Negative Emissions, Positive Returns? How Financial Markets React to Carbon Removal Purchases *

Francesco Lamperti[a,b,c], Massimo Tavoni[b,c,d], and Roberta Terranova[e,b,c]

[a] *Institute of Economics and EMbeDS, Scuola Superiore Sant'Anna, Pisa, Italy*
[b] *CMCC Foundation - Euro-Mediterranean Center on Climate Change, Milan, Italy*
[c] *RFF-CMCC European Institute on Economics and the Environment (EIEE), Milan, Italy*
[d] *Politecnico di Milano, Italy*
[e] *Université Côte D'Azur, CNRS, GREDEG, Nice, France*

January 2026

Abstract

Carbon removal is becoming central to firms' net-zero strategies, yet little is known about how investors perceive its use. We examine financial market reactions to publicly traded firms' purchases of carbon dioxide removal credits. We compile a novel transaction-level dataset of carbon removal purchases and link it to daily stock returns. Using short-horizon event studies, we find that, on average, markets respond negatively to removal purchases, with significantly negative cumulative abnormal returns over event windows from $[-1, +1]$ to $[-1, +7]$. This average masks substantial heterogeneity: negative reactions are stronger for firms with weaker climate-policy quality, in certain sectors—especially energy—and for purchases of novel technical removals relative to nature-based removals. Cross-sectional regressions show that higher per-ton prices are associated with more negative abnormal returns, while larger purchases relative to firm emissions mitigate these effects. Overall, investors appear skeptical of firms' reliance on carbon removal to manage their emissions.

Keywords: Climate finance; carbon dioxide removal credits; cumulative abnormal returns

*Email for correspondence: roberta.terranova@univ-cotedazur.fr. This work has received funding from the European Research Council (ERC) FIND project (Finance and Innovation to couple Negative emissions and sustainable Development) under the Horizon Europe research and innovation programme (grant agreement No. 101117427).

1 Introduction

Limiting global warming to well below 2°C requires not only rapid emissions reductions but, in most mitigation pathways, large-scale deployment of carbon dioxide removal (CDR) to neutralize residual emissions and partially offset legacy emissions (Intergovernmental Panel on Climate Change (IPCC), 2022; National Academies of Sciences, Engineering, and Medicine, 2019). Firms sit at the center of this transition. Beyond regulatory compliance, stakeholders increasingly expect companies to implement climate strategies that are credible, measurable, and aligned with stated net-zero objectives.

A growing literature shows that climate risks are priced in equity markets and that investor attention to climate-related information reallocates valuations across firms (Bolton and Kacperczyk, 2021, 2023; Ardia et al., 2023). Markets also respond to climate-policy milestones and narratives: the Paris Agreement is associated with positive abnormal returns for firms with greater exposure to green opportunities (Kruse et al., 2024), low-carbon assets become more attractive in portfolio choice after major climate-policy announcements (Monasterolo and de Angelis, 2020), and climate-related central bank communication can differentially shift green versus brown equity performance (Campiglio et al., 2025). Taken together, this evidence aligns with a risk-management perspective on climate finance: investors hedge transition risk, demand compensation for bearing it, and increasingly pressure firms to decarbonize (Andersson et al., 2016; Bolton et al., 2022a,b). As a result, pursuing specific climate strategies can affect firms' financial valuations and access to external finance. This shifts the decarbonization problem from a purely technological and regulatory challenge to a strategic and financial one.

The scale-up of CDR implied by climate scenarios is substantial: in "more sustainable" scenario families, global CDR reaches roughly 7–9 GtCO2/year by mid-century, yet current deployment remains far below these levels and is concentrated in conventional methods, with novel CDR contributing only a small share despite rapid growth (Smith et al., 2024). This widening gap between model-consistent needs and observed deployment has catalyzed investment and experimentation in CDR technologies, procurement mechanisms, and supporting market infrastructure, while also intensifying debate over feasibility, cost, integrity, and the credibility of associated claims (McKinsey, 2023; Hickey et al., 2023; Michaelowa et al., 2023). Compared with conventional offsets, CDR methods vary sharply in technological maturity, monitoring and verification challenges, durability of storage, scalability, and governance requirements (Intergovernmental Panel on Climate Change (IPCC), 2022; National Academies of Sciences, Engineering, and Medicine, 2019; Hickey et al., 2023).

CDR purchases can be framed as a way to address residual emissions and to operationalize net-zero commitments, but they have also raised concerns about substituting for internal decarbonization. Recent evidence suggests that voluntary offsetting is not central for most multinational

firms and is often small relative to capital expenditures, though it can compete with internal decarbonization for large-scale offsetters in some sectors (Stolz and Probst, 2025). As a result, the same observed purchase may be interpreted by investors either as credible transition-risk mitigation and strategic positioning, or as an expensive and potentially low-credibility signal.

Against this backdrop, a central question is how investors interpret corporate climate actions when technologies and standards are heterogeneous and still evolving. This paper studies how financial markets respond to firms' CDR purchases. We focus on the short-run stock-market reaction because it provides a disciplined way to measure investors' immediate reassessment of firm value when information about a firm's climate actions becomes salient (Brown and Warner, 1985; MacKinlay, 1997; Kothari and Warner, 2007). In this setting, the sign of the response is theoretically ambiguous. A positive reaction is consistent with investors viewing CDR purchases as strengthening stakeholder legitimacy, improving the firm's transition-risk profile, or positioning the firm for profitable growth opportunities in emerging low-carbon markets (Kruse et al., 2024). A negative reaction is consistent with investors interpreting the purchase as a value-reducing expenditure, an indicator of high marginal abatement costs, or a reputational-risk signal in a context where claims, verification, and standards remain contested (Michaelowa et al., 2023). Because equity valuations affect financing capacity, these reactions can have real consequences: higher valuations can relax financing constraints and reduce the cost of capital, while lower valuations can tighten constraints and raise the cost of external finance (Baker et al., 2003). These mechanisms are consistent with growing evidence that firms' climate strategies respond to their exposure to climate risk and investor pressures (Li and Flammer, 2025b,a).

Our empirical analysis is organized around three testable hypotheses that map directly to these channels. First, if markets systematically reward or penalize CDR purchases on net, we should observe abnormal returns around purchase disclosure that differ from zero:

> *H1 (Financial response to CDR purchases).* CDR purchases generate abnormal stock returns around the event that are systematically different from zero.

Second, because CDR credits vary substantially by durability, quality, level of consolidation – and because buyers differ in the credibility of their broader climate strategies – investors differentiate purchases through observable credit characteristics:

> *H2 (Credibility-screening of CDR purchases).* Investors condition their reaction on observable credibility and uncertainty signals. Abnormal returns are lower for purchases made by weaker climate-policy-quality buyers and for purchases involving worse removal projects and less established removal methods.

Third, conditional on transaction attributes, purchase size relative to the buyer's emissions baseline may proxy for salience and commitment: larger, more material purchases may be interpreted as stronger strategic actions rather than symbolic gestures, mitigating negative valuation effects:

H3 (Salience of the purchase). Conditional on transaction attributes, purchases that are larger relative to the buyer's emissions baseline are more salient commitments and are therefore associated with better abnormal performance.

Empirically, existing evidence on investor responses to corporate climate actions is still limited and often organized around broad policy events or specific financial instruments. Markets respond to climate-policy milestones that shift expectations about green opportunities (Kruse et al., 2024; Monasterolo and de Angelis, 2020), and investors react positively to credible, third-party-verified green finance instruments, consistent with a signaling channel (Flammer, 2021). Yet market responses to CDR purchases have not been studied so far to the best of our knowledge. CDR markets and integrity standards are evolving rapidly, which raises the possibility that investors condition their reactions on who purchases, what is purchased, and how the purchase is intermediated, priced, and verified (Hickey et al., 2023; Dörpmund, 2025; Michaelowa et al., 2023).

To address these gaps, we build a new event-level dataset that links carbon removal purchase transactions to daily equity returns and standard asset-pricing factors in a multi-country setting. We integrate two complementary sources of CDR purchase records, map disclosed buyers to public-equity tickers, and merge the resulting event dataset to firm-day return panels and factor returns. We then implement standard short-horizon event study methods with a pre-event estimation window designed to limit contamination, and we report robustness to alternative expected-return models and inference procedures commonly used when events cluster or when estimation windows overlap (Fama and French, 1993; Aktas et al., 2007; Boehmer et al., 1991; Kolari and Pynnönen, 2010). We align transactions to trading days by setting $\tau = 0$ to the transaction date when it is a trading day for the buyer's stock, and otherwise mapping the event to the first subsequent trading day. This design also connects directly to the literature on firm climate strategy by treating CDR purchases as observable, discretionary actions through which firms manage (or claim to manage) transition risks, stakeholder demands, and decarbonization narratives (Li and Flammer, 2025a).

Our results show that, consistent with *H1*, CDR purchases are associated with statistically significant abnormal stock performance in the days surrounding the transaction. On average, public buyers experience negative cumulative abnormal returns, with the largest adjustment occurring within the week around the purchase. Consistent with *H2*, this average response masks pronounced heterogeneity: market reactions are materially weaker for purchases by lower climate-policy-quality buyers, as well as for transactions linked to lower-quality projects and more uncertain removal methods. Finally, in line with *H3*, we find that transaction salience matters: purchases that are larger relative to the buyer's emissions baseline are associated with stronger subsequent abnormal performance, partially offsetting the weaker reactions observed for less salient transactions. Overall, the evidence indicates that investors, on average, penalise CDR purchases. Further, they condition on both buyer credibility and what is purchased, distinguishing between CDR purchases that plausibly complement a credible transition strategy and purchases that are more likely to be interpreted

as costly or low-credibility signals.

The remainder of the paper proceeds as follows. Section 2 describes the construction of the transaction and financial datasets and the mapping from purchases to publicly traded buyers. Section 3 outlines the event study design. Section 4 presents the baseline market reaction and heterogeneity analyses. Section 5 illustrates the findings about what drives the financial market response to CDR purchases, and the Appendix reports robustness checks and supplementary tests.

2 Data construction and description

Our unit of analysis is a carbon removal purchase event. The empirical analysis requires a dataset that links purchase events to daily equity returns and standard asset-pricing factors. We therefore integrate two transaction-level sources on carbon dioxide removal purchases with public-equity price data, exchange rates, and factor returns, and we harmonize identifiers so that each event can be mapped to a publicly traded firm when applicable.

2.1 Transactions

We assemble purchase records from two complementary sources. The first source is AlliedOffsets, which reports information on buyers, transaction dates, purchased volumes, and project attributes. The second source is CDR.fyi, which reports purchaser-level orders for carbon dioxide removal credits and associated order characteristics. As of April 7, 2025, the AlliedOffsets extract contains 47,454 transactions. We exclude 19 transactions classified under the project sector *Carbon Capture* and 5 transactions with credit status *CDR MoU* (memorandum of understanding), yielding 47,430 AlliedOffsets transactions. Among these, 16,608 transactions feature disclosed purchasers, corresponding to 3,090 unique buyers. The CDR.fyi extract contains 5,020 transactions, of which 1,842 are associated with disclosed purchasers, corresponding to 467 unique buyers.

Because the two sources use different naming conventions, we harmonize buyer identities through a sequential matching procedure. We first match entities using website URLs when available, then match on standardized buyer names, and finally resolve remaining ambiguous cases through manual verification. This procedure identifies 295 buyers that appear in both sources and yields a combined universe of 3,262 unique buyers.

We then merge the two transaction files and remove duplicates to prevent double-counting. The merged file initially contains 52,450 transactions. We flag duplicates based on a concordance of buyer identity, transaction date, and purchased volume (tons) and remove 131 duplicate entries.[1] The resulting consolidated dataset includes 52,319 transactions, of which 18,317 are associated with disclosed purchasers. Figure 1 illustrates the evolution over time of the number of transactions in

[1] Transactions with missing credit size are not treated as duplicates, even when other attributes coincide, because volume is a key discriminator for identifying the same economic event.

