Transaction Costs of Emissions Offsets: Evidence from California

Willard Robinson*

April 8, 2026

This is a working paper; please do not cite.

Abstract

Emissions offsets remain an underutilized compliance instrument in California's cap-and-trade program despite trading at substantial discounts—sometimes less than two-thirds the price of a permit. A majority (50-75%) of firms never purchase an offset, leaving significant cost savings on the table. Using compliance data from 2013-2021, I identify fixed transaction costs as the primary barrier, and it is a substantial one—averaging $220,000 in the energy generation sector to $835,000 for general emitters. But this value is heavily skewed; median transaction cost estimates for the same sectors are about $47,480 and $64,900, respectively, suggesting the average cost is driven by a few high-cost firms. Using the median cost estimates, I estimate that firms implicitly paid over $10 million in total transaction costs during the study period. Moreover, relative to a frictionless market, these transaction costs caused firms to forego over $50 million in potential cost savings. Whether these are search costs, risk premiums, or public relations concerns is still ambiguous; the paper ends by noting potential future directions of research.

*I thank Spencer Banzhaf, Roger von Haefen, Erin Sills, and Zhenxuan Wang for valuable feedback on earlier versions of this paper.
North Carolina State University: wmrobin3@ncsu.edu

1 Introduction

As a program, cap-and-trade harnesses the power of markets to achieve environmental improvements at lower cost than more prescriptive command-and-control regulations. Taking advantage of this opportunity, at least 36 such markets currently operate across the globe (Appunn & Wettengel, 2024). Many of these markets, including Europe (the oldest) and China (the largest), attempt to reduce costs further by allowing regulated firms to purchase "offsets" from outside sources of abatement (Aldy & Halem, 2024). The flexibility of offsets promises substantial cost reductions when unregulated firms can abate pollution more cheaply than their regulated counterparts—and the market prices reflect this: offsets are regularly nearly 20% cheaper than the standard pollution permits. Despite these potential compliance cost savings, offset adoption remains puzzlingly limited in California's cap-and-trade system, where offsets have consistently traded at significant price discounts. This paper identifies fixed transaction costs as the primary barrier to offset adoption, quantifies their magnitude, and estimates both the value that the offsets market generates for Californian firms and the implied costs paid by offset-using firms.

One potential alternative hypothesis is that firms face per-unit costs on each unit traded. As Stavins (1995) discusses, such a case would decrease overall trading in a market; however, conditional on trading offsets, firms appear generally to do the maximum amount of trading. This kind of trading behavior—decreased activity along the extensive margin, but not the intensive margin—is consistent with a fixed transaction cost, not a per-unit cost. This paper uses patterns in the raw data to confirm this interpretation. Using compliance data from 2013-2023, I estimate that a firm in California faces a median implicit transaction cost between $25,000 and $62,000, depending on whether it has used offsets previously. As a further check for per-unit transaction costs, I estimate transaction costs for firms *conditional on participation* in the offsets market using a hurdle model (Cragg, 1971)—these are all statistically indistinguishable from zero. Moreover, I show evidence that the distribution of transaction costs can vary considerably depending on a firm's sector, with median costs

for past offset-using firms in the energy sector as low as $14,700. Whether these are search costs, risk premiums, or public relations concerns is still ambiguous as yet.

In previous work, Naegele (2018) similarly examined transaction costs in the European emission trading system (ETS). She compared semi-parametric binary quantile regression (BQR) models with standard probit models, and found that high average transaction costs were caused by a strongly skewed latent distribution of these costs. Using data on both freely allocated permits and total compliance permits and offsets, she discovered that a majority of the transaction cost was associated with trading compliance instruments in general, and not particularly tied to offsets.

My work differs from Naegele's and builds on it in three ways. First, it applies to California, the fifth-largest world economy which operates a sub-national cap-and-trade system that covers approximately 75% of the state's GHG emissions (*ICAP*, 2024). The California setting has several advantages: a uniform emissions-proportional cap on offsets quantity applies to all firms (unlike the EU, where the cap varies by country), and all eligible firms are covered by a single regulatory authority with consistent rules. Having a single entity keeps regulatory quality constant for all firms, but it does reduce a dimension of potentially useful variation. Second, whereas Naegele used cross-sectional data from Phase II of the EU ETS, I leverage panel data over three compliance periods in California's system. In some models, I use within-firm variation in the gains from using offsets—driven by changes in market-wide offset and permit prices. Third, utilizing California's sector identifiers, I calculate sector-specific heterogeneity in transaction costs. First time offset users in the energy generation sector pay a median transaction cost of approximately $46,810 while those in the oil and gas refinery sector firm pay a median transaction cost of approximately $50,340; all remaining firms pay a median $63,990 transaction cost.[1] Accordingly, firms in all three sectors participate at a rate of approximately 74%, 66%, and 40%.

My empirical analysis reveals four key patterns consistent with fixed transaction costs.

[1]These values differ from those in the abstract because the abstract uses the median cost for both first time and returning offset users, while these costs are for first time users only.

First, larger firms are significantly more likely to adopt offsets than smaller firms within the same sector, a pattern similar to Europe (Jaraitė-Kažukauskė & Kažukauskas, 2015; Naegele, 2018). Second, most firms operate at extremes, either using no offsets or using the maximum legal amount, with few in between. Naegele (2018) shows this same pattern in the EU ETS, except offsets are slightly more widely adopted in her setting, compared to California—nearly 80% of EU firms use offsets while between 40%-60% of Californian firms do, depending on the compliance period. Third, firms persistently use offsets. Taking advantage of the panel nature of the data, I identify a strong likelihood for a firm to continue using offsets after the first time. Fourth, my paper identifies sector-specific heterogeneity in transaction costs, which drives the difference in sector-level participation rates.

California's cap-and-trade program provides a uniquely clean setting to study such transaction costs. The program sets a proportional cap on offset use, meaning a firm can meet just 4% or 8% (depending on the compliance period) of its total compliance obligation with offsets. Large emitters can thus realize greater cost savings from offsets than small emitters, even though both face the same per-unit price discount. The raw data show that large emitters adopt offsets at a somewhat higher rate than small emitters—the first evidence of a fixed transaction cost.

This paper makes three primary contributions. First, it provides the first comprehensive empirical analysis of offset transaction costs in California's cap-and-trade system, quantifying barriers that previous research has acknowledged as important (Naegele, 2018; Pearson et al., 2014; Singh & Weninger, 2017; Stavins, 1995). Second, it demonstrates how policy design—specifically, the proportional offset cap—can unintentionally create distributional consequences by restricting cost-saving opportunities to the largest, highest-emitting firms. Third, it illustrates the usefulness of quantile regression methods used alongside mean regression when the latent variable has substantial heterogeneity (Alejo & Montes-Rojas, 2025; Kordas, 2006; Manski, 1975).

The remainder of the paper proceeds as follows. Section 2 provides institutional back-

ground on California's cap-and-trade program and other environmental policy in the state, provides a brief summary of the literature on transaction costs in offest markets, and discusses some hypothetical explanations for the source and nature of the transaction cost. Section 3 describes the data sources. Section 4 reveals several stylized facts using raw data which support the hypothesis of a fixed transaction cost. Section 5 outlines the empirical strategy and section 6 presents these estimates. Section 7 discusses policy implications and concludes.

2 Background

2.1 Environmental policy in California

California initiated its cap-and-trade system in 2013 with the goal of decreasing the state's GHG emissions to 1990 levels within seven years (AB-32, 2006). The program covers about 75% of the state's entire GHG emissions (about 450 high-emitting firms), and the overall emissions cap is decreasing over time to meet that goal. Regardless of sector, any entity that emits more than 25,000 tons of GHG in a year is covered.[2]

In the first decade, CARB consistently sold less than a quarter of its yearly permit cap. One interpretation of this is that the cap was set too high. Instead, Burtraw (2016) argues that the demand for emissisons was decreasing faster than expected as a result of companion policies that incentivized infrastructure investment and innovation.

California represents the fifth-largest world economy, but its cap-and-trade is a sub-national policy. Caron et al. (2015) find that this originally increased the emissions of out-of-state electricity exporters by up to 45%—a phenomenon known as "leakage." Fowlie and Perloff (2013) find efficiency in trading for Southern California's RECLAIM, a local air pollutant permit trading system: the final distribution of pollution across facilities is not

[2]The main caveat to this statement is that all firms in the energy sector are automatically covered by the cap-and-trade.

strongly determined by the initial allocation of permits. While RECLAIM is much smaller in magnitude than CARB's program, its economic efficiency suggests that cap-and-trade can successfully operate in the state.

To prevent emissions leakage, the California Air Resources Board (CARB) strategically allocates permits to reduce burdens on the firms most likely to relocate.[3] Nearly all electricity generators receive full emission allocations, and energy-intensive, trade-exposed industries receive additional allowances. While allocations are private, the use of compliance instruments to cover emissions is tracked through the Compliance Instrument Tracking System Service (CITSS).

The program is divided into 2-3 year compliance periods. The purpose of a compliance period is to create a timeframe during which a firm is a "regulated entity." All regulated entities must surrender a compliance instrument for each ton of CO_2 (or CO_2-equivalent) emitted during their tenure as a regulated entity. A single entity may operate multiple facilities, with total obligations equalling the sum of all facility emissions.

A firm becomes a regulated entity if it emits 25,000 or more metric tons of CO_2-equivalent in a year. Once regulated, the firm needs to surrender, by the end of the compliance period, compliance instruments equal to all of its emissions during the compliance period. While regulation persists throughout a compliance period once triggered, there is no automatic carryover across periods—firms must exceed the 25,000-ton threshold again to become regulated in subsequent periods. However, most regulated firms do continue to exceed this threshold, creating de facto regulation persistence.

The cap-and-trade program allows regulated entities to cover a small share of their emissions using offsets. During the first three compliance periods (2013-2020), firms could meet up to 8 percent of their total compliance obligation by purchasing and surrendering offsets. Starting in the fourth compliance period, this limit decreased to 4 percent. At the same time, it became the rule that no more than 2 percent of a compliance obligation can be

[3]For more on this, see Bushnell et al. (2014) and Fowlie et al. (2021).

met using offsets that do not provide direct environmental benefits to the state (DEBS) of California. Figure 1 demonstrates the significant incentive to use offsets. The policy changes in the fourth compliance period decrease demand for offsets, but, as figure 1 shows, the price discount simultaneously increased by a large amount.

2.2 Transaction costs in other environmental markets

Payments for ecosystem services (PES) are what generate offsets, and, as Alston et al. (2013) discuss, having formal institutions can ease the transaction costs associated with these markets. Such transaction costs include things that translate to demand-side such as search costs, measurement, enforcement, insurance, and more (Cacho et al., 2005, 2013; Elson, 2011; Galik et al., 2012);[4] other costs are more supply-side specific costs such as training time for new agricultural methods (Merger et al., 2012; Mukama et al., 2012) and protection (from illegal harvesting) costs for newly protected lands (Potvin et al., 2008).

While these papers focus on the supply-side of offsets, many of the same transaction costs can be applied to the demand-side (i.e., regulated Californian firms). Firms acquire offset credits either by entering long-term contracts with suppliers themselves (which must be verified by one of a short list of verifying bodies) or by purchasing offsets on secondary markets (e.g., CNaught, Evolution Markets). Longer-term contracts can provide lower per-unit costs and ease concerns about collecting future offsets, but secondary "spot" markets provide offsets in a more liquid form at a higher price. All offset projects must be registered on one of the two approved offset project registries (American Carbon Reserve and Climate Action Reserve); these registries generate a finite amount of CARB offset "credits" which can be substituted one-for-one with the CARB-issued permits.

The price of offsets has been gradually decreasing each year; Pearson et al. (2014) proposes that many estimates of marginal costs of carbon sequestration are underestimated (sometimes by up to 30%) by ignoring transaction costs. A portion of these costs are paid by the

[4]See Wertz-Kanounnikoff (2008, section 3) for a more in-depth review of these costs as they relate to REDD+.

earliest suppliers of voluntary sequestration, a so-called 'early-actor penalty.' Acknowledging this, California allowed Verified Carbon Standard, a third offset project registry, to operate an "early action offset program" from 2013-2016.[5] We would not expect this early-actor penalty to be so prescient in the California market, as The Climate Action Reserve , Verified Carbon Standard, and the American Carbon Registry all listed offset projects by 2009.

Cacho et al. (2013) discusses that transaction costs have been a major hurdle for suppliers of land-based carbon offsets. Focusing on the supply-side, they show evidence of significant supply-side, annual fixed costs can be the most influential of transaction costs in terms of offset project feasibility.

Ideally, all firms within the same cap-and-trade regime should face the same price for emissions, so emissions are distributed such that aggregate abatement cost is minimized. In such a market, allowing offsets acts as a pure market price reduction. These golden efficiencies can be hampered by transaction costs (Stavins, 1995). According to Hahn and Stavins (2010), these transaction costs arise from search costs, bargaining costs, and enforcement costs. In practice these factors take the form of carbon price uncertainty, doubts over the longevity of a cap-and-trade policy, and learning costs.

There is a substantial body of research on transaction costs in general, but research on transaction costs specifically in environmental policy is somewhat limited (Krutilla, 2011). Jaraitė-Kažukauskė and Kažukauskas (2015) find evidence of transaction costs in the initial phase (2005-2007) of the EU ETS, and suggest that this led to a number of permits being expired at the end of the phase. The broader literature on technology adoption under fixed costs also provides relevant insights: standard adoption models (e.g., Reinganum, 1981) predict that fixed entry costs generate thresholds below which potential adopters remain inactive.

The environmental compliance context adds the complication that the "technology" being adopted (offsets) carries an inherent risk of non-additionality, which may itself be a source

[5]Verified Carbon Standard is now called Verra, a change that took place on its 10-year anniversary in 2018.

of the fixed cost (Hahn & Stavins, 2010; Stavins, 1995). Studies of compliance behavior in emissions markets find that allowance prices are driven by a mix of regulatory, financial, and energy-market factors (Hintermann, 2010), and that the transition to market-based regulation can itself alter industry dynamics in ways that interact with fixed compliance costs (Fowlie et al., 2016).

Sources of the transaction cost Two features of offsets themselves create the transaction cost. First, "offset" is a broad term relating to any number of heterogeneous projects that mitigate greenhouse gas (GHG) emissions. These require evaluation, verification, and ongoing monitoring. While third party verifiers (e.g., Verra, ACR) help, offsets still are not commodified as easily as emissions permits issued by the state, and this complexity can create a barrier for firms looking to leverage offsets to reduce compliance costs. Second, offsets face persistent concerns over "non-additionality" (i.e., whether *credited* emission reductions reflect *true* emission reductions). Recent empirical evidence shows that non-additionality is widespread in most payment-for-ecosystem services markets (Aspelund & Russo, 2025; Badgley, Chay, et al., 2022; Badgley, Freeman, et al., 2022; Calel et al., 2025; West et al., 2023). This influx of empirical critique can potentially create a harmful stigma, causing public relations concerns for offset-adopting firms. Such a stigma (deserved or otherwise) may be particularly acute for firms in environmentally-conscious California.

There is an established literature on the theoretical nature of transaction costs for offsets. Table 1 summarizes three hypothesized sources of transaction costs. The exact nature of the costs is untestable given the current data.

The first is time costs and human capital investment. Transaction costs can include information and compliance expenses. These may be lessened in California where offsets are listed on centralized registries, and sales are brokered by a third party. However, dealing with intermediaries and ensuring a firm is compliant in light of the new instrument type may require a significant amount of employee time.

The second is the risk of potential "invalidation" either by being determined non-additional ex post or by being destroyed in a natural event. In this case, offsets must be replaced with new compliance instruments (AB-32, 2006). Though invalidation is rare, a risk-averse firm may be willing to pay a premium for stable permits provided by the state. Because larger firms typically have more employees and, having more access to capital, are more likely to be risk-neutral, the risk premium can be endogenous across firm size, meaning smaller firms may be more likely to be disproportionately affected.

The third—and most speculative—is the potential for public image concerns. GHG offsets represent voluntary abatement generated by unregulated entities and purchased by regulated firms for compliance. Because they are provided voluntarily, a critical concern is "non-additionality"—whether credited emission reductions would have occurred anyway. Since regulated firms receive credit for these reductions, non-additional offsets effectively increase global emissions. Calculating true additionality requires knowledge of an unregulated entity's business-as-usual behavior, which is fundamentally unobservable, making offsets vulnerable to adverse selection.

Recent empirical evidence demonstrates that non-additionality is pervasive across offset programs. Studies find that only 20% of Conservation Reserve Program contracts are additional (Aspelund & Russo, 2025), just 6.2% of REDD+ forest conservation credits (West et al., 2023), and 52% of EU Clean Development Mechanism wind farms (Calel et al., 2025). These findings suggest that widespread concerns about offset quality are empirically grounded. These doubts about the integrity of offsets *could* potentially create public relations costs for firms using them, particularly in environmentally conscious California.

Table 1: Sources of transaction cost in the compliance market

Transaction cost	Definition
Information/search costs	Cost of finding offsets and verifying compliance within the market
Risk premium	Additional value placed on permits because they cannot be overturned after purchase
Public relations	Firms may be seen as "greenwashing" or the public may have negative feelings toward offsets in general

Note: This table highlights the three hypothesized contributions to a regulated firm's transaction cost.

3 Data

3.1 Emissions, permits, and offsets

The main data for this paper come from CARB's compliance reports.[6] These are annually-published tables which list the emissions compliance obligation, the number of permits retired, and the number of offsets retired for every regulated firm. In effect, they summarize each regulated firm's cap-and-trade history. These are published annually as well as for the full compliance period, with a one-year lag between obligation and compliance report publication.

While the compliance reports cover all *regulated* firms, CARB's Greenhouse Gas Inventory contains emissions for all facilities which produce over 10,000 tons of CO_2-equivalent emissions in a year.[7] This supplementary dataset allows me to assess past emissions history for all regulated firms.

3.2 Compliance instrument prices

CARB publishes annual and quarterly transfer summaries which include the prices and quantities of both allowances and offsets.[8] Each report contains aggregate data on the total

[6]Example: https://ww2.arb.ca.gov/resources/documents/compliance-summary-report
[7]This dataset has been used by Badgley, Freeman, et al. (2022) and Tuladhar et al. (2014).
[8]Example: https://ww2.arb.ca.gov/cap-and-trade-program/summary-market-transfers-report

number of allowances transferred between firms during the period. These reports do not include auctioned permits. Instead, they contain the prices for the allowances and offsets transferred from one firm to another. Each transfer is either a "priced transaction" or an "unpriced transaction." Priced transactions occur when one firm purchases an instrument from another in a single transaction while unpriced transactions consist mainly of transactions between corporate associates, where the instruments are being traded as part of a larger bundle.[9] Each transfer summary provides a weighted average, median, and standard deviations separately for both allowances and offsets. A separate table published by CARB details the prices and quantities of permits bought at auction.[10] Combining these two sources allows me to calculate the prices of offsets and permits across years. Because my analysis is at the compliance period level, I use the last year of a compliance period's prices to reflect that period.[11]

3.3 Analytical sample

Firms are only required to meet a portion of their obligations annually, but they must satisfy *all* obligations by the end of the compliance period. Thus, the compliance period is the time scale along which firms are "fully participating" in cap-and-trade in the sense that they cover all emissions during that time period with compliance instruments. For the sake of interpretability, I focus my analysis at the level of the compliance period.

