Adoption Costs of Emissions Offsets: Evidence from California

Willard Robinson*

October 20, 2025

Abstract

Despite trading at substantial discounts to standard permits, emissions offsets remain underutilized in California's cap-and-trade program, with 50-75 percent of regulated firms never adopting them. This paper provides details this money left "on the table" by the regulated firms, provides institutional background on this policy which has covered the world's fifth largest economy for over a decade, and uses compliance data to identify and quantify a significant fixed transaction cost that is preventing firms from entering the offsets market. Using compliance data from 2013-2023, I employ probit and binary quantile regression methods to estimate that regulated entities face implicit costs for adoption offset use that range from $50,000 to $988,000, varying considerably across sectors. Energy sector firms face negligible barriers, while oil and gas refineries are required to pay approximately $581,000 in fixed costs to use these cost-saving instruments; consumer-facing firms in other sectors face nearly $988,000. These costs reflect information and search expenses, risk premiums from potential offset invalidation, and public relations concerns about greenwashing. These costs are highly heterogeneous, however. Quantile regression methods show that the median value of these costs ranges from $49,000 to $64,000—approximately one-tenth of the averages. This tail-heavy result is in line with Naegele's (2018) trading costs estimates for offset adoption in the European Union's emissions trading system. California's proportional offset cap may have inadvertently concentrated cost-saving opportunities among the largest emitters, creating both efficiency losses and equity concerns. These findings suggest that policies reducing information asymmetries and improving offset credibility could broaden market participation while maintaining environmental integrity, a finding relevant to emissions trading systems globally.

*I thank my advisor Spencer Banzhaf and my dissertation committee members Roger von Haefen, Erin Sills, and Zhenxuan Wang.

North Carolina State University: wmrobin3@ncsu.edu

1 Introduction

By harnessing the power of markets, cap-and-trade programs can achieve environmental improvements at lower cost than more prescriptive command-and-control regulations. Taking advantage of this opportunity, at least 36 such markets are operating across the globe to reduce greenhouse gasses (Appunn & Wettengel, 2024). Many of these markets—including the oldest, Europe's, and the largest, China's—have attempted to further reduce costs by allowing regulated firms to purchase "offsets" from unregulated firms. This flexibility promises substantial cost reductions when unregulated firms can abate pollution more cheaply than their regulated counterparts. Despite these potential efficiency gains, offset adoption remains puzzlingly limited in California's cap-and-trade system where offsets have consistently traded at significant discounts (figure 1) yet 50-75 percent of regulated firms never use them. This paper investigates why.

I identify a substantial "adoption cost" that prevents firms from entering the offsets market, even when there are clear financial benefits. These adoption costs—fixed transaction costs—act as a hurdle that firms must overcome before participating. Using detailed compliance data from California's cap-and-trade program covering 2013-2023, I estimate that regulated entities implicitly pay between $50,000 and $988,000 in expectation to enter the offsets market, though this varies considerably across and within sectors.

Two general features of offset markets create these adoption costs. First, offsets involve heterogeneous projects that require evaluation, verification, and ongoing monitoring, so they are not as easily commodified as standardized emissions permits. This complexity may affect a firm's ability to cheaply leverage offsets for their cost savings. Second, offsets face persistent concerns over "non-additionality": whether credited emission reductions would have occurred anyway. Recent empirical evidence shows that non-additionality is widespread (Aspelund & Russo, 2025; Badgley, Chay, et al., 2022; Badgley, Freeman, et al., 2022; Calel et al., 2025; West et al., 2023), potentially creating a stigma that brings public relations costs for firms that use offsets, particularly in environmentally conscious California.

California's market design provides a uniquely clean setting to identify such adoption costs. The program limits offset use to four or eight percent of each firm's total emissions (depending on the year), creating a proportional cap that scales with firm size. Large emitters can thus realize greater absolute cost savings from offsets than small emitters, even though both face the same per-unit price discount. If adoption costs exist, we should observe larger firms participating more frequently—a prediction strongly supported by the data.

I employ both standard probit and binary quantile regression (BQR) methods to estimate average adoption costs and the distribution of adoption costs, respectively. These methods

were first applied to an application of emissions offset transaction costs in Naegele's (2018) seminal quantification of offset trading costs in the European Union's emissions trading system (EU ETS). Her research highlights significant differences in the magnitude of transaction costs to both participate in the cap-and-trade market and to use offsets. I apply her methods to California's cap-and-trade system, a newer system that more closely reflects a national ETS where all participants are involved in the same market. Like the EU ETS, offsets are cheaper than official emissions permits in California and firms become more likely to use offsets as their compliance costs increase. My paper identifies a new source of heterogeneity in offset adoption costs that is mentioned but not extensively studied in Naegele (2018): certain sectors face higher adoption costs to use offsets.

In particular, firms in the energy production sector face no significant adoption cost and this sector exhibits the highest offset participation rate (72% of firms). Oil and gas refinery firms face a positive, but smaller, adoption cost of around $581,000—half of the $988,000 faced by firms in the "other" sector (comprising universities, manufacturing facilities, and consumer-facing brands). This gap suggests that public relations concerns contribute meaningfully to adoption costs, as consumer-facing entities may fear reputational damage from association with potentially non-additional offsets.

My empirical analysis reveals three key patterns consistent with substantial fixed adoption costs. First, larger firms are significantly more likely to adopt offsets than smaller firms within the same sector—a very similar pattern to the original offset transaction costs study in Naegele (2018). Second, most firms operate at extremes either using no offsets or using the maximum legal amount, with few in between. Naegele (2018) shows this same pattern in the EU ETS, except offsets are slightly more widely adopted in the EU compared to California. Nearly 80% of EU firms use offsets while between 40-60% of Californian firms do. Third, firms persistently use offsets. A firm that uses offsets once will nearly always use them in the future. This last pattern suggests that the adoption cost functions as a one-time barrier rather than a recurring expense. Naegele (2018) focuses on Phase II of the EU ETS, thus her analysis does not specifically address this kind of adoption behavior.

This public relations concern might not be universal among the "other" firms though, as 75% of firms in this sector face an adoption cost below $39,002. The BQR results reveal that the average adoption cost can be influenced by outliers in the error distribution. The median adoption cost for first-time offset users ranges from $49,000 in the oil and gas sectors to $64,000 in all other sectors—far below the probit estimates of $582,000 to $988,000. This finding has important methodological implications for future research on transaction costs in environmental markets.

This paper makes three primary contributions. First, it provides the first comprehensive

empirical analysis of offset adoption costs in California's cap-and-trade system, quantifying barriers that previous research has acknowledged but not measured (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, a point made long ago in Buchanan and Tullock (1975). Third, it illustrates the value of quantile-based methods relative to mean-based approaches when analyzing latent variables with potentially substantial heterogeneity.

These findings carry important policy implications. The proportional offset limit, initially designed to address additionality concerns, has concentrated the benefits of offset use among large emitters while effectively excluding smaller firms. My findings also highlight the potential dangers of relying on status quo probit methods by revealing underlying heterogeneity in the adoption cost that is masked by mean-regression methods, a point originally made in Naegele's (2018) analysis of transaction costs in the European Union Emissions Trading System (EU ETS). The adoption cost creates both efficiency losses and equity concerns, as the policy advantages incumbents over smaller competitors. At the same time, it shows that there may be some socially-motivated hedging against offsets that policy-makers did not consider. Alternative designs, such as flexible offset allowances or measures to improve public information and reduce monitoring costs, could broaden participation while maintaining environmental integrity.

