Essays in Industrial Organization and Environmental Policy

HARVARD UNIVERSITY
Graduate School of Arts and Sciences

DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the Committee for the PhD in Business Economics have examined a dissertation entitled

Essays in Industrial Organization and Environmental Policy

Presented by **Francis Robert Pinter**

candidate for the degree of Doctor of Philosophy and hereby certify that it is worthy of acceptance.

Signature ______________________

Ariel Pakes, Chair

Signature ______________________

Robin Lee

Signature ______________________

Dennis Yao

Date: April 29, 2022

Essays in Industrial Organization and Environmental Policy

A dissertation presented

by

Francis Robert Pinter

to

The Committee on Business Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Business Economics

Harvard University

Cambridge, Massachusetts

April 2022

Dissertation Advisor:
Professor Ariel Pakes

Author:
Francis Robert Pinter

Essays in Industrial Organization and Environmental Policy

Abstract

The first chapter, coauthored with Sarah Armitage, examines the effects of state policy when markets are nationally standardized. When prices are uniform across states, state-level supply-side and demand-side tools, such as producer and consumer subsidies, generate different consumer prices and markups over marginal cost. We study the 2009–2017 electric vehicle industry, which received substantial support from state-level policies. Using a structural model of demand and supply in the new vehicle market, we compare the effects of an important state-level supply-side policy, the zero-emission vehicle mandate in California and nine other states, to those of a counterfactual demand-side policy that instead uses a consumer subsidy and tax. Holding fixed the regulator's stated target, electric vehicle sales in regulated states, the demand-side policy generates higher consumer prices for most electric vehicles, resulting in lower consumer and total surplus. These results persist if electric vehicle product introduction is allowed to adjust.

The second chapter, coauthored with Sarah Armitage, asks how policymakers seeking to support socially beneficial products should account for the costs and benefits of product variety. Using a model of electric vehicle product introduction, and estimates from the first chapter on demand and costs for new vehicles, we measure the social welfare effects of variety in the first generation of electric vehicles and ask whether a state-level policy to encourage electric vehicles aligned these social welfare effects with private incentives. We measure the effects of policy-induced variety on consumers across the income distribution, producers, and the environment. We then apply the model to study the variety effects of the counterfactual policy from the first chapter.

The third chapter examines the credit banking and trading mechanisms used by state and federal vehicle regulations to reduce compliance costs. I document stylized facts from the zero emission vehicle and light-duty vehicle greenhouse gas credit markets: firms acquire credits even when they have large balances, and credit trades result in less concentrated balances across firms. I then develop a dynamic model of automakers who trade credits on a competitive market, but face a risk that the market will be unavailable. I then use the model to explore the effects of changes to credit trading and banking policy.

Contents

List of Tables

List of Figures

Acknowledgments

I would like to thank my advisors, Ariel Pakes, Robin Lee, and Dennis Yao, for their guidance and support over the last several years. Their mentorship has been crucial to my development as an economist and as a researcher. In addition, support and advice from Ed Glaeser, Myrto Kalouptsidi, Jing Li, Mark Shepard, and Elie Tamer have been essential to my progress in graduate school and to the work in this dissertation.

I owe an enormous debt to my fellow graduate students for their encouragement, for their help with research, and for providing a community to fall back on. Collaborating with Sarah Armitage has stretched my intellectual boundaries and has been a motivating and exciting experience. I would also like to thank Angie Acquatella, Alex Albright, Marcos Barrozo, Justin Bloesch, Jake Bradt, Liz Engle, Jack Feng, Ashvin Gandhi, Jeff Gortmaker, Rishab Guha, Stephanie Kestelman, Gabriel Levin Konigsberg, Tianwang Liu, Erica Moszkowski, Sagar Saxena, Gregor Schubert, Yin Wei Soon, Anna Stansbury, Ron Yang, Anthony Yu, and David Zhang, among many others.

Chapters 1 and 2 build on my job market paper, coauthored with Sarah Armitage. That paper benefited from constructive comments at the Harvard Industrial Organization workshop, Environmental Economics workshop, Environmental Economics seminar, and the 2019 Harvard Environmental Economics Program conference. We acknowledge financial support from the Lab for Economic Applications and Policy and Harvard Business School. We also thank three industry experts who generously shared their time and information on the automobile industry. Charles Murry, Mathias Reynaert, and Tom Wollmann also provided helpful comments, as did Sarah's advisors Joe Aldy, Ed Glaeser, Robin Lee, and Jim Stock. Chapter 3 benefited from many useful suggestions and comments from Sarah Armitage and the Harvard Environmental Economics workshop.

Finally, I am grateful to friends and family, near and far, for supporting me in my journey while reminding me of the value of having a life outside of graduate school. I would like to thank my parents and my brother Bobby, who have always been my foundation, and Aarushi Gupta, whose place in my life means everything to me.

Chapter 1

The Consequences of Uniform Pricing for State Regulatory Policy[1]

1.1 Introduction

In the United States, environmental policy is often made in a bifurcated policy landscape. Some states, especially California, have adopted generous subsidies and stringent standards, while other states have adopted few environmental interventions. When policy affects products that are sold in national product markets, evaluating the impact of state-level policies requires understanding not only direct impacts within the regulated state's borders, but also any spillovers to non-covered states.[2] These spillovers are especially important when product prices are standardized across space. We study how instrument choice for state-level environmental policy affects these inter-state spillovers on pricing in the new vehicle market, using the Zero Emissions Vehicle (ZEV) program as a case study.

Standard economic theory predicts that supply-side and demand-side policy instruments are equivalent if the policy covers the entire market. However, in an industry in which

[1]Co-authored with Sarah Armitage.

[2]See, e.g., "Can America's Blue States Tackle Climate Change on Their Own?" (Jonathan Eyer and Matthew E. Kahn, Harvard Business Review, 6/6/17).

firms set national prices, state-level demand-side tools such as consumer subsidies generate price variation across regions. By contrast, state-level supply-side tools such as producer subsidies do not produce the same geographic price differentiation. As a result, the two tools may generate different firm incentives, producing different markups over marginal cost. Furthermore, the pass-through of policy to prices will vary across products based on their relative exposure to the policy. In oligopoly, these heterogeneous exposures will interact with substitution patterns across products to produce differential price impacts. The effects on consumer, producer, and total welfare will then differ, as will the policy's overall effectiveness at achieving environmental goals.

We study the Zero Emissions Vehicle mandate, a prominent state-level policy shaping the electric vehicle industry over 2009 to 2017. Adopted by California and nine other US states, the mandate required the largest automakers to meet a quota of electric vehicles of 0.4–1.5% of their statewide sales. To obtain credit toward the quota, the automaker could sell electric vehicles in the state or buy credits from other automakers. The goal of the mandate was to induce sales of electric vehicles to reach mass-market quantities. Of course, passenger automobiles are a national market, and efforts to spur the development of EVs in certain states are likely to affect automakers' incentives more broadly. We investigate the consequences of policymakers' choice to use a supply-side mandate rather than rely on demand-side policies, focusing on the effect on pricing in oligopoly.

To analyze social welfare and the incentives facing manufacturers, we build and estimate a model of consumer demand and producer price-setting in the market for new passenger vehicles in the United States from 2009 to 2017. We adopt a Berry, Levinsohn, and Pakes (1995)-style model of discrete choice demand, with differentiated products and heterogeneous consumer tastes, and Bertrand competition with multiproduct firms. We adapt the standard model to account for nationwide pricing and heterogeneity across states in environmental regulation.

We then simulate a counterfactual demand-side policy that replaces the ZEV mandate with a budget-neutral combination of a consumer subsidy for electric vehicles and consumer tax on non-electric vehicles.[3] Without any change in the product set, the policy results in $1.4 billion lower consumer surplus and $570 million higher producer surplus across the nationwide market for new vehicles. Because of the low sales of electric vehicles in this period, the environmental effects of both policies are small. Our results suggest that state-level and national policymakers seeking to encourage the development and adoption of nascent socially beneficial products face different problems, and the choice of policy tool matters for state-level policy in ways that would not matter for national policy.

1.1.1 Literature

Industry observers have claimed that the ZEV mandate influenced automakers' electric vehicle programs,[4] but the literature on the early development of electric vehicles has focused on other policies. This literature (reviewed most recently in Rapson and Muehlegger (2021)) has quantified the effects of purchase subsidies (Tal and Nicholas 2016; Jenn, Springel, and Gopal 2018; Muehlegger and Rapson 2018; Muehlegger and Rapson 2020; Remmy 2020; Xing, Leard, and Li 2021); public and private investment in complementary infrastructure, particularly charging stations (Li 2019); and a combination of both (Li, Tong, Xing, and Zhou 2017; Zhou and Li 2018; Springel 2021). Two exceptions are Holland, Mansur, and Yates (2021), which evaluates a hypothetical cap-and-trade system to limit sales of gasoline vehicles over a long horizon, and Cole, Droste, Knittel, Li, and Stock (2021), which evaluates a hypothetical future national ZEV mandate.

To our knowledge, this paper (together with Chapter 2) is the first systematic welfare analysis of the ZEV mandate. Prior literature on the design and effects of the mandate from other perspectives, which has informed our modeling and our discussion of institutional features, includes Dixon, Porche, and Kulick (2002), Bedsworth and Taylor (2007), Vergis

[3]Even a policy directly targeting the environmental externality, such as a carbon tax, may have ambiguous welfare effects when markets are imperfectly competitive and entry is endogenous (Fowlie, Reguant, and Ryan 2016).

[4]See "Automakers question Calif. zero-emission mandate as feds reassess mpg rules" (Eric Kulisch, Automotive News, 12/12/17).

and Mehta (2012), Greene, Park, and Liu (2014), Linn and McConnell (2017), and McConnell and Leard (2021).

Our study of the ZEV mandate broadens an extensive literature on the effects of supply-side environmental policies in the automobile industry, which has focused primarily on fuel economy standards like the Corporate Average Fuel Economy (CAFE) and state and federal greenhouse gas standards. Like the ZEV mandate, fuel economy standards simultaneously target pollution externalities (for which they are less efficient than a fuel tax (Sallee 2011b)) and innovation market failures (Jaffe, Newell, and Stavins 2005). Unlike policies that target electric vehicles, which support entirely new types of products, fuel economy standards alter the mix of existing product types and encourage improvements to existing technologies. The fuel economy literature has documented effects of these standards on vehicle characteristics (Knittel 2011; Klier and Linn 2012; Whitefoot, Fowlie, and Skerlos 2017; Ito and Sallee 2018; Reynaert 2021) and equilibrium prices and quantities (Goldberg 1998; Goulder, Jacobsen, and van Benthem 2012; Jacobsen 2013; Davis and Knittel 2018), and estimated the costs of compliance (Anderson and Sallee 2011). This literature has typically contrasted standards with intensity-based policies like fuel taxes (Knittel 2012; Anderson and Sallee 2016).

The closest work to ours in the fuel economy standards literature is Durrmeyer and Samano (2018), which contrasts supply-side fuel economy standards with a program of consumer taxes and subsidies within a structural model of demand and supply. In their model, a supply-side standard that operates firm-by-firm induces a different shadow cost of regulation at each firm, while consumer subsidies and taxes equalize shadow costs across firms. Standards with credit trading among firms, like CAFE since 2011, are equivalent to consumer subsidies and taxes in their setting.

Other work has examined the effects of purchase subsidies, particularly for hybrid vehicles in the early 2000s (Sallee 2011a; Beresteanu and Li 2011) and consumer subsidies and taxes to encourage fuel efficiency (Durrmeyer and Samano 2018; Durrmeyer forthcoming). Sallee (2011a) documents a difference between national demand-side and supply-side subsidies for the second-generation Toyota Prius, and argues for a mechanism based on dynamics in consumer perceptions.

Our study of the interaction between state-level policy and national prices connects to a literature on uniform pricing across markets, which has largely examined retail chains. This literature has shown that uniform pricing can increase profits under oligopoly (Adams and Williams 2019), and alters the equilibrium effects of local shocks and policies (DellaVigna and Gentzkow 2019; Leung 2021).

Our computations of environmental damages build on a literature comparing the emissions of electric vehicles to those of their closest gas-powered substitutes (Holland, Mansur, Muller, and Yates 2016; Holland, Mansur, Muller, and Yates 2020; Muehlegger and Rapson 2020; Xing, Leard, and Li 2021), which has generally found that, during the early 2010s, electric vehicle adoption had a small effect on short-run pollution damages, and in some regions worsened them. Of these, Xing, Leard, and Li (2021) also uses a random coefficients discrete choice model to study the consequences of substitution between electric and non-electric vehicles.

1.2 Institutional background

We study the period covered by model years 2009 to 2017, which saw the introduction of the first generation of commercially available electric vehicles (EVs) in the US. Most major automakers in the US introduced an electric vehicle during the period, but models varied widely in engineering characteristics and in sales levels.

1.2.1 Policy of interest: ZEV mandate

The Zero Emission Vehicle (ZEV) mandate, adopted by California and nine other states,[5] required the largest automakers to sell a quota of non-fossil-fuel ("zero emission") vehicles, nearly always battery electric vehicles.[6] Each manufacturer's quota was based on a moving

[5]New York, Massachusetts, Vermont, Maine, Connecticut, Rhode Island, Oregon, New Jersey, and Maryland (starting 2011).

[6]Hydrogen fuel cell vehicles also counted generously toward the quota, but few were sold in this period.

average of its past sales of non-electric vehicles in California, and manufacturers could trade credits with each other and bank credits for later use. In our study period, six manufacturers faced the ZEV quota: Chrysler, Ford, GM, Honda, Nissan, and Toyota.

The number of credits earned for a battery electric vehicle was a function of the range the vehicle can travel on a full battery; each manufacturer accumulated credits by selling vehicles, and the manufacturer's credit requirement was a percentage of its statewide sales of non-electric vehicles.[7] For example, in California in model year 2017, Nissan earned three credits for each sale of the Leaf electric vehicle and faced a quota of 3,800 credits (3% of its California sales volume of 127,800), which translated to a quota of 1,300 Leaf vehicles. Nissan well exceeded this quota, selling 4,600 Leaf vehicles in California and 1,100 in the other nine states. (If its sales had fallen short, it could have drawn on its bank of 50,800 credits or purchased credits from another manufacturer.) If an automaker missed its quota in any given year, it had two years to make up the deficit. After that point, in order to return to compliance, it was required to pay a penalty of $5,000 per credit and also make up the deficit.[8]

In this period, the ZEV sales requirement did not strictly apply state-by-state, but instead allowed each state's quota to be met with vehicles sold in any of the regulated states. This rule, called the travel provision, allowed automakers to earn credits in all regulated states for a sale in any of them. As a result, automakers could (and often did) meet their requirements only by selling vehicles in California, which had the largest population of the participating states, better charging infrastructure, and generous government subsidies to consumers.

In addition to the mandate on zero emission vehicles, the ZEV program included a mandate on clean gasoline vehicles, hybrids, and plug-in hybrids, collectively dubbed Partial Zero Emission Vehicles (PZEVs). The mandate applied to automakers that faced

[7] Statewide sales of non-electric vehicles are constructed as a moving average of past years; see Section A.1 for details.

[8] Between 2009 and 2017, no manufacturer was noncompliant. One manufacturer had a deficit that it made up the following year.

the ZEV mandate and an additional group of mid-sized automakers.[9] The ZEV and PZEV mandates were not entirely separate, as excess ZEV credits could count toward the PZEV credit requirement. Nonetheless, each manufacturer's sales of hybrids and clean gasoline vehicles each year was well over PZEV requirements, and there was little trading of PZEV credits among manufacturers. As a result, we assume that the Partial Zero Emission Vehicle mandate was not a binding constraint on any automaker.

Amendments to make the program stricter were announced in 2012 and took effect in model year 2018. The most important changes to the ZEV mandate were to reduce the number of credits earned per vehicle, add to the list of manufacturers who faced the ZEV quota,[10] and replace the travel provision with a cross-state transfer mechanism that did not allow double-counting. In addition, the PZEV mandate was restricted to plug-in hybrids only. Manufacturers anticipating stricter post-2018 regulation may have earned surplus credits before 2018 in order to bank them.

Section A.1 describes the rules of the ZEV mandate in greater detail.

1.2.2 National pricing

Though our setup models consumers as paying a national posted price directly to the manufacturer, only Tesla sells vehicles this way. In practice, other manufacturers sell new vehicles through dealerships, which negotiate adjustments from the posted MSRP separately with each consumer. The bargaining process results in different prices for each consumer, based on factors like consumer demographics, fluctuations in dealer inventory, consumer information, add-on vehicle options, and the potential trade-in value of the consumer's existing vehicle (Murry and Schneider 2016). None of these factors differs systematically across states based on state-level policy. As long as manufacturers cannot respond to cross-state policy differences by manipulating prices separately in each state, the basic

[9]Between 2009 and 2017, this group consisted of BMW, Daimler, Hyundai, Jaguar Land Rover, Kia, Mazda, Mitsubishi (2009 only), Subaru, Volkswagen, and Volvo (2009–11 only).

[10]This change added BMW, Daimler, Hyundai, Kia, and Volkswagen, by reducing the sales threshold at which the quota would apply.

implications of our model of national pricing still hold.[11]

The greatest challenge to our use of national pricing comes through the rebates that manufacturers provide directly to dealers or consumers (Busse, Silva-Risso, and Zettelmeyer 2006). If manufacturers use rebates to adjust prices across states to reflect policy differences, then our model of national pricing would be invalidated. However, we have not found evidence of this behavior for EV pricing. A 2017 California Air Resources Board (CARB) report (California Air Resources Board 2017) found that manufacturer incentives for EVs in February–August 2016 (as reported by AutoNews) were comparable in Seattle, in a non-ZEV state, to a selection of cities in ZEV states. (Manufacturer incentives for hybrids were somewhat lower in Seattle than in the other cities.) Nonetheless, we cannot rule out that manufacturers can vary consumer prices by manipulating dealer inventories, because this behavior would not appear in dealer prices or manufacturer incentives.

Our most direct evidence on relevant transaction prices is from a 2017 analysis of Kelley Blue Book data by the Energy Information Administration (Bratvold and Cleaver 2017), which found that average vehicle sales prices for the 2016 Nissan Leaf SV varied little across metropolitan areas, including between metropolitan areas in ZEV and non-ZEV states.

Though California Air Resources Board (2017) found differences across states in median model sales prices, it did not account for trims. The median price paid for a 2015 Nissan Leaf, across all trims, was $28,900 in California (DMV data) and $25,123 in a group of ten non-ZEV states.[12] The 2015 Nissan Leaf had three trim levels with MSRPs of $29,010, $32,100, and $35,120, respectively (as reported by MSN Autos), so it is possible that California buyers paid more because they opted for higher trim levels. (For this model, median transaction prices were clearly well below MSRP.)

[11] In our model, manufacturers cannot respond to any state-level policy by changing prices separately in each state; this includes both the supply-side ZEV mandate and any demand-side subsidies, which also vary across states.

[12] CARB obtained this number from Experian Automotive data from Colorado, Kentucky, North Carolina, North Dakota, New Mexico, Ohio, Oklahoma, Texas, Virginia, and West Virginia.

1.3 Model

We build and estimate a model of manufacturer pricing and consumer choice each period, which then enables us to run counterfactual simulations of market outcomes under alternative policy instruments. We take as given the set of products available in each model year. Conditional on this product set, firms set prices and consumers choose which products to purchase.

To study consumer surplus and substitution among alternatives, we use a random coefficients logit model of consumer demand for automobiles. To estimate marginal costs, we use a Nash–Bertrand pricing model, with regulatory credits explicitly included in the profit function. The pricing model explicitly combines state-level markets with national pricing decisions, which enables us to explore differences in pass-through for producer and consumer subsidies at the state level.

1.3.1 Demand and pricing

Our analysis is built on a discrete choice model of demand for new vehicles in the vein of Berry, Levinsohn, and Pakes (1995) and Berry, Levinsohn, and Pakes (1999).[13] In each region and model year, there is a population of consumers who each choose one product: either one of the gasoline vehicles, electric vehicles, and hybrids available in that region and model year, or an outside good, which captures the choice not to buy a new vehicle. (The outside good could include driving an existing vehicle for longer, buying a used vehicle, or forgoing car ownership.)

Let the set of geographical regions (US states) be $\mathcal{M}$ and index regions by $m \in \mathcal{M}$. The periods are model years, indexed by $t = 1, \ldots, T$. Let the set of products available in region m and year t be $\mathcal{C}_{mt}$, and index products by j. As described in 1.4.1, products are defined by make, model, technology type, and (within electric vehicles) battery size.

[13] Papers with similar approaches to estimating automobile demand include Petrin (2002), Remmy (2020), Reynaert (2021), and Grieco, Murry, and Yurukoglu (2021), as well as Li (2019) (who only models the plug-in vehicle market).

Demand. Indirect utility for consumer i in region m and model year t from purchasing product j is

$$u_{ijmt} = \alpha_i(p_{jt} - \text{subsidy}_{jmt}) + x'_{jmt}\beta_i + \xi_{jmt} + \varepsilon_{ijmt},$$

where p_{jt} is the price of product j (set nationally), subsidy_{jmt} is the government subsidy for j in region m, x_{jmt} is a vector of observed characteristics, ξ_{jmt} is a quality shock unobserved by the econometrician, and ε_{itmj} is a Type 1 Extreme Value shock distributed independently across consumers, alternatives, and markets. Indirect utility from purchasing the outside good, $j = 0$, is $u_{i0mt} = \varepsilon_{i0mt}$.

We parameterize taste heterogeneity as follows: $\alpha_i = \alpha/y_i$, where α is a parameter and y_i is consumer income;[14] and $\beta_i = \beta + \Pi d_i + \Sigma\nu_i$, where d_i is a vector of observed demographics, $\nu_i \sim N(0, I)$ is a vector of individual taste differences unobserved by the econometrician (independent across consumers and independent of all observed variables), and Π and Σ are matrices of parameters. (We assume that Σ is a diagonal matrix. We also estimate a constrained specification where Σ is set to zero.)

Market shares in region m and year t are then given by

$$s_{jmt} = \int \frac{\exp(\alpha_i(p_{jt} - \text{subsidy}_{jmt}) + x_{jmt}\beta_i + \xi_{jmt})}{1 + \sum_{k \in \mathcal{C}_{mt}}(\alpha_i(p_{kt} - \text{subsidy}_{kmt}) + x_{kmt}\beta_i + \xi_{kmt})}\, dF_\theta(\alpha_i, \beta_i), \tag{1.1}$$

where F_θ is the joint distribution of (α_i, β_i) over the population of consumers in region m and model year t, indexed by the parameter vector $\theta = (\alpha, \beta, \Pi, \Sigma)$.

The product characteristics that enter x_{jmt} are technical characteristics (horsepower-weight ratio, drivetrain), proxies for size (weight, number of doors, wheelbase, footprint), electric range and battery size for electric and plug-in hybrid vehicles, fuel costs per mile, an indicator for the first year a model is available, and fixed effects for fuel type, body style (sedan, SUV, truck, etc.), make, model year, and region.[15] The demographics that enter d_i are a temperature factor capturing the frequency of extreme temperatures and an indicator

[14] This approximation to a Cobb–Douglas-style indirect utility function is taken from Berry, Levinsohn, and Pakes (1999).

[15] Fixed effects for model year and region capture differences in consumer preferences for a new car over time (for example, due to macroeconomic shocks) and across regions.

for college education.

This specification requires that manufacturers set one national price p_{jt} for each product. The final consumer-facing price may then differ across US states due to differences in government subsidies.

In our specification, consumers value a \$1 government subsidy and a \$1 reduction in price equally. This requires that consumers both know about subsidies (subsidies on electric vehicles are typically included in the price dealers advertise) and believe at the time of purchase that they will be able to take advantage of them.

We also assume that consumers make static optimization decisions; that is, they do not respond to beliefs about future product availability or future changes to product characteristics or prices. We do not model state dependence: every consumer enters the market every period, and preferences do not depend on the vehicles the consumer already owns. While these modeling assumptions abstract away from issues such as the strategic timing of vehicle purchases or the effect of owning multiple vehicles on consumer preferences, our approach is consistent with the majority of existing literature on automobile demand.

The discrete choice model we use also rules out capacity constraints or products with fixed production levels, which would induce unobserved variation in consumer choice sets as products become unavailable to late-arriving consumers. Our method is thus imperfect for Tesla, which used waitlists to manage the combination of high demand and production delays (like the second generation Toyota Prius, as documented in Sallee (2011a)). We do not observe waitlist entries, so our estimates subsume this process in the unobserved characteristic.

Pricing. We assume prices form a Nash equilibrium of a Bertrand game among multiproduct firms, who set national prices to maximize model-year profits.[16] We build on earlier models of the US auto industry by using market shares for separate geographic regions

[16] See Section 1.2.2 for evidence that national pricing is an appropriate model of the industry.

within the US, and explicitly modeling the effect of state and national regulations on the manufacturers' pricing decision. Our method assumes that marginal costs are the same across regions, so that the profit from a selling a vehicle only varies geographically due to differences in regulation. We assume that marginal costs do not depend on quantity, which rules out capacity constraints.

Consider a firm f with product set $\mathcal{J}_{ft}$. For each product $j \in \mathcal{J}$, the firm observes marginal cost mc_{jt} and the value of regulatory credits in each region v_{jmt}, then chooses its price p_{jt}. (We define v_{jmt} in the next section.) Let p_t be the vector of prices in year t. The firm's problem is

$$\max_{\{p_{jt}\}_{j\in\mathcal{F}}} \sum_{j\in\mathcal{J}_{ft}} \sum_{m\in\mathcal{M}} (p_{jt} + v_{jmt} - mc_{jt}) s_{jmt}(p_t) M_{mt}, \tag{1.2}$$

where M_{mt} is the market size in region m in year t. The firm's first order condition with respect to p_{jt} is

$$0 = \sum_{m\in\mathcal{M}} \left(s_{jmt} + \sum_{k\in\mathcal{J}_{ft}} (p_{kt} + v_{kmt} - mc_{kt}) \frac{\partial s_{kmt}}{\partial p_{jt}} \right) M_{mt}. \tag{1.3}$$

Under these assumptions, the equilibrium price vector p_t is the joint solution to all firms' first order conditions.

Regulation. Selling an additional vehicle changes the firm's credit holdings under various national and state regulations, which we model using the v_{jmt} term in the profit function. Specifically, selling an alternative fuel vehicle may earn credits under the federal greenhouse gas (GHG) regulation and the state-level ZEV regulation, while selling a gas vehicle may increase the firm's total regulatory requirement.[17] When the regulations are binding upon firms (that is, when credits have non-zero value) changes in credit holdings will enter the firm's per-period profit function.

In order to value this credit effect in firm profits, we take advantage of the data on credit prices described in Section 1.4.3. We treat firms as price takers in the market for regulatory

[17] As described in section 1.4.3, we assume that CAFE standards were not binding in this period because of federal greenhouse gas regulation.

credits, and we assume that accessing this market is costless for all firms. Therefore, in equilibrium, all firms value the marginal credit at its market price.[18] If firms are certain about future credit prices, the mandate will operate as a producer subsidy and tax; with uncertainty, the results will generally differ (as in Weitzman (1974) and related literature). We only consider scenarios in which firms are certain about credit prices. From the details of the GHG and ZEV regulations we obtain the net change in credits from selling one unit of product j ($c_{jmt,GHG}$ and $c_{jmt,ZEV}$, respectively). To determine $c_{jmt,GHG}$, we take the difference between the vehicle's statutory emissions and its regulatory target (a year-specific function of vehicle footprint) and multiply by statutory expected vehicle miles traveled and multipliers for electric or plug-in hybrid vehicles. For electric vehicles sold in ZEV states, $c_{jmt,ZEV}$ is the number of credits earned (generally a function of electric range). For non-EVs sold by large manufacturers in ZEV states, $c_{jmt,ZEV}$ is the additional number of credits that must be surrendered when a non-EV is sold.[19]

Now let $r = (r_{mt,GHG}, r_{mt,ZEV})$ be the vector of credit prices on the two credit markets. (When a regulation does not apply, such as ZEV outside of states with the mandate, we set this to zero.) Then the net value of regulatory credits earned from selling an additional unit of j in region m and model year t is

$$v_{jmt} \equiv c_{jmt,GHG}\, r_{mt,GHG} + c_{jmt,ZEV}\, r_{mt,ZEV}.$$

Under counterfactual simulations where we implement the ZEV program as a set of consumer subsidies and taxes, we set $c_{jmt,ZEV}\, r_{mt,ZEV} = 0$, making the value of regulatory credits

$$\tilde{v}_{jmt} = c_{jmt,GHG}\, r_{mt,GHG}.$$

[18] In reality, the ZEV credit trading market is characterized by one large seller (Tesla, with 83% of sales) and a handful of buyers, including Toyota (37%), Ford (20%) and Fiat Chrysler (14%). We abstract away from the modeling issues that market power may present.

[19] In our model, the number of ZEV credits a regulator requires the manufacturer to surrender is a fixed percentage of its non-EV sales that year in ZEV states. In practice, the number is a fixed percentage of a moving average of California non-EV sales in prior years, as detailed in Section A.1.

1.3.2 Pass-through of supply and demand policy

To illustrate the consequences of uniform pricing for supply- and demand-side policy, consider the pass-through of a small increase in the consumer subsidy or supply-side incentive for a specific product made by a single-product firm. We hold the prices of other products fixed; the result we get can be interpreted as pass-through under monopoly or as the first step in an iterated best response process. Label the product j. The firm first order condition (1.3) reduces to

$$0 = \sum_{m \in \mathcal{M}} \left(s_{jmt} + (p_{jt} + v_{jmt} - mc_{jt}) \frac{\partial s_{jmt}}{\partial p_{jt}} \right) M_{mt}.$$

Suppose the supply-side incentive in market 0, v_{j0t}, increases by a small amount Δ. To a first order approximation, the effect on price p_{jt} is Δ times

$$\frac{\partial p_{jt}}{\partial v_{j0t}} = -M_{0t} \frac{\partial s_{j0t}}{\partial p_{jt}} \left[\sum_{m \in \mathcal{M}} \left(2 \frac{\partial s_{jmt}}{\partial p_{jt}} + (p_{jt} + v_{jmt} - mc_{jt}) \frac{\partial^2 s_{jmt}}{\partial p_{jt}^2} \right) M_{mt} \right]^{-1}.$$

Suppose instead that the demand-side incentive, subsidy_{j0t}, increases by Δ. To a first order approximation, the effect on the national price p_{jt} is Δ times

$$\begin{aligned}
\frac{\partial p_{jt}}{\partial \text{subsidy}_{j0t}} &= -M_{0t} \left(\frac{\partial s_{j0t}}{\partial \text{subsidy}_{j0t}} + (p_{jt} + v_{j0t} - mc_{jt}) \frac{\partial^2 s_{j0t}}{\partial p_{jt} \partial \text{subsidy}_{j0t}} \right) \times \\
&\quad \times \left[\sum_{m \in \mathcal{M}} \left(2 \frac{\partial s_{jmt}}{\partial p_{jt}} + (p_{jt} + v_{jmt} - mc_{jt}) \frac{\partial^2 s_{jmt}}{\partial p_{jt}^2} \right) M_{mt} \right]^{-1} \\
&= M_{0t} \left(\frac{\partial s_{j0t}}{\partial p_{jt}} + (p_{jt} + v_{j0t} - mc_{jt}) \frac{\partial^2 s_{j0t}}{\partial p_{jt}^2} \right) \times \\
&\quad \times \left[\sum_{m \in \mathcal{M}} \left(2 \frac{\partial s_{jmt}}{\partial p_{jt}} + (p_{jt} + v_{jmt} - mc_{jt}) \frac{\partial^2 s_{jmt}}{\partial p_{jt}^2} \right) M_{mt} \right]^{-1},
\end{aligned}$$

and the effect on the consumer-facing price in market 0 is $\Delta \cdot \left(\frac{\partial p_{jt}}{\partial \text{subsidy}_{j0t}} - 1 \right)$.

Proposition 1 *The two policies have the same first order effect on consumer-facing prices in market*

0, $\frac{\partial p_{jt}}{\partial v_{0t}} = \frac{\partial p_{jt}}{\partial subsidy_{j0t}} - 1$, *if and only if*

$$\sum_{m \in \mathcal{M} \setminus \{0\}} \left(2\frac{\partial s_{jmt}}{\partial p_{jt}} + (p_{jt} + v_{jmt} - mc_{jt}) \frac{\partial^2 s_{jmt}}{\partial p_{jt}^2} \right) M_{mt} = 0.$$

Under the types of regularity conditions that are often assumed in these types of models of price-setting, this condition will not be met as long as the product is sold in multiple markets.

1.3.3 Social welfare

We use a social welfare function to evaluate the effects of alternative policy instruments. Total welfare combines consumer surplus, environmental effects, and producer surplus (the sum of variable profit across firms). We are interested in both how total welfare changes under alternative policies and how the distribution of surplus changes across consumers and producers. These welfare calculations depend not only on market outcomes in regulated states, but also on spillovers to other states.

When calculating producer surplus, we depart from the variable profit defined in Section 1.3.1 by excluding the value of regulatory credits. A credit is not a socially useful asset, only a mechanism for firms to internalize environmental externalities that we have already accounted for.

For the exposition that follows, we define present discounted values based on a discount factor δ and a time period $t = 1, \ldots, T$, and condition on the full product set. When taking to data, we use $\delta = 1$ and time periods corresponding to model years 2009 through 2017.

Welfare is broken into consumer surplus, producer surplus, and environmental damages. Let CS be the present discounted value of consumer surplus and let Env be the present discounted value of negative environmental damages. The present discounted value of producer surplus, denoted PS, is defined as follows. Let π_f be the present discounted value of firm f's profits, including regulatory credits:

$$\begin{aligned}\pi_f &= \sum_{t=1}^{T} \delta^t \pi_{ft} \\ &= \sum_{t=1}^{T} \delta^t \sum_{j \in \mathcal{J}_{ft}} \sum_{m \in \mathcal{M}} (p_{jt} + v_{jmt} - mc_{jt}) s_{jmt}(p_t) M_{mt}.\end{aligned}$$

Let v_f be the present discounted value of profits or losses from regulatory credits, similarly defined:

$$v_f = \sum_{t=1}^{T} \delta^t \sum_{j \in \mathcal{J}_{ft}} \sum_{m \in \mathcal{M}} v_{jmt} s_{jmt}(p_t) M_{mt}.$$

Then producer surplus is

$$PS = \sum_f \left(\pi_f - v_f\right), \tag{1.4}$$

and total welfare is

$$W = CS + Env + PS. \tag{1.5}$$

Environmental damages capture the CO_2 emissions from vehicles, and are calculated by assuming that every vehicle purchased is driven 150,000 miles over its lifetime, as in the calculations of Holland, Mansur, Muller, and Yates (2016).[20] We assume no deterioration in the efficiency of the vehicle over time. Since the emissions from electricity use are lower over most of the period than the estimates we use (Holland, Mansur, Muller, and Yates 2020), and electric vehicles are driven fewer miles per year than gas vehicles (Davis 2019; Burlig, Bushnell, Rapson, and Wolfram 2021), we overestimate the climate damages of electric vehicles.

We also ignore consumers' ability to respond to market conditions by driving existing vehicles for longer by assuming that the outside good (not purchasing a new vehicle) has zero emissions. This assumption is only accurate if consumers who choose the outside good do not drive or already own a vehicle that has not reached 150,000 miles. This approach rules out, for example, a consumer responding to price increases by choosing to drive an

[20]This number is lower than the statutory lifetime miles traveled in the GHG regulation, which is 195,264 miles for cars and 225,865 miles for light trucks.

existing car past the 150,000 mile mark.[21]

1.4 Data

To study consumer demand for new vehicles, we combine US new vehicle registrations by model and fuel type, product characteristics collected from various sources, and statistics from consumer surveys.

1.4.1 Product characteristics and sales

To measure vehicle sales, we use the universe of US new vehicle registrations in calendar years 2009 through 2017, obtained from IHS Markit (formerly R.L. Polk). This dataset contains the count of registrations for each model year, make, model, fuel type, and state. (In some cases, the dataset is further broken down by trim.) Our data contain both sales and leases; we treat both as sales. We use state-level data for the ten ZEV states, and aggregate the rest of the US into one region for our analysis.

For product characteristics, we combine trim-level data from MSN Autos, the US Environmental Protection Agency's FuelEconomy.gov dataset, and Ward's Automotive Yearbook, and supplement with additional sources as needed. MSN Autos provides manufacturer suggested retail price (MSRP) and technical specifications, including size, horsepower, weight, and battery capacity. FuelEconomy.gov provides fuel economy data and the battery range for electric vehicles and plug-in hybrids. Ward's Automotive Yearbook provides each model's production location, and additional technical specifications that we use when MSN Autos data are missing. In addition, when battery capacity is not provided, we back it out from the federal IRC 30D subsidy amount (if possible) or obtain it from news sources. Like Reynaert (2021), we summarize fuel economy using the dollar cost (at time of purchase) of driving one mile, constructed by combining fuel economy data with average fuel prices by

[21] An alternative approach that captures the intensity margin is the discrete-continuous model of Goldberg (1998).

state and year from the US Energy Information Administration.[22] (Except for fuel costs and subsidies, product characteristics are the same across states.)

Our main dataset is restricted to gasoline, flex-fuel, electric, and hybrid vehicles classified as cars or light trucks (with a gross vehicle weight rating under 8500 pounds) whose base version has a MSRP under $120,000. We remove products that were only sold to fleet or government buyers. We assume that a product without any sales in a given market and year was not offered in that market.[23] We restrict to model years 2009 through 2017. (Because we lack sales data before January 2009, we assume that the sales of model year 2009 vehicles in calendar year 2009 are a representative sample of all model year 2009 sales.)