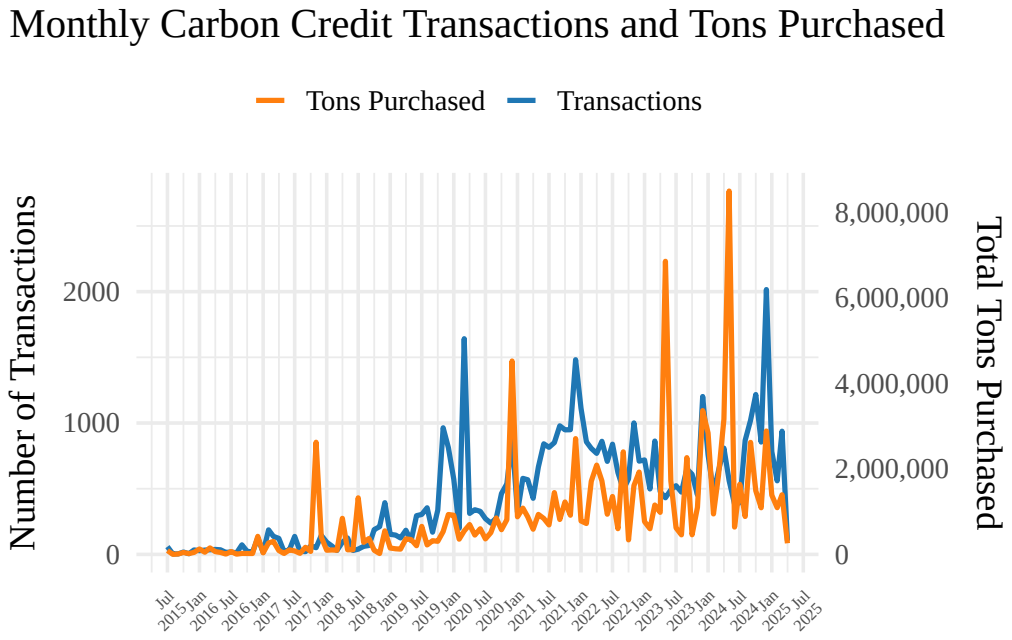


Figure 1: Evolution over time of number of transactions and tons of CO2 purchased. Monthly frequency.

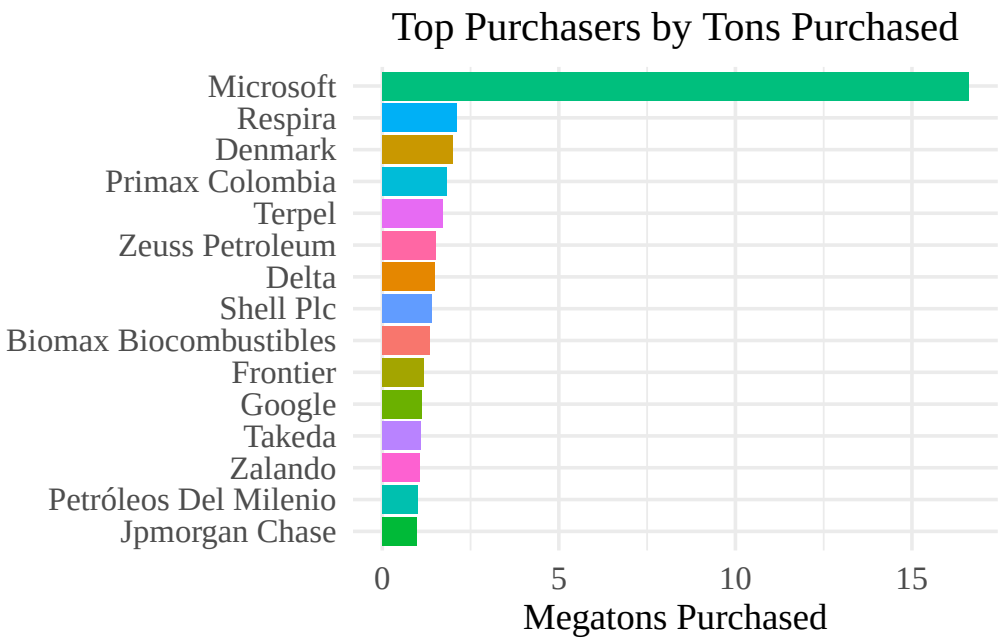


Figure 2: Top purchasers of carbon removal credits by tons purchased.

our dataset (blue line) and the total tons of carbon removals purchased (orange line). Despite fluctuations, the two series have a positive trend, indicating the rising importance of the carbon removal market. Figure 2 gives an overview of who the top purchasers of carbon removal credits are, by ton purchased. The top purchasers come from very different sectors, from Microsoft, which is by far the company buying the most credits, to public entities such as Denmark.

For the return-based analysis, we identify which disclosed buyers are publicly listed and map those entities to stock tickers used by Yahoo Finance. To classify buyers as public or private, we follow source-specific procedures that reflect differences in available metadata. For AlliedOffsets, a buyer is classified as public when an ISIN code or a stock symbol is reported. For CDR.fyi, we implement a multi-stage classification pipeline. We begin with automated queries to Yahoo Finance using the buyer name, and when a candidate match is returned, we extract the ticker and the website associated with the listing. We classify the buyer as public when a ticker is available and the website aligns with the website recorded in CDR.fyi. We then validate and enrich these mappings using automated checks against Wikipedia to confirm the corporate identity and listing status. For buyers not classified through Yahoo Finance or Wikipedia, we use automated searches of Wikipedia and Crunchbase to retrieve additional metadata, including the corporate website and IPO status, and we accept matches when the name and website are consistent with the CDR.fyi record. Finally, we manually review remaining unmatched or ambiguous cases using corporate websites, filings, and

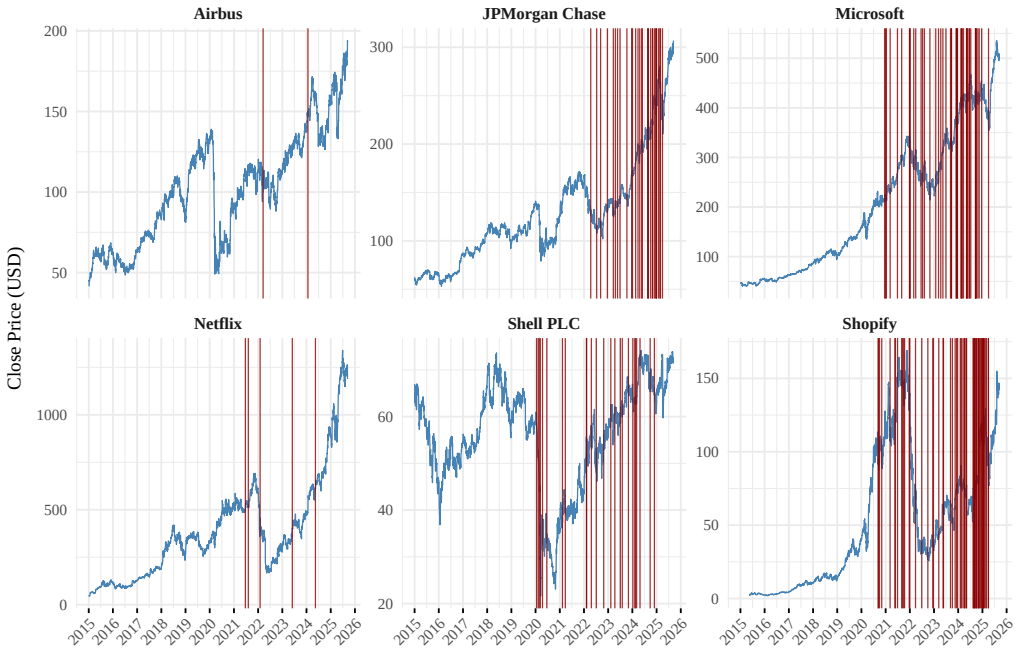


Figure 3: Evolution over time of the stock price (blue line). Red vertical lines represent transactions happening on that day.

other official sources. This process produces a curated mapping from disclosed buyers to public tickers, together with documentation of match quality and the information sources used.

For each purchase, we construct a harmonized set of event-level variables that includes the event date, a canonical buyer identifier, the mapped ticker when available, and purchase characteristics such as volume and price. The sample spans 2010–2025 and includes all publicly traded firms that execute at least one purchase during the sample period.

2.2 Financial returns and factors

To construct daily equity returns, we retrieve historical price series from Yahoo Finance for all mapped tickers. Because firms are listed across multiple currencies, we express returns in a common currency by converting prices to U.S. dollars using daily exchange rates. This yields daily close and adjusted close prices in U.S. dollars and the corresponding return series for each firm.

We then merge asset-pricing factors required to model expected returns. We download daily factor returns from the Fama–French data library and harmonize them across regional groupings

and the United States.[2] To capture systematic variation in pricing related to carbon intensity, we additionally construct a carbon-intensity factor that summarizes the return differential between high- and low-emissions-intensity firms and merge this series into the firm–day panel. The resulting dataset is a daily panel containing firm returns in U.S. dollars, factor returns, and the risk-free rate used to compute excess returns, together with an event file that maps each CDR purchase to the corresponding firm and trading day.

Figure 3 illustrates the evolution over time of the stock price for six public companies included in our dataset: Airbus, JPMorgan Chase, Microsoft, Netflix, Shell, and Shopify. These are among the most active buyers in our dataset. The red vertical lines represent the timing of their purchases of carbon removal credits. For some of them, the frequency increases with time.

2.3 Firm level data

Beyond the information included in the original datasets by AlliedOffsets and CDR.fyi, we need firm-level data, such as CO2 emissions, firms' total assets, and revenues, among others, to study the determinants of the cumulative abnormal returns associated with the purchases of carbon removal credits. We download these data from Refinitiv for the companies included in our dataset of public buyers. The variables we download are: total assets, cash, revenues, emissions (all three scopes), leverage, and profitability variables. We collect firm-level information on 256 firms.[3] Table 1 illustrates the descriptive statistics of the firm-level information we use in the analysis.

Table 1: Descriptive statistics: firm-level fundamentals

variable	N	mean	median	sd	p25	p75
Total assets	378	189066.94	25090.42	510279.77	5448.32	129600.00
Cash	407	15061907507.44	1568806727.07	49955223129.34	285202111.64	7848268302.17
Revenues	404	35292058200.56	10025767603.57	74800092462.64	2298688000.00	32890250000.00
Emissions	302	78965475.99	8128939.98	201503936.91	1282790.66	52545360.78

Notes: The table reports the number of non-missing observations (N), mean, median, standard deviation (SD), and the 25th/75th percentiles (P25/P75). Variables are in original units as recorded in the dataset.

3 Empirical strategy: event study design

We adopt standard short-horizon event study methods to measure how financial markets respond when a publicly listed firm purchases carbon removal credits. Event studies are well-suited to our setting because they isolate the component of realized returns that is not explained by systematic

[2] `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

[3] One of the next steps of our research is to extend the coverage of firms for which we collect these data.

risk, thereby capturing the market's immediate reassessment of firm value around the purchase date (Brown and Warner, 1985; MacKinlay, 1997; Kothari and Warner, 2007).

3.1 Disclosure, registries, and information timing

A key requirement for our event-study design is that transaction information is publicly observable around the recorded purchase date. For a subset of purchases, CDR.fyi reports an "announcement" URL (528 transactions in our sample). These links are publicly accessible and provide an external timestamp at which information about the transaction becomes available. In many cases, the URL redirects to an intermediary—typically a marketplace or registry interface—that displays transaction-level details (buyer name, project, credit type, volume, and delivery/retirement information). Figure 4 illustrates a representative example from the Isometric registry.[4]

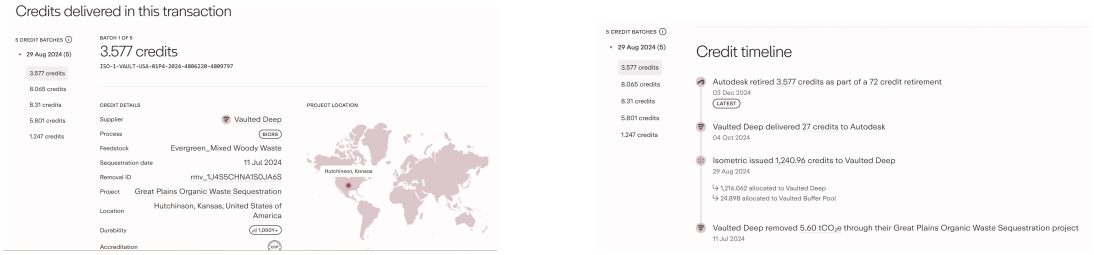


Figure 4: Example of transaction-level information displayed on a public intermediary registry (Isometric).