To keep the data easily interpretable, I drop a small number of firms that acquire or get rid of facilities during a compliance period, as this re-arranging of facilities is a form of adaptation to regulation that could potentially be substituted in place of offsets or permits. Additionally, I remove entity-compliance period pairs that submit compliance reports for only one year. This occurs when firms emit more than 25,000 tons of CO_2 in the last year

[9] 96% of all transferred allowances were transferred in a priced transaction in 2023. This percentage is typical of other years as well.

[10] Example: https://ww2.arb.ca.gov/sites/default/files/2020-08/results_summary.pdf

[11] This is because the compliance instruments are due by the end of the compliance period. I have repeated the analysis using the average price in Appendix B and the qualitative results are unchanged.

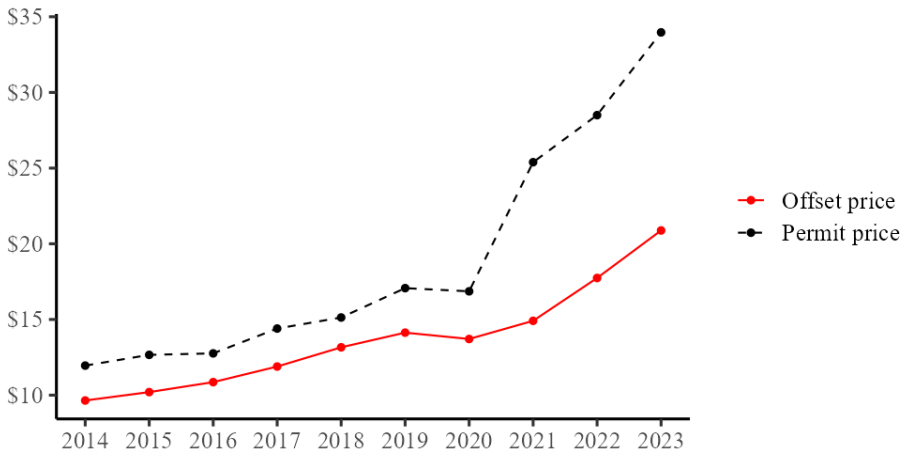


Figure 1: Price evolution of permits and offsets

Note: The price of official CARB permits has been greater than the price of offsets for the entirety of the policy. The price divergence increases greatly in the fourth compliance period. This coincides with both a lowering of the offset cap from 8 to 4 percent and a limiting of offsets not providing direct environmental benefits solely to California. This figure is descriptive and covers all four compliance periods (2013–2023). The econometric analysis in Sections 5–6 is restricted to the first three compliance periods (2013–2020), because the rules governing offsets changed substantially in compliance period 4.

of a compliance period. CARB allows them to surrender compliance instruments over two years, and, as a result, the sum of all compliance instruments surrendered during either compliance period does not match the total obligation. Finally, because the rules around offsets—and their relative price discount—changed dramatically in the fourth compliance period, I restrict my econometric analysis to the first three compliance periods.

4 Patterns in the raw data

Several patterns in the raw data provide initial evidence consistent with fixed transaction costs. These are descriptive patterns, but nonetheless they motivate the modeling approach in section 5.

Offset prices are consistently lower than the official CARB permits throughout the program's history. Figure 1 shows the evolution of permit and offset prices across the four compliance periods. In the first three compliance periods (2013-2020), permits traded at a premium of $2.30 to $3.15 over offsets, representing a 17–19% price discount for offsets.

Hence, there is a meaningful financial incentive to substitute offsets for permits. The price discount increased dramatically in the fourth compliance period (2021-2023), with permits trading at approximately $13.09 more than offsets—more than four times the largest previous premium. This premium increase coincides with both a lowering of the offset cap from 8% to 4% and the new requirement that no more than 2% of a compliance obligation can be met using offsets that do not provide "direct environmental benefit to the state" (DEBS) (AB-398, 2017). As permit prices rise with a falling cap, offset prices have not kept pace, widening the gap and strengthening the economic incentive to use offsets.

Table 2: How prevalent are offsets?

CP	No. firms	Uses any	Uses max	Uses half
1	195	84 (43.1%)	52 (61.9%)	–
2	183	72 (39.3%)	38 (52.8%)	–
3	179	107 (59.8%)	83 (77.6%)	–
4	163	105 (64.4%)	89 (84.8%)	13 (12.4%)

Note: CP stands for compliance period. There is an 8% offsets cap in compliance periods 1-3 and a 4% cap in compliance period 4. The third column shows the number of firms using any offsets, with the percentage of all firms in parentheses. The fourth column shows the number of firms using the maximum amount of offsets, with the percentage using the maximum amount, conditional on using any offsets in parentheses. Finally, the fifth column shows the number of firms using half of their allowable offsets, which is potentially relevant in the fourth compliance period when only half of a firm's offsets can be non DEBS-providing. In parentheses of the fifth column is the percentage of firms using half of their maximum amount, conditional on using any offsets. In this table, using the maximum is defined as having an absolute difference of less than 0.001 between offsets used and the maximum possible; using half is defined similarly.

Table 2 shows the prevalence of offset adoption across firms, by compliance period. Column 1 denotes the compliance period and column 2 the number of regulated firms. Column 3 reports the number of firms using offsets in each period; in parentheses is the percentage of all firms. Well under half of all regulated firms use offsets during the first two compliance periods, but more than half use offsets in the third and fourth compliance period—up to 64% in the fourth. Column 4 shows the number of firms using the maximum amount of offsets, and, in parentheses, the percentage of firms using the maximum amount of offsets (condi-

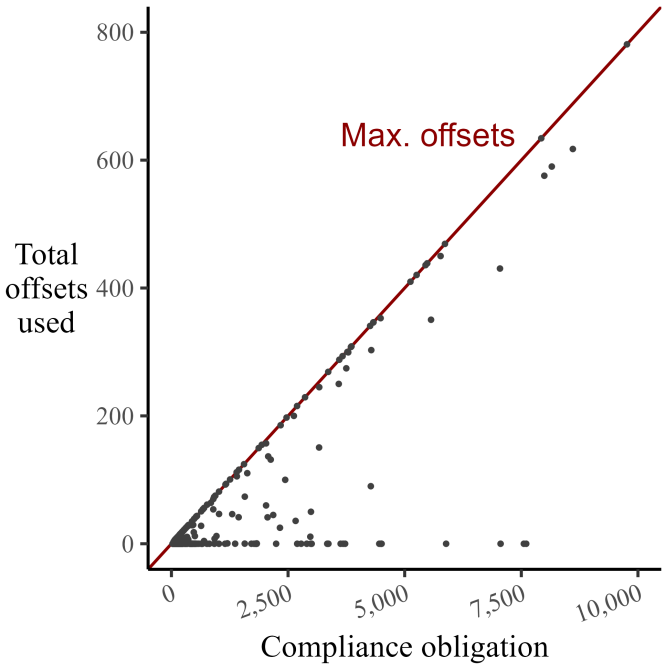


Figure 2: Relationship between emissions and offsets usage during the first three compliance periods

The figure above is limited to firms with 10 million or fewer tons of GHG emissions for visual clarity. The solid red ray coming from the origin represents the maximum amount of offsets that can be used at the an emissions level. Points along the solid red line indicate surrendered offsets equal to 8% of a firm's emissions—the maximum amount. The cluster of points along the red line and along the horizontal axis—with relatively few truly in between—suggests that firms likely are either "all in" or are "out" of the offsets market. This figure is descriptive—not causal—and covers compliance periods 1–3 (2013–2020), which is also the analytic sample for the econometric models.

tional on using any). Consistently, over half of the firms that use offsets, use the maximum amount. The fourth column shows the number of firms using half of their maximum, which becomes relevant in the fourth compliance period when at most half of a firm's maximum offsets can be obtained without providing DEBS—only about 12% of offset-using firms in the fourth compliance period do this.

Figures 2 and 3 show the offset adoption patterns graphically. The solid red line in figure 2 represents the maximum amount of offsets possible at any emissions level during the first three compliance periods (8% of emissions). The clusters of points along this red line and along the horizontal axis—with relatively few truly in between—suggest that firms typically are either "all in" or "completely out" of the offsets market.[12] Figure 3 presents the cumulative distribution functions (CDFs) of the share of allowable offsets used. Figure 3a shows the unconditional CDF across all regulated firms. The tall vertical line at zero indicates that over half of firms use no offsets at all in the first three compliance periods. The long horizontal stretches repeat the clustering-at-limits pattern from figure 2. Figure 3b restricts the sample to firms using any offsets. Doing this reveals that roughly half of firms that use offsets use an amount between 0 and the limit (also shown in column 4 of table 2), with no clear pattern on the use of these "interior" firms. Conditional on using offsets, using the maximum amount is the dominant pattern.

Figure 4 shows the share of firms using offsets, by emissions level. The emissions are split into deciles, and the share of firms using offsets with emissions in that group is represented by the height of the bars. In general, the share of firms using offsets increases as emissions increase, with a majority of above-median emissions firms using offsets. This is consistent with economic theory as the proportional offsets cap means the cost-savings incentive is higher for firms with more emissions.

The proportional offset limit put in place by CARB incentivizes larger firms to use offsets. Because the offset price discount is common to all firms, but the number of offsets a firm

[12]Figure D-1b shows a similar result for the fourth compliance period

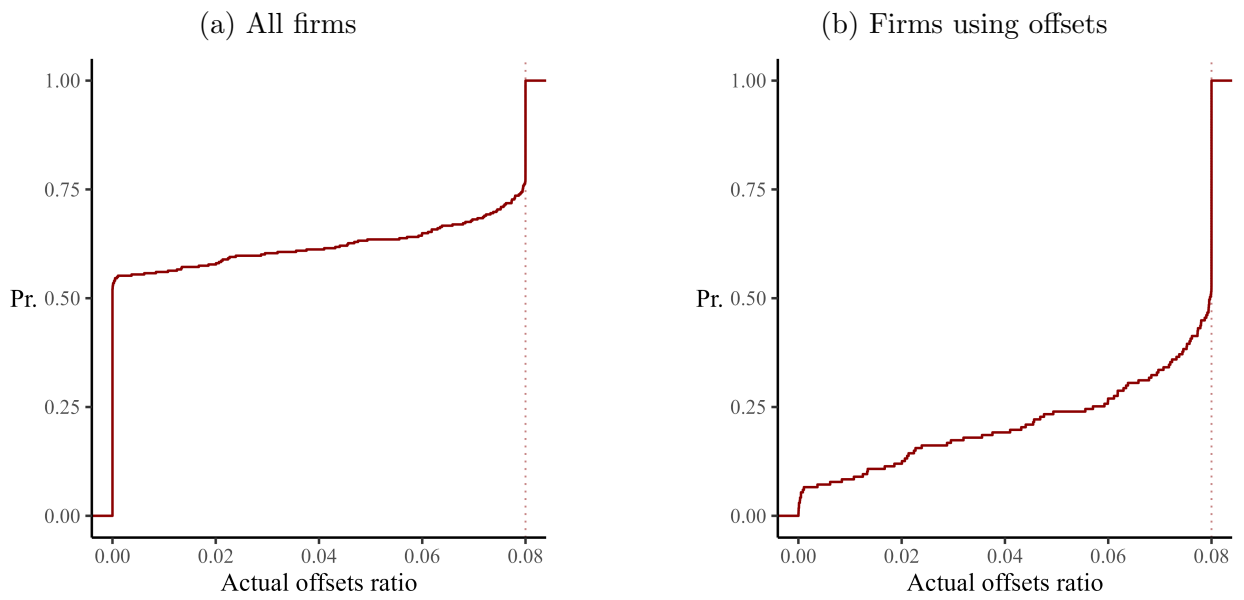


Figure 3: Cumulative density share

This shows the empirical CDF of the share of a compliance obligation met using offsets. The tall vertical line at 0 shows that over half of the firms do not use any offsets. The long horizontal stretch reveals a clustering in the share of offsets used at limit values (0% and 8%).

can use is proportional to its emissions, the incentive to use offsets is increasing in firm emissions. This incentive structure leads to firms selecting into the offset market as revealed in figure 5, which shows that firms using offsets (first order) stochastically dominate firms that do not in the maximum potential gains. The maximum potential gains are calculated by multiplying the offset price discount with the maximum legal amount of offsets. This allows me to compare the dollar-value incentive of all firms to use offsets, independent of whether a firm actually uses offsets. The horizontal axis of figure 5 represents the strength of incentive to use offsets. The stochastic dominance is intuitive: at all values on the horizontal axis, the probability that a firm which uses offsets has a weaker incentive to use them is smaller for the firms that use offsets. This means that the firms observed using offsets are also the firms that stand the most to gain from doing so—firms respond to the market's incentive structure.

Finally, table 3 shows the persistence of offsets use by firms. This table shows the trend that offsets grow in popularity as the program continues, particularly starting in the third

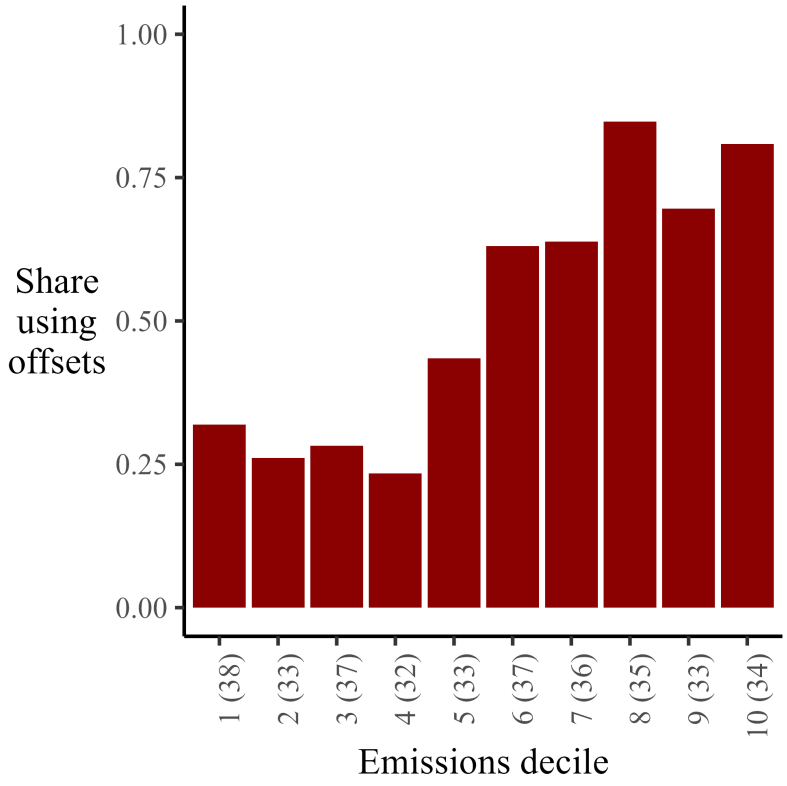


Figure 4: Larger firms are more likely to use offsets

Note: Firms are grouped into deciles based on their total emissions, so the unit of observation is a firm×compliance period. The height of the bar indicates the share of firms within that decile using offsets. The general trend that bars get higher as the decile increases indicates that larger firms are more likely to use offsets.

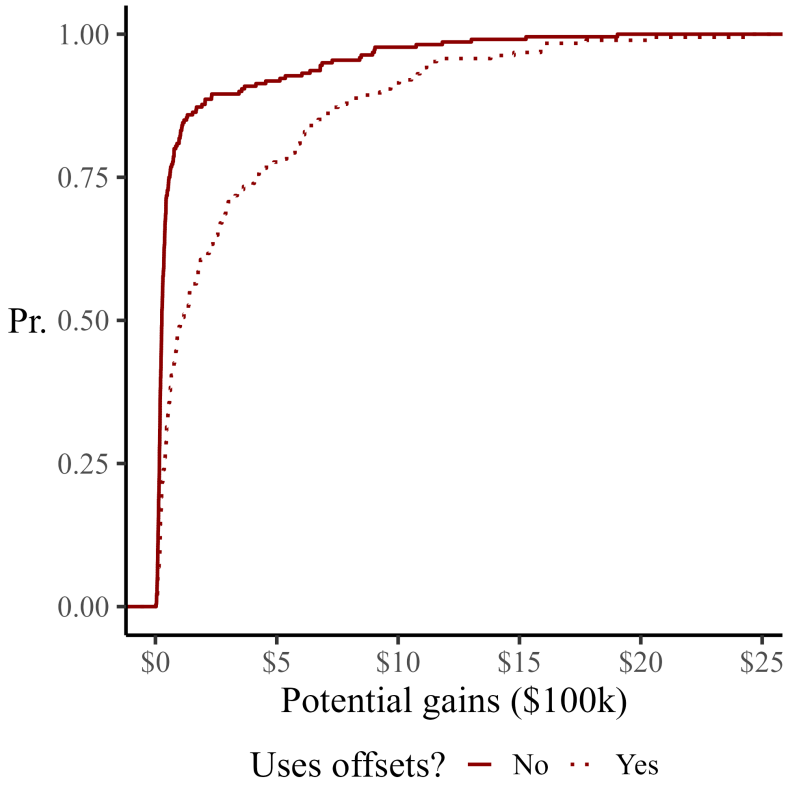


Figure 5: Potential gains from offset adoption

Note: The probability that gains are less than $\$x$ is smaller for firms that use offsets than it is for firms that do not. This true for all x, so the expected gains from offset adoption in the sample of firms using offsets is higher than in the sample that does not. Firms may be selecting into offset adoption based on the potential gains.

Table 3: Do entities continue to use offsets?

		Of those using offsets this CP	
CP	using offsets this CP	uses offsets every CP in future	uses offsets at least once in future
1	43.1%	41.7%	73.8%
2	39.3%	68.1%	68.1%
3	59.8%	100%	100%
4	64.4%	–	–

Note: This table highlights the persistence in offset usage. Once they have "bought in" to the offset market, firms continue using them in future compliance periods.

compliance period. It also highlights a feature that may be missed from the other tables and figures, and is not mentioned in any of the previous literature: firms that use offsets almost certainly continue using them in the future. This persistence could reflect two aspects of the transaction cost. First, it could reflect that the firms have learned how to navigate the offsets market, thus reducing future transaction costs. Second, the persistence could suggest that a public relations concern may drive the transaction cost. Once a firm is revealed to have used offsets, it will likely continue to use them to meet future compliance obligations.

5 Estimation strategy

Motivated by the raw data, I attempt to measure the actual cost.

5.1 Decision framework and latent variable model

I assume that firms in California maximize profit by making a two-step decision regarding offsets: (1) whether to use any offsets; and (2) if using any, how many offsets to buy. The raw data shows that several firms choose not to use any offsets, so this decision process has many corner solutions which need to be accounted for in an econometric model.

A firm i in sector j uses offsets during compliance period t if and only if the cost-savings outweigh the fixed transaction cost. Let $Gains_{ijt} = p_t \cdot \tilde{q}_{ijt}$ be the maximum potential

gains from offset adoption where p_t is the observed per-ton offset price discount and $\tilde{q}_{ijt}$ is the calculated maximum possible offset quantity. Because offsets are cheaper than permits everywhere in the data, p_t is always positive. The maximum possible emissions offset is the interaction of the firm's compliance period emissions with the maximum share of emissions that can be met with offsets. For example, a firm that emits 100 tons of GHG, $\tilde{q}_{ijt} = 8$ tons of emissions that can be offset for $t \in \{1, 2, 3\}$. Note that $Gains_{ijt}$ is in units of dollars. Because the policy regarding offsets use changed dramatically in the fourth compliance period, the estimation for the remainder of the paper uses only data from the first three compliance periods.[13]

The underlying latent payoff to firm i of using offsets in period t is:

$$y^*_{ijt} = \beta\, Gains_{ijt} - \sum_{j=1}^{3} \rho_j \mathbb{1}_i^j + \epsilon_{ijt}, \tag{1}$$

where $\mathbb{1}_i^j$ indicates that firm i is in sector j (there are 3), and ϵ_{ijt} is an idiosyncratic error.[14] The parameters of interest are the marginal value of cost savings, β, and sector-specific transaction costs, ρ. The average transaction cost for a firm in sector j is ρ_j/β, while that firm's idiosyncratic transaction cost is $(\rho_j + \epsilon_{ijt})/\beta$.

