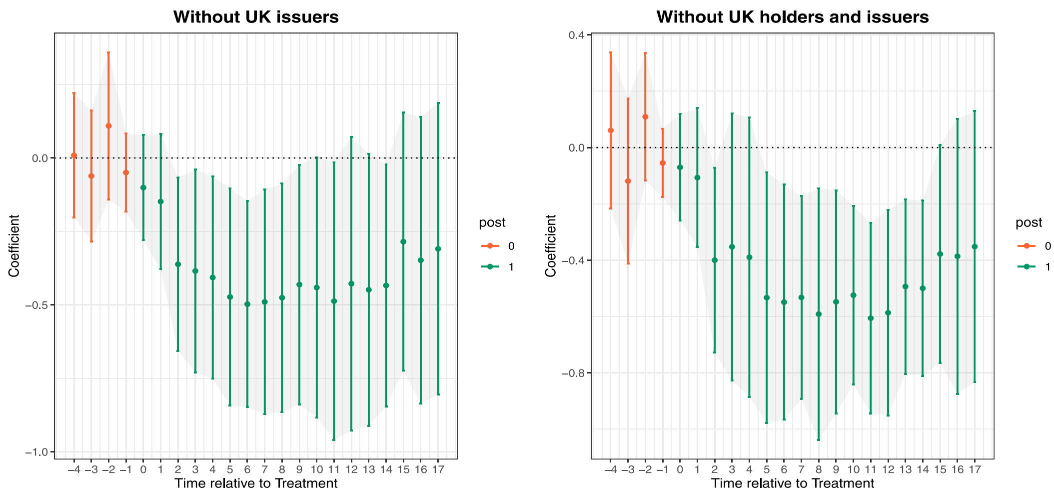

Fig. 13. Estimated ATTs without the UK. Both figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. In the left panel, only the UK issuers are excluded, whereas both the UK issuers and holders are omitted in the right panel. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

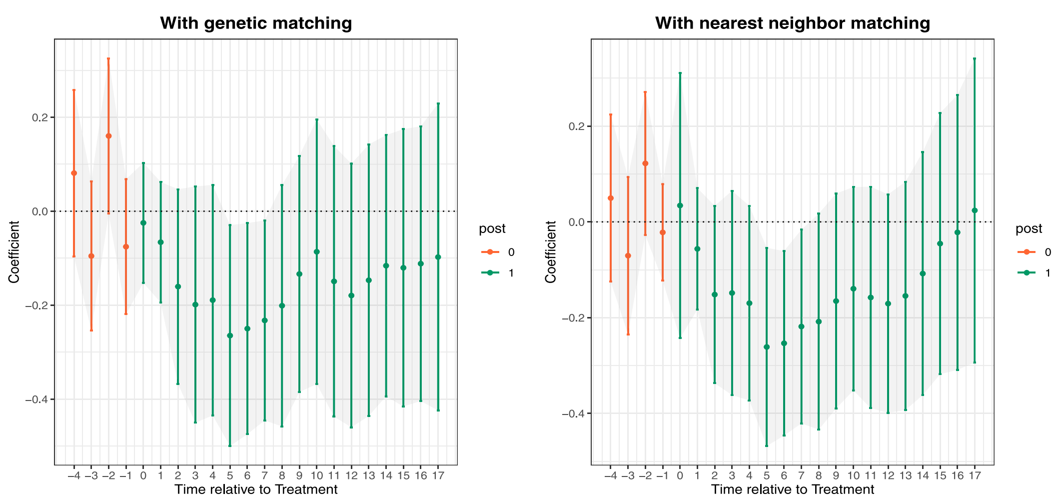

Fig. 14. Estimated ATTs with alternative matching methods. Both figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. The left panel uses the genetic matching, whereas the nearest neighbor matching is applied in the right panel. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

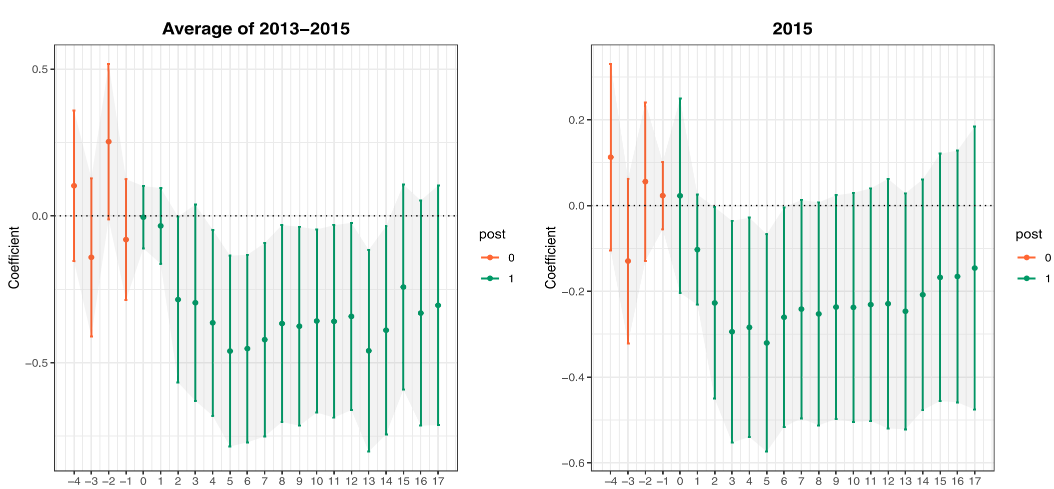

Fig. 15. Estimated ATTs with alternative periods used for matching. Both figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. The panels are linked to different periods/aggregation used for the matching procedure. The average 2013–2015 values of the matching variables are employed in left panel, whereas the values of 2015 are used in the right panel. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