We are interested in product differentiation that is technologically significant and relevant to consumers, not small differences between trims of the same model. Therefore, we aggregate products to the level of model year, make, model, technology type (electric, plug-in hybrid, hybrid, gas) and battery size; within each group, we use the characteristics of the lowest-priced trim that accounted for at least 1% of the group's national sales total.

For each vehicle's price to consumers, we use MSRP minus a measure of average manufacturer rebates derived from Automotive News[24] and federal and state government incentives. In the period we study, state government incentives generally went only to electric vehicles, plug-in hybrids, and hybrids; the only relevant federal program was the tax credit for plug-in hybrids and electric vehicles (IRC 30D). The IRS provides the federal tax credit amount for each vehicle model, and we collect state incentive amounts from government websites and the Alternative Fuels Data Center. We use nominal dollars throughout (including for household income, described below).

Table 1.1 shows summary statistics of our product characteristics data.

---

[22]For plug-in hybrids, we weight electric and gas modes using the EPA utility factor.

[23]This contrasts with Li (2019), who uses more granular markets and thus encounters products that were offered but had zero sales.

[24]We take the maximum consumer rebate observed in each quarter and average across quarters, and ignore manufacturer rebates to dealers. This roughly approximates the findings by Busse, Silva-Risso, and Zettelmeyer (2006) of high pass through for customer rebates and low pass through for dealer rebates. The Automotive News data do not distinguish between national and regional rebates.

Table 1.1: *Summary statistics, vehicle characteristics and sales data*

	Min	Max	Mean	Wgt. mean	SD
Model year	2009.00	2017.00	2013.14	2013.52	2.61
Products per market	240.00	332.00	286.42	296.81	20.94
Inside good share	0.05	0.18	0.12	0.12	0.03
Price	7969.17	145000.00	36687.69	24948.22	21394.87
Govt subsidy	0.00	12500.00	305.02	47.33	1442.06
Doors	2.00	5.00	3.72	3.71	0.81
Weight (lbs)	1808.00	6000.00	3827.84	3649.65	821.03
Footprint (sq ft)	25.78	68.70	48.04	48.37	5.60
Wheelbase (in)	73.50	145.70	110.16	110.47	8.73
New model	0.00	1.00	0.09	0.05	0.28
HP/Weight	0.03	0.19	0.06	0.06	0.02
Displacement (cc)	0.00	8390.00	2957.67	2769.29	1242.60
Hybrid	0.00	1.00	0.09	0.03	0.28
PHEV	0.00	1.00	0.02	0.00	0.14
EV	0.00	1.00	0.02	0.00	0.14
Battery (kWh)	0.00	100.00	1.00	0.17	6.49
Electric range (mi)	0.00	315.00	3.17	0.54	21.43
Gas cost (dollars/mi)	0.00	0.33	0.13	0.12	0.05
Electric cost (dollars/mi)	0.00	0.09	0.00	0.00	0.01

Note: *Compiled from data from MSN Autos, FuelEconomy.gov, Ward's Automotive Yearbook, and IHS. Columns are the minimum, maximum, unweighted mean, sales-weighted mean, and unweighted standard deviation across products. Prices are MSRP after average manufacturer rebates. All years are model years. (Includes content supplied by R. L. Polk; Copyright © R. L. Polk, 2019. All rights reserved)*

1.4.2 Demographics

Household demographics are taken from American Community Survey microdata for 2009 to 2017.[25] Demographics from a given calendar year t are matched to our market for model year t. The variables we use include income (in nominal dollars), location (Public Use Microdata Area), and college education. We exclude households with income below $10,000. The market size we use is the number of households in each state and year, from American Community Survey 1-year estimates.[26]

We also include county-level temperature factor as an additional demographic, which proxies for consumers' perceptions of the effects of extreme weather on battery performance. To construct this measure, we use the formula from Holland, Mansur, Muller, and Yates (2016) and historical temperature data from the North American Land Data Assimilation System (via Holland, Mansur, Muller, and Yates (2016)'s replication files). The temperature factor ranges from 0 to 1; a constant temperature of 68° F yields 1, while a constant temperature of 19.4° F or 116.6° F yields 0.67. Among US counties, the temperature factor ranges from 0.825 (Cavalier County, North Dakota) to 0.993 (Orange County, California).

Micro-moments, which match consumers' observable demographics to their vehicle choices, are drawn from surveys of vehicle purchasers. Data related to consumer location (for California buyers only) are drawn from the California Vehicle Rebate Program's survey of rebate recipients, which matches the make and fuel type of a plug-in hybrid or electric vehicle to the census tract of the buyer. (We only use data before the program introduced an income requirement in March 2016.) Data on the education of electric vehicle buyers is drawn from summary tables in Xing, Leard, and Li (2021), which use the MaritzCX survey of new vehicle buyers.

[25] Data obtained via IPUMS (Ruggles et al. 2019).

[26] Because, on average, 77% of sales in model year t are made in calendar year t, we take the market size for model year 2009 (for which we are missing data from calendar year 2008) to be only 77% of US households in 2009.

1.4.3 Regulation

We estimate average ZEV credit prices by dividing the quantity of credits sold by Tesla, obtained from state regulators, by the revenue Tesla earned from those sales, as reported in quarterly filings and shareholder letters.[27] This accounts for a large part of the credit market: during this period, Tesla was the seller for 83% of the credits that were traded overall. We are unable to observe if the price paid per credit varies across buyers. We weight credits from different states equally because they were interchangeable under the travel provision. Table 1.2 shows our estimates of the prices of ZEV credits, along with Tesla's share of total credit sales in the corresponding period.

Table 1.2: *Estimated credit prices*

Window	Revenue (m)	Credits	Avg credit price	Tesla share
2010Q4—2013Q3	$166	45,617	$3,630	67%
2013Q4—2014Q3	$86	35,869	$2,400	71%
2014Q4—2015Q3	$170	87,243	$1,950	70%
2015Q4—2016Q3	$204	85,098	$2,390	92%
2016Q4—2017Q3	$120	82,584	$1,460	89%

Note: *This table shows the computation of the average price of ZEV credits sold by Tesla in each year. Tesla's revenue from credit sales comes from Tesla's quarterly reports and shareholder letters. The number of credits sold by manufacturer by year was obtained from state regulatory agencies in the ten ZEV states.*

The number of ZEV credits earned by each vehicle comes from public data from the California Air Resources Board and New Jersey Department of Environmental Protection.

We estimate the number of greenhouse gas regulation (GHG) credits earned by each vehicle using formulas from the regulation, fuel consumption from FuelEconomy.gov, and vehicle size data from MSN Autos. We assume a constant credit price of $40 per megagram of CO_2, based on estimates summarized in Leard and McConnell (2017). In order to approximate the findings of Leard and McConnell (2017) that Corporate Average Fuel Economy (CAFE) standards were not binding on automakers while the GHG regulation was in effect, we set the CAFE credit price to zero during our study period. Prior to 2012,

[27] This method has also been used by bloggers and (independently) by McConnell, Leard, and Kardos (2019).

the GHG regulation applied differently in twelve states (the ZEV states plus Pennsylvania and Washington) than in the rest of the country; given our data limits, we approximate this by modeling GHG as applying only in the ZEV states.

1.4.4 Environmental damages

We use estimates from prior literature to determine CO_2 emissions per mile for different types of vehicles, and apply a social cost of carbon of $41 per megagram. For electric vehicles, we use estimates by Holland, Mansur, Muller, and Yates (2016) of marginal emissions from electricity use by North American Electric Reliability Corporation (NERC) region, in megagrams of CO_2 per kilowatt-hour, which they estimate from 2010–2012 data. (When NERC regions do not line up with our geographical regions, we take an average, weighting each NERC region using average vehicle miles traveled from EPA MOVES.) We then convert to vehicle emissions using vehicle-specific electricity consumption per mile (in FuelEconomy.gov data). We determine CO_2 emissions per mile from gas vehicles by multiplying gasoline consumption per mile (in FuelEconomy.gov data) by 8887 grams of CO_2 per gallon (from EPA).

Section 1.3.3 describes how these emissions enter our welfare calculations.

1.5 Estimating the demand model

We estimate demand parameters using a generalized method of moments estimator following Berry, Levinsohn, and Pakes (1995), Petrin (2002), and Conlon and Gortmaker (2020). Our estimator relies only on the model of consumer demand, not on the model of pricing. The estimator simultaneously matches model-predicted to observed market shares, matches model-predicted micro-moments to observed survey data, and fits unobserved quality ξ to be uncorrelated with instrumental variables.

The micro-moments, as in Petrin (2002), recover heterogeneity in consumer tastes based on observed demographics, Π. Instrumental variables based on cost and price shifters (government subsidies, cost shifters based on production locations, and a proxy for lithium ion battery costs) identify consumer preference for price, α, in a way that accounts for the dependence of price and unobserved quality ξ. As described in Gandhi and Houde (2019), instrumental variables based on characteristics of competing products help to estimate consumer preference heterogeneity, Σ.

1.5.1 Demand estimation method

We estimate demand parameters using a demand side-only version of the generalized method of moments estimator from Petrin (2002), implemented using PyBLP (Conlon and Gortmaker 2020). This estimator relies on two sets of assumptions: the exogeneity of the instrumental variables and the assumption that the model-predicted micro-moments match empirical values from survey data. Notation is as follows. For product j in region m and model year t, let z_{jmt} be a vector of instruments; stack the row vectors z'_{jmt} to form a matrix Z. Let $\xi(\theta)$ be the vector ξ which solves the system of market share equations (1.1) given parameters θ.[28] Defining $\mathbf{G}_1(\theta) = Z'\xi(\theta)$, the key assumption is that $E[\mathbf{G}_1(\theta)] = 0$. The micro-moments simply match the model's predictions of the moments from Section 1.5.3 to the values from data; defining $\mathbf{G}_2(\theta)$ to be the difference between model-implied moments and moments from data, the key assumption is that $E[\mathbf{G}_2(\theta)] = 0$.

Define $\mathbf{G}(\theta) = (\mathbf{G}_1(\theta), \mathbf{G}_2(\theta))$, and denote the sample analog of $\mathbf{G}$ by $\hat{\mathbf{G}}$. The estimates are then computed by solving

$$\min_{\theta} \hat{\mathbf{G}}(\theta)' \, W \, \hat{\mathbf{G}}(\theta),$$

where W is a positive definite weight matrix.

We adopt many of the best practices in PyBLP recommended in Conlon and Gortmaker (2020). We solve the system of equations (1.1) using SQUAREM with the Berry, Levinsohn, and Pakes (1995) contraction map and tolerance 10^{-14}. We compute the integral over (d_i, ν_i) using Monte Carlo integration with 200 draws per market. The weight matrix is determined by two-step GMM, clustering observations at the make-model level, and the minimization problem is solved using L-BFGS-B.

[28] We assume that markets are large enough that observed market shares equal choice probabilities.

Standard errors are computed by the GMM formula, clustering observations at the make-model level to allow for within-model correlation in unobserved quality across regions, across time, and across fuel type and battery size variants.

1.5.2 Instrumental variables

We assume that characteristics of all products are mean-independent of all products' quality shocks ξ.[29] We implement this assumption as follows: for each product j, region m, and model year t, the product's quality ξ_{jmt} is restricted to be mean-independent of both the product's own characteristics x_{jmt} and the local differentiation instruments of Gandhi and Houde (2019),[30] which measure the availability of close substitutes along observable dimensions. Specifically, for each product j and a subset of characteristics k,[31] the local differentiation instruments are the count of products j' for which $x_{j'k}$ is within one standard deviation of x_{jk}, separated into products by the same firm as j and products by other firms. As described by Gandhi and Houde (2019), the primary role of the differentiation instruments is to estimate random coefficient parameters.

We assume that government subsidies are exogenous.[32] Subsidy levels are deterministic functions of vehicle characteristics, typically range or battery capacity; though the exact function varies over time, this is a response to changes in funding and not a response to the unobserved quality of particular vehicles. Subsidies do not depend on vehicle price, except some programs that exclude vehicles whose MSRP is above a cutoff level. (In practice, products are rarely priced close to the cutoff.)

[29] According to industry experts, the technical characteristics of a vehicle are determined early in the design process. The electric range of non-native first-generation electric vehicles, for example, was typically determined by space constraints and the energy density of the battery chemistry used. See "What a teardown of the latest electric vehicles reveals about the future of mass-market EVs" (Antoine Chatelain, Mauro Erriquez, Pierre-Yves Moulière, and Philip Schäfer, McKinsey, 3/21/18).

[30] We use PyBLP (Conlon and Gortmaker 2020) to construct the local differentiation instruments.

[31] We use all characteristics except make, model year, and region fixed effects.

[32] This instrument is also used by Li (2019), which also features a national price with local variation due to exogenous government subsidies.

For additional cost shifter instruments, we use the average manufacturing wage in the area of production (the 2009–17 average from the Quarterly Census of Employment and Wages),[33] and (for vehicles with a lithium ion battery) a proxy for the cost of the battery obtained by multiplying the battery size (in kilowatt-hours) and BloombergNEF's measure of the industry-wide average battery pack price (in dollars per kilowatt-hour).

1.5.3 Micro-moments

We use two sets of micro-moments that help to identify demographic parameters in Π. The first matches electric vehicle purchases to county-level climate, using data from the California rebate program prior to 2016. We calculate a temperature factor, which proxies for consumers' perceptions of the effects of extreme weather on battery performance, for each county. Then, for each model year (inferred based on purchase date) we calculate the mean temperature factor among buyers of electric vehicles. The second micro-moment, taken from Xing, Leard, and Li (2021), is the statistic that across model years 2010 through 2014, 81% of electric vehicle purchasers nationally had a college degree.

1.5.4 Demand estimates

Table 1.3 shows the estimated linear (β) and nonlinear (Π, Σ) parameters (except for the magnitudes of fixed effects). All else equal, consumers prefer vehicles that are gas-efficient, but are relatively insensitive to electric efficiency. Consumers on average dislike hybrids and electric vehicles, but have heterogeneous preferences.

Because we obtain estimates of Σ that are close to zero, we also estimate a constrained model that sets $\Sigma = 0$, so that all heterogeneity across consumers comes from observed demographics. We use the constrained model for the estimates and counterfactuals that follow.

The estimates imply own-price elasticities in line with prior literature for electric vehicles, but not non-electric vehicles. In the constrained model, the estimated sales-weighted average

[33]This instrument is also used by Wollmann (2018).

Table 1.3: *Estimates of demand parameters*

	Demographics		Full	
	Estimate	SE	Estimate	SE
Linear parameters (β)				
Displacement (L)	0.24	0.07	0.23	0.07
Doors	-0.21	0.09	-0.20	0.09
Wheelbase	-0.03	0.02	-0.03	0.02
HP/Weight	-4.17	5.17	-4.29	5.23
Weight (tons)	-0.60	0.40	-0.65	0.41
Gas cost (dollars/mi)	-14.82	2.13	-14.46	2.16
Electric cost (dollars/mi)	-41.04	1.88	-0.41	2.40
PHEV Range	0.01	0.02	0.01	0.02
EV Range	-0.02	0.02	-0.01	0.02
Battery (kWh)	0.06	0.05	0.02	0.06
Footprint	0.12	0.04	0.12	0.04
New model	-0.01	0.07	-0.01	0.07
Hybrid	-2.33	0.18	-2.33	0.18
PHEV	-1.73	0.40	-2.55	0.41
EV	-6.42	8.91	-10.62	10.02
Coupe	0.67	0.24	0.66	0.25
Hatchback	0.56	0.33	0.57	0.33
SUV	1.88	0.31	1.87	0.32
Sedan	1.38	0.30	1.37	0.31
Truck	1.00	0.39	0.99	0.40
Van	1.12	0.42	1.06	0.44
Wagon	0.94	0.39	0.93	0.39
FWD	0.46	0.17	0.47	0.17
RWD	0.06	0.14	0.09	0.15
Demographics (Π)				
(price-subsidy)/income	-8.27	3.29	-8.09	3.36
EV*College	1.08	0.22	1.08	0.47
EV*Temp. Factor	4.75	9.52	7.18	10.61
Unobserved heterogeneity (Σ)				
(constant)			0.00	2.92
EV			0.00	4.62

Note: *Estimates from random coefficients logit demand system, except for magnitudes of fixed effects (on make, model year, and region). The coefficient on characteristic k for consumer i is $\beta_k + \Pi_k d_i + \Sigma_k v_i$, where d_i is a vector of demographics and v_i is unobserved heterogeneity. The specification labeled Demographics sets $\Sigma = 0$ so that heterogeneity only comes from demographic variation. The specification labeled Full allows Σ to be nonzero. Standard errors are clustered at the make-model level.*

own-price elasticity of demand for electric vehicles is -2.22; estimates from other studies using US data (surveyed in Cole, Droste, Knittel, Li, and Stock (2021)) range from -1.0 to -2.7. The estimated average own-price elasticity for non-electric vehicles is -1.69, much more inelastic than other estimates from the US market. For example, Beresteanu and Li (2011) finds an average own-price elasticity of -8.4 in the 1999–2006 period, and Grieco, Murry, and Yurukoglu (2021) finds own-price elasticities between -6.5 and -9.4 (depending on income group) in 2018.

1.6 Estimating marginal costs

We estimate marginal costs by solving the Nash–Bertrand first order conditions (1.3) given demand parameters. This method adapts existing techniques for sequential demand and supply estimation (e.g., Nevo (2001)) to accommodate national pricing and state variation in regulation, and relies only on prices, quantities, and demand elasticities.

According to industry sources, electric vehicles have higher marginal costs than gas-powered vehicles, but the difference is falling throughout the period. The marginal cost of a vehicle consists primarily of the labor and parts involved in assembly. Although electric vehicles have simpler powertrains than gas-powered vehicles, the resulting savings are dwarfed by the per-vehicle cost of the battery pack.[34] This difference changed over time, however, as cost of a battery pack of a fixed size fell rapidly through this period. Figure 1.1 shows estimates of industry averages from BloombergNEF, which show a decline of 80% (in real terms) from 2011 to 2018.[35]

The demand estimates and the additional assumption of the pricing model together give estimates of marginal cost. Figure 1.2 shows trends in sales-weighted average prices and estimated marginal costs for non-Tesla electric vehicles. Marginal costs fall until 2015, and

[34]See "Making electric vehicles profitable" (Yeon Baik, Russell Hensley, Patrick Hertzke, and Stefan Knupfer, McKinsey, 3/8/19), which breaks down the difference in average cost.

[35]Estimates from other sources are roughly similar; see Nykvist and Nilsson (2015) and Ziegler and Trancik (2021) for comparisons.

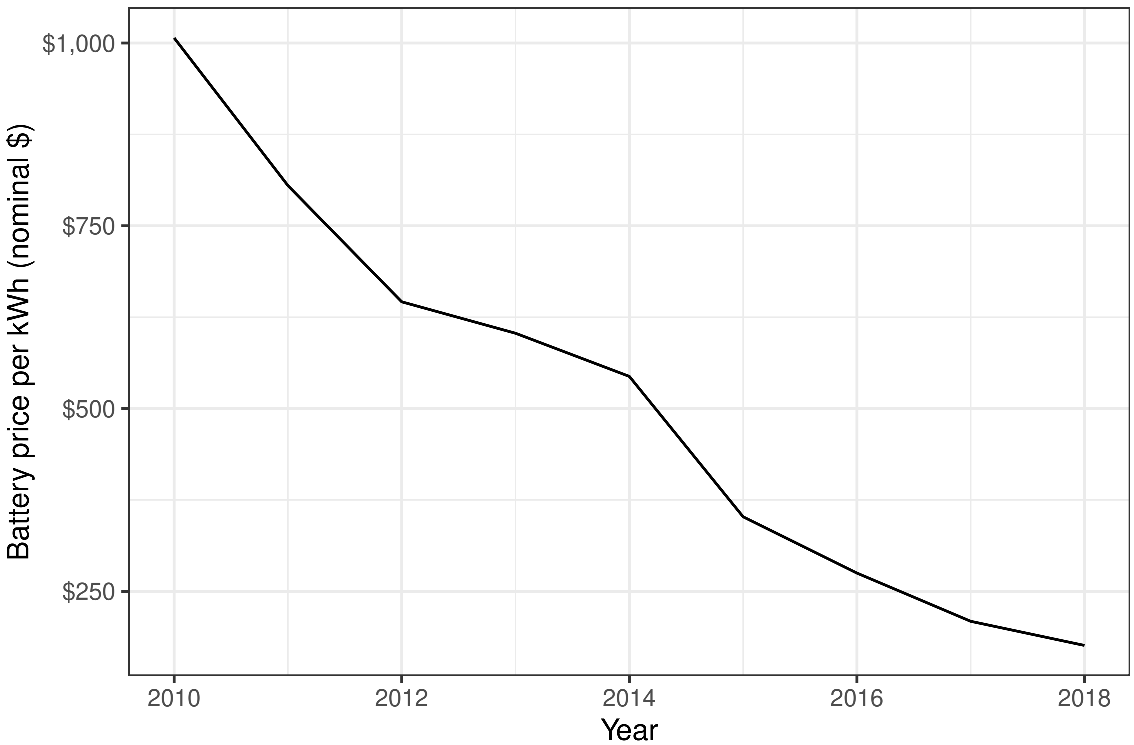

Figure 1.1: *Lithium-ion battery pack prices, 2010–18, from BloombergNEF surveys of the EV industry*

Note: *Volume-weighted averages from "A Behind the Scenes Take on Lithium-ion Battery Prices" (Logan Goldie-Scot, BloombergNEF, 3/5/19), converted back from 2018 dollars to nominal dollars.*

then rise; the rise is mainly attributable to increases in the battery pack size in later-model EVs. The trend of declining marginal cost is consistent with the falling battery prices in this period.

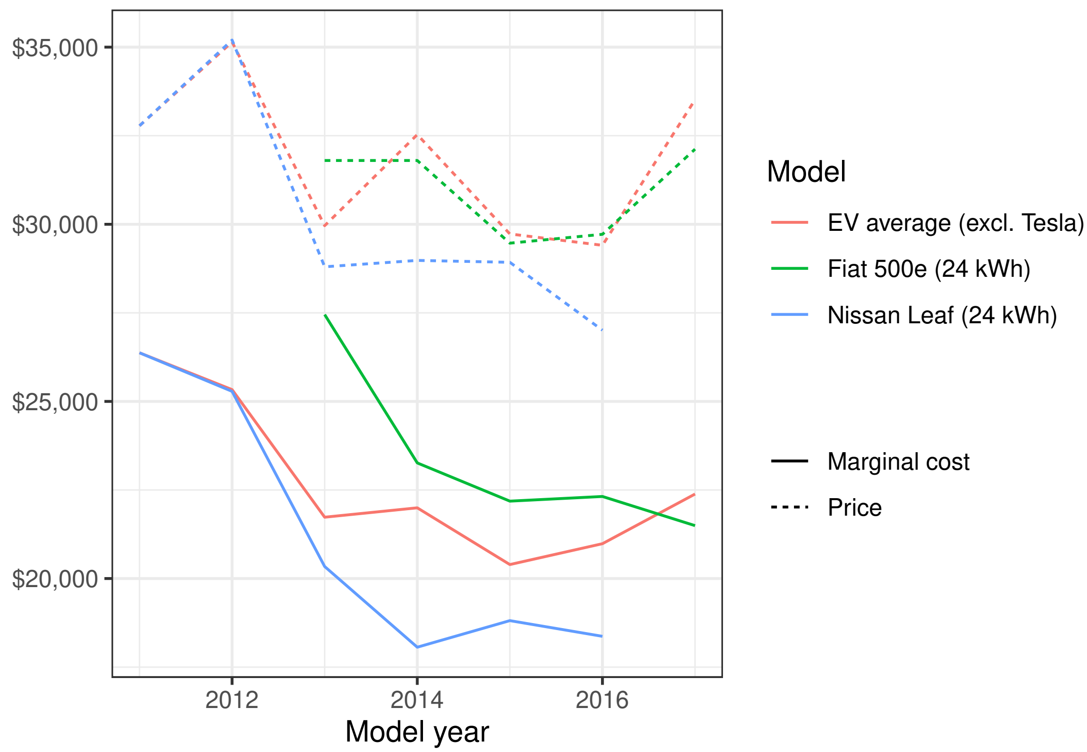

Figure 1.2: *Prices and estimated marginal costs, selected EVs*

Note: *Prices are MSRP after average manufacturer rebate (from MSN Autos and Automotive News). Marginal costs are estimated from our demand system and national Nash–Bertrand pricing. EV average is sales-weighted.*

1.7 Counterfactual demand-side policy

What are the welfare effects of achieving the same electric vehicle quantity target by a demand-side policy? We simulate the effects of replacing the ZEV mandate with a counterfactual policy that only uses consumer incentives.[36] In particular, we eliminate ZEV

[36]This type of policy is known as a 'feebate', and is commonly used to control greenhouse gas emissions from cars. A greenhouse gas-based feebate system was proposed in California in 2008, but not passed (Durrmeyer

credits on the producer side and add (1) a consumer subsidy for electric vehicles in the ten regulated states, plus (2) a consumer tax on non-electric vehicles in the ten regulated states. By simulating the new equilibria of pricing, we measure the welfare effects of changing to the counterfactual policy.

Under national pricing, a subsidy for consumers in one region will introduce a difference between the consumer price in that region and the consumer price elsewhere, while a subsidy for producers in a region will not. A firm's incentive to adjust prices in response to the regional subsidy is dampened by the demand response among consumers in other regions.

We set the specific level of the consumer subsidy and tax by targeting two outcomes: electric vehicle quantity and budget balance. First, the policy must achieve the same total electric vehicle sales each year across the ten regulated states as observed data. Second, the policy must be budget-balanced: total expenditure on the subsidy must equal total collection through the tax, across all of the regulated states. We adopt total electric vehicle sales as the target because it is easily measured and explicitly mentioned as a goal in regulator reports.[37]

To account for the effect of the subsidy and tax on consumer prices, we recompute the Nash–Bertrand price equilibrium by solving equations (1.3).[38] This computation requires the subsidy and tax to be set simultaneously with prices, and the regulator to have the same amount of information as the firms.

The structure of the counterfactual policy mimics that of the ZEV mandate as closely as possible, but minor differences remain. The subsidy amount follows the ZEV formula for credits per vehicle, generating a larger subsidy for vehicles with a longer battery range. Unlike the ZEV mandate's quota, which only applies to large automakers, the consumer tax

---

and Samano 2018).

[37] As in Durrmeyer and Samano (2018), we wish to compare total surplus under two different policy instruments, holding fixed the outcome targeted by the baseline policy. This approach allows us to abstract away from the relative weights that the regulator places on the stated policy goals versus total surplus.

[38] We use SciPy's implementation of Steffensen's Method with Aitken's Δ^2 convergence acceleration, applied to the sequence generated by iterating the pricing equations. Conlon and Gortmaker (2020) document that this sequence does not always converge; in our case, it does.

applies to all non-EVs sold in the ten regulated states. In addition, while the ZEV mandate allows firms to smooth out the policy over multiple years using credit banking, there is no similar mechanism to smooth out the consumer subsidy.

Formally, the setup is as follows. The policymaker controls the EV subsidy τ_t^{EV} and the non-EV tax $\tau_t^{\text{non-EV}}$ for each model year t. For product j in region m and model year t, let q^0_{jmt} be the quantity sold in the data. Let $e_j = 1$ if product j is an electric vehicle and 0 otherwise, let c_j be a subsidy multiplier (equal to the number of credits j earns under the ZEV mandate) and let $z_m = 1$ if region m has the regulation and 0 otherwise. The net consumer subsidy for purchasing j in m and t (in addition to existing subsidy programs) is $z_m(\tau_t^{\text{EV}} c_j e_j - \tau_t^{\text{non-EV}}(1 - e_j))$. Let $q_{jmt}(\tau_t^{\text{EV}}, \tau_t^{\text{non-EV}})$ be the quantity sold under the demand-side policy, which depends implicitly on price adjustment and requires full knowledge of the demand model. The sales constraint is therefore

$$\sum_m z_m \sum_j e_j (q_{jmt}(\tau_t^{\text{EV}}, \tau_t^{\text{non-EV}}) - q^0_{jmt}) = 0,$$

the budget balance constraint is

$$\sum_m z_m \sum_j (\tau_t^{\text{EV}} c_j e_j - \tau_t^{\text{non-EV}}(1 - e_j)) q_{jmt}(\tau_t^{\text{EV}}, \tau_t^{\text{non-EV}}) = 0,$$

and the policymaker chooses $(\tau_t^{\text{EV}}, \tau_t^{\text{non-EV}})$ to solve this system of equations.

1.7.1 Results and discussion

We begin by measuring the effects of the demand-side policy if the product set is held fixed. The required consumer subsidy is larger than the value of a corresponding ZEV credit, and relative prices adjust in a way that lowers consumer surplus and raises total variable profit.

Table 1.4 shows the magnitude of the consumer subsidy and tax that implement this policy. The subsidy amount varies from \$2,300–\$3,600 before range multipliers, which is substantial and larger than the ZEV credit prices in Table 1.2. This amount translates to \$7,000–\$10,700 per vehicle for a Nissan Leaf. (Because of the small market share of electric vehicles, the tax on non-EVs is under $150 per vehicle.) This additional subsidy is on top of the $7,500 federal tax credit and existing state rebates ($2,500 in California).

Table 1.4: *Counterfactual consumer subsidy and tax amounts*

Model year	Subsidy	Tax
2009	$0	$0
2010	$0	$0
2011	$3,566	$17
2012	$2,616	$19
2013	$3,697	$83
2014	$2,884	$68
2015	$2,356	$76
2016	$2,993	$126
2017	$2,316	$110

Note: *This table shows the magnitude of consumer subsidy and tax, within the ten regulated states only, that achieves the same EV sales within the regulated states each year when the ZEV mandate is removed. (Greenhouse gas credits remain with prices unchanged.) The consumer subsidy shown is the amount before range multipliers, and is comparable to the ZEV credit price in Table 1.2. The tax per non-EV is set by constraining the subsidy outflows and tax inflows to balance. Firms respond by resetting prices in Nash–Bertrand equilibrium.*

The main effect of the counterfactual demand-side policy is to raise EV prices outside the regulated states, as the ZEV mandate no longer holds down national prices. In the ZEV states, net consumer prices after subsidies are higher for the products that are only sold in ZEV states, and lower for the products that are sold nationally. Price comparisons for EVs, in California and the non-ZEV states, are shown in Figure 1.3.

Without any change in the product set, the price changes resulting from the switch to a demand-side policy result in $1.4 billion lower consumer surplus and $570 million higher producer surplus.

1.8 Conclusion

We examine the effects of the ZEV mandate, an influential state-level supply-side environmental policy in early generations of the US electric vehicle market. Because of the interaction between national pricing and regional policy variation, the mandate generated

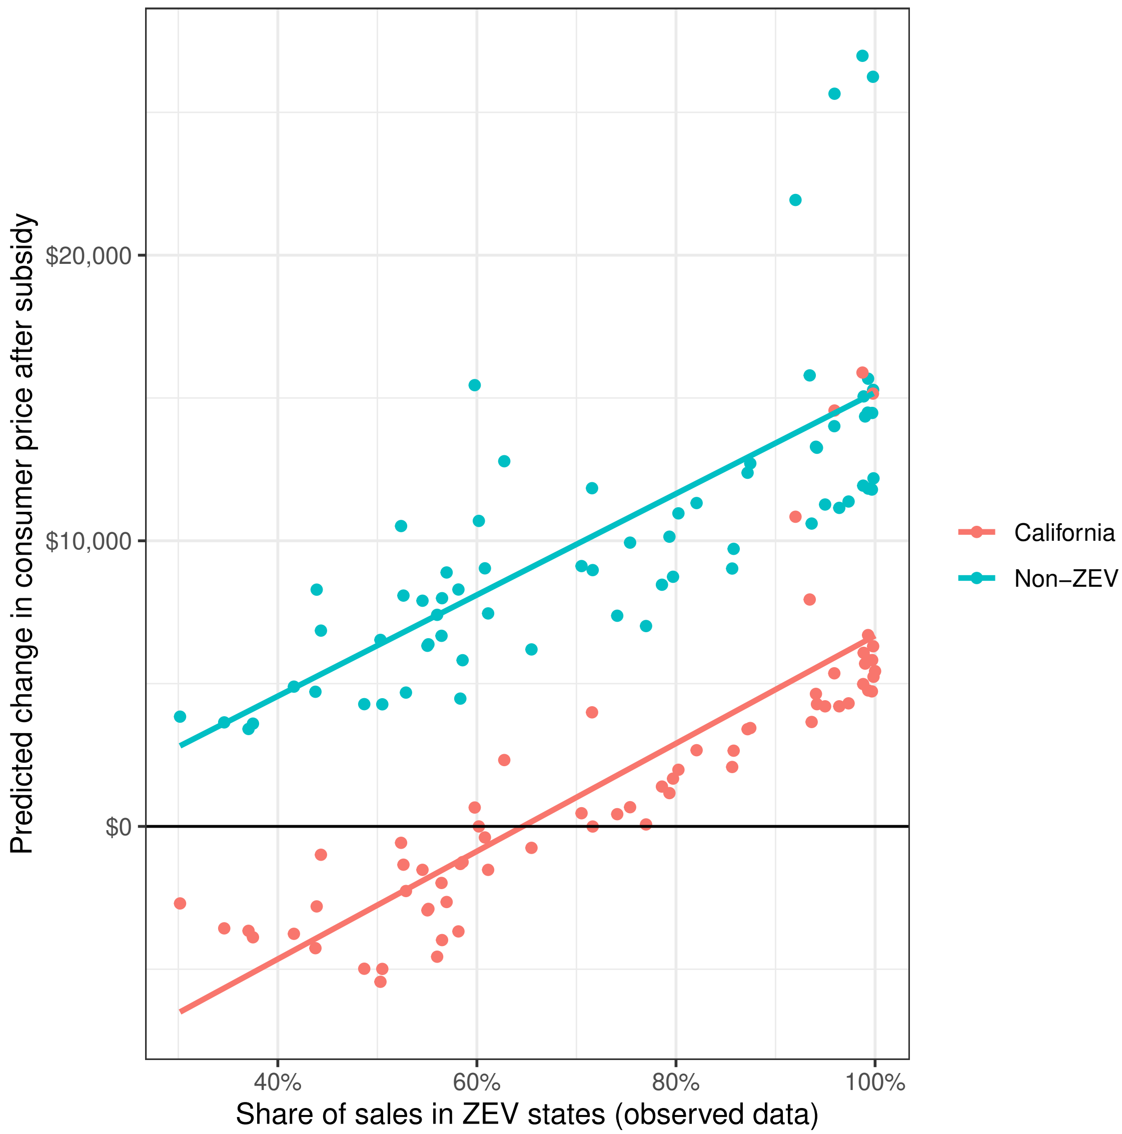

Figure 1.3: *Predicted change in consumer price (including subsidies) for EVs, California and outside regulated states.*

Note: *For all electric vehicles, this plot shows the predicted change in consumer price (national price minus subsidies) from switching from the ZEV mandate to the counterfactual demand-pull policy, for California and the rest-of-country region only. Products are arranged by the share of their sales that are in the ten regulated states in observed data (under the ZEV mandate). Consumer prices rise for all EVs outside the ten regulated states; consumer prices in California rise for products that are predominantly sold in regulated states, but fall for products that are sold nationally.*

higher consumer surplus and lower producer surplus than a comparable demand-side policy would have. These findings have consequences for future state policies to encourage the sales of clean products in their states in the context of nationally standardized markets. Though electric vehicle characteristics, costs, and quantities have evolved significantly since the period we study, our findings may also have implications for the welfare consequences of current and future electric vehicle policy.

Chapter 2

When Should Policy Encourage Product Variety? Evidence from Electric Vehicles[1]

2.1 Introduction

Starting in the 2010s, environmental policy has increasingly focused on accelerating large-scale technological efforts to replace carbon-emitting technologies with new non-carbon alternatives. This change, often called the energy transition, entails significant investment to develop and deploy new products.[2] Through their choice of policy design, regulators may influence the number of new products introduced, and hence the level of product variety, in these nascent markets.

In this paper we ask how policymakers should account for short-run variety effects when evaluating policies, and who gains and who loses from policy-induced increases in variety. We develop simple tools to study equilibrium product variety under alternative policy

[1] Co-authored with Sarah Armitage.

[2] Such policies have been studied by, among many others, Acemoglu, Akcigit, Hanley, and Kerr (2016) with a general framework, Stock and Stuart (2021) on the power sector and Holland, Mansur, and Yates (2021) and Cole, Droste, Knittel, Li, and Stock (2021) on the transition to a fully-electric vehicle fleet.

environments, and apply those tools to study the effects of an influential technology-forcing policy.

As in Chapter 1, we examine this question in the context of the electric vehicle industry in its first generation, between 2009 and 2017, and study the effects of the zero-emission vehicle (ZEV) mandate adopted by California and nine other states. During this period, 18 electric vehicle models were introduced in the US, and policymakers and some industry observers credited the ZEV mandate for much of this entry.[3] The electric vehicle industry is a useful case study because clean products are substitutes for existing, dirty products; are made by many of the same firms; and compete in differentiated products oligopoly. Studying the first generation of electric vehicles is particularly useful because all products are the result of research and development programs from the early 2000s, which occurred simultaneously across manufacturers, and because the design choice to build an electric vehicle on an existing gas vehicle platform provides an indication that a product that may have been introduced as a direct consequence of the regulation.

Using a model of product entry and the estimated model of demand and supply from Chapter 1, we examine how product variety in the electric vehicle market would vary if the ZEV mandate were eliminated or changed. We measure the short-run effects of policy-induced product variety on consumer choices and the resulting effects on consumer surplus, producer surplus, and the environment. Using these estimates, we augment the counterfactual simulation of Chapter 1, which only allows prices and quantities to change when the ZEV mandate is replaced by a consumer subsidy and tax, to include changes in the equilibrium product set.

