NC STATE UNIVERSITY

Center for Environmental and Resource Economic Policy
College of Agriculture and Life Sciences
https://cenrep.ncsu.edu

Transaction Costs of Emissions Offsets:

Evidence from California

Willard Robinson

Center for Environmental and Resource Economic Policy
Working Paper Series: No. 26-003
June 2026

Suggested citation: Robinson, Willard, "Transaction Costs of Emissions Offsets: Evidence from California," CEnREP Working Paper No. 26-003, June 2026, https://go.ncsu.edu/cenrep.wp.26.003.

Transaction Costs of Emissions Offsets: Evidence from California*

Willard Robinson†

June 26, 2026

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. Many firms never purchase an offset, leaving significant cost savings on the table. Using compliance data from 2013-2020, I identify fixed transaction costs (e.g., risk premia or public relations concerns) as the primary barrier, and it is a substantial one: averaging between \$220,000 and \$835,000 (respectively, 10% and 38% of average spending on offsets). But this value is heavily skewed; median transaction cost estimates are between \$37,000 and \$56,000 (respectively, 13% and 19% of median spending on offsets), suggesting the average cost is driven by a few high-cost firms. Using the median cost estimates, I estimate that, in all, firms incurred between \$19 and \$95 million in transaction costs over the sample period; during the same period, the offsets market generated between \$250 and \$326 million.

Key words: Carbon offsets, transaction costs, emissions trading, compliance market

JEL Codes: Q28, Q38, Q52, Q56

*I thank Spencer Banzhaf, Roger von Haefen, Erin Sills, Zhenxuan Wang, and Longji Xia for valuable feedback.

†Department of Economics, North Carolina State University; wmrobin3@ncsu.edu

1 Introduction

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" attributed to actions like avoided deforestation or voluntary methane capture (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: for example, offsets are regularly nearly 20% cheaper than the standard pollution permits in California. 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 (e.g., the time cost of learning to navigate the offsets market) as the primary barrier to offset adoption, quantifies their magnitude, and estimates both the value that the offsets market generates for Californian firms (between \$250 and \$326 million) and the implied costs paid by offset-using firms (between \$19 and \$95 million).

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 decreased activity along the extensive margin, but not the intensive margin is consistent with a fixed compliance period-level transaction cost, not a per-unit cost. I show in Section 4 that the raw data confirms this interpretation. Using compliance data from 2013–2020, I estimate that a firm in California faces a median implicit transaction cost between \$37,000 and \$56,000 each compliance period, depending on whether it has used offsets previously.[1] As a further check for per-unit transaction costs, I estimate transaction costs for

[1] I discuss the median estimates as the distribution of firms in California's compliance market has "fat

firms *conditional on participation* in the offsets market using a hurdle model (Cragg, 1971) and find that all are statistically indistinguishable from zero. Moreover, I show evidence that the distribution of these barriers to adoption can vary considerably depending on a firm's sector, with median costs for offset-using firms in the energy sector can be as low as $37,690 while the median cost for firms outside of this sector can be up to $56,140. Whether these are search costs, risk premiums, or public relations concerns is still ambiguous, however.

Naegele (2018) similarly examined transaction costs in the European Union's emission trading system (EU ETS). Using data on both freely allocated permits and total compliance permits and offsets, she compared semi-parametric binary quantile regression (BQR) models with standard probit models, finding that high average transaction costs were caused by a strongly skewed latent distribution of these costs.

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 greenhouse gas (GHG) emissions (*ICAP*, 2024). This particular setting has an emissions-proportional cap on the number of offsets a firm can purchase, which applies to all firms whereas offset caps in the EU ETS vary by country. Furthermore, the single regulatory entity in California keeps regulatory quality constant for all firms (Lessmann & Kramer, 2024). 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. Third, utilizing California's sector classification for each firm, I measure the degree of heterogeneity in transaction costs across sectors. As such, I find that firms in the energy-producing sector have the smallest transaction cost, followed by the oil and gas refinery sector, with all other firms combined having the highest transaction cost. Sector adoption of offsets follows this pattern with firms in these sectors participating at a rate of approximately 74%, 66%, and 40%, respectively.[2]

tails" with a small number of firms producing exceptionally large levels of emissions, and, therefore, facing exceptionally high compliance costs. As such, medians—rather than means—tend to represent the marginal average participant.

[2]The energy-producing sector is responsible for about 12.5% of total emissions while the oil and gas sector

My empirical analysis reveals key patterns consistent with fixed transaction costs. 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). Additionally, 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—nearly 80% of EU firms use offsets while between 40%-60% of Californian firms do, depending on the compliance period. Finally, 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 once it has done so for the first time.

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; Naegele, 2018).

The remainder of the paper proceeds as follows. Section 2 provides institutional background on California's cap-and-trade program and discusses potential sources for the participation frictions. Section 3 describes the data sources. Section 4 reveals several stylized facts using raw data to support the fixed transaction cost hypothesis. Section 5 outlines the empirical strategy and section 6 presents these estimates. Section 7 discusses policy implications and concludes.

is responsible for about 35.1% of total emissions. All other sectors make up the remaining 52.4% of total emissions.

2 Background

2.1 Cap-and-trade in California

The California Global Warming Solutions Act initiated California's cap-and-trade system in 2013, administered and monitored by the California Air Resources Board (CARB) (AB-32, 2006), with the goal of decreasing GHG emissions to 1990 levels. The program covers roughly three-quarters of the state's emissions. To limit emissions "leakage" (Caron et al., 2015; Fell & Maniloff, 2018), CARB allocates additional allowances to energy-intensive, trade-exposed industries (Fowlie et al., 2021; Johnson, 2023); these allocations are not publicly observed, though the surrender of compliance instruments, or permits, is tracked through the Compliance Instrument Tracking System Service (CITSS).

The program is divided into 2–3 year compliance periods. A firm becomes a "regulated entity" in any year it emits 25,000 or more metric tons of CO_2-equivalent, after which it must surrender one compliance instrument per ton emitted by the end of the 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. During the first three compliance periods (2013–2020), firms could meet up to 8 percent of their total compliance obligation with offsets; this cap fell to 4 percent in the fourth period (2021–2023), half of which had to provide direct environmental benefits to the state (DEBS) of California.[3]

2.2 Transaction costs in environmental markets

The offsets market is a voluntary payments for ecosystem services market whose demand side—the focus of this paper—is comparatively understudied. All offset projects generate a

[3]From 2026 through 2030, the offsets cap is raised to 6 percent of compliance obligations, half of which must provide DEBS.

finite amount of CARB offset “credits” which must be registered on one of the two approved offset project registries (American Carbon Registry and Climate Action Reserve), and can be substituted one-for-one with the CARB-issued permits. Californian firms acquire these credits either by entering (CARB-approved) long-term contracts with suppliers themselves, or by purchasing offsets on secondary markets (e.g., CNaught, Evolution Markets).

In a cap-and-trade market with no frictions, introducing offsets acts as a pure price reduction; a transaction cost dampens this reduction, and can possibly overturn it (Hahn & Stavins, 2011; Stavins, 1995). Such costs can arise from search and learning costs, bargaining costs, doubts over policy longevity, and enforcement costs (Cacho et al., 2013; Galik et al., 2012; Mayr et al., 2025; Pan et al., 2022). The evidence closest to this paper comes from the EU ETS: Jaraitė-Kažukauskė and Kažukauskas (2015) find that smaller firms face higher transaction costs than larger firms—this holds in California as well—while Naegele (2018) estimates that the bulk of the transaction cost is attributed to compliance instrument trading in general, with a smaller additional transaction cost tied to using offsets.

There are three hypothesized sources of transaction cost for Californian firms. The first is the time costs and human capital investment. Offsets are listed on centralized registries, but a firm must still ensure its compliance, which may require a significant amount of employee time. The second is the potential “invalidation” risk. If an offset is determined non-additional ex post or destroyed in a natural event, its use for compliance reasons is nullified. Offsets being invalidated are an extremely uncommon event due to the rigorous verification process California requires. During the first three compliance periods, less than 0.1% of the more than 159 million offsets were invalidated (Burtraw et al., 2022).[4] Though rare, a risk-averse firm might be willing to pay a premium for the stability of government-created permits. The third—and most speculative—is the potential for public image concerns. Recent empirical evidence demonstrates that non-additionality is pervasive across offset programs (Aspelund & Russo, forthcoming; Calel et al., 2025; West et al., 2023). These doubts about the integrity

[4]These come from just four projects. The investigation of each is documented and published on CARB’s website: https://ww2.arb.ca.gov/our-work/programs/compliance-offset-program/offset-credit-invalidation.

of offsets could potentially create public relations costs for firms using them, particularly in environmentally conscious California.

3 Data

3.1 Emissions, permits, and offsets

The main data for this paper come from CARB's compliance reports.[5] These are annual tables that list the emissions compliance obligation, the number of permits retired, and the number of offsets retired for every regulated firm. 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, regardless of their regulation status.[6] 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.[7] Each report contains aggregate data on the total 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 trans-

[5] Example: https://ww2.arb.ca.gov/resources/documents/compliance-summary-report

[6] This dataset has been used by Badgley, Freeman, et al. (2022) and Tuladhar et al. (2014). Every regulated firm is listed in the Greenhouse Gas Inventory, but not every firm in the Greenhouse Gas Inventory is necessarily regulated.

[7] Example: https://ww2.arb.ca.gov/cap-and-trade-program/summary-market-transfers-report

actions between corporate associates, where the instruments are being traded as part of a larger bundle.[8] Each transfer summary provides a weighted average price, median price, and standard deviation of price separately for both allowances and offsets. A separate table published by CARB details the prices and quantities of permits bought at auction.[9] 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.[10]

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 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.

[8] 96% of all transferred allowances were transferred in a priced transaction in 2023. This percentage is typical of other years as well.

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

[10] 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 A and the qualitative results are unchanged.