The remainder of the paper proceeds as follows. Section 2.1 describes California's cap-and-trade program and the institutional features relevant to offset adoption and 2.2 discusses the source and nature of the adoption cost. Section 3 first describes my data sources and presents several stylized facts about offset adoption using raw data. Section 4 outlines the empirical strategy and section 5 presents the results. Section 6 discusses policy implications and concludes.

2 Background

While both offsets and permits are used in California's cap-and-trade regime, certain qualities about offsets makes them susceptible to being treated differently from permits in California in general.

2.1 Cap-and-trade in California

California initiated its cap-and-trade system in 2013 with the goal of decreasing the state's GHG emissions to 1990 levels within seven years (AB-32, 2006).

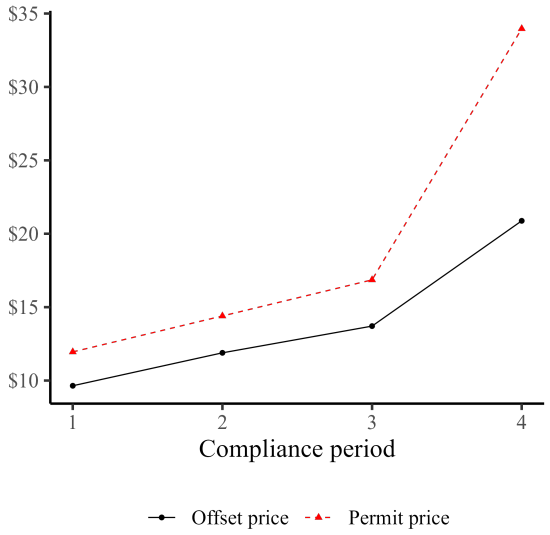


Figure 1: Price evolution of permits and offsets

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.

To prevent emissions leakage, the California Air Resources Board (CARB) strategically allocates permits to reduce compliance burdens on the firms most likely to relocate. Nearly all electricity generators receive full emission allocations, and energy-intensive, trade-exposed industries receive additional allowances. CARB sets an annual emissions cap and uses a systematic process to determine initial permit distributions without grandfathering. While allocation decisions are private, all compliance instrument surrenders (permits, offsets, and allocations) are tracked through the Compliance Instrument Tracking System Service (CITSS).

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

A firm becomes a regulated entity by emitting 25,000 or more metric tons of CO_2 in a year. Once regulated, the firm needs to surrender, by the end of the compliance period, compliance instruments equal to all of its emissions during the compliance period. A compliance period begins on the January 1 of its first year and ends on December 31 of its last year. For example, the second compliance period began on January 1, 2015 and ended on December 31, 2017, with final instrument surrenders completed in November 2018. 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 continue exceeding this threshold, creating de facto regulation persistence.

The cap-and-trade program allows regulated entities to cover a small share of their emissions using offsets. During the first three compliance periods (2013-2020), firms could meet up to 8 percent of their total compliance obligation by purchasing and surrendering offsets. Starting in the fourth compliance period, this limit decreased to 4 percent. At the same time, it became the rule that no more than 2 percent of a compliance obligation can be met using offsets that do not provide direct environmental benefits to the state (DEBS) of California. Figure 1 demonstrates the significant incentive to use offsets. The policy changes in the fourth compliance period decrease demand for offsets, but, as figure 1 shows, the price discount increased greatly.

2.2 Sources of the adoption cost

GHG offsets represent voluntary abatement generated by unregulated entities and purchased by regulated firms for compliance. A critical concern is "non-additionality"—whether credited emission reductions would have occurred anyway. Since regulated firms receive credit for these reductions, non-additional offsets effectively increase global emissions. Calculating true additionality requires knowledge of an unregulated entity's business-as-usual behavior, which is fundamentally unobservable, making offsets vulnerable to adverse selection.

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

Beyond reputational concerns, adoption costs include information and compliance expenses. These may be lessened in California where offsets are listed on centralized registries, and sales are brokered by a third party. However, dealing with intermediaries and ensuring a firm is compliant in light of the new instrument type may require a significant amount of employee time. Further, offsets face potential "invalidation" either by being determined non-additional ex post or by being destroyed in a natural event. In this case, offsets must be replaced with new compliance instruments (AB-32, 2006). Though invalidation events are uncommon, risk-averse firms may view them as an additional cost. Because larger firms

typically have more employees and, having more access to capital, are more likely to be risk-neutral, their adoption costs may comprise a larger share of public relations concerns relative to operational compliance costs. Table 1 summarizes the three primary sources of adoption costs.

Table 1: Sources of adoption cost in the compliance market

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

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

3 Data

3.1 Emissions, permits, and offsets

The main data for this paper comes from CARB's compliance reports.[1] These are annually-published tables which list the emissions compliance obligation, the amount of permits retired, and the amount of offsets retired for every regulated firm. In effect, they summarize each regulated firm's cap-and-trade history. These are published annually, but this paper uses the full period report which contains only full compliance period-level data.

There is 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. 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.[2] Each report contains aggregate data on the total

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

[2] Example: https://ww2.arb.ca.gov/our-work/programs/cap-and-trade-program/program-data/summary-market-transfers-report

number of allowances transferred between firms during the period. These reports do not include auctioned permits. Instead, they contain the prices for the allowances and offsets transferred from one firm to another. Each transfer is either a "priced transaction" or an "unpriced transaction." Priced transactions occur when one firm purchases an instrument from another in a single transaction while unpriced transactions consist mainly of transactions between corporate associates, where the instruments are being traded as part of a larger bundle.[3] Each transfer summary provides a weighted average, median, and standard deviations for allowances and offsets. A separate table published by CARB details the prices and quantities of permits bought at auction.[4] 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. This is because the compliance instruments are due by the end of the compliance period. I have repeated the analysis using the average price and the qualitative results are unchanged.

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. I have re-run the analysis using annual data instead and found qualitatively similar results.

To keep the data easily interpretable, I drop a small number of firms that acquire or get rid of facilities during a compliance period. 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. This means the firm becomes regulated in the last year of a compliance period, and, as a result, the sum of all compliance instruments surrendered during either compliance period does not match the total compliance period's obligation. Finally, because the rules around offsets and their price discount relative to permits changed dramatically in the fourth compliance period, I restrict my econometric analysis to the first three compliance periods.

[3] 96% of all transferred allowances were a transferred in a priced transaction in 2023. This percentage is typical of other years as well.

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

3.4 Stylized facts about offset use

Table 2: How prevalent are offsets?

