

Cap-and-Trade Emissions Regulation: A Strategic Analysis

Frank Cheng¹*, Yagil Engel² and Michael P. Wellman¹

¹University of Michigan, Ann Arbor

²Ford Analytics

fcheng@umich.edu, yengel@ford.com, wellman@umich.edu

Abstract

Cap-and-trade schemes are designed to achieve target levels of regulated emissions in a socially efficient manner. These schemes work by issuing regulatory credits and allowing firms to buy and sell them according to their relative compliance costs. Analyzing the efficacy of such schemes in concentrated industries is complicated by the strategic interactions among firms producing heterogeneous products. We tackle this complexity via an agent-based microeconomic model of the US market for personal vehicles. We calculate Nash equilibria among credits-trading strategies in a variety of scenarios and regulatory models. We find that while cap-and-trade results improves efficiency overall, consumers bear a disproportionate share of regulation cost, as firms use credit trading to segment the vehicle market. Credits trading volume decreases when firms behave more strategically, which weakens the segmentation effect.

1 Introduction

In 1992, the UN established cap-and-trade as the standard regulatory approach to achieve emissions reduction of greenhouse gases (GHG) [Boswall and Lee, 2002]. This approach allocates pollution rights to large-scale polluters according to a desired emissions cap. Polluters are then permitted to trade these rights, thereby achieving a more efficient economic outcome. Several studies commissioned at the time, as well as numerous subsequent studies, found that this approach resulted in great cost savings as compared to a simple hard-cap approach that does not allow the trading of pollution rights [Ogden, 1966; Perl and Dunbar, 1982].

One domain under the purview of GHG regulation is the vehicle market. The auto market in the US alone represents hundreds of \$billions per year in transactions, and environmental regulation has a significant impact on the vehicles produced and sold. A technical challenge in analyzing this market is that vehicles are differentiated in the eyes of consumers, and the oligopolistic nature of the industry means that each firm has individual pricing power. The demand for each

vehicle line depends not only on its own price, but on the price of every other vehicle. In general these prices are set by different firms, resulting in profit functions that are strategically entwined. This is in contrast to the heavily studied power industry, where the product (electricity) is homogeneous. There, the typical scenario has prices that are fixed by an independent party that attempts to maximize social welfare [Limpaitoon *et al.*, 2011]. In the vehicle market, prices are determined as a result of a game that firms play with each other. This complicates the determination of the effects of cap-and-trade, both on overall efficiency and the distribution of costs among firms and consumers.

The innovation in our approach is that we cast the cost of pollution to firms as lost opportunities to increase profit due to regulation. We model firms in a vehicle market with many vehicle types, tied together by consumer demand, modeled using a nested-logit form combining a comprehensive set of vehicle features. This demand model is the product of industrial research, borne from the need to make accurate predictions rather than to build analytical models. We compute strategic equilibria among firms in the vehicle market through *iterative best-response* (IBR), which is guaranteed to converge for several of the models that we compare. Because our model allows computation of strategic dynamics explicitly, we are able to explore counterfactual scenarios such as alternative industrial organization and different regulation setups.

The market for trading pollution rights (or *credits*) introduces another level of strategic behavior for firms, as they consider their utilization of credits to shape their vehicle production, along with their bidding behavior in the market for credits. We model the credits market as an iterative double auction mechanism. Equilibrium analysis or even best-response calculation is not apparently feasible for the credits-market behavior alone, let alone considering the coupling of credits allocation to strategic behavior in the vehicle market. We therefore explore heuristic strategies for credit trading, and employ *empirical game-theoretic analysis* (EGTA), a simulation-based game reasoning framework, to derive equilibrium behavior in the heuristic strategy space.

Our main contributions are twofold. First, we construct and solve novel models for studying emissions trading in the vehicle market. Second, we produce economic insights for this domain, most notably the finding that the welfare gained by moving from hard-cap to cap-and-trade is disproportion-

*Contact Author

ately accrued by firms. In aggregate, firm profits under cap-and-trade approach that of the unregulated case, thanks to market segmentation opportunities created by emissions regulation. In contrast, consumers absorb the brunt of costs through higher prices for polluting vehicles. This effect is more pronounced with a less fragmented industrial organization. Strategic credit trading reduces this effect slightly, as firms lose some credit trading volume to an inability to coordinate trades, and thus are not able to segment as efficiently.