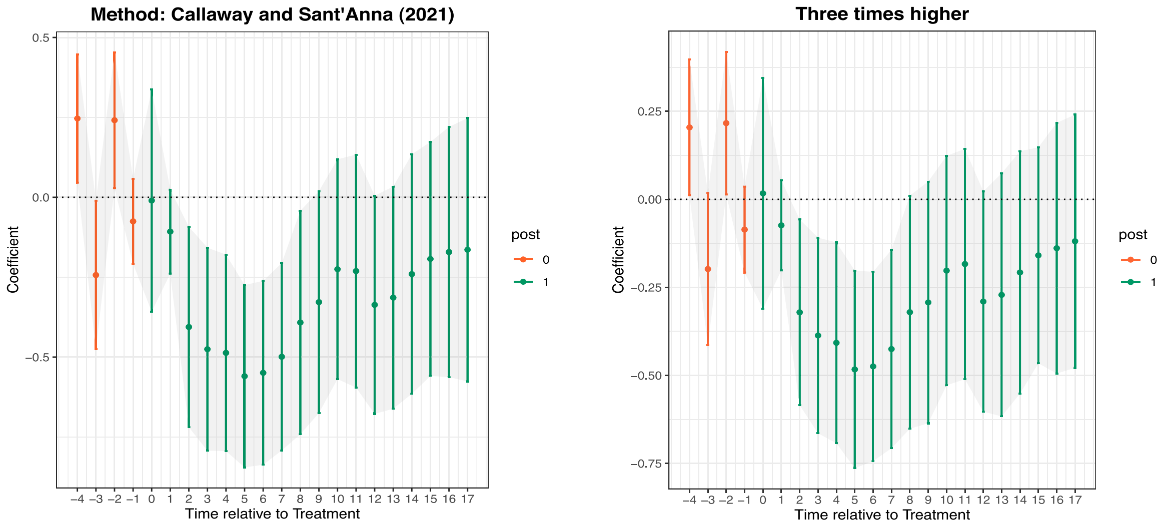

Fig. 16. Estimated ATTs with alternative minimum distance of ratios of emissions to assets and sales in the matched treated and comparison groups. Both figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. No restriction is imposed on emissions-to-assets and emissions-to-sales in the left panel, whereas at least three times higher ratios for treated than comparison units are required in the right panel. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

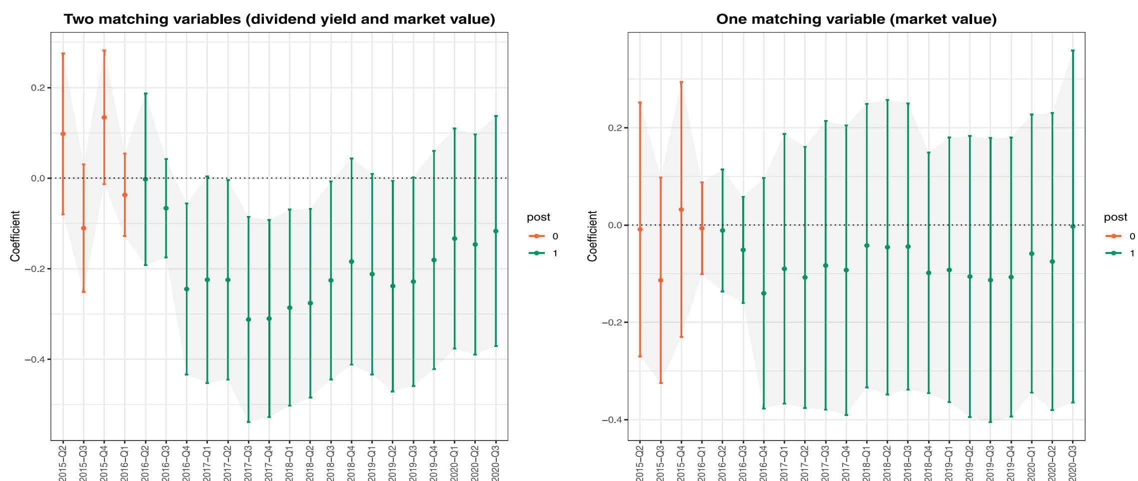

Fig. 17. Results with a smaller number of matching variables. The figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. Two matching variables are used instead of three in the left panel (market value and dividend yield), and only market value in the right panel. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

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.[55]

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.[56] 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.

[55] 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.

[56] 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.