Our study of product entry builds on the established theoretical result that imperfectly competitive product markets can feature more or fewer products than the social optimum, because firm gains from product introduction come partly at the expense of other firms (the business stealing effect of Mankiw and Whinston (1986)) and do not include the

[3]See "Automakers question Calif. zero-emission mandate as feds reassess mpg rules" (Eric Kulisch, Automotive News, 12/12/17). More generally, the ZEV mandate is often described as a technology-forcing policy (Collantes and Sperling 2008; Vergis and Mehta 2012; McConnell and Leard 2021).

inframarginal consumer surplus from added product variety (Spence 1976). Policies that subsidize consumers or producers based on quantity will exacerbate both effects, so their net result is theoretically ambiguous and requires empirical analysis.

We begin by showing descriptive evidence that the ZEV mandate has induced product entry, by identifying a set of electric vehicle models that are sold primarily in states with the mandate and not in other states. We document that electric vehicle models that used existing vehicle platforms were sold in low quantities and almost entirely in states with the mandate; using existing gas vehicle platforms was a design choice which lowered manufacturers' entry costs and therefore the sunk costs of complying with the regulation. By contrast, most models that were designed from the ground up as an electric vehicle sold in larger quantities across the US. We also show descriptive evidence that product variety has benefited middle-income consumers more than high-income consumers. In California, the state with the greatest electric vehicle product availability, electric vehicle models designed on existing vehicle platforms were purchased primarily by consumers in lower- and middle-income areas, while models that were designed from the ground up were purchased across the income distribution (and the higher-end models mainly by consumers in high-income areas).

Next, we construct a model of product entry that allows us to predict firm behavior under alternative policies and evaluate the welfare consequences of changes in the product set. We model the first generation of electric vehicles as a game among firms, who simultaneously choose whether to introduce an electric vehicle, pay an upfront cost, and then collect profits from sales over several periods. An additional term, which firms collect at the end of the study period, captures the long-run benefits that result from the development of a first-generation electric vehicle.

Taking the observed pattern of electric vehicle entry as an equilibrium of this static entry model, we combine accounting estimates of the upfront cost of entry with the estimated model of demand and supply from Chapter 1 to estimate a lower bound on long-run benefits. In order to explain observed entry, long-run benefits need to be large: at least three-quarters of entry cost in our preferred specification.

While simplified, our static entry model allows us to shed light on policy-induced product entry in an important new technology market, in which the number of new products is small and useful cross-market variation is absent. While we cannot definitively predict which vehicles would have been absent without the ZEV program, we select a set of possible entry equilibria that are consistent with the observed data.

We use the model of product entry together with our demand and supply estimates to address three questions about the welfare effects of product variety in the early electric vehicles market. First, we ask whether the ZEV mandate encourages welfare-enhancing or welfare-reducing product entry. We compare the private benefits and social welfare effects of product entry, accounting for three types of externalities not internalized by individual firms making product entry decisions: the effect on competitors' profits, consumer benefits from variety, and environmental effects. We find that the social benefits of product entry for electric vehicles are positive and substantial, as the gain from inframarginal consumers' tastes for variety exceeds the negative effect on competitor profits and dwarfs the short-run environmental effects. By encouraging electric vehicle sales, the mandate makes individual firm incentives better aligned with these social benefits and allows firms to capture up to half of these social benefits. This exercise gives insight into the direction of the welfare effect of the increased product variety induced by the ZEV program.

Second, to shed light on the magnitudes of welfare changes, we apply the product entry model and estimates to ask which electric vehicles would not have entered without the ZEV mandate, and what the resulting effects would be on quantities and welfare. We find that although consumer surplus and electric vehicle sales would be substantially reduced by the loss of all the vehicles that were designed on existing gas vehicle platforms, such an outcome is not an equilibrium under the estimated magnitude of long-run benefits. Instead, the equilibria that are plausible under our estimates feature modest negative effects on consumer surplus and electric vehicle sales and smaller positive effects on producer surplus, driven by avoided entry costs. For this reason, we conclude that the ZEV mandate encouraged welfare-enhancing product entry, in contrast to industry observers' claims that the ZEV program encouraged wasteful entry.

Third, we use our model estimates to study the effects of alternative policy design on product variety. We revisit the counterfactual demand-side policy from Chapter 1, which replaces the ZEV mandate with a consumer subsidy and tax; for reasons explored in Chapter 1, this demand-side policy changes markups over marginal cost and therefore product entry incentives. We simulate the effect of this alternative policy on product entry, accounting for the possibility of multiple equilibria. Across a set of plausible scenarios, firms capture less of the social benefits of product entry under the demand-side policy. Allowing product entry to respond exacerbates the policy's negative effect on consumer surplus by reducing product variety. Across entry scenarios that are compatible with our estimated product entry model, consumer surplus is between \$1.8 billion and \$1.4 billion lower under the demand-side policy. The effect on producer surplus is ambiguous, because of avoided upfront cost and foregone long-run benefits: between \$420 million and \$680 million higher. These results imply that in the uniform pricing setting of the auto market, a supply-side policy can induce producers to incur greater entry costs than a demand-side policy would, but can also encourage product variety from which consumers benefit.

2.1.1 Literature

In studying the effect of regulatory mandates on the variety of products available in the market in equilibrium, we relate regulatory policy to a literature on the determinants of equilibrium product variety and welfare effects of changes in product variety, both theoretical and empirical. Economic theory is ambiguous about whether equilibrium product variety is higher or lower than socially optimal when products are differentiated (Mankiw and Whinston 1986), as consumer gains from variety (Spence 1976) may be countered by business stealing. One strain of empirical work has focused on measuring the difference between observed and optimal variety (Berry, Eizenberg, and Waldfogel 2016; Fan and Yang 2020). Another, closer to our setting, has studied how product variety responds to changes in technology (Eizenberg 2014; Brand 2020) and market structure (Wollmann 2018; Fan and Yang 2020) and the welfare effects of those changes.

We focus on the product variety generated by new products that are substitutes for, but differentiated from, existing products. Prior literature has examined new products in other differentiated product markets, such as cereal (Hausman 1996; Nevo 2003). Like Petrin (2002), which studies the introduction of the minivan in the established US automobile market, we measure how much of the social surplus created by new products is captured by the firms that produce them.

This paper relates to other studies of the non-price effects of environmental policy in the auto industry. In electric vehicles, Remmy (2020) studies the effects of European consumer subsidies on battery range. Prior work has studied the effects on the characteristics of gasoline vehicles of fuel economy and greenhouse gas regulations (Knittel 2011; Klier and Linn 2012; Whitefoot, Fowlie, and Skerlos 2017; Ito and Sallee 2018; Reynaert 2021) and regulations on local pollutants (Bresnahan and Yao 1985). Focusing on upstream innovation, Aghion, Dechezleprêtre, Hémous, Martin, and Van Reenen (2016) documents the effects of fuel taxes on the types of patents filed by automakers.

Together with Chapter 1, this paper provides the first systematic welfare analysis of the consequences of the ZEV mandate, and adds to a growing literature on the effects and design of electric vehicle policy (reviewed in Chapter 1).

2.2 Institutional background

We study the period covered by model years 2009 to 2017, which saw the introduction of the first generation of commercially available electric vehicles (EVs) in the US. Industry observers have claimed that some low-sales electric vehicles, dubbed 'compliance cars', entered the market primarily to comply with the ZEV mandate. Most major automakers in the US introduced an electric vehicle during the period, but models varied widely in engineering characteristics and in sales levels. Table 2.1 shows selected data for the models available in the US in this period, including the timing of product introduction, sales in the study period, manufacturer suggested retail price, and battery range.

Table 2.1: *Summary of electric vehicles in the US in 2009–17*

Model	Native?	Intro	Sales	Price	Range (mi)
Ford Focus		2012	9,000	\$26,437–\$38,138	76–115
Mitsubishi i-MiEV		2012	2,000	\$22,995–\$29,125	59–62
Toyota RAV4		2012	2,000	\$48,967–\$49,717	103
Fiat 500e		2013	24,000	\$29,467–\$32,120	84–87
Honda Fit		2013	1,000	\$36,625	82
smart fortwo		2013	6,000	\$23,488–\$25,000	58–68
Chevrolet Spark		2014	7,000	\$22,870–\$26,685	82
Mercedes-Benz B-Class		2014	4,000	\$39,900–\$41,450	87
Kia Soul		2015	5,000	\$30,450–\$33,700	93
Volkswagen e-Golf		2015	12,000	\$28,538–\$33,450	83–125
Nissan LEAF	X	2011	114,000	\$27,010–\$35,200	73–107
Tesla Model S	X	2012	113,000	\$57,400–\$135,000	139–315
Tesla Model X	X	2016	32,000	\$74,000–\$145,000	200–289
Chevrolet Bolt EV	X	2017	23,000	\$36,453	238
Tesla Model 3	X	2017	1,000	\$35,000	310
BMW i3	X*	2014	8,000	\$40,692–\$42,400	81
Honda Clarity	X*	2017	1,000	\$36,620	89
Hyundai Ioniq	X*	2017	<1,000	\$29,500	124

Note: *Compiled from data from MSN Autos and IHS. Prices are MSRP after average manufacturer rebates. All years are model years. Native designation assigned by authors based on press reports. Columns are the model name; whether the model is a native electric vehicle (* = EV and hybrid/plug-in hybrid versions developed jointly on the same platform); the model year of introduction; sales through the end of calendar year 2017; MSRP; and battery range in miles. (Includes content supplied by R. L. Polk; Copyright © R. L. Polk, 2019. All rights reserved)*

The first generation of models fell mainly into two groups: models that were designed from the ground up to be electric vehicles (called "native" within the industry) and models that used existing platforms from gas-powered vehicles ("non-native"). The native vehicles received upgrades during our study period, with longer ranges and increased efficiency, and usually had high sales. The non-native electric vehicles typically were not upgraded, had low sales, and dropped out of the market by 2021. (The one exception is the BMW i3, a low-sales model whose platform was designed from the ground up to support electric and plug-in hybrid versions.)

Native electric vehicle platforms offer a superior combination of engineering characteristics, but cost more to develop and require dedicated production lines. In particular, a native electric vehicle has dedicated space for the battery pack, allowing for greater battery capacity and a more spacious interior than a vehicle that must fit a battery pack in a space designed for a internal combustion engine.[4] Reasons manufacturers gave for opting for non-native electric vehicles included the lower upfront cost and the flexibility from making gas-powered and electric vehicles on the same production line.[5] Non-native models also adopted the branding and design of the gas-powered vehicles they were based on.

The timing of the first generation of commercially available electric vehicles, most of which entered the US market between 2010 and 2015, has been attributed to rapid cost declines in lithium ion batteries driven by consumer technology applications, such as laptops and phones. Compared to nickel-metal hydride batteries, used in the late 1990s and early 2000s by the Toyota Prius and the short-lived GM EV1, lithium ion batteries offer much more usable charge for the same amount of battery weight, allowing for a longer electric range, and can charge and discharge more quickly.[6] Tesla engineers have been credited with demonstrating that lithium ion batteries were feasible for electric vehicle applications.[7]

In this paper, we classify the major passenger vehicles available in this period into four technology types: conventional gas-powered vehicles with combustion engines (including flex-fuel ethanol); hybrids, which combine a combustion engine with a battery pack that cannot be charged externally; plug-in hybrids, whose battery packs are larger and can be charged externally; and battery electric vehicles. (Less commonly used technologies for passenger vehicles in this period include diesel, natural gas, and hydrogen.) Industry sources sometimes refer to plug-in hybrids and battery electric vehicles collectively as plug-in vehicles.

---

[4]See "What a teardown of the latest electric vehicles reveals about the future of mass-market EVs" (Antoine Chatelain, Mauro Erriquez, Pierre-Yves Moulière, and Philip Schäfer, McKinsey, 3/21/18).

[5]See "The Battery-Driven Car Just Got a Lot More Normal" (Bradley Berman, The New York Times, 5/4/12).

[6]See "Car Industry: Charging up the Future" (Jeff Tollefson, Nature, 11/26/08).

[7]See "Plugged In" (Tad Friend, The New Yorker, 8/17/09).

2.2.1 Entry costs for electric vehicles

The upfront costs of introducing a new vehicle include the design of the vehicle platform, body, and powertrain; the engineering work required to pass safety tests and improve driver comfort; upstream research and development; and the costs of retooling production lines. In addition to final assembly, most electric vehicle makers, native and non-native, make the battery pack and motor in-house. Manufacturers vary in their choices to make or buy the other components.[8] Because electric vehicles are mechanically simpler than internal combustion engines, requiring fewer parts and fewer assembly steps, the cost of developing a powertrain and setting up production may be lower compared to a gas-powered vehicle.

Industry estimates of the upfront cost of product entry are around $1 billion for native electric vehicles, comparable to new gas-powered models, and $100–$400 million for non-native electric vehicles. The upfront cost of a new gas-powered model and platform is at least $1 billion, and as high as $6 billion if the engine and transmission are also new.[9] Examining only redesigns of existing vehicle models, Blonigen, Knittel, and Soderbery (2017) estimate redesign costs for conventional vehicles ranging from $850 million to $3 billion, depending on vehicle class.[10]

Publicly available engineering estimates also put the upfront cost of a native electric vehicle near $1 billion.[11] According to engineering estimates, the upfront cost of an electric vehicle built on an existing platform is much lower. We draw on three accounting estimates: first, the cost of the powertrain for the Toyota RAV4 EV; second, the research and development (R&D) costs and capital investment for the Tesla Roadster (a non-native

[8]See "Trends in electric-vehicle design" (Mauro Erriquez, Thomas Morel, Pierre-Yves Moulière, and Philip Schäfer, McKinsey, 10/25/17).

[9]See "Why Does It Cost So Much For Automakers To Develop New Models?" (Terry Shea, Translogic, 6/27/10).

[10]Wollmann (2018) estimates sunk costs for commercial trucks using entry and exit behavior and finds much smaller amounts ($5 to $25 million). New truck models are usually adaptations of existing models and require less design work than new car models.

[11]See, for example, "Making electric vehicles profitable" (Yeon Baik, Russell Hensley, Patrick Hertzke, and Stefan Knupfer, McKinsey, 3/8/19), which compares the fixed costs of native and non-native electric vehicles.

model which used the Lotus Elise chassis); and third, projections from Ford in 2015 of its future investments in its EV program. These estimates provide only rough guidance — in particular, they are from different times and represent different quantities — but together form an approximate range of reasonable figures, all of them much lower than the $1 billion benchmark for a new gas-powered model. (These figures exclude earlier investments in the vehicle platform when it was only used for gas-powered cars.)

The cost of the powertrain for the Toyota RAV4 EV comes from a contract between Toyota and Tesla from 2010–12, in which Toyota agreed to pay Tesla $60 million to develop the RAV4 EV powertrain.[12] Tesla would then manufacture the powertrain components in its own facilities and ship them to Toyota for final assembly. This suggests that it would have cost Toyota more than $60 million to produce the RAV4 EV powertrain in-house, and provides a lower bound on the total cost to a major automaker of developing a non-native electric vehicle during this period.

The R&D costs and capital investment for the Tesla Roadster prior to its introduction in 2010, as measured by accounting standards, totaled $125 million.[13] Tesla was one of a group of early electric vehicle startups developing non-native vehicles, and the first to sell a highway-capable vehicle to consumers. This estimate of upfront costs includes R&D costs that were likely unnecessary once electric vehicle design was more established, but also excludes the cost of setting up manufacturing beyond Tesla's initial production rate of 50 cars per quarter.[14]

Finally, in 2015, Ford announced it would invest $4.5 billion between 2015 and 2020 to develop 13 EVs, of which some or all would use existing platforms, in addition to the EV and hybrid program it already had.[15] This translates to an average upfront investment of

[12]Source: 2012 Tesla 10-K.

[13]Source: 2010 Tesla 10-K.

[14]See "Tesla Motors Announces 2008 Roadster Production Schedule and Achievement of Critical Milestones on Crash Tests and Range Testing" (press release, 4/20/10).

[15]Source: "Ford Investing $4.5 Billion in Electrified Vehicle Solutions, Reimagining How to Create Future Vehicle User Experiences" (press release, 12/10/15).

$350 million per model.

We argue that it is reasonable to extrapolate these estimates to other non-native electric vehicles in this generation for two reasons. First, general technological improvement over the study period occurred mainly in batteries, and so principally affected marginal cost, not upfront cost. Second, these vehicles all had production volumes below mass-market levels, so manufacturers would have made similar choices for production scale. (Facing a product with a much higher volume, a manufacturer may choose to pay more upfront for optimizations that reduce marginal cost.)

2.2.2 Descriptive evidence

What is the relationship between the ZEV mandate and product entry? The sales data reveals a stark discrepancy between the distribution of native versus non-native EV sales across ZEV and non-ZEV states, as shown in Figure 2.1. Between their respective introduction dates and model year 2017, non-native EVs had 91% of their sales in states with the mandate; the figure for native EVs was 54%. In addition, except for the BMW i3, native EVs had much higher sales — 15,000 per year or more — while non-native EVs sell between 500 and 5,000 vehicles per year during the study period.

Though the decision to sell a vehicle primarily in states with the mandate suggests that the vehicle was likely not profitable without the mandate, the mandate was not the only difference between these groups of states. States with the mandate had denser charging stations and, often, more generous consumer incentives for EVs. By modeling profits explicitly, we disentangle the impact of the mandate from these demand-side factors.

In states with the mandate, did the buyers of native and non-native EVs differ along demographic characteristics? Figure 2.2 shows data from California buyers of electric vehicles who claimed rebates from the California Vehicle Rebate Program between 2010 and March 2016, grouped into equal-sized bins by census tract median income.[16] We

[16]Census tract median income from 2011–15 is drawn from the American Community Survey via the National Historical Geographic Information System (Manson, Schroeder, Van Riper, Kugler, and Ruggles 2020). More recent data are difficult to interpret because the program instituted an income requirement in March 2016.

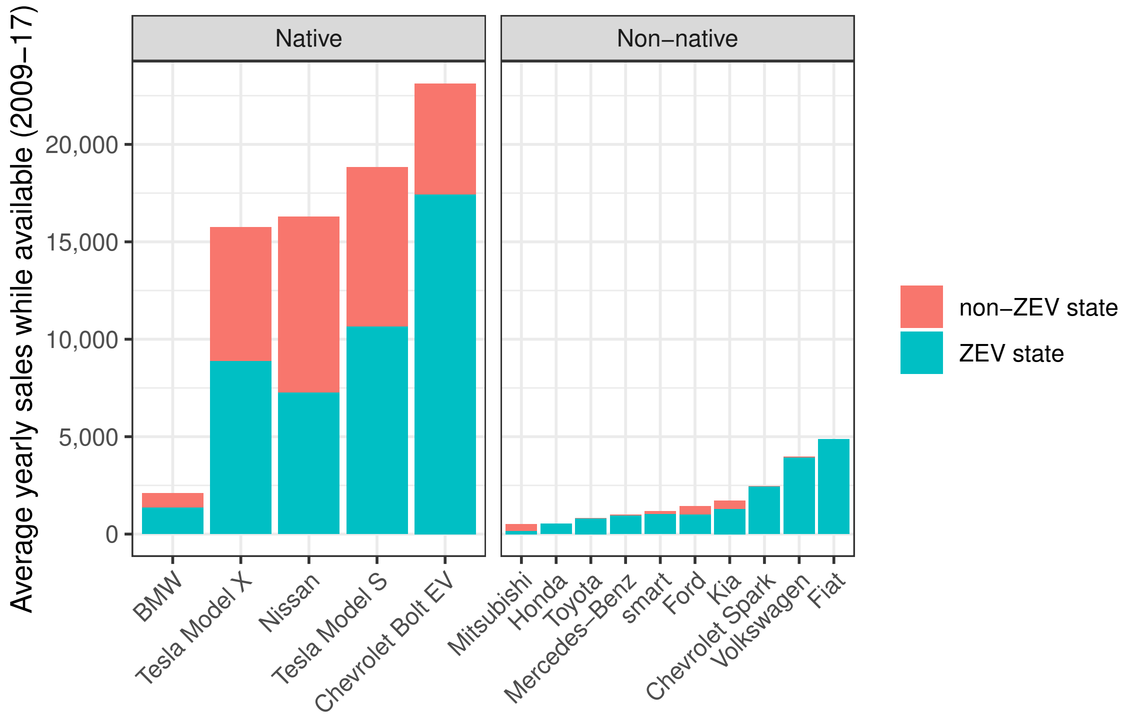

Figure 2.1: *Average yearly sales by model, selected EVs, by ZEV states and non-ZEV states*

Note: *Derived from IHS data. Average is only taken over the years the model was available. Native vs. non-native classification assigned by authors from news reports. (Includes content supplied by R. L. Polk; Copyright © R. L. Polk, 2019. All rights reserved)* divide vehicles into non-native, native non-Tesla (predominantly the Nissan Leaf), and Tesla. Though the share of EVs that are non-Tesla and native is roughly the same in higher- and lower-income areas, 47% of the EVs purchased by residents of areas in the lowest-income group (with a median income below $56,000 per year) are non-native. This percentage falls to 20% in the highest-income group, with the difference made up by Tesla vehicles. The difference in Tesla adoption is not surprising — in this period, Tesla prices were $60,000 or more — but the Nissan Leaf and most non-native vehicles were priced similarly in the low $30,000 range. These results suggest that, if the ZEV mandate increased the entry of non-native electric vehicles, the benefits accrued mostly to consumers in lower- and middle-income areas.

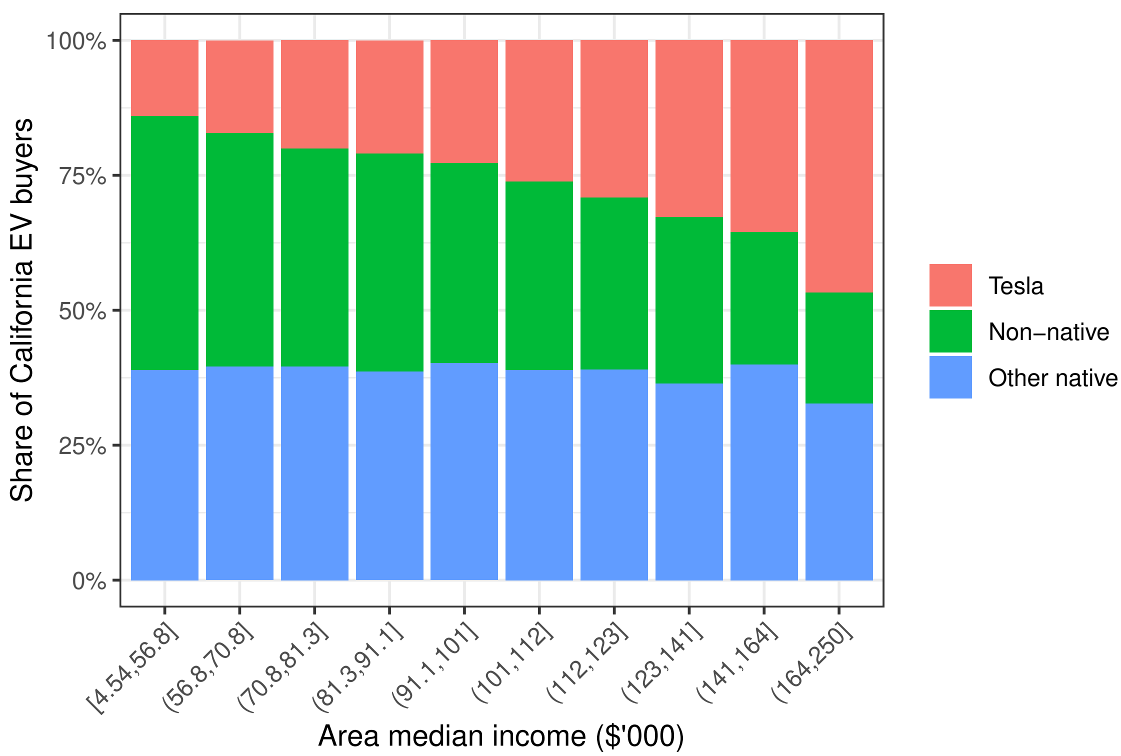

Figure 2.2: *Breakdown of EV sales by non-native, Tesla, and other native, by census tract median income (California, 2010–16)*

Note: *Derived from California Vehicle Rebate Program survey data from inception to March 2016. Buyers are divided into ten equally sized groups by census tract median income.*

2.3 Model of product entry

We adopt a simple static model of firm product introduction decisions as a framework for measuring private entry incentives and comparing them to social welfare. At the beginning of the study period, prior to the first appearance of electric vehicles in the data, firms simultaneously choose which electric vehicles to enter. Then, in each period, firms set prices and consumers choose which products to purchase, using the demand and pricing models from Chapter 1. We later use the product introduction model to predict the equilibrium product set under counterfactual policies.

2.3.1 Incentives for product entry

We examine entry incentives using a static model in which a firm compares an upfront entry cost to two types of benefits: variable profit until some fixed date and a term capturing the the long-run benefit of entry. We apply this framework specifically to the introduction of electric vehicles, holding fixed non-electric vehicle offerings by all firms. We later use the conditions for equilibrium in the model and the observed pattern of entry, together with accounting data on entry cost, to estimate lower bounds on the long-run benefit of electric vehicle entry. We then use this model to predict electric vehicle entry under counterfactual policies and determine the contribution of entry costs and benefits to welfare.

We use a static model of simultaneous product introduction decisions to capture the generational nature of electric vehicle development. That is, we assume that firms decide at the start of the study period whether to commercialize a first-generation electric vehicle. Most electric vehicles introduced during our study period were the outcome of development programs that started in the early 2000s. According to industry experts, early engineering decisions constrained the characteristics of these vehicles. Few automakers devoted resources to producing more than one EV at a time, and second-generation EVs only entered the market starting in 2017. Therefore, we consider it unlikely that firms chose the timing of product introduction or adjusted their decisions significantly after observing other firms' decisions.[17]

Our model of firm behavior depends crucially on three terms: the upfront costs of commercializing a first-generation electric vehicle, firm variable profits during the study period, and the long-run benefits of product entry. The entry cost captures the upfront investments that must be made when commercializing a new model, as described in Section 2.2.1. These include product design, production line setup, and compliance with government safety standards. The long-run benefit of entry captures variable profit after the end of our study period, spillovers across projects within a firm, and the scrap value when the product exits. In rapidly developing technologies like electric vehicles, investments in early projects may lower the entry costs of later products. In addition, introducing a product in the US may lower that product's entry costs in non-US markets. Both the long-run benefit and the entry cost are allowed to vary across products, but do not depend on quantity or on the product set in the study period. As a result, the expected long-run benefit does not capture quantity-based outcomes, like learning by doing, or the effect of entry decisions on competition in later periods.[18]

By assuming observed entry decisions form a pure strategy Nash equilibrium of the product entry game, we obtain an upper bound on each product's entry cost net of long-run benefit. This bound is not sharp: we only consider the condition that no firm can profitably deviate by removing an electric vehicle from its product set, ignoring deviations that remove multiple EVs by the same firm (for the few firms with more than one electric vehicle) or alter offerings of non-electric vehicles.

In order to combine these bounds with accounting data on entry cost, we assume that all non-native EVs have the same ratio of long-run benefit to entry cost. This assumption captures that higher upfront investment in first-generation EVs may translate to larger cost

[17]Our setting contrasts with the modular truck manufacturing studied by Wollmann (2018), in which new product introduction only requires relatively minor, well-defined changes to existing products, and thus can be done quickly.

[18]For example, our model does not capture mechanisms by which first-generation product introduction, by lowering future costs, can induce greater concentration in later generations (Dasgupta and Stiglitz 1988).

savings in future generations of electric models.[19]

The timing of the model is as follows. Firms simultaneously decide at the beginning of the study period which electric vehicles to introduce, products enter and exit on known and fixed timelines, and profits are collected from entry until exit or the end of the study period. Specifically, each firm begins the study period with a (possibly empty) set of potential electric vehicle products. At $t = 0$, each firm chooses which electric vehicles to introduce, and pays an entry cost for each introduction. Then, in each period $t = 1, \ldots, T$, products enter and exit according to their fixed timelines, firms set prices conditional on the product set, and firms collect variable profits. At the end ($t = T$), each firm collects the long-run benefit for each electric vehicle it introduced.

All firms share an information set $\mathcal{I}_0$ at $t = 0$. There is no private information, and all unknown variables are revealed to firms at the beginning of $t = 1$. Across all specifications, $\mathcal{I}_0$ includes the observable characteristics, marginal costs, and entry and exit timeline of all potential products, all of which are exogenously determined.[20] By treating each product's entry and exit timeline as fixed and known to the firm when it makes the entry decision, we rule out a product exiting during the study period in response to an unexpected shock. For products that were not sold in every state, we also treat the state-level pattern of entry as fixed.

In our main specification, $\mathcal{I}_0$ does not contain residual product quality ξ, which was defined in our model of product demand in Chapter 1. Firms instead form expectations of profits using the unconditional distribution of ξ across products. (In our alternative specification, $\mathcal{I}_0$ contains residual product quality.)

In our model, firms have the same belief about the time path of credit prices r, and we examine three scenarios for beliefs. In all scenarios, firms are fully certain about credit

[19] In a different setting, Wollmann (2018) models the scrap value of product exit as proportional to the sunk cost of entry.

[20] The assumption that firms are fully informed about marginal costs is a departure from prior literature (e.g., Wollmann (2018)). According to industry experts, firms believed in 2009 that battery prices were going to fall, though they may have disagreed on how quickly.

prices (even if incorrect ex post). This is a simplification, because credit prices clear a market whose supply and demand sides both depend on vehicle sales, and thus depend on the realization of ξ. We adopt this approach because firm valuations of credits depend on beliefs about the long-run trajectory of sales and future ZEV policy, which we do not observe. (Unfortunately, this assumption also means we cannot model the effect of product entry decisions on credit prices. Instead, we assume that removing one product from the product set does not change credit prices.) As discussed in Chapter 1, the vector of credit prices includes both ZEV credit prices and federal greenhouse gas (GHG) credit prices. We look at three scenarios for the ZEV component of r: no ZEV mandate ($r_{mt,ZEV} = \$0$), the observed realization of $r_{mt,ZEV}$, and a scarce-credit scenario ($r_{mt,ZEV} = \$5000$). In all scenarios, we assume the credit price for GHG remains fixed at \$40 per megagram of CO_2.

Variable profits. Variable profits, obtained from the per-period profit function introduced in Chapter 1, explicitly account for the effects of product entry on other products by the same firm. Let $\mathcal{J}_{ft}$ be the set of products by firm f in period t and let $\mathcal{J}_{-f,t}$ be the set of products by other firms. Then variable profit for firm f in period t is

$$\pi_{ft}(\mathcal{J}_{ft}, \mathcal{J}_{-f,t}) = \sum_{j \in \mathcal{J}_{ft}} \sum_{m \in \mathcal{M}} (p^*_{jt}(\mathcal{J}_{ft} \cup \mathcal{J}_{-f,t}) + v_{jmt} - mc_{jt}) s_{jmt}(p^*_t(\mathcal{J}_{ft} \cup \mathcal{J}_{-f,t})) M_{mt},$$

where $p^*_t(\mathcal{J}_{ft} \cup \mathcal{J}_{-f,t})$ is the Nash equilibrium of the pricing game given the product set; as in Chapter 1, j indexes products and m indexes regions, defined as US states.

We will adopt the following notation to make exposition of the product entry model simpler. Write firm f's product set across all periods as $\mathcal{J}_f = \bigcup_{t=1,\dots,T} \mathcal{J}_{ft}$ and the set of other firms' products across all periods as $\mathcal{J}_{-f} = \bigcup_{t=1,\dots,T} \mathcal{J}_{-f,t}$. Then write the present discounted value of variable profit as

$$\pi_f(\mathcal{J}_f, \mathcal{J}_{-f}) = \sum_{t=1}^{T} \delta^t \pi_{ft}(\mathcal{J}_{ft}, \mathcal{J}_{-f,t})$$

and its expectation as $\overline{\pi}_f(\mathcal{J}_f, \mathcal{J}_{-f}) = E[\pi_f(\mathcal{J}_f, \mathcal{J}_{-f}) \mid \mathcal{I}_0]$. Now, write the present discounted value of the incremental variable profit from introducing product j as $\Delta_j \pi_f(\mathcal{J}_f, \mathcal{J}_{-f})$:

$$\Delta_j \pi_f(\mathcal{J}_f, \mathcal{J}_{-f}) = \begin{cases} \pi_f(\mathcal{J}_f, \mathcal{J}_{-f}) - \pi_f(\mathcal{J}_f \setminus \{j\}, \mathcal{J}_{-f}), & j \in \mathcal{J}_f \\ \pi_f(\mathcal{J}_{ft} \cup \{j\}, \mathcal{J}_{-f,t}) - \pi_f(\mathcal{J}_f, \mathcal{J}_{-f}), & j \notin \mathcal{J}_f \end{cases}$$
$$= \begin{cases} \sum_{t=1}^{T} \delta^t \left(\pi_{ft}(\mathcal{J}_{ft}, \mathcal{J}_{-f,t}) - \pi_{ft}(\mathcal{J}_{ft} \setminus \{j\}, \mathcal{J}_{-f,t}) \right), & j \in \mathcal{J}_f \\ \sum_{t=1}^{T} \delta^t \left(\pi_{ft}(\mathcal{J}_{ft} \cup \{j\}, \mathcal{J}_{-f,t}) - \pi_{ft}(\mathcal{J}_{ft}, \mathcal{J}_{-f,t}) \right), & j \notin \mathcal{J}_f \end{cases},$$

and its expectation as $\Delta_j \overline{\pi}_f(\mathcal{J}_f, \mathcal{J}_{-f}) = E[\Delta_j \pi_f(\mathcal{J}_f, \mathcal{J}_{-f}) \mid \mathcal{I}_0]$.

Payoff from entry. The payoff to firm f is the sum of variable profits until T and, for each product the firm introduces, the product's long-run benefit net of entry cost. For each product j, let $SC_j \geq 0$ be entry cost and $\delta^{-T} \pi_j^0 \geq 0$ be the expected long-run benefit collected at T. (The value of this benefit at $t = 0$ is therefore π_j^0.) Both terms are fixed and do not depend on firm actions or the actions of other firms. Summing these components, we arrive at the total payoff.

Assumption 1 *Firm f's total expected payoff, given $\mathcal{J}_f$ and $\mathcal{J}_{-f}$, is*

$$\overline{\pi}_f(\mathcal{J}_f, \mathcal{J}_{-f}) + \sum_{j \in \mathcal{J}_f} \left(\pi_j^0 - SC_j \right).$$

Now we examine $\mathcal{J}_f$ as (pure strategy) Nash equilibrium play in the simultaneous-move product entry game. In equilibrium, no firm expects removing any one EV from its product set to be a profitable deviation. Therefore, if the firm's product set includes product j in equilibrium, the sum of its incremental variable profit and long-run benefit must be larger than its entry cost.

Proposition 2 *Suppose a pure strategy Nash equilibrium of the product entry game exists. Let firm f's product set in this equilibrium be $\mathcal{J}_f$ and the set of other firms' products be $\mathcal{J}_{-f}$. Then for each $j \in \mathcal{J}_f$,*

$$\Delta_j \overline{\pi}_f(\mathcal{J}_f, \mathcal{J}_{-f}) + \pi_j^0 - SC_j \geq 0.$$

We will call $\Delta_j \overline{\pi}_f(\mathcal{J}_f, \mathcal{J}_{-f}) + \pi_j^0 - SC_j$ the private entry incentive.

As described in Section 2.4.1, we use this condition to obtain an upper bound on the net entry cost, $SC_j - \pi_j^0$, for each non-native EV that entered. We use a proportionality assumption in order to interpret the net entry cost using industry estimates of SC_j.

Assumption 2 *There exists κ such that for all non-native EVs j, $\pi_j^0 / SC_j = \kappa$.*

We therefore write the net entry cost as $(1-\kappa)SC_j$, and the private entry incentive as $\Delta_j \overline{\pi}_f(\mathcal{J}_f, \mathcal{J}_{-f}) - (1-\kappa)SC_j$.

2.3.2 Product entry and welfare

Using the product entry model, we ask whether the private incentives to introduce a product are aligned with the effect of product entry on social welfare. Each firm's product entry decision affects consumers, other firms, and the environment, and the total of these externalities can be positive or negative in equilibrium. By comparing these externalities and the incentives created by the mandate, we can determine how well the policy is aligning private incentives with social welfare.

We use the social welfare function described in Chapter 1, augmenting producer surplus to include entry costs and benefits, to compare the equilibrium of the product entry game to the socially efficient pattern of entry. Total welfare combines consumer surplus, environmental effects, and producer surplus (which itself combines variable profit, long-run benefit, and entry cost).