1.1 Relation to Prior Work

The standard tool for analyzing cap-and-trade emissions regulation is the marginal abatement cost (MAC) curve [Goulder, 2013]. At competitive equilibrium, the marginal cost for abating one more unit of GHG pollution is the same for all firms—otherwise, firms could gain from trading the right to pollute. This facilitates analysis of cap-and-trade in the electric power industry [Limpitoo *et al.*, 2011; Carlson *et al.*, 2000; Benz and Trück, 2009], where firms lack pricing power and therefore marginal costs directly determine equilibrium. Prior studies on vehicles [Karplus *et al.*, 2010; Kesicki, 2012] likewise apply MAC analysis in a non-strategic manner, by focusing on the cost of improving vehicle efficiency. In our setting, the more relevant abatement cost is foregone opportunity to sell profitable high-polluting vehicles, which calls for a more explicitly strategic analysis.

Previous work has followed different approaches to estimate MAC curves for vehicle markets. Firm behavior has been modeled as a fixed and non-optimizing [Fullerton and Gan, 2005; Mannering and Winston, 1985] as well as attempting to maximize fuel economy for consumers [Rubin *et al.*, 2009]. Our approach models firms as strategic profit maximizers, and reflects consumer utility in the demand function.

Studies using simple analytical models have examined the effect of industrial organization, specifically oligopoly, on emissions regulation [Sartzetakis, 1997; Malueg, 1990]. These models suggest that oligopolies are detrimental to overall welfare outcomes under cap-and-trade regulation. We duplicate these results under a more complex model, and in addition find that consumers bear the brunt of this welfare loss.

Strategic credit trading has been studied in experimental laboratory settings as well as through stylized analytical models [Westskog, 1996; Muller *et al.*, 2002; Godby, 2002]. In general, prior studies have found that market power in the emissions credit market leads to a decrease in efficiency. We see the same in our study, with total welfare falling when we allow strategic credit trading, but we find in addition that consumers benefit while firms make less profit.

This work adds to a growing library of AI research that seeks to uncover economic insights using tools from computational game theory [Reddy and Veloso, 2011; Jordan *et al.*, 2007; Li and Das, 2016; Wah and Wellman, 2013].

1.2 Background

The US government regulates auto emissions primarily through two programs. The *CAFE* program, administered by the National Highway Traffic Safety Administration, mandates corporate average fuel economy standards. The second program, maintained by the Environmental Protection

Agency (*EPA*), directly regulates GHG emissions. We focus on the EPA regulations, which broadly operate as follows:

1. Each vehicle is assigned a rating: negative for polluting, positive for non-polluting.
2. Each firm’s annual *credit balance*, defined as sum of ratings weighted by sales volume, must exceed zero.
3. Positive balances may be carried forward to the next year.
4. Credits may be traded between firms within each year.

Due to a combination of technological, consumer-preference, and other economic factors, non-polluting vehicles tend to be much less profitable than polluting vehicles like SUVs and trucks. Acquiring credits therefore allows a firm to produce a larger fraction of more profitable vehicles.

1.3 Overview of Models and Assumptions

Cap-and-trade is implemented by a market in compliance credits, secondary to the auto market itself. We are interested in understanding the dynamics of the credits market, and in particular to examine possible strategic behavior by automakers, and how the special dynamics of this secondary market affect the auto market.

For this purpose, we model four different environments:

1. *Unregulated*. This represents a world where the auto market operates without constraint on emissions.
2. *Hard-cap*. This serves as a baseline model, where each firm must satisfy EPA requirements individually, with no trading of compliance credits.
3. *Credits trading with price takers*. In this model we derive a single clearing price for emission credits, under competitive (Walrasian) assumptions. The assumption is similar in spirit to the MAC models which dominate the literature.
4. *Trading with strategic bidders*. This model includes a concrete trading mechanism for compliance credits, where firms interact through a bidding process, iterating between actions in the credits market and decisions about vehicle production. Including strategic bidding captures issues of market power and dynamic information reflective of cap-and-trade in practice.

In each setting the solution concept is Nash equilibrium with respect to the decision variables in that setting.

We assume *complete information* in the vehicle market game; that is, each firm knows the set of vehicles and costs of its competitors, and also has an accurate consumer demand model. Though not exactly realistic, we believe this provides a reasonable approximation. Firms can make an educated guess of their competitors’ costs using their own costs, and by observing equilibrium prices over time. Furthermore, consumer demand models can be built using historic transaction data available from commercial sources.

Another assumption we make is that the cost of producing vehicles, the consumer demand for vehicles, and the way regulators assign credit balances do not change over time. This assumption is not made because the methods in this paper

rely on it; rather, we do this mainly for convenience since we do not have the information required to forecast these quantities. If forecasts were provided, the analysis could be updated without any model changes.