The latent y^*_{ijt} represents the real cost savings earned from using offsets, but is unobservable. However, firm offset adoption is observed, and one can make inferences based on which firms adopt, and which do not. A firm uses offsets (denoted $\mathbb{1}^O_{ijt} = 1$) if and only if the payoff is positive:

$$\mathbb{1}^O_{ijt} = \mathbb{1}\left\{y^*_{ijt} > 0\right\} = \mathbb{1}\left\{Gains_{ijt} \geq \frac{1}{\beta}\left(\sum_{j=1}^{3} \rho_j \mathbb{1}_i^j - \alpha_i - \epsilon_{ijt}\right)\right\}.$$

Thus, offset adoption occurs when the potential gains exceed the net transaction cost $(\rho_j - \epsilon_{ijt})/\beta$, which is the sum of the sector-level fixed cost and an unobservable shock. As such,

[13] Descriptive patterns in the raw data are presented in appendix D.

[14] There is the potential for an unobserved heterogeneous effect as well. I discuss the implication of this, and describe robustness tests to address it in section 5.3.

the probability of observing a firm using offsets is a function of the potential cost-savings.

The potential gains variable, $Gains_{ijt}$, provides two sources of variation. First, *time-series* variation comes from changes in the offset price discount p_t across compliance periods: as the statewide cap tightens, permit prices rise while offset prices lag behind, widening the discount. Because p_t is a market-wide price, it is plausibly exogenous to any individual firm's decisions in the short run. However, this does not change much in the first 3 compliance periods.

The second source of variation is *cross-sectional*, coming from differences in firm emissions, which determine the maximum allowable offset quantity $\tilde{q}_{ijt}$. An important caveat is that firm emissions may not necessarily be exogenous—a firm's scale and abatement investments are jointly determined with its compliance strategy. If, conditional on sector, residual variation in emissions is uncorrelated with unobserved propensity to adopt offsets. This is plausible if sector indicators absorb the main structural differences in compliance environments across firm types, but readers should note that within-sector size variation may still capture characteristics—such as compliance staffing or capital access—that also directly affect adoption probability. The panel dimension (three compliance periods) allows within-firm variation in p_t to contribute to identification; this variation is more cleanly exogenous, though it is limited to three points in time.

5.2 Estimating equations

Taking the framework to the data lends itself naturally to standard binary choice techniques for the opting-in decision (Jaraitė-Kažukauskė & Kažukauskas, 2015; Naegele, 2018). I estimate separate transaction costs for the three main sectors identified by CARB: "energy", "oil and gas", and "other." The energy sector consists of cogeneration, electricity generation, and hydrogen plants; the oil and gas sector consists of oil and gas production sites and refineries; all remaining firms—e.g. transportation, warehousing, manufacturing, construction,

and education—are in the "other" sector.[15]

5.2.1 Probit model

Assuming the error follows a normal distribution with constant standard deviation σ, the probability that a firm i in sector j uses offsets in compliance period t is

$$\Pr\left[\mathbb{1}_{ijt}^{O} \mid Gains_{ijt}, j\right] = \Phi\left(\frac{1}{\sigma} Gains_{ijt} - \sum_{j=1}^{3} \frac{\rho_j}{\sigma} \mathbb{1}_i^j\right). \tag{2}$$

The sector-specific intercept ρ_j/σ represents the baseline propensity for firms in sector j to use offsets. The coefficient on $Gains_{ijt}$ is normalized to 1; this implicitly assumes that there is no discounting of the cost savings earned by using offsets—i.e., the value of a dollar is a dollar. If firms systematically discount these savings (e.g., due to invalidation risk), the transaction cost estimates will be inflated. This logic is akin to the "travel cost" coefficient in random utility models in the recreation demand literature (Haab & McConnell, 2002).

Under this assumption, the recovered coefficient is just $1/\hat{\sigma}$, identifying the scale parameter. In this model, $\dfrac{\hat{\rho}_j/\hat{\sigma}}{1/\hat{\sigma}} = \hat{\rho}_j$ is the estimated minimum lump-sum cost-savings that will induce the average firm in sector j to enter the offsets market—i.e., the sector's average fixed participation cost. Because this is an *average* cost, it can be biased by outliers.

5.2.2 Binary quantile response

The probit estimates parameters fit the data at the mean under the assumption of a zero conditional mean error (Cameron & Trivedi, 2005). Naegele (2018) reveals substantial heterogeneity at the upper and lower ends of the error distribution, thus a mean-based regression model may mask important heterogeneity in fixed transaction costs in an emissions trading system. If Californian firms have the same latent distributions of their transaction costs, the

[15] CARB assigns each participating firm one of these sectors. I use these definitions because they directly represent a firm's compliance regulation.

BQR will highlight this.[16]

Following her approach, I employ binary quantile response (BQR) econometric methods based on Kordas (2006), which instead fit the conditional *quantile* (e.g. median) of the distribution of costs. This is a useful exercise especially if there are concerns about heteroskedasticity in the error terms. The semi-parametric BQR is more robust to non-symmetric error distributions and outliers by estimating the conditional distribution of transaction costs rather than the average.

The BQR model also allows me to determine the degree to which the distribution of transaction costs are asymmetric. Perfectly symmetric distributions have equal averages and medians. As such, if the predicted mean and median transaction costs are significantly different, this suggests an asymmetric transaction cost distribution.

Originally proposed for the median by Manski (1975) and generalized by Horowitz (1992), the method maximizes the following objective function

$$\max_{\beta(\tau),\rho(\tau)} \sum_{i=1}^{N} \left[\mathbb{1}_{ijt}^{O} - (1-\tau)\right] \cdot K\left(\frac{\beta(\tau)p_t\tilde{q}_{ijt} - \sum_{j=1}^{3}\rho_j(\tau)\cdot\mathbb{1}_i^j}{h_N}\right) \tag{3}$$

where $\tau \in (0,1)$ is the targeted quantile, K is the Gaussian kernel and $h_N = \frac{0.9}{N^{1/5}}$ is a rule-of-thumb bandwidth (Alejo & Montes-Rojas, 2025) with N observations (firm×period pairs).

The coefficients are functions of τ because the objective function is separately optimized for every quantile. The coefficients are defined so that they consistently predict offset adoption for observations with the $\tau \cdot N$-highest errors. At the median ($\tau = 0.5$), the transaction cost represents the implicit cost threshold for a firm with median unobserved characteristics.

A firm at the 75$^{\text{th}}$ quantile ($\tau = 0.75$) is more intrinsically disposed toward offset use than the median firm, perhaps due to unobserved organizational factors. Because these firms

[16] However, the nature of the econometric problem is slightly different. Naegele (2018) used the amount of compliance instruments needed, as the difference between a firm's allocation and its emissions to calculate her righthand side variables. Allocation data is not available for California firms, so I use their total emissions.

are already favorably disposed, a higher transaction cost is needed in order to rationalize their position at the margin of adoption—i.e., to switch from $\mathbb{1}_{ijt}^{O} = 0$ to $\mathbb{1}_{ijt}^{O} = 1$. In general, one would anticipate that transaction costs should increase with τ, reflecting that higher quantiles correspond to firms with increasingly offset-favorable unobserved characteristics.

The BQR allows me to study heterogeneity in transaction costs, but its true necessity for Naegele (2018) may be driven by certain institutional aspects of the EU ETS that are absent in the California cap-and-trade regime. First, the proportion of EU emissions that can be met with offsets varies at the country-level. In effect, two EU firms with equal emissions can have different maximum possible savings depending on their country, but this might raise endogeneity concerns if firms relocate based on the attractiveness of offsets. This is likely an issue only in the long run. If higher-emitting firms are also located in countries with less restrictive offset policies, there will be a concentration of the maximum cost savings in those countries and the transaction cost from mean-based regression models will be biased toward their transaction cost, which may be different from the median firm's transaction cost. If firm size is endogenous to firm location, this may be an issue. Her data also show a more significant jump in offset usage from smaller to larger quantiles of maximum possible offset use. In the CARB setting, all firms face an identical maximum offset proportion, so maximum possible cost savings vary only on the price and emissions dimension. Figure 4 shows that offset use is increasing in emissions, but the jump in adoption is not as drastic as in the EU (seen in figure 4a in Naegele (2018)).

5.3 Threats to exogeneity

Unobserved heterogeneity The payoff to firm i could be

$$y_{ijt}^{*} = \beta Gains_{ijt} - \sum_{j=1}^{3} \rho_j \mathbb{1}_i^j + \alpha_i + \epsilon_{ijt},$$

where α_i represents the firm's baseline propensity to use offsets, as influenced by relationships with offset suppliers among other things.

If $\alpha_i = 0$, there is no firm-specific heterogeneity in cost savings gained from using offsets. In this case the expected transaction cost reduces to ρ_j/β, as before. If, however, $\alpha_i \neq 0$, this is not the case. For instance, if maximum cost savings (a function of emissions) are correlated with a firm's compliance infrastructure, one would need to rely on within-firm heterogeneity to estimate the parameters.

In a linear model, this heterogeneity can be "differenced out" by including firm fixed effects; in a nonlinear model, estimating firm fixed effects leads to the incidental parameters problem. If such a correlation exists in the offsets market, the coefficient on $Gains_{ijt}$ will be biased. I explore this issue by running auxiliary "Mundlak" correlated random effects (CRE) models, and comparing the resultant transaction costs to the standard probit models.

The tests work as follows. Following Mundlak (1978), I relax the independence assumption using a correlated random effects (CRE) framework and model the heterogeneous effect as a linear function of a vector of the firm's characteristics averaged over time:

$$\alpha_i = \delta_0 + \delta_1 \overline{\boldsymbol{x}}_i + u_i$$

where $\overline{x}_i = \frac{1}{T}\sum_{t=1}^{T} x_{ijt}$ is the time average of the explanatory variable x, $\overline{\boldsymbol{x}}_i$ is the vector of these, and $\eta_i \sim N(0, \sigma_u^2)$ is uncorrelated with firm sector or offset price. Under this structural assumption, the influence of unobserved heterogeneity on β is removed. One interpretation of the random effects model is that each firm has its own intercept which is drawn from a random distribution with mean 0 and variance σ_u^2, where σ_u is estimated from the data.

The transaction costs implied by these auxiliary models (presented in table A-2) are larger than those implied by the standard probit. At the same time, the estimated variance attributed to the heterogeneous random error (i.e., u_i) has a value ranging between 0.002 and 1.1. As such, the gains in terms of identification are relatively small. Since the expected

value of u_i is zero by assumption, a small $\hat{\sigma}_u$ and insignificant coefficient on the firm-average variable together imply that any unobserved heterogeneity is relatively ineffectual. Further, because these CRE estimates of transaction cost are larger in magnitude than the standard probit, the estimated welfare losses calculated using the smaller probit estimates can be interpreted as a lower bound.

Endogeneity of lagged offset adoption The persistence in offset use shown in table 3 suggests that firms continue using offsets once using them initially. Because of this, including lagged offset adoption, $\mathbb{1}^O_{ij,t-1}$, as an explanatory variable improves predictive performance. At the same time, this opens up an endogeneity problem if the errors are serially correlated.

If this is the case, the lagged indicator will be mechanically correlated with the idiosyncratic error. For example, if the serial correlation follows an auto-regressive process with just one lag term, $\text{Cov}\left(\mathbb{1}^O_{ij,t-1}, \epsilon_{ijt}\right) > 0$—endogenous. The result is that the error in the starting period affects the parameters of all future errors; Wooldridge (2005) calls this an "initial conditions problem."

To explore the prevalence of this problem, I include each firm's initial offset use indicator, $\mathbb{1}^O_{ij1}$, in the Mundlak specification in an auxiliary model. The estimated transaction costs from this model are in table A-4 and the coefficients are in table A-5. The coefficients on both initial offset use and on previous offset use are statistically significant at the 1% level. Previous offset use has a negative coefficient, but initial offset use is positive and about twice as large in magnitude. The transaction cost estimates in this model are significantly larger than the standard probit, suggesting again that the baseline model's welfare losses should be a lower bound.

6 Results

6.1 Transaction cost estimates

Baseline specification The baseline specification is a standard probit model where the outcome is a binary indicator for whether a firm uses offsets, and the key explanatory variable is the maximum potential gains from offset use. I augment this baseline with sector indicators, a returning user indicator (equal to one if the firm has used offsets in a prior compliance period), and their interaction, to assess heterogeneity in transaction costs across firm types.

I emphasize that this econometric exercise is not intended to establish causal identification—its sole purpose is to produce a reasonable point estimate of transaction costs for back-of-the-envelope welfare calculations. The qualitative conclusions of the paper rest on the descriptive patterns in the raw data, and the welfare estimates are presented as illustrative magnitudes. The transaction cost is identified from the ratio of the intercept to the coefficient on maximum possible savings, and I report separate estimates for first-time and returning users using the coefficient on returning users.

While the returning user coefficient may partly reflect persistent unobserved heterogeneity rather than the true effects of returning, both estimates are identified from the same maximum cost savings coefficient—the offset-incumbency shifts the intercept but does not affect the slope. The welfare calculations are therefore robust to alternative interpretations of the incumbency effect.

With only three compliance periods, a fully dynamic panel approach that would separately identify true state dependence from unobserved heterogeneity is infeasible. As a robustness check I include firm-level means of maximum cost savings following Mundlak (1978) to partial out time-invariant firm heterogeneity. The resulting transaction cost estimates are somewhat smaller, suggesting the baseline may modestly overstate the true cost. I therefore treat the baseline as the preferred specification and note that the Mundlak results

imply the welfare calculations are, if anything, an upper bound on the gains from trade and a lower bound on the inefficiencies caused by transaction costs.

Average transaction costs Table 4 presents average transaction cost, split by whether a firm has used offsets before.[17] The most basic specification, in column 1, does not separately identify first time offset users from returning users. As such, the estimated transaction cost for both is the same. This implies that a representative firm pays almost $260,000 on average every compliance period that it uses offsets. Column 2 estimates separate transaction costs for each of the three CARB-identified sectors. The estimates imply that there is not strong evidence that, without controlling for previous offset use, the energy sector pays any substantial transaction cost. Oil and gas firms pay an average $267,852 in transaction costs, and other firms pay $562,347; however, the oil and gas sector's transaction cost is only significant at the 10% level.

The third column estimates transaction costs separately for returning offsets users, for the combined sectors. The transaction cost for returning offset users is negative (-$400,009) and significant at the 5% level, meaning returning users would need to be paid this much in order to be convinced *not* to use offsets. This suggests that firms that have used offsets before incur an additional cost by not using them again, either as sunk human capital investments, (potential) reputational costs, or something else. That the transaction cost for first time offset users doubles to $510,111 when this variable is included is a composition effect—the column 1 estimate is a weighted average of first time and returning user costs, and, since returning user costs are negative, the average is pulled down.

The fourth column estimates transaction costs for each sector, separately for returning and first time offset users. The transaction cost for first time offset users is $220,604 (significant at the 10% level) in the energy sector, $426,849 (significant at the 5% level) in the oil and gas sector, and $835,677 (significant at the 1% level) for all other firms. While first time users show an increasing transaction cost across these sectors, respectively, returning users

[17] The coefficients underlying these calculations are presented in tables A-1

Table 4: Mean transaction cost by sector ($1k)

	(1)	(2)	(3)	(4)
Combined (first time)	259.622***		510.111***	
	(54.127)		(119.157)	
Energy (first time)		-59.929		220.604*
		(104.723)		(125.180)
Oil and gas (first time)		267.858*		426.849**
		(148.727)		(190.025)
Other (first time)		562.347***		835.677***
		(132.808)		(227.756)
Combined (returning)	259.622***		-400.009**	
	(54.127)		(196.118)	
Energy (returning)		-59.929		-697.256**
		(104.723)		(306.769)
Oil and gas (returning)		267.858*		-491.012*
		(148.727)		(296.963)
Other (returning)		562.347***		-82.183
		(132.808)		(185.473)
Obs.	556	556	556	556
Model	Probit	Probit	Probit	Probit
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	No	No	No	No
p returning			0.002	0.004
p sector heterogeneity		0.009		0.039

These are average transaction costs calculated using a standard probit specification. Not accounting for unobserved heterogeneity via a Mundlak CRE significantly increases the transaction cost; see table A-2 for these transaction cost calculations, and see table A-1 for the coefficients underlying this table's calculations. Columns 1 and 3 calculate economy-wide average transaction costs. These models partial out firm fixed effects, but otherwise column 1 provides non other controls and column 3 provides a dummy variable indicating whether a firm has previously used offsets. Column 2 includes just sector dummy variables, and column 4 includes sector indicators and returning offset-user indicator variables. The transaction cost in column 1 is calculated by dividing the (negative of the) sum of the intercept by the price coefficient. To calculate returning offset-user and sector transaction costs, divide the relevant intercept(s) by the price coefficient. In models with sector heterogeneity, there is no excluded sector as there is no intercept.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

show a decreasing transaction cost, but with less statistical power. In particular, returning offset users in the energy sector, the oil and gas sector, and all other firms have transaction costs of -$697,256 (significant at the 5% level), -$491,012 (significant at the 10% level), and -$82,183 (insignificant at the 10% level), respectively.

One explanation for the negative transaction costs for returning users is that they represent a form of sunk cost investment in offset market participation. This may be the case if firms that have previously used offsets may have accumulated human capital—e.g. expertise in navigating this market—that makes future non-participation costly. Alternatively, it may reflect reputational commitments to offset suppliers or internal organizational investments. While firms can buy fungible offsets on the spot market, getting favorable prices by making long-term contracts with the offset project registries or the offset suppliers themselves could lead to additional fees for breaking contract. While none of these are observable in the current data, it remains a promising direction for future research.

Quantile transaction cost estimates Table 5 presents transaction cost estimates across the distribution of offset propensity, comparing the probit average to BQR estimates at the 10^{th} and 90^{th} quantiles as well as the interquartile range. The results suggest some heterogeneity that the probit average obscures.[18]

At the median ($\tau = 0.5$), first-time and returning user transaction costs are $62,360 and $25,530 respectively, both significant at the 1% level. While the returning user transaction cost remains smaller than the first-time cost at the median—qualitatively consistent with the probit—they are positive rather than negative, suggesting the probit's negative average is driven by a small number of firms with strongly negative costs rather than a systematic pattern across the distribution.

At the first quartile ($\tau = 0.25$), the estimated transaction costs are large in magnitude but statistically insignificant, and should not be given a structural interpretation. This is

[18]The coefficients and transaction costs for the full distribution are presented graphically, in figures A-1 and A-2, respectively.

an artifact of the coefficient on maximum possible savings being near zero at this quantile, meaning the transaction cost ratio—which divides by this coefficient—is undefined in the limit. For the same reason, the inflated standard errors reflect this near-zero denominator rather than genuine uncertainty about a large transaction cost.