Specifically, to producer surplus we add π_j^0, long-run benefits of a product to the firm that introduced it, net of entry cost SC_j. We exclude benefits or costs for other firms outside product-market competition (such as through information spillovers), effects on consumer surplus or environmental externalities in future technological generations, and complementarities in entry costs across products. (We now condition explicitly on the choice set $\mathcal{J} = \{\mathcal{J}_f\}_f$.)

$$PS(\mathcal{J}) = \sum_f \left(\pi_f(\mathcal{J}_f, \mathcal{J}_{-f}) - v_f(\mathcal{J}_f, \mathcal{J}_{-f}) + \sum_{j \in \mathcal{J}_f} (\pi_j^0 - SC_j) \right). \quad (2.1)$$

Total welfare again combines consumer surplus, producer surplus, and environmental effects; we now condition explicitly on the product set and write

$$W(\mathcal{J}) = CS(\mathcal{J}) + Env(\mathcal{J}) + PS(\mathcal{J}). \tag{2.2}$$

The consumer and producer externalities that arise from our product entry model and social welfare function parallel Mankiw and Whinston (1986); we also include environmental externalities. For a given firm f, rewrite (2.1) to separate firm f's profit from other firms:

$$\begin{aligned} PS(\mathcal{J}) = \pi_f(\mathcal{J}_f, \mathcal{J}_{-f}) - v_f(\mathcal{J}_f, \mathcal{J}_{-f}) + \sum_{j \in \mathcal{J}_f} (\pi_j^0 - SC_j) \\ + \sum_{g \neq f} \left(\pi_g(\mathcal{J}_g, \mathcal{J}_{-g}) - v_g(\mathcal{J}_g, \mathcal{J}_{-g}) + \sum_{j \in \mathcal{J}_g} (\pi_j^0 - SC_j) \right). \end{aligned}$$

We can therefore compare the private entry incentive for a particular product j with the effect of j's entry on total welfare. Consider a product j to be produced by firm f and two product sets, $\mathcal{J}^0$ and $\mathcal{J}^1$, that only differ by the inclusion of j (that is, $j \notin \mathcal{J}_f^0$, $\mathcal{J}_f^1 = \mathcal{J}_f^0 \cup \{j\}$, and $\mathcal{J}_g^1 = \mathcal{J}_g^0$ for $g \neq f$). Then

$$\begin{aligned} PS(\mathcal{J}^1) - PS(\mathcal{J}^0) = \Delta_j \pi_f(\mathcal{J}_f^1, \mathcal{J}_{-f}^1) - (v_f(\mathcal{J}_f^1, \mathcal{J}_{-f}^1) - v_f(\mathcal{J}_f^0, \mathcal{J}_{-f}^1)) + (\pi_j^0 - SC_j) \\ + \sum_{g \neq f} \Big(\pi_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^1) - \pi_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^0) - (v_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^1) - v_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^0)) \Big). \end{aligned}$$

Since we assumed they do not depend on the product set, the $\pi^0 - SC$ terms for all other products drop out.

The effect on welfare thus combines the ex-post analog of firm f's private entry incentive, the effect on firm f's regulatory credits, and the uncaptured externalities for other firms, consumers, and the environment:

$$
\begin{aligned}
W(\mathcal{J}^1) - W(\mathcal{J}^0) = & \underbrace{\Delta_j \pi_f(\mathcal{J}_f^1, \mathcal{J}_{-f}^1) + (\pi_j^0 - SC_j)}_{\text{private entry incentive (ex post)}} \\
& - \underbrace{v_f(\mathcal{J}_f^1, \mathcal{J}_{-f}^1) - v_f(\mathcal{J}_f^0, \mathcal{J}_{-f}^1)}_{\text{regulatory credits}} \\
& + \sum_{g \neq f} \Big(\pi_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^1) - \pi_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^0) - (v_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^1) - v_g(\mathcal{J}_g^1, \mathcal{J}_{-g}^0)) \Big) \\
& + CS(\mathcal{J}^1) - CS(\mathcal{J}^0) + Env(\mathcal{J}^1) - Env(\mathcal{J}^0).
\end{aligned}
$$

This expression demonstrates that product introduction has externalities, whose total may be positive, negative, or zero. These include the negative effect of entry on other firms' profits, which is related to the business stealing effect (Mankiw and Whinston 1986); the consumer benefit from variety (Spence 1976); and the net environmental effect of consumer substitution from cleaner and dirtier alternatives.

Therefore, we assess the effects of a supply-side policy on entry incentives by testing whether the value of regulatory credits aligns private incentives with social welfare. This test does not require any estimates of the magnitudes of entry costs and long-run benefits, but it is only a test of incentives, not whether there is over- or under-entry in equilibrium. If the value of regulatory credits is less the sum of the externalities, for example, the firm does not fully internalize the benefits to others of product introduction, but when accounting for entry cost, entry on the whole may be welfare-enhancing or welfare-reducing.

2.4 Estimating product entry incentives

Using the static product entry framework described in Section 2.3.1, we compare accounting estimates of entry costs with observed variable profit in order to estimate a lower bound on long-run benefit. We estimate a long-run benefit for non-native electric vehicles of at least $0.78 \times$ entry cost.

To calculate firm-level variable profits, consumer surplus, and environmental externalities under a variety of scenarios and product sets, we use the model and estimates from Chapter 1. The only data we add are the entry cost estimates from Section 2.2.1.

The model of product entry described above is conditional on firms competing in prices given the product set. Due to the computational burden of recalculating equilibrium prices with different product sets and policy environments, the exercises that follow typically hold product prices fixed even as the product set changes. This assumption may bias estimates of consumer surplus (though the direction may depend on substitution patterns), and will generally understate counterfactual firm profits.[21]

2.4.1 Estimation of product entry parameters

We estimate bounds on product entry parameters in order to predict product introduction under counterfactual policies and measure total welfare. By assuming the entry pattern observed in the data is a Nash equilibrium of the product entry game described in Section 2.3.1, we can estimate an upper bound on the entry cost net of long-run benefit.[22] To separate these terms, we use the accounting estimates from Section 2.2.1 to bound entry costs.

For any non-native electric vehicle model j, let $f(j)$ be the firm that produces product j. Proposition 2 gives a necessary condition for equilibrium play that depends on unobserved variables and variables that can be estimated from the demand model:

$$\text{introduce if } (1-\kappa)SC_j \leq \Delta_j \pi_{f(j)}.$$

To estimate expected incremental variable profit $\Delta_j \pi_{f(j)}$, we compare firm profits in the observed equilibrium (obtained from the demand model) with profits in the scenario in which j is dropped and market shares adjust. As discussed earlier, we hold prices fixed. We try three paths for r, as described in Section 2.3.1, and ignore the potential effect of product

[21] Work to incorporate equilibrium price changes, using computational methods adopted in Chapter 1, is ongoing.

[22] Our approach is an application of the framework in Pakes, Porter, Ho, and Ishii (2015), and has elements in common with Wollmann (2018).

introduction on credit prices.

In our main specification, the shared firm information set $\mathcal{I}_0$ does not contain product quality ξ. Firms instead form expectations of profits using the distribution of ξ across products. We approximate this expectation by calculating average profits over 50 draws of the ξ vector. Specifically, for each j, we draw ξ_j i.i.d. from a normal distribution with mean and variance calibrated to the distribution of estimated ξs across products. In our alternative specification, $\mathcal{I}_0$ contains product quality.

Using these estimates for all non-native EVs, we can obtain a lower bound on κ, the ratio of expected long-run benefit to entry cost. This estimate uses only the incremental variable profit of the least profitable product and the lower bound on entry cost. When incremental variable profit falls short of entry cost, the long-run benefit required to justify entry is larger, giving a higher implied lower bound on κ. By contrast, if incremental variable profit is higher than entry cost, entry can be rationalized by any value of $\kappa \geq 0$.

Consider $j \in \mathcal{J}_{NN}$, the set of non-native ("NN") electric vehicle models that enter before 2017. Let $[\underline{SC}_{NN}, \overline{SC}_{NN}]$ be the range of accounting estimates. Then, for each $j \in \mathcal{J}_{NN}$,

$$(1-\kappa)\underline{SC}_{NN} \leq (1-\kappa)SC_j \leq \Delta_j \pi_{f(j)},$$

giving a lower bound on κ, which we denote $\underline{\kappa}$:

$$\kappa \geq \underline{\kappa} \equiv 1 - \left(\underline{SC}_{NN}\right)^{-1} \min_{j \in \mathcal{J}_{NN}} \Delta_j \pi_{f(j)}.$$

That is, the long-run benefit for the least profitable product observed in the data must be large enough to recover the lowest possible entry cost.

Because we do not have data on products that were not introduced, we are unable to use entry behavior to estimate an upper bound on κ. For the non-native models introduced early in our study period, we assume that $\kappa \leq 1$. For products with positive $\Delta_j \pi_{f(j)}$, this assumption does not restrict entry in equilibrium: the model generates the same prediction for any $\kappa \geq 1$. (The exact value of κ matters for welfare calculations.)

Given a particular value of κ, the entry rule also implies an upper bound on SC_j for each product j that entered: $SC_j \leq (1-\kappa)^{-1}\Delta_j \pi_{f(j)}$.

2.4.2 Product entry parameter estimates

We begin by estimating incremental variable profit for each electric vehicle j using the demand and supply model. These estimates are shown in Figure 2.3, where they are compared to accounting estimates of entry cost. While all native models except the BMW i3 recover their entry cost without long-run benefits, about half of non-native models do not, with expected incremental variable profits as low as $50 million (as compared to entry costs that start at $100 million). In our preferred specification, we estimate a long-run benefit for non-native electric vehicles of at least 0.78 $\times$ entry cost.

These figures are generally robust to the inclusion of unobserved product quality ξ in firm information sets, but the identities of the lowest-profit products depends on which information set is assumed. For example, the Mercedes and Ford models have high expected incremental variable profits based on observable characteristics and costs, but low realized values of ξ that cause the realized incremental variable profits to be much lower.

We next estimate a lower bound on κ for non-native EVs. From accounting estimates, we use $\underline{SC}_{NN}$ =$100 million. Under the realized credit price path, the lowest expected incremental variable profit is $47 million when ξ is not conditioned on, and $22 million when ξ is conditioned on. (We exclude products that entered during model year 2017.) Therefore, we obtain $\underline{\kappa} = 0.53$ when ξ is not conditioned on, and 0.78 when ξ is conditioned on. The full set of results is summarized in Table 2.2.

The results from the $0 credit scenario imply that the mandate only affects entry in a relatively narrow set of values of κ: if long-run benefit is high enough, observed entry can be explained by non-regulatory factors alone.

2.4.3 Comparing social welfare and entry incentives

Using the framework for product entry incentives and welfare in Section 2.3.2, we compare our estimates of credit value to the effects of product entry on the profits of other

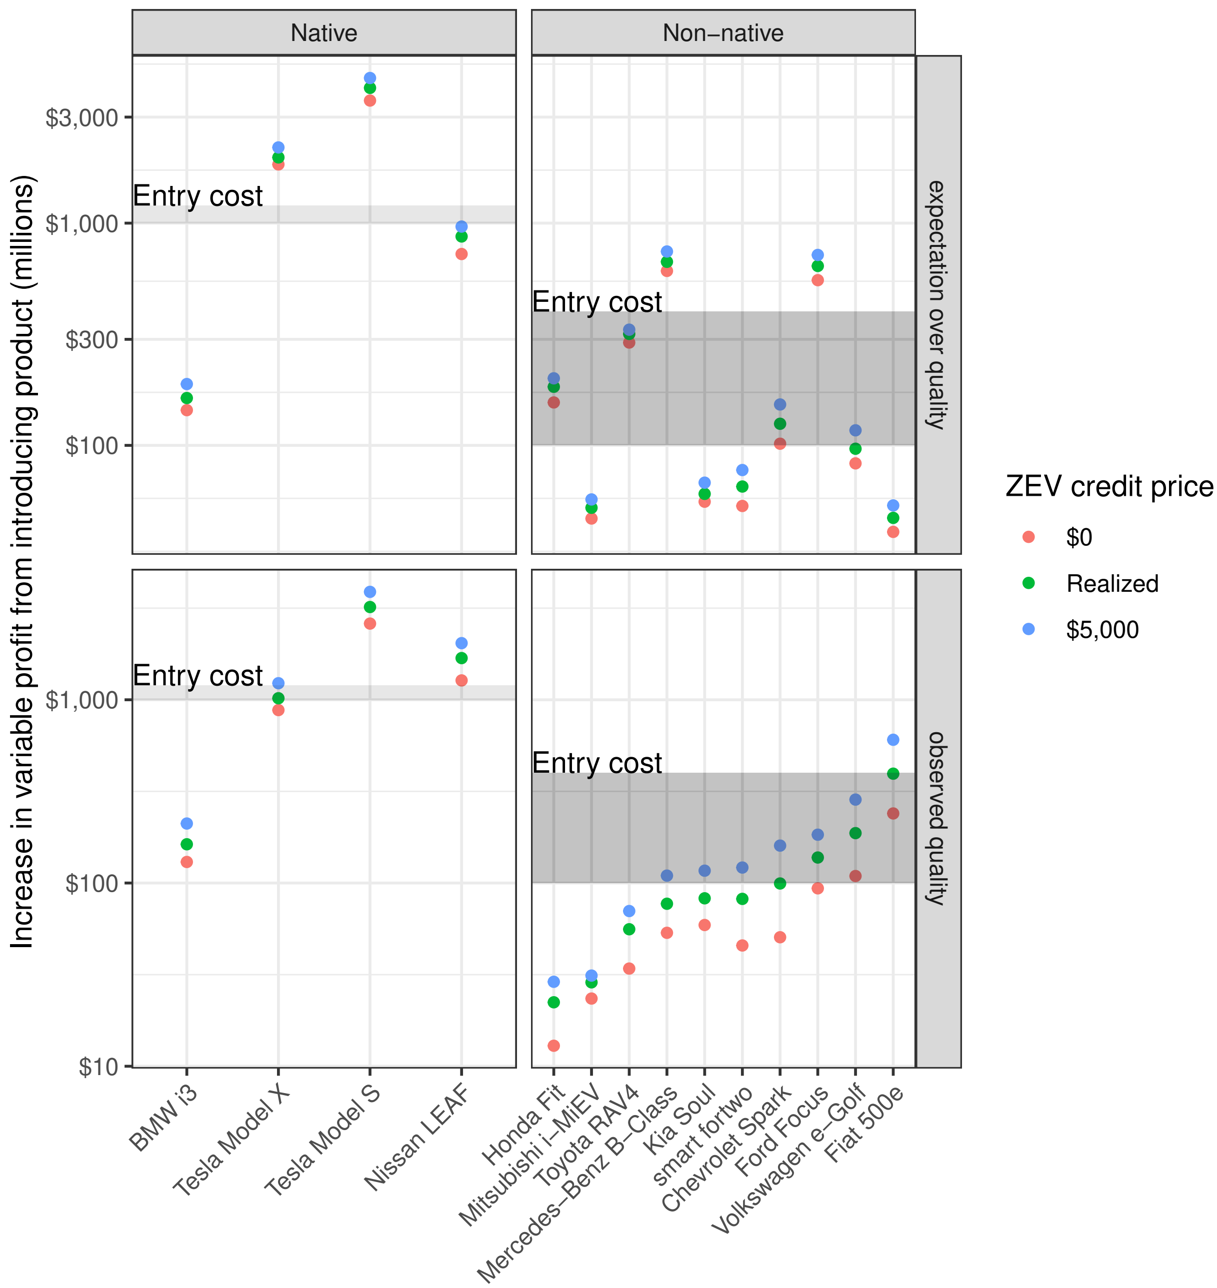

Figure 2.3: *Comparison of incremental firm variable profit with accounting estimates of entry cost*

Note: *Incremental firm variable profit is the difference between variable profit when the product is in and when it is out, including regulatory credits. It gives an upper bound on entry cost net of unobserved benefit. Entry cost is based on industry accounting estimates. Top panel shows average profits over draws of product quality ξ; bottom panel shows profits using estimated product quality $\hat{\xi}$. All profits are shown under three scenarios for the ZEV credit price: $0, $5,000, and the path of realized prices.*

Table 2.2: *Estimates of the lower bound on κ*

Credit scenario	Observed quality	Expectation over quality
\$0	0.871	0.592
Realized	0.777	0.528
\$5,000	0.711	0.463

Note: *Estimates of $\underline{\kappa}$, the lower bound on long-run benefit as a fraction of entry cost, for non-native electric vehicles. Obtained under three scenarios for anticipated credit prices (\$0, \$5,000, and the path of realized prices) and two firm information sets about the vector of product quality ξ (certainty and an expectation over i.i.d. draws from the unconditional distribution from our demand estimates). Estimates are the lowest level of long-run benefit that rationalizes the observed entry pattern for non-native electric vehicles within our model, assuming entry costs of at least \$100 million.*

firms, consumers, and the environment. We find that the mandate allowed firms to capture 25–50% of the externalities of entry, as opposed to less than 10% under the greenhouse gas program alone.

As discussed in Section 2.3.2, we can use the demand and pricing model to test whether the credit value is properly aligning private entry incentives with the social welfare effects of entry. If the incremental value of credits is close to the sum of the product entry externalities, then the firm is internalizing those externalities when choosing product entry.

We quantify the three externalities — the effect on other firms, consumer surplus, and environmental effects — and compare them to incremental credit value $\Delta_j v_f$ under various assumptions about credit prices. To do so, for each electric vehicle j, we predict changes in consumer surplus and market shares if j is dropped. We then compute other firms' profits, using marginal cost estimates and the predicted shares, and environmental effects from predicted changes in shares (as described in Chapter 1).

As discussed above, we hold product prices fixed to simplify calculations. In reality, other firms will respond to a product removal by changing prices, which will also affect quantities for the firm's other products. Therefore, accounting properly for price changes could increase or decrease the sum of the externalities and incremental credit value.

We compute welfare with and without information about the realization of ξ. In our ex post specification, we assess the incentives and externalities under full information about

ξ, using each product's estimated ξ from the demand model. In our ex ante specification, which more closely aligns with our model of firm information, we compute incentives and externalities as an expectation over ξ, which we approximate using a normal distribution calibrated to the distribution of estimated ξs. We use three scenarios for credit prices, ignoring uncertainty, as described in Section 2.3.1.

We estimate that the sum of the externalities is positive for all products. That is, there are social benefits from product entry that are not internalized by private firms in the absence of regulatory credits. In all cases, consumer surplus is the positive externality, the effect on other firms' profits is negative but smaller, and the environmental benefit is small in absolute value and sometimes negative. The small environmental benefit reflects both the low sales of electric vehicles and the implication of our demand specification that the closest substitutes for electric vehicles are other electric or clean vehicles.

As noted by Ackerberg and Rysman (2005) and Fan and Yang (2020), the positive effect on consumer surplus may be partially driven by an unwanted feature of logit demand systems: high draws of the independent Type 1 Extreme Value shocks for some consumers for the products being considered.[23]

The percentage of the externality captured by credit values is given in Figure 2.4 for each product. In the ex post specification, using estimated values of ξ, GHG credits alone make up about 10% of the entry externalities; under realized ZEV prices, 25–50% of the externality is captured by credit values. Even in the \$5000 ZEV credit scenario, most firms do not internalize the full benefits of product entry.

In the ex ante specification, which averages over the distribution of ξ, the gap between credit values and the sum of externalities is wider. GHG credits alone make up 9–17% of entry externalities, while in the \$5000 ZEV credit scenario, credit value is still only 25–45% of the sum of entry externalities.

[23] Work to calculate consumer surplus with this effect removed, following the method of Fan and Yang (2020), is ongoing.

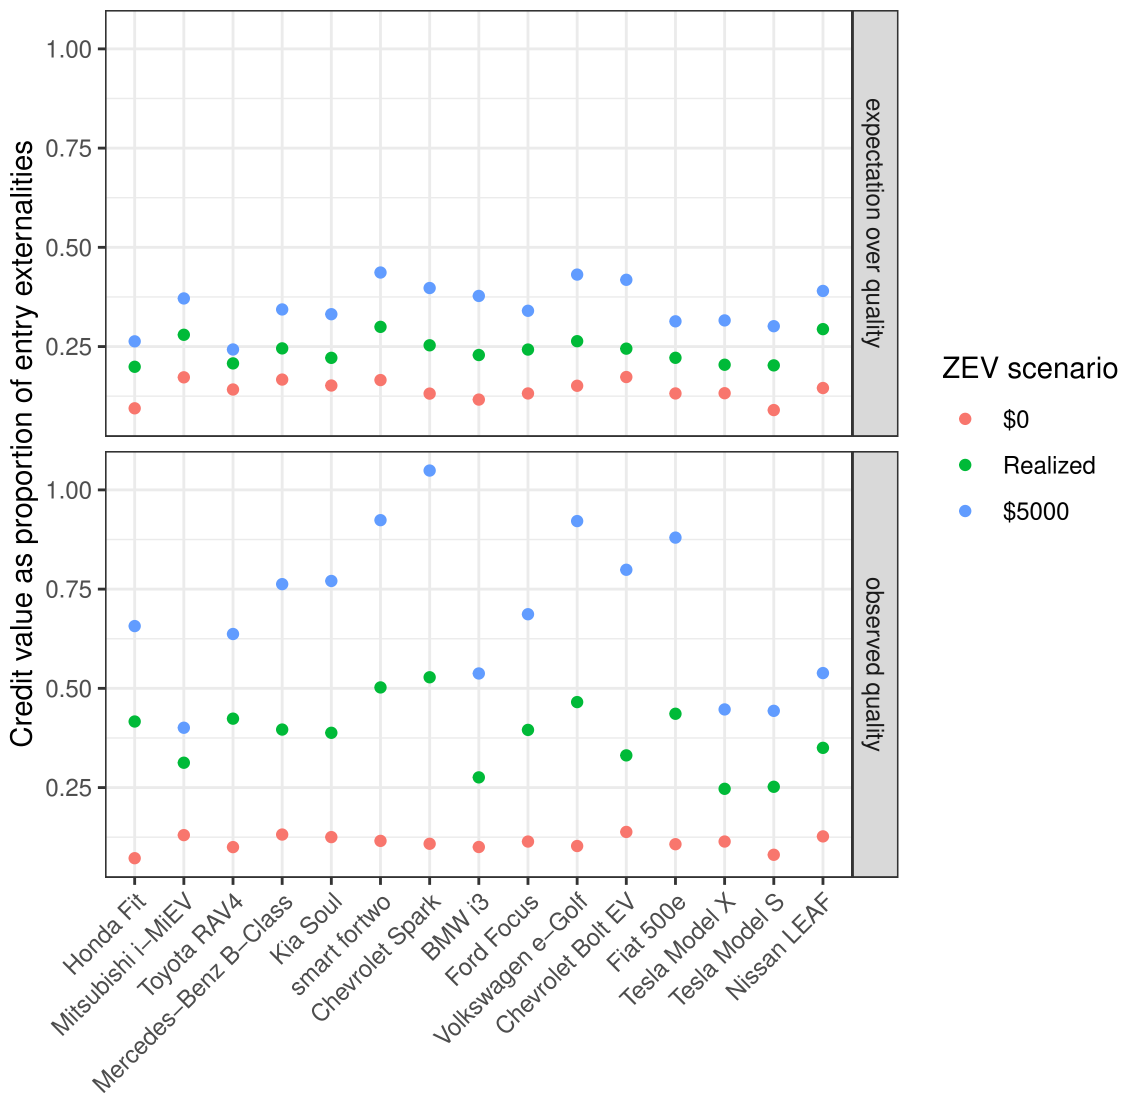

Figure 2.4: *Ratio of credit value to sum of entry externalities, selected EVs*

Note: *Credit value includes GHG and ZEV credits. Entry externalities are the change in consumer surplus, negative environmental damages, and other-firm profit resulting from entry. Top panel shows averages over draws of product quality; bottom panel shows calculations when product quality is known.*

2.5 Policy-induced product variety

We now examine what product entry would look like in the absence of the ZEV regulation, and the welfare consequences of that change. We focus on the loss of non-native electric vehicles; as a result, welfare losses are concentrated in the regulated states.

We start by documenting that reduced product variety, by itself, has a negative impact on consumers and sales, but that impact is much larger if higher-sales products are removed than if lower-sales products are removed. Nationally, the consumer surplus effect is constant across the income distribution (in relative terms) but within the regulated states, the loss of high-sales products lowers consumer surplus among lower-income consumers the most.

The effect of removing the ZEV regulation is to raise the cost of selling electric vehicles and lower the cost to large manufacturers of selling gas vehicles. Removing this incentive will then lower the returns to EV product introduction, and the vehicles providing the lowest incremental variable profit might now become unprofitable to enter. Because these are also the lowest-sales products, we find that most plausible product entry equilibria also feature relatively muted welfare effects; in all of them, however, total surplus is lowered by the lost variety.

Using the product entry model and estimates, we compare the lost incremental variable profit to estimated entry cost and long-run benefits. We implement this method by selecting a subset of product entry configurations — specifically, those in which all EVs enter except the the n lowest-selling non-native EVs, for various n — and testing in each whether any firm has an incentive to deviate. As above, we do this exercise holding product prices fixed; this will generally lead us to overstate the effect of the policy change. In reality, firms can recoup part of the lost incentives from the removal of ZEV by increasing consumer prices, so this assumption may understate firm profits without ZEV and lead us to predict less entry than there actually would be. We also ignore that lower EV sales will lower the supply of GHG credits, potentially increasing the price of GHG credits and counteracting some of the lost electric vehicle incentives from ZEV. (In extreme cases, ZEV may not be a binding constraint on firms due to the GHG regulation, and so removing ZEV will have no effect;

see Leard and McConnell (2021) for an exploration of this case.)

Across scenarios, we explore the effect of the lost product variety on sales and welfare, finding a modest negative effect on electric vehicle sales and a negative effect on consumer surplus but a smaller, positive effect on producer surplus due to avoided entry costs. Holding prices fixed isolates the product variety channel of the policy, and excludes the effects of the pass-through of the ZEV regulation to prices. In general, the price increases that result from the removal of ZEV will add to consumer losses from foregone variety, and increase producer gains.

2.5.1 Effects of reduced entry

We illustrate the welfare effects of reduced product entry in multiple scenarios: when only the four lowest-selling non-native EVs avoid entry, when the six lowest-selling non-native EVs avoid entry, and when all ten non-native EVs avoid entry. As we show in Section 2.5.2, the final scenario is not an equilibrium given our estimates, but it is nonetheless useful for comparison. The dollar effect of removing these products on consumer surplus and variable profit, and the percent effect on electric vehicle sales, are shown in Table 2.3. The effect of reducing the product set by the four lowest sellers is relatively modest, while the effect of eliminating non-native electric vehicles is many times larger: about seven times larger for consumer surplus and EV sales and four times larger for variable profits. The consumer surplus loss is focused on the ZEV states, which is (as shown in Figure 2.1) the source of most of their sales.

We also break down the estimated change in consumer surplus by income group, in order to test the hypothesis from Section 2.2.2 that the variety induced by the ZEV mandate benefits lower- and middle-income consumers more than high-income consumers. In our demand specification, differences in preferences across income groups are largely driven by the tight link between income and price sensitivity, so the result is mainly driven by price differences between native and non-native EVs (shown in Table 2.1). (Some of the effect is also driven by the estimated relationship between education and the preference for an EV.)

Table 2.3: *Welfare and sales effects of reduced product variety, selected scenarios*

	Scenario 4	Scenario 6	Scenario 10
Consumer surplus, national ($m)	−$260	−$500	−$1,710
Consumer surplus, ZEV states ($m)	−$230	−$430	−$1,580
Variable profit, national ($m)	−$50	−$80	−$230
EV sales, national (%)	−2.6%	−5.6%	−19.9%
EV sales, ZEV states (%)	−3.5%	−7.5%	−29.4%

Note: *The effects on electric vehicle sales, consumer surplus, and variable profit of reducing the set of electric vehicles while leaving prices fixed and letting consumers adjust. In Scenario 4, the four non-native electric vehicles with the lowest sales (Mitsubishi, Honda, Toyota, and Mercedes-Benz) are removed from the product set; Scenario 6 adds the two next-lowest (Kia and Smart) and Scenario 10 adds the remainder (Ford, Chevrolet, Volkswagen, Fiat). We report the dollar change in consumer surplus and percent change in electric vehicle sales at the national level and across the ten ZEV states; we report variable profit across firms (not including entry cost or long-run benefit) at the national level.*

The predicted change in consumer surplus by income quintile are shown in Figure 2.5 for both scenarios.

We find that the hypothesis only holds when variety is significantly reduced, and only holds in the regulated states. We show relative effects, compared to the baseline of the observed product set, because higher-income groups are less price sensitive and purchase more new vehicles, and thus see larger absolute effects in dollar terms. Eliminating the four lowest-sales non-native vehicles (including the inexpensive Mitsubishi but also the more expensive Mercedes-Benz) has the largest impact on higher-income consumers, especially in regulated states. By contrast, eliminating all non-native vehicles also has a substantial negative effect on low-income consumers. In regulated states, the effect is muted for high-income consumers, but when the rest of the US is included, the effect is relatively constant across the income distribution.

2.5.2 Evaluating entry scenarios

We now consider a selection of entry scenarios in which the lowest-selling non-native electric vehicles do not enter. Other equilibria are possible, but variable profits are tightly linked to sales and so the lowest-profit products are also the ones with the lowest sales.

Which scenarios are possible equilibria is sensitive to the magnitude of long-run benefits.

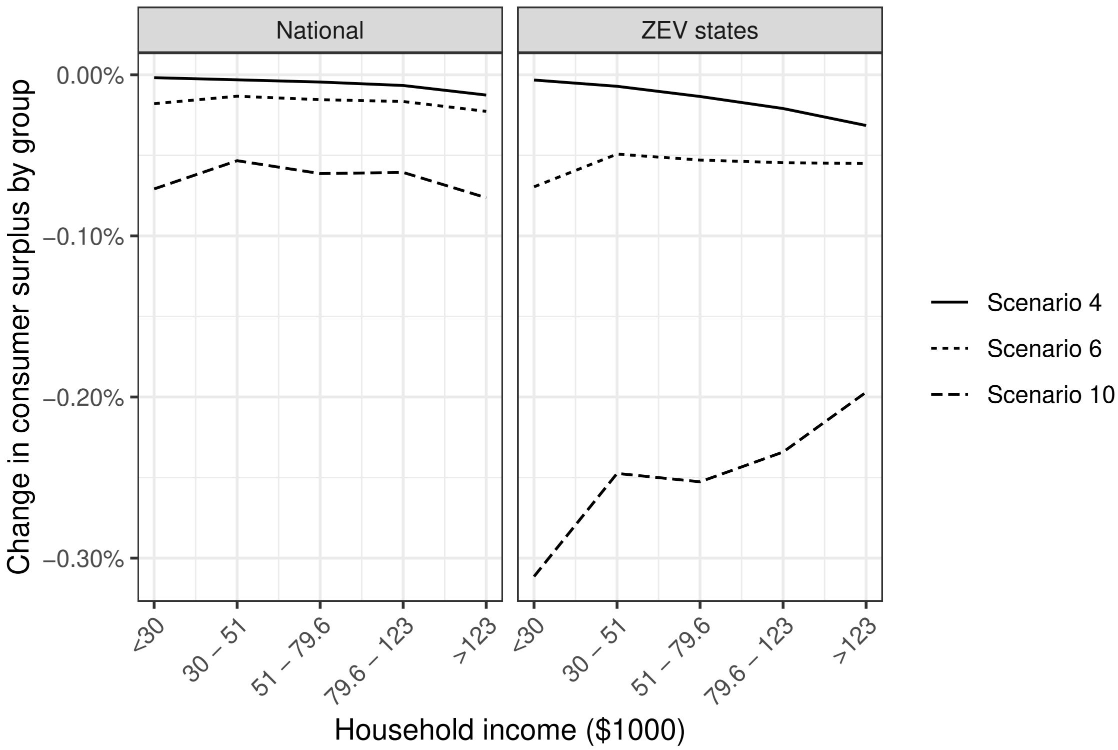

Figure 2.5: *Welfare effect of reduced variety across the income distribution*

Note: *This figure shows the change in consumer surplus in the new vehicle market from reducing the set of electric vehicles while leaving prices fixed. In Scenario 4, the four non-native electric vehicles with the lowest sales (Mitsubishi, Honda, Toyota, and Mercedes-Benz) are removed from the product set; Scenario 6 adds the two next-lowest (Kia and Smart) and Scenario 10 adds the remainder (Ford, Chevrolet, Volkswagen, Fiat). Households are grouped into quintiles of nominal income and aggregated over model years 2012–2017. Results are reported separately for the entire US and only for the states with the ZEV regulation.*

The results presented in Table 2.2 already imply that if long-run benefits are high, the observed product set can be explained even without the ZEV program. Specifically, if firms believed credit prices would be zero, the observed product set can be an equilibrium as long as long-run benefits are greater than 87% of sunk cost.

The results of our equilibrium tests are shown in Figure 2.6, which counts the number of firms with an incentive to deviate in each scenario. Our estimates reject the hypothesis that the non-native EVs were only introduced because of the ZEV regulation. In fact, for any value of κ supported by our estimates, scenarios where seven or more of the ten non-native EVs avoid entering are rejected. Among the remaining scenarios, those with more entry are only supported if long-run benefit is high, and those with less entry are only supported if long-run benefit is low.

To implement the equilibrium test for each $\kappa \in [\underline{\kappa}, 1)$, we determine whether any firm has an incentive to deviate by entering or removing a single product. (We do not consider the hypothetical entry of products that were not observed in the data, and ignore incentives to deviate by products that were only introduced in 2017.) We do so by comparing the counterfactual incremental variable profit of the product, which we write $\Delta_j \pi^{cf}_{f(j)}$, against the entry cost net of long-run benefit, $SC_j(1-\kappa)$. As in the estimation section, we hold product prices fixed for this comparison. Because we do not observe SC_j, we use the bounds derived in Section 2.4.2, and use the data to assign $d_j(\kappa) \in \{\text{in, out, unknown}\}$ as follows:

$$\begin{cases} \text{out,} & \Delta_j \pi^{cf}_{f(j)} < \underline{SC}(1-\kappa), \\ \text{in,} & \Delta_j \pi^{cf}_{f(j)} > \min\{\overline{SC}(1-\kappa), \Delta_j \pi_{f(j)}\}, \\ \text{unknown,} & \text{otherwise} \end{cases} \cdot$$

For each κ, we then say that the scenario is rejected if any product is out and has $d_j(\kappa) = \text{in}$, or the product is in and has $d_j(\kappa) = \text{out}$.

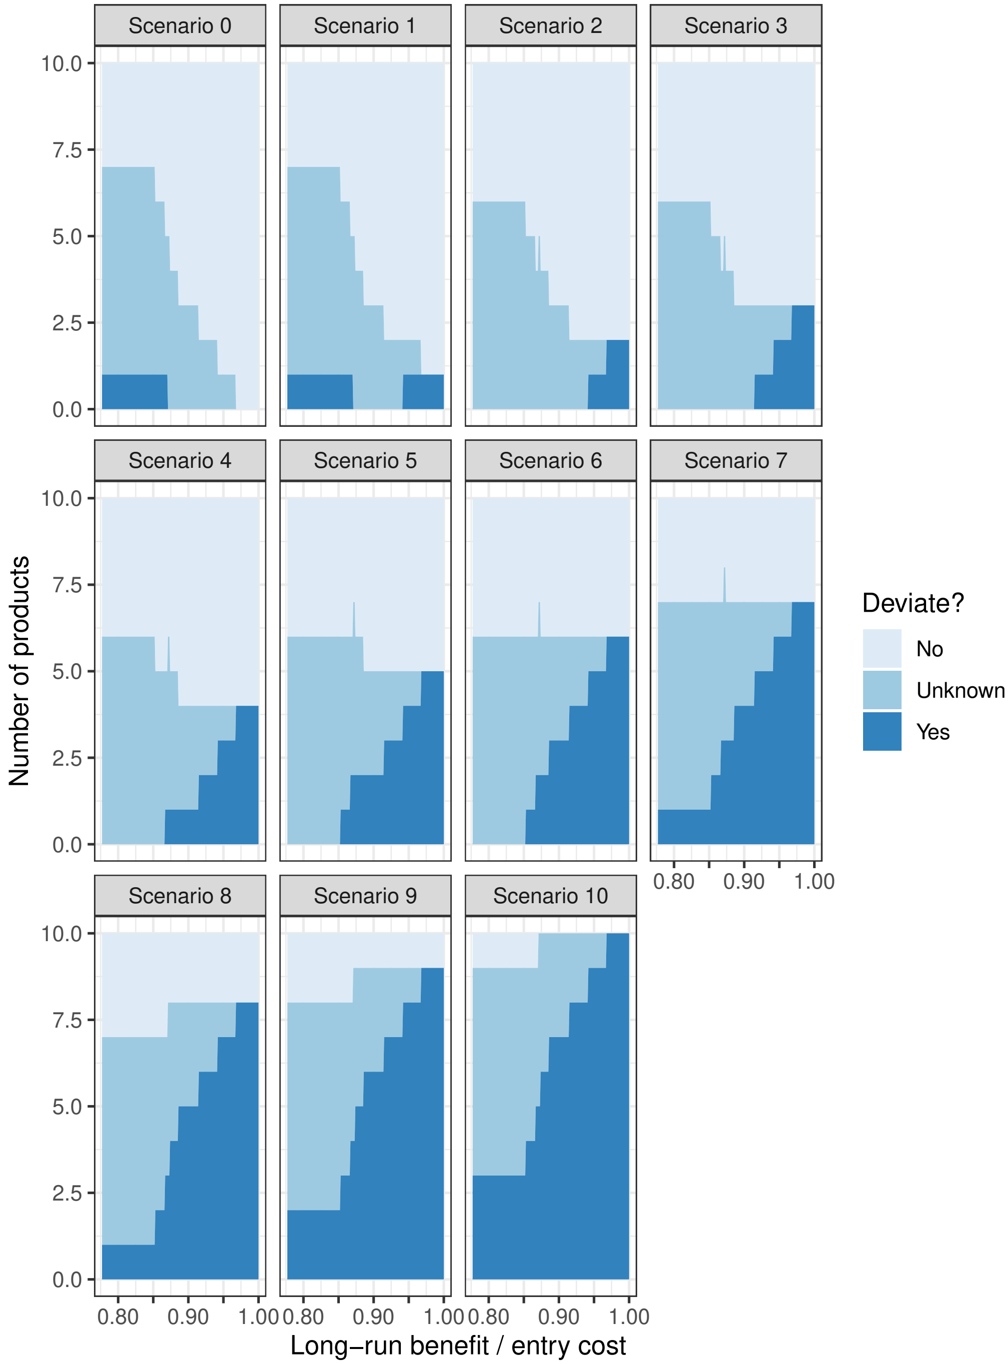

Figure 2.6: *Number of firms with incentives to deviate in each entry scenario, given product entry parameters*

Note: *Given a value of the ratio of long run benefit to entry cost κ, compute the profit each firm obtains by introducing or removing each non-native electric vehicle, and report whether it is positive. If the sign is ambiguous because of unknown entry cost, report unknown. Scenarios are labeled by the number of models that do not enter.*

2.5.3 Welfare effects

We now combine the scenario-specific welfare results with the equilibrium tests. For each value of the long-run benefit term κ, we show the range of outcomes across scenarios that are not rejected. (Except when κ is high, there are usually multiple.) EV sales and consumer surplus are higher in scenarios with more product variety; producer surplus, because of the entry cost term, is higher in scenarios with less product variety. The maximum consumer surplus loss is \$500 million and the maximum producer surplus gain is \$260 million. The maximum loss in total surplus, found at $\kappa = 0.85$, is \$490 million.