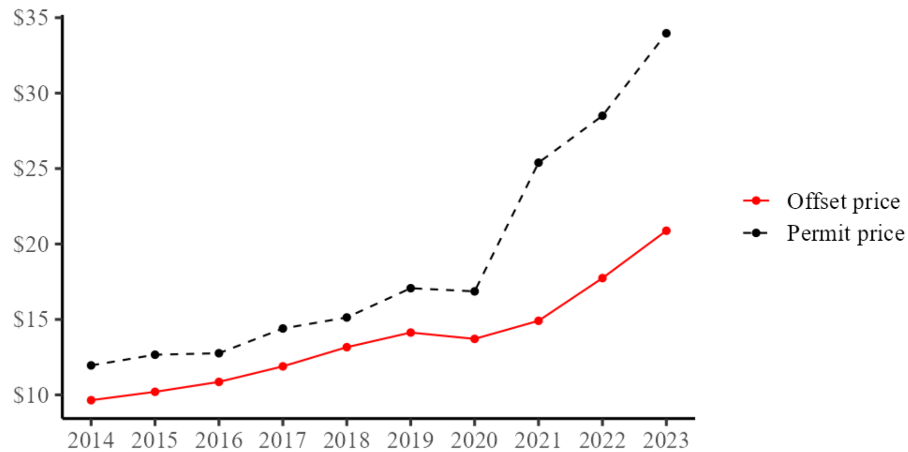

Figure 1: Price evolution of permits and offsets

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.

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 stark increase coincides with both a lowering of the offset cap from 8% to 4% and the introduction of a 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 1: 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 1 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. Less than half of all regulated firms used 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, conditional on using any. Consistently, over half of the firms that use offsets use the maximum amount. The fifth 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 13% 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

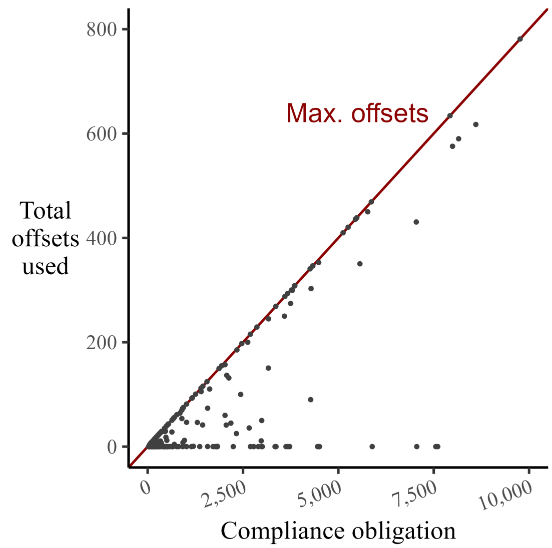

Figure 2: Relationship between emissions and offsets usage during the first three compliance periods

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 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 covers compliance periods 1–3 (2013–2020), which is also the analytic sample for the econometric models.

the first three compliance periods (8% of emissions). The clusters of points along this red line and along the horizontal axis—with relatively few in between—indicate firms that are either "all in" to or "completely out" of the offsets market. Figure 3 presents the cumulative distribution functions (CDFs) of the share of compliance met with offsets. Figure 3a shows the unconditional CDF across all regulated firms. The tall vertical line at zero indicates that over half of firm-period observations use no offsets at all in the first three compliance periods. The long horizontal stretch repeats 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 do use offsets use an amount between 0 and the limit (also shown in column 4 of table 1), with a roughly uniform adoption pattern of these "interior" firms. Conditional on using offsets, using the maximum amount appears to be the dominant pattern.

Because the offset price discount is common to all firms, but the number of allowable offsets is proportional to emissions, the incentive to use offsets is increasing in firm emissions.

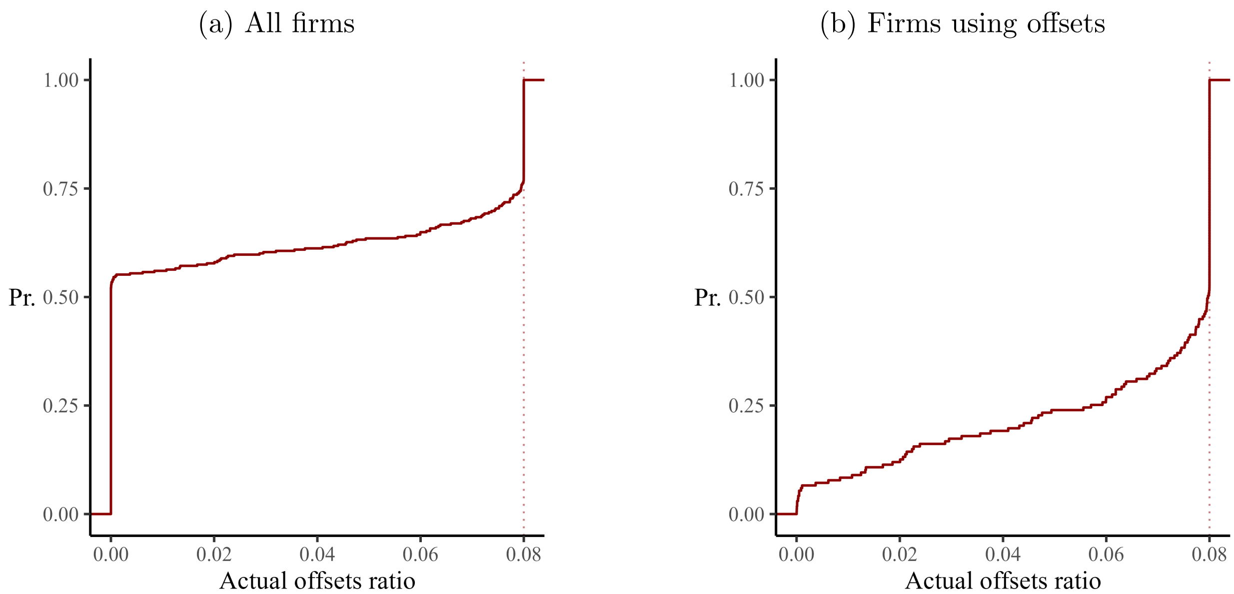

Figure 3: Cumulative density share

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%).

This incentive structure leads to firms selecting into the offset market as revealed in figure 4, 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 4 represents the strength of incentive to use offsets. The stochastic dominance is intuitive: at all values on the horizontal axis, the probability of having a weaker incentive to use offsets is always smaller for firms that use them. In other words, a randomly chosen firm that uses offsets is more likely to have higher incentive than another randomly chosen firm that does not.

Finally, table 2 shows the persistence of offsets use by firms. This table shows the trend that offsets grow in popularity as the program continues. It also reveals something which may be missed from the other tables and figures, and is not mentioned in any of the previous literature: firms that use offsets are highly likely to continue using them in the future. This

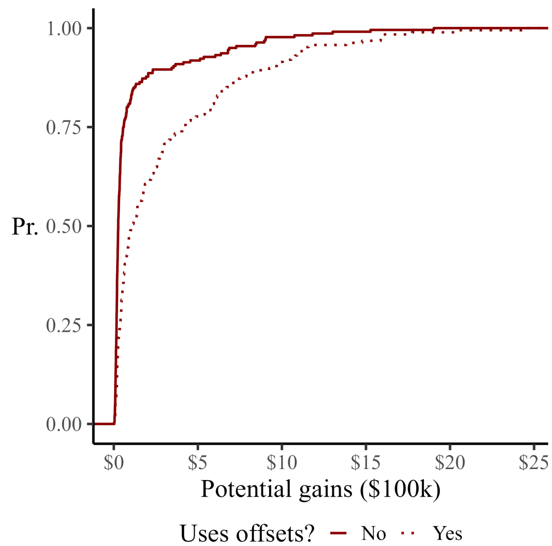

Figure 4: Potential gains from offset adoption

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 is true for all x, so the expected gains are higher among firms that use offsets than among those that do not. Firms may be selecting into offset adoption based on the potential gains.

persistence may reflect two features of a participation wedge. First, it could reflect that the firms have learned how to navigate the offsets market, thus reducing future barriers to offset adoption. Second, the persistence could suggest 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

This section outlines an empirical framework for characterizing firms' offset adoption decisions and recovering the level of gains associated with observed participation patterns. The goal is not necessarily to estimate structural transaction cost parameters, but rather to infer the magnitude of the participation wedge that rationalizes firms' behavior under a revealed-preference interpretation: the wedge acts as a reduced-form measure of transaction cost severity. I estimate separate adoption barriers for the three main sectors identified by CARB: "energy", "oil and gas", and "other." The energy sector consists of cogenera-

Table 2: Do entities continue to use offsets?

CP	using offsets this CP	Of those using offsets this CP: uses offsets at least once in future	uses offsets every CP in future
1	43.1%	77.4%	29.8%
2	39.3%	73.6%	52.8%
3	59.8%	77.6%	77.6%
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.

tion, 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.[11]

5.1 Decision framework

Following similar literature on the EU ETS (Jaraitė-Kažukauskė & Kažukauskas, 2015; Naegele, 2018), 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 decision model closely follows Naegele's (2018). Transaction costs can be different across sectors, particularly if the public relations concerns are relevant. 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 contemporaneous compliance period emissions with the maximum share of emissions that

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

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.

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

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

where $\mathbb{1}_i^j$ indicates that firm i is in sector j (there are 3), and ϵ_{ijt} is an idiosyncratic error.[12] The above equation says that the true payoff to using offsets is the monetary cost savings ($Gains_{ijt}$) less the average transaction cost in the sector (ρ_j) and an idiosyncratic factor (ϵ_{ijt}). The average transaction cost for a firm in sector j is ρ_j, while that firm's specific transaction cost is $\rho_j + \epsilon_{ijt}$.

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:

$$\mathbb{1}^O_{ijt} = \mathbb{1}\left\{y^*_{ijt} > 0\right\} = \mathbb{1}\left\{Gains_{ijt} > \sum_{j=1}^{3} \rho_j \, \mathbb{1}_i^j - \epsilon_{ijt}\right\}.$$

Thus, offset adoption occurs when the potential gains exceed the net transaction cost $\rho_j - \epsilon_{ijt}$, which is the sum of the sector-level fixed cost and an unobservable shock. As such, 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. The first 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 three compliance periods.

[12] 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 6.3.

The second source of variation comes 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. The key assumption is that, 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. I address this using within-firm variation in the potential gains from offset adoption and discuss more in section 6.3.

In appendix B, I test for additional transaction costs (fixed or per-unit) to using offsets, *conditional on* the firm choosing to use offsets in the first step following Jaraitė-Kažukauskė and Kažukauskas (2015). The results show strong evidence that, once a firm enters the offsets market, it uses very close to the maximum amount of offsets and there are no significant barriers in this second stage. As such, I focus on the first step in the main text, but view the results in appendix B as evidence that the implied barriers truly are fixed costs.