At $\tau = 0.1$, the transaction costs are negative, but statistically insignificant. A negative transaction cost suggests that these low-propensity firms make decisions based on factors beyond the maximum cost savings, such as existing business relationships or (un)willingness to adopt new cost-minimizing technologies. At the upper quantiles ($\tau = 0.75$ and $\tau = 0.9$), transaction costs are small and statistically indistinguishable from zero. This suggests firms with the highest propensity to use offsets face lower transaction costs. Such an interpretation is consistent with the earlier notion that high-propensity firms are experienced participants for whom the fixed costs of offset market participation have already been absorbed. At both tails of the distribution, firms appear to make participation decisions based on factors beyond observable cost savings—at $\tau = 0.1$, these factors push *against* participation despite low costs, while at $\tau = 0.75$ and $\tau = 0.9$ they *facilitate* participation at negligible cost.

Together, these results suggest that transaction costs are most relevant for the firm with a moderate propensity to use offsets, and the probit average is driven primarily by this middle portion of the distribution. The welfare calculations therefore use the median estimates as a conservative alternative to the probit, and note that the probit average likely reflects a small number of high-cost firms rather than a representative participation cost.

While a formal test of homogeneity across quantiles fails to reject equality ($p = 0.302$), this test has low power at $\tau = 0.25$ due to the near-zero denominator inflating standard errors, making the overall test conservative. The pattern of significant estimates at the median and near-zero estimates at the tails nonetheless means one cannot reject the existence of heterogeneity across the transaction cost distribution.

The combined estimates mask heterogeneity across sectors that the probit results suggest

Table 5: Transaction cost by sector ($1k)

	(1) Probit	(2) $\tau = 0.1$	(3) $\tau = 0.25$	(4) $\tau = 0.5$	(5) $\tau = 0.75$	(6) $\tau = 0.9$
Combined (first time)	510.11***	-3.96*	-362.65	62.29***	21.75	0.90
	(119.16)	(2.31)	(3,809.25)	(9.86)	(24.94)	(0.75)
Combined (returning)	-400.01**	-0.98	-171.67	26.40***	2.22	0.12
	(196.12)	(2.27)	(1,819.47)	(8.69)	(7.59)	(1.01)
Obs.	556	556	556	556	556	556

Probit uses robust standard errors; BQR standard errors recovered through bootstrap with 500 repitions. Testing for significantly different transaction costs at the targeted quantiles yields a p-value of 0.302.

*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

is meaningful. Table B-9 disaggregates the quantile estimates by sector.[19]

At the median ($\tau = 0.5$), the sector-specific results broadly mirror the combined estimates. First-time transaction costs are significant at the 1% level across all three sectors, at $48,960, $52,000, and $66,220 for the energy, oil and gas, and other sectors respectively. The ordering is consistent with the probit results—other firms face the highest transaction costs—though the magnitudes are considerably smaller than the probit averages, again consistent with the probit mean being pulled upward by a small number of high-cost firms. Returning user transaction costs at the median are smaller than first-time costs in all three sectors, but are only statistically significant for other firms ($31,430, significant at the 1% level). The energy and oil and gas returning user estimates are positive but insignificant, suggesting the negative returning user costs in the probit reflect tail behavior rather than a systematic effect distinguishable at the median of the offset propensity distribution.

At the third quartile ($\tau = 0.75$), the pattern is more nuanced. First-time transaction costs remain positive for other firms ($35,330, significant at the 5% level) but are statistically indistinguishable from zero for energy and oil and gas firms. Returning user costs at this quantile are negative for energy and oil and gas firms, though insignificant, while other firms

[19]The coefficients and transaction costs for the full distribution are presented graphically, in figures A-3 and A-4, respectively.

show a positive and significant returning user cost ($16,200, significant at the 5% level). This suggests that the "other" sector drives much of the heterogeneity in the upper tail of the distribution.

The first quartile's estimates ($\tau = 0.25$) should again not be given a structural interpretation due to the near-zero denominator. At $\tau = 0.1$ and $\tau = 0.9$, most estimates are small and insignificant, with the exception of returning user costs for energy ($1,630, significant at the 5% level) and oil and gas firms ($970, significant at the 10% level) at $\tau = 0.1$. These are the only cases where returning user costs are positive and significant in the lower tail, suggesting that in these sectors a small number of low-propensity returning users face genuine positive costs to re-entering the offset market.[20]

The formal tests of homogeneity across quantiles tell a cleaner story than the combined results. One can reject the null hypothesis that transaction costs are homogeneous for the energy and oil and gas sectors across quantiles ($p = 0.002$ and $p = 0.003$ respectively), while the other sector fails to reject homogeneity ($p = 0.163$). This is somewhat counterintuitive given that the other sector shows significant estimates at both the median and the third quartile—the failure to reject for other firms likely reflects the wide standard errors at $\tau = 0.25$ dominating the test statistic. Together, these sector-specific results confirm that the median is the most informative point of the distribution for welfare calculations, and that the probit sector averages mask meaningful heterogeneity particularly in the energy and oil and gas sectors.

6.2 Welfare

Table 7 presents the components underlying the welfare calculations. For each sector, I report the total transaction costs incurred by offset users, the total transaction costs avoided by non-users, the gross savings earned by users, and the gross savings missed by non-users.

[20] A potential explanation for this is that in such capital-intensive sectors, even low-propensity returning users face non-trivial administrative or compliance costs associated with re-entering the offset market, costs that are absorbed by high-propensity firms but remain binding for occasional participants. However, it could be statistical noise.

Table 6: Transaction cost by sector ($1k)

	Probit	$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$	$\tau = 0.75$	$\tau = 0.9$
Energy (first time)	220.60*	-0.40	1,172.07	47.44***	-1.14	-1.02
	(125.18)	(1.22)	(23,759.94)	(9.34)	(8.10)	(1.94)
Oil/gas (first time)	426.85**	-1.08	1,341.39	50.64***	4.26	0.14
	(190.03)	(0.99)	(27,207.17)	(10.45)	(10.73)	(2.41)
Other (first time)	835.68***	-2.51*	1,555.65	64.72***	34.75**	8.55
	(227.76)	(1.46)	(31,591.67)	(9.67)	(14.02)	(9.18)
Energy (returning)	-697.26**	1.63**	428.26	13.90	-19.26	-1.56
	(306.77)	(0.70)	(8,626.73)	(13.19)	(20.62)	(2.52)
Oil/gas (returning)	-491.01*	0.95*	597.58	17.10	-13.86	-0.41
	(296.96)	(0.56)	(12,074.67)	(13.24)	(17.47)	(2.76)
Other (returning)	-82.18	-0.48	811.84	31.18***	16.63***	8.00
	(185.47)	(1.06)	(16,455.95)	(7.71)	(6.22)	(9.81)
Obs.	556	556	556	556	556	556

Probit uses robust standard errors; BQR standard errors recovered through bootstrap with 500 repitions. Testing for significantly different transaction costs at the targeted quantiles for the energy, oil/gas, and other sectors yields a p-value of 0.001, 0.003, and 0.159, respectively. The transaction cost estimates and coefficients are presented at each ventile in figures A-4 and A-3, respectively. Figure A-6 presents the estimated offset adoption rates compared with the observed rates as a goodness-of-fit exercise (Kordas, 2006).

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table 7: Numbers used in welfare calculations

	Energy	Oil and gas	Other
Incurred transaction cost (−)	$4,269,730	$1,874,800	$4,123,560
Avoided transaction cost (+)	$293,340	$2,031,830	$10,685,410
Gross earned savings (+)	$58,574,056	$99,465,039	$14,535,722
Gross missed savings (−)	$9,932,814	$7,894,877.75	$10,452,469
Net achieved savings (+)	$54,304,326	$97,590,239	$10,412,162
Net money left on the table (−)	$6,993,474	$5,863,048	-$232,941
Value of offsets market	**$47,310,852**	**$91,727,191**	**$10,645,103**
Frictionless cost savings (hypothetical)	$68,506,870	$107,359,917	$24,988,191

Next to each welfare component is an indicator in parantheses. A (+) indicates that the component contributes positively to overall cost savings, and a (−) indicates that the component contributes negatively to overall cost savings. Incurred transaction cost is the total transaction cost paid by all firms that use offsets within a sector; avoided transaction cost is the total value of the transaction cost that *would have* been paid by firms, but was avoided because they did not use offsets; gross earned savings is the total value of cost savings gained by users of offsets; and gross missed savings is the gross cost savings that firms could have earned, but did not because they did not use offsets. The net achieved savings is the difference between gross earned savings and incurred transaction costs, and net money left on the table is the difference between gross missed savings and avoided transaction cost. The total value of a market is the difference between net achieved savings and net money left on the table. The frictionless cost savings assumes zero transaction cost and full participation, so it is the sum of gross earned savings and gross missed savings.

From these components I construct net achieved savings, net money left on the table, and the total value of the offset market, which I compare to the frictionless counterfactual with full market participation.

Across all three sectors, firms that used offsets realized substantial gross savings—approximately \$58.6M, \$99.5M, and \$14.5M for the energy, oil and gas, and other sectors respectively. After subtracting incurred transaction costs, net achieved savings are \$54.3M, \$97.6M, and \$10.4M. The net money left on the table—gross missed savings minus avoided transaction costs—is \$7.0M and \$5.9M for the energy and oil and gas sectors, but is negative for the other sector (-\$232,941). This last figure suggests that non-participating firms outside of the energy and oil and gas sectors made the correct decision given the median estimated transaction costs they would have faced, since the savings they missed were smaller than the transaction costs they avoided.

The total value of the offset market—which can be interpreted as the total net cost savings—is \$47.3M, \$91.7M, and \$10.6M for the energy, oil and gas, and other sectors respectively, for a combined \$149.7M.

A natural counterfactual is a market with no transaction costs. In this setting, because there are no frictions, all firms participate in the offsets market, and retain their full gross savings. This yields \$68.5M, \$107.4M, and \$25M for the energy, oil and gas, and other sectors, respectively. Removing transaction costs and inducing full participation—ceteris paribus—can only (weakly) increase total gross savings. As such, the gap between the frictionless cost savings and the actual cost savings can be interpreted as the inefficiency induced by the transaction costs: \$21.2M in energy, \$15.6M in oil and gas, and \$14.4M for all others.

The size of this efficiency cost varies across sectors, reflecting two offsetting forces. On one hand, transaction costs reduce net savings of participating firms. On the other hand, they also reduce the gross missed savings attributed to non-participants, since non-participants do not pay transaction costs. The net efficiency cost therefore depends on the relative size of

these two forces in a given sector. In the other sector, the avoided transaction cost ($10.7M) is actually larger than the incurred transaction cost ($4.1M), meaning non-participants actually avoided more in transaction costs than they nominally left on the table. This partially offsets the efficiency cost, which is why the money left on the table is negative and the gap between frictionless and actual welfare ($14.4M) is smaller than the raw transaction cost figures may suggest.

The offset market generates substantial positive welfare in all three sectors, and transaction costs represent a meaningful, but not dominant, share of total potential gains. Approximately 31%, 15%, and 58% of frictionless welfare in the energy, oil and gas, and other sectors, respectively, is lost due to transaction costs. These figures can be interpreted as an upper bound on the efficiency lost from transaction costs, since the frictionless counterfactual assumes both full participation and fixed offset prices. It is possible that there may not be full participation in the frictionless market, especially if the quantity of offsets is limited or prices are allowed to vary. The welfare calculations nonetheless confirm that the California offset market generates large net benefits even accounting for participation costs, and that reducing transaction costs—particularly in the energy and other sectors where the efficiency cost is largest—would yield meaningful welfare gains.

Several caveats apply to these welfare calculations. First, the frictionless counterfactual assumes that all non-participating firms would enter the offset market if transaction costs were zero, and that they would earn their maximum potential savings. Both assumptions are likely too strong. Some firms may face barriers to participation beyond transaction costs—regulatory constraints, lack of awareness, or organizational inertia—and potential savings are estimated from observable firm characteristics rather than realized offset purchases, so that gross savings represent an upper bound on what firms could actually earn. Together these mean the frictionless welfare figures should be interpreted as upper bounds on the gains from eliminating transaction costs.

Second, the transaction cost estimates themselves carry uncertainty. The welfare calcu-

lations use the median binary quantile regression estimates, which are identified from the middle of the offset propensity distribution and may not be representative of transaction costs faced by firms at the tails. In particular, the large number of non-participating other sector firms that drive much of the other sector welfare result are likely in the lower tail of the propensity distribution, where transaction costs are poorly identified. Sensitivity of the welfare results to alternative transaction cost estimates—including the probit averages and the Mundlak-corrected estimates—is therefore an important check. That the range of transaction cost estimates varies somewhat substantially across specifications provides a sense of the uncertainty in the welfare figures; using the median is the most conservative estimate, as table C-1's significantly larger values illustrates.

Third, the welfare calculations treat potential cost savings as exogenous and fixed across the actual and counterfactual markets. In practice, if transaction costs were eliminated and more firms entered the offset market, the price of offsets will likely adjust, changing the cost savings available to all firms. The calculations therefore implicitly assume a partial equilibrium in which offset prices are unaffected by changes in participation—an assumption that is likely to fall short in practice.

Fourth, the analysis focuses exclusively on the cost savings side of the welfare calculation and does not account for the environmental integrity of offsets. If some offsets do not represent genuine emissions reductions—a concern literature frequently mentioned in the literature—then the gross savings figures overstate the true social value of offset market participation, since firms are purchasing credits that do not deliver the claimed environmental benefit. This caveat applies equally to the actual and frictionless markets and so does not affect the comparison between them, but it does affect the interpretation of the absolute welfare figures.

Finally, the three-period panel covers only the first decade of California's cap-and-trade program. Transaction costs may decline over time as firms accumulate experience with the offset market, which would imply that the efficiency cost estimated here is an upper bound

on the long-run cost. Conversely, if the firms most likely to benefit from offsets have already entered the market by the third compliance period, the pool of potential entrants in future periods may have lower average savings, which would reduce the frictionless welfare gains from eliminating transaction costs going forward.

7 Discussion

The results of this paper establish that fixed transaction costs are a meaningful barrier to offset adoption in California's cap-and-trade program, costing firms a combined $10.3 million in median transaction costs over the first three compliance periods and generating roughly $51 million in efficiency losses relative to a frictionless market. But the nature of these costs remains an open question—and the answer has real policy implications.

Three candidate explanations—search costs, invalidation risk, and public relations concerns—have different policy remedies. Search costs can be reduced through better market infrastructure, such as centralized registries and standardized contracts. Invalidation risk can be mitigated through improved verification standards or buffer pool mechanisms. Public relations costs, however, are not easily engineered away. If firms are reluctant to use offsets because they fear being perceived as "greenwashing," then even a perfectly functioning offset market may underperform.

Recent empirical evidence lends credence to the public relations explanation. Studies documenting widespread non-additionality in offset programs—including California's own forest offset protocol (Badgley, Chay, et al., 2022)—have attracted significant media attention. In an environmentally conscious state like California, the reputational risk of being associated with low-quality offsets may be a rational concern for regulated firms. The persistence of offset use documented in table 3 can represent such public relations costs as well as it can sunk search costs or employee hiring. Once a firm's offset purchasing is publicly revealed, the reputational cost is sunk, and continued use carries no additional stigma.

This interpretation suggests a somewhat uncomfortable conclusion: efforts to publicize the limitations of offsets, however empirically justified, may have the unintended consequence of reducing participation in a market that still generates substantial net welfare gains. Moreover, the transaction cost is incurred by any offset-using firm, regardless of the environmental benefits gained. As such, the purchase of non-additional offsets when a transaction cost exists, acts as pure welfare-decreasing action. If some of the transaction cost is a result of true non-additionality affecting social image, then it can be a socially-beneficial aspect of the market. On the other hand, it is still paid by participants. Policymakers seeking to improve offset market efficiency may therefore first focus on integrity reforms with active communication efforts—this improves additionality and lowers transaction costs jointly.

References

Aldy, J. E., & Halem, Z. M. (2024). The evolving role of greenhouse gas emission offsets in combating climate change. *Review of Environmental Economics and Policy*, *18*(2), 212–233. https://doi.org/10.1086/730982

Alejo, J., & Montes-Rojas, G. (2025). Quantile regression under limited dependent variable in *Stata*. *Journal of Statistical Software*, *114*(8). https://doi.org/10.18637/jss.v114.i08

Alston, L. J., Andersson, K., & Smith, S. M. (2013). Payment for environmental services: Hypotheses and evidence. *Annual Review of Resource Economics*, *5*(1), 139–159. https://doi.org/10.1146/annurev-resource-091912-151830

Appunn, K., & Wettengel, J. (2024). Understanding the European Union's Emissions Trading Systems (EU ETS) [magazine]. *Clean Energy Wire: Journal for the energy transition*. https://www.cleanenergywire.org/factsheets/understanding-european-unions-emissions-trading-system#:~:text=The%20EU%20ETS%20is%20the%20oldest%20and,report%20by%20the%20International%20Carbon%20Action%20Partnership.

Aspelund, K. M., & Russo, A. (2025). Additionality and asymmetric information in environmental markets: Evidence from conservation auctions. *Working paper*.

Badgley, G., Chay, F., Chegwidden, O. S., Hamman, J. J., Freeman, J., & Cullenward, D. (2022). California's forest carbon offsets buffer pool is severely undercapitalized. *Frontiers in Forests and Global Change*, *5*, 930426. https://doi.org/10.3389/ffgc.2022.930426

Badgley, G., Freeman, J., Hamman, J. J., Haya, B., Trugman, A. T., Anderegg, W. R. L., & Cullenward, D. (2022). Systematic over-crediting in California's forest carbon offsets program. *Global Change Biology*, *28*(4), 1433–1445. https://doi.org/10.1111/gcb.15943

Burtraw, D. (2016). *Why California's cap-and-trade program works*.

Bushnell, J., Chen, Y., & Zaragoza-Watkins, M. (2014). Downstream regulation of CO2 emissions in California's electricity sector. *Energy Policy*, *64*, 313–323. https://doi.org/10.1016/j.enpol.2013.08.065

Cacho, O. J., Lipper, L., & Moss, J. (2013). Transaction costs of carbon offset projects: A comparative study. *Ecological Economics*, *88*, 232–243. https://doi.org/10.1016/j.ecolecon.2012.12.008

Cacho, O. J., Marshall, G. R., & Milne, M. (2005). Transaction and abatement costs of carbon-sink projects in developing countries. *Environment and Development Economics*, *10*(5), 597–614.

Calel, R., Colmer, J., Dechezleprêtre, A., & Glachant, M. (2025). Do carbon offsets offset carbon? *American Economic Journal: Applied Economics*, *17*(1), 1–40. https://doi.org/10.1257/app.20230052

California Global Warming Solutions Act of 2006 (2006, September 27). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=200520060AB32

California Global Warming Solutions Act of 2006: Market-Based Compliance Mechanisms: Fire Prevention Fees: Sales and Use Tax Manufacturing Exemption. (2017, July 25). https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201720180AB398

Cameron, A. C., & Trivedi, P. K. (2005). *Microeconometrics: Methods and applications*. Cambridge University Press.