The product variety effect of removing the ZEV mandate translates into a modest decline in national EV sales, as shown in Figure 2.7. Among values of long-run benefit supported by our estimates, the maximum effect is a 6% decrease in EV sales (about 21,000 vehicles). As the long-run benefit increases and scenarios with fewer products are ruled out, the effect on sales decreases. The estimated impact on sales would be lower if we had estimated a large EV random coefficient in the demand model of Chapter 1: we would predict more consumers switching from the products that are no longer sold to other EVs, rather than to gas vehicles or the outside good.

We quantify the welfare effects of removing the ZEV mandate conditional on κ. The bounds across scenarios are shown in Figure 2.8. The signs of the effects do not depend on κ, even for values below the estimated lower bound: consumer surplus falls, due to reduced product variety; producer surplus rises due to avoided entry cost; and total surplus falls. As long-run benefit increases, increasing product variety toward its observed level and reducing the effective entry cost, consumer and producer surplus tend monotonically toward zero.

Producer surplus combines variable profit, entry cost, and long-run benefit. Since we only have bounds on entry cost, and therefore on producer surplus, we calculate the maximum across scenario-specific upper bounds and minimum across scenario-specific lower bounds. In addition, the results depend on κ in two ways: directly through the long-run benefit estimated for producers, and indirectly through the selection of which

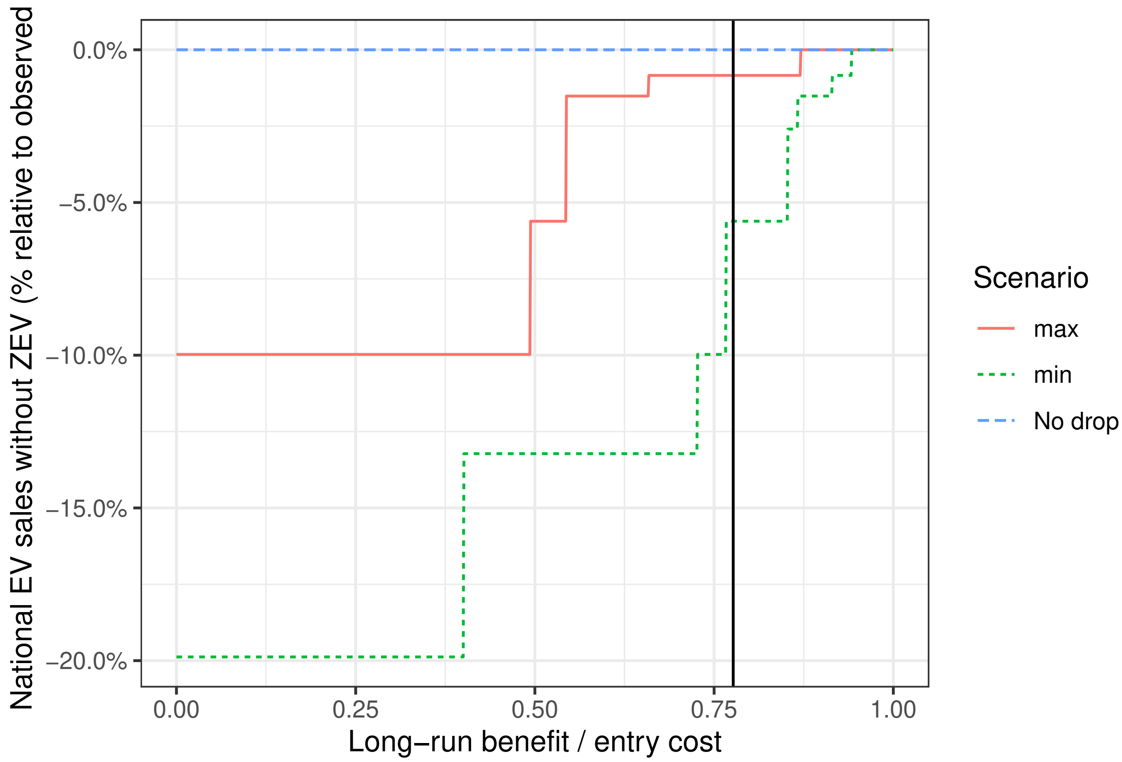

Figure 2.7: *Effect of removing ZEV mandate on national EV sales, across entry scenarios*

Note: *Given a value of the ratio of long run benefit to entry cost κ, compute total EV sales across scenarios where the n lowest-selling non-native models don't enter. Across the entry scenarios where no firm has an incentive to deviate (or the sign of the incentive to deviate is unknown), take the maximum and minimum. Report as a percentage difference from national EV sales in the observed equilibrium. Results are divided by the vertical line into values of κ rejected by observed entry (left) and values of κ compatible with observed entry (right).*

scenarios are potential equilibria. Because we have held prices fixed, variable profit is always lower when fewer products enter.

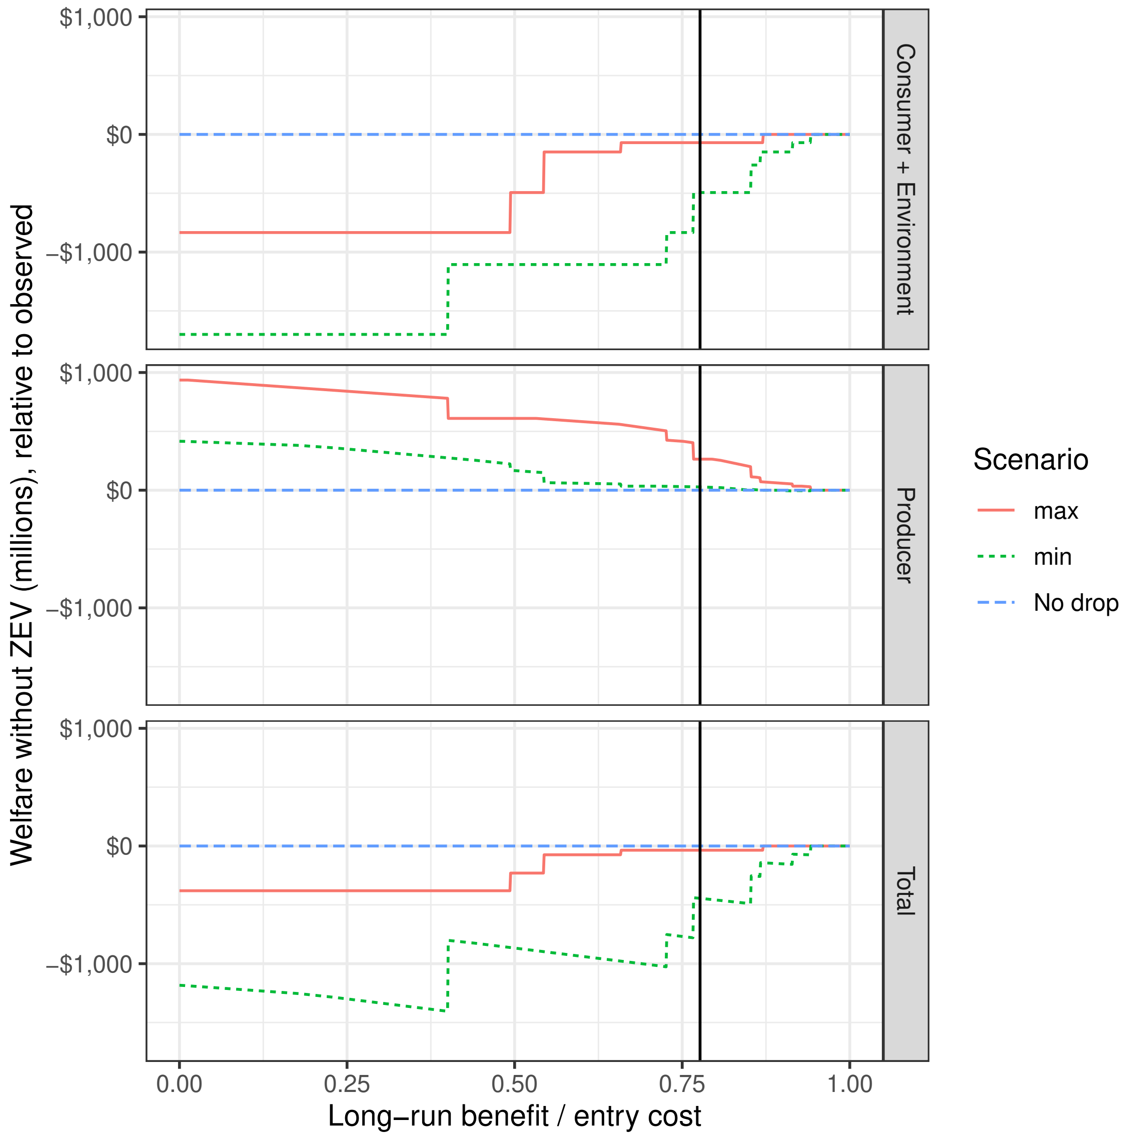

Figure 2.8: *Welfare effect of removing ZEV mandate, across entry scenarios*

Note: *Given a value of the ratio of long run benefit to entry cost κ, compute welfare across scenarios where the n lowest-selling non-native models don't enter. Across the entry scenarios where no firm has an incentive to deviate (or the sign of the incentive to deviate is unknown), take the maximum of the upper bound on welfare and the minimum of the lower bound. Report welfare as a difference from welfare in the observed equilibrium. The scenario where all products stay in is also shown. Results are divided by the vertical line into values of κ rejected by observed entry (left) and values of κ compatible with observed entry (right).*

Our calculation of producer surplus combines the difference in variable profit, avoided entry costs, and foregone long-run benefits. Let π_f^{cf} be firm f's counterfactual variable profit, and let $\mathcal{J}_f^{cf}$ be firm f's counterfactual product set. Then the difference in producer surplus is

$$\sum_f \left(\pi_f^{cf} - \pi_f + \sum_{j \in \mathcal{J}_f \setminus \mathcal{J}_f^{cf}} SC_j(1-\kappa) \right).$$

Because we do not have point estimates of SC_j or κ, we hold κ fixed and compute bounds on the difference in consumer surplus. For each $\kappa \in [\underline{\kappa}, 1]$, we use

$$SC_j \in [\underline{SC}, \min\{\overline{SC}, (1-\kappa)^{-1}\Delta_j \pi_{f(j)}\}].$$

2.6 Counterfactual demand-side policy

We now turn to the effects of the counterfactual demand-side policy from Chapter 1 on entry incentives, and ask how the welfare effect of the counterfactual policy varies when the product set is allowed to adjust. Consumer surplus falls further because of reduced product variety. Across the equilibria that are supported by our estimates of product entry parameters, the effect of endogenous entry on producer surplus ranges from a small negative effect to a small positive effect. In some equilibria, foregone long-run benefits exceed avoided entry costs, but in others they do not.

To assess the effect of product introduction, we compute these subsidy and tax levels under a variety of plausible entry scenarios. We then restrict attention to scenarios that could be equilibria in the product entry game, by checking if any firm can profitably deviate in its EV entry decision. In computing producer surplus, we report a range of estimates given the bounds on entry cost and long-run benefit from Section 2.4.2.

2.6.1 Social welfare and entry incentives

We next measure how much the firm captures of the externalities of product entry (consumer surplus, avoided environmental externalities, and changes in other firms' surplus, as described in Section 2.4.3). For each product j, we drop the product from the choice set, let market shares adjust (holding fixed GHG credit prices, the subsidy, and product prices), and compute the difference in consumer surplus and profits. The sum of entry externalities stays constant or rises slightly; because the credit term no longer includes ZEV credits, a smaller portion of the externalities is being captured.

We repeat this calculation across scenarios in which the n lowest-selling non-native electric vehicles avoid entering the market. Figure 2.9 shows the uncaptured social benefit, measured as the sum of the externalities of product entry minus the private gain from GHG credits. Across all the entry scenarios we consider, the uncaptured benefit is larger under the demand-side policy than under the mandate.

2.6.2 Effect on product entry

If the demand-side policy alters the variable profits earned from electric vehicle sales, it is possible that fewer products will enter. Figure 2.10 shows the change in product-level incremental variable profits under the baseline counterfactual in which only prices and quantities change. While some products become more profitable as a result of equilibrium price changes, the three non-native EVs with the lowest sales become less profitable, suggesting that they may exit in equilibrium.

We assess this hypothesis formally by examining the equilibria of the product entry game. We test a variety of product entry scenarios by computing variable profits and assessing whether any firm has an incentive to deviate: a product which is out but faces a positive entry incentive, or a product which is in and faces a negative entry incentive. We use the simplified model in Section 2.4.2, and adopt the specification in which product quality ξ is in all firms' information sets.

Product prices adjust to reflect the pass-through of the consumer subsidy and tax (as in Chapter 1) and account for the product set in the given entry scenario. The counterfactual subsidy and tax satisfy the quantity and budget-balance constraints under the new equilibrium prices and quantities.

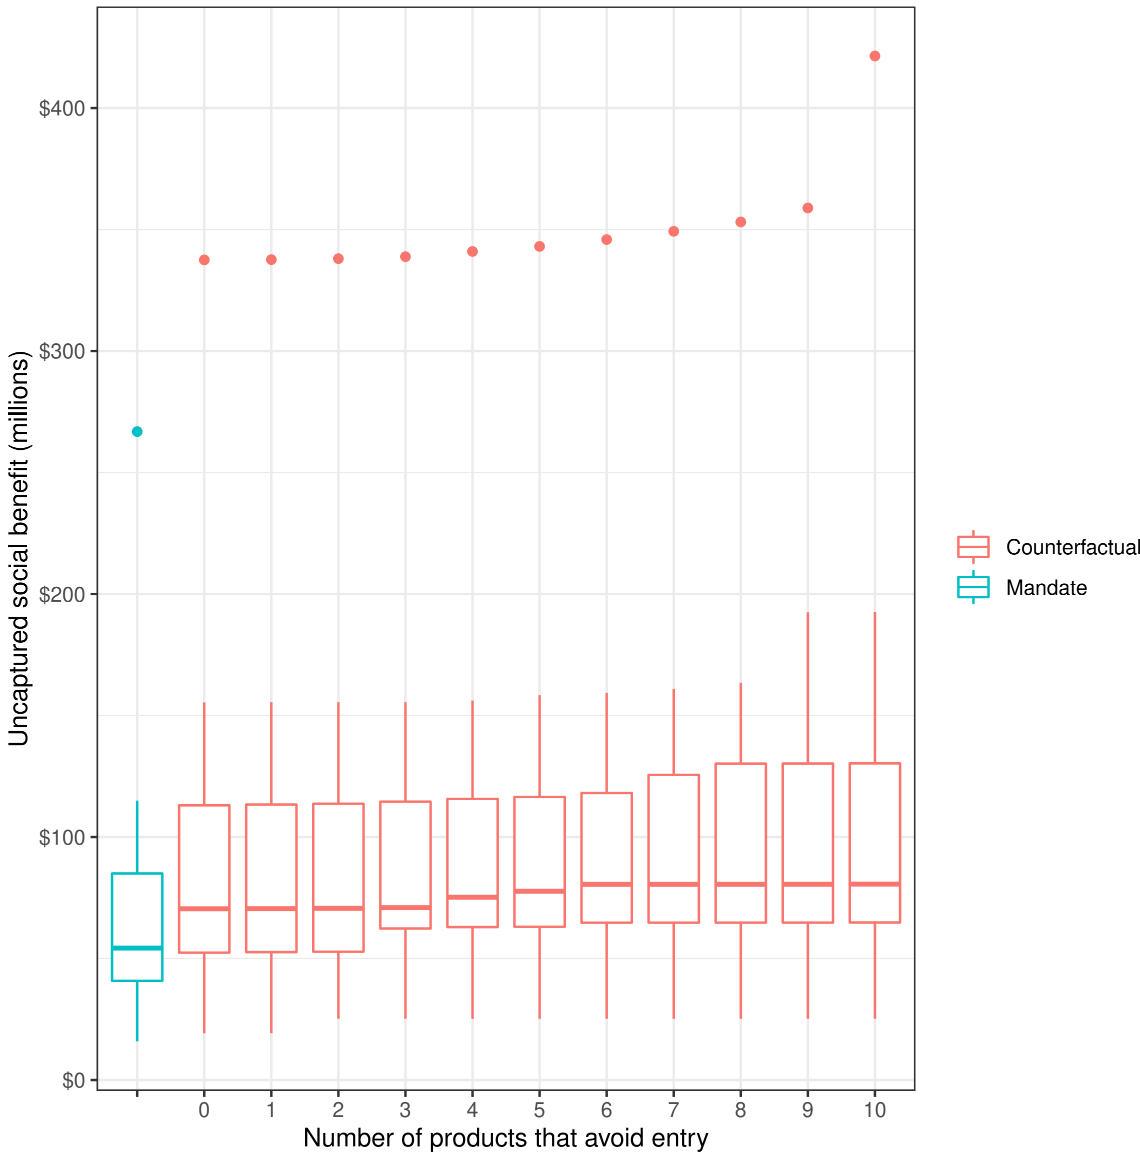

Figure 2.9: *Sum of entry externalities minus credit value, selected EVs, under counterfactual demand-side policy*

Note: *Credit value consists of GHG credits. Entry externalities are the change in consumer surplus, negative environmental damages, and other-firm profit resulting from entry. All calculations use estimated values of product quality ξ.*

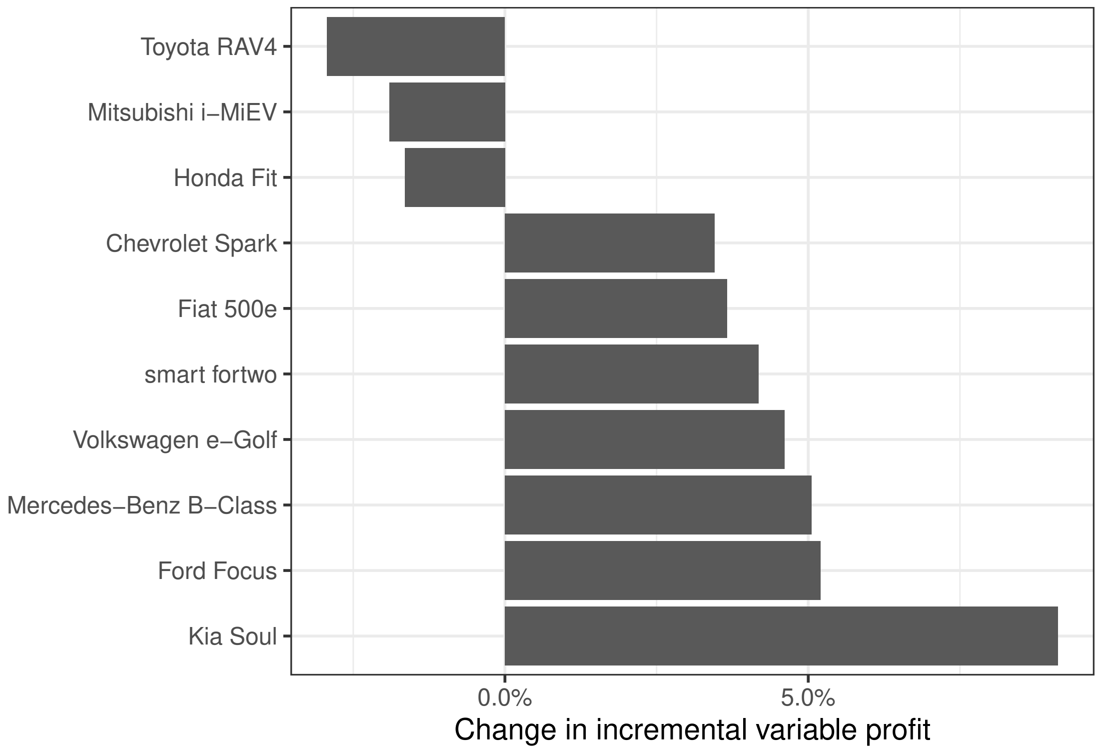

Figure 2.10: *Changes in incremental variable profit under counterfactual demand-side policy, baseline*

Note: *This plot shows how the incremental variable profit of non-native electric vehicles changes when ZEV is replaced by the counterfactual demand-side policy and the product set does not change. Some products become less profitable, suggesting that their manufacturers may have an incentive to avoid introducing them. Others become more profitable, as a result of the changes in equilibrium prices.*

Subsidy and tax amounts. In order to hold EV sales in the regulated states constant, the incentive for EVs needs to increase to replace the lost sales from the models that stay out. The right panel of Table 2.4 shows the subsidy and tax in the scenario in which four (of ten) non-native electric vehicles stay out of the market. The subsidy is up to 27% higher than in the scenario in which the product set is unchanged.

Table 2.4: *Counterfactual consumer subsidy and tax amounts*

	Product set unchanged		Product set reduced	
Model year	Subsidy	Tax	Subsidy	Tax
2009	\$0	\$0	\$0	\$0
2010	\$0	\$0	\$0	\$0
2011	\$3,566	\$17	\$3,566	\$17
2012	\$2,616	\$19	\$3,323	\$25
2013	\$3,697	\$83	\$3,850	\$87
2014	\$2,884	\$68	\$3,593	\$86
2015	\$2,356	\$76	\$2,479	\$80
2016	\$2,993	\$126	\$3,059	\$129
2017	\$2,316	\$110	\$2,402	\$114

Note: *This table shows the magnitude of consumer subsidy and tax, within the ten regulated states only, that achieves the same EV sales within the regulated states each year when the ZEV mandate is removed. (Greenhouse gas credits remain with prices unchanged.) The consumer subsidy shown is the amount before range multipliers, and is comparable to observed ZEV credit prices (see table in Chapter 1). The tax per non-EV is set by constraining the subsidy outflows and tax inflows to balance. Firms respond by resetting prices in Nash–Bertrand equilibrium. Run under both the scenario where the product set is unchanged and the scenario where the four lowest-selling non-native models drop out.*

Evaluating scenarios. As in Section 2.5, we consider a selection of entry scenarios and evaluate which are possible equilibria given a value of κ. Figure 2.11 shows the results. Entry scenarios in which the n lowest-selling non-native electric vehicles avoid entering, the scenarios with $n = 7, 8, 9, 10$ (of ten) are rejected for all $\kappa \in [\underline{\kappa}, 1)$. The remaining scenarios are rejected for some values of κ.

Effects on welfare. Reduced product variety decreases consumer surplus, variable profit for producing firms, and total entry costs, but increases variable profit for non-producing firms. The environmental effects of reduced variety depend on whether consumers who

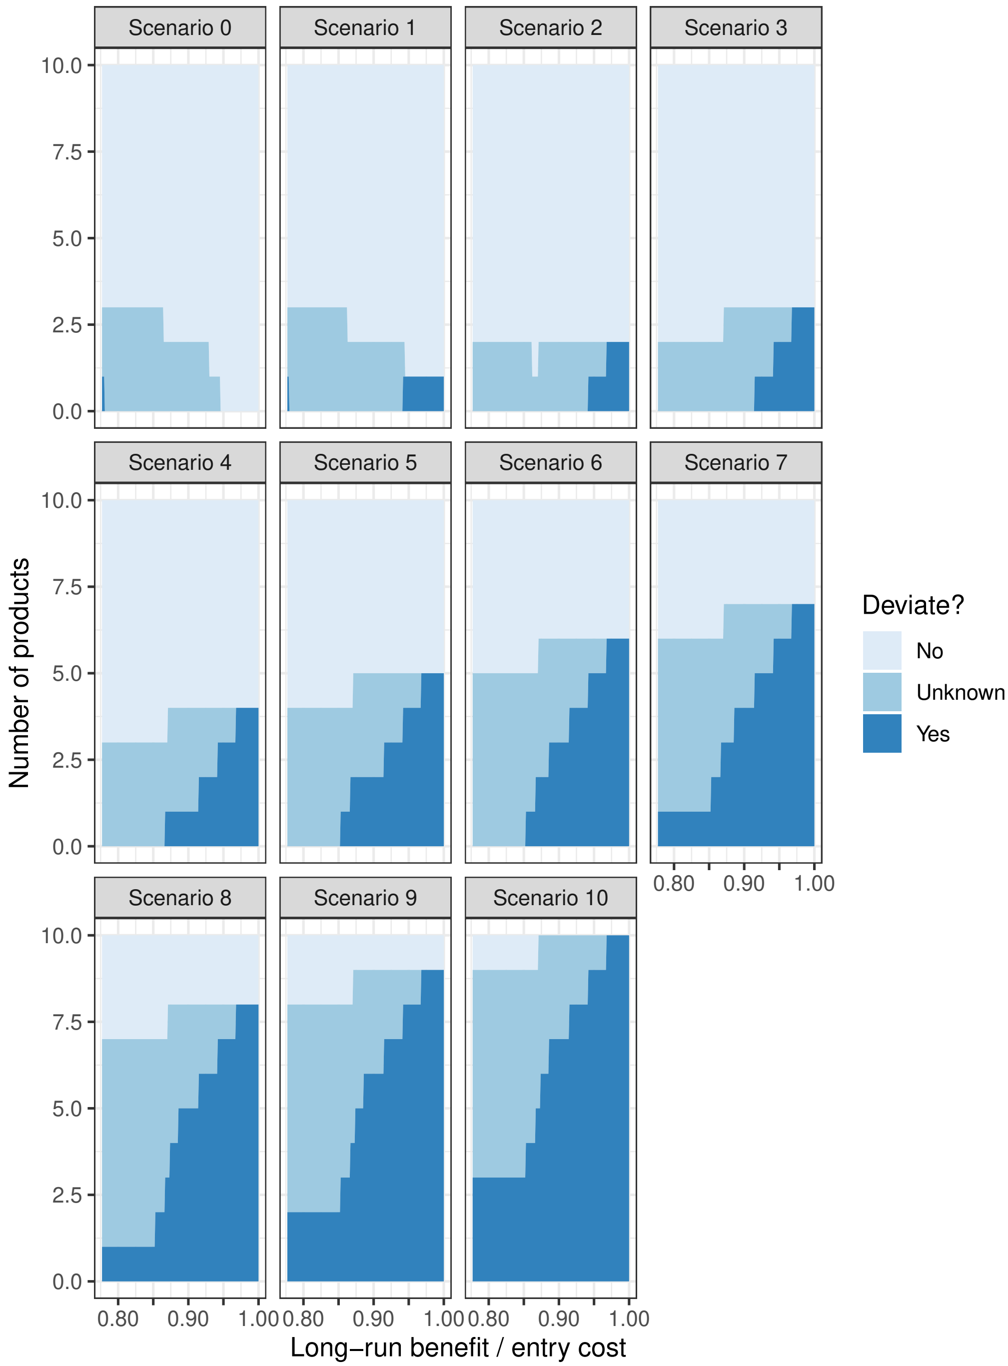

Figure 2.11: *Number of firms with incentives to deviate in each entry scenario, given product entry parameters*

Note: *Given a value of the ratio of long run benefit to entry cost κ, compute the profit each firm obtains by introducing or removing each non-native electric vehicle, and report whether it is positive. If the sign is ambiguous because of unknown entry cost, report unknown. Scenarios are labeled by the number of models that do not enter.*

would have purchased the electric vehicles that did not enter instead purchase more- or less-polluting alternatives. The effects of the subsidy and tax on welfare are theoretically ambiguous.

We quantify welfare effects using the model of product entry with the ratio of long-run benefit to entry cost ranging from $\underline{\kappa} = 0.78$ to 1. As in Section 2.5, we assume that entry costs and long-run benefits do not change in the counterfactual, and calculate the maximum upper bound and minimum lower bound on consumer and producer surplus across the scenarios that are not rejected. The results across the range of entry scenarios we consider are shown in Figure 2.12.

When long-run benefits are very low, producers gain by avoiding entry costs, while consumers lose from reduced variety. When long-run benefits are at the level implied by observed entry, $\kappa \geq \underline{\kappa}$, consumer surplus is between \$1.8 billion and \$1.4 billion lower and producer surplus is between \$420 million and \$680 million higher, for a total loss of between \$1.3 billion and \$750 million from switching to the demand-side policy. When long-run benefits are high, there is no response on the entry margin.

2.7 Conclusion

We examine the effects of the ZEV mandate, an influential state-level supply-side environmental policy in early generations of the US electric vehicle market, on product variety in electric vehicles. The mandate allowed entrant firms to capture more of the social surplus generated by their products. We find evidence that a number of additional electric vehicles entered the market as a result of the mandate, with a modest positive effect on electric vehicle sales and consumer surplus. These findings have consequences for future state policies to encourage new, socially beneficial types of consumer products within national markets, and for understanding the effects of current and future electric vehicle policy.

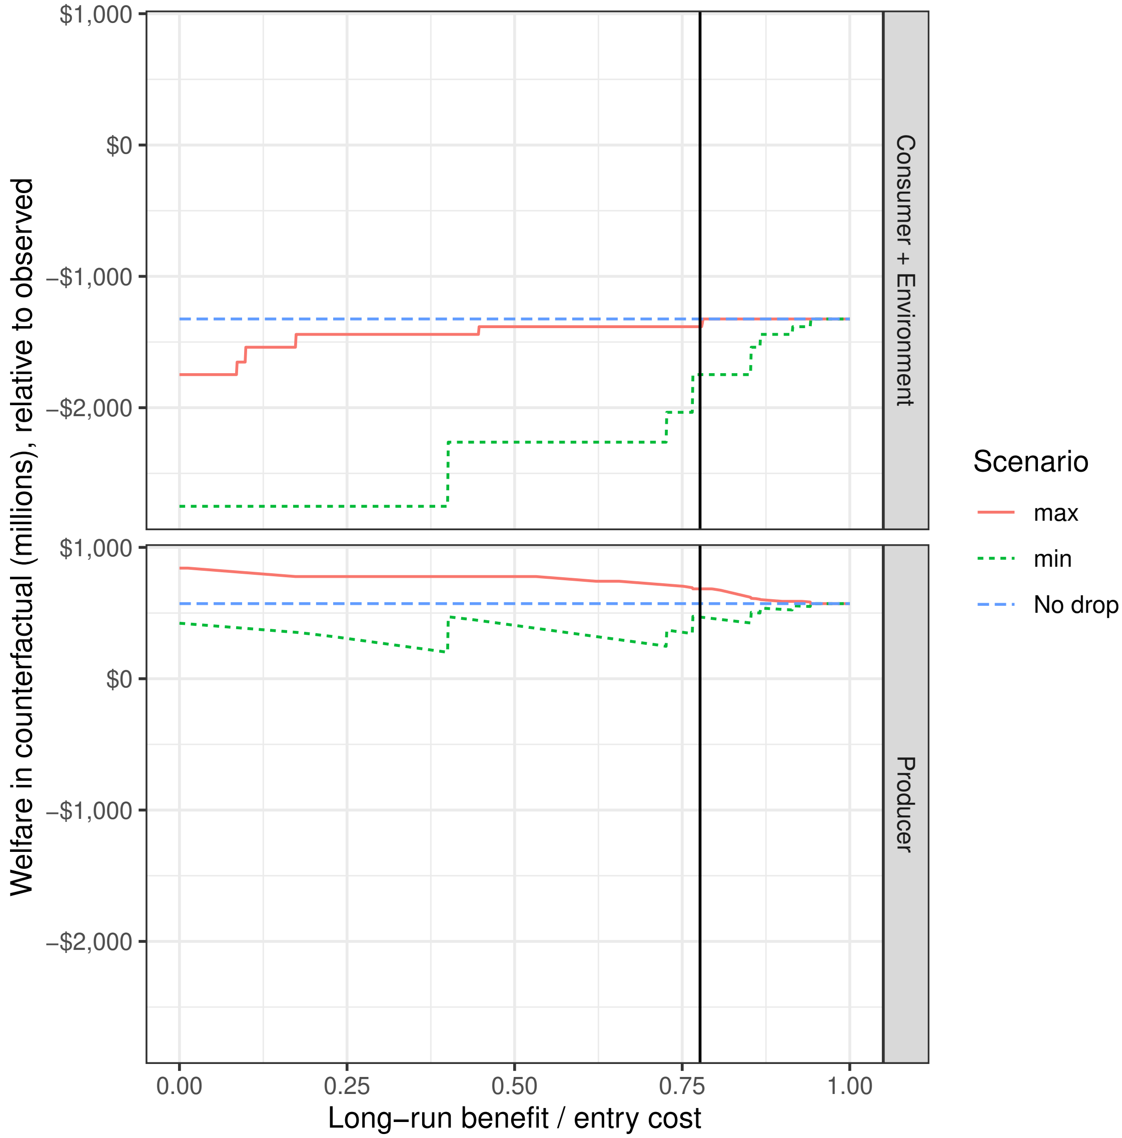

Figure 2.12: *Welfare effect of replacing ZEV mandate with demand-side policy, across entry scenarios*

Note: *Given a value of the ratio of long run benefit to entry cost κ, compute welfare across scenarios where the n lowest-selling non-native models don't enter. Across the entry scenarios where no firm has an incentive to deviate (or the sign of the incentive to deviate is unknown), take the maximum of the upper bound on welfare and the minimum of the lower bound. Report welfare as a difference from welfare in the observed equilibrium. The scenario where all products stay in is also shown. Results are divided by the vertical line into values of κ rejected by observed entry (left) and values of κ compatible with observed entry (right).*

Chapter 3

Credit Trading Markets for Vehicle Emissions Regulation: Theory and Evidence

3.1 Introduction

Contemporary environmental policies in the US are often structured using credit-based systems, in which a regulator awards firms credits for environmentally beneficial actions and takes credits away for environmentally destructive ones. Such systems use market mechanisms to improve cost efficiency, take advantage of firms' private information about costs, and avoid the political roadblocks to explicit subsidy and tax policies. Prominent examples include the supply-side regulations on the vehicle market from Chapters 1 and 2: the state-level Zero Emission Vehicle (ZEV) mandate for electric vehicle sales and the federal greenhouse gas (GHG) vehicle regulation.

The market mechanisms that set these regulations apart from traditional standards are credit trading and credit banking. Credit trading addresses heterogeneity across firms, allowing the firms with the lowest abatement costs to bear emissions reductions and be compensated for it by the highest-cost firms. Credit banking addresses heterogeneity across time, allowing firms to undertake additional emissions reductions in times when abatement costs are low.[1]

In Chapters 1 and 2, I use a simple model of firm interactions with credit trading and banking, in which all firms take a credit's market price from the data as the shadow price when making pricing decisions. I do not model the determinants of the credit price itself. Without more structure, this approach severely limits my ability to study counterfactual supply-side policy designs, ignores interactions among policies, and limits my ability to study the effects of large-scale changes (for example, eliminating high-sales electric vehicle models from the product set) on credit markets.

Unlike traditional targets of credit-based regulations, vehicle markets are characterized by a small number of multiproduct firms, whose products typically vary in pollution intensity. Investments, such as in new products and major technological changes, often affect markups and consumer surplus in addition to traditional abatement costs. In the context of electric vehicle policy, the design of credit markets has at times been controversial, as observers have claimed that large credit accumulations by firms with minimal electric vehicle programs undermine policy goals.[2]

In this chapter, I address this gap by introducing a model that explains the prices and quantities at which credits trade and how these outcomes depend on policy design. The model is dynamic to capture credit banking. Firms remain price takers, but face uncertainty about future production conditions and a possibility that the credit market will be temporarily unavailable in the future. As a result, heterogeneous future production conditions become an important determinant of firm strategies in the credit market and therefore of credit prices.

I begin by describing the Zero Emission Vehicle (ZEV) mandate and greenhouse gas

[1]Credit *banking* usually refers to a firm's ability to undertake additional emissions reductions early on and keep the credits for later use. Full flexibility across time is provided when credit banking is combined with credit *borrowing*, which allows a firm to run credit deficits in order to defer emissions reductions.

[2]See "California's zero-emission vehicle program is stuck in neutral" (Rory Carroll, Alexandria Sage, Reuters, 9/1/16).

(GHG) regulation, the structure of firm requirements and credit markets, and trends in firm compliance strategies from data collected from regulatory reports. I then document two stylized facts shared by both policies. First, firms purchase credits even when they already have large balances on hand, and larger balances are associated with larger purchases. Second, credit trading, in aggregate, results in a less concentrated distribution of credit balances across firms.

I then describe the model and show how the prices and quantities of credits are determined as a function of product-market profit functions and firm beliefs about the future. Within the model, the stylized facts are explained by firm heterogeneity in production conditions (either costs or demand), which give firms incentives to choose different quantities of credits to sell or buy.

I conclude by outlining the implications of the model for policy design. I consider two types of policies: one that affects credit trading, by providing a safety valve, and one that affects credit banking, by introducing credit appreciation or depreciation. I find that the effects of both policies depend on beliefs about future policies and production conditions.

3.1.1 Literature

Credit-based environmental policies were proposed in the 1960s and 1970s as alternatives to command-and-control policies. Instead of regulators choosing standards firm-by-firm, market mechanisms would choose the allocation of emissions across firms subject to a fixed total level of emissions (Tietenberg 1980). An early large-scale example was the 1990 program to control fixed-source sulfur dioxide emissions, such as from power plants (Schmalensee and Stavins 2013). Stavins (2003) surveys early credit-based policies.