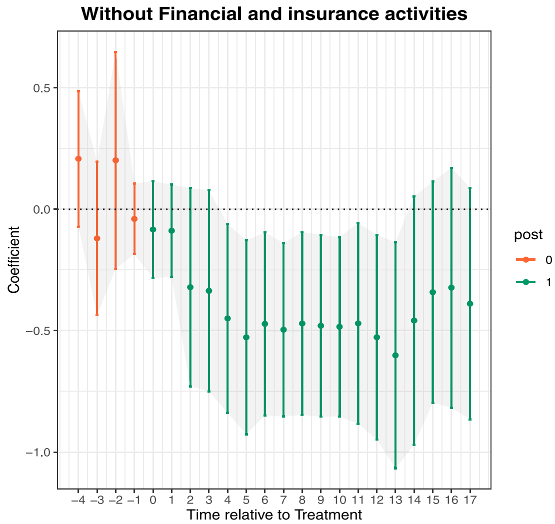

Fig. 18. Estimated ATTs without Financial activities (sector K of NACE). The figure plots the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands, when companies from the Financial and insurance activities are excluded. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).).

Table 5
Multivariate imbalance, local common support, and number of matched units. The table characterizes the matching samples for different matching methods: CEM—coarsened exact matching (the base method), GEN1—genetic matching; NN2—nearest neighbor matching.

	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

For each covariate, Figs. 7 and 8 illustrate the performance of the performed matching procedures in terms of the empirical Quantile–Quantile (eQQ) adequacy between the treated and comparison (control) units. Fig. 7 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 Fig. 7), 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 Fig. 8, 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 5). 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.

The resulting distributions of treated and comparison firms by the broad NACE activity sectors are reported in Tables 6 and 7.

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

Table 6
Activity sector of treated firms. The table characterizes the distribution of matched treated firms across NACE activities.

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 7
Activity sector of comparison firms. The table characterizes the distribution of matched comparison companies across NACE activities.

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

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 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 Fig. 9 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 gray lines, respectively), the support is cut at 500 in the figures on the right side that, otherwise, present the same information.

Appendix C. 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 Fig. 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 (2021) 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

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

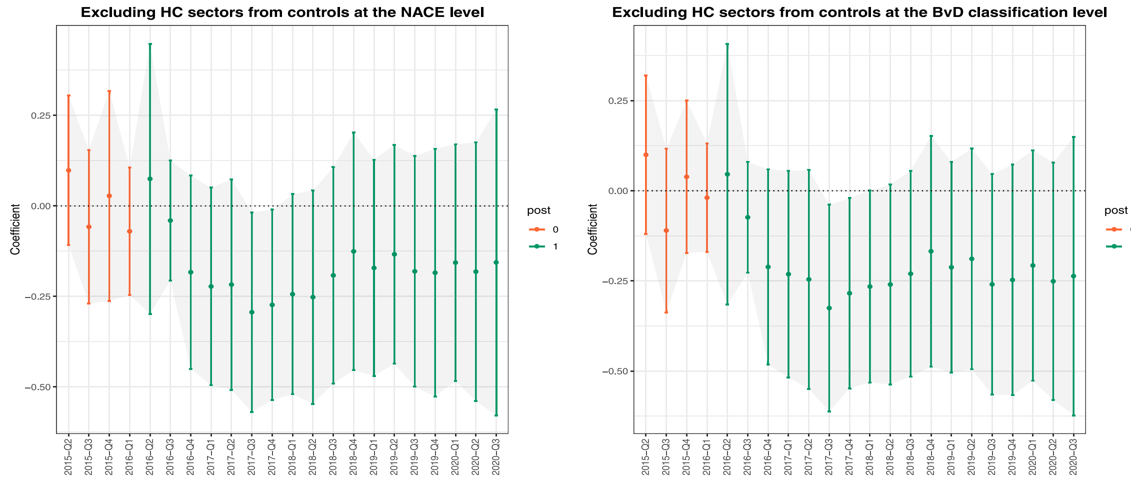

Fig. 19. Different levels of exclusion of potentially affected from controls. The figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands, when companies from the following sectors are excluded from the potential control group. Left panel: “B - Mining and quarrying”, “C - Manufacturing”, and “D - Electricity, gas, steam and air conditioning supply”, “F - Construction”, and “H - Transportation and storage” from NACE. Right panel: “Construction”, “Chemicals, Petroleum, Rubber & Plastic”, “Miscellaneous Manufacturing”, “Mining & Extraction”, “Transport, Freight & Storage”,“Metals & Metal Products”, “Transport Manufacturing”, and “Utilities” using the Bureau van Dijk (BvD) classification of sectors.

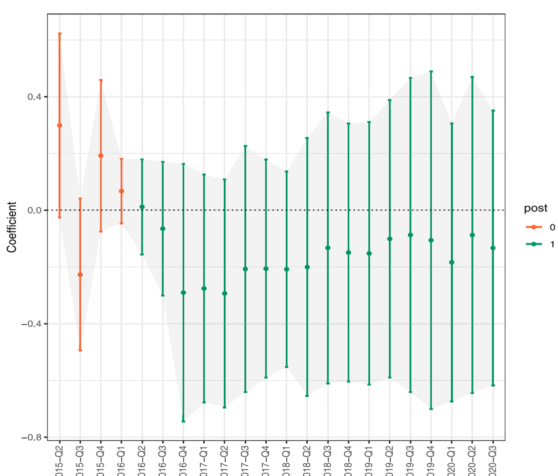

Fig. 20. ESG scores-based evaluation (high ESG-scores as controls and low ESG-scores as treated). The figure plots the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands, whenever the ESG scores-based split is used instead of the levels of emissions. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