CP	No. firms	Uses any	Uses max	Uses half	Share using max
1	195	84	52	0	61.9%
2	183	72	38	0	52.8%
3	179	107	83	0	77.6%
4	163	105	89	13	84.8%

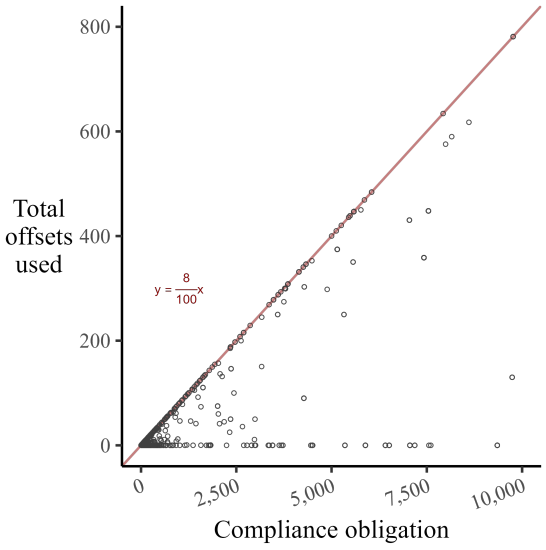
Note: CP stands for compliance period. There is an 8 percent cap in compliance periods 1-3 and a 4 percent cap in compliance period 4. The third column shows the number of firms using any offsets and the fourth column shows the number of firms using the maximum amount of offsets. The fifth column shows the number of firms using half of the offsets cap, presumably using only non-DEBS providing offsets. The sixth column shows the percentage of all firms using offsets that use the maximum amount. Here, using the maximum is defined as having an absolute value of the difference between offsets used and the maximum possible of less than 0.001. Using half is defined similarly.

In periods 1-3 about 84% of available offsets were used, and about 77% of offsets are used in period 4. Thus, there are offsets leftover at the end of each compliance period. Table 2 shows the prevalence of offset adoption across firms of offset use. This same data is shown graphically in figures 2 and 3.

Figure 2 shows that firms will frequently use the maximum amount of offsets possible, or will use no offsets at all. The solid red line in figure 2a represents the maximum amount of offsets possible at any emissions level, for the first three compliance periods. The solid blue line in figure 2b similarly represents the maximum amount of offsets possible at any emissions level during the fourth compliance period while the dashed orange line represents half of this level. The first three compliance periods are separated from the fourth compliance period because the rules around offset use changes in the fourth compliance period. Starting in the fourth compliance period, *no more than* 2 percent of a firm's obligation can be met using offsets that do not provide "direct environmental benefits to the state" (DEBS). Table 2 shows that there are firms both along this DEBS limit and along the total offsets limit.

Figure 3 allows for a closer examination of offset adoption patterns. It shows CDFs across the distribution of the share of allowable offsets used. Panel a shows the unconditional CDF and panel b shows the CDF conditional on some use of offsets. The CDF in figure 3a uses the exact same sample as figure 2, but makes clearer that there is a near-discrete jump in the share of offsets used from 0 to the upper bound. The solid red line is the CDF of the ratio of offsets to emissions for the first three compliance periods. It is clear that over half of firms in these compliance periods use no offsets. The nearly flat CDF between horizontal axis values 0 and 0.08 reflects the relative dearth of firms using a positive amount of offsets less than the

(a) Offset usage during the first three compliance periods



(b) Offset usage in the fourth compliance period

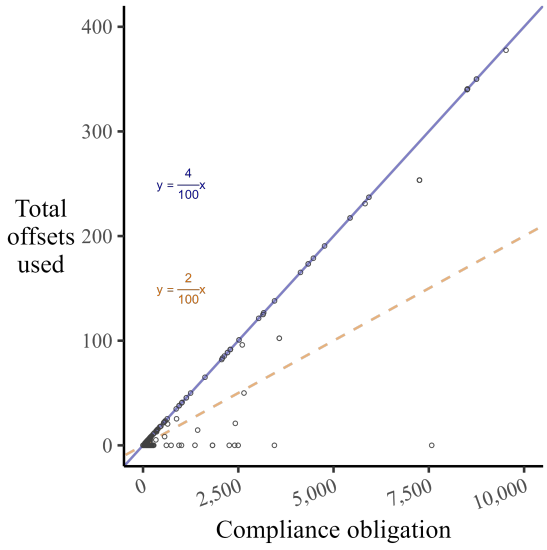


Figure 2: Relationship between offsets used and potential offsets allowed

These figures show the relationship between total offsets used and compliance obligation. The figures are limited to firms with 10 million tons of emissions for clarity. In both graphs, the solid ray is the maximum offsets for a given emissions level. The orange dashed line in 2b is half of this. Points on the solid red line in 2a indicate surrendered offsets equal to 8% of a firm's emissions—the maximum amount. Points on the solid blue line in 2b 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.

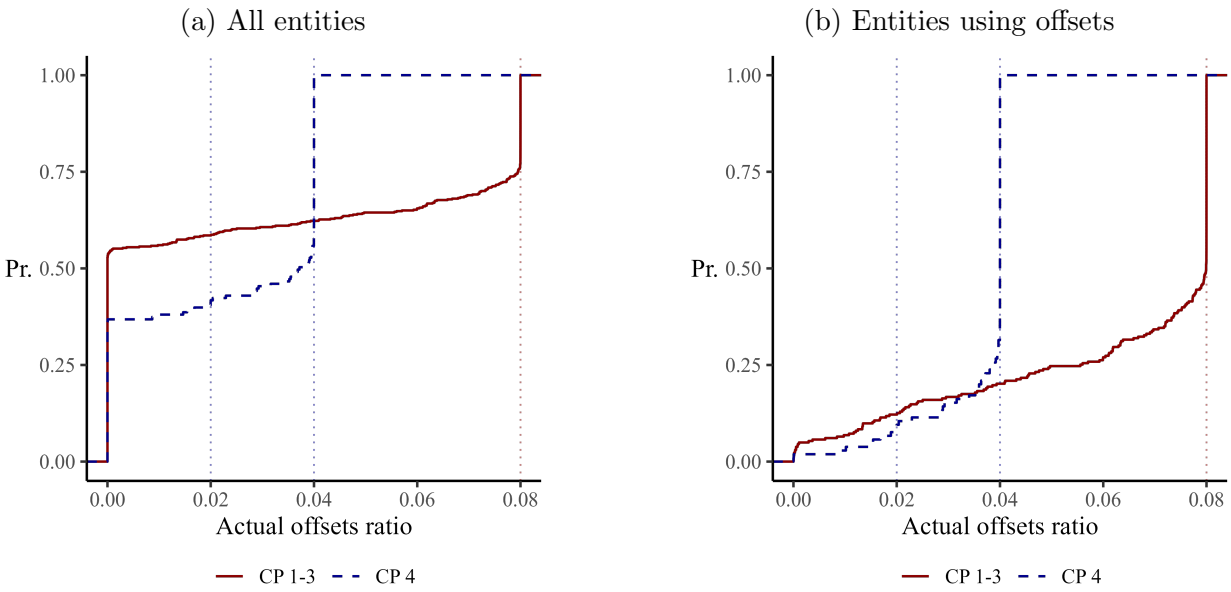


Figure 3: Cumulative density offset share

Note: These are the full compliance-period data. The tall vertical line of the CDFs at 0 in figure 3a 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 adoption cost related to locating DEBS-providing offsets. Figure 3b 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.

limit. The dashed blue line is defined similarly for the fourth compliance period. Over half of firms now use offsets (likely due to the higher cost savings), but we see the same pattern of a flat CDF with a discrete jump at the upper limit, 0.04. For clarity, figure 3 restricts the sample to all firms that use any offsets. Here, the horizontal lines are revealed to be slightly upward sloping, but they still show the pattern of discrete jumps at key points along the horizontal axis. Figures 2 and 3 show that a majority of firms lie on the extreme ends of the offset quantity spectrum.