Surveys of credit trading markets in the greenhouse gas and zero emission vehicle regulations include Leard and McConnell (2017), which studies federal Corporate Average Fuel Economy (CAFE) and GHG credit markets in the 2010s; Leard and McConnell (2021), which discusses the interactions between GHG and ZEV; and McConnell and Leard (2021), an overview of the ZEV program. Work on potential future credit-based policies to phase out gas-powered vehicles includes Holland, Mansur, and Yates (2021) and Cole, Droste, Knittel, Li, and Stock (2021), which models a future nationwide ZEV mandate.

The role of uncertainty in determining outcomes in credit-based systems, and the relative advantage of quantity-based credits over emissions taxes, is the subject of a theoretical literature originating in Weitzman (1974). Newell, Pizer, and Zhang (2005) study enhancements to credit systems that better accommodate cost uncertainty. Empirical work on uncertainty includes Borenstein, Bushnell, Wolak, and Zaragoza-Watkins (2019), which examines the California greenhouse gas cap-and-trade market (which applies to fuel but not directly to vehicle sales).

Theoretical work has shown that deviations from perfect markets typically reduce, and sometimes reverse, the benefits of allowing credit trading. If the relevant product market is not competitive, allowing credits to be traded can be welfare-reducing (Malueg 1990; Sartzetakis 1997; Sartzetakis 2004). If transacting in the credit market is costly, credit trading may be less efficient than it would otherwise be (Stavins 1995). If firms have market power in the credit market, dominant firms with surplus credits will restrict credit sales in order to raise the equilibrium credit price (Hahn 1984) and may also restrict credit sales to raise the costs of product market rivals (Misiolek and Elder 1989; von der Fehr 1993); however, private information about costs mitigates these market power effects (Malueg and Yates 2009).

A separate theoretical literature has addressed credit banking. Rubin (1996) and Cronshaw and Kruse (1996) provide conditions under which credit banking and trading are optimal. Kling and Rubin (1997) shows that fully flexible credit banking may produce a dynamically inefficient path of emissions.

Banking and trading have both been studied empirically, though usually separately. For the US sulfur dioxide program, Carlson, Burtraw, Cropper, and Palmer (2000) evaluates credit trading and Ellerman and Montero (2007) evaluates credit banking. Within the vehicle context, Rubin and Kling (1993) uses estimated cost functions to simulate the effect of credit banking within California regulations on smog-forming pollutants. Rubin, Leiby, and Greene (2009) calibrates a Cournot model of the US auto industry and finds that allowing trading of CAFE credits reduces costs, even if the credit trading market is imperfectly competitive. Durrmeyer and Samano (2018) uses a structural model of the French automobile industry to compare a perfectly competitive fuel economy credit market to firm-level standards and a consumer subsidy-and-tax ("feebate") scheme.

3.2 Background on the ZEV and GHG programs

In both the ZEV and GHG programs, automakers earned and lost credits according to simple formulas that are easy to interpret and incorporate into a model of firm behavior. Specifically, each vehicle was assigned a credit value based on its characteristics (electric/non-electric and battery range for ZEV; fuel economy and vehicle footprint for GHG), and earned (or lost) that value for every vehicle that was sold. The value was positive for some vehicles and negative for others, so a manufacturer could comply using only its own products in each year, or take advantage of inter-firm and inter-temporal heterogeneity by trading and banking credits.

The largest difference between the programs is that the ZEV mandate underwent major change in 2018, while the statutory changes to the GHG regulation over time were relatively minor.[3] The programs also differ in the exact values they assign to each vehicle and in minor rules relating to banking and borrowing over time. In this paper, I consider each program in isolation, though they interact in potentially meaningful ways (Leard and McConnell 2021).

3.2.1 State-level Zero Emission Vehicle mandate

As described in Chapter 1, the state-level Zero Emission Vehicle mandate provided a credit trading system to encourage the sales of electric vehicles in the ten ZEV states, even among firms that were too small to be subject to the requirement. Firms earned

[3] Automakers faced uncertainty about the future paths of both programs. The ZEV mandate would have been blocked, and the GHG regulation significantly loosened, if the Trump administration's Safer and Affordable Fuel Efficient (SAFE) Vehicles Rules had taken full effect. See "Trump administration to freeze fuel-efficiency standards and fight states" (Juliet Eilperin, Brady Dennis and Michael Laris, The Washington Post, 7/24/18).

credits by selling electric vehicles, which they could bank for the future or trade with other manufacturers.[4] The largest manufacturers then used the credits to meet requirements each year, which were set as a function of their sales in the ZEV states. Chapter 1 contains an overview of the structure of credit earning and requirements.

Firms faced penalties for carrying deficits for extended periods of time. Credit balances could run negative for two years without penalty; after this point deficits were subject to a penalty of $5,000 per credit. (From 2018 onward, credit balances could only run negative for one year without a penalty.) In the period for which data are available (2009–2019), no firms faced penalties. Only one firm had a deficit, which it made up the following year.[5]

Credits were typically traded in ad-hoc bilateral transactions between manufacturers. (By law, only vehicle manufacturers could participate in credit trading.) Tesla signed longer-term contracts in its early years to sell credits for fixed prices as they were earned, but those agreements ended around 2012. For example, Honda agreed to buy credits from the first 650 vehicles Tesla sold between 2009 and 2011 for a fixed price.[6] Afterward, credits were traded in blocks for prices negotiated at the time of sale.

Some further complexities of the regulation are important for understanding patterns in the credit data. Most importantly, automakers could use excess ZEV credits toward the requirements of the partial zero emission vehicle (PZEV) mandate, a similarly designed policy that encouraged low-emission gas vehicles, hybrids, and plug-in hybrids. The PZEV mandate applied both to the large manufacturers that faced the ZEV mandate and to a group of intermediate-volume manufacturers. For most manufacturers, ZEV credits were scarce while PZEV credits were abundant, and PZEV credits were rarely traded. Nonetheless, the credit data contain examples of intermediate-volume manufacturers purchasing and using ZEV credits, which can only make sense if they used the credits toward their PZEV

[4] In this chapter, I use "electric vehicle" to refer to the regulatory classification of "zero emission vehicle" (ZEV), which covers both battery electric vehicles and hydrogen fuel cell vehicles.

[5] This firm was Jaguar Land Rover in 2014, which had deficits in California and New Jersey. Because Jaguar Land Rover was an intermediate volume manufacturer, this deficit came from its use of ZEV credits to comply with the partial-ZEV mandate, as discussed below.

[6] The text of the Toyota–Honda agreement, with prices redacted, was obtained from Tesla's 2010 IPO filings.

requirements.

Another related complexity is the mechanisms by which the regulations interacted across states. Formally, the regulation applied state-by-state: each automaker faced a requirement based on its statewide sales, and could only meet that requirement using ZEV credits earned in the state.[7] As a result of the travel provision, however, the regulation effectively grouped the states together until 2017: an EV sold in any ZEV state earned credits in all of them. As a result, California data gave a complete picture of compliance across all states, and credit prices had no reason to vary based on where the credit was earned. From 2018 on, the pooling provision allowed some credits to be moved between states, but not into or out of California. As a result, the ZEV states were effectively separated into two groups: California and the other states.

Timeline of mandate

The ZEV mandate can be divided into three phases.[8] The first phase, covering model years 2005 to 2008, preceded the widespread availability of electric vehicles; automakers earned credits for low-speed neighborhood vehicles, demonstrations, and fleet sales. The second phase, from 2009 to 2017, saw automakers introduce first-generation electric vehicles, and total EV sales generally exceeded regulatory requirements. The third phase, from 2018 onward, features tighter rules and applies to more manufacturers.

In the first phase, automakers received many more credits for their limited EV projects than were required for compliance. In addition, credit balances appreciated over time until 2010.[9] As a result, automakers entered the second phase with large credit balances. However, a one-time provision of the regulation forced credits earned up to 2011 to expire,

[7] Under Section 177 of the Clean Air Act, only California can design its own vehicle emissions regulation. Other states can choose to adopt a California regulation, but cannot alter it.

[8] The rules for the ZEV mandate are in Title 13 of the California Code of Regulations. The first phase corresponds to Section 1962, the second phase to Section 1962.1, and the third phase to Section 1962.2.

[9] This appreciation was accomplished by setting the unit of account for a credit to be the California local pollution requirement ("g/mi NMOG"), which changed over time. For example, a credit earned in 2005 was worth 0.049 g/mi NMOG, while a credit earned in 2010 was worth only 0.035 g/mi NMOG.

in phases, between 2012 and 2014. (After this point, ZEV credits did not expire.) As a result, in order to meet requirements from 2015 on, automakers needed to sell new EVs or buy credits from other manufacturers.

Changes that came into effect in model year 2018, which were proposed in 2012 as part of California's Advanced Clean Cars Program, made credits harder to earn, raised credit requirements, and ended the travel provision, all while leaving balances intact. As a result, automakers likely had an incentive in 2013–2017 to build up balances to smooth the transition. The largest effect was on the five automakers that were reclassified from intermediate-volume to large-volume, which were now subject to the ZEV credit requirement. For all manufacturers, changes to the earning formula reduced the number of credits a given EV would earn, and the end of the travel provision (except for fuel cell vehicles) meant that manufacturers needed to sell EVs in ZEV states outside California.[10]

In addition, the PZEV mandate was removed and replaced with a mandate specifically for plug-in hybrids ("transitional zero emission vehicles" or TZEVs). Because not all intermediate-volume manufacturers produced plug-in hybrids before 2018, this change created an incentive for intermediate-volume manufacturers to purchase ZEV or TZEV credits and bank them.

One provision in the 2018 changes made compliance easier for manufacturers with low EV sales and a low-emissions gas fleet. Between 2018 and 2021, an automaker that over-complied with California's greenhouse gas regulation (which was equivalent to, but administered separately from, the EPA greenhouse gas regulation described in Section 3.2.2) could convert excess credits from that program into ZEV credits. According to regulator disclosures, in model year 2020, Honda was the only automaker to take advantage of this provision.

[10] The credit requirement was also redefined from a percentage of non-EV sales to a percentage of total sales, effectively lowering the credit earning from a given EV even further.

3.2.2 Federal greenhouse gas regulation

The second policy I study is the US Environmental Protection Agency's requirement on auto manufacturers to limit the greenhouse gas emissions of their new vehicles. The EPA GHG regulation was deployed starting in 2012, with an optional phase-in period from 2009–2011, and was developed in coordination with an overhaul to federal CAFE standards to allow credit trading among firms.[11] A manufacturer earned or lost credits for each vehicle sold according to the difference between the vehicle's greenhouse gas emissions and a target based on the vehicle's footprint.[12] A manufacturer aggregated these credit gains and losses, denominated in units of megagrams of CO_2, across its entire fleet. Credit balances could temporarily run negative (that is, firms could borrow from the future). Firms could buy and sell credits freely in bilateral transactions, but a firm could not sell more credits than it had available in its balance.

Specifically, a product j in model year t had an associated target target_{jt}, calculated as a function of vehicle footprint, and an emissions rating emissions_{jt}, calculated as a function of fuel economy. (Electric vehicles were assigned emissions of zero.) Denoting its sales as sales_{jt}, the manufacturer's credit earning or loss in year t was therefore

$$\sum_j (\text{emissions}_{jt} - \text{target}_{jt}) \cdot \text{sales}_{jt}.$$

The allowance for larger footprints makes the GHG regulation an example of attribute-based regulation, as discussed in Anderson and Sallee (2016). A supplementary process allowed manufacturers to earn credits for other reductions in greenhouse gas emissions not reflected in fuel economy, particularly improvements to air conditioning.

Compared to the ZEV mandate, the GHG regulation applied to a larger set of manufacturers, targeted a metric that was closer to the CAFE standards that had existed for decades, and featured more limits on credit banking. All manufacturers with over 5000 vehicles sold

[11] The EPA GHG and post-2012 CAFE regulations are described in detail in Leard and McConnell (2017). The two policies are tightly linked because tailpipe greenhouse gas emissions are proportional to gasoline consumption.

[12] Footprint, defined as wheelbase times track width, is a measure of vehicle size.

in the US faced the regulation, though some manufacturers that sold up to 400,000 vehicles received additional allowances. Automakers could borrow credits from up to three years in the future (that is, could run deficits for up to three years). Credits generally expired five years after being earned, although a one-time exception allowed credits earned in 2010–2015 to be valid until 2021.

Tesla was a major seller of GHG credits, but not as dominant as it was in the ZEV credit market. (Like the ZEV market, only vehicle manufacturers could participate.) According to Tesla financial statements, it sold GHG credits on long-term contracts, rather than in ad-hoc transactions. In most years, all its sales were to one automaker, Fiat Chrysler; in 2019, it also sold credits to GM.

For estimates of GHG credit prices I rely on Leard and McConnell (2017), which use 2012–2014 data to place the value of a credit at $35–40 per megagram of CO_2.

The requirements of the GHG regulation progressively tightened on a fixed schedule that was announced well in advance. Nonetheless, market participants had some uncertainty over future requirements for political reasons. For example, GM purchased credits in 2019 as a "hedge" against a Democratic administration tightening requirements in 2021.[13]

3.3 Patterns in regulatory credit data

The regulators that enforce the ZEV and GHG programs publish comprehensive data on manufacturer compliance, particularly credit balances and trades between automakers. I use those data first to describe the major trends and automaker compliance methods, and then collect a set of stylized facts.

3.3.1 Sources of credit data

Data on the ZEV mandate comes from state regulators. I use public reports from the California Air Resources Board, which contain a full record of credit balances from model

[13]See "Tesla's Secret Source of Cash Unmasked as GM, Fiat Chrysler" (Miles Weiss and David Welch, Bloomberg, 6/3/19).

years 2009 through 2019, credit trades between manufacturers from 2010 through 2019, and vehicle production volume (excluding ZEVs) from 2009 through 2019.

Regulators only observe quantities of credits traded, not prices. To obtain average prices I combine Tesla's credit sales by model year, across all ZEV states, with Tesla's quarterly reports of the revenue earned from credit sales. To get credit sales outside of California, I incorporate comparable data from regulators in other ZEV states, which combine public reports and public records requests to the relevant agencies. These data cover almost all states and years but exclude some smaller states in earlier years. The calculation of ZEV credit prices is provided in in Chapter 1.

The data I use on the EPA GHG regulation comes from annual publications by the EPA, entitled "Greenhouse Gas Emission Standards for Light-Duty Vehicles: Manufacturer Performance Report" (model years 2012–2016) and "The EPA Automotive Trends Report" (model years 2017–2019). Of the data provided in those publications, I use manufacturer-level credit balances and the aggregate trades by each manufacturer. By tracking reported aggregate trades across editions of the report, I can estimate when the trades occurred. I often exclude data from Hyundai and Kia, whose GHG data were unavailable from 2012 to 2014 while they were under investigation for reporting inaccurate emissions test results.[14] Data on small manufacturers is also limited in some years.

In both California's and the EPA's public disclosures, regulators only reveal aggregate credit trades by manufacturer and year, not the identities of the buyer and seller in each bilateral transaction. Therefore, in years with multiple buyers and multiple sellers, it is impossible to know for sure which pairs transacted with each other.[15] That said, the quantities found in the data generally allow for reasonable guesses; reconstructions on those lines can be found in Appendix B of McConnell, Leard, and Kardos (2019) for ZEV, and

[14]The 2014 settlement that ended the investigation required both firms to forfeit a portion of their GHG credits.

[15]For a simple example, suppose the disclosures state that Tesla sold 15 credits, Nissan sold 5 credits, Toyota purchased 12 credits, and Honda purchased 8 credits. The exact bilateral transactions that generate these aggregates cannot be uniquely identified.

Table 3 of Leard and McConnell (2017) for GHG.

To provide context on trends in the market for new vehicles, I use registration data purchased from IHS Markit. I give more information on these data in Chapter 1.

3.3.2 Trends in ZEV credit data

Trends in the ZEV regulation were dominated by the effects of the rapid increase in EV sales: as shown in Figure 3.1, the EV market share went from near zero in 2009 to over 1% in the ZEV states by 2017. Over the same period, as shown in Figure 3.2, aggregate credit balances rose rapidly as EV sales growth outstripped the industry's ZEV requirement. Though Tesla's sales of EVs were responsible for much of this aggregate credit accumulation, the same trend was present without including Tesla's credits. In fact, the combined balance of the six largest manufacturers would still have grown even if their credit trades did not occur, as these manufacturers' own sales of EVs outpaced requirements.[16]

The six largest automakers' firm-level balances are shown in Figure 3.3. All firms see a general pattern of rising balances. Within this group, firms take heterogeneous approaches to the credit market. Ford and Honda purchase credits in order to stay in compliance. Fiat Chrysler and GM take the same approach in early years, but diverge once their EV sales rise: Fiat Chrysler continues purchasing credits (and then sells some), while GM ends credit transactions. Toyota slowly runs down its earlier balances, and supplements its credit bank with additional credit purchases, including a substantial purchase in 2017 from Tesla.

The firm-level balances for the next five automakers, which were reclassified from intermediate to large and subject to the ZEV regulation starting in 2018, are shown in Figure 3.4. Unsurprisingly, the balances for these firms are much smaller.[17] In contrast to the largest automakers, this group of automakers uses the credit market rarely or not at all,

[16]The measure I construct of credit balances without trade, shown in Figure 3.2, only eliminates trades of credits earned in California. It may still include credits that a manufacturer purchased outside of California and then added to its California balance using the travel provision.

[17]The credit balance for BMW understates its ability to comply with the regulation. BMW was the only automaker to use a provision of the regulation that allowed credits earned from certain plug-in hybrids ("Range Extended Battery Electric Vehicle" or "BEVx") to offset its ZEV requirements.

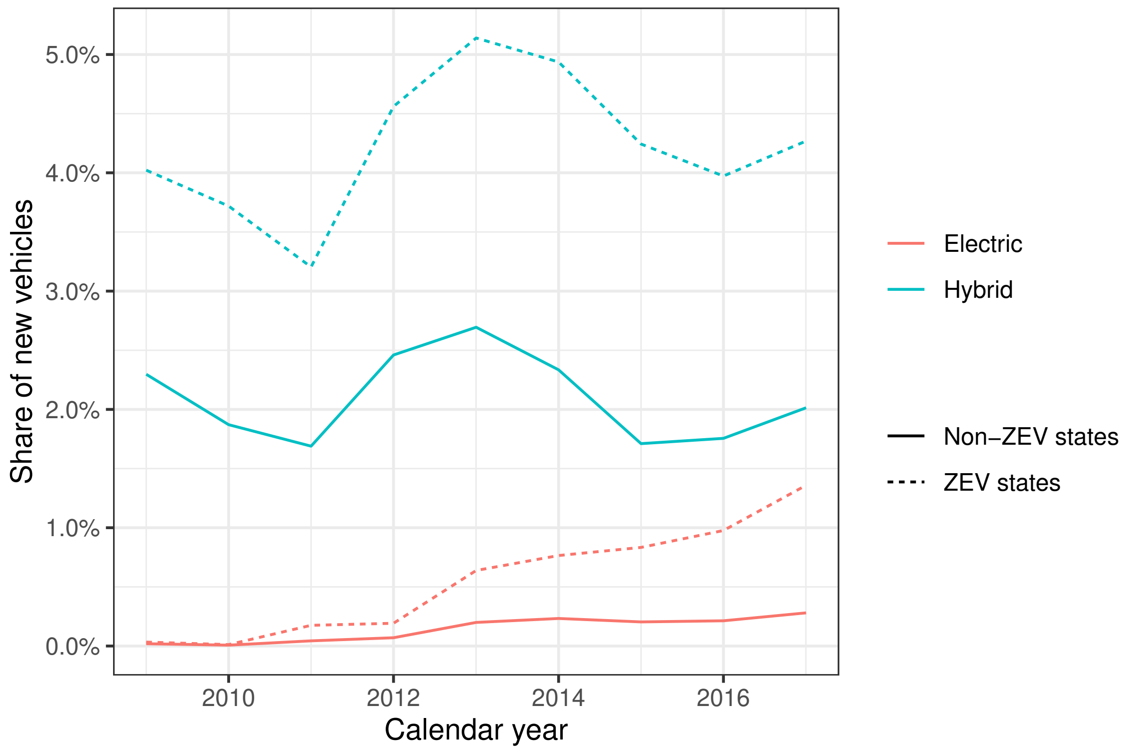

Figure 3.1: *Market share of EVs and hybrids by region, 2009–2017*

Note: *This figure shows the share of new vehicle registrations that were EVs or hybrids, over time from 2009 to 2017, broken out by the ten ZEV states vs. the rest of the US. Data come from IHS. (Includes content supplied by R. L. Polk; Copyright © R. L. Polk, 2019. All rights reserved)*

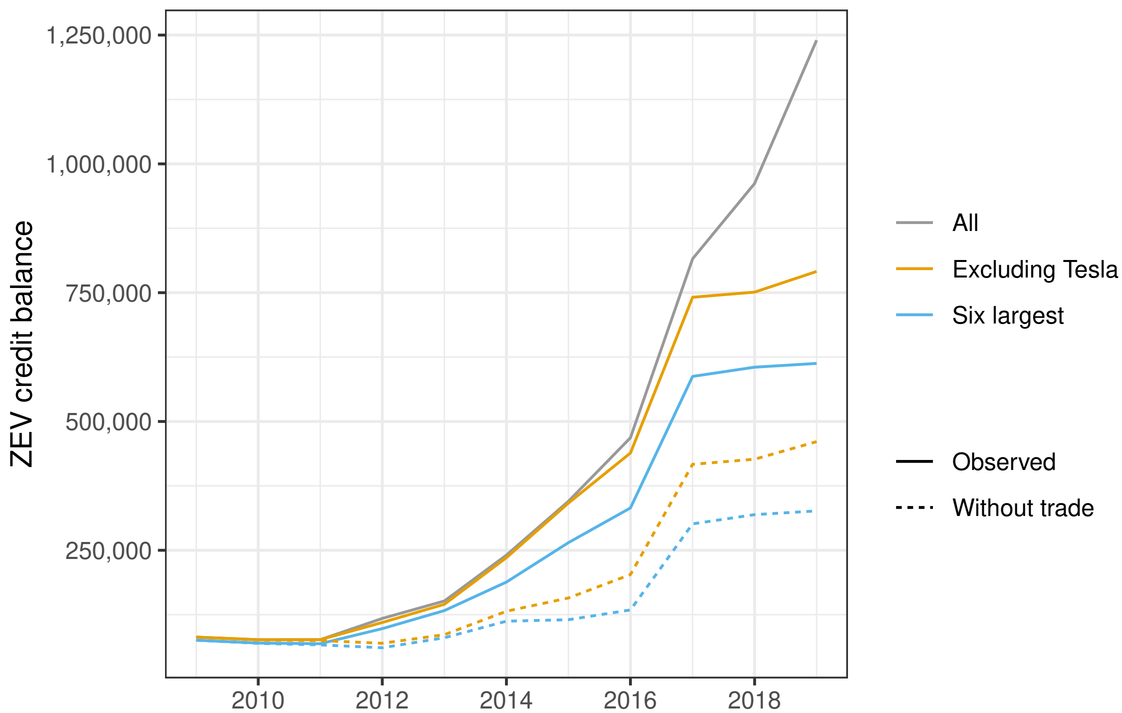

Figure 3.2: *Aggregate ZEV credit balances over time*

Note: *This figure shows the California ZEV credit balances left at the end of each model year, across groups of manufacturers. The 'six largest' group consists of the six manufacturers that were subject to the mandate for the entire period (Chrysler/Fiat Chrysler, Ford, GM, Honda, Nissan, and Toyota). The solid line shows observed balances, and the dashed line shows balances with trades of California credits taken out. (Trades in other states, which can affect California balances through the travel provision, are ignored.) Data come from California Air Resources Board disclosures; a typical electric vehicle earned two credits and a long-range Tesla earned four.*

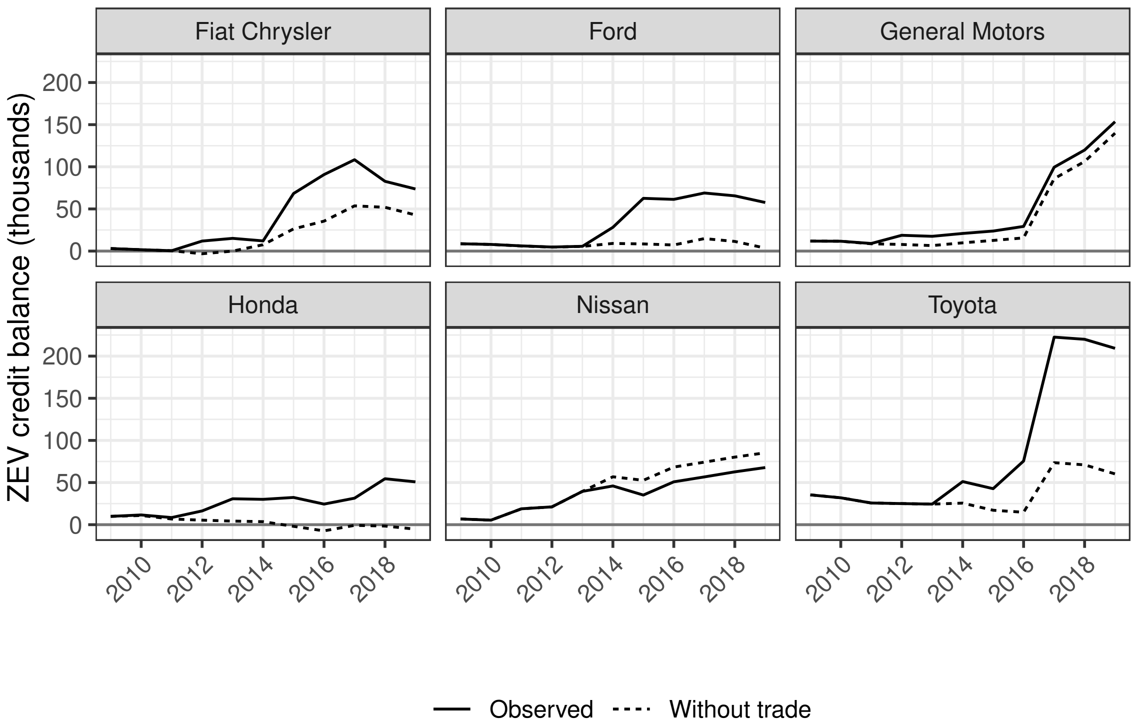

Figure 3.3: *ZEV credit balances with and without trading for large manufacturers*

Note: *The solid line shows California ZEV credit balances at the end of each model year, 2009–2019, for the six large automakers. The dashed line eliminates trades between model years 2010 and 2019 (that is, removes credits that were purchased and restores credits that were sold) to show the contribution of trades to the observed credit balances. For context, a credit sold for $1500–3500; a typical electric vehicle earned two credits while a long-range vehicle (such as a Tesla) earned four. Data come from California Air Resources Board disclosures.*

and builds up a modest bank of credits ahead of 2018 by introducing and selling EVs.

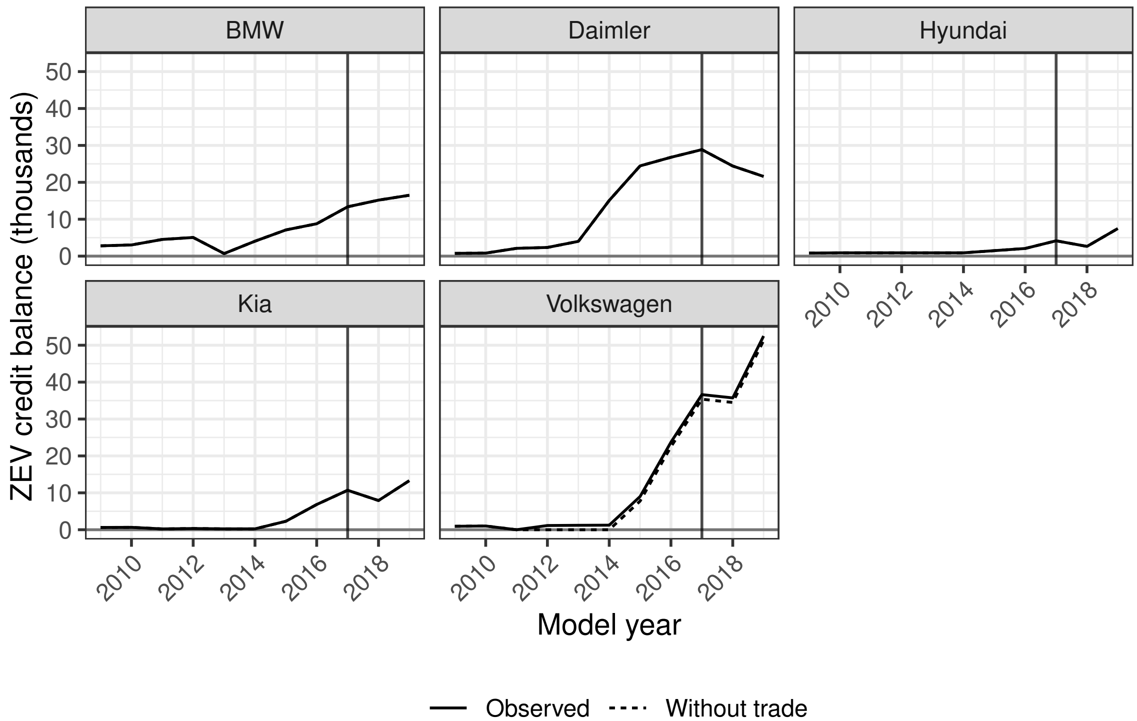

Figure 3.4: *ZEV credit balances with and without trading for manufacturers whose requirements change*

Note: *The solid line shows California ZEV credit balances at the end of each model year, 2009–2019, for the five large automakers that were subject to the ZEV mandate starting in 2018. The dashed line eliminates trades between model years 2010 and 2019 to show the contribution of trades to the observed credit balances. (The start of the 2018 model year is shown using a vertical line.) Data come from California Air Resources Board disclosures.*

While manufacturers' credit balances rose sharply, credit prices fell by 40–60% between 2012 and 2017, as shown in Figure 3.5. (Due to changes in Tesla's accounting reports, prices after 2017 are not observable.) Automakers with low EV sales continued to purchase credits. Tesla was the dominant seller of credits, but not the only one (as shown in Figure 3.6).

The largest credit buyers, shown in Figure 3.7, were the large automakers with low or zero EV sales (Fiat Chrysler, Ford, GM, Honda, and Toyota). About 10% of overall ZEV credit purchases were by intermediate volume automakers (mostly Subaru), which may have used the credits toward the PZEV mandate.

Though Tesla's EV sales grow rapidly, giving it more credits to sell, its actual sales of

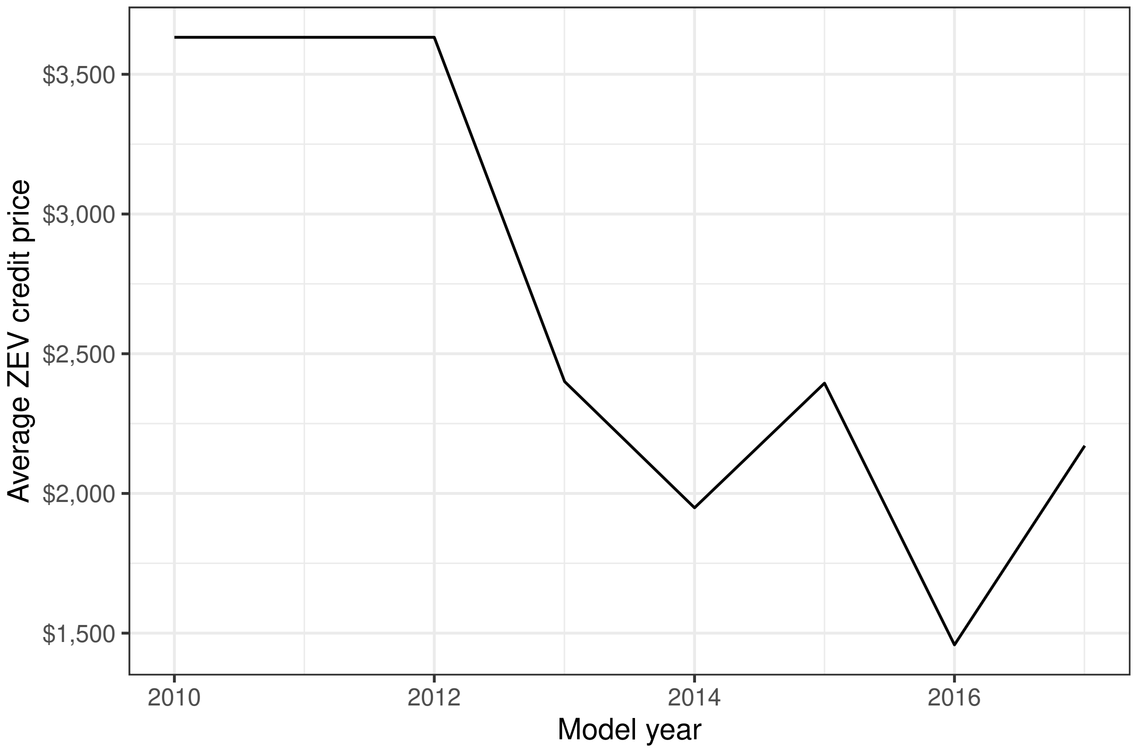

Figure 3.5: *Average prices for Tesla's ZEV credits over time*

Note: *This figure shows the average prices for Tesla ZEV credits, computed by the method in Chapter 1 using state regulatory disclosures and Tesla financial statements. This computation divides the quantity of credits sold by Tesla (across all ZEV states, not just California) to the revenues Tesla reported from those sales. Because of data limitations, I compute a single average price for the 2010–2012 period, and year-by-year average prices for 2013–2017. For context, a typical electric vehicle earned two credits and a long-range vehicle (such as a Tesla) earned four.*

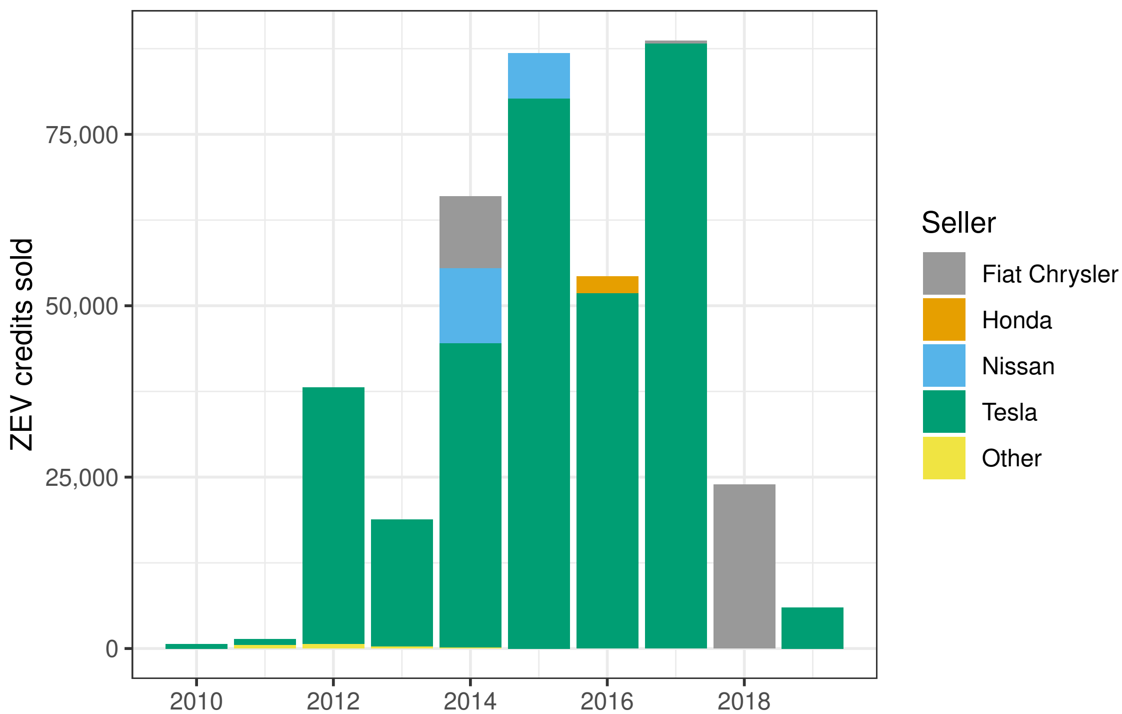

Figure 3.6: *ZEV credit sales over time*

Note: *This figure shows ZEV credit sales by manufacturer and model year (California credits only). Data come from California Air Resources Board disclosures; a typical electric vehicle earned two credits and a long-range Tesla earned four.*

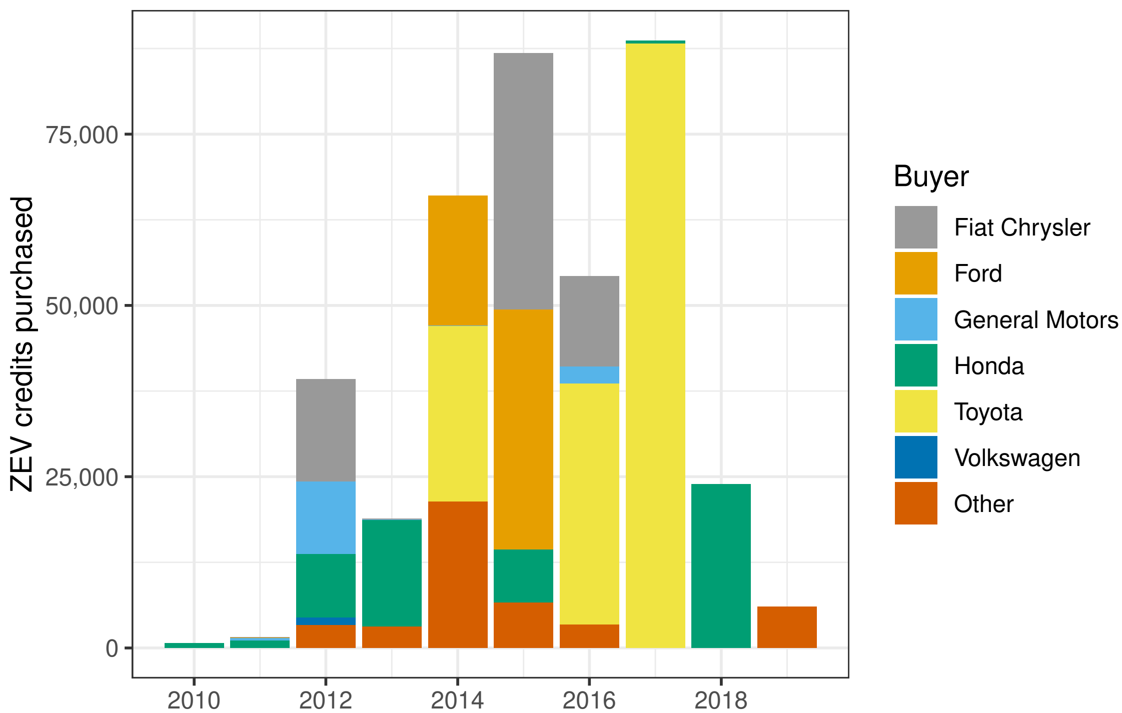

Figure 3.7: *ZEV credit purchases over time*

Note: *This figure shows ZEV credit purchases by manufacturer and model year (California credits only). The group labeled 'Other' combines the three manufacturers that were classified as intermediate-volume throughout the period (and thus not subject to the ZEV mandate): Subaru, Jaguar Land Rover, and Mazda. Data come from California Air Resources Board disclosures; a typical electric vehicle earned two credits and a long-range Tesla earned four.*

credits stagnated before falling to zero. Figure 3.8 compares Tesla's credit sales with the stock of credits available. Until 2015, Tesla sold almost all the credits it had. In 2016 and 2017, Tesla's credit earning began to outstrip its credit sales, and Tesla's credit balance grew rapidly.