Beyond the transaction-specific links, we systematically verify the public registries of all intermediaries represented in our dataset. These registries typically disclose the buyer, the project and credit attributes, and the relevant delivery or retirement date, and they are updated frequently. To illustrate the refresh frequency close to the end of our sample, as of January 19, 2026, Isometric reports transactions through January 17; Gold Standard reports transactions through January 18;[5] Puro.earth through January 15;[6] Riverse through January 19;[7] ACR through January 16;[8] Climate Action Reserve through January 16;[9] and the Woodland Carbon Code registry reports its most recent update (for an undelivered credit) on January 16.[10] Together, these checks support the view that core transaction attributes are publicly available through intermediary registries and can be observed at a relatively high frequency around the transaction date.

[4] https://registry.isometric.com/delivery/dlv_1J9CFM6A91S012PX
[5] https://registry.goldstandard.org/credit-blocks
[6] https://registry.puro.earth/retirements
[7] https://registry.rainbowstandard.io/ledger/retired-credits
[8] https://acr2.apx.com/myModule/rpt/myrpt.asp?r=206
[9] https://thereserve2.apx.com/myModule/rpt/myrpt.asp
[10] https://registry.spglobal.com/uklandcarbonregistry/public/wcc

3.2 Event-time alignment and abnormal return construction

Let e index purchase events and let $i(e)$ denote the buyer firm associated with event e. Transaction dates are recorded in calendar time. We align each event to the firm's trading calendar by setting the event day $\tau = 0$ to the transaction date when it is a trading day for the buyer's stock; when it falls on a non-trading day (e.g., weekend or holiday), we map the event to the first subsequent trading day. We then define event time τ in trading days relative to $\tau = 0$.

For each event e, we compute abnormal returns by subtracting a model-implied "normal" return from the realized excess return. We then aggregate abnormal returns across events to obtain average abnormal returns by relative day and sum them to obtain the cumulative average abnormal return (CAAR) over an event window $[\tau_1, \tau_2]$:

$$CAAR[\tau_1, \tau_2] = \sum_{\tau=\tau_1}^{\tau_2} \frac{1}{N_\tau} \sum_{e=1}^{N_\tau} AR_e(\tau), \tag{1}$$

where $\frac{1}{N_\tau} \sum_{e=1}^{N_\tau} AR_e(\tau)$ is the average abnormal return (AAR) at relative day τ, and N_τ denotes the number of events with non-missing abnormal returns at relative day τ.

Normal returns are obtained from a five-factor model (Fama and French, 1993) estimated in a pre-event window that runs from 200 to 41 trading days prior to the event (i.e., $[-200, -41]$). This choice reduces the risk that parameter estimates reflect information related to the purchase itself (Brown and Warner, 1985; MacKinlay, 1997). The market factor is the local market return for the buyer's listing country. This is natural in a multi-country setting where systematic risk is priced relative to country benchmarks and return comovement varies across markets.

We examine five short-run windows around the purchase date:

$$[-1, +1], \quad [-1, +3], \quad [-1, +5], \quad [-1, +7], \quad [-1, +9].$$

These windows serve two purposes. First, allowing $\tau = -1$ accommodates modest information leakage or anticipatory trading. Second, extending the post-event horizon allows us to study how market reaction evolves over time.

We winsorize abnormal returns at the 1st and 99th percentiles. The Appendix reports robustness to alternative winsorization thresholds. We also verify robustness to alternative expected return models, including multi-factor specifications (e.g., CAPM, market model, and three-factor models) and extensions incorporating a carbon-related factor. Finally, in addition to the full event sample, we analyse a "clean estimation window" subsample that excludes events whose $[-200, -41]$ estimation window overlaps with another purchase event by the same firm.

We test the null hypothesis that carbon removal credit purchases do not generate abnormal

performance:

$$H_0 : \ CAAR[\tau_1, \tau_2] = 0.$$

Because event clustering and event-induced variance can affect classical inference in some settings, the Appendix also reports standardized cross-sectional tests and correlation-robust adjustments commonly used in event studies (Boehmer et al., 1991; Kolari and Pynnönen, 2010).

3.3 Heterogeneity

Beyond the pooled average, our economic question is when and for whom carbon-credit purchases affect firm value. Purchases differ along multiple dimensions that shape credibility and informational salience to investors. For example, higher-quality projects or more established removal methods may be viewed as more reliable inputs into a firm's compliance and reputational risk management, whereas unusually high transaction prices may be interpreted as a costly signal (or, alternatively, as reflecting scarcity or higher-quality supply). These considerations motivate a structured heterogeneity analysis.

We therefore examine whether abnormal returns vary systematically with (i) who purchases credits, (ii) what type of carbon removal is purchased, and (iii) how the transaction is intermediated and priced. For each subgroup g, we compute subgroup-specific $CAAR_g[\tau_1, \tau_2]$ using (1) and test $H_0 : CAAR_g[\tau_1, \tau_2] = 0$.

Who purchases credits. We begin by characterising transactions by observable buyer characteristics. First, we study heterogeneity across buyer sectors. The raw sector labels in the underlying data are harmonized into GICS sectors.[11] We then compare CAARs across sectors to assess whether the market response differs systematically across types of firms.

Second, we explore whether investors react differently depending on the buyer's climate-policy posture. To do so, we construct a buyer-quality index summarizing third-party information on firms' climate commitments and target-setting practices. The index captures two dimensions: the firm's *governance and commitments*(reflecting the existence and status of formal commitments) and *ambition and timing* (reflecting target stringency and target horizons). We identify low and high buyer-quality groups using the sample median of the index.[12]

What type of carbon removal is purchased. Each purchase is linked to a specific removal project, which varies in technology, country, and verification pathways. We first examine heterogeneity across removal methods. We map project-level method labels into a set of standardized method categories (e.g., afforestation/reforestation/revegetation, biochar, direct air capture with

[11] We group buyers into a common set of sector categories (e.g., Industrials, Consumer Staples, Consumer Discretionary, Information Technology, Energy, Communication Services, among others).

[12] Full definitions and scoring rules are provided in Appendix A.1.

storage, bioenergy with carbon capture and storage, enhanced weathering, and direct ocean removal), and estimate CAARs separately for each category.

We also consider a coarser taxonomy that distinguishes *nature-based* from *technical* removals. Nature-based removals include ARR and soil carbon sequestration; all other methods are classified as technical removals.[13]

Beyond removal type, we study heterogeneity by project quality. We build a project-quality index that aggregates available third-party assessments and registry information capturing both credibility and maturity. Intuitively, the index rewards projects that are more likely to satisfy integrity benchmarks, that receive affirmative assessment signals when available, and that have progressed further in the registry lifecycle (from early-stage registration and validation to verification, crediting/issuance, and retirement). We compute the index as the average of available components and split projects into *low* and *high* quality groups using the sample median.[14]

Transaction features. Finally, we evaluate whether market reactions vary with observable transaction characteristics. First, we stratify purchases by the intermediary or marketplace through which the transaction is executed and estimate CAARs by venue. This analysis assesses whether the market response depends on the type of platform facilitating the purchase.

Second, we study heterogeneity by transaction size. We measure size as the volume of credits (in tons) purchased in the transaction and classify transactions as *small* or *large* using the 70th percentile of the transaction-volume distribution in the full sample: purchases at or below this threshold are labeled small, and those above are labeled large.

Third, we examine heterogeneity by transaction price. We construct a per-ton transaction price and assign transactions to *low*, *mid*, and *high* price groups based on terciles of the observed price distribution. We then estimate CAARs within each tercile to evaluate whether the market response varies with the implied cost of the purchased credits.

4 Market reaction

This section reports the cumulative average abnormal returns (CAARs) around public companies' purchases of carbon removal credits, for the full sample and for key subsamples.

Figure 5 summarizes the average stock-market response to these transactions. Across all buyers and all observed transactions, the estimated CAAR turns negative starting from the first time interval, although the immediate effect is not significantly different from zero. The magnitude of the decline increases and becomes significantly negative over subsequent days. It reaches its largest

[13] The mapping from raw method labels to standardized categories is documented in Appendix A.1.

[14] Appendix A.1 reports component definitions, the scoring rubric (including the registry-status mapping), and diagnostics on the number of non-missing components used in the index.

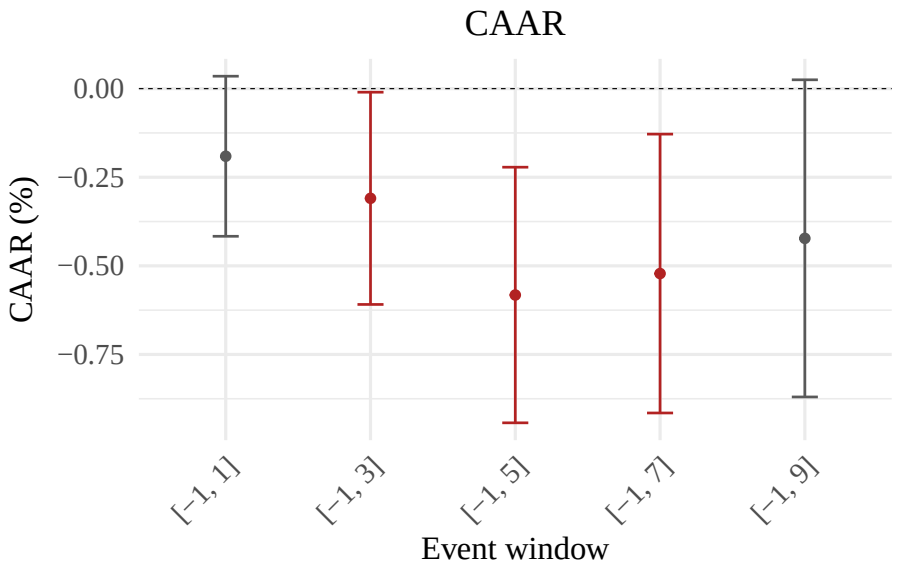


Figure 5: CAAR over different time windows

value in the $[-1, +5]$ window: the CAAR is lower than -0.5%. This estimate implies that, over the five trading days surrounding the transaction, buyers underperform by about 0.5 percentage points relative to the counterfactual performance predicted by the return model absent the event.

The negative reaction remains statistically significant in the $[-1, +7]$ window, but attenuates thereafter; by $[-1, +9]$ the CAAR is no longer statistically different from zero. Appendix D reports a battery of robustness checks, which corroborate the baseline finding of a negative average market reaction, with the clearest evidence concentrated in intermediate windows – particularly $[-1, +5]$ and $[-1, +7]$.

Our estimated CAAR aligns with the documented market reaction to various corporate actions. In particular, sustainability index inclusion (DJSI) is associated with announcement-window losses of about -0.5% on $[0, 0]$ and approximately -1.8% to -2.0% over $[0, 5]$ (Oberndorfer et al., 2013). The magnitude is also comparable to the negative abnormal returns reported for standard financing events: MacKinlay (1997) summarizes the literature showing that straight debt issues average roughly -0.23% over a two-day window.

By contrast, other climate-finance events often generate positive valuation effects of a similar order of magnitude. For example, corporate green bond issuance is associated with a positive market response of about $+0.49\%$ over $[-5, +10]$ in Flammer (2021), and matched comparisons of green versus conventional bonds yield differentials around $+1.08\%$ over $[-2, +2]$ in Schneider (2024).