Finally, we assume that firms do not price their vehicles below cost. This assumption is salient for non-polluting vehicles whose production generate emissions credits that could be sold. We make this assumption mainly as a reasonable baseline expectation for firm behavior. Our models and methodology do not depend on this assumption, and only the empirical results are affected by it.

2 A Model of Auto Emissions Trading

2.1 Notation

Consider a set of m **products**, denoted $V = \{v_1, \dots, v_m\}$. These products are sold in each of T time periods. Each product is produced by a single **firm** from set $F = \{f_1, \dots, f_n\}$ where $|F| = n$. The firm producing a particular product is given by function $f : V \rightarrow F$. The products produced by firm i are denoted $V_i \subset V$. The set of indices corresponding to V_i are $I_i \subset \{1, \dots, m\}$. Product **prices**, which are set by their producers, are given by vector $\Psi = \{\psi_1, \dots, \psi_m\} \in \mathbb{R}^m$. The **demand function** $\phi : \Psi \in \mathbb{R}^m \rightarrow \mathbb{R}^m$ returns the quantity of each product that is sold to a mass of consumers given Ψ . **Costs**, which are fixed and given per marginal unit of production, are denoted $\Sigma = \{\sigma_1, \dots, \sigma_m\}$. **Regulation credits** are assigned according to a function $\kappa : \phi(\Psi) \in \mathbb{R}^m \rightarrow \mathbb{R}^n$. Roughly speaking, these are non-negative for production mixes that achieve a regulatory goal and negative otherwise.

For convenience, we use V_j to denote a vector selecting V 's elements with indices I_j . V_{-j} denotes a vector selecting elements of V with indices from the complement of I_j . It is understood that the union operation, \cup , preserves indices such that, for example, $\Sigma_{-j} \cup \Sigma_j = \Sigma$. To select a particular element of vector V with index i , we use $V_{(i)}$. We use superscripts to denote an instantiation of a vector. These instantiations can be of functions or vectors. So for period t , Ψ_j^t is the vector of product prices actually set by firm j . Meanwhile, κ_j^t is equivalent to $\kappa(\phi(\Psi_j^t))$. This is the actual amount of regulation credits assigned to firm j in time t .

To denote elements of a function that are given, we use the | symbol in the following way: $F(x) | y, G = x + G(y)$ means that x is a variable, G is a given function, and y is fixed value. It should be clear from context whether functions or constants are being given.

Finally, let θ_j^t denote **credits retained** without being sold (by firm j). These credits will roll over into the next period, $t + 1$, and must be non-negative because deficits cannot be rolled over. And let λ_j^t denote **credit obligations** held by firm j at time t . Negative means obligations owed (sold) while positive means obligations held (purchased).

2.2 The Vehicle Market Model

We model the vehicle market as a Bertrand pricing game, for n multi-product firms with differentiated products. Firm j at time t strategizes over the prices it controls (each element of Ψ_j^t). The profit for firm j made by selling cars at time t

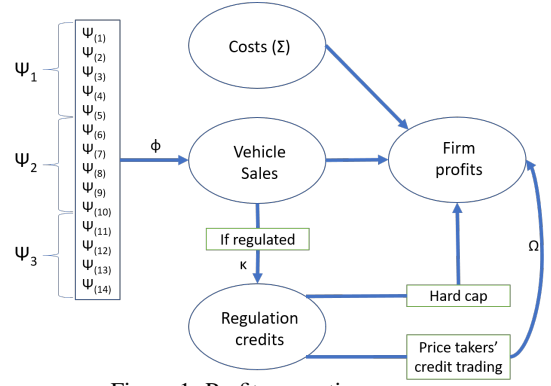


Figure 1: Profit generation process

is fully determined given a set of vehicle prices $\{\cup \Psi^t, t \in \{1, \dots, T\}\}$, as follows:

$$\Pi_j^{cars}(\Psi^1, \dots, \Psi^T) | \Sigma_j, \phi = \sum_{t=1}^T [\Psi_j^t - \Sigma_j]' \phi_j^t. \quad (1)$$

The demand function ϕ is a consumer choice model with (nested) logit demand, loosely based on real industry data. Naturally, vehicles in this model are **gross substitutes** (increasing the price of a product can only increase the demand for other products). We assume throughout that ϕ is known to all firms.

2.3 Unregulated Vehicle Market

In the unregulated vehicle market, firms set prices in an attempt to maximize Π_j^{cars} . In Equation 1, every component of firm payoffs is constant over time. This means that firms treat each period independently and equivalently. Their decision on price is independent of T . Call this constant decision Ψ^0 . Then the payoff function can be reduced to

$$\Pi_j^U(\Psi^0) | \Sigma_j, \phi = T \times [\Psi_j^0 - \Sigma_j]' \phi_j^0.$$

We can now define a Nash equilibrium for the pricing game in the unregulated market.