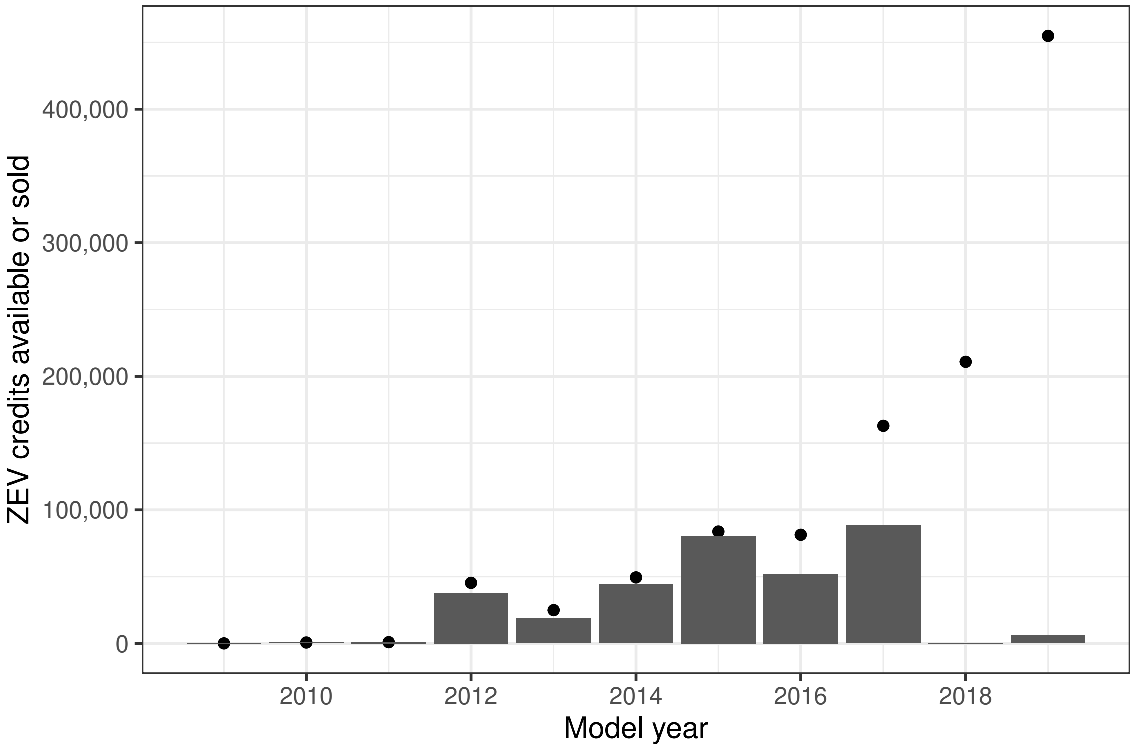

Figure 3.8: *Tesla ZEV credit sales over time*

Note: *This figure shows Tesla's ZEV credit sales and credits available at the end of each model year (California credits only). The dots show the credits available to sell, including earnings in the year and balances carried from previous years. The bars show the credits actually sold. Data come from California Air Resources Board disclosures; a typical electric vehicle earned two credits and a long-range Tesla earned four.*

Figure 3.9 compares prices and quantities for Tesla's credit sales by year; credit sales are generally rising as prices fall, but there is not a direct relationship.

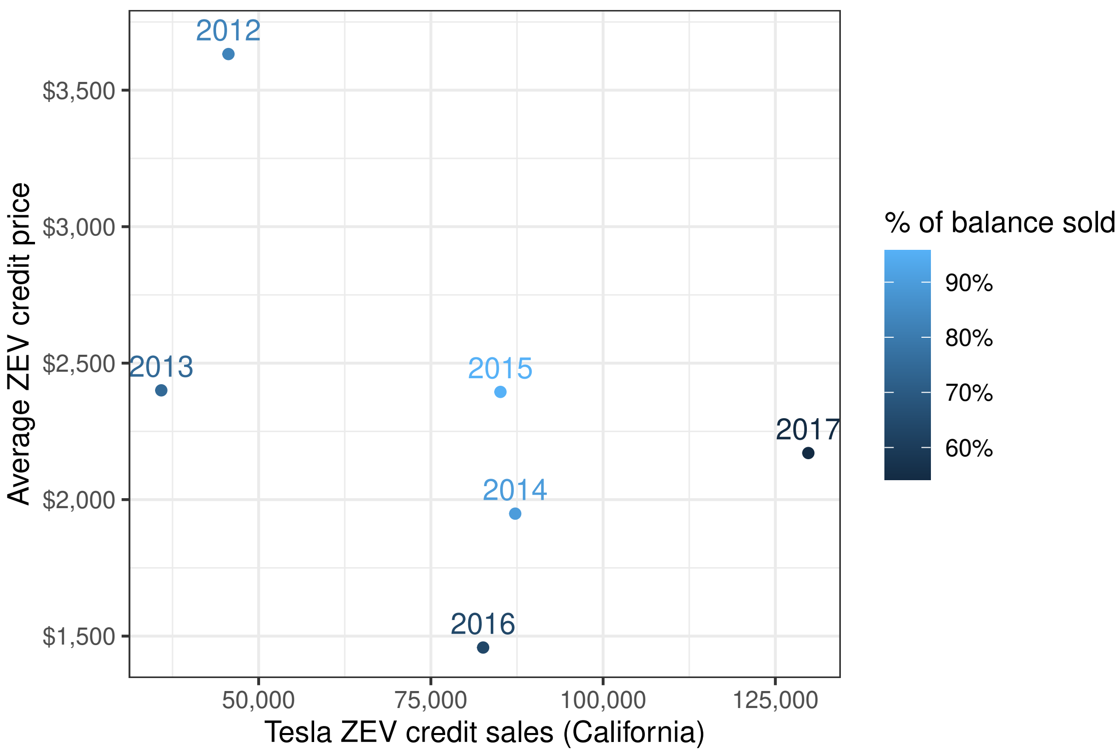

Figure 3.9: *Prices and quantities for Tesla ZEV credit sales by year*

Note: *This figure shows the average prices and quantities for Tesla credit sales by model year. The horizontal axis is Tesla's sales of California ZEV credits, in units of credits (4 credits = 1 Tesla Model S). The vertical axis is the average price Tesla obtained for ZEV credits that year, across all states. (The price for 2012 is an average from 2010–12 data.) Points are colored by the proportion of Tesla's credit inventory that it sold that year.*

3.3.3 Trends in GHG credit data

Compliance with the GHG regulation was affected by the growth in EV sales, but also by offsetting growth in sales of high-emissions vehicles and regulation-specific factors like the tightening of requirements and expiration of older-vintage credits. As a result, aggregate credit balances were declining over the same period, as shown in Figure 3.10. Without the GHG credits earned by Tesla's sales, the decline after 2016 would be steeper.

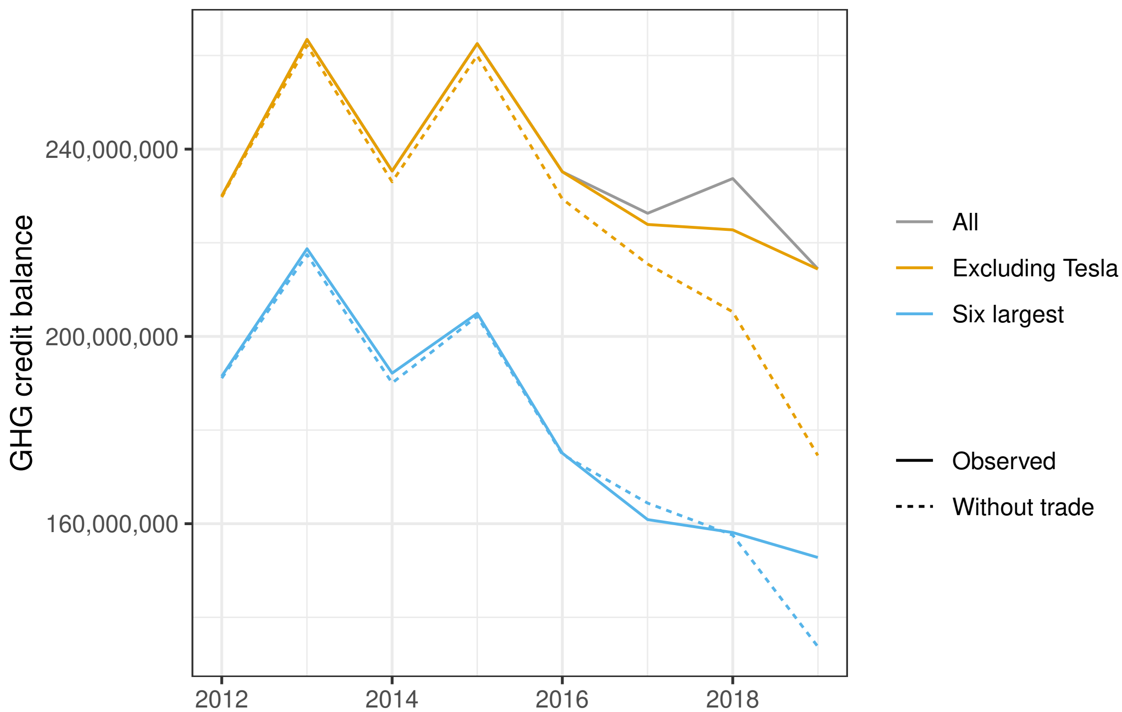

Figure 3.10: *Aggregate GHG credit balances over time*

Note: *This figure shows the GHG credit balances left at the end of each model year, across groups of manufacturers. Hyundai and Kia are excluded from all data. Data from some small manufacturers are missing in some years; missing values are assumed to be zero. The 'six largest' group consists of Chrysler/Fiat Chrysler, Ford, GM, Honda, Nissan, and Toyota. The solid line shows observed balances, and the dashed line removes credit trades among manufacturers. Data come from EPA reports.*

Tesla was a major seller of GHG credits but did not dominate the supply side of the market the way it did for the ZEV program. Honda and Toyota, which had large volumes of hybrid and clean gas car sales, also earned many GHG credits that they then sold. As shown in Figure 3.11, Honda and Toyota were larger sellers in the earlier years, but Tesla had caught up by 2018.

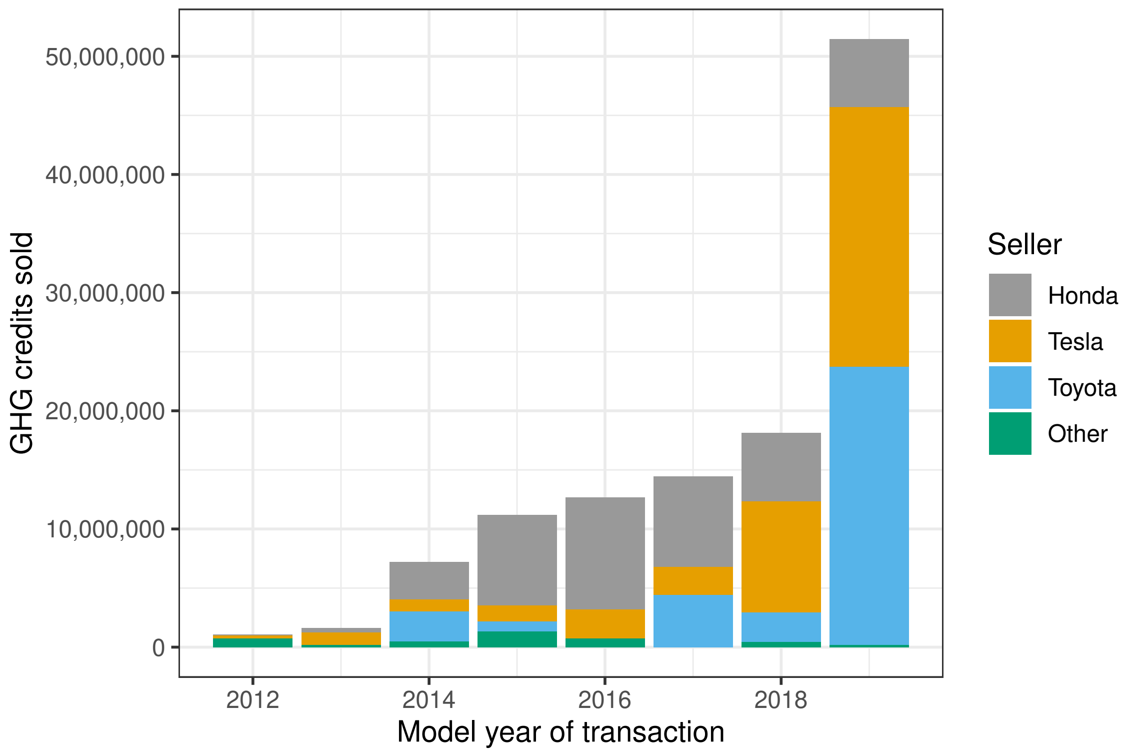

Figure 3.11: *GHG credit sales over time*

Note: *This figure shows GHG credit sales, by the selling manufacturer and the model year in which the transaction occurred. (This may differ from the model year in which the transacted credit was originally earned.) Data come from yearly EPA reports. One credit corresponds to a megagram of carbon dioxide emissions and was valued at $35–40.*

Purchasers of credits are shown in Figure 3.12; the largest was Fiat Chrysler, followed by Daimler, GM, BMW, and Volkswagen.

Under its long-term contracts, Tesla sold all the GHG credits it had available by 2019, as shown in Figure 3.13. Tesla's credit sales fell short of its overall available credits in 2017 and 2018, but Fiat Chrysler purchased all of Tesla's accumulated credits in 2019.

Figure 3.14 shows firm-level balances for the six largest automakers (the same group subject to the ZEV mandate). The largest driver of the declining trend is Toyota, which started 2012 with the largest balance and gradually used it up and sold part of the remainder.

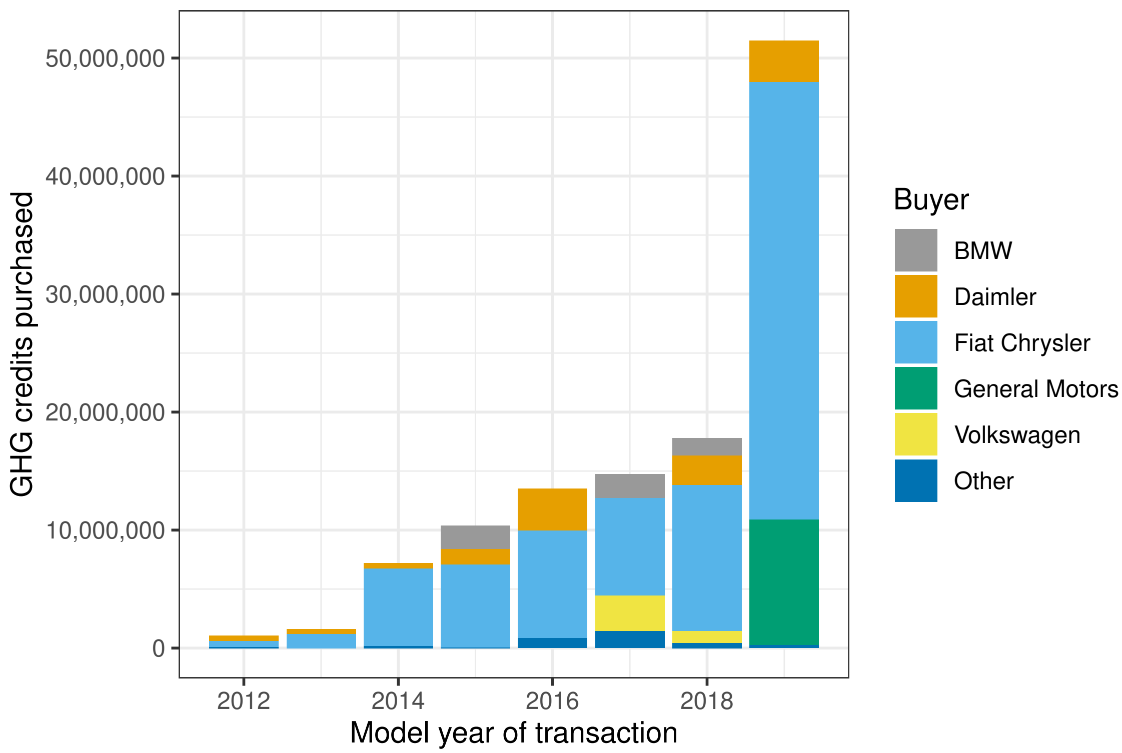

Figure 3.12: *GHG credit purchases over time*

Note: *This figure shows GHG credit purchases, by the purchasing firm and the model year in which the transaction occurred. (This may differ from the model year in which the transacted credit was originally earned.) Data come from yearly EPA reports. One credit corresponds to a megagram of carbon dioxide emissions and was valued at $35–40.*

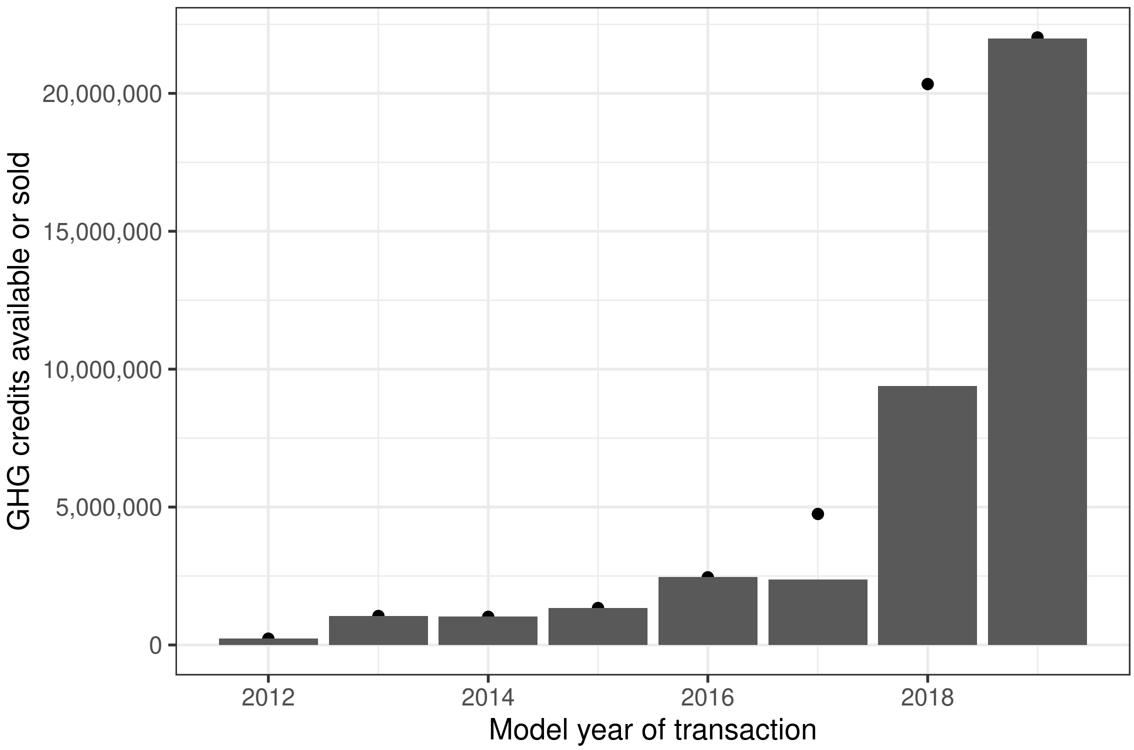

Figure 3.13: *Tesla GHG credit sales over time*

Note: *This figure shows Tesla's GHG credit sales and credits available at the end of each model year. The dots show the credits available to sell, including earnings in the year and balances carried from previous years. The bars show the credits actually sold. All Tesla vehicles are counted as zero emissions; larger-footprint vehicles earn more credits. Data come from EPA reports.*

Most of the other firms' balances were flat throughout the period, as their fleets' emissions were close to their target levels.[18] Honda used the credit market to sell off its excess earnings and maintain a balance of about 40 million credits. Fiat Chrysler relied solely on the credit trading market in order to comply: without its purchases it would have had an escalating deficit (up to 35 million credits in 2019), but instead it continued buying credits (presumably under its contract with Tesla) to build up a large balance of over 45 million credits by 2019.

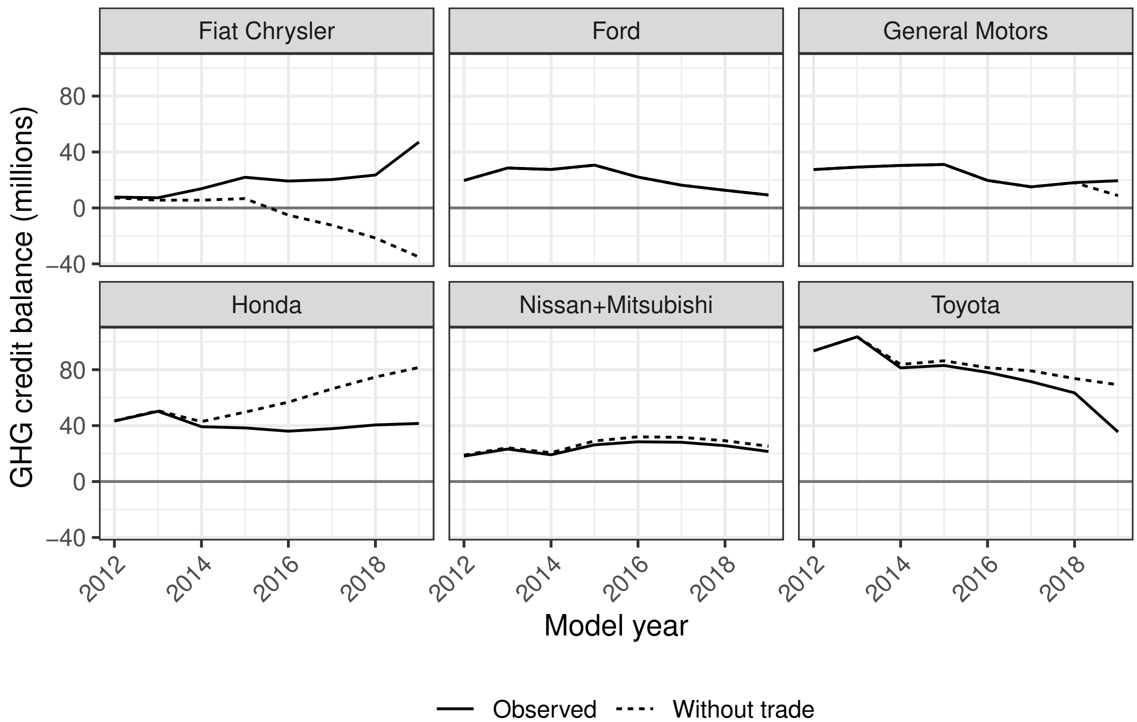

Figure 3.14: *GHG credit balances with and without trading, selected large manufacturers*

Note: *The solid line shows EPA GHG credit balances at the end of each model year, 2012–2019, for a set of large automakers. The dashed line eliminates trades between manufacturers to show the contribution of trades to the observed credit balances. Nissan and Mitsubishi are combined to make data comparable before and after their 2017 alliance. Credits are shown in units of millions of credits; one credit corresponds to a megagram of carbon dioxide emissions and was valued at $35–40. Data come from EPA reports.*

[18]The credit purchase GM made in 2019 brought its balance back to 2016 levels, and was explained as insurance against future regulatory uncertainty. See "Tesla's Secret Source of Cash Unmasked as GM, Fiat Chrysler" (Miles Weiss and David Welch, Bloomberg, 6/3/19).

3.3.4 Stylized facts about credit trading

In this section, I describe patterns in the credit data that shed light on the determinants of firm valuations for credits. Firms buy credits for three potential reasons: to meet contemporaneous compliance requirements, to bank in anticipation of future compliance requirements, or to bank in order to resell later on (if they expect credit prices to rise).[19] By observing when firms buy credits, what variables correlate with purchase decisions, and the effects of credit trades, I find evidence consistent with banking in anticipation of future compliance requirements. Credit trading thus occurs among firms with heterogeneous future compliance needs, keeping prices positive even as most firms' credit balances are enough to comply with current requirements.

Stylized fact 1 *Firms acquired credits even when their balances were far from zero. Conditional on buying credits, larger balances were associated with larger purchases.*

The relationship between existing balances and credit transactions, excluding EV-only manufacturers, is shown in Figure 3.15 for both the ZEV and the GHG programs. Firms sometimes acquired credits even when their existing balances were high; among firm-years with credit acquisitions, larger purchases were associated with higher existing balances.

This pattern suggests firms had heterogeneous demands for credits. If firms were all targeting the same fixed balance, for example, larger balances would be associated with smaller purchases. Similarly, if firms purchased credits solely to meet contemporaneous compliance requirements, they would have no need to purchase credits once their balances were sufficiently high.

One driver of this effect is that negative balances were rare (especially under the ZEV mandate, but also under GHG) and, when present, they were small in magnitude. Firms rarely needed to use credit purchases to meet same-year requirements, and could instead rely on earned or previously purchased credits.

[19] Firms with market power in the credit market may also buy and hold credits to prop up credit prices.

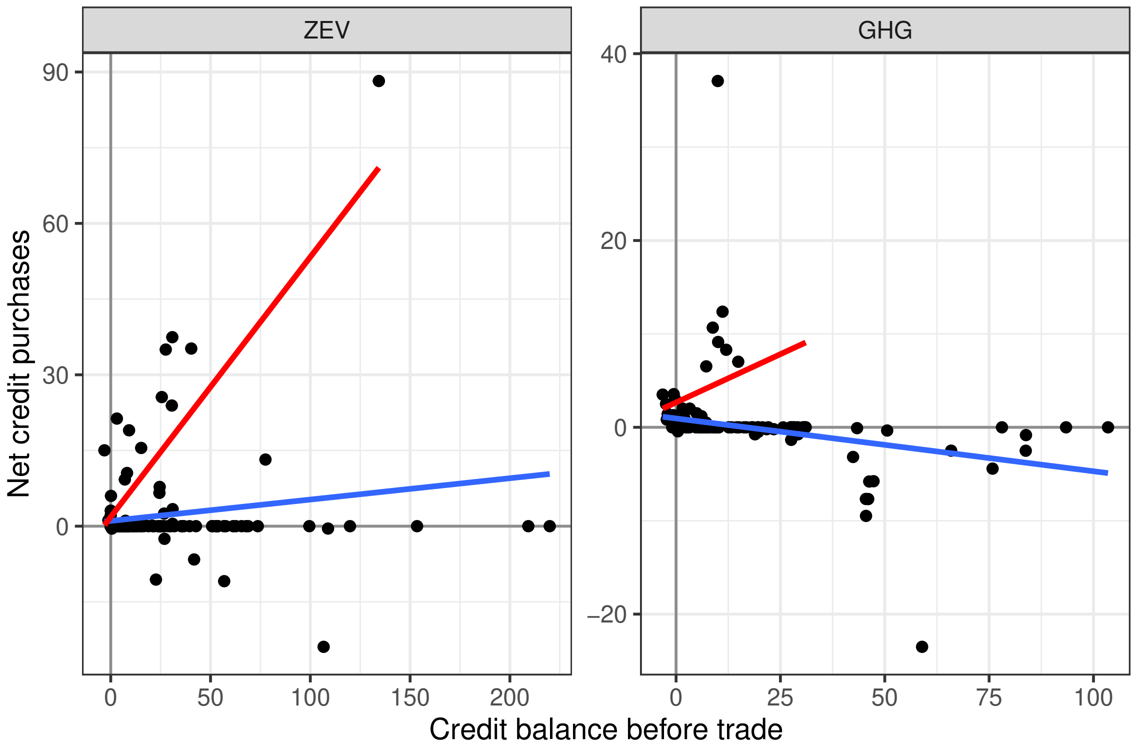

Figure 3.15: *Relationship of credit balances and purchases, ZEV and GHG*

Note: *This figure shows the relationship between end-of-year credit balances, defined as the balance a firm would have had if not for that year's credit trades, and its credit trades; positive values indicate credit purchases, negative values indicate credit sales, and zeros indicate no trades. EV-only manufacturers, such as Tesla, are excluded. The blue line is the overall trend, while the red line is the trend conditional on purchasing credits. The left panel shows ZEV credits (California credits only) in units of thousands; data come from California Air Resources Board disclosures. The right panel shows GHG credits in units of millions; data come from EPA reports.*

When the sample is expanded from credit purchases to all credit transaction decisions (accounting for credit sales and non-trades), the relationship between net trades and balances is flat or negative, as larger balances are also associated with greater sales and a lower likelihood of purchasing credits.

Figure 3.16 replaces the stock of balances, which reflect overall ability to meet requirements, with the flow of contemporaneous credit earning from vehicle sales, which reflects the firm's current vehicle portfolio. For ZEV, the trend is the same for credit flows as for credit stocks: higher EV sales and higher credit purchases often went together. For GHG, the trend reverses: firms purchased credits to offset recent fleet emissions, and conditional on purchasing credits, greater credit earning was associated with smaller credit purchases. In the GHG data, firms typically purchased credits when their earnings were negative (fleet emissions above the standard), suggesting that they targeted a net-zero balance change from the previous year.

Stylized fact 2 *In aggregate, credit trades resulted in less concentration of credit balances across firms.*

The supporting data is in Figure 3.17. For both policies, credit balances became less concentrated over time, but this trend would be mitigated or reversed if none of the observed credit trades had occurred.

This fact suggests that firms primarily use credit trading to smooth out heterogeneity between a few credit suppliers and a larger group of manufacturers in need of credits. It suggests that firms do not have incentives to accumulate large stockpiles of credits by consolidating other firms' balances. For example, if firms had heterogeneous beliefs about the path of future credit prices, bullish firms could buy credits from bearish ones; likewise, if a dominant credit supplier could gain from propping up the credit price, it might attempt to buy credits from smaller suppliers. (If the industry instead featured a large number of credit suppliers and a small number of manufacturers in need of credits, trades to smooth out heterogeneity would also make balances more concentrated.)

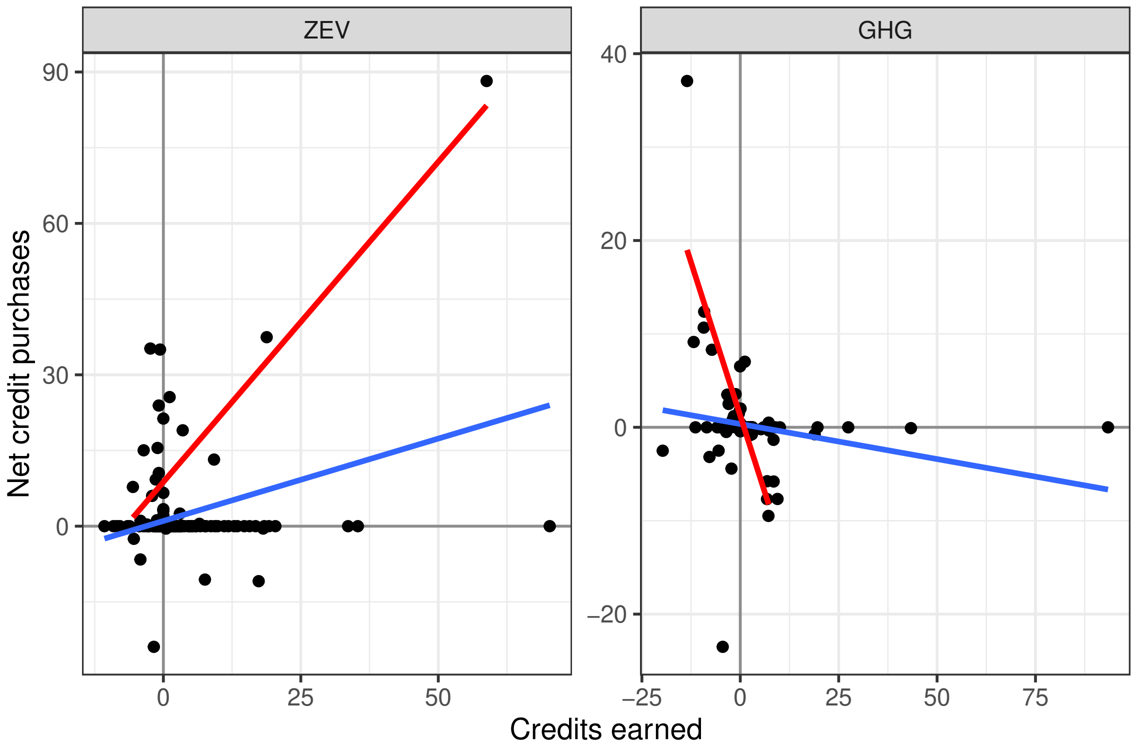

Figure 3.16: *Relationship of credit earning and purchases, ZEV and GHG*

Note: *This figure shows the relationship between credit earning (usually from vehicle sales) and credit trades; positive values indicate credit purchases, negative values indicate credit sales, and zeros indicate no trades. EV-only manufacturers, such as Tesla, are excluded. The blue line is the overall trend, while the red line is the trend conditional on purchasing credits. The left panel shows ZEV credits (California credits only) in units of thousands; data come from California Air Resources Board disclosures. The right panel shows GHG credits in units of millions; data come from EPA reports.*

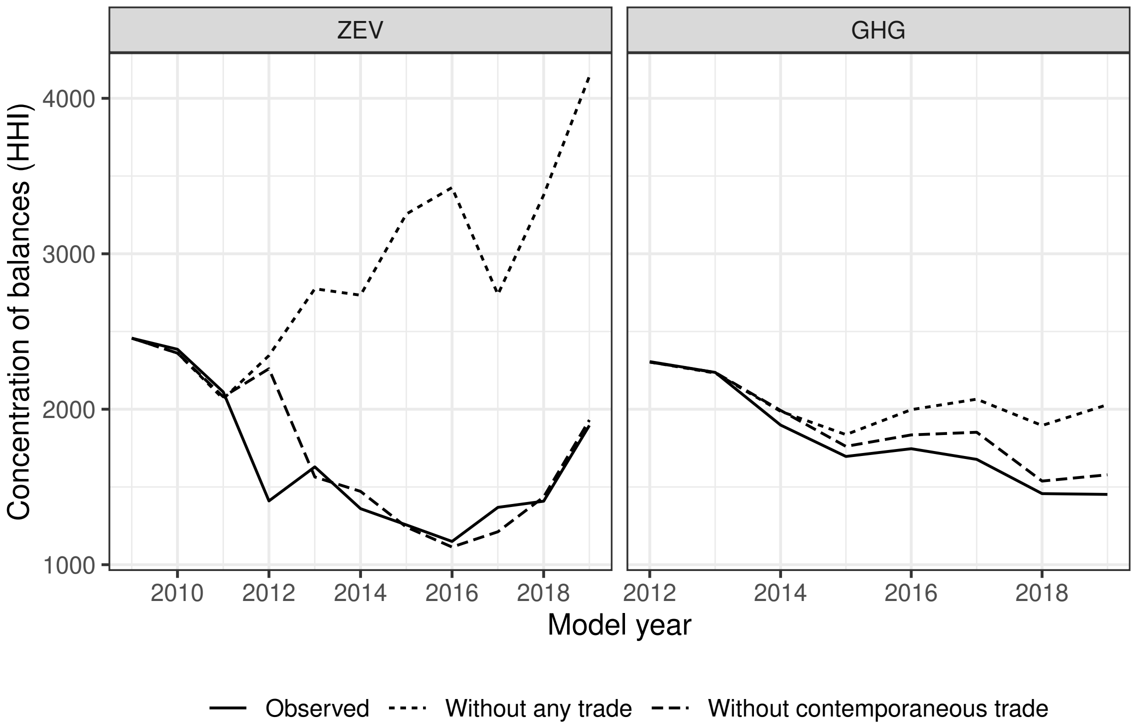

Figure 3.17: *Concentration of credit balances over time, ZEV and GHG*

Note: *This figure shows the concentration of ZEV and GHG credit balances across firms in each year, as measured by the Herfindahl–Hirschman index with negative values censored at zero. The solid (Observed) line shows the observed balances; the light dashed line (Without any trade) shows balances with all cumulative trades removed; and the heavy dashed line (Without contemporaneous trade) shows balances each period with that period's trades removed. The left panel shows ZEV credits (California credits only); data come from California Air Resources Board disclosures. The right panel shows GHG credits; data come from EPA reports.*

It is less clear that credit trades reduced credit balance concentration in every period. For ZEV, the credit trades that occurred in 2012 reduced the concentration of balances, but the trades that occurred later had a minimal effect. For GHG, the credit trades that occurred in each year reduced balance concentration.