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). 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]

[13]Descriptive patterns in the raw data are presented for the fourth compliance period in appendix C.

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

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

where Φ is the CDF of the standard normal distribution. 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 that fit the data at the mean under the assumption of a zero conditional mean error (Cameron & Trivedi, 2005). Naegele's 2018 analysis of the EU ETS 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, a binary quantile response (BQR) offers more flexibility.[14]

[14] 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.

Following her approach, I employ BQR econometric methods based on Kordas (2006), which instead fit the conditional *quantile* (e.g. median) of the distribution of costs. In addition to capturing heterogeneity, 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 these barriers 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 the 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 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.

6 Results

I present the estimated participation thresholds in section 6.1 and use these in an exercise to calculate welfare in section 6.2. I emphasize that this econometric exercise is not intended to establish *causal identification* of the transaction cost: the purpose is to produce *reasonable* point estimates to use in back-of-the-envelope welfare calculations. The paper's qualitative conclusions rest mainly on the descriptive patterns in section 4, and the welfare estimates are presented as illustrative magnitudes. Throughout, I recover the transaction cost as the ratio of the intercept to the coefficient on maximum possible savings. Sector-specific transaction costs are recovered by the ratio of the sum of the intercept and the sector indicator with the coefficient on maximum possible savings. Recovering transaction costs for first-time and returning users is done similarly. All coefficients are reported in appendix D.

6.1 Implied transaction cost estimates

Average transaction costs 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.

Table 3 presents average transaction cost, split by whether a firm has used offsets before.[15] 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 suggest that oil and gas firms pay an average $267,000 in transaction costs, and that all other firms pay $562,000; I cannot reject the hypothesis that, failing to account for previous offsets use, energy sector firms pay no significant transaction cost to use offsets.

The third column again combines all sectors and estimates transaction costs separately for returning users. The transaction cost for returning offset users is negative (-$400,000), meaning returning users would need to be paid this much to be convinced *not* to use offsets. That the transaction cost for first-time offset users doubles to $510,000 when this variable is included reflects that the average transaction cost in column 1 is pulled down by returning offset users when this is not included as an explanatory variable.

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 $220,000 in the energy sector, $426,000 in the oil and gas sector, and $835,000 for all other firms. While first-time users show an increasing transaction cost across these sectors, respectively, returning

[15] The coefficients underlying these calculations are presented in table D.1.

Table 3: 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
Sector indicators?	No	Yes	No	Yes
Previous offset use control?	No	No	Yes	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 D.3 for these transaction cost calculations, and see table D.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 no 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.

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

users show a decreasing transaction cost. In particular, returning offset users in the energy sector, the oil and gas sector, and all other firms have transaction costs of -\$697,000, -\$491,000, and -\$82,000, 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.[16] This may be the case if firms that previously used offsets accumulated human capital (e.g. hiring a new compliance officer) that makes future non-participation costly. Alternatively, it may reflect commitments to offset suppliers requiring fees to cancel. 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. None of these potential explanations are observable in the current data, so decomposing the nature of the transaction costs remains a promising direction for future research.

There are several caveats to the transaction cost estimates. In theory, the returning user coefficient may partly reflect persistent unobserved heterogeneity rather than a true transaction cost differential from returning to the market. With only three compliance periods, a fully dynamic panel approach separately identifying true path dependence is infeasible. As a robustness check, I include firm-level means of maximum cost savings following Mundlak (1978) to partial out time-invariant unobserved firm heterogeneity; transaction cost estimates from this model are in table D.3 and the coefficients in table D.4. These transaction costs are smaller in magnitude, but reflect the same pattern as those above.

As another robustness check, I include the firm-level means as well as an indicator for whether the firm used offsets in the first compliance period. This controls for path-dependence which may bias the previous offset use indicator (Wooldridge, 2005); known as the "initial conditions problem." The transaction costs from this check and the underlying coefficients are in tables D.5 and D.6, respectively. The coefficient on initial offset use and

[16]One might reasonably debate whether sunk costs reflect *true* transaction costs. In the sense that they represent costs incurred by offsets users, these are genuine transaction costs; however, in the sense that these are costs incurred *in order to use* offsets, these are only genuine transaction costs if one assumes that firms are forward-looking and consider these costs before using offsets the first time.

on previous offset use are both statistically significant at the 1% level. The coefficient on initial offset use is always positive, but it increases when controlling for previous offset use. This implies that the firms that used offsets in the first period are more likely to use them in future periods; the fact that the coefficient on previous use is negative suggests that the firms that use offsets for the first time in the second period are slightly less likely to use them in the third. Taken together, this suggests that the baseline transaction costs in table 3 may modestly overstate the true cost, thus representing an upper bound on the transaction cost (hence, a lower bound on the inefficiencies caused by transaction costs).

Table 4: Transaction cost estimates at different quantiles ($1k)

	Probit	$\tau = 0.1$	$\tau = 0.25$	$\tau = 0.5$	$\tau = 0.75$	$\tau = 0.9$
Combined	259.62***	-43.12	164.76	53.51***	15.58	1.72
	(54.13)	(42.84)	(101.90)	(6.21)	(13.67)	(1.44)
Energy	220.60*	1.21	-24.64	37.69***	0.35	-1.11
	(125.18)	(1.19)	(91.57)	(9.42)	(6.25)	(2.38)
Oil/gas	426.85**	-0.65	-38.39	49.37***	0.26	0.21
	(190.03)	(1.52)	(131.88)	(14.61)	(8.67)	(3.07)
Other	835.68***	-1.94	-44.33	56.14***	28.07***	10.63***
	(227.76)	(2.97)	(151.05)	(8.71)	(9.16)	(4.00)

Probit uses robust standard errors; BQR standard errors recovered through bootstrap with 500 repetitions. Testing for significantly different transaction costs at the targeted quantiles yields F statistics of 80.33 for the combined model and 19.51, 13.24, and 49.99 for the energy, oil/gas, and other sectors, respectively. The transaction cost estimates (coefficient values) for the combined model and for the sector-specific model are presented at each ventile in figures D.1a and D.1b (D.2a and D.2b), respectively.

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

Quantile transaction cost estimates The BQR estimates from two models are presented in table 4. The leftmost column contains the average transaction cost estimates (from the probit) compared to the 10$^{\text{th}}$, 25$^{\text{th}}$, 50$^{\text{th}}$, 75$^{\text{th}}$, and 90$^{\text{th}}$ quantiles.[17] The first row of estimates is from the combined model with no controls for offset history or for a firm's

[17] The transaction cost estimates across the full error distribution for the combined model and for the sector-specific model are presented graphically in figures D.1a and D.1b, respectively. The coefficients underlying these estimates are presented graphically in figures D.2a and D.2b, respectively.

sector. Below this, the transaction cost estimates for the three main sectors designated by CARB are presented. The combined average estimates come from the first column, and the sector average estimates come from the second column of table 3.

In the combined model, the median ($\tau = 0.5$) firm incurs a transaction cost of about \$53,000. The sector-specific transaction costs increase following the same pattern as the average, with the median energy transaction cost about \$37,000, the median oil/gas transaction cost about \$49,000, and the median transaction cost for all other firms about \$56,000. The main difference between the probit and median transaction cost estimates is the magnitude. That the probit estimates transaction costs that follow the same pattern in sector heterogeneity, but are much larger in magnitude points to large tails in the latent transaction cost distribution.

At $\tau = 0.1$ and $\tau = 0.25$, the estimated transaction costs are negative but insignificant. Thus, I cannot reject the hypothesis that there are no transaction costs at this point in the distribution. A negative transaction cost could suggest that these low-propensity firms make decisions based on factors beyond the maximum cost savings. For example, existing business relationships or (un)willingness to adopt new compliance instruments. At $\tau = 0.75$ and $\tau = 0.9$, the transaction costs are all positive, but only significant for firms in the "other" sector. These transaction costs are smaller than at the median, suggesting firms at this point in the distribution have fewer barriers to using offsets.

I calculate an F-statistic as a formal test of varying transaction costs at different quantiles. The null hypothesis in this test is that the transaction costs are equivalent at the 5 quantiles presented here. I do this for the combined model and separately for each sector. All of these tests reject the null hypothesis at the 1% significance level.

6.2 Welfare

Table 5 collects the components of the welfare calculation, reported separately for each of the three sectors and for the economy as a whole. Every component that depends on transaction cost is presented twice: once using the mean (probit, row 1 column 1 of table 4) and using the median (BQR, row 1 column 4 of table 4). The remaining components—gross earned savings, gross missed savings, and the frictionless benchmark—are invariant to the choice of transaction cost estimate. Gross earned savings is the sum of realized cost savings ($p_t\tilde{q}_{it}$) across all firms that used offsets; gross missed savings is the analogous sum across firms that did not. The frictionless cost savings is the total savings that would be earned if there were no transaction costs and there was full participation: all firms use offsets and earn the total savings. I use this full-participation world as the benchmark against which the inefficiency induced by transaction costs is measured.

The gross figures demonstrate substantive savings potential. Regulated firms realized \$346.4 million in gross cost savings over the first three compliance periods, against \$53 million in savings that non-participating firms left unrealized; full frictionless participation would have generated \$399.4 million. Further, the savings potential is not uniform across sectors. Oil and gas firms account for the bulk of realized savings (\$221.4 million), followed by energy (\$94.5 million) and all other firms (\$30.6 million). This ranking mirrors the sectors' emissions and hence their proportional offset entitlements.

Netting out the transaction costs from the gross savings yields the net achieved savings, which can be interpreted as the realized value of the offsets market. Under the median estimate, firms incurred \$19.9 million in transaction costs and retained \$326.5 million in net savings; under the mean transaction cost estimate, these estimates are \$95.5 million in costs and retained \$250.9 million. The gap between the estimates using the mean and the median is large because the difference in transaction cost estimates appears to be heavily right-skewed, as suggested by the difference in the magnitude of the transaction cost estimates. The sectoral pattern is preserved under both estimates: oil and gas captures roughly two-thirds of net savings (\$217.2 million at the median), with energy contributing \$85.2 million.