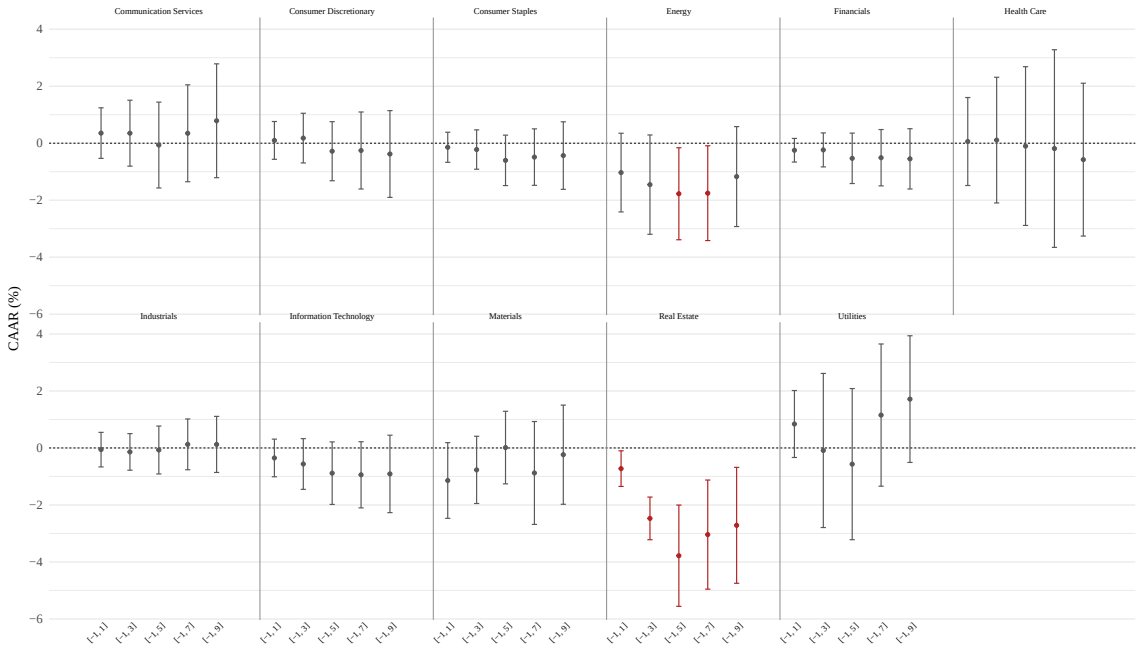


Figure 6: CAAR by buyer's sector

In oil and gas acquisitions, Sabet and Heaney (2016) document positive announcement effects for both reserve deals (+0.99% on $[-1, +1]$) and acreage deals (+0.35%), whereas bidder returns in broader M&A settings are typically mildly negative (about −1.2% to −0.7%) (Martynova and Renneboog, 2008). In sharp contrast, targets experience much larger gains: the literature summarized by MacKinlay (1997) often reports announcement-period target returns exceeding 20%, and Martynova and Renneboog (2008) reports CAARs for targets in the range of 14% to 44% over an approximately two-week horizon.[15]

Taken together, these benchmarks place our average CAAR of roughly −0.5% over $[-1, 5]$ in the lower-to-moderate portion of these event-study effects: it is far smaller than valuation changes typical of corporate control events, yet comparable to a set of announcement types for which the literature finds economically small but statistically meaningful market responses.

4.1 Heterogeneity analysis

We now turn to the heterogeneity results. Figure 6 reports CAARs by buyer sector across the event windows. For most sectors, the estimated responses are economically small and statistically

[15] See Table 6 in the Appendix for details.

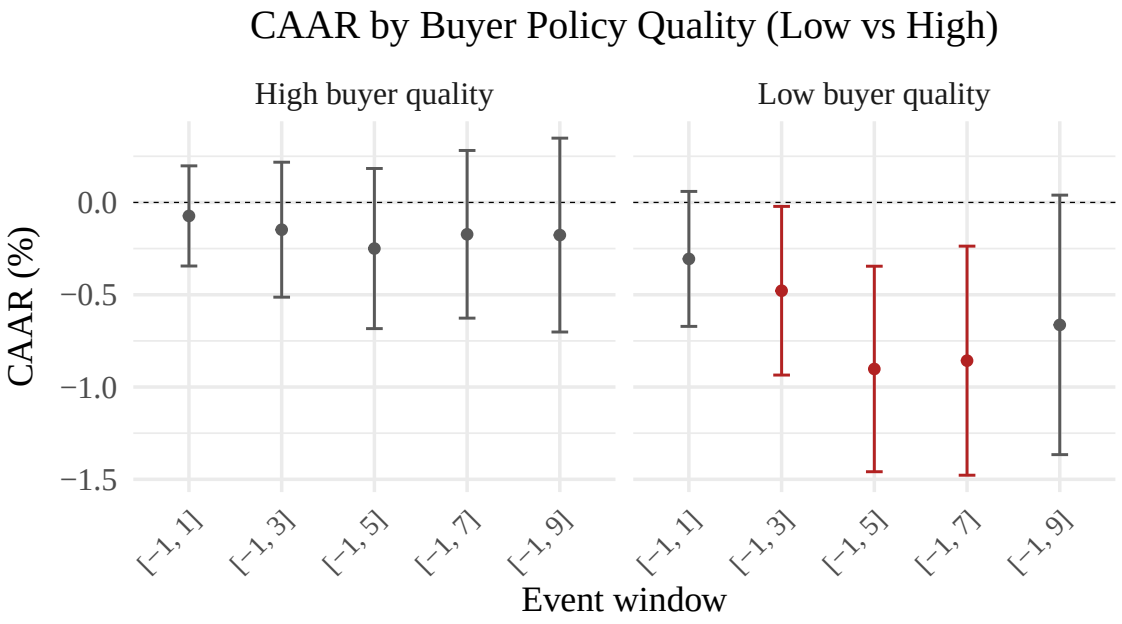


Figure 7: CAAR by buyer's quality

indistinguishable from zero. Two sectors stand out. In real estate, CAARs are negative across all windows considered, with the largest decline approaching 4% in the $[-1, 5]$ window. In other words, for real-estate buyers, credit purchases are followed by a sizable short-horizon underperformance relative to the counterfactual return absent the transaction. We also observe a negative response in the energy sector, where CAARs reach roughly -2% in selected windows.

Buyer climate-policy quality is an equally salient dimension. Figure 7 contrasts transactions by the buyer-quality index. Purchases by high-quality buyers are not associated with a statistically significant market reaction. By contrast, purchases by low-quality buyers generate a negative and statistically significant response from $[-1, 3]$ through $[-1, 7]$, with CAARs reaching approximately -1%. This pattern suggests that markets condition their interpretation of credit purchases on the buyer's broader climate strategy: purchases by firms with weaker commitments and less ambitious targets are more likely to be viewed unfavourably. We next examine heterogeneity in what is purchased. Figure 8 relate market reactions to the underlying removal method.[16] The broader comparison between nature-based and technical removals yields a clearer pattern. Transactions involving technical removals are followed by more negative CAARs, reaching about -1.5% in the $[-1, 5]$ window, whereas nature-based removals exhibit weaker responses. One interpretation is

[16] Here we show the CAARs for the nature versus technical removal methods, whereas in Figure 15 in the Appendix **??** we show the CAARs related to the specific removal methods.

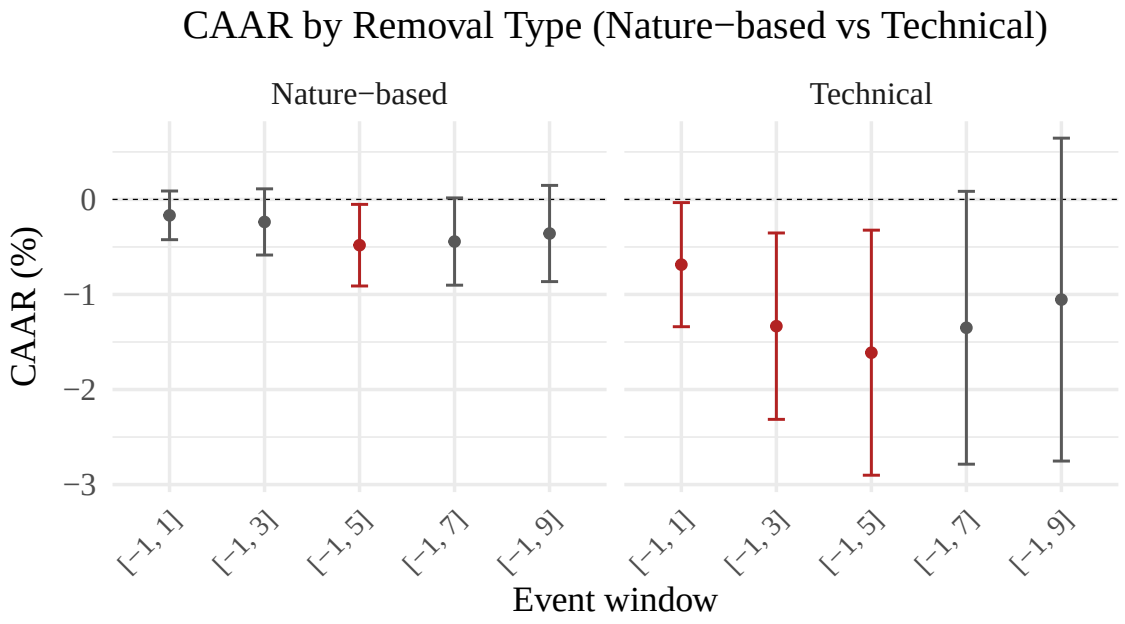


Figure 8: CAAR by removal type (nature vs technical)

that investors respond less favourably when purchased credits rely on methods that are newer or perceived as more uncertain in terms of verification, durability, or scalability. Project quality further differentiates market responses. Figure 9 shows that purchases linked to high-quality projects display a small and statistically insignificant negative reaction. In contrast, purchases associated with lower-quality projects are followed by more negative CAARs. This evidence is consistent with markets discounting transactions that are tied to projects with weaker third-party signals of credibility or lower registry maturity.

Finally, we study heterogeneity in transaction characteristics. Figure 10 reports CAARs by intermediary. Only one, Gold Standard Marketplace, is associated with positive CAARs in some windows, while significantly negative are reported for Supercritical in the $[-1, 9]$ window. Transaction size also matters. Figure 11 shows that smaller purchases are followed by significantly negative CAARs from $[-1, 1]$ through $[-1, 9]$, whereas larger purchases elicit weaker and insignificant responses. Figure 12 reports results by transaction price. The pattern is more mixed, but high price transactions display the most negative response, particularly in the mid window.