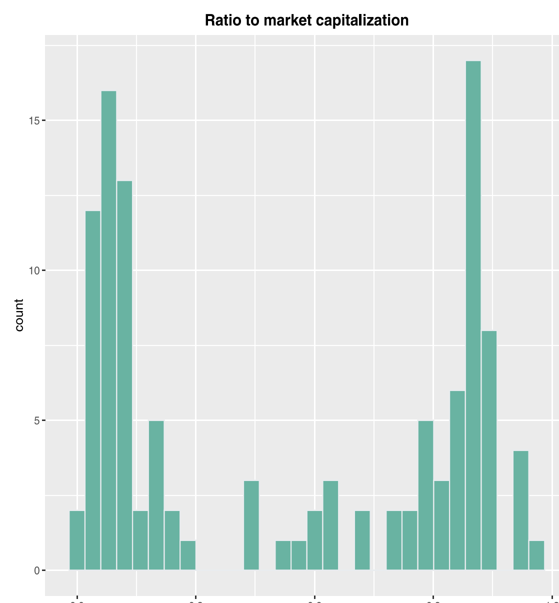

Fig. 21. Histogram of the ratio of total sum of holdings in the ECB SHS to the market capitalization for the sample of matched companies in the baseline estimation.

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 (2021) 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:

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

Appendix D. Reaction of NFCs

See Fig. 10.

Appendix E. Robustness plots

See Figs. 11–26.

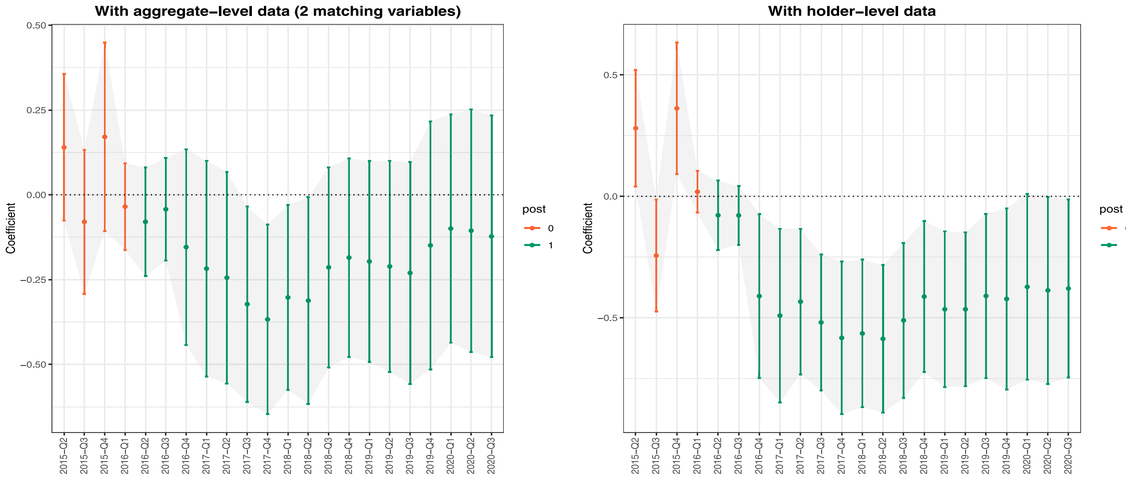

Fig. 22. Estimated ATTs without issuers for which the total sum of holdings in the SHS is greater than the market capitalization. Both figures plot the average treatment effects on the treated in terms of the (log-) participation with the ninety percent confidence bands. The left and right panels use the aggregate and holder-level data, correspondingly. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