The skew is most consequential for firms in the "other" sector, where the net achieved savings have different signs. Under the mean transaction cost, the transaction costs incurred by participating other-sector firms ($30.9 million) exceed the gross savings those firms earned ($30.6 million), implying a net achieved savings of essentially zero (-$287,000). Taken literally, this says that the firms observed using offsets collectively spent more to do so than they saved—thus, they participated inefficiently. Using the median transaction cost instead, incurred costs fall to $6.4 million and net achieved savings turn solidly positive at $24.2 million. This is because the median is robust to skewed error terms. This sector has the greatest variance in the types of participants, which may be a part of the explanation for this.

The implied inefficiency is frictionless benchmark less net achieved savings. This is equal to the sum of foregone savings by non-participants and the transaction costs that participants actually incurred. Both parts of this decomposition are "losses" relative to a frictionless world with full participation. Combining the two, transaction costs generated $72.9 million in inefficiency under the median transaction cost estimate ($53 million in foregone savings plus $19.9 million in incurred costs) and $148.5 million under the mean transaction cost estimate ($53 million plus $95.5 million). The median inefficiency is $33.2 million in energy, $16.1 million in oil and gas, and $23.5 million for all other firms. The avoided transaction costs reported in Table 5 (those that non-participating firms did not pay) do not enter this calculation since the frictionless benchmark contemplates no transaction costs at all. For the "other" sector in particular, however, the median avoided cost ($11.5 million) falls short of the savings those firms missed ($17 million), indicating that some non-participation reflects genuine inefficiency rather than an efficient response to high costs.

Several caveats apply to this welfare analysis. The frictionless counterfactual assumes that every non-participating firm would both enter the market and capture its full potential savings if there were no transaction costs; further, it assumes that prices would remain fixed despite the resulting surge in offset demand. Both of these assumptions individually increase the counterfactual cost savings, which leads to higher inefficiency estimates. The calculation also values offsets at their compliance-cost savings, ignoring their environmental integrity. Thus, to the extent that some credits do not represent non-additional offsets, the gross savings overstate the social value of participation. Finally, the three-period window captures only the program's first decade, and whether transaction costs rise or fall as the market matures is ex ante ambiguous.

6.3 Threats to exogeneity

Unobserved heterogeneity The payoff to firm i could be

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

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 $\rho_j + \epsilon_{ijt}$, as before. If, however, $\alpha_i \neq 0$, this is not the case. For instance, if $Gains_{ijt}$ (a function of emissions) is 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 (Neyman & Scott, 1948). 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 (Wooldridge, 2010).

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:

$$\alpha_i = \delta_0 + \delta_1 \overline{Gains}_{ij} + u_i$$

where $\overline{Gains}_{ij} = \frac{1}{T}\sum_{t=1}^{T} Gains_{ijt}$ is the time average of the cost savings from using offsets, and $u_i \sim N(0, \sigma_u^2)$ is uncorrelated with firm sector or offset price after controlling for this. 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 the variance is estimated from the data.

The transaction costs implied by these auxiliary models (presented in table D.3) are comparable to those implied by the standard probit. At the same time, the estimated variance attributed to the heterogeneous random error (i.e., $\hat{\sigma}_u^2$) has a value ranging between 0.5 and 1. Since the expected value of u_i is zero by assumption, a small estimated correlated random error variance and insignificant coefficient on the Mundlak controls together imply that the gains in terms of identification from this are relatively small.

Endogeneity of lagged offset adoption The persistence in offset use shown in table 2 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 may create 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, we have that $\text{Cov}\big(\mathbb{1}^O_{ij,t-1}, \epsilon_{ijt}\big) > 0$. 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 predicted transaction costs for this specification (table D.5) are much larger than the main probit estimates. As such, the estimated welfare losses calculated using the smaller probit estimates can be interpreted as a lower bound.

7 Conclusion

The results of this paper suggest that fixed transaction costs are a meaningful barrier to offset adoption in California's cap-and-trade program. Conservative estimates suggest firms incur a combined \$19 million in transaction costs over the first three compliance periods, generating roughly \$72 million in efficiency losses relative to a frictionless market. The nature of these costs remains an open question.

I propose three candidate explanations, each with different policy solutions. Search costs can be reduced through better market infrastructure, such as more efficient centralized registries and standardized contracts. Invalidation risk can be mitigated through improved verification standards or increased safety nets such as buffer pool mechanisms for avoided deforestation offset projects. Public relations concerns, however, do not have a clearcut fix. 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 (Aspelund & Russo, forthcoming; Calel et al., 2025; West et al., 2023) lends credence to the possibility of a public relations explanation. Studies documenting widespread non-additionality in offset programs (e.g. being credited with avoided deforestation by paying to preserve a forest that was not at risk of deforestation) have attracted significant media attention—including California's own forest offset protocol (Badgley, Chay, et al., 2022; Liu & Cui, 2017). In an environmentally conscious state like California, the reputational risk of being associated with non-additional offsets (even if the association is unjustified) may be a rational concern for firms.

The persistence of offset use documented in table 2 can represent such public relations costs as well as it can sunk capital investment costs. 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 positive welfare gains.

Moreover, the transaction cost is incurred by any offset-using firm, regardless of the environmental benefits gained. This has immediate welfare implications. A truly additional offset generates environmental benefit while a non-additional offset generates none. As illustrated in this paper, Californian firms incur the same transaction cost to purchase either one of these. As such, purchasing non-additional offsets is a pure welfare-decreasing action: the transaction cost is transferred to nobody, and no environmental benefit is accrued. If some of the transaction cost arises from true average 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 those participants purchasing entirely additional offsets, decreasing the motivation to participate in the market. Policymakers seeking to improve offset market efficiency may therefore first focus on integrity reforms with active communication efforts, thus improving additionality and lowering 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

Appunn, K., & Wettengel, J. (2024). Understanding the European Union's Emissions Trading Systems (EU ETS) [magazine]. *Clean Energy Wire: Journal for the energy transi-*

tion. https://www.cleanenergywire.org/factsheets/understanding-european-unions-emissions-trading-system

Aspelund, K. M., & Russo, A. (forthcoming). Additionality and asymmetric information in environmental markets: Evidence from conservation auctions. *American Economic Review*.

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., Cullenward, D., Fowlie, M., Sutter, K. R., & Brown, R. (2022, February 4). *Annual report of the independent emissions market advisory committee*. https://calepa.ca.gov/wp-content/uploads/2022/02/2021-IEMAC-Annual-Report.pdf

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

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

Fell, H., & Maniloff, P. (2018). Leakage in regional environmental policy: The case of the regional greenhouse gas initiative. *Journal of Environmental Economics and Management*, *87*, 1–23. https://doi.org/10.1016/j.jeem.2017.10.007

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

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. (2011). The effect of allowance allocations on cap-and-trade system performance. *The Journal of Law and Economics*, *54*(S4), S267–S294. https://doi.org/10.1086/661942

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

Johnson, K. C. (2023). California's ambitious greenhouse gas policies: Are they ambitious enough? *Energy Policy*, *177*, 113545. https://doi.org/10.1016/j.enpol.2023.113545

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

Lessmann, C., & Kramer, N. (2024). The effect of cap-and-trade on sectoral emissions: Evidence from California. *Energy Policy*, *188*, 114066. https://doi.org/10.1016/j.enpol.2024.114066

Liu, X., & Cui, Q. (2017). Baseline manipulation in voluntary carbon offset programs. *Energy Policy*, *111*, 9–17. https://doi.org/10.1016/j.enpol.2017.09.014

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

Mayr, S., Pokorny, B., Montero-de-Oliveira, F.-E., & Reinecke, S. (2025). Scaling agroforestry through payments for ecosystem services: A scoping review. *Climate Policy*, 1–20. https://doi.org/10.1080/14693062.2025.2490205

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

Neyman, J., & Scott, E. L. (1948). Consistent Estimates Based on Partially Consistent Observations. *Econometrica*, *16*(1), 1. https://doi.org/10.2307/1914288

Pan, C., Shrestha, A., Innes, J. L., Zhou, G., Li, N., Li, J., He, Y., Sheng, C., Niles, J.-O., & Wang, G. (2022). Key challenges and approaches to addressing barriers in forest carbon offset projects. *Journal of Forestry Research*, *33*(4), 1109–1122. https://doi.org/10.1007/s11676-022-01488-z

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

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

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

Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data* (2. ed). MIT Press.

Table 5: Numbers used in welfare calculations

	Combined	Energy	Oil and gas	Other
Gross earned savings (+)	\$346,446,674	\$94,484,645	\$221,354,236	\$30,607,793
Gross missed savings (−)	\$52,954,438	\$23,931,604	\$11,976,328	\$17,046,506
Frictionless cost savings (hypothetical)	\$399,401,112	\$118,416,249	\$233,330,564	\$47,654,300
Incurred transaction cost (−)				
mean:	\$95,540,797	\$44,654,938	\$19,990,873	\$30,894,986
median:	\$19,905,681	\$9,303,742	\$4,165,047	\$6,436,891
Avoided transaction cost (+)				
mean:	\$91,386,849	\$24,664,064	\$11,682,978	\$55,039,807
median:	\$19,040,217	\$5,138,695	\$2,434,119	\$11,467,403
Net achieved savings				
mean:	\$250,905,877	\$49,829,707	\$201,363,362	-\$287,193
median:	\$326,540,993	\$85,180,903	\$217,189,188	\$24,170,902
Implied inefficiency				
mean:	\$148,495,235	\$68,586,542	\$31,967,201	\$47,941,492
median:	\$72,860,120	\$33,235,347	\$16,141,375	\$23,483,398

Note: Next to each welfare component is an indicator in parentheses. A (+) indicates that the component contributes positively to overall cost savings, and a (−) indicates that the component contributes negatively to overall cost savings. For each estimate that includes a transaction cost (i.e., all but gross values and the frictionless hypothetical), the estimated value calculated using the average transaction cost is presented in the row above the estimated value calculated using the median transaction cost. Gross earned savings is the dollar value of total cost savings gained by users of offsets; gross missed savings is the dollar value of total cost savings that firms could have earned if they used offsets, but did not; The frictionless cost savings assumes zero transaction cost and full participation, so it is the sum of gross earned savings and gross missed 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, and the net achieved savings is the difference between gross earned savings and incurred transaction costs. The inefficiency implied by the existing transaction costs is the difference between the frictionless cost savings and the net achieved savings.