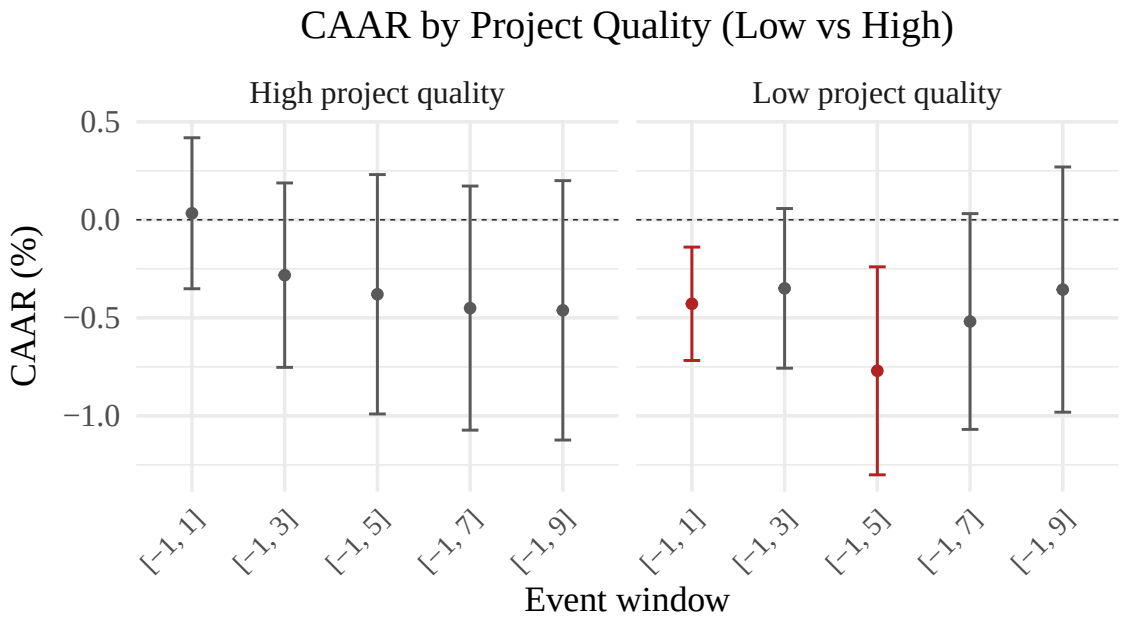


Figure 9: CAAR by project quality

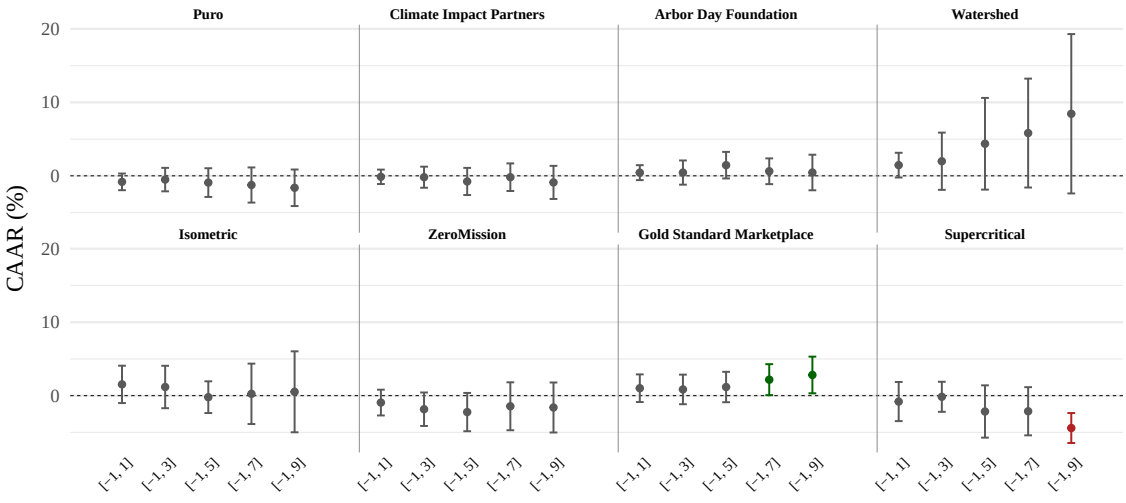


Figure 10: CAAR by broker

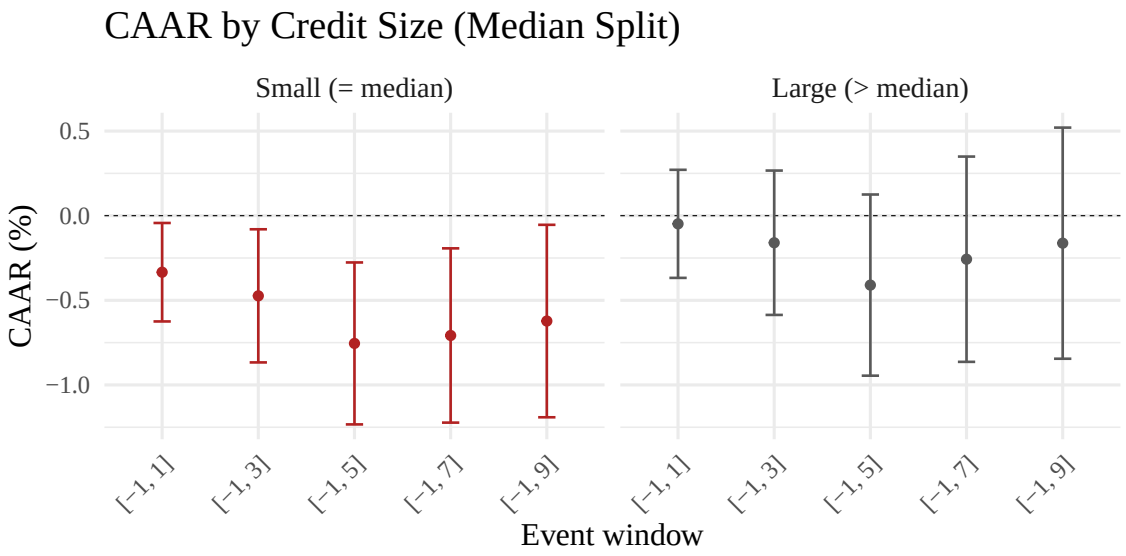


Figure 11: CAAR by credit size

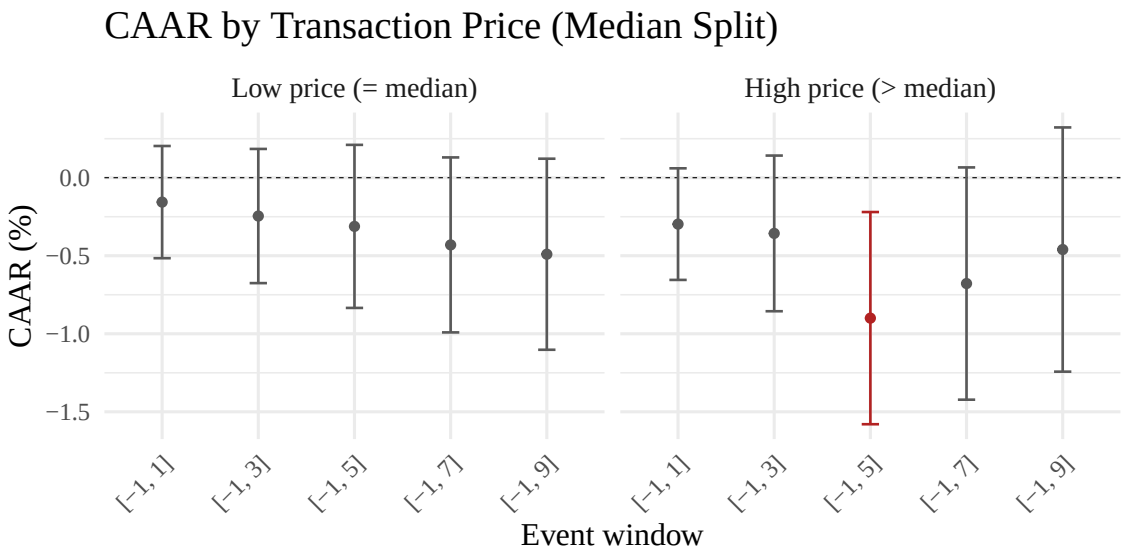


Figure 12: CAAR by transaction price

5 What drives financial market reaction?

We next examine which features of carbon dioxide removal purchases are associated with stronger (or less negative) stock-market reactions. To do so, we merge the transaction-level purchases data with firm characteristics, which yields a cross-section of 256 publicly listed buyers for whom we can compute cumulative abnormal returns (CARs) around the announcement of a purchase.

Our baseline specification relates the CAR around event e for firm i to transaction characteristics, buyer history, and firm observables:

$$CAR_{i,e} = \beta_1 \log(p_e) + \beta_2 \log\Big(1 + \tfrac{Vol_e}{CO2_{i,t-1}}\Big) + \beta_3 \log\Big(1 + \tfrac{PriorTons_{i,e}}{CO2_{i,t-1}}\Big) + \gamma' X_{i,t-1} + \delta' Z_e + \alpha_{ym(e)} + \varepsilon_{i,e}.$$

The vector $X_{i,t-1}$ collects standard firm controls measured prior to the event (log assets, cash-to-assets, and emissions intensity). The vector Z_e includes transaction-side covariates that proxy for buyer and project quality, liquidity, and sectoral differences in CDR use. We include month-year fixed effects $\alpha_{ym(e)}$ to absorb common shocks to both financial markets and the CDR market (e.g., aggregate ESG news, macro conditions, and time-varying salience of climate disclosures). Standard errors are clustered at the firm level, and in robustness we allow for two-way clustering at the firm and event-day level.

Table 2 reports the estimates for our main specification (FF5 benchmark model and a $[-1, 5]$ event window). Two patterns stand out.

First, the CAR is decreasing in the transaction price. The coefficient on $\log(p)$ is negative and statistically significant (-0.0073, s.e. 0.0036). Interpreting magnitudes, a doubling of the per-ton price ($\Delta \log p \approx \log 2 \approx 0.69$) is associated with a decline in the CAR of about 0.5 percentage points ($0.69 \times 0.0073 \approx 0.005$). This pattern is consistent with investors treating higher unit costs as a direct cash-flow drag, particularly in the absence of clear evidence that more expensive credits reliably proxy for stronger reputational benefits or lower regulatory risk. Importantly, this result is estimated conditional on month-year fixed effects and a rich set of controls, so it reflects cross-sectional differences in pricing within the same broad market environment.

Second, the CAR is increasing in the scale of the purchase relative to the buyer's own emissions. The coefficient on $\log\Big(\frac{CO2 \text{ purchased}}{CO2 \text{ firm}}\Big)$ is positive and statistically significant (0.598, s.e. 0.251). This elasticity implies economically meaningful variation across events: moving from a purchase that is one log-point larger relative to the firm's emissions baseline is associated with an increase in the CAR of roughly 0.60 percentage points. In other words, conditional on price and controls, the market response appears more favourable when the transaction is large enough—relative to the buyer's footprint—to plausibly signal a salient shift in climate strategy or to meaningfully affect expected future compliance and transition costs. This pattern is difficult to reconcile with purely symbolic interpretations of CDR purchases and instead suggests that investors distinguish between marginal and material transactions.

Table 2: Cumulative abnormal returns and CDR transaction characteristics (FF5 baseline).

Dependent Variable:	CAR
Model:	(1)
Variables	
$\log(p)$	-0.0073**
	(0.0036)
$\log\left(\frac{CO2 \quad purchased}{CO2 \quad firm}\right)$	0.5981**
	(0.2509)
Buyer sector	yes
Buyer sector = Energy	-0.0418**
	(0.0207)
Buyer sector = Real estate	-0.0361*
	(0.0193)
Buyer sector = Utilities	-0.0716***
	(0.0193)
Buyer history	yes
Firm controls	yes
Buyer quality	yes
Project quality	yes
Fixed-effects	
event year month	Yes
Fit statistics	
Observations	256
R^2	0.39008
Adjusted R^2	0.12130

Clustered (firm) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The remaining covariates are included to address compositional differences across buyers and transactions. Buyer-sector indicators reveal substantial heterogeneity: relative to the omitted sector (communication services), Energy, Real Estate, and especially Utilities exhibit more negative CARs (e.g., Utilities: -0.0716, s.e. 0.0193). These sectoral patterns are consistent with investors applying different priors about the informativeness and credibility of voluntary climate actions in high-emitting and regulated sectors, or with differences in how such announcements interact with expected future compliance obligations and political exposure. Buyer-history and quality controls are included (and absorbed in the table for brevity) to ensure that the estimated price and scale relationships are not mechanically driven by repeat buyers, maturity in procurement, or systematic sorting into higher-quality projects.

Overall, Table 2 indicates that investors respond more negatively to purchases that are more expensive per ton, but more positively to purchases that are larger relative to the buyer's emissions baseline, with substantial heterogeneity across sectors.