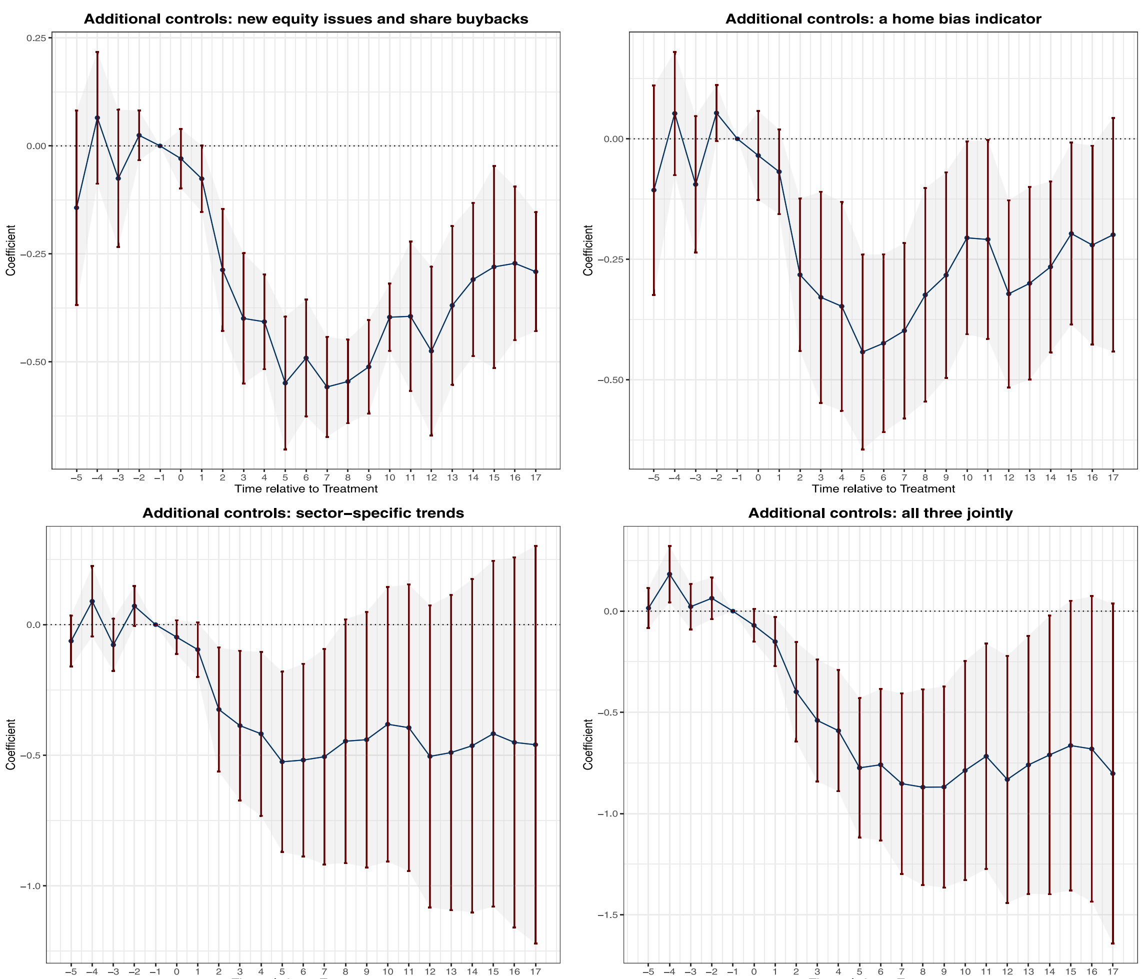

Fig. 23. Estimations with additional control variables. The figures plot the estimated average treatment effect on the treated in terms of holder-level (log-) participation with the ninety-five percent bootstrap confidence bands using the de Chaisemartin et al. (2021) implementation with three additional control variables: (i) new equity issues and share buybacks; (ii) a home bias indicator; and (iii) sector-specific trends (at NACE two digit level). The last (bottom-right) panel includes these controls jointly.

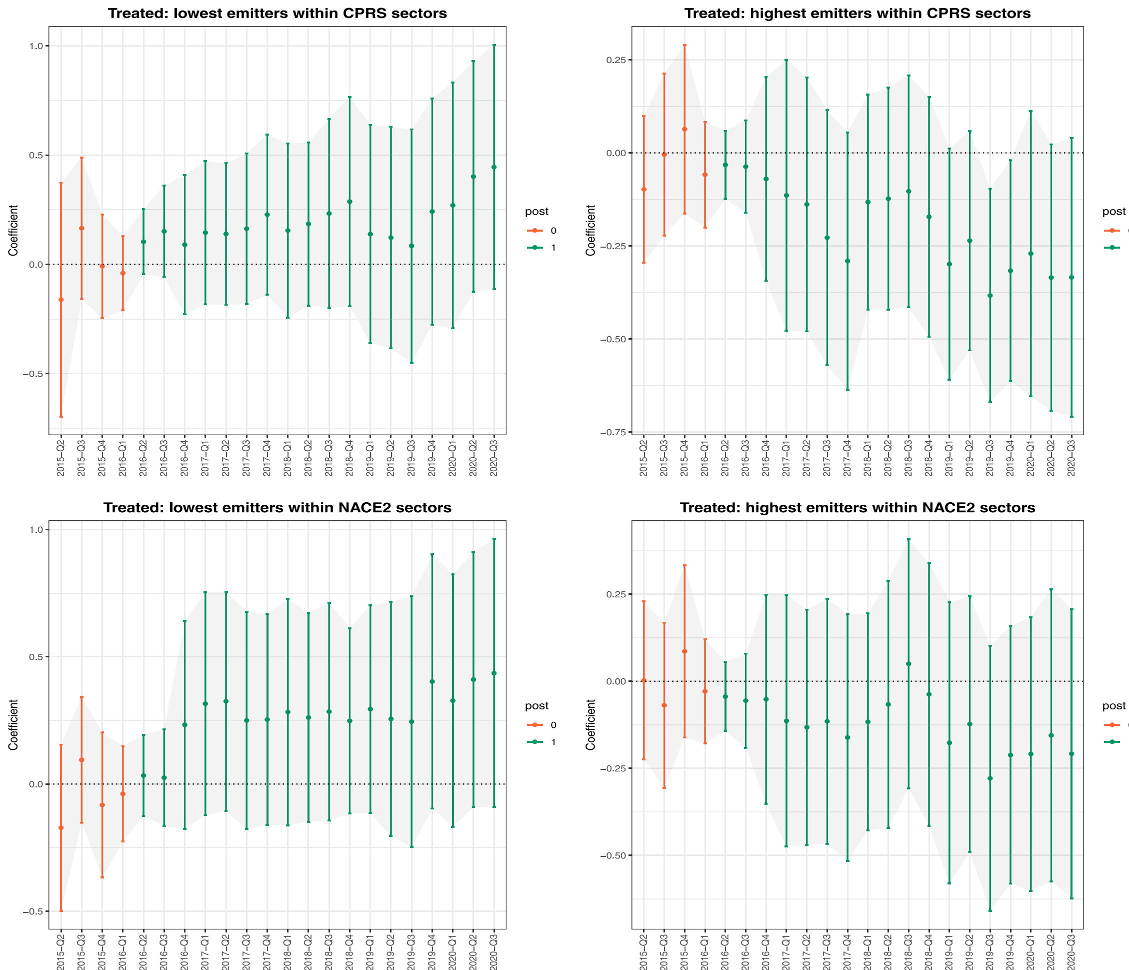

Fig. 24. Estimated ATTs for relatively low-/high-carbon companies within a sector. The figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively). The top and bottom rows are based on the classification of companies within the CPRS and NACE (two digit), respectively. The left and right columns of figures use the lowest and the highest thirds of emitters within a sector as the treatment group. The comparison group always remains the same, i.e., companies from the relatively unaffected sectors (by the Paris Agreement) also having the lowest emissions.