A 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 A.1 matches table 3, table A.2 matches table D.3, table A.3 matches the first row of table 4, and table A.4 matches the bottom three rows of table 4. Figure A.1 matches figure D.2a, figure A.2 matches figure D.1a, figure A.3 matches figure D.2b, figure A.4 matches figure D.1b, figure A.5a matches figure D.3a, and figure A.5b matches figure D.3b.

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 A.1 and A.3) exhibit patterns nearly identical to their counterparts in appendix D, and the transaction cost estimates at the median (Figures D.1a and D.1b) are statistically indistinguishable from the baseline results. This exercise suggests that the choice of price measure does not materially affect the conclusions of the paper.

Table A.1: 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 D.3 for these transaction cost calculations, and see table D.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 no 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 A.2: 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 A.1 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 no 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: 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 repetitions. 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 A.4: 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 repetitions. 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 D.1 and D.2, respectively. Figure D.3 presents the estimated offset adoption rates compared with the observed rates as a goodness-of-fit exercise (Kordas, 2006).

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

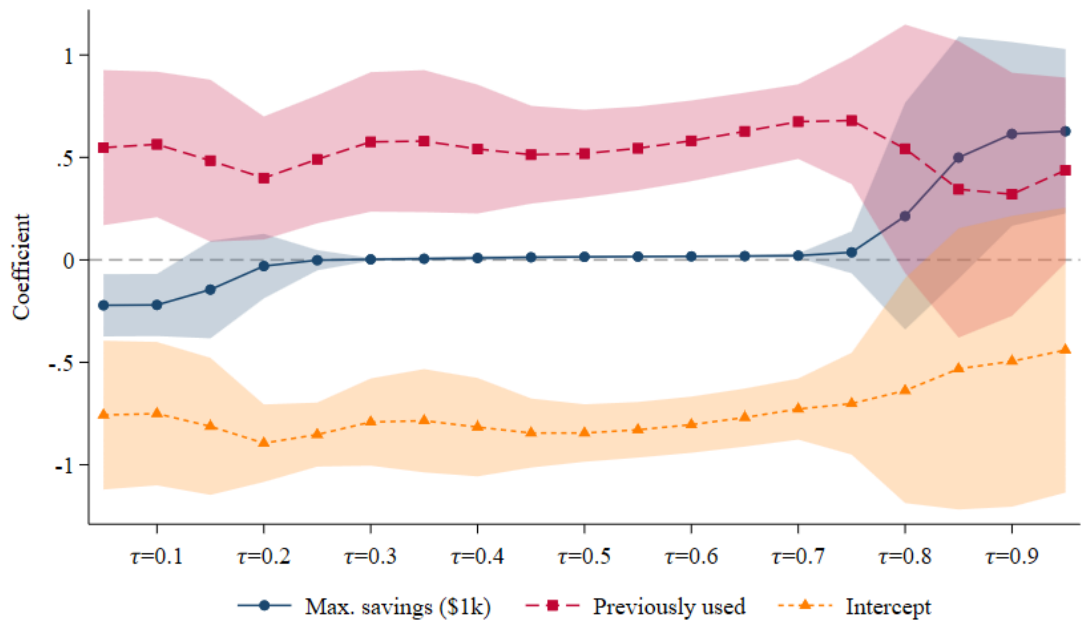

Figure A.1: BQR coefficients—combined model, using average compliance period price

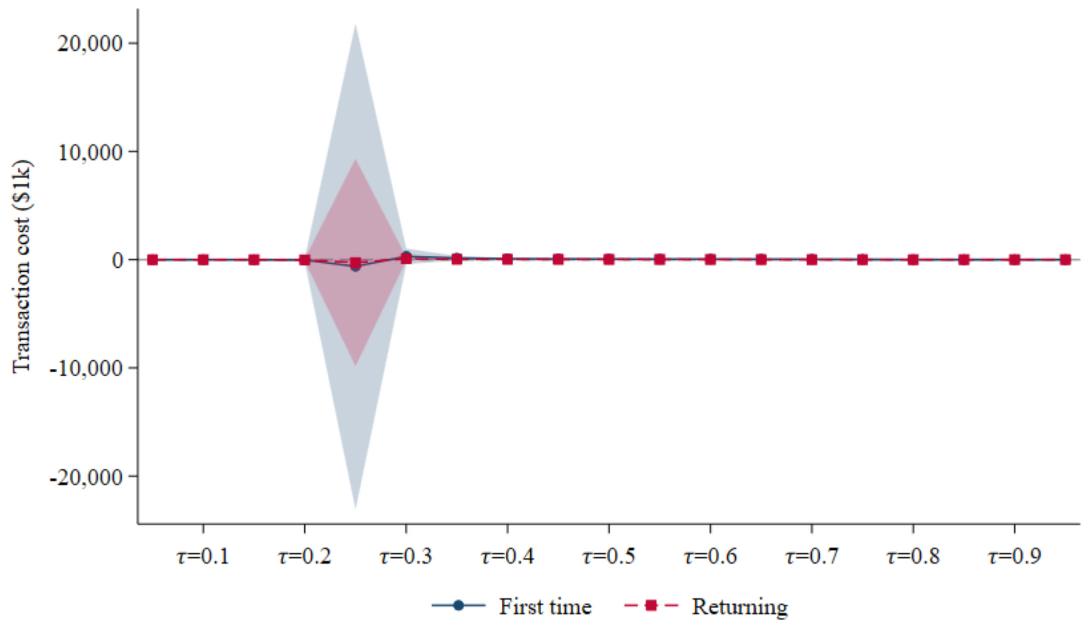

Figure A.2: Combined transaction costs at different quantiles, using average compliance period price

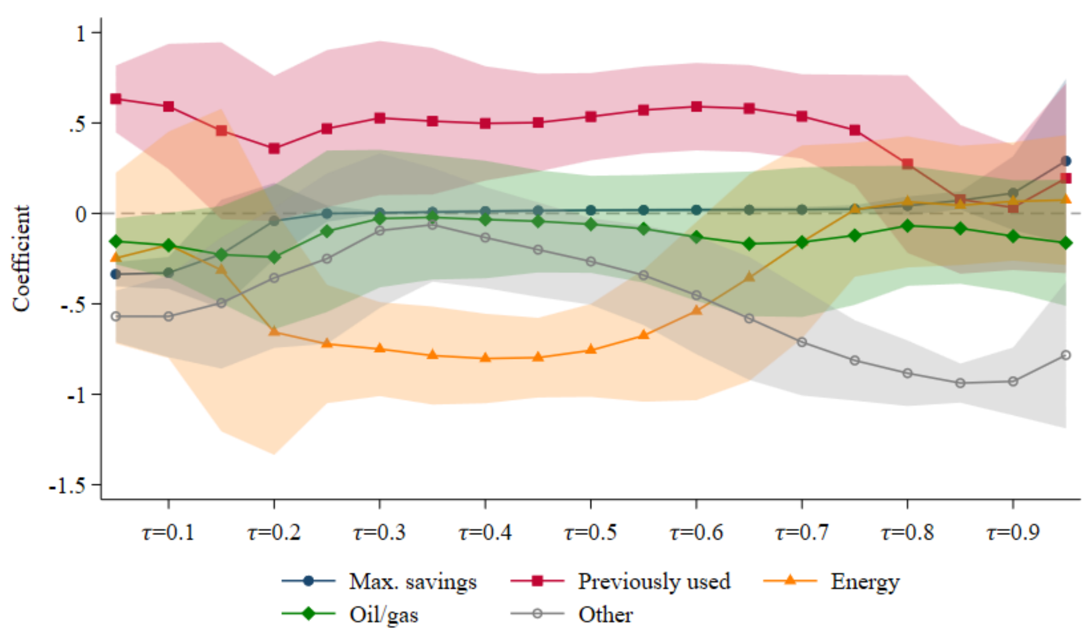

Figure A.3: BQR coefficients—sector-specific model, using average compliance period price

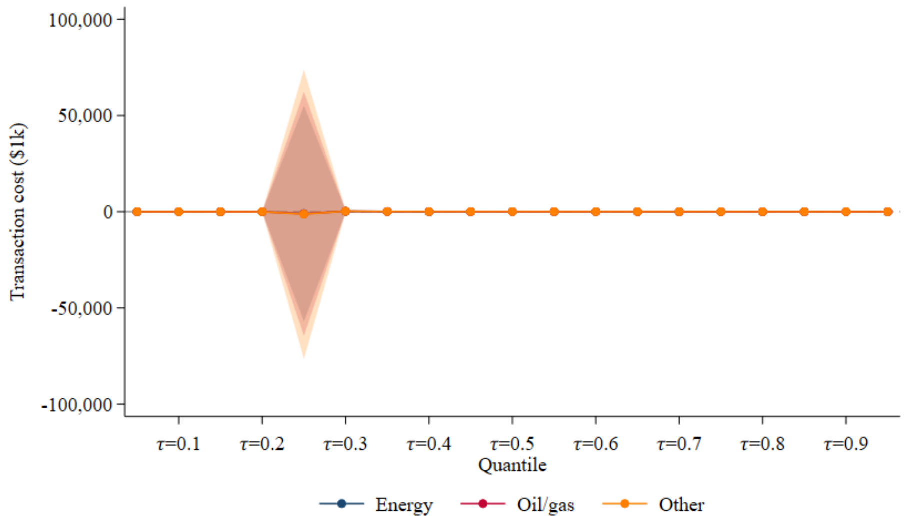

(a) First-time offset users

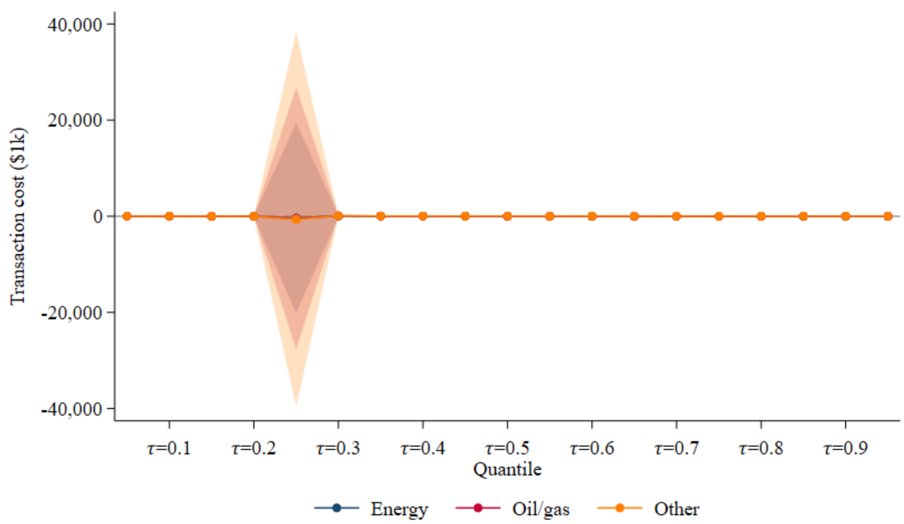

(b) Returning offset users

Figure A.4: Sector-specific transaction costs at different quantiles, using average compliance period price