Definition 1 (Unregulated price equilibrium). *A set of prices Ψ^* is an unregulated price equilibrium if, for all $j \in \{1, \dots, n\}$:*

$$\Pi_j^U(\Psi^*) | \Sigma_j, \phi \geq \Pi_j^U(\Psi_j \cup \Psi_{-j}^*) | \Sigma_j, \phi, \forall \Psi_j \in \mathbb{R}^{|I_j|}.$$

Our simulations find Nash Equilibrium in this market using IBR - simulating each automaker's optimal response to the current prices of other automakers. In general games IBR might not converge and cycle forever; but since the game model we employ is a Bertrand pricing game with logit demand model and linear cost function, we can show that IBR converges:

Theorem 1 (Convergence of Iterative Best Response). *IBR on the market described above converges to a unique (unregulated) price equilibrium. [Milgrom and Roberts, 1990]*

2.4 Emissions Regulation: Hard-cap

Under hard-cap emissions regulation, firms need to satisfy a cap on emissions per vehicle while maximizing vehicle profits. Again letting Ψ^0 be the firms' constant decision, this leads to a constrained profit function in the following form:

$$\Pi_j^{cap}(\Psi^0) | \Sigma_j, \phi, \kappa = \begin{cases} -\infty, & \text{if } \kappa_j^0 < 0 \\ T \times [\Psi_j^0 - \Sigma_j]'\phi_j^0, & \text{otherwise} \end{cases}$$

Recall that κ_j^0 is the overall amount of emissions credits assigned to firm j by regulators based on its vehicle sales, and it is a weighted sum of ϕ_j^0 where the coefficients are the EPA rating for each vehicle. We can now define Nash equilibrium:

Definition 2 (Hard-cap price equilibrium). *A set of prices Ψ^* is a hard-cap price equilibrium if, for all $j \in \{1, \dots, n\}$,*

$$\Pi_j^{cap}(\Psi^*) | \Sigma_j, \phi, \kappa \geq \Pi_j^{cap}(\Psi_j \cup \Psi_{-j}^*) | \Sigma_j, \phi, \kappa, \\ \forall \Psi_j \in \mathbb{R}^{|\mathcal{I}_j|}.$$

Although the hard-cap breaks the linearity of the cost function, we can still show (by employing results by [Vives, 1990; Vives, 2005]) that this market maintains the same useful property of the unregulated case.

Theorem 2 (Convergence of iterative best response under hard cap). *IBR on the market described above converges to a unique (hard-cap) price equilibrium.*

2.5 A Price-Takers Emissions Credits Market

In an effort to allow firms to meet regulatory constraints as efficiently as possible, credits can be traded on a *credits market*. This is intended to incentivize firms that have a lower cost for producing credits to take on more of the burden.

In the price-takers market, firms maximize the following profit function:

$$\Pi_j^W(\Psi^0) | \Sigma_j, \phi, \Omega, \kappa = \begin{cases} -\infty, & \text{if } \kappa_j^0 + \lambda_j^0 < 0 \\ T \times [\Psi_j^0 - \Sigma_j]'\phi_j^0 + \Omega\lambda_j^0, & \text{otherwise} \end{cases}$$

Recall that λ_j^0 is the number of credit obligations bought or sold. These are credits that transfer hands from a firm who produces them by producing non-polluting vehicles to a firm who needs them to satisfy regulatory requirements.

There are two key assumptions embedded in this payoff function. First, firms take credits prices as given externally. Second, firms do not attempt to manipulate the credits market by retaining credits for the next period, as is allowed by regulation. We will loosen both of these assumptions in our strategic credit trading model.

We can specify a Nash equilibrium given this payoff function. Note that although λ_j^0 is a choice made by firms, it is determined mechanically given Ψ_j^0 via the constraint $\kappa_j^0 + \lambda_j^0 < 0$. Thus, it is not a strategic variable.

Definition 3 (Price taker's credit market equilibrium). *A set of prices Ψ^* is a price taker's credit market equilibrium if, for all $j \in \{1, \dots, n\}$, $\Psi_j \in \mathbb{R}^{|\mathcal{I}_j|}$:*

$$\Pi_j^W(\Psi^*) | \Sigma_j, \phi, \kappa, \Omega \geq \Pi_j^W(\Psi_j \cup \Psi_{-j}^*) | \Sigma_j, \phi, \kappa, \Omega.$$

The convergence result of Iterative Best Response can easily be extended to this setting, by adding the cost of compliance to the cost of each vehicle.