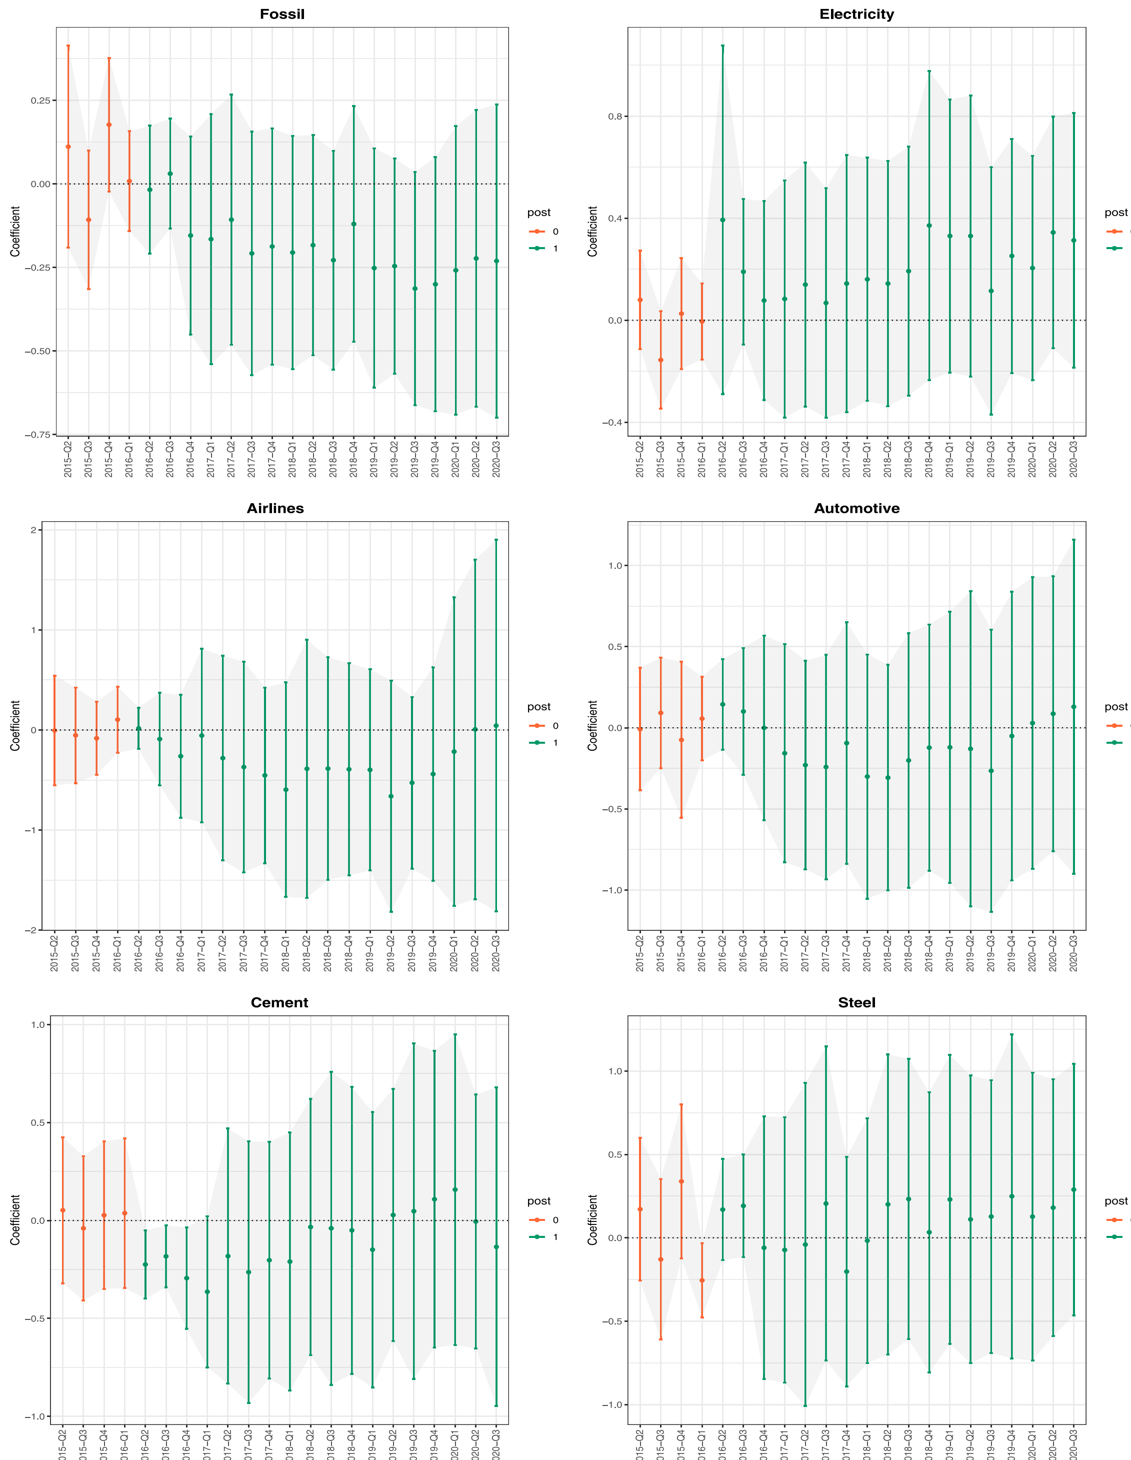

Fig. 25. Estimated ATTs by industry. The figure plots the average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

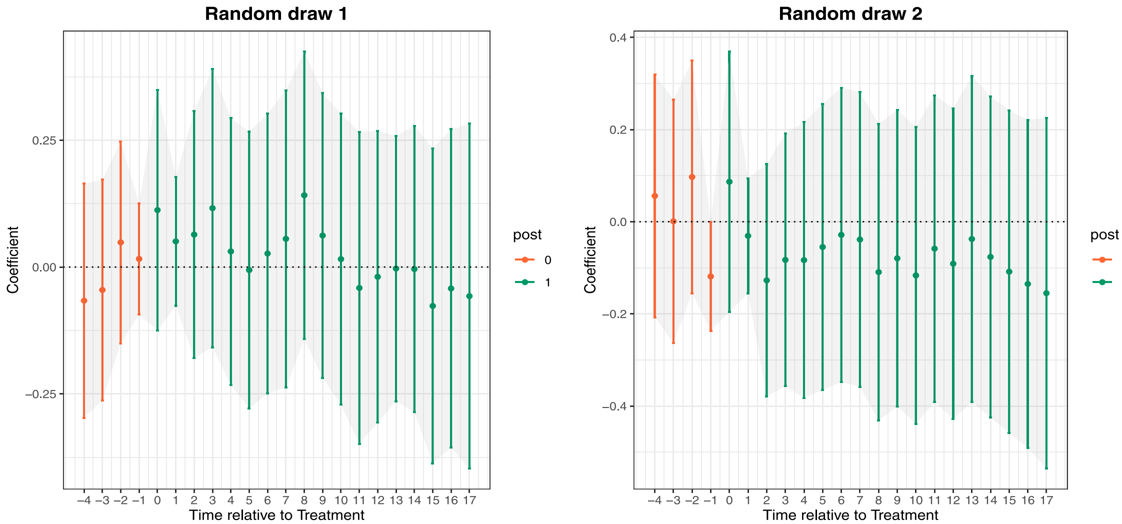

Fig. 26. Estimated ATTs with a random split of firms between treated and control groups. Both figures plot the estimated average treatment effect on the treated in terms of the (log-) participation with the ninety percent bootstrap confidence bands. The left and right panels just correspond to different random splits of companies into pseudo 'treated' and 'control' groups. The orange and green colors signify the pre- and post-Paris Agreement periods (also identified by 0 and 1, respectively).

References

Abadie, A., Alexis, D., Jens, H., 2010. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. J. Amer. Statist. Assoc. 105 (490), 493–505.

Alessi, L., Di Girolamo, F., Pagano, A., Petracco-Giudici, M., 2022. Accounting for Climate Transition Risk in Banks' Capital Requirements. European Commission - Joint Research Centre, Working papers in economics and finance 2022/8, JRC129221.

Alessi, L., Ossola, E., Panzica, R., 2021a. What greenium matters in the stock market? The role of greenhouse gas emissions and environmental disclosures. J. Financ. Stab. 54 (10086), 9.