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

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

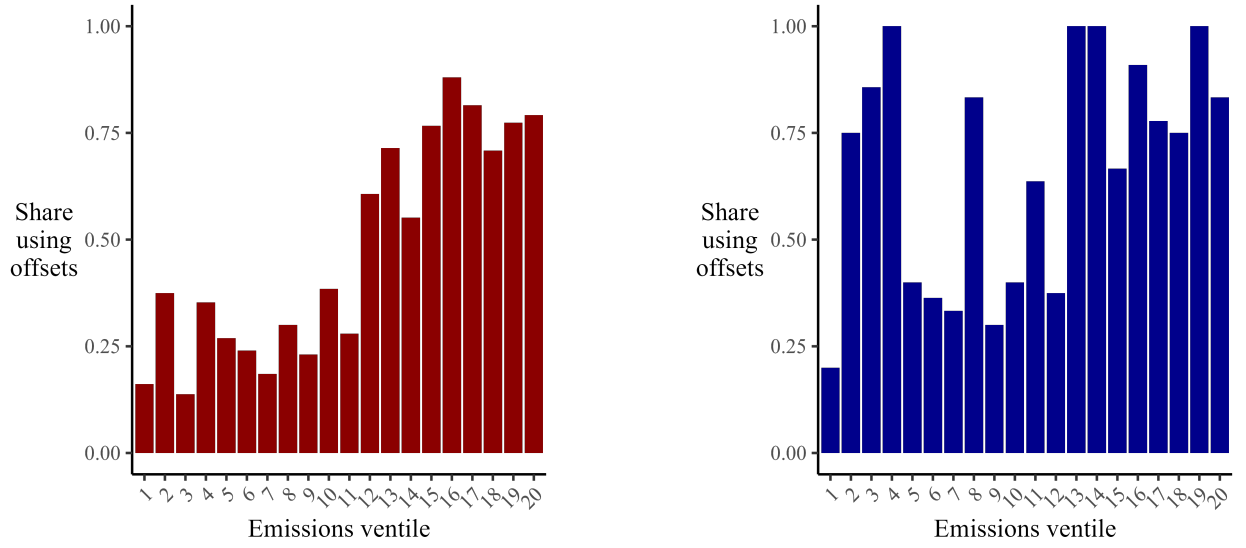


Figure 4: Larger firms are more likely to use offsets

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

Figure 4 shows the share of firms using offsets, by emissions level. It shows that the distribution of offset adoption is not uniform across emissions-level. Figure 4a shows relatively small offset adoption for firms with fewer emissions. Offset adoption is more common among high-emissions firms, with a majority of above-median emissions firms using offsets. Starting in the fourth compliance period, figure 4b shows a different overall distribution of offset adoption, but the pattern of higher-emitting firms being more likely to use offsets remains. Recall that the premium paid for permits increased greatly at the start of the fourth compliance period, more than doubling the incentive to adopt offsets. This is only descriptive, but the less pronounced skew suggests that the policy's favoring of large firms may be less dramatic under the new offsets policy.

Two important factors in a firm's offset decision changed at the start of the fourth

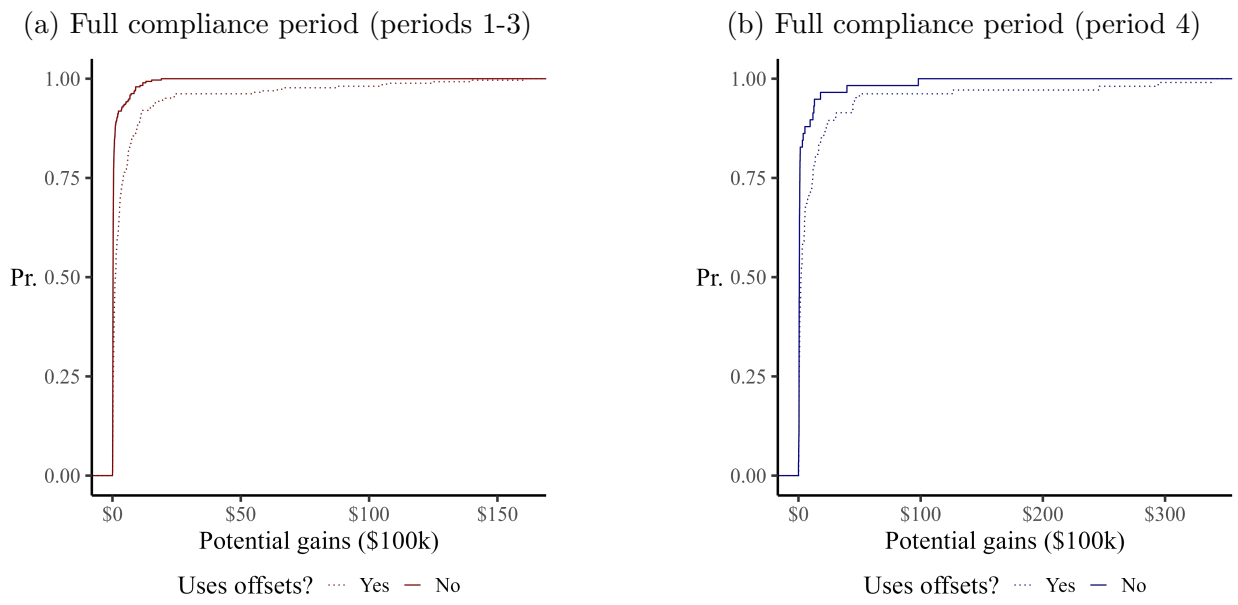


Figure 5: Potential gains from offset adoption

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

compliance period. First, the offset cap was halved from 8 to 4 percent. This limits the maximum benefits that a firm can gain by using offsets. On its own, this decrease in the offsets cap disincentivizes their use. The second factor was the the increase of the price discount of offsets. As figure 1 illustrates, the premium on permits increased from \$3.15 in the third compliance period to \$13.1 in the fourth. While the quantity restriction was cut in half, the substitution benefit per unit increased by over four times. So the fact that offset market participation increased in the fourth compliance period is consistent with the idea that firms must overcome an adoption cost in order to benefit from using offsets.

The proportional offset limit put in place by CARB incentivizes larger firms to use offsets. Because the offset price discount is common to all firms, but the number of offsets a firm can use is proportional to its emissions, the incentive to use offsets is increasing in firm emissions. This incentive structure leads to firms selecting into the offset market as revealed in figure 5. This shows that firms that do not use offsets stochastically dominate (to a first order) firms that do 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. In both figures, at all values on the horizontal axis, the probability that a firm who uses offsets has a weaker incentive to use them is smaller for the firms that

use offsets. This means that the firms observed using offsets are also the firms that stand the most to gain from doing so—firms respond to the market's incentive structure.