6 Conclusion

This paper studies how equity markets respond when publicly traded firms purchase carbon dioxide removal (CDR) credits. CDR has become a visible component of corporate climate strategies, yet it remains contested in both policy and financial circles because removal methods differ substantially in credibility, permanence, and technological maturity. These features make the valuation implications of a purchase theoretically ambiguous. A purchase may be interpreted as value-enhancing if it strengthens the firm's transition strategy and reduces regulatory or reputational exposure, or value-reducing if it is perceived as an expensive, low-credibility substitute for operational decarbonization. By bringing transaction-level evidence to the stock market, our analysis connects the literature on climate risk pricing and investor attention to climate change with the emerging literature on firm climate strategies and the use of carbon markets as a managerial response to decarbonization pressure.

Empirically, we construct a new event-level dataset that links firm-level CDR purchase events to daily equity returns and standard asset-pricing factors in a multi-country setting. The dataset integrates two complementary sources of CDR purchases, harmonizes buyer identities across sources, maps disclosed purchasers to publicly listed firms, and aligns transactions to trading days. This integrated structure allows us to evaluate market responses not only on average, but also across economically relevant dimensions of heterogeneity that reflect credibility and materiality, including buyer climate-policy posture and project characteristics. The resulting dataset and matching infrastructure provide a foundation for future work on voluntary carbon markets and firm-level climate actions.

Our baseline evidence indicates that, on average, equity markets react negatively to CDR pur-

chases. The decline is concentrated in the week around the purchase and is robust across a range of robustness analyses. The average effects conceal substantial heterogeneity. Negative reactions are concentrated among purchases by firms with weaker climate-policy postures and among transactions linked to less established removal methods. By contrast, purchases by higher-quality buyers are not associated with statistically significant abnormal performance. These patterns are consistent with the view that investors interpret carbon removal purchases through a credibility lens rather than treating them as homogeneous climate actions. In other words, markets appear to distinguish between purchases that plausibly complement a coherent transition strategy and purchases that are more likely to be interpreted as costly or exposed to scepticism about the type of credit.

Our third set of results provides cross-sectional evidence on correlates of the market response. We find that CARs are decreasing in transaction price: higher per-ton prices are associated with more negative abnormal performance, consistent with investors treating unit cost as a cash-flow drag in the short run. Further, CARs are increasing in the scale of the purchase relative to the buyer's emissions baseline: conditional on price and controls, larger purchases (scaled by firm emissions) are associated with less negative (or more positive) CARs. This pattern suggests that investors distinguish between marginal actions and purchases that are plausibly material for a firm's transition profile. Sector controls reveal additional heterogeneity, with systematically more negative CARs in high-emitting or politically exposed sectors, consistent with differences in investor priors about the informativeness and credibility of voluntary climate actions across industries.

References

Aktas, N., E. de Bodt, and J.-G. Cousin (2007): "Event studies with a contaminated estimation period," *Journal of Corporate Finance*, 13, 129–145.

Andersson, M., P. Bolton, and F. Samama (2016): "Hedging Climate Risk," *Financial Analysts Journal*, 72, 13–32, publisher: Routledge _eprint: https://doi.org/10.2469/faj.v72.n3.4.

Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2023): "Climate Change Concerns and the Performance of Green vs. Brown Stocks," *Management Science*, 69, 7607–7632, publisher: INFORMS.

Baker, M., J. C. Stein, and J. Wurgler (2003): "When Does the Market Matter? Stock Prices and the Investment of Equity-Dependent Firms," *The Quarterly Journal of Economics*, 118, 969–1005.

Boehmer, E., J. Masumeci, and A. B. Poulsen (1991): "Event-study methodology under conditions of event-induced variance," *Journal of financial economics*, 30, 253–272.

Bolton, P., Z. Halem, and M. Kacperczyk (2022a): "The Financial Cost of Carbon," *Journal of Applied Corporate Finance*, 34, 17–29, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jacf.12502.

Bolton, P. and M. Kacperczyk (2021): "Do investors care about carbon risk?" *Journal of Financial Economics*, 142, 517–549.

——— (2023): "Global Pricing of Carbon-Transition Risk," *The Journal of Finance*, 78, 3677–3754, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.13272.

Bolton, P., M. Kacperczyk, and F. Samama (2022b): "Net-Zero Carbon Portfolio Alignment," *Financial Analysts Journal*, 78, 19–33, publisher: Routledge _eprint: https://doi.org/10.1080/0015198X.2022.2033105.

Brown, S. J. and J. B. Warner (1985): "Using Daily Stock Returns: The Case of Event Studies," *Journal of Financial Economics*, 14, 3–31.

Campiglio, E., J. Deyris, D. Romelli, and G. Scalisi (2025): "Warning words in a warming world: central bank communication and climate change," .

Dörpmund, F. (2025): "Motivations and challenges for carbon dioxide removal development: empirical evidence from market practitioners," *Environmental Research Letters*, 20, 054066, publisher: IOP Publishing.

Fama, E. F. and K. R. French (1993): "Common risk factors in the returns on stocks and bonds," *Journal of financial economics*, 33, 3–56.

Flammer, C. (2021): "Corporate Green Bonds," *Journal of Financial Economics*, 142, 499–516.

Hickey, C., S. Fankhauser, S. M. Smith, and M. Allen (2023): "A review of commercialisation mechanisms for carbon dioxide removal," *Frontiers in Climate*, 4, publisher: Frontiers.

Intergovernmental Panel on Climate Change (IPCC) (2022): *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, UK and New York, NY, USA: Cambridge University Press.

Kolari, J. W. and S. Pynnönen (2010): "Event Study Testing with Cross-sectional Correlation of Abnormal Returns," *The Review of Financial Studies*, 23, 3996–4025.

Kothari, S. P. and J. B. Warner (2007): "Econometrics of event studies," in *Handbook of empirical corporate finance*, Elsevier, 3–36.

Kruse, T., M. Mohnen, and M. Sato (2024): "Do Financial Markets Respond to Green Opportunities?" *Journal of the Association of Environmental and Resource Economists*, 11, 549–576, publisher: The University of Chicago Press.

Li, X. and C. Flammer (2025a): "Climate Risk Exposure and Firms' Climate Strategies," NBER Working Paper 34276, National Bureau of Economic Research.

——— (2025b): "Climate Risk Exposure and Firms' Climate Strategies," Tech. rep., National Bureau of Economic Research.

MacKinlay, A. C. (1997): "Event Studies in Economics and Finance," *Journal of Economic Literature*, 35, 13–39.

Martynova, M. and L. Renneboog (2008): "A Century of Corporate Takeovers: What Have We Learned and Where Do We Stand?" *Journal of Banking & Finance*, 32, 2148–2177.

McKinsey (2023): "Carbon removals: How to scale a new gigaton industry," Tech. rep., McKinsey Sustainability: New York, NY, USA.

Michaelowa, A., M. Honegger, M. Poralla, M. Winkler, S. Dalfiume, and A. Nayak (2023): "International carbon markets for carbon dioxide removal," *PLOS Climate*, 2, e0000118, publisher: Public Library of Science.

Monasterolo, I. and L. de Angelis (2020): "Blind to carbon risk? An analysis of stock market reaction to the Paris Agreement," *Ecological Economics*, 170, 106571.

National Academies of Sciences, Engineering, and Medicine (2019): *Negative Emissions Technologies and Reliable Sequestration: A Research Agenda*, Washington, DC: The National Academies Press.

Oberndorfer, U., P. Schmidt, M. Wagner, and A. Ziegler (2013): "Does the Stock Market Value the Inclusion in a Sustainability Stock Index? An Event Study Analysis for German Firms," *Journal of Environmental Economics and Management*, 66, 497–509.

Sabet, A. H. and R. Heaney (2016): "An Event Study Analysis of Oil and Gas Firm Acreage and Reserve Acquisitions," *Energy Economics*, 57, 215–227.

Schneider, B. (2024): "Green Bonds: Commitment to Sustainability under Asymmetric Information," Working paper.

Smith, S. M., O. Geden, M. J. Gidden, W. F. Lamb, G. F. Nemet, J. C. Minx, H. Buck, J. Burke, E. Cox, M. R. Edwards, S. Fuss, I. Johnstone, F. Müller-Hansen, J. Pongratz, B. S. Probst, S. Roe, F. Schenuit, I. Schulte, and N. E. Vaughan (2024): "The state of carbon dioxide removal: a global, independent scientific assessment of carbon dioxide removal," Num Pages: 221 Publisher: University of Oxford.

Stolz, N. and B. S. Probst (2025): "The negligible role of carbon offsetting in corporate climate strategies," *Nature Communications*, 16, 7963, publisher: Nature Publishing Group.

Appendix A Construction and Abnormal Return Estimation

This section summarizes the choices underlying the dataset construction and the computation of abnormal returns.

A.0.1 Buyer harmonization and de-duplication

Buyer names are standardized through a sequence of cleaning operations that remove punctuation, harmonize case, and normalize common legal suffixes. Cross-source entity consolidation proceeds by first matching on standardized website domains when available and then applying cleaned-name matching for remaining observations. Transactions are de-duplicated using a composite key that combines the canonical buyer identifier, the mapped ticker (when available), the harmonized event date, and core transaction attributes (such as volume and broker/marketplace fields). When multiple records share the same composite key, the retained record is selected deterministically based on information completeness, and indicators are preserved to track whether the record was observed in one or both sources.

A.0.2 Event date mapping and trading-day alignment

Let d_i denote the calendar event date for event i as defined in the standardized transaction file. For each firm, we construct an ordered sequence of observed trading dates and assign a trading-day index. The event is assigned to the first observed trading date weakly after d_i within a bounded roll-forward window; events that cannot be mapped within this tolerance are excluded from the abnormal-return computation. This mapping ensures that day 0 corresponds to an actual trading day for the security.

A.0.3 Returns, currency conversion, and data quality filters

Daily returns are computed from both close prices and adjusted close prices. When returns are constructed from prices expressed in local currency, we convert prices to U.S. dollars using daily foreign exchange rates and then compute U.S.-dollar returns. To avoid spurious observations, we apply two filters. First, returns are not computed across large calendar gaps beyond a prespecified threshold, which limits the influence of stale prices around long non-trading periods. Second, we flag and remove extreme one-day returns outside conservative bounds to reduce the influence of clear data errors. These filters are applied symmetrically across the alternative return series so that robustness comparisons are not driven by inconsistent cleaning.

A.0.4 Factor construction and excess returns

Factor data are harmonized to a daily frequency and expressed in decimal units. We merge regional factor series to firms based on the firm's listing region inferred from ticker conventions and merge U.S. factor series as a common benchmark specification. Excess returns are computed by subtracting the relevant risk-free rate from the firm return series and by using market excess return factors (e.g., Mkt_RF) when estimating factor models.

A.0.5 Abnormal return models, estimation window, and event window

Abnormal returns are computed from standard linear expected-return models estimated by ordinary least squares at the firm–event level. Let $R_{i,t}$ denote the firm's daily return in U.S. dollars on trading day t, and let $R_{f,t}$ denote the risk-free rate. For excess-return specifications, the left-hand-side variable is $R_{i,t} - R_{f,t}$. We estimate parameters over a pre-event trading-day estimation window that starts 200 trading days before the event and ends 41 trading days before the event, requiring a minimum number of non-missing observations to ensure stable estimation. Abnormal returns are then computed over a wide event window from -20 to $+30$ trading days around the mapped event date.

For the factor specifications, expected returns are given by variants of

$$R_{i,t} - R_{f,t} = \alpha_i + \beta_i^m (Mkt_RF_t) + \beta_i^s SMB_t + \beta_i^h HML_t + \beta_i^r RMW_t + \beta_i^c CMA_t + \beta_i^b BMG_t + \varepsilon_{i,t}, \quad (2)$$

where the included factors depend on the model (CAPM, three-factor, five-factor, or an augmented specification that adds the carbon-intensity factor). For market-model robustness specifications, expected returns are modeled using a market index return rather than factor portfolios. For each event, we store the estimated coefficients, the number of estimation-window observations, and fit diagnostics, and we output a firm–event–day panel of abnormal returns that serves as the input for subsequent cumulative abnormal return calculations.

A.1 Project and buyer quality indices

Project quality index (PQI) We measure the quality of the carbon-removal project underlying each transaction using a composite project-quality index. The index summarizes third-party signals of project credibility and maturity that are available in the AlliedOffsets data and commonly used by market participants to screen carbon removal credits.