Alessi, L., Ossola, E., Panzica, R., 2023. When Do Investors Go Green? Evidence from a Time-Varying Asset-Pricing Model. Int. Rev. Fin. Anal. 90, 102898.

Allevi, E., Basso, A., Bonenti, F., Oggioni, G., Riccardi, R., 2019. Measuring the environmental performance of green SRI funds: A DEA approach. Energy Econ. 79, 32–44.

Athey, S., 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. Am. J. Epidemiol. 172 (9), 1092–1097.

Baiardi, D., Morana, C., 2021. Climate change awareness: Empirical evidence for the European Union. Energy Econ. 96, 2021.

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

Battiston, S., Mandel, A., Monasterolo, I., Schütze, F., Visentin, G., 2017. A climate stress-test of the financial system. Nature Clim. Change 7 (4), 283–288.

Battiston, S., Monasterolo, I., Riahi, K., van Ruijven, B.J., 2021. Accounting for finance is key for climate mitigation pathways. Science 372 (6545), 918–920.

Benz, L., Jacob, A., Paulus, S., Wilkens, M., 2020. Herds on green meadows: The decarbonization of institutional portfolios. J. Asset Manag. 21 (1), 13–31.

Berg, T., Reisinger, M., Streitz, D., 2021. Spillover effects in empirical corporate finance. J. Financ. Econ. 142, 1109–1127.