Table 3: Do entities continue to use offsets?

		Of those using offsets this CP	
CP	using offsets this CP	uses offsets every CP in future	uses offsets at least once in future
1	44.2%	47.8%	77.2%
2	41.4%	74.4%	74.4%
3	62.2%	100%	100%
4	65.7%	–	–

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.

Finally, table 3 shows the persistence of offset use by firms. This table shows the trend that offsets grow in popularity as the program continues, particularly starting in the third compliance period. It also highlights a feature that may be missed from the other tables and figures, and is not mentioned in any of the previous literature: firms that use offsets almost certainly continue using them in the future. This persistence could reflect two aspects of the adoption cost. First, it could reflect that the firms have learned how to navigate the offsets market, thus reducing future adoption costs. Second, the persistence could suggest that a public relations concern may drive the adoption cost. Once a firm is revealed to have used offsets, it will likely continue to use them to meet future compliance obligations.

4 Estimation strategy

This empirical strategy is intended to (i) describe the barrier to entry in the compliance market for offsets and (ii) highlight any heterogeneity in this barrier to entry across sector and firm size. 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.

A firm i in sector j uses offsets during compliance period t if and only if the cost-savings outweigh the fixed adoption cost. This adoption cost is the latent variable ρ^*_{ijt}. The empirical exercise links the latent variable to the known per-ton offset price discount p_t and maximum possible emissions coverage offset $\tilde{q}_{ijt}$. Because offsets are everywhere cheaper than permits in the data, p_t is always positive. The maximum possible emissions offset is the interaction of the firm's compliance period emissions with the maximum share of emissions that can be

met with offsets. So for 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 $p_t \tilde{q}_{ijt}$ is in units of dollars.

I make the assumption that each sector has a common adoption cost, but idiosyncratic factors affecting firm decisions to use offsets. For the sake of estimation, the coefficient on maximum possible savings from offset adoption, $p_t \tilde{q}_{ijt}$, is normalized to 1. Formally, the link is

$$\begin{aligned} \mathbb{1}_{ijt}^{O} &= \mathbb{1}\left\{p_t \tilde{q}_{ijt} \geq \rho_{ijt}^{*}\right\} \\ &= \mathbb{1}\left\{p_t \tilde{q}_{ijt} \geq \sum_{j=1}^{3} \rho_j \mathbb{1}_i^j + \epsilon_{ijt}\right\} \end{aligned} \tag{1}$$

where $\mathbb{1}$ is the indicator function, $\mathbb{1}_{ijt}^{O}$ indicates that firm i uses offsets in compliance period t, and $\mathbb{1}_i^j$ indicates that firm i is in sector j. The unknown parameters are ρ_j, the sector-level common adoption cost. The distribution of the unobservable error ϵ_{ijt} determines the kind of econometric model to be used. These terms are both measured in units of dollars.

Sector-specific adoption costs I estimate separate adoption costs for the three main sectors identified by CARB: energy, oil and gas, and other. The energy sector consists of cogeneration, electricity generation, and hydrogen plants while the oil and gas sector consists of oil and gas production sites and refineries—all other firms are in the "other" sector.[5] There are two reasons that sectors may have different adoption costs. First, some sectors—like energy—will necessarily have higher emissions; hence a stronger incentive to use offsets. Controlling for firm size omits this factor from confounding the adoption cost estimates, allowing me to focus on the second reason, which is novel to this work.

The second reason is that certain sectors may be more hesitant to use offsets because they have a "clean" reputation and are worried using offsets could affect this. This concern is due to the public's ideas about offsets, not necessarily whether the firm's offsets are revealed to be non-additional. Thus, it is not the case that the firm must intentionally purchase non-additional offsets to incur this cost, but that the lack of formal (or reliable) additionality measurements creates this. Knowing this, firms in more "client-facing" sectors such as education or transportation might have higher adoption costs compared to those either further upstream in the manufacturing process such as energy or those already labelled "dirty" such as oil and gas companies.

[5]Other includes transportation, warehousing, manufacturing, construction, and education. I run the same analysis, defining sectors manually using two-digit NAICS codes and the results are qualitatively the same as using the CARB-defined codes.

Probit model Assuming the error follows a normal distribution, with constant standard deviation σ, equation (1) can be estimated by probit methods. The estimating equation here is

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

The intercept sector-specific intercept $\frac{\rho_j}{\sigma}$ represents the baseline propensity for firms in sector j to use offsets. The coefficient $\frac{1}{\sigma}$ is akin to the "travel cost" coefficient in the recreation demand literature (Haab & McConnell, 2002). In this model, $\frac{(\rho_j)/\sigma}{1/\sigma} = \rho$ measures the lump-sum minimum cost-savings that will induce the average firm in sector j to enter the offsets market.

Hurdle model The raw data in figures 2 and 3 show that many firms use zero offsets. Thus, determining the quantity of offsets used should be modelled conditional on using any offsets at all. I model the demand for offsets as a two-step process where firms first decide whether to enter the offsets market, and then 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 adoption 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 (2) under the probit assumption. The second step, conditional on using offsets, models the quantity of offsets used as a function of the maximum possible offsets savings and the sector identifier. The complete hurdle model is defined in equation (3):

$$\begin{aligned} \mathbb{1}_{ijt}^{O} &= \Phi\left(\frac{1}{\sigma}p_t\tilde{q}_{ijt} - \sum_{j=1}^{3}\frac{\rho_j}{\sigma}\mathbb{1}_i^j\right) && \text{(first step)} \\ o_{ijt} &= \mathbb{1}_{ijt}^{O} \cdot \exp\left(\beta p_t\tilde{q}_{ijt} + \sum_{j=1}^{3}\alpha_j\mathbb{1}_i^j + \eta_{ijt}\right) && \text{(second step)} \end{aligned} \tag{3}$$

The interpretation of the coefficients in the first step is unchanged from the probit. The coefficients in the second step require some interpretation. The intercept α_j represents the average offset quantity for sector j *conditional on using any offsets*. 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. The second step has ann error term η_{ijt} that has a conditional mean of zero.

Binary quantile response Both the probit and hurdle models are "mean-based" regression models, meaning they estimate parameters which fit the data at the mean under the assumption that the error term has a zero conditional mean (Cameron & Trivedi, 2005). Naegele's (2018) analysis 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. Following her approach, I employ binary quantile responose (BQR) econometric methods inspired by Kordas (2006), which instead fit the conditional *quantile* (e.g. median) of the distribution of costs. This is a useful exercise especially if there are concerns about heteroskedasticity in the error terms.