Conceptually, it captures three aspects. First, it reflects whether a project is assessed as likely to satisfy recognized integrity criteria under the Core Carbon Principles (CCP) framework. CCP is a widely used benchmark for "high-integrity" carbon removal credits; the dataset reports an ordinal assessment of the likelihood that a project would meet CCP requirements. Second, it

records whether the project has received an explicit CCP-related approval signal (when such a flag is available), which provides a discrete indication that the project has passed an external eligibility threshold. Third, it captures the project's registry status, which we interpret as a measure of operational maturity and verification progress. Registries track projects through a standardized lifecycle—typically moving from registration and validation, to verification, to credit issuance (or "crediting"), and ultimately to retirement of issued credits. A project with issuances, crediting, verification, or retirement status reflects more advanced documentation and third-party validation than a project that is merely registered or undergoing verification; conversely, projects that are cancelled, de-registered, withdrawn, or put on hold signal elevated implementation or compliance risk.

Because these indicators are reported on different scales (ordinal labels, binary flags, and registry-status strings), we convert each component to a common 0–1 scale, where higher values indicate higher assessed quality. We then compute the index as the average of the available component scores for each transaction. If none of the three components is available for a transaction, the index is missing.

Buyer quality index (BQI) We measure a buyer's climate-policy quality using a composite index built from the buyer-level indicators reported by AlliedOffsets. The underlying information captures two distinct dimensions of a firm's climate strategy. First, it reflects governance and commitments—whether the buyer has made credible climate commitments and whether those commitments are supported by observable organizational and third-party signals. Second, it captures ambition and timing—how stringent the buyer's targets are and how quickly the buyer intends to meet them.

Because these inputs are reported on different scales (ratings, categorical labels, and quantitative targets), we convert each component to a common 0–1 score, where higher values represent stronger climate-policy quality. We then aggregate the components into two block scores and a final index.

Governance and commitments (GOV). This block summarizes indicators that speak to whether the buyer's climate strategy is formalized, verifiable, and supported by institutional commitments. It combines: (i) the buyer's overall AlliedOffsets rating, (ii) whether the buyer is flagged as an unregistered carbon dioxide removal buyer, (iii) whether the buyer is committed to recognized target-setting initiatives (such as SBTi and/or net-zero target frameworks), and (iv) the status of the buyer's near-term and long-term targets (e.g., whether targets are set, only committed, or removed). We compute GOV as the average of the available component scores after rescaling them to the 0–1 range.

Ambition and timing (AMB). This block captures the stringency of the buyer's targets and the planned speed of decarbonization. It combines: (i) the temperature alignment classification of the buyer's near-term target, (ii) the near-term target year, (iii) the net-zero end year (using the earliest

Table 3: Construction of the project-quality index (PQI): components and rescaling rules

Input (raw variable)	**Rescaling rule**	**Score**
CCP likelihood	Minimal	0.10
	Low	0.25
	Medium	0.50
	Highest	1.00
	Registry not under assessment	Missing
CCP approval	Approved (TRUE)	1.00
	Not approved (FALSE)	0.00
	Missing / not reported	Missing
Registry status	*High maturity: verified issuance/crediting or equivalent*	1.00
	GOLD_STANDARD_CERTIFIED_PROJECT; Registered with issuances; Certified; CREDITED; Verified; Retired; End of crediting period	
	Intermediate: registered/active, verification ongoing, or not yet credited	0.60
	Registered; Validated (including capitalization variants); Active; Verification; Verification approval requested; Late to verify; Units Transferred from Approved GHG Program	
	Low: withdrawn, cancelled, de-registered, on hold, or denied	0.20
	Withdrawn; De-registered/Deregistered; Canceled; On Hold – see notification letter; Verification approval request denied	
Aggregation (project-quality index)	For each transaction, compute the simple average of the available component scores (CCP likelihood, CCP approval, registry status). Let K be the number of non-missing components; the index equals the sum of available scores divided by K. If all three components are missing, the index is missing. We record K as the number of quality components for diagnostics.	$[0, 1]$

available end-year information across long-term targets and net-zero frameworks), (iv) the type of interim target (e.g., absolute, intensity-based, or net-zero-like), (v) the type of end target (e.g., net zero versus carbon neutrality), and (vi) the stated interim and end target percentage reductions when reported. We compute AMB as the average of the available component scores after placing them on the 0–1 scale.

Finally, we define the buyer-quality index as the mean of the two block scores, using all available information:

$$BQI = \text{mean}(\texttt{GOV},\ \texttt{AMB}),$$

.

In the heterogeneity analysis, we split transactions into "low" and "high" buyer-quality groups using the sample median of BQI.

Table 4 reports the complete list of inputs used in GOV and AMB and documents the rescaling rules used to construct the final index.

Table 4: Scoring rules for the buyer-quality index: Governance & Commitments (GOV) and Ambition & Timing (AMB)

Input (raw variable)	**Mapping rule**	**Score**
Panel A. Governance & Commitments (GOV)		
`AlliedOffsets.buyer.rating`	`A+`	1.00
	`A`	0.95
	`A-`	0.90
	`B+`	0.80
	`B`	0.75
	`B-`	0.70
	`C+`	0.60
	`C`	0.55
	`C-`	0.50
	`D`	0.35
	Missing rating	Missing
`CDR.buyer.flag`	Labeled `Unregistered CDR buyer`	0.40

Continued on next page

Input (raw variable)	Mapping rule	Score
	Flag is missing (no evidence of being unregistered)	1.00
	Any other non-missing value	Missing
SBTi.or.NZT.commitments	un-committed	0.00
	SBTi committed	0.70
	NZT committed	0.70
	NZT & SBTi committed	1.00
SBTi.near.term.target.status	Targets set	1.00
	Committed	0.60
	Commitment removed	0.00
SBTi.long.term.target.status	Targets set	1.00
	Non-missing status indicating targets are not set	0.00
	Missing status	Missing
Panel B. Ambition & Timing (AMB)		
SBTi.near.term.target.classification	1.5 C	1.00
	1.5 C/Well-below 2 C	0.95
	1.5 C/2 C	0.90
	Well-below 2 C	0.85
	Well-below 2 C/2 C	0.80
	2 C/2 C	0.65
	2 C	0.60
SBTi.near.term.target.year	Year $\leq$ 2025	1.00
	2026–2030	0.85
	2031–2035	0.65
	Year $>$ 2035	0.50
SBTi.long.term.target.year, NZT.end.target.year	Year $\leq$ 2040	1.00

Continued on next page

Input (raw variable)	Mapping rule	Score
	2041–2050	0.75
	Year > 2050	0.50
NZT.interim.target	Net zero or similar	0.90
	Absolute emissions target	0.80
	Emissions reduction target	0.70
	Emissions intensity target	0.60
	Other	0.50
	No target	0.00
NZT.end.target	Net negative / Carbon negative	1.00
	Net zero / Zero emissions / Zero carbon	0.95
	Science-based target	0.90
	Carbon neutral(ity) / Climate neutral	0.80
	Emissions reduction target	0.70
	Emissions intensity target	0.60
	Other	0.50
NZT.interim.target.percentage.reduction	Scale to $[0, 1]$	Continuous
NZT.end.target.percentage.reduction	Scale to $[0, 1]$	Continuous

Appendix B Additional descriptive statistics

We now illustrate additional stylised facts from our dataset. Figure 13 shows the number of transactions featuring a disclosed purchaser in AlliedOffsets, CDR.fyi and in the combined dataset. In the combined dataset, there are around 19 thousand transactions with a disclosed purchaser. Of these, some are carried out by individual, private, or public buyers. Figure 14 how the different types of buyers are distributed by total tons purchased. A small share of tons is bought by individual purchasers. Private companies have purchased slightly more than 30 million tons of CO2, while public companies are closer to 40 million tons purchased.

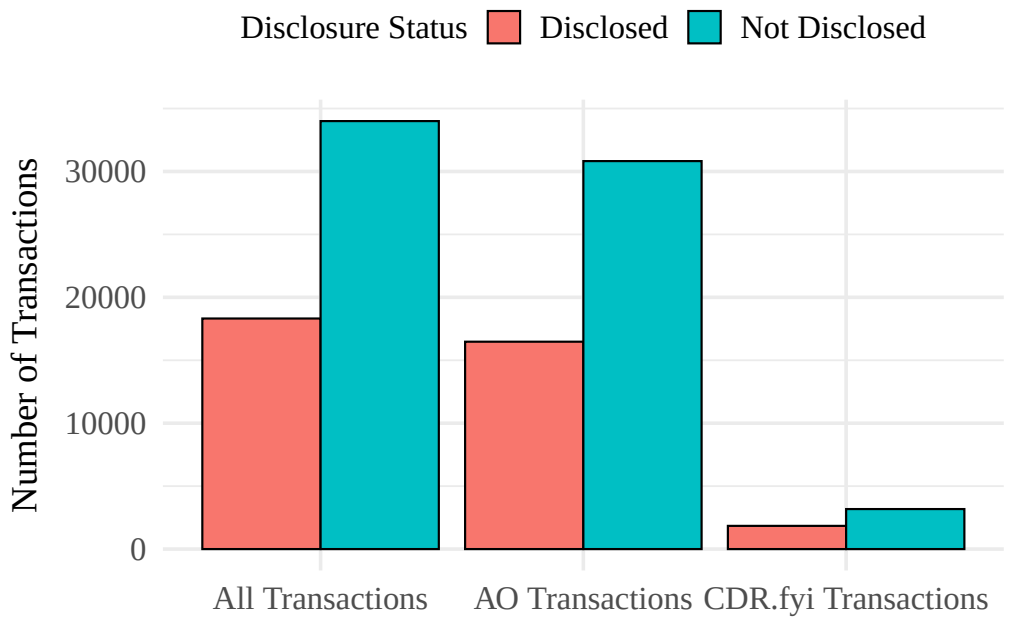


Figure 13: Number of transactions with a disclosed purchaser, from AlliedOffset, CDR.fyi and both datasets.

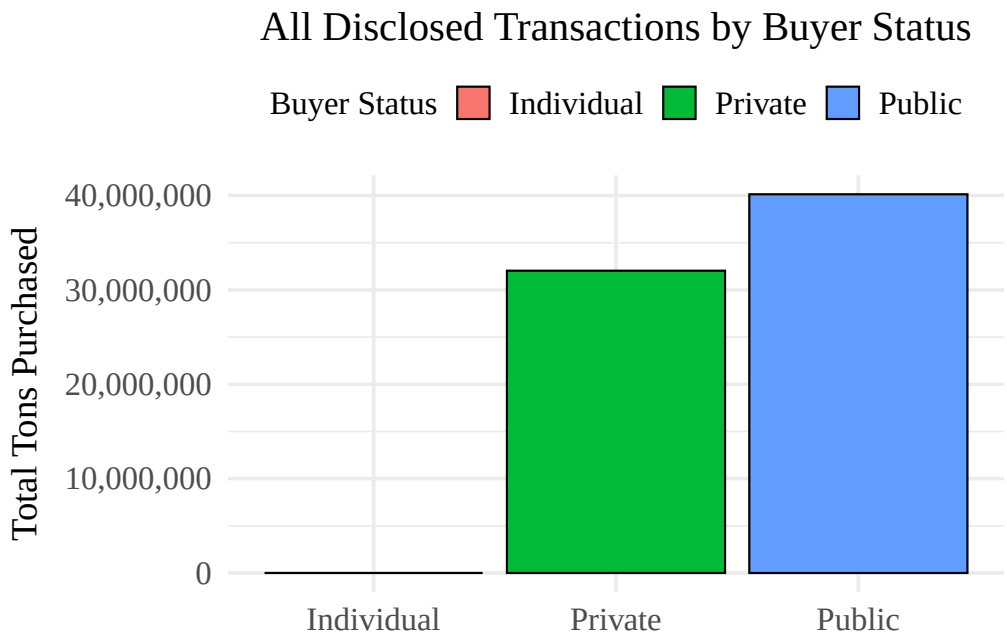


Figure 14: Purchaser type by total tons purchased.