Caron, J., Rausch, S., & Winchester, N. (2015). Leakage from sub-national climate policy: The case of California's cap-and-trade program. *The Energy Journal*, *36*(2), 167–190. http://www.jstor.org/stable/24695752

Cragg, J. G. (1971). Some Statistical Models for Limited Dependent Variables with Application to the Demand for Durable Goods. *Econometrica*, *39*(5), 829. https://doi.org/10.2307/1909582

Elson, D. (2011, February 27). *Cost-benefit analysis of a shift to a low carbon economy in the land use sector in Indonesia* (Independent report). The UK Climate Change Unit of the British Embassy. Jakarta. https://seventythree.org.uk/wp-content/uploads/2024/03/Elson_Shift-to-Low-Carbon-Economy_UK-Embassy_FINAL.pdf

Fowlie, M., & Perloff, J. M. (2013). Distributing pollution rights in cap-and-trade programs: Are outcomes independent of allocation? *The Review of Economics and Statistics*, *95*(5), 1640–1652. https://doi.org/10.1162/REST_a_00345

Fowlie, M., Petersen, C., & Reguant, M. (2021). Border carbon adjustments when carbon intensity varies across producers: Evidence from California. *AEA Papers and Proceedings*, *111*, 401–405. https://doi.org/10.1257/pandp.20211073

Fowlie, M., Reguant, M., & Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, *124*(1), 249–302. https://doi.org/10.1086/684484

Galik, C. S., Cooley, D. M., & Baker, J. S. (2012). Analysis of the production and transaction costs of forest carbon offset projects in the USA. *Journal of Environmental Management*, *112*, 128–136. https://doi.org/10.1016/j.jenvman.2012.06.045

Haab, T. C., & McConnell, K. E. (2002, May 28). *Valuing Environmental and Natural Resources: The econometrics of non-market valuation*. Edward Elgar Publishing. https://doi.org/10.4337/9781843765431

Hahn, R. W., & Stavins, R. N. (2010). The effect of allowance allocations on cap-and-trade system performance. *NBER Working Paper*. http://www.nber.org/papers/w15854

Hintermann, B. (2010). Allowance price drivers in the first phase of the EU ETS. *Journal of Environmental Economics and Management*, *59*(1), 43–56. https://doi.org/10.1016/j.jeem.2009.07.002

Horowitz, J. L. (1992). A Smoothed Maximum Score Estimator for the Binary Response Model. *Econometrica*, *60*(3), 505. https://doi.org/10.2307/2951582

Jaraitė-Kažukauskė, J., & Kažukauskas, A. (2015). Do transaction costs influence firm trading behaviour in the European Emissions Trading System? *Environmental and Resource Economics*, *62*(3), 583–613. https://doi.org/10.1007/s10640-014-9831-7

Kordas, G. (2006). Smoothed binary regression quantiles. *Journal of Applied Econometrics*, *21*(3), 387–407. https://doi.org/10.1002/jae.843

Krutilla, K. (2011). Transaction costs and environmental policy: An assessment framework and literature review. *International Review of Environmental and Resource Economics*, *4*(3–4), 261–354. https://doi.org/10.1561/101.00000035

Manski, C. F. (1975). Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics*, *3*(3), 205–228. https://doi.org/10.1016/0304-4076(75)90032-9

Merger, E., Held, C., Tennigkeit, T., & Blomley, T. (2012). A bottom-up approach to estimating cost elements of REDD+ pilot projects in Tanzania. *Carbon Balance and Management*, *7*(1), 9. https://doi.org/10.1186/1750-0680-7-9

Mukama, K., Mustalahti, I., & Zahabu, E. (2012). Participatory forest carbon assessment and REDD+: Learning from Tanzania. *International Journal of Forestry Research*, *2012*, 1–14. https://doi.org/10.1155/2012/126454

Mundlak, Y. (1978). On the Pooling of Time Series and Cross Section Data. *Econometrica*, *46*(1), 69. https://doi.org/10.2307/1913646

Naegele, H. (2018). Offset credits in the EU ETS: A quantile estimation of firm-level transaction costs. *Environmental and Resource Economics*, *70*(1), 77–106. https://doi.org/10.1007/s10640-017-0111-1

Pearson, T. R. H., Brown, S., Sohngen, B., Henman, J., & Ohrel, S. (2014). Transaction costs for carbon sequestration projects in the tropical forest sector. *Mitigation and Adaptation Strategies for Global Change*, *19*(8), 1209–1222. https://doi.org/10.1007/s11027-013-9469-8

Potvin, C., Guay, B., & Pedroni, L. (2008). Is reducing emissions from deforestation financially feasible? A Panamanian case study. *Climate Policy*, *8*(1), 23–40. https://doi.org/10.3763/cpol.2007.0386

Reinganum, J. F. (1981). On the diffusion of new technology: A game theoretic approach. *The Review of Economic Studies*, *48*(3), 395. https://doi.org/10.2307/2297153

Singh, R., & Weninger, Q. (2017). Cap-and-trade under transactions costs and factor irreversibility. *Economic Theory*, *64*(2), 357–407. https://doi.org/10.1007/s00199-016-0991-2

Stavins, R. N. (1995). Transaction costs and tradeable permits. *Journal of Environmental Economics and Management*, *29*(2), 133–148. https://doi.org/10.1006/jeem.1995.1036

Tuladhar, S. D., Mankowski, S., & Bernstein, P. (2014). Interaction effects of market-based and command-and-control policies. *The Energy Journal*, *35*, 61–88. https://doi.org/10.5547/01956574.35.SI1.4

USA—California cap-and-trade program. (2024). International Carbon Action Partnership. https://icapcarbonaction.com/en/ets/usa-california-cap-and-trade-program

Wertz-Kanounnikoff, S. (2008, November). *Estimating the costs of reducing forest emissions*. https://webdoc.sub.gwdg.de/ebook/serien/yo/CIFOR_WP/42.pdf

West, T. A. P., Wunder, S., Sills, E. O., Börner, J., Rifai, S. W., Neidermeier, A. N., Frey, G. P., & Kontoleon, A. (2023). Action needed to make carbon offsets from forest conservation work for climate change mitigation. *Science*, *381*(6660), 873–877. https://doi.org/10.1126/science.ade3535

Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics*, *20*(1), 39–54. https://doi.org/10.1002/jae.770

A Regression estimates and alternative specifications

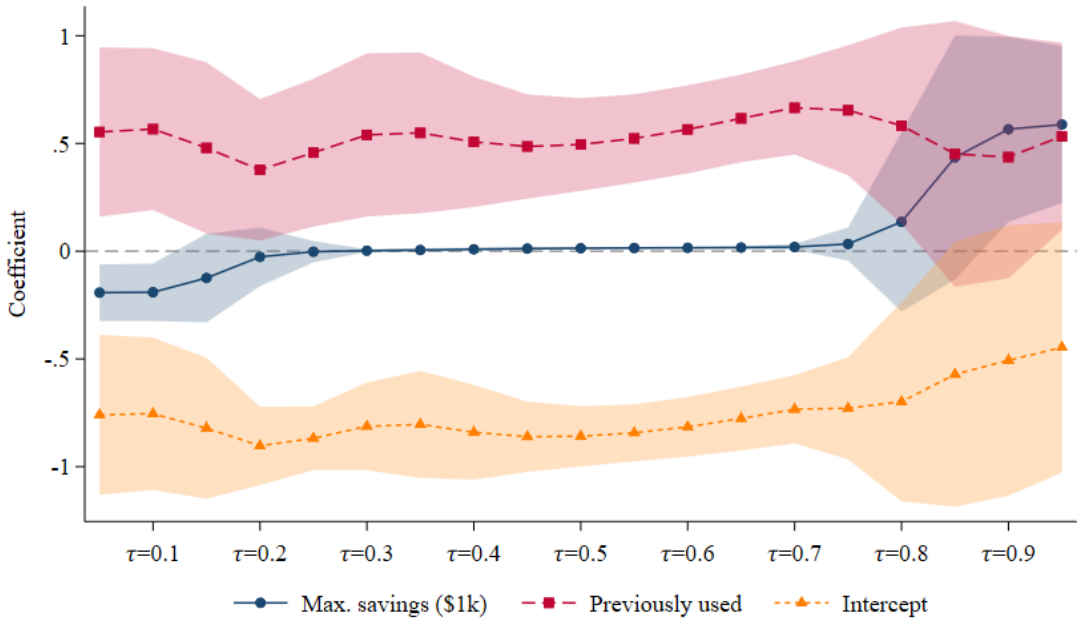


Figure A-1: BQR coefficients—combined model

Figures A-5 and A-6 aggregate firms along maximum cost savings ventiles and compare the predicted probabilities of using offsets against the observed share of firms using offsets within that quantile: A-5 shows the combined model and A-6 shows the predictions by sector.

Table A-1: Coefficients: used in table 4

	(1)	(2)	(3)	(4)
Uses offsets this period				
Constant	-0.403***		-0.451***	
	(0.112)		(0.096)	
Max. benefit ($1k)	0.002***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Firm average max. profit (1k)	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Energy		0.132		0.006
		(0.161)		(0.163)
Oil and gas		-0.307		-0.322
		(0.253)		(0.223)
Other		-0.818***		-0.792***
		(0.156)		(0.137)
Previously used			0.478**	0.345
			(0.225)	(0.224)
$\ln \sigma_u^2$	0.092	-0.052	-0.606	-0.518
	(0.298)	(0.308)	(0.559)	(0.503)
Obs.	556	556	556	556
Model	CRE	CRE	CRE	CRE
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	Yes	Yes	Yes	Yes

These coefficients are used to calculate the average combined and average sector transaction costs presented in table 4.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table A-2: Mean transaction cost by sector ($1k)

	(1)	(2)	(3)	(4)
Combined (first time)	234.948***		337.220***	
	(66.959)		(103.311)	
Energy (first time)		-88.739		-4.918
		(116.465)		(131.903)
Oil and gas (first time)		206.839		259.715
		(167.394)		(182.792)
Other (first time)		551.614***		639.222***
		(151.617)		(189.253)
Combined (returning)	234.948***		-20.455	
	(66.959)		(167.789)	
Energy (returning)		-88.739		-283.018
		(116.465)		(212.081)
Oil and gas (returning)		206.839		-18.386
		(167.394)		(249.901)
Other (returning)		551.614***		361.121*
		(151.617)		(195.804)
Obs.	556	556	556	556
Model	CRE	CRE	CRE	CRE
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	Yes	Yes	Yes	Yes
p returning			0.118	0.217
p sector heterogeneity		0.014		0.022

These are average transaction costs calculated using the Mundlak CRE specification. Accounting for unobserved heterogeneity in this way only dampens the estimated transaction costs; see table 4 for calculations using standard probit methods and table A-3 for the coefficients underlying this table's calculations. Columns 1 and 3 calculate economy-wide average transaction costs. These models partial out firm fixed effects, but otherwise column 1 provides non other controls and column 3 provides a dummy variable indicating whether a firm has previously used offsets. Column 2 includes just sector dummy variables, and column 4 includes sector indicators and returning offset-user indicator variables. The transaction cost in column 1 is calculated by dividing the (negative of the) sum of the intercept by the price coefficient. To calculate returning offset-user and sector transaction costs, divide the relevant intercept(s) by the price coefficient. In models with sector heterogeneity, there is no excluded sector as there is no intercept. The second to last row reports the p-values from tests with a null hypothesis that the transaction costs are identical across all sectors, and the last row reports p-values from tests with a null hypothesis that returning users face identical transaction costs for first time users. Columns 1 and 2 have identical 'first time' and 'returning' transaction costs because these models do not include any controls for returning users.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table A-3: Coefficients: used in table A-2

	(1)	(2)	(3)	(4)
Uses offsets this period				
Constant	-.403***		-.451***	
	(.112)		(.0956)	
Max. benefit ($1k)	.00172***	.00148***	.00134***	.00124***
	(.00039)	(.00037)	(.00036)	(.00036)
Firm average max. profit (1k)	-.00011	-.00012	-5.9e-05	-8.1e-05
	(.00021)	(.0002)	(.00017)	(.00017)
Energy		.132		.00609
		(.161)		(.163)
Oil and gas		-.307		-.322
		(.253)		(.223)
Other		-.818***		-.792***
		(.156)		(.137)
Previously used			.478**	.345
			(.225)	(.224)
$\ln \sigma_u^2$	.0922	-.0518	-.606	-.518
	(.298)	(.308)	(.559)	(.503)
Obs.	556	556	556	556
Model	CRE	CRE	CRE	CRE
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	Yes	Yes	Yes	Yes

These coefficients are used to calculate the average combined and average sector transaction costs presented in table A-2.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table A-4: Mean transaction cost by sector ($1k)

	(1)	(2)	(3)
Energy	1,520***	1,124***	908.09**
	(581.32)	(410.23)	(377.7)
Oil and gas	968.88**	730.35**	1,027**
	(402.87)	(298.32)	(433.95)
Other	1,396***	1,035***	1,201**
	(532.17)	(374.64)	(476.85)
Energy previously used	1,520***	1,822***	2,296***
	(581.32)	(630.94)	(875.79)
Oil and gas previously used	968.88**	1,428***	-258.52
	(402.87)	(510.65)	(782.85)
Other previously used	1,396***	1,733***	1,556**
	(532.17)	(601.63)	(628.38)
Obs.	556	556	556
Previous offset use control?	No	Yes	Yes
Previous use sector interaction	No	No	Yes
Mundlak controls?	Yes	Yes	Yes
Initial conditions?	Yes	Yes	Yes
p-value: Energy		0.014	0.015
p-value: Oil/gas		0.014	0.180
p-value: Other		0.014	0.281

These are all correlated random effects (CRE) models withh initial conditions control included (Wooldridge, 2005). They are calculated by dividing the (negative of the) coefficients for sector by the price coefficient. I test whether previous offset users in each sector have a significantly different transaction cost compared to first-time users. The p-values for these tests for each model are presented in the last three rows fo the table. The energy sector only has a significantly different transaction cost when the previously used identifier is additive (i.e. columns (2) and (3)), but not when the sector is interacted with the identifier. Both oil/gas and other sectors are significantly different transaction costs in both specifications.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table A-5: Estimated coefficients from probit models

	(1)	(2)
Uses offsets this period		
Max. benefit ($1k)	.00106***	.00097**
	(.00039)	(.00039)
Energy	-1.1869***	-.87756***
	(.1722)	(.17671)
Oil and gas	-.77131***	-.99222***
	(.19856)	(.21706)
Other	-1.093***	-1.1609***
	(.12899)	(.1366)
Firm average max. profit (1k)	-.00029	-.00024
	(.00032)	(.00032)
Previously used	-.73696***	
	(.22874)	
Initial offsets use	2.4914***	2.5082***
	(.27759)	(.28718)
Energy previously used		-1.3409***
		(.29731)
Oil and gas previously used		1.242
		(.75603)
Other previously used		-.34317
		(.31505)
$\ln \sigma_u^2$	-2.9409	-3.9169
	(2.2653)	(5.7938)
Obs.	556.00	556.00
Previous offset use control?	Yes	Yes
Previous use sector interaction	No	Yes
Mundlak controls?	Yes	Yes
Initial conditions?	Yes	Yes

These estimated coefficients correspond by column to the transaction cost estimates in table A-4.
$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

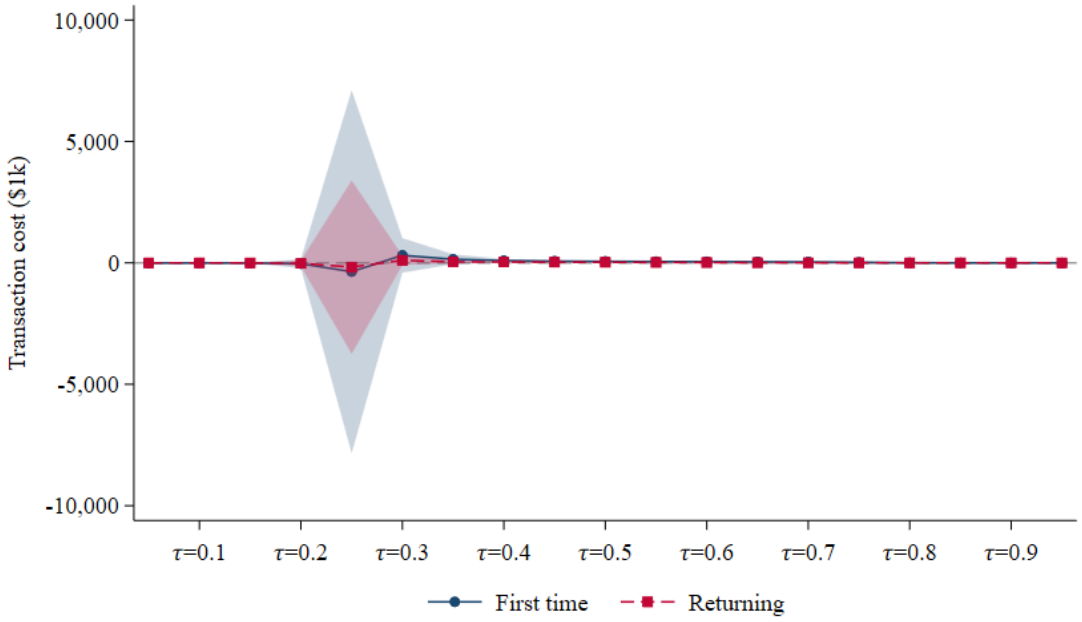


Figure A-2: Combined transaction costs at different quantiles

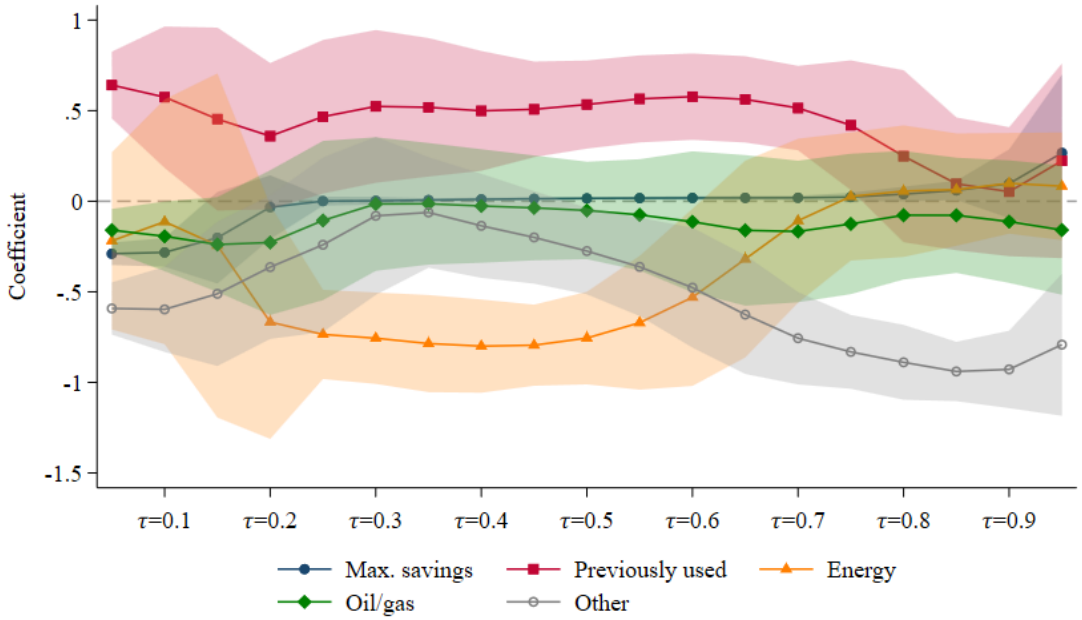
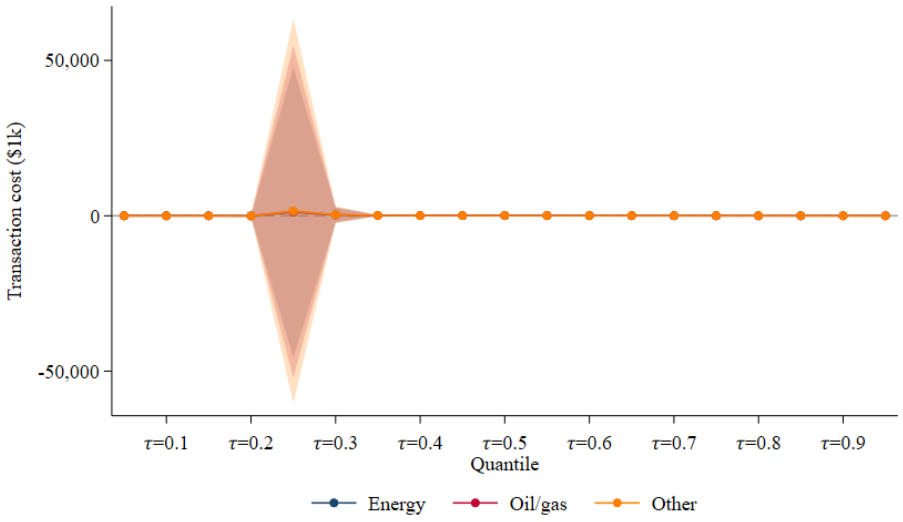
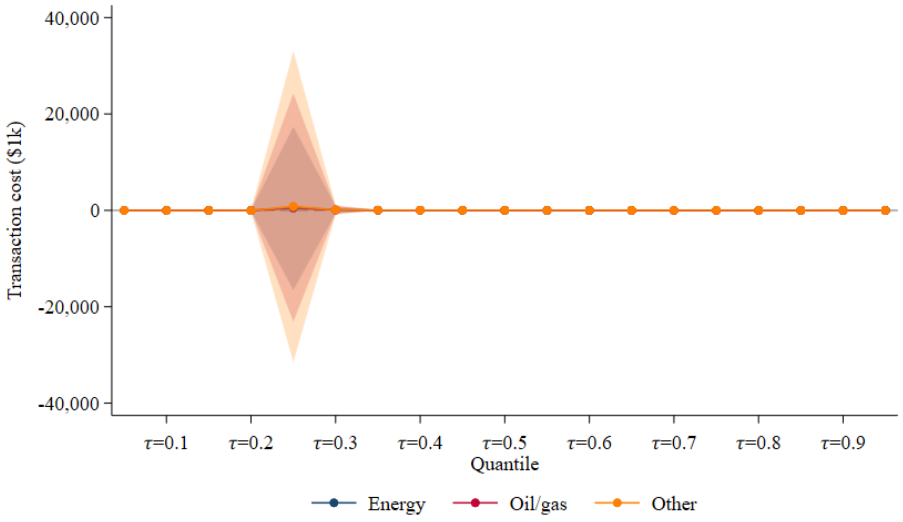