The BQR estimation developed in Kordas (2006) was recently implemented in `Stata` by Alejo and Montes-Rojas (2025). The method maximizes the following objective function to estimate adoption cost parameters at the τ^{th} quantile:

$$\max_{\beta(\tau),\rho(\tau)} \sum_{i=1}^{N} \left[\mathbb{1}_{ijt}^{O} - (1-\tau)\right] \cdot K\left(\frac{\beta(\tau)p_t\tilde{q}_{ijt} - \sum_{j=1}^{3}\rho_j(\tau)\cdot\mathbb{1}_i^j}{h_n}\right) \tag{4}$$

where K is the Gaussian kernel and $h_n = \frac{0.9}{n^{1/5}}$ is a rule-of-thumb bandwidth (Alejo & Montes-Rojas, 2025). The coefficients are functions of τ because the objective function is separately optimized for every quantile. The coefficients are defined so that they correctly predict offset adoption for observations with the $\tau \cdot N$ highest errors. The interpretation fo the τ^{th} adoption cost, $\frac{\rho(\tau)}{\beta(\tau)}$ is then the fixed cost that $\tau \cdot 100$ percent of firms can cover—thus they will use offsets. As a positive error term means a firm is more likely to adopt offsets, we expect ex ante that the adoption cost will decrease as τ increases. That is, the adoption cost estimated at the median will be smaller than the adoption cost estimated at the 75^{th} quantile.

The BQR allows me to study heterogeneity in adoption 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 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. This might raise endogeneity concerns if we believe that a firm will relocate based on the attractiveness of offsets. If higher-emitting firms are also located in more lenient offset countries, there will be a concentration of the maximum cost savings in those countries and the adoption cost from mean-based regression models will be biased toward their adoption cost, which may be different from the median firm's adoption cost. Naegele (2018) does not explicitly mention this. Her data also show a more significant jump in offset usage from smaller to larger quantiles of maximum possible offset use (seen in figure 4a in Naegele (2018)). In

the California setting, all firms face an identical maximum offset proportion, so maximum possible cost savings vary only on the price and emissions dimension. Figure 4 shows that offset use is increasing in emissions, but the jump in adoption is not as drastic as in the EU.

5 Results

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

	(1)	(2)	(3)
Energy	-59.929	220.6	-52.808
	(104.72)	(125.18)	(149.68)
Oil and gas	267.86	426.85*	581.84*
	(148.73)	(190.03)	(227.45)
Other	562.35***	835.68***	988.29***
	(132.81)	(227.76)	(278.17)
Energy previously used		-697.26*	-238.34
		(306.77)	(227.59)
Oil and gas previously used		-491.01	-1,620*
		(296.96)	(799.65)
Other previously used		-82.183	-776.8*
		(185.47)	(365.58)
Obs.	556	556	556
Previously used control?	No	Yes	Yes
Previously used interaction?	No	No	Yes

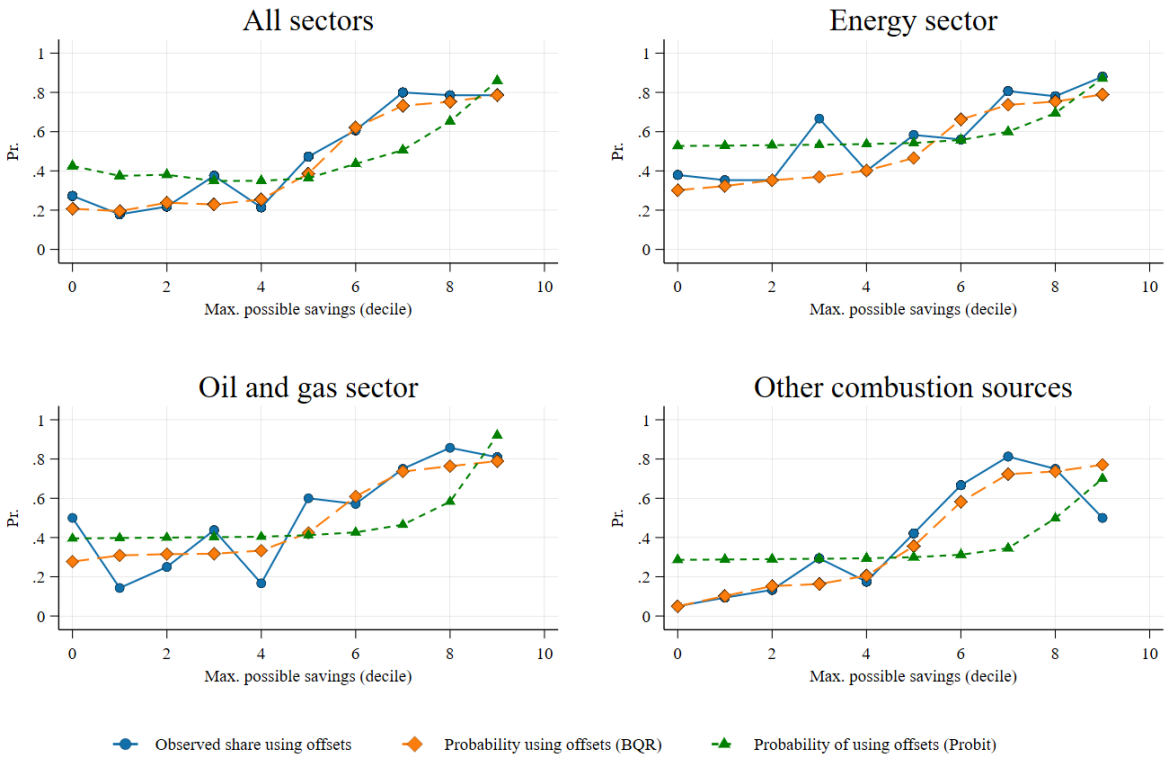
Robust standard errors in parentheses

These are sector-level mean adoption costs estimated by probit. They are calculated by dividing the sector indicator coefficients by the price coefficient. There is no excluded sector as there is no intercept. The equation in column (1) includes only cost savings and sector indicators as explanator variables. The equation in column (2) includes cost savings, sector indicators, and a dummy variable indicating whether the firm had used offsets before. The equation in column (3) includes cost savings, sector indicators, the previously used indicator, and an interaction between previously used and sector.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 6 compares the predicted rate of offset adoption for the two models compared to the observed rate—Kordas (2006) recommends this comparison for the BQR predicted probabilities to check model fit.. Note that the maximum cost savings are grouped in bins for visual clarity, but these bins *are not* related to the error quantiles (τ) in equation 4.

Figure 6: Predicted v. observed offset adoption



Note: This figure compares observed offset adoption rates with the average predicted offset-use probability from probit and quantile estimation binned across maximum savings. Note that the BQR estimates are not binned To get the observed adoption rate within a maximum savings bin, I split maximum savings into 20 bins and calculated the share of firms in that bin that use offsets, for each sector. To get the average predicted offset-use probability, I calculate the mean predicted probability of using offsets for firms within each maximum savings bin, for each sector. The predicted probabilities are the average of the kernel function used in equation 4 over $\tau \in \{1, 2, \ldots, 19\}$ for the firms within each maximum savings bin. In both the probit and BQR estimation, the only explanatory variables are cost savings and sector indicators (the same as in the first column of table 4). The coefficient on cost savings is positive (lower panel of table 7), so the probability of offset adoption is non-decreasing in cost savings within sectors. The upper left panel shows the combined predicted probability for all sectors. It shows some decreasing probability in cost savings that is a result of the different sector composition within each decile—this is because the sector makeup can vary across ventile, something the BQR nonparametrically controls for by estimating parameters at a specific quantile.