Theorem 3 (Convergence of iterative best response under credit market price). *IBR on the market described above converges to a (credit market) equilibrium.*

Firms in this setting assume they will be able to buy and sell λ_j^0 credits without considering who their trading partner will be. In reality, these trades are not all feasible, as setting a particular Ω will usually induce a shortage or surplus in the credits market. So the price taker's Nash equilibrium is only feasible if a *market clearing condition* holds:

Definition 4 (Walrasian credit equilibrium). *A tuple (Ψ^*, Ω) is a Walrasian credit equilibrium if Ψ^* is a Price taker's credit market equilibrium, and in addition the following condition holds:*

$$\sum_{j \in \{1, \dots, n\}} \lambda_j^* = 0$$

With credits in hand, a firm has the flexibility to produce more polluting vehicles and earn more profit. By selling credits, a firm has constrained themselves to losing profits in the vehicle market since it must produce enough non-polluting vehicles to honor the sale. At a Walrasian credit equilibrium, buyers have determined that their collective profit gain per credit in the Bertrand pricing game from buying a set number of credits is equal to the given price of credits. Similarly, the collective profit loss per credit experienced by sellers by selling a set number of credits is equal to the given price of credits. Thus, the price of credits is analogous to the cost of abatement, and the market clearing condition is analogous to equality of marginal costs of abatement across firms.

3 Strategic Emissions Market

The assumption that firms are strictly price takers in the credits market is a simplification of reality. Firms have an incentive to buy and sell credits at prices favorable to themselves, and they have different levels of bargaining power to achieve those prices. In the strategic emissions market, we allow firms to play out this bargaining process using heuristic strategies while trying to maximize estimated profit from both credit obligation transactions as well as the vehicle market.

The mechanism for exchanging credits is a single price multi-unit double auction. Firms submit (price, quantity) bids to either sell or buy credits. After all bids clear, firms observe the outcome of their bids and submit a new round of bids. This process iterates until a fixed horizon is reached. This mechanism does nothing to ensure truth telling, and rich strategic behavior is expected.

Firms in this model strategize over two bidding parameters: χ_j , an *impatience parameter*, and μ_j , a *shading parameter*. Roughly speaking, χ_j corresponds to a bid pricing strategy, while μ_j corresponds to a bid quantity strategy. A baseline bidding strategy for each firm is to offer around the historically observed market price of credits at a quantity given by the amount of credits transacted at the associated price taker's credit market equilibrium. This strategy, utilized by choosing a large μ_j and large χ_j , signifies a firm's willingness to be a price taker.

Conversely, a low χ_j and μ_j signifies a firm's willingness to wait for a high surplus credit transaction. These firms will offer prices above (for sellers) or below (for buyers) the historical market average, and they will only offer the price taker's credit market equilibrium quantity if the surplus per

credit is very high. This strategy signifies a firm’s belief that it has market power in the credit’s market.

Firms choose their bidding parameters strategically by observing an overall payoff, generated as follows. First, firms submit bids in the credit market. They observe the resulting clearing price and execute any bids that have been matched. They then adjust their bids based on this information, and submit them to the double auction again. They do this for ν *bidding iterations*, coming away with either positive or negative revenue and, correspondingly, a negative or positive credits balance. Using what they have observed in the credits market, they form a belief about what the future credits price is. Armed with their actual credit balance and their estimated credit price, they enter the vehicle market where iterated best response runs and converges. They then carry forward any unused credits into the next period, and repeat the process. The total *number of periods* is an environment setting, called T . The total payoff is the sum of all credit revenues added to the sum of all vehicle profits.

We need to simulate payoffs for many strategy profiles before being able to find a Nash equilibrium in χ and μ . To do this search, we use empirical game-theoretic analysis, which is a tool for sampling payoffs intelligently [Cassell and Wellman, 2013].

The purpose of this model is not to incorporate every possible bidding strategy, but to complement our price-takers credit market analysis. By allowing firms to exercise market power to some degree, we address a gap in our analysis of the credit market for vehicles since we know that in reality there are firms who can bargain successfully. For a deeper, reproducible explanation of the strategic credit trading model, see the supplemental materials¹.

4 Experiments and Results

4.1 Scenarios and Setup

The *baseline scenario* in our experiments has $n = 3$ and $m = 14$. ϕ , Σ , and κ are given by our industry sponsor and are based on real data and government guidance. Each firm can produce a mix of emissions credit generating (mostly electric) vehicles and credit consuming (gas) vehicles, at different profitability levels. Firm 1 can produce credit generating vehicles at a lower cost than the others, and will be the seller in the credits market. The other two firms produce highly profitable gas vehicles, and will be buyers in the credits market. Unlike firms 1 and 2, firm 3 does not produce trucks, and produces a much smaller total number of vehicles than the other two.