Beyene, W., De Greiff, K., Delis, M.D., Ongena, S., 2021. Too-big-to-strand? Bond versus bank financing in the transition to a low-carbon economy.

Boermans, M.A., Galema, R., 2019. Are pension funds actively decarbonizing their portfolios? Ecol. Econom. 161, 50–60.

Boermans, M., Galema, R., 2020. Carbon home bias of European investors. Available at SSRN 3632723.

Bolton, P., Kacperczyk, M., 2021. Do investors care about carbon risk? J. Financ. Econ. 142 (2), 517–549.

Callaway, B., Sant'Anna, P.H.C., 2021. Difference-in-Differences with multiple time periods. J. Econometrics 225 (2), 200–230.

Dai, H.C., Zhang, H.B., 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. Adv. Clim. Change Res. 8 (4), 226–234.

Darvas, Z., Schoenmaker, D., 2017. Institutional Investors and Home Bias in Europe's Capital Market Union. Technical report.

de Chaisemartin, C., d'Haultfoeuille, X., 2020a. Difference-in-Differences Estimators of Intertemporal Treatment Effects. No. w29873, National Bureau of Economic Research, 2022.

de Chaisemartin, C., d'Haultfoeuille, X., 2020b. Two-way fixed effects estimators with heterogeneous treatment effects. Amer. Econ. Rev. 110 (9), 2964–2996.

de Chaisemartin, C., d'Haultfoeuille, X., 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.

Diamond, A., Sekhon, J.S., 2013. Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. Rev. Econ. Stat. 95 (3), 932–945.

Dunz, N., Naqvi, A., Monasterolo, I., 2021. Climate sentiments, transition risk, and financial stability in a stock-flow consistent model. J. Financ. Stab. 54, 100872.

Ellis, J., Moarif, S., 2017. Enhancing Transparency of Climate Finance under the Paris Agreement: Lessons from Experience. OECD/IEA Climate Change Expert Group Papers 2016/3. Technical report, OECD Publishing.

Fatica, S., Panzica, R., 2021. Green bonds as a tool against climate change? Bus. Strategy Environ. 30 (5), 2688–2701.

Ferrara, F.M., Sattler, T., 2018. The political economy of financial markets. Oxf. Res. Encycl. Politics 2018.

Ho, D.E., Imai, K., King, G., Stuart, E.A., 2011. `MatchIt`: Nonparametric preprocessing for parametric causal inference. J. Stat. Softw. 42, 8.

Iacus, S.M., King, G., Porro, G., 2009. `cem`: Software for coarsened exact matching. J. Stat. Softw. 30, 9.

Iacus, S.M., King, G., Porro, G., 2011. Multivariate matching methods that are monotonic imbalance bounding. J. Amer. Statist. Assoc. 106 (493), 345–361.

Iacus, S.M., King, G., Porro, G., 2012. Causal inference without balance checking: Coarsened exact matching. Political Anal. 20 (1), 1–24.

Iacus, S.M., King, G., Porro, G., 2019. A theory of statistical inference for matching methods in causal research. Political Anal. 27 (1), 46–68.

Law, C., Zhang, H., 2019. Implementing provisions on climate finance under the Paris agreement. Clim. Law 9 (1–2), 21–39.

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

Monasterolo, I., de Angelis, L., 2020. Blind to carbon risk? An analysis of stock market reaction to the Paris Agreement. Ecol. Econom. 170 (10657), 1.

NGFS, 2019. Network for Greening the Financial System. Report: "A Call for Action". Technical report.

OECD, 2020. Developing Sustainable Finance Definitions and Taxonomies, Green Finance and Investment. Technical report, OECD Publishing, Paris.

Pindyck, R.S., 2007. Uncertainty in environmental economics. Rev. Environ. Econ. Policy 1 (1), 45–65.

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

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

Reins, L., van Calster, G., 2021. Introduction – The Paris Agreement on climate change. In: Reins, L.S., van Calster, G. (Eds.), A Commentary on the Paris Agreement on Climate Change. Edward Elgar, Cheltenham.

Roncoroni, A., Battiston, S., Escobar-Farfán, L.O.L., Martinez-Jaramillo, S., 2021. Climate risk and financial stability in the network of banks and investment funds. J. Financ. Stab. 54, 10087.

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

Sekhon, J.S., 2011. Multivariate and propensity score matching software with automated balance optimization: The matching package for R. J. Stat. Softw. 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. Technical report, EU Diplomacy Paper 04/2018.

Stuart, E.A., Rubin, D.B., 2008. Best practices in quasi–experimental designs: Matching methods for causal inference. In: Osborne, J. (Ed.), Best Practices in Quantitative Methods. pp. 155–176, Ch. 11.

TEG, 2020. Taxonomy: Final Report of the Technical Expert Group on Sustainable Finance. Technical report.

Xu, Y., 2017. Generalized synthetic control method: Causal inference with interactive fixed effects models. Political Anal. 25 (1), 57–76.

Zhang, H.B., Dai, H.C., Lai, H.X., Wang, W.T., 2017a. US. withdrawal from the Paris Agreement: Reasons, impacts, and China's response. Adv. Clim. Change Res. 8 (4), 220–225.

Zhang, H.B., Dai, H.C., Lai, H.X., Wang, W.T., 2017b. The withdrawal of the US. From the Paris Agreement and its impact on global climate change governance. Adv. Clim. Change Res. 8 (4), 213–219.