The blue dots show the raw data's observed offset adoption rates within each cost savings quantile while the orange diamonds (BQR) and green triangles (probit) show the average predicted probability (predicted share of offset adopters) within each quantile. Given the hypothesis that firms are more likely to adopt offset use if their cost savings outweighs their adoption cost, it makes sense to see a higher adoption rate for larger cost savings bins.

Table 4 shows the estimated average adoption cost using the probit model. There are three specifications included in this table. The leftmost specification includes only sector indicators and maximum cost savings on the righthand side of equation (2). This specification returns the smallest adoption cost estimates assuming as it combines firms that have used offsets before with those that have not.

The middle specification includes sector indicators, maximum cost savings, and an indicator variable equal to 1 if a firm has used offsets in a previous compliance period. The adoption cost estimates for all sectors increases with this inclusion, but only energy faces a statistically significant adoption cost for firms that have used offsets before. This sector's cost is negative, suggesting that firms that have used offsets before incur an additional cost by not using them again. That the initial adoption cost estimates grow by including the previous offset user indicator suggests suggests that firms incur some additional cost by not continuing to use offsets once they've entered the market, either as sunk or as reputational costs. The rightmost column contains adoption cost estimates for that allow for different intercepts for each sector and for different slope coefficients on maximum cost savings. Firms in the oil/gas and other sectors have the highest adoption cost under this specification and they both have large, negative, and significant adoption costs for the firms that have used offsets before. The negative adoption costs for firms returning to the offsets market helps explain the persistence seen in table 3 and are discussed more in section 6.

The estimated adoption costs are presented in table 5 where each model specification includes sector indicator variables and an indicator for firms that have used offsets in the past. The probit estimated adoption costs are identical to the middle column of table 4. The probit specification estimates significantly larger adoption costs compared to median and 75^{th} percentile estimates. Recall that the adoption cost for $\tau = 0.25$ is the cost that will induce only the 25 percent of observations that are the *most likely* to adopt offsets conditional on the explanatory variables under the true distribution of error terms. As such, the adoption cost for these firms are extremely large at around \$2,054, \$2594, and \$3,817 for the energy, oil/gas, and other sectors, respectively—though these estimates are statistically different from zero at all conventional levels of significance. This could suggest that the determinants of the highest adoption costs are more idiosyncratic than the model allows.

At the median, the energy and oil/gas adoption costs match each other very closely both

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

	Probit		$\tau = 0.25$		$\tau = 0.5$		$\tau = 0.75$	
	First time	Previously used	First time	Previously used	First time	Previously used	First time	Previously used
Energy	220.6	-697.26*	2,054	-683.1	48.943***	17.224	-5.3965	-29.722
	(125.18)	(306.77)	(1.0e+05)	(65072)	(10.069)	(11.689)	(11.944)	(29.89)
Oil and gas	426.85*	-491.01	2,594	-142.56	49.551***	17.831	2.8754	-21.45
	(190.03)	(296.96)	(1.3e+05)	(13854)	(10.34)	(12.428)	(12.422)	(25.134)
Other	835.68***	-82.183	3,817	1,080	63.605***	31.886***	39.875*	15.549*
	(227.76)	(185.47)	(1.9e+05)	(1.0e+05)	(9.2086)	(7.1452)	(17.121)	(7.2719)
Obs.	556	556	556	556	556	556	556	556

Standard errors in parentheses

Probit standard errors are heteroskedasticity robust. Quantile point estimates are estimated separately for each quantile and standard errors are recovered through bootstrap with 500 repititions. All models include cost savings, sector indicators, and an indicator variable indicating that a firm has used offsets in the past—the same explanatory variables used in column (2) of table 4.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

for first-time offset use and for incumbents to the offset market. The adoption cost that will induce about half of the energy firms to adopt offsets is about $48,943 while the same for oil and gas firms is about $49,551. Half of the firms in all other sectors will use offsets when the adoption cost is around $63,605 while half will not. Half of the firms in both the energy and the oil and gas sectors which have used offsets before face an adoption cost that is statistically insignificant, while firms in the other sector that have used offsets before still face an adoption cost of about $31,886. An adoption cost of $39,875 will bring 75 percent of firms in the other sectors to use offsets for the first time, meaning that only the 25 percent of firms with the *lowest* adoption cost in this sector face a cost higher than $39,875. An adoption cost of $15,549 will allow 75 percent of firms in these sectors to participate in the market, conditional on having used offsets before.

The quantile results demonstrate that the typical firm faces substantially lower barriers than probit estimates suggest—median adoption costs are approximately one-tenth of the mean estimates. This divergence reveals that a small number of firms have exceptionally high adoption costs (potentially due to unique reputational concerns) that drive up the average, while the majority of firms face more modest but still economically significant adoption costs.

Table 6: Grouped adoption cost by sector ($1k)

	Probit	$\tau = 0.25$	$\tau = 0.5$	$\tau = 0.75$
Energy	-59.929	902.34	38.948***	-4.5357
	(104.72)	(87588)	(8.4725)	(8.824)
Oil and gas	267.86	1,720	47.125***	-6.5651
	(148.73)	(1.7e+05)	(12.214)	(12.498)
Other	562.35***	1,468	54.475***	31.642*
	(132.81)	(1.4e+05)	(6.9917)	(12.326)
Obs.	556	556	556	556

Standard errors in parentheses

Probit standard errors are heteroskedasticity robust. Quantile point estimates are estimated separately for each quantile and standard errors are recovered through bootstrap with 500 repititions. The probit model is the same as column (1) of table 4. The quantile estimates are estimated using only cost savings and sector indicators—there is no distinction for previously using offsets.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6 shows the adoption costs by sector for the specification only including sector indicators and maximum cost savings as explantory variables. These costs are notably lower than in the specifications which account for previously adopting offset, but the general qualitative conclusion is the same. The firms with errors in the 25^{th} quantile have very high,

but insignificant adoption costs while the median range between $38,000 and $54,000. The 75^{th} quantile estimates of adoption cost in the other sector are about $10,000 lower when not accounting for previous offset adoption.

The offset quantity predictions are presented in the top panel of table 7. Conditional on entering the offsets market, firm sector and whether a firm has used offsets before plays no role in the quantity decision. The only statistically significant predictors in the second stage of the hurdle model are the constant and the maximum cost savings. This is consistent with the idea that there is a fixed adoption cost that firms need to overcome, but that any firm that uses offsets treats them the same.

The first stage of the hurdle model is identical to the probit model, and the predicted coefficients are presented in the lower panel of table 7. This stage highlights that all of the heterogeneity in the use of offsets for California's regulated firms takes place at the selection stage. That is, whether a firm uses offsets depends strongly on its sector and whether it has used offsets before.

6 Discussion

This paper provides empirical evidence of substantial adoption costs in California's offsets market, with important implications for both policy design and our understanding of how transaction costs shape environmental markets. The findings reveal important aspects for policy-makers designing offsets markets. First, there is a cost to adopting a new compliance instrument that varies across broadly defined sectors. Second, the decision to use offsets in the current period affects the decision in all future periods; there are dynamic implications of past adoption decisions.