Figure A-3: BQR coefficients—sector-specific model



(a) First time offset users



(b) Returning offset users

Figure A-4: Sector-specific transaction costs at different quantiles

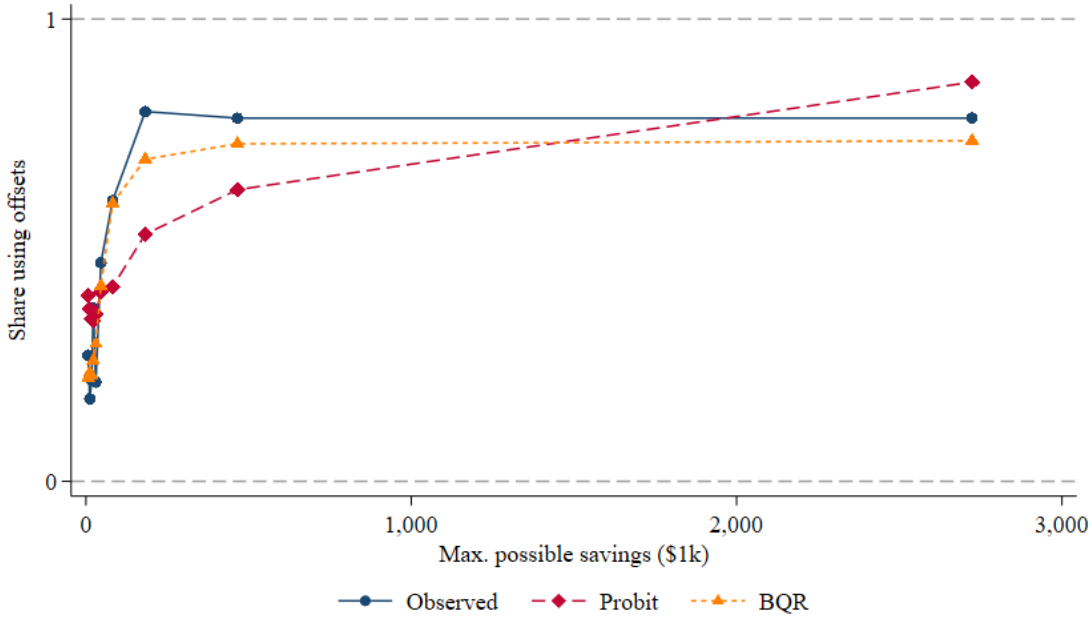


Figure A-5

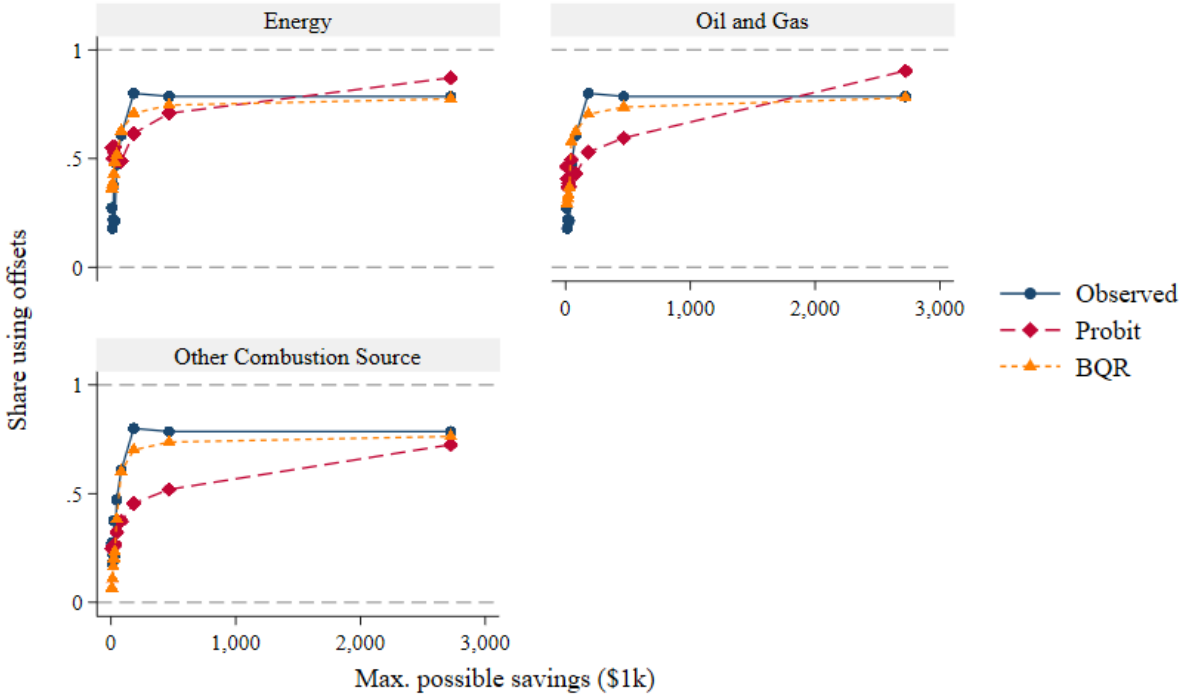


Figure A-6

B Using average prices

As a robustness check, I replicate the main analysis using average compliance period prices rather than end-of-period prices to construct the maximum potential gains variable. Table B-6 matches table 4, table B-7 matches table A-2, table B-8 matches table 5, and table B-9 matches table B-9. Figure B-7 matches figure A-1, figure B-8 matches figure A-2, figure B-9 matches figure A-3, figure B-10 matches figure A-4, figure B-11 matches figure A-5, and figure B-12 matches figure A-6.

Across all specifications, the qualitative findings are unchanged: fixed transaction costs remain the primary barrier to offset adoption, the distribution of costs is right-skewed with the probit mean driven by a small number of high-cost firms, and the median estimates are of similar magnitude to those in the main analysis. Sector-specific heterogeneity in transaction costs is also preserved, with the "other" sector continuing to face the highest costs and the energy sector the lowest. The BQR coefficient plots (figures B-7 and B-9) exhibit patterns nearly identical to their counterparts in appendix A, and the transaction cost estimates at the median (Figures A-2 and A-4) are statistically indistinguishable from the baseline results. This exercise suggest that the choice of price measure does not materially affect the conclusions of the paper.

Table B-6: Mean transaction cost by sector ($1k)

	(1)	(2)	(3)	(4)
Combined (first time)	237.897***		467.010***	
	(49.433)		(106.612)	
Energy (first time)		-55.556		207.374*
		(95.942)		(113.458)
Oil and gas (first time)		246.386*		393.321**
		(136.528)		(172.254)
Other (first time)		515.445***		760.066***
		(121.175)		(203.413)
Combined (returning)	237.897***		-368.104**	
	(49.433)		(175.464)	
Energy (returning)		-55.556		-636.506**
		(95.942)		(274.462)
Oil and gas (returning)		246.386*		-450.559*
		(136.528)		(266.901)
Other (returning)		515.445***		-83.814
		(121.175)		(167.608)
Obs.	556	556	556	556
Model	Probit	Probit	Probit	Probit
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	No	No	No	No
p returning			0.001	0.003
p sector heterogeneity		0.008		0.038

These are average transaction costs calculated using a standard probit specification. Not accounting for unobserved heterogeneity via a Mundlak CRE significantly increases the transaction cost; see table A-2 for these transaction cost calculations, and see table A-1 for the coefficients underlying this table's calculations. Columns 1 and 3 calculate economy-wide average transaction costs. These models partial out firm fixed effects, but otherwise column 1 provides non other controls and column 3 provides a dummy variable indicating whether a firm has previously used offsets. Column 2 includes just sector dummy variables, and column 4 includes sector indicators and returning offset-user indicator variables. The transaction cost in column 1 is calculated by dividing the (negative of the) sum of the intercept by the price coefficient. To calculate returning offset-user and sector transaction costs, divide the relevant intercept(s) by the price coefficient. In models with sector heterogeneity, there is no excluded sector as there is no intercept.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table B-7: Mean transaction cost by sector ($1k)

	(1)	(2)	(3)	(4)
Combined (first time)	212.832***		314.135***	
	(61.409)		(95.204)	
Energy (first time)		-80.526		8.155
		(105.801)		(119.462)
Oil and gas (first time)		188.633		242.587
		(151.844)		(165.888)
Other (first time)		499.250***		584.744***
		(140.130)		(174.561)
Combined (returning)	212.832***		-43.495	
	(61.409)		(152.868)	
Energy (returning)		-80.526		-276.949
		(105.801)		(194.298)
Oil and gas (returning)		188.633		-42.517
		(151.844)		(227.534)
Other (returning)		499.250***		299.640*
		(140.130)		(177.809)
Obs.	556	556	556	556
Model	CRE	CRE	CRE	CRE
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	Yes	Yes	Yes	Yes
p returning			0.089	0.170
p sector heterogeneity		0.016		0.025

These are average transaction costs calculated using the Mundlak CRE specification. Accounting for unobserved heterogeneity in this way only dampens the estimated transaction costs; see table B-6 for calculations using standard probit methods and average prices. Columns 1 and 3 calculate economy-wide average transaction costs. These models partial out firm fixed effects, but otherwise column 1 provides non other controls and column 3 provides a dummy variable indicating whether a firm has previously used offsets. Column 2 includes just sector dummy variables, and column 4 includes sector indicators and returning offset-user indicator variables. The transaction cost in column 1 is calculated by dividing the (negative of the) sum of the intercept by the price coefficient. To calculate returning offset-user and sector transaction costs, divide the relevant intercept(s) by the price coefficient. In models with sector heterogeneity, there is no excluded sector as there is no intercept. The second to last row reports the p-values from tests with a null hypothesis that the transaction costs are identical across all sectors, and the last row reports p-values from tests with a null hypothesis that returning users face identical transaction costs for first time users. Columns 1 and 2 have identical 'first time' and 'returning' transaction costs because these models do not include any controls for returning users.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table B-8: Transaction cost by sector ($1k)

	(1) Probit	(2) $\tau = 0.1$	(3) $\tau = 0.25$	(4) $\tau = 0.5$	(5) $\tau = 0.75$	(6) $\tau = 0.9$
Combined (first time)	467.01***	-3.41*	-622.82	57.06***	18.99	0.80
	(106.61)	(1.97)	(11,446.97)	(9.48)	(27.39)	(0.78)
Combined (returning)	-368.10**	-0.85	-264.15	22.04***	0.58	0.28
	(175.46)	(1.91)	(4,881.19)	(8.07)	(7.09)	(1.06)
Obs.	556	556	556	556	556	556

Probit uses robust standard errors; BQR standard errors recovered through bootstrap with 500 repitions. Testing for significantly different transaction costs at the targeted quantiles yields a p-value of 0.419.

*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$

Table B-9: Transaction cost by sector ($1k)

	Probit	$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$	$\tau = 0.75$	$\tau = 0.9$
Energy (first time)	207.37*	-0.53	-950.31	43.70***	-0.86	-0.59
	(113.46)	(0.98)	(28,516.99)	(8.71)	(8.12)	(1.67)
Oil/gas (first time)	393.32**	-1.06	-1,079.76	47.16***	4.24	0.53
	(172.25)	(0.83)	(32,335.63)	(9.89)	(10.20)	(2.06)
Other (first time)	760.07***	-2.26*	-1,279.79	59.04***	33.26***	7.71
	(203.41)	(1.19)	(38,321.31)	(8.99)	(12.45)	(7.71)
Energy (returning)	-636.51**	1.27**	-333.57	12.83	-20.21	-0.88
	(274.46)	(0.58)	(10,063.34)	(11.81)	(18.93)	(2.45)
Oil/gas (returning)	-450.56*	0.73	-463.02	16.29	-15.11	0.24
	(266.90)	(0.50)	(13,882.66)	(12.45)	(16.00)	(2.61)
Other (returning)	-83.81	-0.46	-663.05	28.17***	13.90**	7.42
	(167.61)	(0.86)	(19,865.78)	(7.19)	(6.16)	(8.21)
Obs.	556	556	556	556	556	556

Probit uses robust standard errors; BQR standard errors recovered through bootstrap with 500 repitions. Testing for significantly different transaction costs at the targeted quantiles for the energy, oil/gas, and other sectors yields a p-value of 0.001, 0.005, and 0.226, respectively. The transaction cost estimates and coefficients are presented at each ventile in figures A-4 and A-3, respectively. Figure A-6 presents the estimated offset adoption rates compared with the observed rates as a goodness-of-fit exercise (Kordas, 2006).

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

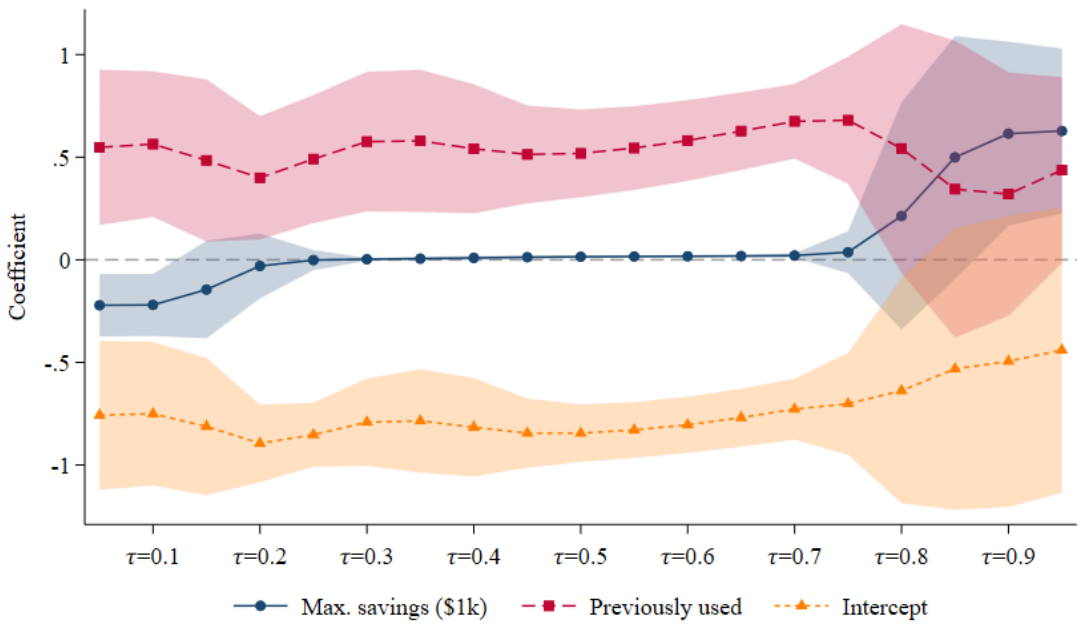


Figure B-7: BQR coefficients—combined model, using average compliance period price

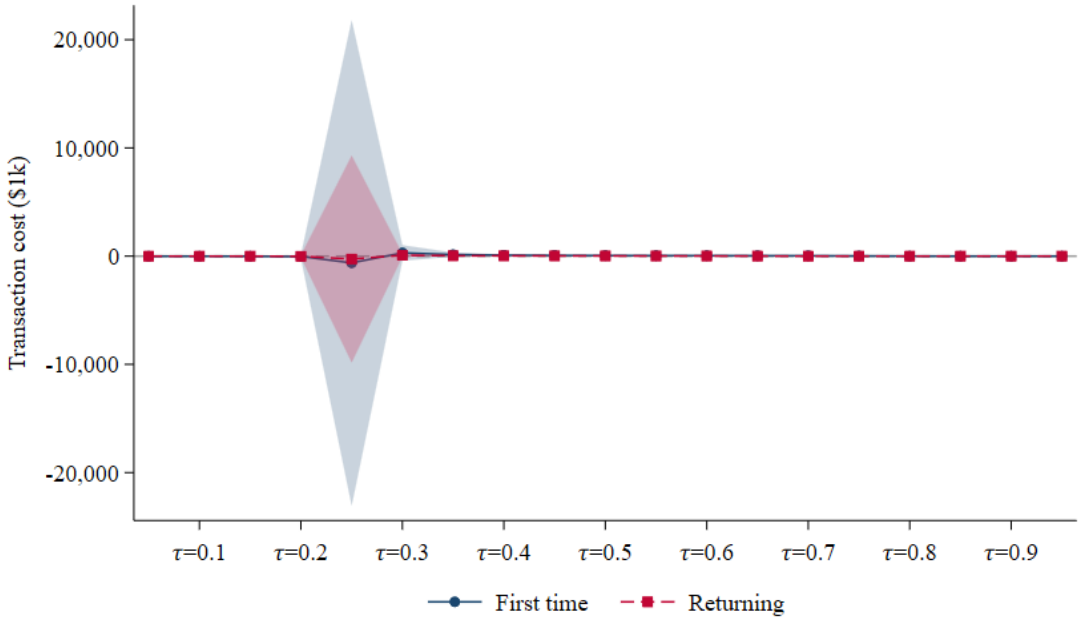


Figure B-8: Combined transaction costs at different quantiles, using average compliance period price