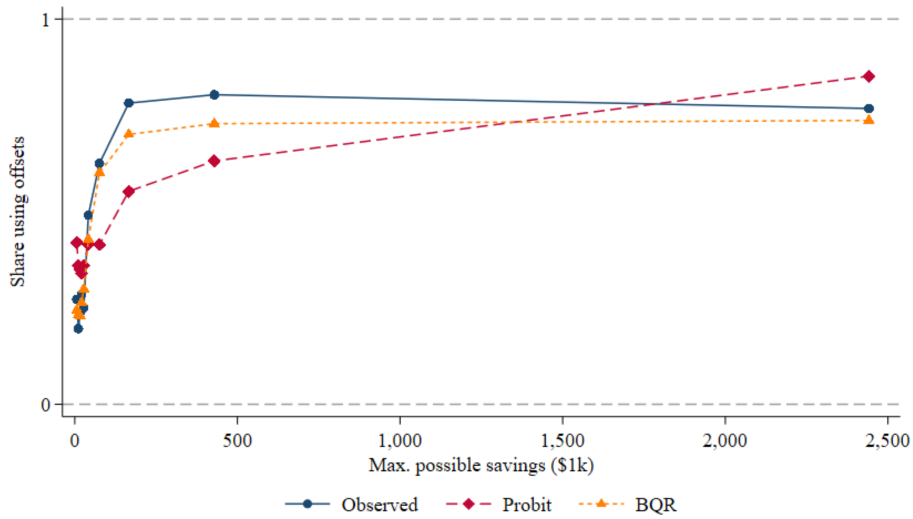

(a) Combined sector model

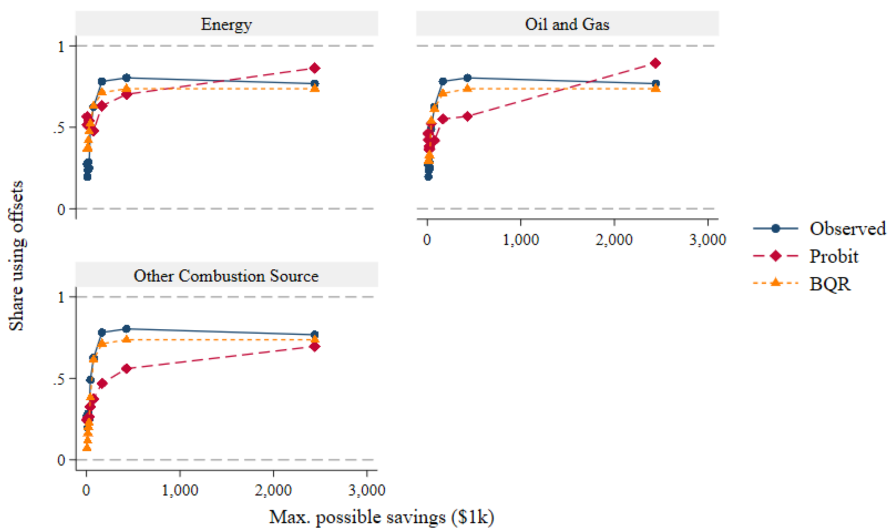

(b) Sector-specific model

Figure A.5: Predicted v. realized shares using offsets

B Hurdle model

It can be useful to try to understand the mechanism driving "interior" offset quantities. 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 step will face consistency issues due to the large number of zeros.

To address this, 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. If they enter, they 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 log-linear function of the maximum possible offsets limit and the sector identifier. The complete hurdle model is defined in equation (B.1):

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

where O_{ijt} is the quantity 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 rule out per-unit transaction costs in two ways. First, given that the second stage of the model returns transaction costs 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 estimate the relationship between the offset limit and the number of offsets purchased. The second stage is an exponential function, so including the natural log of the maximum offset limit as a right-hand 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 B.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 B.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 B.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 B.1: Transaction costs conditional on participation

	(1)	(2)	(3)	(4)
Combined (first time)	218,951		189,926	
	(996,982)		(742,511)	
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)
Combined (returning)	218,951		217,228	
	(996,982)		(851,197)	
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 transaction cost is absorbed in the first stage. This insignificance suggests the transaction cost affects only the extensive margin.

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

Table B.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 B.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.

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

C Including the fourth compliance period

(a) Offset usage during the first three compliance periods

(b) Offset usage in the fourth compliance period

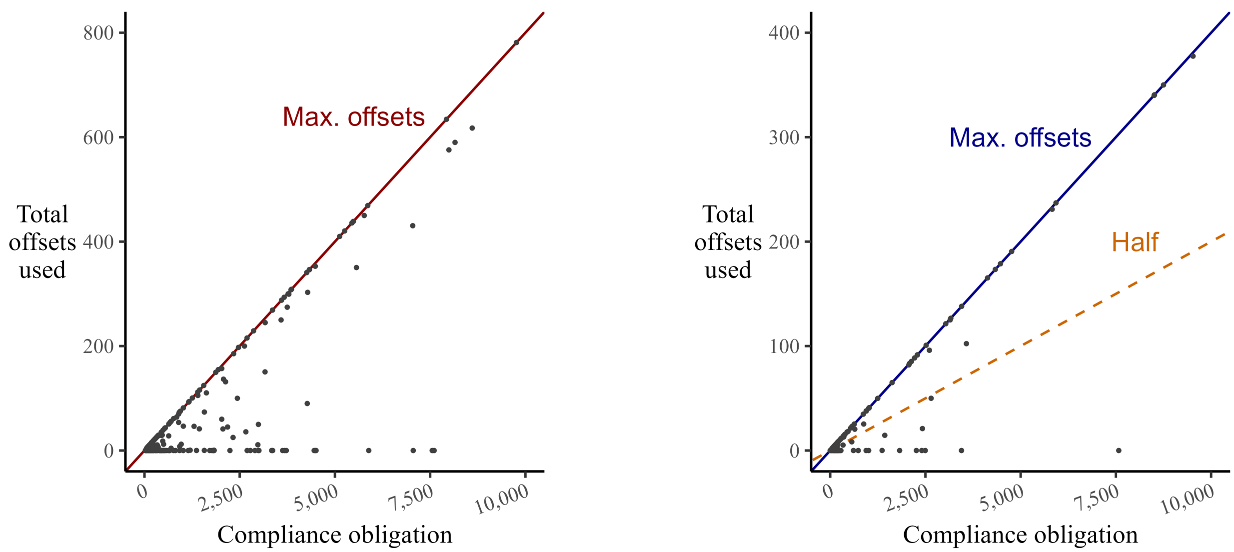

Figure C.1: Relationship between offsets used and compliance obligation

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 C.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 C.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.

Figure C.1a is identical to figure 2; figure C.1b is the analogous figure using data only from the fourth compliance period. The solid blue line in figure C.1b represents the maximum amount of offsets that can be used (4% of the compliance obligation) while the orange dashed line represents the maximum amount of non-DEBS-providing offsets that can be used (2% of the compliance obligation). As in the first three compliance periods, firms will frequently use the maximum amount of offsets possible, or will use no offsets at all. There does not appear to be a clear clustering around the non-DEBS limit, which would be expected if there were an additional fixed cost to use DEBS-providing offsets.

Figure C.2 allows for a closer examination of offset adoption patterns. It shows CDFs across the distribution of the share of allowable offsets used. In particular, C.2a shows the

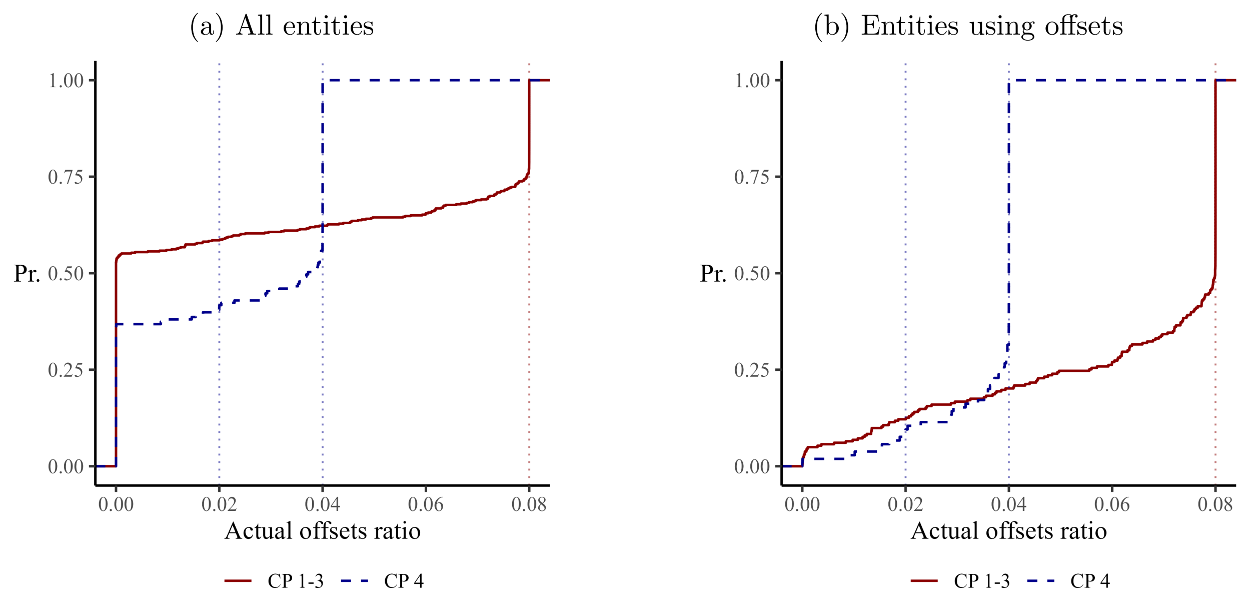

Figure C.2: Cumulative density offset share

Note: The tall vertical line of the CDFs at 0 in figure C.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 C.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.

unconditional CDF and C.2b shows the CDF conditional on some use of offsets. The solid red lines are identical to 3, and the dashed blue lines are the CDFs using only data from the fourth compliance period. The fourth compliance period mirrors the near-discrete jump in the share of offsets used from 0 to the upper bound that occurs in the first three periods. The lower "vertical intercept" of the dashed blue line compared to the solid red line shows that a higher share of firms use offsets in the fourth compliance period. 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. 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.