Two other scenarios examine how extreme types of industrial organization affect the model. The first scenario is a *monopoly scenario* where all products are produced by a single firm ($n = 1, m = 14$). The second scenario is a *fragmented scenario* where every product is produced by a different firm ($n = 14, m = 14$). Although the ownership relation from firms to vehicles changes, we keep ϕ , Σ , and κ constant throughout.

¹https://www.dropbox.com/s/a2dfjzkfp25ppgv/ijcai2019_supplemental.pdf?dl=0

	Unregulated			Hard Cap			Walrasian credits			Strategic credits (long)			Strategic credits (short)		
Monopoly	12.4	6.3	18.7	10.8	5.1	15.8	-	-	-	-	-	-	-	-	-
Baseline	10.5	8.9	19.4	7.3	4.8	12.1	10.5	5.9	16.5	10.1	6.0	16.1	10.2	6.0	16.2
Fragmented	10.3	9.1	19.3	3.1	2.8	5.91	10.7	7.8	18.5	-	-	-	-	-	-

Figure 2: Aggregate welfare in millions. Profit in green followed by consumer surplus in red followed by total welfare in blue

For the strategic credit market model, we examine two settings. First, we set $T = 4$ and $\nu = 20$. We call this model setting *strategic credits (long)*. The other strategic model setting, *strategic credits (short)*, has $T = 5$ and $\nu = 10$. The firms in strategic credits models draw bidding strategy (μ_j, χ_j) from the set $(\chi, \mu) | \chi \in \{0.5, 0.7, 0.9, 0.999\}, \mu \in \{0.01, 0.4, 0.7\}$.

4.2 Evaluating Welfare

The measure of value we use to judge equilibria is *total welfare*. Welfare in our case is defined as the sum of *consumer surplus* and firm profit. We do not attempt to quantify the benefit to the environment, assuming that emissions regulation is doing its job.

Consumer surplus is a measure of the total savings (compared to their own valuation) that consumers get from buying vehicles at an equilibrium price. Given a set of equilibrium prices Ψ^* , consumer surplus for vehicle i (denoted γ_i) is calculated as follows:

$$\gamma_i = \int_{\Psi_{(i)}^*}^{\infty} \phi(\Psi_{(i)} | \Psi_{(-i)}^*) d\Psi_{(i)}$$

Total consumer surplus is just $\sum_{i=1}^m \gamma_i$. In general, consumer surplus decreases with higher equilibrium prices. Calculation is done via numerical integration.

Welfare results are shown in Figure 2. Omitted entries for the monopoly row are due to the impossibility of trading with oneself. Instead, since the monopoly controls all vehicle markets, it generates credits in fuel-efficient segments for use in fuel inefficient ones. Thus, the hard-cap model for monopoly is comparable to the credit trading models for other scenarios.

Omitted entries in the competitive setting were due to computational constraints. Since none of these games are symmetric, expanding to 14 firms using the full strategy set requires evaluating payoffs for $12^{14} = 1.28e+15$ profiles.

We call attention to several phenomena in Figure 2:

1. Monopoly profit is higher than all other scenarios in the unregulated model. In addition, monopoly profit under the hard-cap scenario beats profit under all other regulated scenarios. Both of these results are expected, as monopolies can coordinate between the different vehicle markets more efficiently, both to avoid cannibalization and to navigate regulatory requirements.
2. The unregulated monopoly beats all other unregulated scenarios by much more than the hard-cap monopoly beats the regulated scenarios with credit trading. The

credits market allows firms to collectively achieve a near-monopolistic profit level under regulation.

- In the baseline and fragmented scenarios, credit trading results in profit levels at or above those of the unregulated model. Meanwhile, consumer surplus always decreases greatly when regulation with credit trading is applied. Consumers are hurt more than firms by emissions regulation. This effect is least pronounced, though still significant, in the fragmented scenario.
- Strategic bidding decreases profits and raises consumer surplus slightly.

The overarching theme is that when credit trading is introduced, firms are able to offset and even overturn losses due to hard-cap emissions constraints by raising prices in tandem. Consumer surplus suffers as a result.

It is the unique industrial organization in the baseline scenario that allows firms to maintain high prices and profits given the ability to trade credits. Namely, each firm has a very diverse set of vehicles. This allows each firm to strategically exit and enter certain markets based on their credit costs. Taking the unregulated scenario as a baseline, firm 1, which must maintain a high credits balance, abandons high-pollution trucks. Firm 2, which can buy low-cost credits from firm 1, abandons low-pollution cars. This segmentation gets rid of competitive forces within each vehicle type’s market and keeps prices high.