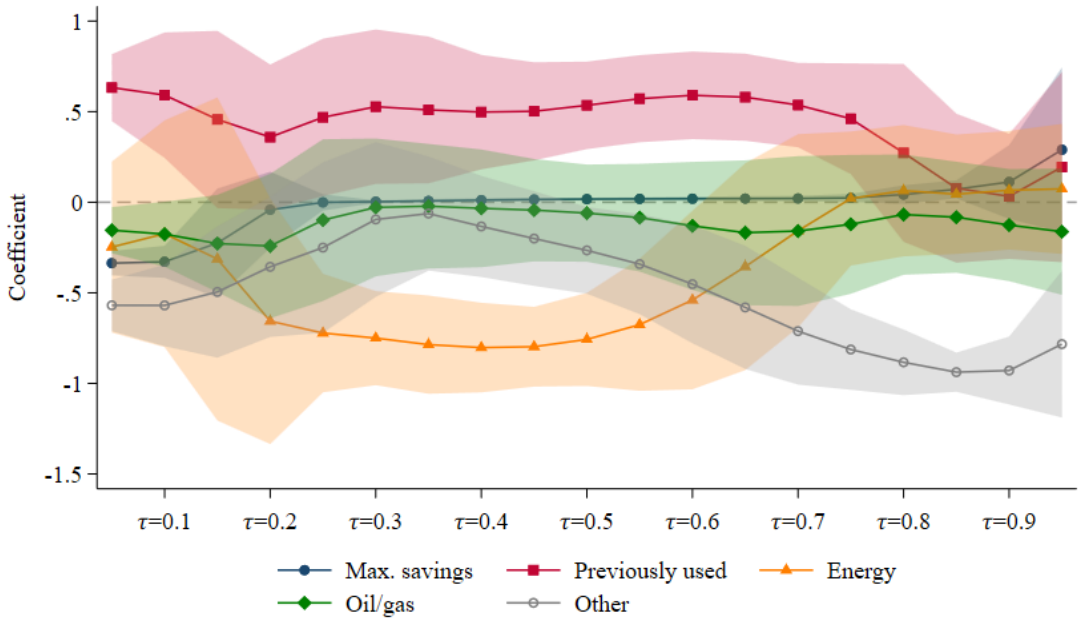
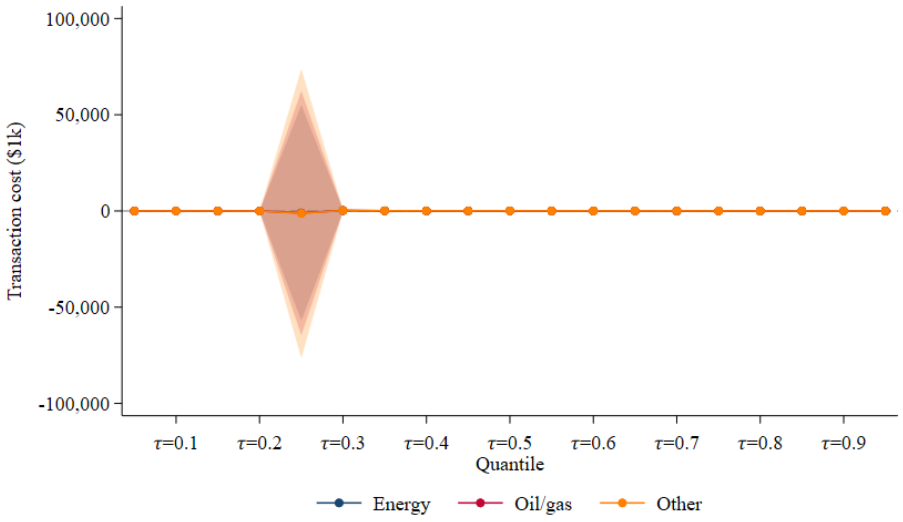
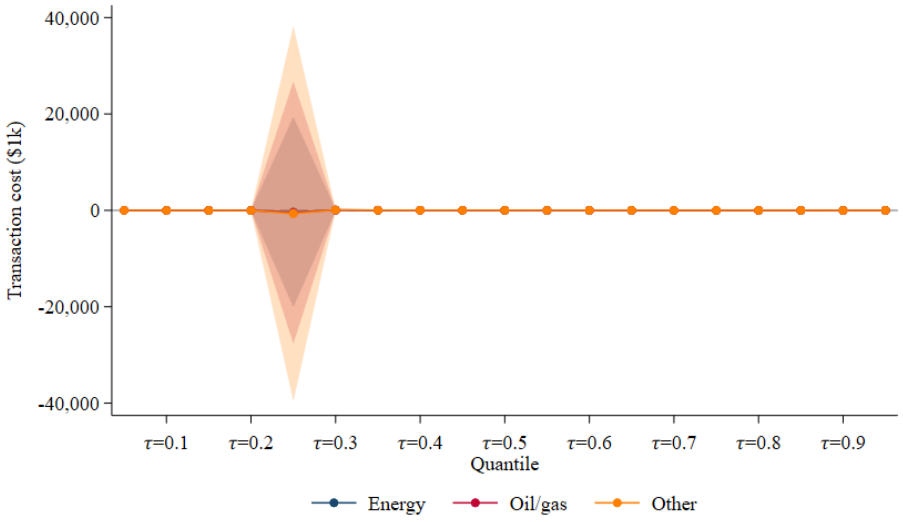


Figure B-9: BQR coefficients—sector-specific model, using average compliance period price



(a) First time offset users



(b) Returning offset users

Figure B-10: Sector-specific transaction costs at different quantiles, using average compliance period price

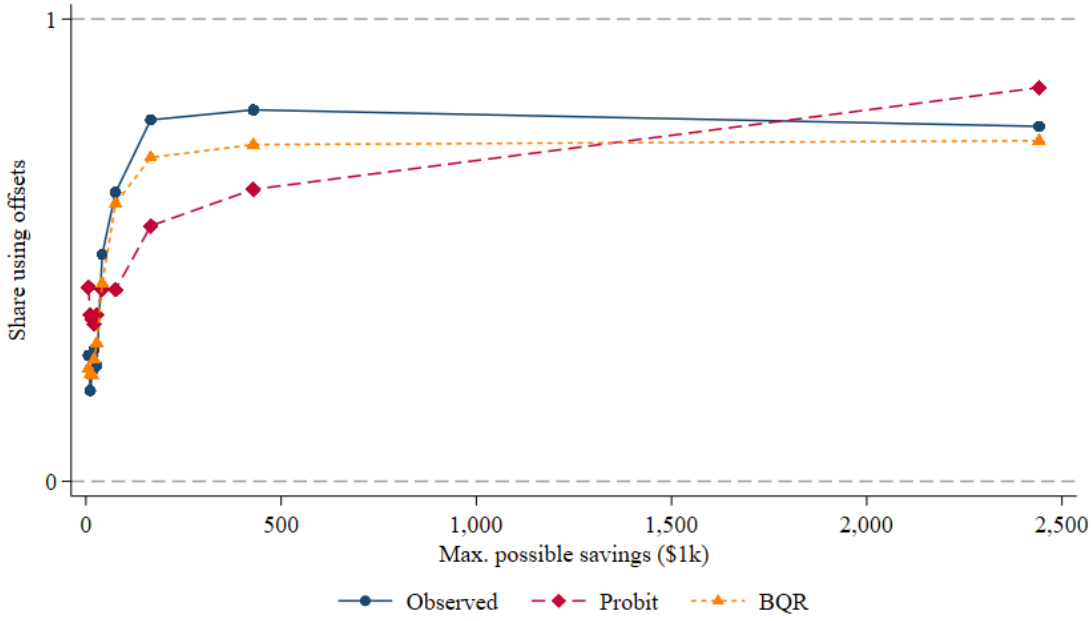


Figure B-11

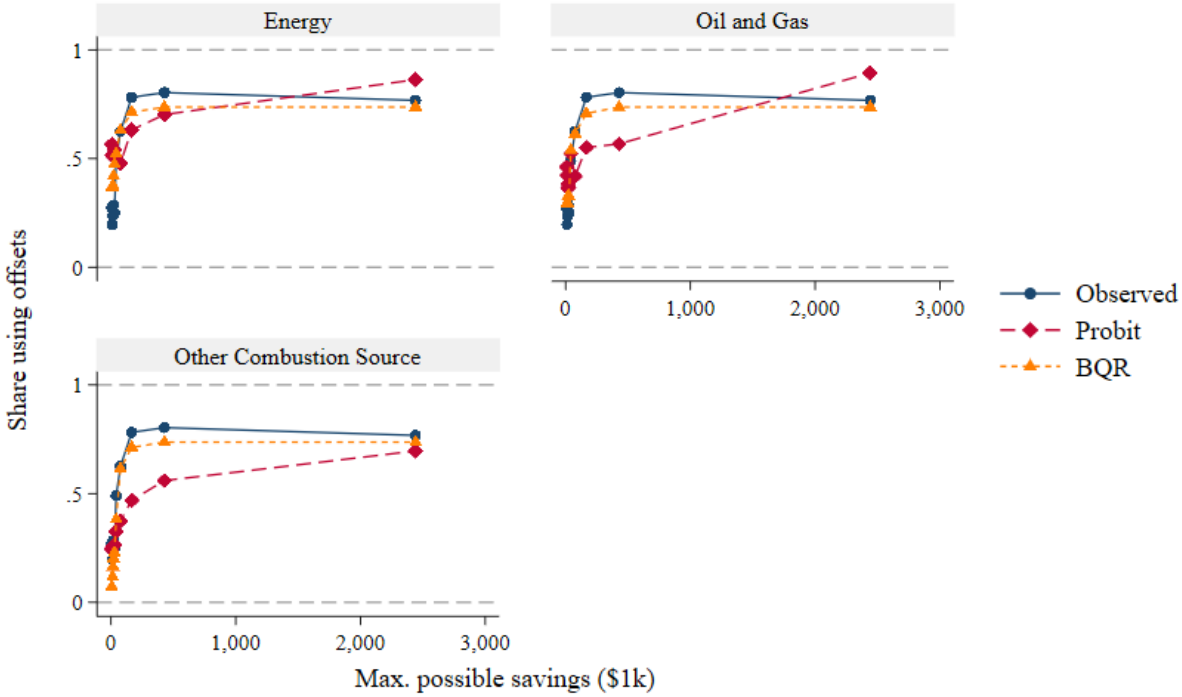


Figure B-12

C Welfare comparisons

Tables C-1 through C-3 present welfare calculations using alternative transaction cost estimates as robustness checks on the median-based results in the main text. The results illustrate why median BQR estimates are the preferred basis for welfare calculations.

Table C-1: Welfare calculations using mean transaction costs

	Energy	Oil and gas	Other
Incurred transaction cost (−)	-$24,336,756	$2,985,230	$37,400,457
Avoided transaction cost (+)	-$7,610,552	$16,582,948	$135,557,887
Gross earned savings (+)	$58,574,056	$99,465,039	$14,535,722
Gross missed savings (−)	$9,932,814	$7,894,877.75	$10,452,469
Net achieved savings (+)	$82,910,812	$96,479,809	-$22,864,735
Net money left on the table (−)	$17,543,366	-$8,688,070	-$125,105,418
Value of offsets market	**$65,367,446**	**$105,167,879**	**$102,240,683**
Frictionless cost savings (hypothetical)	$68,506,870	$107,359,917	$24,988,191

Next to each welfare component is an indicator in parantheses. A (+) indicates that the component contributes positively to overall cost savings, and a (−) indicates that the component contributes negatively to overall cost savings. Incurred transaction cost is the total transaction cost paid by all firms that use offsets within a sector; avoided transaction cost is the total value of the transaction cost that *would have* been paid by firms, but was avoided because they did not use offsets; gross earned savings is the total value of cost savings gained by users of offsets; and gross missed savings is the gross cost savings that firms could have earned, but did not because they did not use offsets. The net achieved savings is the difference between gross earned savings and incurred transaction costs, and net money left on the table is the difference between gross missed savings and avoided transaction cost. The total value of a market is the difference between net achieved savings and net money left on the table. The frictionless cost savings assumes zero transaction cost and full participation, so it is the sum of gross earned savings and gross missed savings.

Table C-1 uses the probit mean transaction cost estimates from table 4. Several components of this table are difficult to interpret economically. In the energy sector, for example, the incurred transaction cost is negative (-$24.3M), implying that offset users were on net made better off by paying transaction costs—a result that follows directly from the negative returning user transaction cost in the probit mean along with the many returning offset users in the energy sector, but there is no convincing structural interpretation. Similarly, the other sector market value ($102.2M) substantially exceeds the frictionless benchmark ($24.9M), implying that transaction costs more than quadruple the value of the offset market relative

to a world with no frictions—again a result driven by the large negative avoided transaction costs for other sector non-participants. When the frictionless market has full participation, this is implausible.

Tables C-2 and C-3 repeat this exercise for the Mundlak CRE models and the models controlling for initial conditions in tables A-2 and A-4, respectively. Table C-4 compares the total implied cost savings and efficiency losses across all four specifications. The median estimates yield the lowest implied cost savings, and, therefore, the most implied inefficiency. As such, these represent an upper bound on the efficiency losses attributed to transaction costs.

Table C-2: Welfare calculations using mean transaction costs in table A-2

	Energy	Oil and gas	Other
Incurred transaction cost (−)	-$16,774,058	$7,423,730	$42,599,649
Avoided transaction cost (+)	-$8,190,076	$10,370,214	$107,082,154
Gross earned savings (+)	$58,574,056	$99,465,039	$14,535,722
Gross missed savings (−)	$9,932,814	$7,894,877.75	$10,452,469
Net achieved savings (+)	$75,348,114	$92,041,309	-$28,063,927
Net money left on the table (−)	$18,122,890	-$2,475,336	-$96,629,685
Value of offsets market	**$57,225,224**	**$92,516,645**	**-$28,063,927**
Frictionless cost savings (hypothetical)	$68,506,870	$107,359,917	$24,988,191

Next to each welfare component is an indicator in parantheses. A (+) indicates that the component contributes positively to overall cost savings, and a (−) indicates that the component contributes negatively to overall cost savings. Incurred transaction cost is the total transaction cost paid by all firms that use offsets within a sector; avoided transaction cost is the total value of the transaction cost that *would have* been paid by firms, but was avoided because they did not use offsets; gross earned savings is the total value of cost savings gained by users of offsets; and gross missed savings is the gross cost savings that firms could have earned, but did not because they did not use offsets. The net achieved savings is the difference between gross earned savings and incurred transaction costs, and net money left on the table is the difference between gross missed savings and avoided transaction cost. The total value of a market is the difference between net achieved savings and net money left on the table. The frictionless cost savings assumes zero transaction cost and full participation, so it is the sum of gross earned savings and gross missed savings.

Table C-3: Welfare calculations using mean transaction costs in table A-4

	Energy	Oil and gas	Other
Incurred transaction cost (−)	-$199,458,570	$25,639,600	$108,996,000
Avoided transaction cost (+)	-$113,324,860	$40,821,480	$208,211,000
Gross earned savings (+)	$58,574,056	$99,465,039	$14,535,722
Gross missed savings (−)	$9,932,814	$7,894,877.75	$10,452,469
Net achieved savings (+)	-$140,884,514	$73,825,439	-$94,460,278
Net money left on the table (−)	$103,392,046	-$32,926,602	-$197,758,531
Value of offsets market	**-$37,492,468**	**$106,752,041**	**$103,298,253**
Frictionless cost savings (hypothetical)	$68,506,870	$107,359,917	$24,988,191

Next to each welfare component is an indicator in parantheses. A (+) indicates that the component contributes positively to overall cost savings, and a (−) indicates that the component contributes negatively to overall cost savings. Incurred transaction cost is the total transaction cost paid by all firms that use offsets within a sector; avoided transaction cost is the total value of the transaction cost that *would have* been paid by firms, but was avoided because they did not use offsets; gross earned savings is the total value of cost savings gained by users of offsets; and gross missed savings is the gross cost savings that firms could have earned, but did not because they did not use offsets. The net achieved savings is the difference between gross earned savings and incurred transaction costs, and net money left on the table is the difference between gross missed savings and avoided transaction cost. The total value of a market is the difference between net achieved savings and net money left on the table. The frictionless cost savings assumes zero transaction cost and full participation, so it is the sum of gross earned savings and gross missed savings.

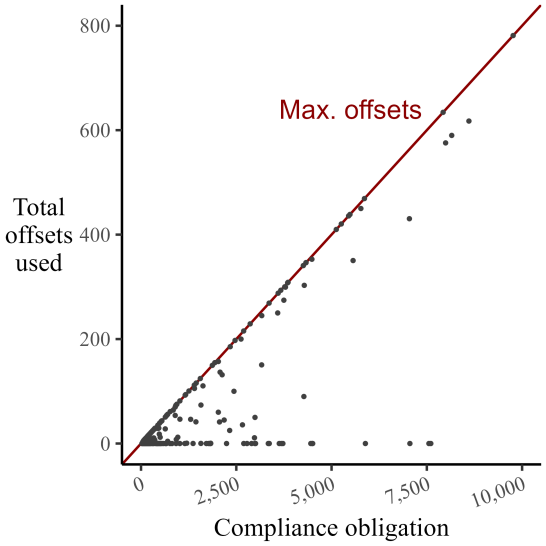
Table C-4: Welfare comparison

	Median	Baseline	Mundlak CRE	Initial conditions
Total implied cost savings	$149,683,146	$272,776,008	$220,307,627	$172,557,826
Implied inefficiency	$51,171,832	-$71,921,031	-$19,452,650	$28,297,152

Next to each welfare component is an indicator in parantheses. A (+) indicates that the component contributes positively to overall cost savings, and a (−) indicates that the component contributes negatively to overall cost savings. Incurred transaction cost is the total transaction cost paid by all firms that use offsets within a sector; avoided transaction cost is the total value of the transaction cost that *would have* been paid by firms, but was avoided because they did not use offsets; gross earned savings is the total value of cost savings gained by users of offsets; and gross missed savings is the gross cost savings that firms could have earned, but did not because they did not use offsets. The net achieved savings is the difference between gross earned savings and incurred transaction costs, and net money left on the table is the difference between gross missed savings and avoided transaction cost. The total value of a market is the difference between net achieved savings and net money left on the table. The frictionless cost savings assumes zero transaction cost and full participation, so it is the sum of gross earned savings and gross missed savings.

D Including the fourth compliance period

(a) Offset usage during the first three compliance periods



(b) Offset usage in the fourth compliance period

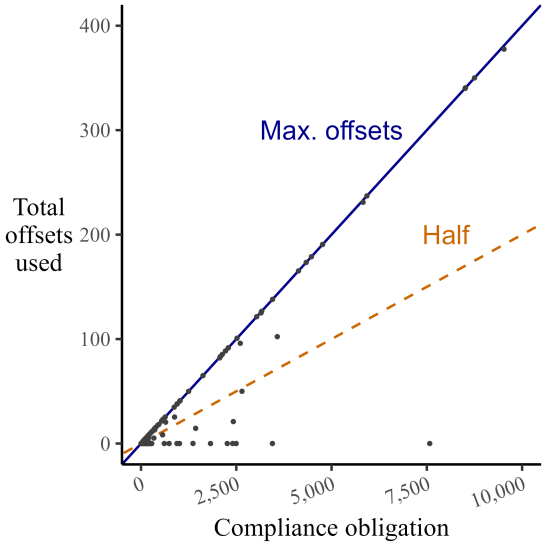


Figure D-1: Relationship between offsets used and potential offsets allowed

These figures show the relationship between total offsets used and compliance obligation. The figures are limited to firms with 10 million tons of emissions for clarity. In both graphs, the solid ray is the maximum offsets for a given emissions level. The orange dashed line in D-1b is half of this. Points on the solid red line in the left panel indicate surrendered offsets equal to 8% of a firm's emissions—the maximum amount. Points on the solid blue line in D-1b are doing the same while those on the dashed orange line are surrendering half (the non DEBS-providing limit). Taken together, these figures show that firms are likely to be either "out" of the offsets market or "all in" the market, with few firms acting in the middle.