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

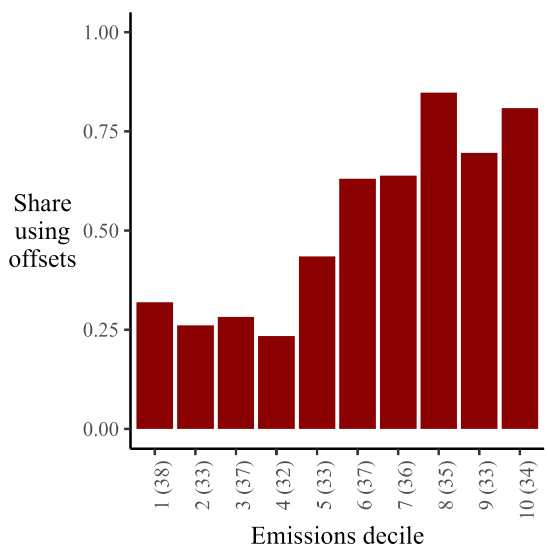

(b) Share of firms using offsets in the fourth compliance period

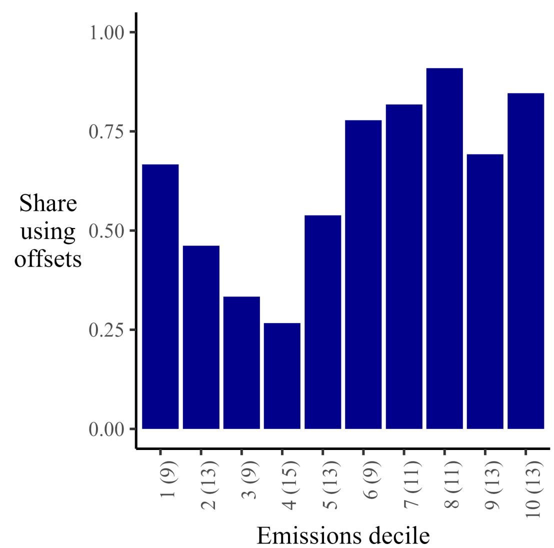

Figure C.3: Larger firms are more likely to use offsets

Note: Firms are grouped into deciles based on their total emissions. The height of the bar indicates the share of firms in 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 C.3 shows offset adoption at different emissions deciles for the first three compliance periods (figure C.3a) and for the fourth compliance period (figure C.3b). In general, the fourth compliance period has a higher adoption rate at the lowest decile, but otherwise shows the same patterns as the first three periods. The pattern of higher-emitting firms being more

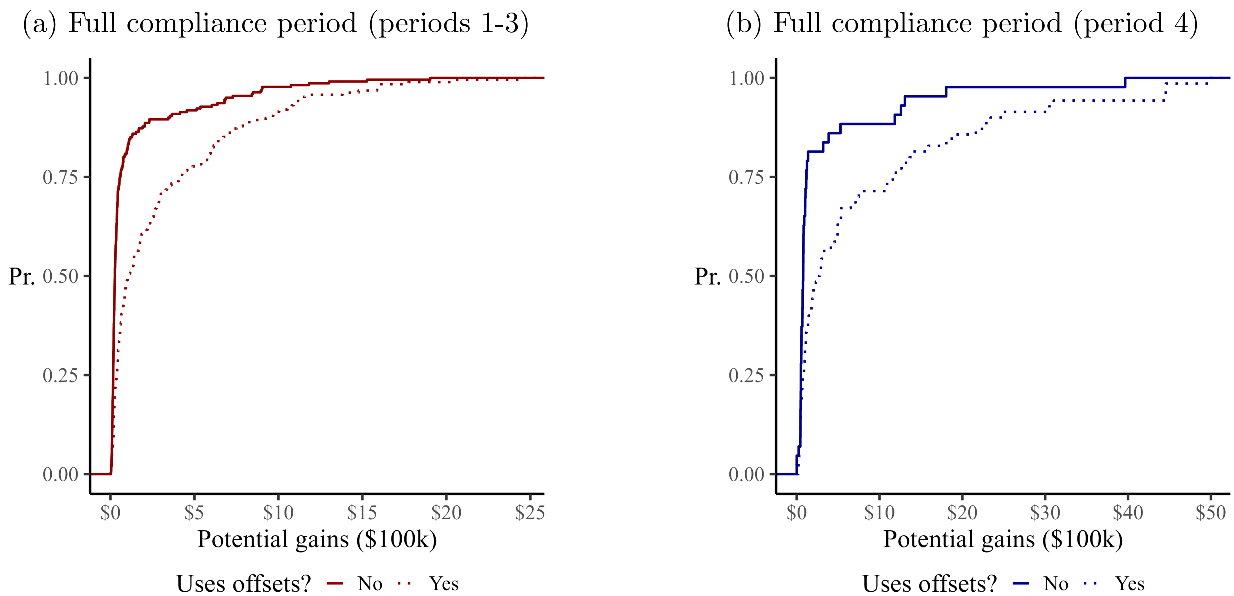

Figure C.4: Potential gains from offset adoption

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 is true for all x, so the expected gains are higher among firms that use offsets than among those that do not. Firms may be selecting into offset adoption based on the potential gains.

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.

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. However, the price of permits increased substantially. 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.

Figure C.4 compares the CDFs of potential gains by offset-using firms with those that do not. The solid lines are the CDFs for firms that do not use any offsets while the dotted lines are the CDFs of firms that do. The pattern of offset-using firms first order stochastically dominating firms that do not is robust across the two timeframes.

D Regression estimates and alternative specifications

Table D.1: Coefficients: used in table 3

	(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 3.

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

Figures D.3a and D.3b 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: D.3a shows the combined model and D.3b shows the predictions by

Table D.2: Coefficients used in average transaction cost calculation

	(1)	(2)	(3)	(4)
Uses offsets this period				
Constant	-0.305***		-0.479***	
	(0.064)		(0.069)	
Max. benefit ($1k)	0.001***	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Energy		0.061		-0.183*
		(0.099)		(0.109)
Oil and gas		-0.272*		-0.354**
		(0.150)		(0.149)
Other		-0.571***		-0.693***
		(0.085)		(0.086)
Previously used			0.855***	0.761***
			(0.136)	(0.141)
Obs.	556	556	556	556
Sector heterogeneity?	No	Yes	No	Yes
Previous offset use control?	No	No	Yes	No
Likelihood ratio	21.92	63.73	69.88	97.26
p-value likelihood ratio	0.000	0.000	0.000	0.000

These coefficients are used to calculate the average combined and average sector transaction costs presented in table 3.

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

Table D.3: Mean transaction cost by sector from CRE models ($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
Sector dummy?	No	Yes	No	Yes
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	Yes	Yes	Yes	Yes
Initial conditions control?	No	No	No	No
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 3 for calculations using standard probit methods and table D.4 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 no 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 D.4: Estimated coefficients from CRE models

	(1)	(2)	(3)	(4)
Uses offsets this period				
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		.13152		.00609
		(.1615)		(.16341)
Oil and gas		-.30655		-.32178
		(.25266)		(.22261)
Other		-.81753***		-.79199***
		(.156)		(.13746)
Previously used			.47802**	.34456
			(.22542)	(.22362)
Constant	-.40296***		-.45069***	
	(.11244)		(.09561)	
$\ln \sigma_u^2$	.09223	-.05184	-.60578	-.51828
	(.29821)	(.30823)	(.55894)	(.50258)
Obs.	556	556	556	556
Sector dummy?	No	Yes	No	Yes
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	Yes	Yes	Yes	Yes
Initial conditions control?	No	No	No	No
Likelihood ratio	24.81	41.42	38.30	54.49
p-value likelihood ratio	0.000	0.000	0.000	0.000

These estimated coefficients correspond by column to the transaction cost estimates in table D.3.

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

Table D.5: Mean transaction cost by sector from CRE models with initial conditions controls ($1k)

	(1)	(2)	(3)	(4)
Combined (first time)	1,295***		969.69***	
	(447.96)		(327.57)	
Energy (first time)		1,520***		1,124***
		(581.32)		(410.23)
Oil and gas (first time)		968.88**		730.35**
		(402.87)		(298.32)
Other (first time)		1,396***		1,035***
		(532.17)		(374.64)
Combined (returning)	1,295***		1,659***	
	(447.96)		(543.26)	
Energy (returning)		1,520***		1,822***
		(581.32)		(630.94)
Oil and gas (returning)		968.88**		1,428***
		(402.87)		(510.65)
Other (returning)		1,396***		1,733***
		(532.17)		(601.63)
Obs.	556	556	556	556
Sector dummy?	No	No	No	Yes
Previous offset use control?	No	No	No	Yes
Mundlak controls?	Yes	Yes	Yes	Yes
Initial conditions control?	Yes	Yes	Yes	Yes
p returning			0.010	0.014
p sector heterogeneity		0.290		0.316

These are all correlated random effects (CRE) models with 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 of 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 D.6: Estimated coefficients from CRE models with initial conditions controls

	(1)	(2)	(3)	(4)
Uses offsets this period				
Max. benefit (\$1k)	.00082***	.00078***	.0011***	.00106***
	(.00029)	(.0003)	(.00039)	(.00039)
Firm average max. profit (1k)	-9.8e-05	-.00014	-.00025	-.00029
	(.00022)	(.00023)	(.00032)	(.00032)
Initial offsets use	1.909***	1.9903***	2.4369***	2.4914***
	(.13269)	(.14945)	(.26621)	(.27759)
Energy		-1.1853***		-1.1869***
		(.14894)		(.1722)
Oil and gas		-.75544***		-.77131***
		(.18299)		(.19856)
Other		-1.0887***		-1.093***
		(.10567)		(.12899)
Previously used			-.75916***	-.73696***
			(.22858)	(.22874)
Constant	-1.0564***		-1.0686***	
	(.09027)		(.11499)	
$\ln \sigma_u^2$	-12.756	-12.341	-2.7478	-2.9409
	(227.04)	(21.794)	(1.9107)	(2.2653)
Obs.	556	556	556	556
Sector dummy?	No	No	No	Yes
Previous offset use control?	No	No	Yes	Yes
Mundlak controls?	Yes	Yes	Yes	Yes
Initial conditions control?	Yes	Yes	Yes	Yes
Likelihood ratio	232.98	237.63	117.81	120.73
p-value likelihood ratio	0.000	0.000	0.000	0.000

These estimated coefficients correspond by column to the transaction cost estimates in table D.5.