Appendix C Group differences

Table 5: Differences in CAAR across subsamples within each event window. Cells report the CAAR difference in percentage points and the p-value for equality of CAAR across groups. For two-group splits, Δ is computed as High−Low (or Large−Small per split definition). For multi-group categorisations, we report Max−Min across group CAARs and the omnibus p-value. Significance: † $p < 0.10$, $*$ $p < 0.05$, $**$ $p < 0.01$, $***$ $p < 0.001$.

family label	0	[-1, 1]	[-1, 3]	[-1, 5]	[-1, 7]	[-1, 9]	[-3, 3]	[-7, 7]
Buyer policy quality (low/high)	Δ 0.10 (0.7866)	Δ 0.23 (0.5946)	Δ 0.33 (0.4256)	Δ 0.65 (0.2561)	Δ 0.68 (0.4533)	Δ 0.49 (0.5065)	Δ 0.32 (0.8861)	Δ 0.39 (0.2719)
Credit size (median split)	Δ 0.06 (0.9801)	Δ 0.29 (0.7657)	Δ 0.31 (0.4596)	Δ 0.34 (0.1312)	Δ 0.45 (0.4030)	Δ 0.46 (0.6390)	Δ -0.04 (0.3179)	Δ 0.09 (0.2540)
Credit size (p70 split)	Δ 0.17 (0.4608)	Δ 0.41 (0.7030)	Δ 0.35 (0.7889)	Δ 0.17 (0.1417)	Δ 0.16 (0.2610)	Δ 0.40 (0.7209)	Δ -0.30 (0.1685)	Δ -0.26 (0.1557)
Liquidity (low/high)	Δ -0.29 (0.1339)	Δ -0.08 (0.1681)	Δ **-0.31** (0.0431*)	Δ 0.05 (0.0810†)	Δ -0.04 (0.1750)	Δ 0.01 (0.3099)	Δ **-0.64** (0.0371*)	Δ **-0.60** (0.0491*)
Project quality (low/high)	Δ 0.11 (0.6079)	Δ 0.46 (0.8635)	Δ 0.07 (0.2368)	Δ 0.39 (0.2203)	Δ 0.07 (0.1541)	Δ -0.11 (0.1691)	Δ 0.35 (0.6836)	Δ -0.14 (0.1293)
Removal type (Nature vs Technical)	Δ -0.14 (0.8992)	Δ -0.52 (0.1991)	Δ -1.10 (0.1817)	Δ **-1.13** (0.0281*)	Δ **-0.91** (0.0585†)	Δ -0.70 (0.1646)	Δ -1.32 (0.3415)	Δ **-0.73** (0.0885†)
Transaction price (median split)	Δ 0.16 (0.5582)	Δ -0.14 (0.1020)	Δ -0.11 (0.1592)	Δ **-0.59** (0.0095**)	Δ **-0.25** (0.0733†)	Δ 0.03 (0.2469)	Δ -0.03 (0.3309)	Δ **-0.46** (0.0749†)

Appendix D Robustness

We conduct several robustness exercises designed to address model misspecification and outlier sensitivity, both of which are first-order concerns in event-study implementations (Brown and Warner,

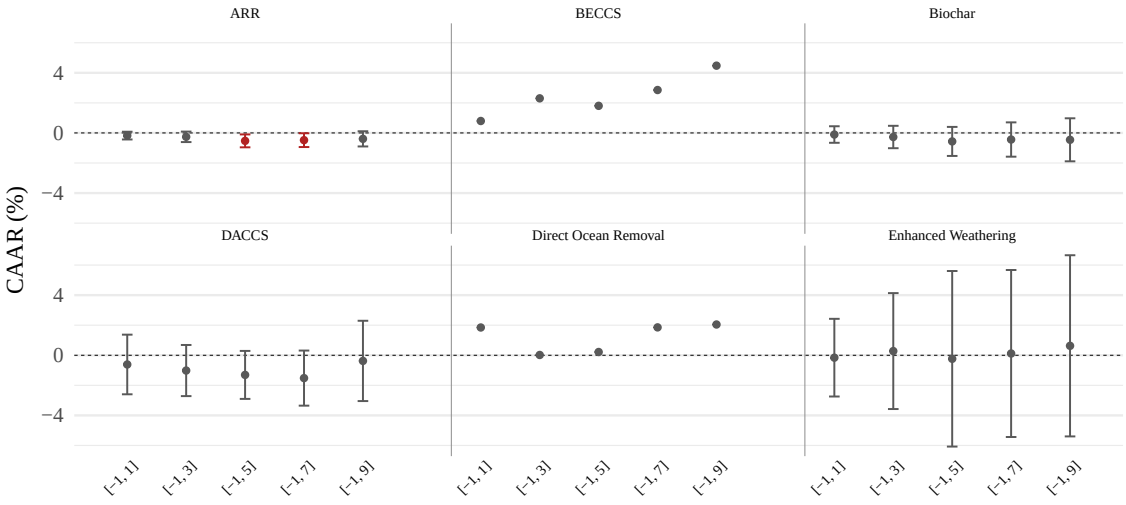


Figure 15: CAAR by removal method

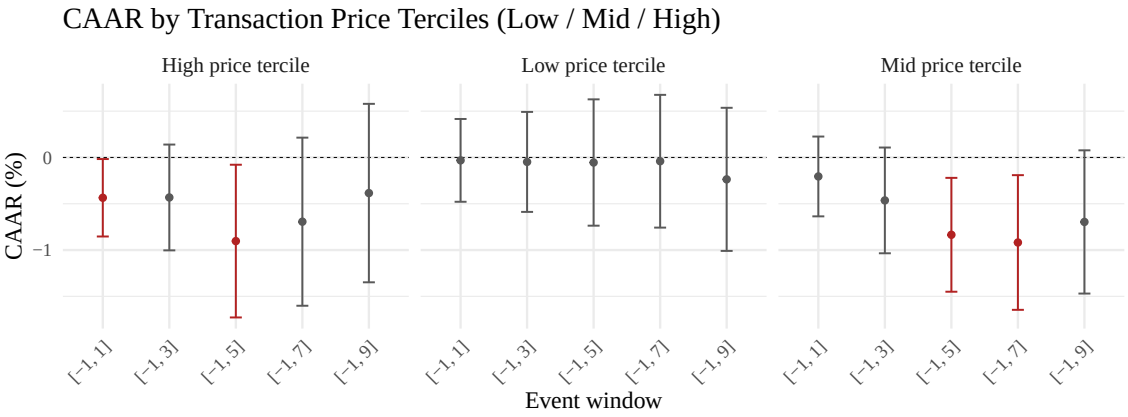


Figure 16: CAAR by transaction price

1985; Kothari and Warner, 2007).

Alternative expected return models. Beyond the one-factor specification, we re-estimate abnormal returns using multifactor models, including (i) three-factor and five-factor Fama–French specifications (Fama and French, 1993), and (ii) six-factor specifications that augment standard factors with a carbon-related factor motivated by evidence that carbon-transition risk may be priced (Bolton and Kacperczyk, 2021). In each case, we repeat the full pipeline—factor estimation in $[-200, -41]$, computation of ARs, and construction/testing of CAARs in the short windows.

Winsorization of abnormal returns. To mitigate the influence of extreme observations, we winsorize *abnormal returns* (not raw returns) as a robustness check. Specifically, within each relative day τ we winsorize $AR_{i,e}(\tau)$ at three conventional levels: 1–99, 2.5–97.5, and 5–95 percentiles. We then recompute AARs and CAARs using winsorized ARs. This procedure preserves the time-series information in the estimation window while limiting cross-sectional tail influence in the event window.

Clean-estimation-window subsample. Multiple purchases by the same firm can create overlapping estimation or event windows, mechanically inducing cross-event dependence and potentially contaminating the normal-return benchmark. To address this concern, we construct (i) the full sample of events and (ii) a "clean" subsample that retains only events for which the estimation window $[-200, -41]$ is free of other purchase events by the same firm. We replicate all CAAR estimates and inference in both samples.

Appendix E Comparison of CAARs with other event studies

Table 6: Selected Event-Study Estimates of Abnormal Returns (CAR/CAAR) Across Firm Events

Paper (Outlet, Year)	Event / Variant	Window	Metric	Estimate (%)	95% CI
Bänziger-Aiba et al. (2023)	Earnings announcements: good news portfolio	$[0, 0]$	AAR	1.084	[0.941, 1.227]
	Earnings announcements: neutral news portfolio	$[0, 0]$	AAR	-0.247	[-0.431, -0.063]
	Earnings announcements: bad news portfolio	$[0, 0]$	AAR	-1.745	[-1.929, -1.561]

Continued on next page

Paper (Outlet, Year)	Event / Variant	Window	Metric	Estimate (%)	95% CI
Oberndorfer et al. (JEBEO, 2013)	DJSI inclusion: t-GARCH specification	$[0, 5]$	ACAR	-1.83	n/a
	DJSI inclusion: non-GARCH specification	$[0, 5]$	ACAR	-1.95	n/a
	DJSI inclusion: t-GARCH specification	$[0, 0]$	AAR	-0.53	n/a
	DJSI inclusion: non-GARCH specification	$[0, 0]$	AAR	-0.52	n/a
Sabet & Heaney (Energy Econ., 2016)	Oil & gas acquisitions: reserve deals	$[-1, +1]$	CAR	0.99	[0.58, 1.40][a]
	Oil & gas acquisitions: acreage deals	$[-1, +1]$	CAR	0.35	[-1.02, 1.72][a]
	Reserve minus acreage (difference in CAR)	$[-1, +1]$	ΔCAR	-0.64	n/a
	Reserve deals	$[0, 0]$	CAR(0,0)	0.65	[0.40, 0.90][a]
	Acreage deals	$[0, 0]$	CAR(0,0)	0.37	[-0.15, 0.89][a]
Wang & Chen (JEIC, 2017)	CSR awards: awarded firms (reported peak)	$[0, +10]$	CAR (peak)	0.47	n/a
MacKinlay (JEL, 1997)	Earnings announcements: good news (example)	$[0, 0]$	AAR	0.965	[0.761, 1.169]
	Earnings announcements: bad news (example)	$[0, 0]$	AAR	-0.679	[-0.871, -0.487]
	Equity issue announcements (literature summary)	2-day	CAR	-2.70	n/a
	Straight debt issue announcements (literature summary)	2-day	CAR	-0.23	n/a
	Successful takeovers: targets (literature summary)	around annc.	CAR	>20	n/a
Flammer (JFE, 2021)	Corporate green bond issuance announcement	$[-5, +10]$	CAR	0.489	[0.017, 0.961]
	Pre-trend / placebo window (same issuers)	$[-20, -11]$	CAR	-0.129	[-0.437, 0.179]
Schneider et al. (Working paper)	Green vs conventional bond (matched): DAR	$[-2, +2]$	DAR	1.078	[0.110, 2.046][a]
Li (2015)	Bank loan announcements: pre-crisis	$[0, +1]$	CAR	0.47	n/a
	Bank loan announcements: crisis	$[0, +1]$	CAR	1.91	n/a
	Bank loan announcements: pre-crisis	$[-2, +2]$	CAR	0.88	n/a
	Bank loan announcements: crisis	$[-2, +2]$	CAR	3.45	n/a
Martynova & Renneboog (2008)	M&A bid announcement: targets (survey range)	$\approx$ two-week	CAAR range	14 to 44	n/a

Continued on next page

Paper (Outlet, Year)	Event / Variant	Window	Metric	Estimate (%)	95% CI
	M&A bid announcement: bidders (survey range)	≈ annc.	CAAR range	-1.2 to -0.7	n/a

Notes: Estimates are reported in percentage points. "n/a" indicates that a CI cannot be constructed from the reported information. [a] CI computed as estimate $\pm 1.96 \times \widehat{SE}$, with $\widehat{SE}$ backed out from the reported t-statistic when standard errors are not reported (approximation).