This same data is shown graphically in figures D-1 and D-2. Figure D-1 shows that firms will frequently use the maximum amount of offsets possible, or will use no offsets at all. The solid red line in the left panel represents the maximum amount of offsets possible at any emissions level, for the first three compliance periods. The solid blue line in figure D-1b similarly represents the maximum amount of offsets possible at any emissions level during the fourth compliance period while the dashed orange line represents half of this level. The first three compliance periods are separated from the fourth compliance period because the rules around offset use change in the fourth compliance period. Table 2 shows that there are firms both along this DEBS limit and along the total offsets limit.

Figure D-2 allows for a closer examination of offset adoption patterns. It shows CDFs

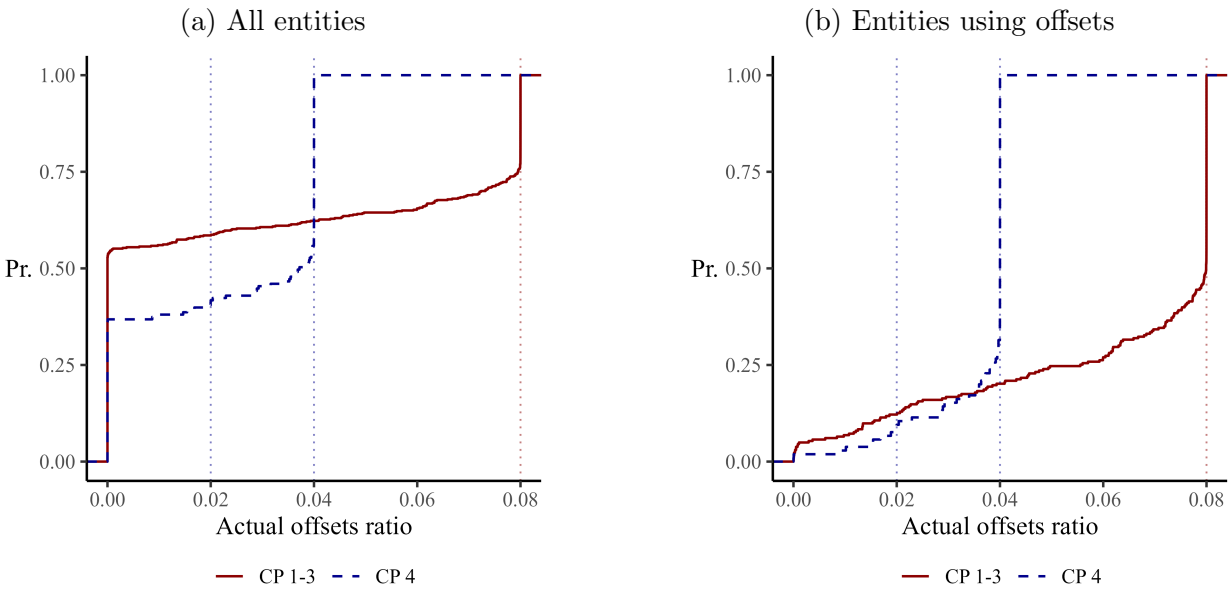
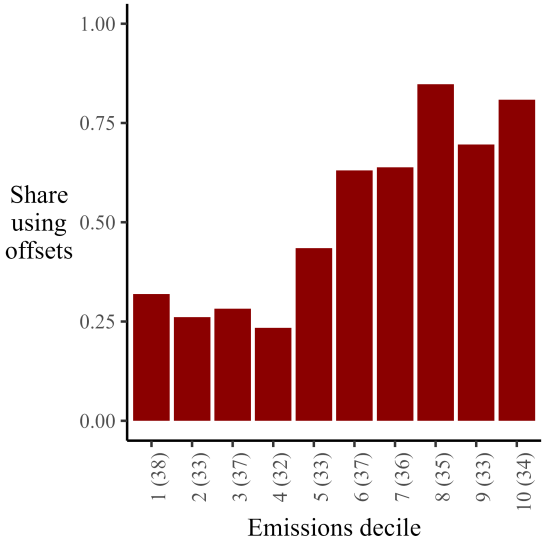


Figure D-2: Cumulative density offset share

Note: These are the full compliance-period data. The tall vertical line of the CDFs at 0 in figure D-2a indicates that a large portion of firms use zero offsets. The long horizontal stretches reveal a clustering in the share of offsets used at limit values. In particular, the first three compliance periods have a cluster at share 0.08, and the fourth compliance period has a cluster at 0.04. If there were a discrete jump in the blue dashed CDF at 0.02 (the non DEBS-providing limit), we might interpret this as a second transaction cost related to locating DEBS-providing offsets. Figure D-2b shows that there are some firms in between zero and the maximum possible offsets, but the discrete jumps in the CDFs are where most movement occurs.

(a) Share of firms using offsets in the first three compliance periods



(b) Share of firms using offsets in the fourth compliance period

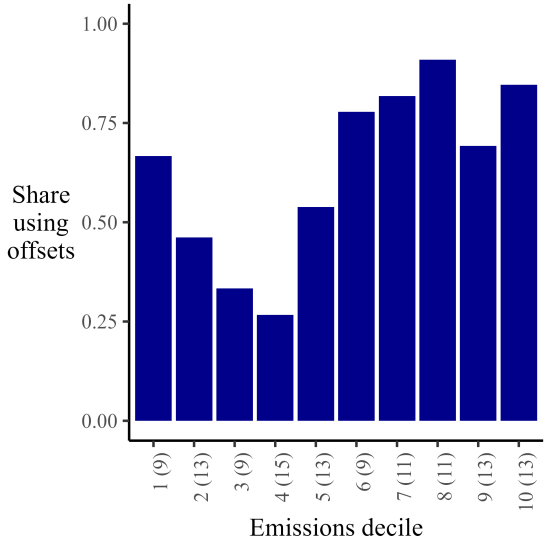
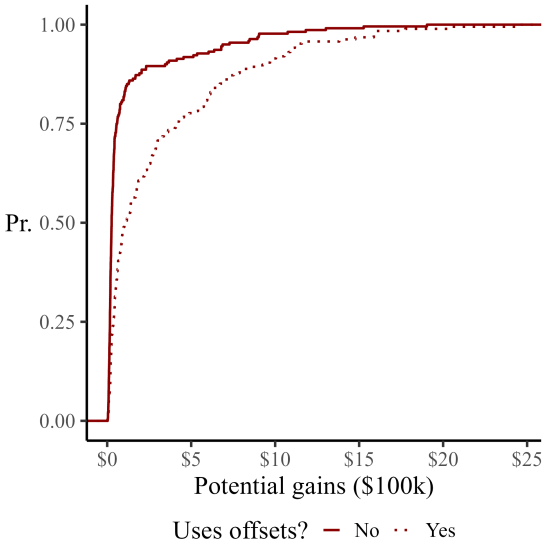


Figure D-3: Larger firms are more likely to use offsets

Note: Firms are grouped into quintiles based on their total emissions. The height of the bar indicates the share of firms in that quintile using offsets. The general trend that bars get higher as the quintile increases indicates that larger firms are more likely to use offsets.

(a) Full compliance period (periods 1-3)



(b) Full compliance period (period 4)

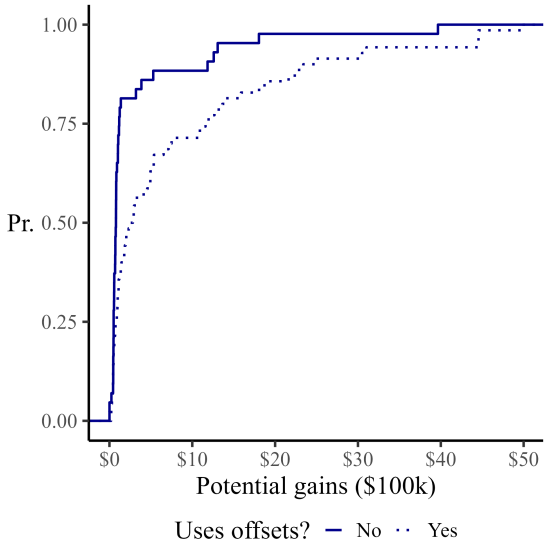


Figure D-4: Potential gains from offset adoption

Note: The probability that gains are less than $\$x$ is smaller for firms that use offsets than it is for firms that do not. This true for all x, so the expected gains from offset adoption in the sample of firms using offsets is higher than in the sample that does not. Firms may be selecting into offset adoption based on the potential gains.

across the distribution of the share of allowable offsets used. Panel a shows the unconditional CDF and panel b shows the CDF conditional on some use of offsets. The CDF in figure D-2a makes clear that there is a near-discrete jump in the share of offsets used from 0 to the upper bound. The solid red line is the CDF of the ratio of offsets to emissions for the first three compliance periods. It is clear that over half of firms in these compliance periods use no offsets. The nearly flat CDF between horizontal axis values 0 and 0.08 reflects the relative dearth of firms using a positive amount of offsets less than the limit. The dashed blue line is defined similarly for the fourth compliance period. Over half of firms now use offsets (likely due to the higher cost savings), but we see the same pattern of a flat CDF with a discrete jump at the upper limit, 0.04. For clarity, figure D-2 restricts the sample to all firms that use any offsets. Here, the horizontal lines are revealed to be slightly upward sloping, but they still show the pattern of discrete jumps at key points along the horizontal axis. Figures D-1 and D-2 show that a majority of firms lie on the extreme ends of the offset quantity spectrum.

Starting in the fourth compliance period, figure D-3b shows a different overall distribution of offset adoption, but the pattern of higher-emitting firms being more likely to use offsets remains. Recall that the premium paid for permits increased greatly at the start of the fourth compliance period, more than doubling the incentive to adopt offsets. This is only descriptive, but the less pronounced skew suggests that the policy's favoring of large firms may be slightly less dramatic under the new offsets policy; perhaps in line with this, figure D-4 shows a wider gap between the CDF of potential gains in the fourth compliance period, relative to the first three. Two important factors in a firm's offset decision changed at the start of the fourth compliance period. First, the offset cap was halved from 8 to 4 percent. This limits the maximum benefits that a firm can gain by using offsets. On its own, this decrease in the offsets cap disincentivizes their use. The second factor was the increase of the price discount of offsets. As figure 1 illustrates, the premium on permits increased from \$3.15 in the third compliance period to \$13.1 in the fourth. While the quantity restriction

was cut in half, the substitution benefit per unit increased by over four times. So the fact that offset market participation increased in the fourth compliance period is consistent with the idea that firms must bear a transaction cost in order to benefit from using offsets, but the benefits of doing so are higher.

E Hurdle model

It can be useful to try to understand the mechanism driving "interior" offset quantities, but the raw data in figures 2 and 3 show that many firms use zero offsets or the maximum. Thus, any model naively regressing offset quantity on explanatory variables in one equation will face consistency issues due to the large number of zeros.

As such, I use a hurdle model to specify the quantity of offsets, conditional on using any offsets at all (Cragg, 1971). This is a two-step process where firms first decide whether to enter the offsets market, and, conditional on entering, choose how many offsets to purchase. The hurdle model works as follows. A firm first realizes its maximum possible cost savings for using offsets. If the cost savings do not cover the "hurdle" of the transaction cost then the firm's offsets quantity is zero. If the firm does clear the hurdle, then it will have a positive number of offsets, the quantity of which is modelled via a truncated distribution (Cameron & Trivedi, 2005).

The first step is modelled exactly as in equation (2). The second step, conditional on using offsets, models the quantity of offsets used as a function of the maximum possible offsets savings and the sector identifier. The complete hurdle model is defined in equation (E-1):

$$\begin{aligned} \Pr\left[\mathbb{1}_{ijt}^{O} \mid Gains_{ijt}, \rho_j\right] &= \Phi\left(\frac{1}{\sigma} Gains_{ijt} - \sum_{j=1}^{3} \frac{\rho_j}{\sigma} \mathbb{1}_i^j\right) && \text{(first step)} \\ O_{ijt} &= \exp\left(\beta \ln\left(\tilde{q}_{ijt}\right) + \delta_0 Gains_{ijt} + \sum_{j=1}^{3} \delta_j \mathbb{1}_i^j + \eta_{ijt}\right) && \text{(second step)} \end{aligned} \tag{E-1}$$

where O_{ijt} is the quantitiy of offsets purchased, and η_{ijt} is an error term that, conditional on the variables, is assumed to be independent of the error in the first stage.

The main reason for this exercise is to have the opportunity to rule out per-unit transaction costs in two ways. First, given that the second stage of the model returns transaction costs that are conditional on participation, these are essentially transaction costs on the *intensive margin*. That is, once a firm has incurred the extensive margin fixed cost (discussed in the body of the paper), this additional transaction cost is paid for each unit.

Second, exploiting the functional form of the second stage, I can try to measure the relationship between the offset limit and the number of offsets purchased. The second stage is an exponential function (because the outcome is always positive), so including the natural log of the maximum offset limit as a righthand side variable allows me to recover the impact of this limit on the number of offsets purchased. For example, if all coefficients are zero, but β is exactly 1, the equation reduces to $O_{ijt} = \tilde{q}_{ijt}$—the quantity of offsets purchased is exactly equal to the maximum amount. Such an outcome, combined with the evidence that the first stage has a real transaction cost would suggest that the transaction cost is purely realized on the extensive margin. That is, the transaction cost is purely a fixed cost.

The coefficient β measures the conditional impact of compliance cost savings on the firm's quantity of offsets and can be used to calculate conditional cost savings elasticity of offsets. That is, conditional on the firm's entry into the market and on its offset cap, how meaningful is an additional dollar of cost savings? The second step has an error term η_{ijt} that follows a truncated normal distribution.

Table E-1 presents the conditional transaction costs for using offsets. These transaction costs are calculated in the same way as in the first stage, but they use coefficients that are conditional on participating in the offsets market. The estimates are large in magnitude, but statistically insignificant. As such, it appears that all meaningful transaction costs are borne along the extensive margin. Table E-2 presents the coefficients used in the conditional transaction cost calculations; in particular, the top section. The constants and the sector

coefficients are statistically significant, and they are statistically different from each other, suggesting that there is a sector-level average conditional offset quantity.

Finally, table E-3 presents the coefficients for the same specification, but excludes maximum cost savings as an explanatory variable. The coefficient on the natural log of offset quantity is robust to this, suggesting that the near-one-to-one relationship between the offset limit and the offset quantity for participating firms is not driven by any mechanical relationship between the two variables.

Table E-1: Transaction costs conditional on participation

	(1)	(2)	(3)	(4)
Combined (first time)	218,951		189,926	
	(996,982)		(742,511)	
Combined (returning)	218,951		217,228	
	(996,982)		(851,197)	
Energy (first time)		64,663		64,907
		(73,597)		(73,687)
Oil and gas (first time)		52,354		53,136
		(60,465)		(61,230)
Other (first time)		54,214		54,325
		(61,271)		(61,235)
Energy (returning)		64,663		72,622
		(73,597)		(82,995)
Oil and gas (returning)		52,354		60,851
		(60,465)		(70,530)
Other (returning)		54,214		62,040
		(61,271)		(70,544)
Obs.	556	556	556	556
p returning			0.803	0.440
p sector heterogeneity		0.688		0.687

Transaction costs estimated from second stage hurdle model. These are the ratio of the intercept (in columns 1 and 3) or the sector dummy variables (in columns 2 and 4) to the coefficient on maximum cost savings, *conditional on participating* in the offsets market. The estimates are large in magnitude, but statistically insignificant, implying that all significant role of transaction cost is absorbed in the first stage. This insignificance suggests the transactionn cost affects only the extensive margin.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$

Table E-2: Hurdle model coefficients

	(1)	(2)	(3)	(4)
Second stage				
ln Emissions (1k MTe)	0.921***	0.951***	0.946***	0.975***
	(0.063)	(0.063)	(0.063)	(0.063)
Max. benefit ($1k)	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Energy		-3.029***		-3.024***
		(0.410)		(0.407)
Oil and gas		-2.452***		-2.476***
		(0.430)		(0.427)
Other		-2.539***		-2.531***
		(0.389)		(0.386)
Previously used			-0.373**	-0.359**
			(0.177)	(0.174)
Constant	-2.599***		-2.597***	
	(0.389)		(0.386)	
First stage				
Max. benefit ($1k)	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Energy		0.061		-0.183*
		(0.097)		(0.107)
Oil and gas		-0.272*		-0.354**
		(0.150)		(0.154)
Other		-0.571***		-0.693***
		(0.085)		(0.090)
Previously used			0.855***	0.761***
			(0.135)	(0.138)
Constant	-0.305***		-0.479***	
	(0.063)		(0.070)	
First stage $\ln \hat{\sigma}$	0.331***	0.314***	0.323***	0.306***
	(0.044)	(0.044)	(0.044)	(0.044)
Obs.	556	556	556	556
Sector F-statistic		4.62		4.52

The first stage mirrors the probit models, and the second stage coefficients are interpreted as log-linear effects offset quantity *conditional on market participation*. If firms used exactly the maximum allowed amount of offsets and had no other influences, the coefficient on log emissions would be 1, and all other coefficients would be zero. In each second stage, the maximum cost savings is insignificant while the intercepts are significant but small in magnitude. The F-statistics are from a test with a null hypothesis that all sector intercepts are equal; this hypothesis cannot be rejected, suggesting that each sector does play a role in determining conditional baseline offset quantity.

Table E-3: Hurdle model coefficients (excluding max. cost savings)

	(1)	(2)	(3)	(4)
Second stage				
ln Emissions (1k MTe)	0.913***	0.921***	0.937***	0.945***
	(0.050)	(0.050)	(0.051)	(0.051)
Energy		-2.849***		-2.846***
		(0.349)		(0.347)
Oil and gas		-2.329***		-2.353***
		(0.404)		(0.401)
Other		-2.369***		-2.361***
		(0.331)		(0.328)
Previously used			-0.373**	-0.360**
			(0.177)	(0.175)
Constant	-2.554***		-2.546***	
	(0.329)		(0.326)	
First stage				
Max. benefit ($1k)	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Energy		0.061		-0.183*
		(0.097)		(0.107)
Oil and gas		-0.272*		-0.354**
		(0.150)		(0.154)
Other		-0.571***		-0.693***
		(0.085)		(0.090)
Previously used			0.855***	0.761***
			(0.135)	(0.138)
Constant	-0.305***		-0.479***	
	(0.063)		(0.070)	
First stage $\ln \hat{\sigma}$	0.331***	0.315***	0.323***	0.307***
	(0.044)	(0.044)	(0.044)	(0.044)
Obs.	556	556	556	556
Sector F-statistic		4.29		4.19

The first stage mirrors the probit models, and the second stage coefficients are interpreted as log-linear effects offset quantity *conditional on market participation*. The purpose of this table is to show that excluding the maximum cost savings from the second stage does not change the conclusion that the coefficient on maximum offset use is significant and close to 1. The F-statistics are from a test with a null hypothesis that all sector intercepts are equal; this hypothesis cannot be rejected, suggesting that each sector does play a role in determining conditional baseline offset quantity.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$