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

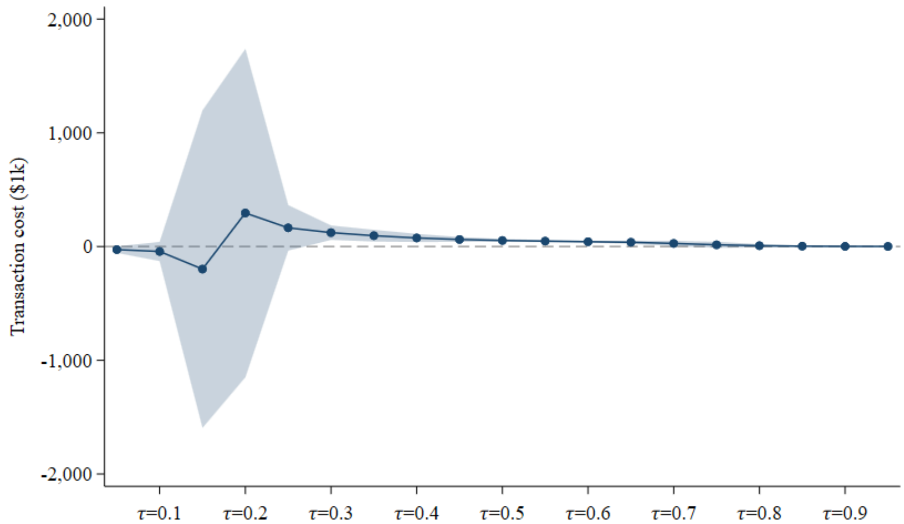

(a) Combined transaction costs at different quantiles

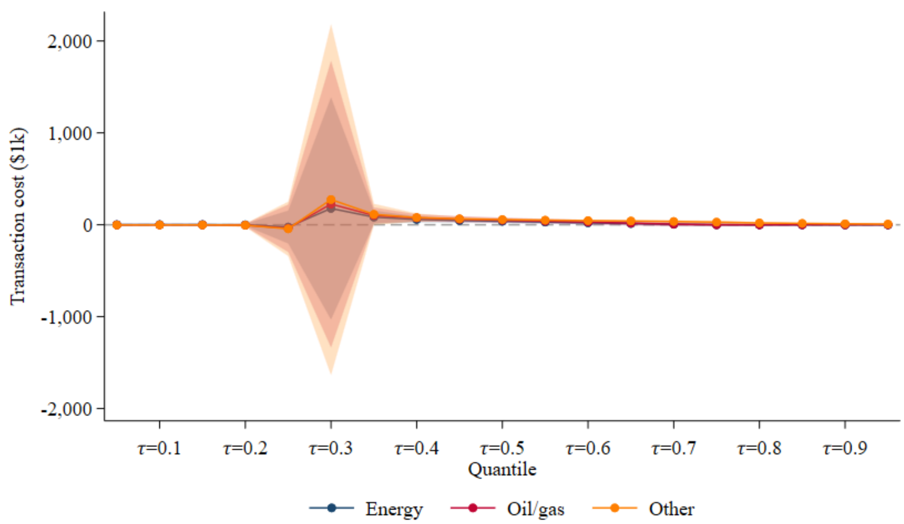

(b) Sector-specific transaction costs at different quantiles

Figure D.1: Sector-specific transaction costs at different quantiles

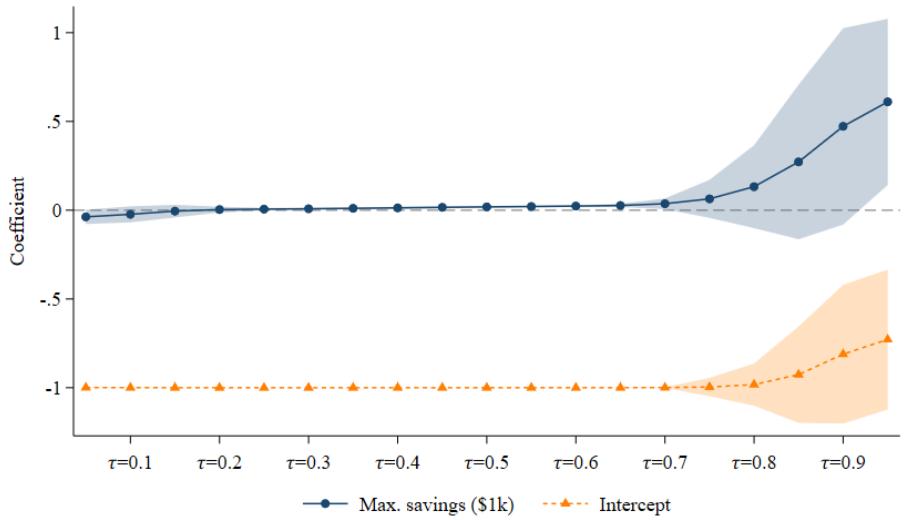

(a) BQR coefficients—combined model

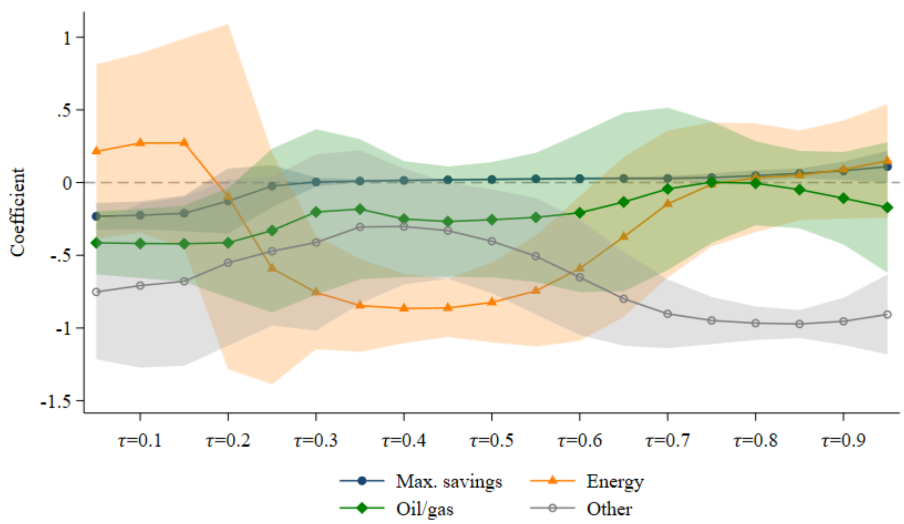

(b) BQR coefficients—sector-specific model

Figure D.2: BQR coefficients at different quantiles

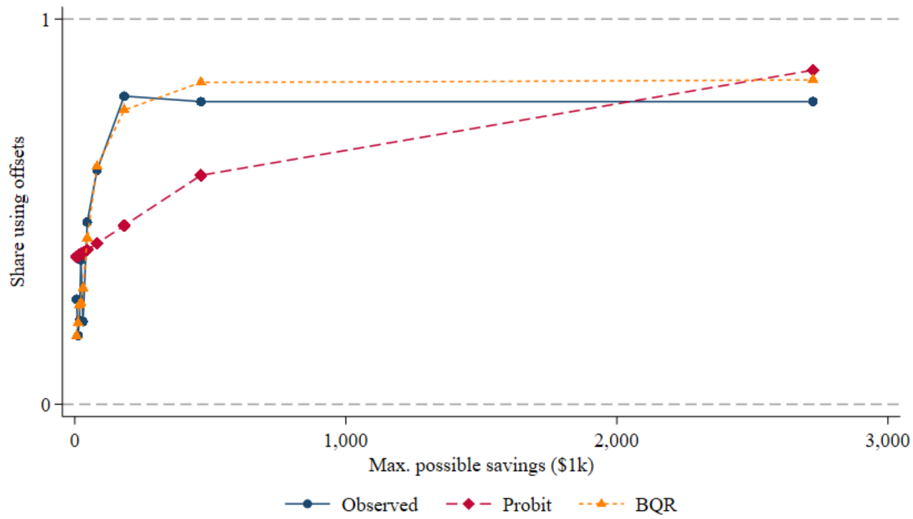

(a) Combined sector model

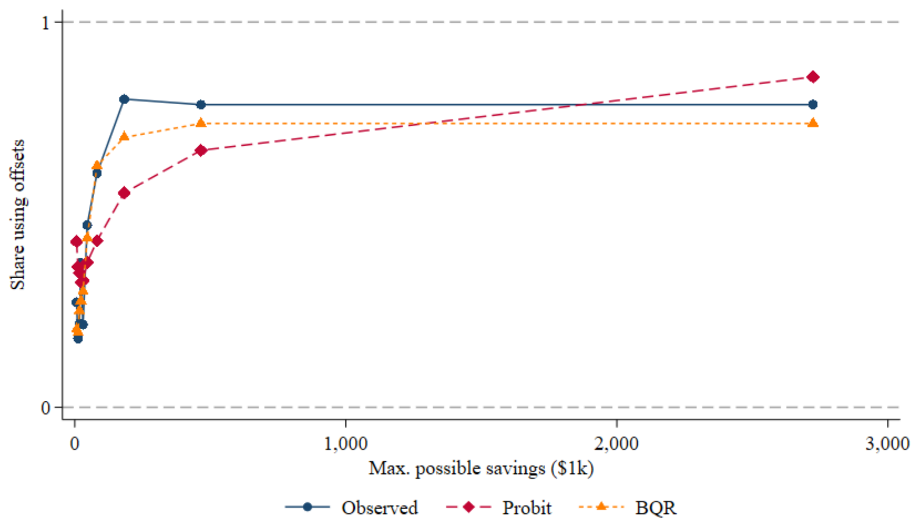

(b) Sector-specific model

Figure D.3: Predicted v. realized shares using offsets sector.