Table 3 demonstrates that firms adopting offsets in one compliance period almost always continue using them going forward: 77% of first-period adopters use offsets in all future periods while 100% of third-period adopters use offsets in the fourth compliance period. The econometric results reinforce this pattern by revealing that all sectors have lower—sometimes negative—adoption costs simply for having used offsets previously. One interpretation of this persistence result is that adoption costs genuinely function as one-time investments in creating internal compliance processes and establishing vendor relationships that can be amortized over multiple periods. Another interpretation is that it indicates reputational concerns as a real part of the adoption cost. That is, once a firm publicly uses offsets, the marginal reputational cost of continued use approaches zero as the firm has already "revealed" its position.

The binary quantile regression results shed additional light on the distribution of adoption

Table 7: Hurdle model

	(1) Offsets (1k MTe)	(2) Offsets (1k MTe)
Second stage		
Constant	2.9304***	2.8738***
	(.16644)	(.19468)
Max. benefit ($1k)	.00046***	.00046***
	(6.2e-05)	(6.3e-05)
Previously used		.13259
		(.23705)
Oil and Gas	.28935	.30236
	(.32666)	(.32729)
Other Combustion Source	-.0267	-.02329
	(.266)	(.26591)
First stage		
Constant	.0609	-.18299
	(.09709)	(.10678)
Max. benefit ($1k)	.00102***	.00083***
	(.00019)	(.00019)
Previously used		.76136***
		(.13764)
Oil and Gas	-.33309*	-.17108
	(.16822)	(.17401)
Other Combustion Source	-.63234***	-.5102***
	(.12361)	(.12757)
$\ln(\hat{\sigma})$		
Constant	.6303***	.62971***
	(.04369)	(.04369)
Observations	556	556

Standard errors in parentheses

The excluded sector in the first step is the energy sector. All models include sector intercepts to account for different average offset adoption rates across sectors.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

cost magnitudes. While probit estimates suggest average adoption costs exceeding half a million dollars for oil and gas firms and nearing one million dollars for other sectors, median estimates are substantially lower: approximately $50,000 and $64,000, respectively. Thus, if policy tries to correct for transaction costs using mean-based estimation methods alone, there could be a significant misallocation of funds. For the majority of firms, adoption costs remain economically significant but more modest than probit alone implies.

Firms in the energy and oil/gas sectors face smaller adoption costs compared to all other firms. This ordering provides suggestive evidence about the composition of adoption costs. Energy firms operate in wholesale markets with limited consumer interaction, potentially insulating them from public relations concerns about offset use. Oil and gas firms may face lower reputational costs from offset adoption because they are already viewed as "dirty" or as high emitters—the incremental reputational penalty is small. In contrast, entities like universities and food manufacturers may be either more reliant on brand identity or have "clean" public images that could be jeopardized by association with potentially non-additional offsets, even if they make all efforts at ensuring the quality of their offsets. The public relations component of adoption costs is subtly ironic: firms most concerned about their environmental reputation may be the least willing to use these (carefully chosen) cost-saving instruments which bring about previously unincentivized environmental benefits. Thus, non-additionality concerns about offsets create real economic costs by deterring participation. Policy interventions that enhance offset credibility—such as more rigorous additionality verification or real-time monitoring systems—could both encourage market participation by reducing these barriers and improve the quality of offsets that are brought to market. This paper puts a dollar target on the cost that these interventions should aim to cover.

Back of the envelope calculations suggest that, on average, the total adoption cost paid across the first three compliance periods is at least $51,092,650. This number comes from multiplying the number of firms using offsets in each sector (131, 50, and 81 in energy, oil/gas, and other) by the average adoption cost estimates in column 1 of table 4. There are 294 firms that do not use offsets during this time, and the sum of the total cost savings they would have achieved if using offsets had zero trade frictions is$172,457,300; this is the sum of each individual firm's potential cost savings. Using the same adoption cost estimates, if each of the firms that did not use offsets (82, 41, and 171 in energy, oil/gas, and other) was given a subsidy equal to the adoption cost, it would require $102,229,930 in subsidies. The lost savings from these firms not participating in the offsets market outweighs this subsidy by around $70,000 suggesting significant deadweight loss is incurred by the existence of the adoption cost.

Several limitations merit attention. In particular, the analysis cannot separately identify

the relative contributions of information costs, compliance costs, risk aversion, and reputational concerns to total adoption costs. Surveys or interviews with regulated entities could provide valuable qualitative evidence about which factors matter most, informing targeted interventions. Second, the unavailability of permit allocation data prevents full analysis of how firms' net positions (short vs. long) affect adoption decisions, though Naegele's (2018) work suggests this may not be a primary driver of transaction costs. Third, the analysis treats adoption costs as time-invariant within sectors, though institutional learning or changing social norms may have shifted these over the study period. Finally, this paper cannot assess whether reduced offset adoption represents a welfare loss or gain. If adoption costs primarily reflect firms conducting due diligence on additionality—essentially screening out non-additional offsets—then these costs serve a valuable verification function. The welfare implications depend on the relationship between adoption costs and offset quality, a topic for future research.

Offsets provide access to lower cost abatement, but concerns over non-additionality have dampened their global acceptance. The seminal models of offset markets assume negligible transaction costs (Aldy & Halem, 2024; Bento et al., 2015; Van Benthem & Kerr, 2013). Future work may benefit from revisiting the theoretical models and accounting for a transactions cost. Specifically, Bento et al. (2015) states that any supplier who was going to provide ecosystem services in absence of a payment will not be affected by the price of the offset. That is, changing the price of offsets will not affect this agent's decision to sell their non-additional offsets. But it can be shown that when firms face a transaction cost to supply offsets, there is a barrier that prevents participation from this group—that is, changing price *can* affect additionality. My and Naegele's (2018) findings point out that there are transaction costs in these markets, perhaps motivating a need to re-visit this literature.

These findings contribute to broader debates about transaction costs in environmental markets. Coase's (1960) foundational insight that well-defined property rights allow for efficient allocation depends critically on low transaction costs—an assumption often violated in practice. This paper demonstrates that even in sophisticated, well-designed markets, frictions do exist and should be considered by policy-makers. Taken together with Naegele (2018), this problem may be particularly salient in cap-and-trade regimes with offsets.

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 transition*. https://www.cleanenergywire.org/factsheets/understanding-european-unions-emissions-trading-system#:~:text=The%20EU%20ETS%20is%20the%20oldest%20and,report%20by%20the%20International%20Carbon%20Action%20Partnership.

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

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

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

Bento, A. M., Kanbur, R., & Leard, B. (2015). Designing efficient markets for carbon offsets with distributional constraints. *Journal of Environmental Economics and Management*, *70*, 51–71. https://doi.org/10.1016/j.jeem.2014.10.003

Buchanan, J. M., & Tullock, G. (1975). Polluters' profits and political response: Direct controls versus taxes. *American Economic Review*, *65*(1), 139–147. https://www.jstor.org/stable/1806402

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

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

Coase, R. H. (1960). The problem of social cost. *The Journal of Law and Economics*, *3*, 1–44. https://doi.org/10.1086/466560

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

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

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

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

Van Benthem, A., & Kerr, S. (2013). Scale and transfers in international emissions offset programs. *Journal of Public Economics*, *107*, 31–46. https://doi.org/10.1016/j.jpubeco.2013.08.004

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