This is in contrast to the fragmented scenario, where firms do not have the option of exiting a vehicle segment. Thus, competitive forces are maintained in each vehicle segment, and consumer surplus is highest in this scenario out of all regulated scenarios. Profit is actually almost at monopoly levels as well. This is because of the firms producing polluting vehicles, the most cost-efficient ones become more dominant as they can afford to buy more credits at the same price. They do this because they know that they can maintain margins even with the credit price factored in. This efficiency results in the highest total welfare out of any regulated scenario. Overall, we find that a more fragmented industrial organization is beneficial to cap-and-trade regulation.

4.3 Strategic Credit Trading Analysis

The shift from profits to consumer surplus in the strategic trading case stems from the fact that fewer credits are traded. Figure 4 shows how profit and credit market activity varies among firms, compared to the Walrasian equilibrium baseline. In all cases, total credit obligations trading volume decreases in the strategic case. The one exception is for Firm 3 in the “long” model, where the credit balance is similar to the Walrasian case. Both large firms (1 and 2) suffer profit losses, while the smaller firm 3 comes out ahead. For an analogous table for the “short” setting, see the supplemental materials.

The root cause of this is the large amount of shading applied at equilibrium (see Figure 3). The two large firms are engaged in restricting credits volume for fear of the credit price moving against them if they yield. Thus, they both shade to the maximum level. This leads them to compete with each other in the vehicle market, as firm 1 owes less credit obligations so prefers to sell more polluting vehicles.

	v = 20, T = 4		v = 10, T = 5	
	μ^*	χ^*	μ^*	χ^*
Firm 1	0.5	0.01	0.5	0.4
Firm 2	0.5	0.01	0.5	0.4
Firm 3	0.9	0.01	0.7	0.7

Figure 3: Nash equilibrium bidding parameters

Firm	Vehicle Profit	Total Profit	Absolute Credits Balance	Average Credit Price
1	-2.83%	-4.16%	-4.26%	-1.41%
2	-8.53%	-9.74%	-5.09%	-1.54%
3	3.28%	4.85%	-0.69%	-0.87%

Figure 4: Aggregate changes in strategic trading (long) baseline vehicle market as a percentage from Walrasian model

A secondary effect is that firm 2 focuses on selling its most profitable polluting vehicle, trucks. This creates an opening for firm 3, who chooses to shade less in order to buy more credits and sell more SUVs. Overall, there is less segmentation of the market between the two large firms, resulting in a transfer of welfare from firms to consumers. For evidence of this summary of results at the vehicle level, see the supplemental materials.

The biggest welfare effect of the strategic maneuvering in the credits market is more competition in the truck sector in the (long) model. Since firm 3 has relatively many credits, it focuses on selling SUVs, crowding the other two firms out into the truck market. The subsequent price war transfers welfare from firms to consumers.

5 Conclusion

By analyzing how firms adapt strategically to emissions regulation, we give insight into the question: who pays for emissions reduction in the US vehicles market? The answer is, resoundingly, that consumers pay as firms segment themselves by vehicle type and raise prices to protect themselves against production constraints. Trading emissions credits provides a remarkably efficient way for firms to preserve profits, as even a monopoly facing emissions regulation cannot do appreciably better. Strategic trading provides some relief to consumers, as firms tend to limit credits trading volume and thus compete more. But with this redistribution of welfare comes some deadweight loss. The best way we have found to bring consumer surplus up, and with it overall welfare, is to fragment the production of vehicles so that no firm has an incentive to focus on one vehicle at the expense of another. Our results suggest that cap-and-trade is definitely an improvement over a hard-cap, but that caution must be used to ensure that price hikes do not become an issue for consumers in the oligopolistic vehicle market.

Acknowledgements

This work was partially funded by Ford’s GDI&A university collaboration program. The authors would like to thank Ford collaborators Drs. Suzhou Huang, Shaozhong Zhao, and Alan Xu.

References

- [Benz and Trück, 2009] Eva Benz and Stefan Trück. Modeling the price dynamics of co2 emission allowances. *Energy Economics*, 31(1):4 – 15, 2009.
- [Boswall and Lee, 2002] J. Boswall and R. Lee. *Economics, Ethics and the Environment*. Taylor & Francis Group, 2002.
- [Carlson *et al.*, 2000] Curtis Carlson, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer. Sulfur dioxide control by electric utilities: What are the gains from trade? *Journal of Political Economy*, 108(6):1292–1326, 2000.
- [Cassell and Wellman, 2013] Ben-Alexander Cassell and Michael P. Wellman. EGTAOnline: An experiment manager for simulation-based game studies. In *Multi-Agent Based Simulation XIII*, volume 7838 of *Lecture Notes in Artificial Intelligence*, pages 85–100. 2013.
- [Fullerton and Gan, 2005] Don Fullerton and Li Gan. Cost-effective policies to reduce vehicle emissions. *American Economic Review*, 95(2):300–304, 2005.
- [Godby, 2002] Robert Godby. Market power in laboratory emission permit markets. *Environmental and Resource Economics*, 23(3):279–318, 2002.
- [Goulder, 2013] Lawrence H. Goulder. Markets for pollution allowances: What are the (new) lessons? *Journal of Economic Perspectives*, 27(1):87–102, 2013.
- [Jordan *et al.*, 2007] Patrick R. Jordan, Christopher Kiekintveld, and Michael P. Wellman. Empirical game-theoretic analysis of the tac supply chain game. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS '07*, pages 193:1–193:8, New York, NY, USA, 2007. ACM.
- [Karplus *et al.*, 2010] Valerie J. Karplus, Sergey Paltsev, and John M. Reilly. Prospects for plug-in hybrid electric vehicles in the united states and japan: A general equilibrium analysis. *Transportation Research Part A: Policy & Practice*, 44(8):620 – 641, 2010.
- [Kesicki, 2012] Fabian Kesicki. Intertemporal issues and marginal abatement costs in the uk transport sector. *Transportation Research Part D: Transport and Environment*, 17(5):418 – 426, 2012.
- [Li and Das, 2016] Zhuoshu Li and Sanmay Das. An agent-based model of competition between financial exchanges: Can frequent call mechanisms drive trade away from cdas? In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, AAMAS '16*, pages 50–58, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.
- [Limpaitoon *et al.*, 2011] Tanachai Limpaitoon, Yihsu Chen, and Shmuel S. Oren. The impact of carbon cap and trade regulation on congested electricity market equilibrium. *Journal of Regulatory Economics*, 40(3):237–260, 2011.
- [Malueg, 1990] David A Malueg. Welfare consequences of emission credit trading programs. *Journal of Environmental Economics & Management*, 18(1):66 – 77, 1990.
- [Mannering and Winston, 1985] Fred Mannering and Clifford Winston. A dynamic empirical analysis of household vehicle ownership and utilization. *RAND Journal of Economics*, 16(2):215–236, 1985.
- [Milgrom and Roberts, 1990] Paul Milgrom and John Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990.
- [Muller *et al.*, 2002] R. Andrew Muller, Stuart Mestelman, John Spraggon, and Rob Godby. Can double auctions control monopoly and monopsony power in emissions trading markets? *Journal of Environmental Economics & Management*, 44(1):70 – 92, 2002.
- [Ogden, 1966] Delbert C. Ogden. Economic analysis of air pollution. *Land Economics*, 42(2):137–147, 1966.
- [Perl and Dunbar, 1982] Lewis J. Perl and Frederick C. Dunbar. Cost effectiveness and cost-benefit analysis of air quality regulations. *American Economic Review*, 72(2):208–213, 1982.
- [Reddy and Veloso, 2011] Prashant P Reddy and Manuela M Veloso. Learned behaviors of multiple autonomous agents in smart grid markets. In *Twenty-fifth AAAI conference on Artificial intelligence*, 2011.
- [Rubin *et al.*, 2009] Jonathan Rubin, Paul N. Leiby, and David L. Greene. Tradable fuel economy credits: Competition and oligopoly. *Journal of Environmental Economics & Management*, 58(3):315 – 328, 2009.
- [Sartzetakis, 1997] Eftichios Sophocles Sartzetakis. Tradeable emission permits regulations in the presence of imperfectly competitive product markets: Welfare implications. *Environmental and Resource Economics*, 9(1):65–81, 1997.
- [Vives, 1990] Xavier Vives. Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19(3):305 – 321, 1990.
- [Vives, 2005] Xavier Vives. Games with strategic complementarities: New applications to industrial organization. *International Journal of Industrial Organization*, 23(7):625 – 637, 2005.
- [Wah and Wellman, 2013] Elaine Wah and Michael P. Wellman. Latency arbitrage, market fragmentation, and efficiency: A two-market model. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC '13*, pages 855–872, New York, NY, USA, 2013. ACM.
- [Westskog, 1996] Hege Westskog. Market power in a system of tradeable carbon dioxide quotas. *Energy Journal*, 17:85–104, 1996.