3.4 A model of the credit market

In this section I describe a dynamic model of firms interacting in the product and credit markets. Firms compete in quantities in the product market, and earn and give up regulatory credits based on their product sales. In any period, the regulation is assigned by an exogenous process to one of two regimes: a credit trading regime, featuring a perfectly competitive market that clears at a single price, and a no-trading regime in which the credit market is unavailable.

Though the no-trading regime is not observed in our data, only the risk of a no-trading regime gives firms incentives to accumulate credit banks. The no-trading regime can be considered an approximation to the interaction between market thinness and transaction costs (which I do not model), which in some years may make credit transactions technically feasible but prohibitively expensive to carry out.

Firms are risk neutral and face two types of uncertainty about future conditions. The first is idiosyncratic noise in the profit function, reflecting uncertainty over marginal costs or demand, which creates uncertainty over the future credit price. The second is the possibility of a shift between the trading and no-trading regimes. Firms anticipate the cost of adjusting production to meet compliance requirements in a future no-trading regime, and bank credits to reduce that adjustment cost.

The policy is a simplified version of the ZEV and GHG regulations. Most importantly, there is only one regulation. Firms earn and give up credits based on current-year sales of their products, there is no option to comply by paying a fixed fine, credits do not expire, and (unlike GHG) credits may not be borrowed from the future.

The main divergence between my model and industry accounts of the actual ZEV and

GHG credit markets is that I do not allow firms to have market power in the credit market, which is a known source of inefficiency in credit trading markets and an explanation industry observers have offered for Tesla's growing credit banks after 2015.[20] Because the credit market is thin and has few players (which also interact in the product market), market power cannot be ruled out. As I show below, however, the possibility of a no-trading regime also explains the stylized facts I observe.

3.4.1 Model setup

There is a set of multiproduct firms $\mathcal{J}$. Firm j begins each period t with some balance of credits; it then competes in the product market with a product set $\mathcal{K}_j$, chooses the number of credits to buy or sell, earns or loses regulatory credits as a result of its product market sales, and carries forward its remaining credit balance to the next period $t+1$. The firm's balance carried forward cannot be negative (that is, I do not allow for credit borrowing).

Firms compete in quantities in the product market. Firm j chooses q_{jt}, the vector of quantities of its products. In the trading regime, firm j also chooses x_{jt}, the net amount of credits it will purchase at the fixed price of p_{xt}. (Negative values of x_{jt} indicate credit sales.) In the no-trading regime, $x_{jt} = 0$. Outcomes in the product market are summarized by a differentiable profit function $\pi_{jt}(q_{jt}, q_{-j,t}; \omega_{jt})$, where $q_{-j,t}$ is the vector of quantities of products by other firms and ω_{jt} is a random variable drawn independently across firms. Firm j's flow payoff, taking other firms' quantities as given, is

$$\pi_{jt}(q_{jt}, q_{-j,t}; \omega_{jt}) - p_{xt}x_{jt}.$$

Firm j enters the period with the balance b_{jt}, which the regulator updates by a weighted sum of quantities $q'_{jt}\phi_{jt}$. The regulator seeks to encourage products j whose weights ϕ_{jt} are positive, and discourage products whose ϕ_{jt} are negative. The balance at the end of the

[20] See, e.g., "Hyundai and Kia EVs may just avoid 'compliance car' stigma, but is that enough?" (Bengt Halvorson, Green Car Reports, 2/6/19).

period is thus

$$b_{j,t+1} \equiv b_{jt} + q'_{jt}\phi_{jt} + x_{jt}.$$

Putting these parts together, the firm maximizes its expected discounted stream of profits subject to the constraint that balances are non-negative. Let δ be the discount rate and let θ be the probability of a no-trading regime. The Bellman equation corresponding to the firm's problem is V_{jt}^{tr} for the trading regime and V_{jt}^{ntr} for the no-trading regime, where

$$\begin{aligned}
V_{jt}^{tr}(b_{jt}, q_{-j,t}) &= \max_{q_{jt}, x_{jt}} \pi_{jt}(q_{jt}, q_{-j,t}; \omega_{jt}) - p_{xt}x_{jt} + \delta E[V_{j,t+1}(b_{jt} + q'_{jt}\phi_{jt} + x_{jt}, q_{-j,t+1})] \\
&\quad \text{subject to } b_{jt} + q'_{jt}\phi_{jt} + x_{jt} \geq 0, \\
V_{jt}^{ntr}(b_{jt}, q_{-j,t}) &= \max_{q_{jt}} \pi_{jt}(q_{jt}, q_{-j,t}; \omega_{jt}) + \delta E[V_{j,t+1}(b_{jt} + q'_{jt}\phi_{jt}, q_{-j,t+1})] \\
&\quad \text{subject to } b_{jt} + q'_{jt}\phi_{jt} \geq 0, \\
E[V_{jt}(b_{jt}, q_{-j,t})] &= (1-\theta)E[V_{jt}^{tr}(b_{jt}, q_{-j,t})] + \theta E[V_{jt}^{ntr}(b_{jt}, q_{-j,t})].
\end{aligned}$$

The dynamic channel I wish to focus on is the effect of a firm's actions on its own credit bank. Therefore, I assume that firms have no incentive to alter their behavior today in order to affect other firms' future product market choices.

Assumption 3 *Firms do not believe that their choices of (q_{jt}, x_{jt}) will affect other firms' decisions next period, $q_{-j,t+1}$.*

Definition 1 *Suppose that the regulation is in the trading regime in period t. A trading regime equilibrium is a vector of quantities and credit-market positions $(q_{jt}, x_{jt})_{j \in \mathcal{J}}$ and a credit price p_{xt} such that:*

- *For each firm $j \in \mathcal{J}$, (q_{jt}, x_{jt}) solves the firm's dynamic optimization problem subject to the non-negative balance constraint, taking $q_{-j,t}$ and p_{xt} as given.*
- *The credit market clears, $\sum_{j \in \mathcal{J}} x_{jt} = 0$.*

These two conditions generate an aggregate non-negative balance constraint,

$$\sum_j b_{jt} + \sum_j (q'_{jt}\phi_{jt}) \geq 0. \tag{3.1}$$

Because of the flexibility of the credit market, any vector of quantities that satisfies this constraint is feasible.

Definition 2 *Suppose that the regulation is in the no-trading regime in period t. A no-trading regime equilibrium is a vector of quantities $(q_{jt})_{j\in\mathcal{J}}$ such that, for each firm $j \in \mathcal{J}$, q_{jt} solves its dynamic optimization problem subject to the non-negative balance constraint, taking $q_{-j,t}$ as given.*

With the credit market removed, the set of feasible quantity vectors is severely restricted, as the non-negative balance constraint must hold firm-by-firm.

3.4.2 Firm policy functions

How do firms' dynamic incentives affect their decisions? I show that in the trading regime, production decisions are only affected by the current credit price, but credit trading decisions take firms' future compliance into account.

In the trading regime, opportunities for credit market arbitrage place strong restrictions on how firms obtain and use their credits. Letting $\lambda_{jt} \geq 0$ be the multiplier on firm j's constraint, the first order condition on x_{jt} is

$$-p_{xt} + \delta E\left[\frac{\partial V_{j,t+1}}{\partial b_{j,t+1}}(b_{jt} + q'_{jt}\phi_{jt} + x_{jt}, q_{-j,t+1})\right] + \lambda_{jt} = 0.$$

At the optimum, buying an additional credit will cost p_{xt}, relax the constraint (if it binds), and marginally increase the incoming balance next period.

Proposition 3 *In a trading regime equilibrium, firm production $q^*_{jt}(p_{xt}, \omega_{jt}, q_{-j,t})$ satisfies*

$$\frac{\partial \pi_{jt}}{\partial q_{jtk}}(q^*_{jt}(p_{xt}, \omega_{jt}, q_{-j,t}), q_{-j,t}; \omega_{jt}) + p_{xt}\phi_{jtk} = 0 \text{ for all } k \in \mathcal{K}_j.$$

That is, when making choices in the product market, firms value the credits earned (or lost) at their market price. (The analogous result for competition in prices is used to construct the profit functions in Chapter 1.) At the optimum, each firm is indifferent between buying a credit and generating it through the product market.

I now examine the firm's choice of x_{jt}.

Proposition 4 *The marginal value of increasing firm j's incoming balance, b_{jt}, in the trading regime is*

$$\frac{\partial V_{jt}^{tr}}{\partial b_{jt}} = \delta E\left[\frac{\partial V_{j,t+1}}{\partial b_{j,t+1}}(b_{jt} + q'_{jt}\phi_{jt} + x_{jt}, q_{-j,t+1})\right] + \lambda_{jt} = p_{xt}.$$

This result is a direct application of the envelope theorem. In the trading regime, the value of an additional credit to the firm is its market price. This means that the expected marginal value of banking a credit is

$$\begin{aligned} E&\left[\frac{\partial V_{j,t+1}}{\partial b_{j,t+1}}(b_{jt} + q'_{jt}\phi_{jt} + x_{jt}, q_{-j,t+1})\right] \\ &= (1-\theta)E[p_{x,t+1}] + \theta E\left[\frac{\partial V_{j,t+1}^{ntr}}{\partial b_{j,t+1}}(b_{jt} + q'_{jt}\phi_{jt} + x_{jt}, q_{-j,t+1})\right]. \end{aligned}$$

Rewriting the first order condition on x_{jt}, and using the complementary slackness condition, gives a condition for equilibrium in the credit market.

Proposition 5 *In a trading equilibrium, at least one of the following must be true for each $j \in \mathcal{J}$:*

$$b_{jt} + q'_{jt}\phi_{jt} + x_{jt} = 0, \quad \textit{or}$$

$$p_{xt} - \delta(1-\theta)E[p_{x,t+1}] - \delta\theta E\left[\frac{\partial V_{j,t+1}^{ntr}}{\partial b_{j,t+1}}(b_{jt} + q'_{jt}\phi_{jt} + x_{jt}, q_{-j,t+1})\right] = 0.$$

Proposition 5 gives a condition on the credit bank, $b_{j,t+1}$, that a firm will carry forward following a trading equilibrium. It also gives a condition on equilibrium p_{xt} when the aggregate balance constraint (3.1) does not bind. (And when the aggregate balance constraint does bind, the policy functions for quantities place restrictions on equilibrium p_{xt}.)

Therefore, in the trading regime, any two firms with a positive credit bank (if they have the same information) must have the same marginal value of a credit if next period is a no-trading regime, $E[\partial V_{j,t+1}^{ntr}/\partial b_{j,t+1}]$. If the firms have different future profit functions $\pi_{j,t+1}$ that produce different marginal values for a given credit bank, they will use credit trading to adjust their credit banks and compensate for those differences. Furthermore, that marginal value of a credit under a no-trading regime is closely related to firm beliefs about credit price evolution.

If firms do not anticipate the no-trading regime, $\theta = 0$, then the model delivers no predictions about x_{jt} except when the aggregate balance constraint (3.1) binds. If all firms have the same information about future prices, they are indifferent between selling and banking any number of excess credits.

3.4.3 Relationship to stylized facts

How are balances and credit purchases related in the model? Proposition 5 states that firms target a given (perhaps unique) optimal credit bank $b^*_{j,t+1} = b_{jt} + q'_{jt}\phi_{jt} + x_{jt}$. All else equal, increasing a firm's balance b_{jt} reduces its equilibrium purchases x_{jt} one for one. But the balance is endogenous because it was chosen by the firm the previous period. Suppose the optimal credit bank is positively correlated within a firm across subsequent years due to, for example, persistence in ω_{jt} or cross-firm heterogeneity in production costs. Then, in the cross section, b_{jt} and x_{jt} may also be positively correlated:

$$\mathrm{Cov}(b_{jt}, x_{jt}) = \mathrm{Cov}(b_{jt}, b^*_{j,t+1}) - \mathrm{Cov}(b_{jt}, q'_{jt}\phi_{jt}).$$

How is the concentration of balances determined? Ignoring negative balances, the difference between concentration with and without trade is

$$HHI(b_{jt} + q'_{jt}\phi_{jt} + x_{jt}) - HHI(b_{jt} + q'_{jt}\phi_{jt}) = \frac{\sum_j 2(b_{jt} + q'_{jt}\phi_{jt})x_{jt} + \sum_j x_{jt}^2}{\sum_j (b_{jt} + q'_{jt}\phi_{jt})^2}.$$

If $b_{jt} + q'_{jt}\phi_{jt}$ and x_{jt} are sufficiently negatively correlated, this difference can be negative, as observed in the data.

3.5 Policy implications

Using the model, I consider two policy changes that target the trading and banking mechanisms. First, I alter the no-trading regime by allowing firms to purchase credits from the regulator for a fixed price. Second, I alter credit banking by introducing exogenous credit appreciation (or depreciation).

Allowing the regulator to sell credits to firms is a type of safety valve suggested for the ZEV mandate by McConnell and Leard (2021) and the GHG regulation by Leard and

McConnell (2017). The CAFE standards effectively have such a policy, by allowing firms to comply by paying a fine, and California regulatory documents are unclear about whether the existing ZEV mandate also considers a firm to be in compliance once it has paid the fine.

To implement this change, I assume that the regulator's price or fine F is higher than any possible equilibrium price in the trading regime. Therefore, the change only affects the no-trading regime. Firm j's dynamic optimization problem in the no-trading regime thus becomes

$$V_{jt}^{ntr}(b_{jt}, q_{-j,t}) = \max_{q_{jt}} \pi_{jt}(q_{jt}, q_{-j,t}; \omega_{jt}) + \delta E[V_{j,t+1}(\max(0, b_{jt} + q_{jt}'\phi_{jt}), q_{-j,t+1})] + \\ + F \min(0, b_{jt} + q_{jt}'\phi_{jt})$$

and the marginal value of increasing firm j's incoming balance is

$$\frac{\partial V_{jt}^{ntr}}{\partial b_{jt}} = \begin{cases} F, & b_{jt} + q_{jt}'\phi_{jt} < 0 \\ \delta E\left[\frac{\partial V_{j,t+1}}{\partial b_{j,t+1}}(b_{jt} + q_{jt}'\phi_{jt}, q_{-j,t+1})\right], & b_{jt} + q_{jt}'\phi_{jt} > 0 \end{cases}.$$

If any firm j chooses $b_{j,t+1} > 0$ and knows with certainty ahead of time that it will pay the fine next period ($b_{j,t+1} + q_{j,t+1}'\phi_{j,t+1} < 0$), then the equilibrium price will satisfy

$$p_{xt} = \delta(1-\theta)E[p_{x,t+1}] + \delta\theta F.$$

If F is less than the marginal value of a credit in the no-trading regime under current policy, the change will generally lead to lower credit prices.

Next, I look at the effect of altering credit banking through exogenous credit depreciation or appreciation. As described in Section 3.2.1, the ZEV mandate featured credit appreciation in the very beginning. Neither policy used credit depreciation, but both had a form of credit expiration: ZEV in the early 2010s and GHG throughout.

To implement this change, let the appreciation or depreciation factor be $\rho > 0$ and define

$$\tilde{V}_{jt}^{tr}(b_{jt}, q_{-j,t}) = V_{jt}^{tr}(\rho b_{jt}, q_{-j,t}), \quad \tilde{V}_{jt}^{ntr}(b_{jt}, q_{-j,t}) = V_{jt}^{ntr}(\rho b_{jt}, q_{-j,t}),$$

Firm j's new dynamic optimization problem has the value function $\tilde{V}$, rather than V. Note that

$$\frac{\partial \tilde{V}_{jt}^{tr}}{\partial b_{jt}}(b_{jt}, q_{-j,t}) = \rho \frac{\partial V_{jt}^{tr}}{\partial b_{jt}}(\rho b_{jt}, q_{-j,t}) = \rho\, p_{xt}, \quad \frac{\partial \tilde{V}_{jt}^{ntr}}{\partial b_{jt}}(b_{jt}, q_{-j,t}) = \rho \frac{\partial V_{jt}^{ntr}}{\partial b_{jt}}(\rho b_{jt}, q_{-j,t}),$$

so the second part of Proposition 5 (the condition that must hold if firm j carries a positive balance) becomes (holding b_{jt} fixed, but indicating new endogenous quantities as $\tilde{p}_{xt}, \tilde{q}_{jt}, \tilde{x}_{jt}$):

$$\tilde{p}_{xt} - \rho\delta(1-\theta)E[\tilde{p}_{x,t+1}] - \rho\delta\theta E\left[\frac{\partial V_{j,t+1}^{ntr}}{\partial b_{j,t+1}}(\rho(b_{jt} + \tilde{q}_{jt}'\phi_{jt} + \tilde{x}_{jt}), \tilde{q}_{-j,t+1})\right] = 0.$$

As a consequence, the change in prices can be written

$$\begin{aligned}\tilde{p}_{xt} - p_{xt} = (\rho - 1)p_{xt} + \rho\delta(1-\theta)E[\tilde{p}_{x,t+1} - p_{x,t+1}] + \\ + \rho\delta\theta E\left[\frac{\partial V_{j,t+1}^{ntr}}{\partial b_{j,t+1}}(\rho(b_{jt} + \tilde{q}_{jt}'\phi_{jt} + \tilde{x}_{jt}), \tilde{q}_{-j,t+1})\right] + \\ - \rho\delta\theta E\left[\frac{\partial V_{j,t+1}^{ntr}}{\partial b_{j,t+1}}(\rho(b_{jt} + q_{jt}'\phi_{jt} + x_{jt}), q_{-j,t+1})\right].\end{aligned}$$

The terms depending on beliefs are required to make exact predictions.

3.6 Conclusion

A major limitation of the models of the vehicle market in Chapters 1 and 2 is that they do not model the formation of market-clearing credit prices for the two major supply-side regulations. As a result, the models there cannot accommodate counterfactuals that involve changes to policy design, while keeping the credit market intact. In addition, they ignore a potentially important channel by which large changes to the vehicle market, such as product entry and exit by other firms, can affect firm decisions. In this chapter, I fill that gap by developing a model of the credit market. The main determinant of firm demand for credits is future compliance needs, consistent with stylized facts about the credit markets. When coupled with a model of firm profits and firm belief formation, the model could be taken to data and used to quantify the effects of policy changes.

References

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr (2016). "Transition to Clean Technology". In: *Journal of Political Economy* 124.1, pp. 52–104.

Ackerberg, Daniel A. and Marc Rysman (2005). "Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects". In: *The RAND Journal of Economics* 36.4, pp. 771–788.

Adams, Brian and Kevin R. Williams (2019). "Zone Pricing in Retail Oligopoly". In: *American Economic Journal: Microeconomics* 11.1, pp. 124–156.

Aghion, Philippe, Antoine Dechezleprêtre, David Hémous, Ralf Martin, and John Van Reenen (2016). "Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry". In: *Journal of Political Economy* 124.1, pp. 1–51.

Anderson, Soren T. and James M. Sallee (2011). "Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards". In: *American Economic Review* 101.4, pp. 1375–1409.

— (2016). "Designing Policies to Make Cars Greener". In: *Annual Review of Resource Economics* 8.1, pp. 157–180.

Bedsworth, Louise Wells and Margaret R. Taylor (2007). "Learning from California's Zero-Emission Vehicle Program". In: *California Economic Policy* 3.4.

Beresteanu, Arie and Shanjun Li (2011). "Gasoline Prices, Government Support, and the Demand for Hybrid Vehicles in the United States". In: *International Economic Review* 52.1, pp. 161–182.

Berry, Steven, Alon Eizenberg, and Joel Waldfogel (2016). "Optimal Product Variety in Radio Markets". In: *The RAND Journal of Economics* 47.3, pp. 463–497.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium". In: *Econometrica* 63.4, pp. 841–890.

— (1999). "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy". In: *The American Economic Review* 89.3, pp. 400–430.

Blonigen, Bruce A., Christopher R. Knittel, and Anson Soderbery (2017). "Keeping It Fresh: Strategic Product Redesigns and Welfare". In: *International Journal of Industrial Organization* 53, pp. 170–214.

Borenstein, Severin, James Bushnell, Frank A. Wolak, and Matthew Zaragoza-Watkins (2019). "Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design". In: *American Economic Review* 109.11, pp. 3953–3977.

Brand, James (2020). *Differences in Differentiation: Rising Variety and Markups in Retail Food Stores*. Job Market Paper, p. 75.

Bratvold, Delma and Matthew Cleaver (2017). *Analysis of the Effect of Zero-Emission Vehicle Policies: State-Level Incentives and the California Zero-Emission Vehicle Regulations*. Tech. rep.

Bresnahan, Timothy F. and Dennis A. Yao (1985). “The Nonpecuniary Costs of Automobile Emissions Standards”. In: *The RAND Journal of Economics* 16.4, pp. 437–455.

Burlig, Fiona, James B. Bushnell, David S. Rapson, and Catherine Wolfram (2021). *Low Energy: Estimating Electric Vehicle Electricity Use*. Tech. rep. w28451. National Bureau of Economic Research.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer (2006). “$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions”. In: *American Economic Review* 96.4, pp. 1253–1270.

California Air Resources Board (2017). *Appendix B: Consumer Acceptance of Zero Emission Vehicles and Plug-in Hybrid Electric Vehicles*. Tech. rep., p. 151.

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer (2000). “Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?” In: *Journal of Political Economy* 108.6, pp. 1292–1326.

Cole, Cassandra, Michael Droste, Christopher R. Knittel, Shanjun Li, and James H. Stock (2021). *Policies for Electrifying the Light-Duty Vehicle Fleet in the United States*. Working Paper 2021-014.

Collantes, Gustavo and Daniel Sperling (2008). “The Origin of California’s Zero Emission Vehicle Mandate”. In: *Transportation Research Part A: Policy and Practice* 42.10, pp. 1302–1313.

Conlon, Christopher and Jeff Gortmaker (2020). “Best Practices for Differentiated Products Demand Estimation with PyBLP”. In: *The RAND Journal of Economics* 51.4, pp. 1108–1161.

Cronshaw, Mark B. and Jamie Brown Kruse (1996). “Regulated Firms in Pollution Permit Markets With Banking”. In: *Journal of Regulatory Economics* 9.2, pp. 179–189.

Dasgupta, Partha and Joseph Stiglitz (1988). “Learning-by-Doing, Market Structure and Industrial and Trade Policies”. In: *Oxford Economic Papers* 40.2, pp. 246–268.

Davis, Lucas W. (2019). “How Much Are Electric Vehicles Driven?” In: *Applied Economics Letters* 26.18, pp. 1497–1502.

Davis, Lucas W. and Christopher R. Knittel (2018). “Are Fuel Economy Standards Regressive?” In: *Journal of the Association of Environmental and Resource Economists* 6.S1, S37–S63.

DellaVigna, Stefano and Matthew Gentzkow (2019). “Uniform Pricing in U.S. Retail Chains”. In: *The Quarterly Journal of Economics* 134.4, pp. 2011–2084.

Dixon, Lloyd, Isaac R. Porche III, and Jonathan Kulick (2002). *Driving Emissions to Zero: Are the Benefits of California’s Zero Emission Vehicle Program Worth the Costs?* Santa Monica, CA: RAND Corporation.

Durrmeyer, Isis (forthcoming). “Winners and Losers: The Distributional Effects of the French Feebate on the Automobile Market”. In: *The Economic Journal*.

Durrmeyer, Isis and Mario Samano (2018). “To Rebate or Not to Rebate: Fuel Economy Standards Versus Feebates”. In: *The Economic Journal* 128.616, pp. 3076–3116.

Eizenberg, Alon (2014). “Upstream Innovation and Product Variety in the U.S. Home PC Market”. In: *The Review of Economic Studies* 81.3, pp. 1003–1045.

Ellerman, A. Denny and Juan-Pablo Montero (2007). “The Efficiency and Robustness of Allowance Banking in the U.S. Acid Rain Program”. In: *The Energy Journal* 28.4, pp. 47–71.

Fan, Ying and Chenyu Yang (2020). “Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market”. In: *American Economic Journal: Microeconomics* 12.2, pp. 99–134.

Fowlie, Meredith, Mar Reguant, and Stephen P. Ryan (2016). “Market-Based Emissions Regulation and Industry Dynamics”. In: *Journal of Political Economy* 124.1, pp. 249–302.

Gandhi, Amit and Jean-François Houde (2019). *Measuring Substitution Patterns in Differentiated Products Industries*. Working Paper 26375. National Bureau of Economic Research.

Goldberg, Pinelopi Koujianou (1998). “The Effects of the Corporate Average Fuel Efficiency Standards in the US”. In: *The Journal of Industrial Economics* 46.1, pp. 1–33.

Goulder, Lawrence H., Mark R. Jacobsen, and Arthur A. van Benthem (2012). “Unintended Consequences from Nested State and Federal Regulations: The Case of the Pavley Greenhouse-Gas-per-Mile Limits”. In: *Journal of Environmental Economics and Management* 63.2, pp. 187–207.

Greene, David L., Sangsoo Park, and Changzheng Liu (2014). “Public Policy and the Transition to Electric Drive Vehicles in the U.S.: The Role of the Zero Emission Vehicles Mandates”. In: *Energy Strategy Reviews*. US Energy Independence: Present and Emerging Issues 5, pp. 66–77.

Grieco, Paul L. E., Charles Murry, and Ali Yurukoglu (2021). *The Evolution of Market Power in the US Auto Industry*. Working Paper 29013. National Bureau of Economic Research.

Hahn, Robert W. (1984). “Market Power and Transferable Property Rights”. In: *The Quarterly Journal of Economics* 99.4, pp. 753–765.

Hausman, Jerry (1996). “Valuation of New Goods under Perfect and Imperfect Competition”. In: *The Economics of New Goods*. Ed. by Timothy F. Bresnahan and Robert J. Gordon. University of Chicago Press, pp. 207–248.

Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates (2016). “Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors”. In: *The American Economic Review* 106.12, pp. 3700–3729.

— (2020). “Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation”. In: *American Economic Journal: Economic Policy* 12.4, pp. 244–274.

Holland, Stephen P., Erin T. Mansur, and Andrew J. Yates (2021). “The Electric Vehicle Transition and the Economics of Banning Gasoline Vehicles”. In: *American Economic Journal: Economic Policy* 13.3, pp. 316–344.

Ito, Koichiro and James M. Sallee (2018). “The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel Economy Standards”. In: *The Review of Economics and Statistics* 100.2, pp. 319–336.

Jacobsen, Mark R. (2013). “Evaluating US Fuel Economy Standards in a Model with Producer and Household Heterogeneity”. In: *American Economic Journal: Economic Policy* 5.2, pp. 148–187.

Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins (2005). “A Tale of Two Market Failures: Technology and Environmental Policy”. In: *Ecological Economics*. Technological Change and the Environment 54.2, pp. 164–174.

Jenn, Alan, Katalin Springel, and Anand R. Gopal (2018). “Effectiveness of Electric Vehicle Incentives in the United States”. In: *Energy Policy* 119, pp. 349–356.

Klier, Thomas and Joshua Linn (2012). “New-Vehicle Characteristics and the Cost of the Corporate Average Fuel Economy Standard”. In: *The RAND Journal of Economics* 43.1, pp. 186–213.

Kling, Catherine and Jonathan Rubin (1997). “Bankable Permits for the Control of Environmental Pollution”. In: *Journal of Public Economics* 64.1, pp. 101–115.

Knittel, Christopher R. (2011). “Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector”. In: *American Economic Review* 101.7, pp. 3368–3399.

— (2012). “Reducing Petroleum Consumption from Transportation”. In: *Journal of Economic Perspectives* 26.1, pp. 93–118.

Leard, Benjamin and Virginia McConnell (2017). “New Markets for Credit Trading Under U.S. Automobile Greenhouse Gas and Fuel Economy Standards”. In: *Review of Environmental Economics and Policy* 11.2, pp. 207–226.

— (2021). *Interpreting Tradable Credit Prices in Overlapping Vehicle Regulations*. Working Paper 20-07. Resources for the Future, p. 22.

Leung, Justin H. (2021). “Minimum Wage and Real Wage Inequality: Evidence from Pass-Through to Retail Prices”. In: *The Review of Economics and Statistics* 103.4, pp. 754–769.

Li, Jing (2019). *Compatibility and Investment in the U.S. Electric Vehicle Market*. Working Paper.

Li, Shanjun, Lang Tong, Jianwei Xing, and Yiyi Zhou (2017). “The Market for Electric Vehicles: Indirect Network Effects and Policy Design”. In: *Journal of the Association of Environmental and Resource Economists* 4.1, pp. 89–133.

Linn, Joshua and Virginia McConnell (2017). *The Role of State Policies under Federal Light-Duty Vehicle Greenhouse Gas Emissions Standards*. Report. Resources for the Future, p. 29.

Malueg, David A (1990). “Welfare Consequences of Emission Credit Trading Programs”. In: *Journal of Environmental Economics and Management* 18.1, pp. 66–77.

Malueg, David A. and Andrew J. Yates (2009). “Bilateral Oligopoly, Private Information, and Pollution Permit Markets”. In: *Environmental and Resource Economics* 43.4, pp. 553–572.

Mankiw, N. Gregory and Michael D. Whinston (1986). “Free Entry and Social Inefficiency”. In: *The RAND Journal of Economics* 17.1, pp. 48–58.

Manson, Steven, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles (2020). *IPUMS National Historical Geographic Information System: Version 15.0 [Dataset]*.

McConnell, Virginia and Benjamin Leard (2021). “Pushing New Technology into the Market: California’s Zero Emissions Vehicle Mandate”. In: *Review of Environmental Economics and Policy* 15.1, pp. 169–179.

McConnell, Virginia, Benjamin Leard, and Fred Kardos (2019). *California’s Evolving Zero Emission Vehicle Program: Pulling New Technology into the Market*. Working Paper 19-22. Resources for the Future, p. 51.

Misiolek, Walter S and Harold W Elder (1989). “Exclusionary Manipulation of Markets for Pollution Rights”. In: *Journal of Environmental Economics and Management* 16.2, pp. 156–166.

Muehlegger, Erich and David S. Rapson (2018). *Subsidizing Mass Adoption of Electric Vehicles: Quasi-Experimental Evidence from California*. Working Paper 25359. National Bureau of Economic Research.

— (2020). *Measuring the Environmental Benefits of Electric Vehicles (Relative to the Car That Wasn’t Bought)*. Working Paper 27197. National Bureau of Economic Research.

Murry, Charles and Henry S. Schneider (2016). “The Economics of Retail Markets for New and Used Cars”. In: *Handbook on the Economics of Retailing and Distribution*.

Nevo, Aviv (2001). “Measuring Market Power in the Ready-to-Eat Cereal Industry”. In: *Econometrica* 69.2, pp. 307–342.

— (2003). “New Products, Quality Changes, and Welfare Measures Computed from Estimated Demand Systems”. In: *The Review of Economics and Statistics* 85.2, pp. 266–275.

Newell, Richard, William Pizer, and Jiangfeng Zhang (2005). “Managing Permit Markets to Stabilize Prices”. In: *Environmental and Resource Economics* 31.2, pp. 133–157.

Nykvist, Björn and Måns Nilsson (2015). “Rapidly Falling Costs of Battery Packs for Electric Vehicles”. In: *Nature Climate Change* 5.4, pp. 329–332.

Pakes, A., J. Porter, Kate Ho, and Joy Ishii (2015). “Moment Inequalities and Their Application”. In: *Econometrica* 83.1, pp. 315–334.

Petrin, Amil (2002). “Quantifying the Benefits of New Products: The Case of the Minivan”. In: *Journal of Political Economy* 110.4, pp. 705–729.

Rapson, David S. and Erich Muehlegger (2021). *The Economics of Electric Vehicles*. Working Paper 29093. National Bureau of Economic Research.

Remmy, Kevin (2020). *Subsidy Design When Firms Can Adjust Product Attributes: The Case of Electric Vehicles*. Working Paper, p. 58.

Reynaert, Mathias (2021). “Abatement Strategies and the Cost of Environmental Regulation: Emission Standards on the European Car Market”. In: *The Review of Economic Studies* 88.1, pp. 454–488.

Rubin, Jonathan and Catherine Kling (1993). “An Emission Saved Is an Emission Earned: An Empirical Study of Emission Banking for Light-Duty Vehicle Manufacturers”. In: *Journal of Environmental Economics and Management* 25.3, pp. 257–274.

Rubin, Jonathan, Paul N. Leiby, and David L. Greene (2009). “Tradable Fuel Economy Credits: Competition and Oligopoly”. In: *Journal of Environmental Economics and Management* 58.3, pp. 315–328.

Rubin, Jonathan D. (1996). “A Model of Intertemporal Emission Trading, Banking, and Borrowing”. In: *Journal of Environmental Economics and Management* 31.3, pp. 269–286.

Ruggles, Steven, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek (2019). *IPUMS USA: Version 9.0 [Dataset]*.

Sallee, James M. (2011a). “The Surprising Incidence of Tax Credits for the Toyota Prius”. In: *American Economic Journal: Economic Policy* 3.2, pp. 189–219.

— (2011b). “The Taxation of Fuel Economy”. In: *Tax Policy and the Economy* 25.1, pp. 1–38.

Sartzetakis, Eftichios Sophocles (1997). “Tradeable Emission Permits Regulations in the Presence of Imperfectly Competitive Product Markets: Welfare Implications”. In: *Environmental and Resource Economics* 9.1, pp. 65–81.

— (2004). “On the Efficiency of Competitive Markets for Emission Permits”. In: *Environmental and Resource Economics* 27.1, pp. 1–19.

Schmalensee, Richard and Robert N. Stavins (2013). “The SO_2 Allowance Trading System: The Ironic History of a Grand Policy Experiment”. In: *Journal of Economic Perspectives* 27.1, pp. 103–122.

Spence, Michael (1976). “Product Selection, Fixed Costs, and Monopolistic Competition”. In: *The Review of Economic Studies* 43.2, pp. 217–235.

Springel, Katalin (2021). “Network Externality and Subsidy Structure in Two-Sided Markets: Evidence from Electric Vehicle Incentives”. In: *American Economic Journal: Economic Policy* 13.4, pp. 393–432.

Stavins, Robert N. (1995). “Transaction Costs and Tradeable Permits”. In: *Journal of Environmental Economics and Management* 29.2, pp. 133–148.

— (2003). “Chapter 9 - Experience with Market-Based Environmental Policy Instruments”. In: *Handbook of Environmental Economics*. Ed. by Karl-Göran Mäler and Jeffrey R. Vincent. Vol. 1. Environmental Degradation and Institutional Responses. Elsevier, pp. 355–435.

Stock, James H. and Daniel N. Stuart (2021). *Robust Decarbonization of the US Power Sector: Policy Options*. Tech. rep. w28677. National Bureau of Economic Research.

Tal, Gil and Michael Nicholas (2016). “Exploring the Impact of the Federal Tax Credit on the Plug-In Vehicle Market”. In: *Transportation Research Record* 2572.1, pp. 95–102.

Tietenberg, Thomas H. (1980). “Transferable Discharge Permits and the Control of Stationary Source Air Pollution: A Survey and Synthesis”. In: *Land Economics* 56.4, pp. 391–416.

Vergis, Sydney and Vishal Mehta (2012). “Technology Innovation and Policy: A Case Study of the California ZEV Mandate”. In: *Paving the Road to Sustainable Transport*. Routledge, pp. 157–179.

von der Fehr, Nils-Henrik Mørch (1993). “Tradable Emission Rights and Strategic Interaction”. In: *Environmental and Resource Economics* 3.2, pp. 129–151.

Weitzman, Martin L. (1974). “Prices vs. Quantities”. In: *The Review of Economic Studies* 41.4, pp. 477–491.

Whitefoot, Kate S., Meredith Fowlie, and Steven J. Skerlos (2017). “Compliance by Design: Influence of Acceleration Trade-offs on CO2 Emissions and Costs of Fuel Economy and Greenhouse Gas Regulations”. In: *Environmental Science & Technology* 51.18, pp. 10307–10315.

Wollmann, Thomas G. (2018). “Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles”. In: *American Economic Review* 108.6, pp. 1364–1406.

Xing, Jianwei, Benjamin Leard, and Shanjun Li (2021). “What Does an Electric Vehicle Replace?” In: *Journal of Environmental Economics and Management* 107, p. 102432.

Zhou, Yiyi and Shanjun Li (2018). “Technology Adoption and Critical Mass: The Case of the U.S. Electric Vehicle Market”. In: *The Journal of Industrial Economics* 66.2, pp. 423–480.

Ziegler, Micah S. and Jessika E. Trancik (2021). “Re-Examining Rates of Lithium-Ion Battery Technology Improvement and Cost Decline”. In: *Energy & Environmental Science* 14.4, pp. 1635–1651.

Appendix A

Appendix to Chapter 1

A.1 Details of the ZEV mandate

As shown in Table A.1 (source: 13 CCR §1962.1(d)(5)(A)), the number of credits earned per vehicle was a function of its range. (The regulation used the UDDS urban driving range, which is about 40% higher than the EPA range.) In almost all cases, electric vehicles earned between two and four credits and did not qualify for fast refueling. (An exception is the longer-range versions of the Tesla Model S, which qualified in 2012 and 2013 on the basis of an experimental battery swap program.[1])

Table A.1: *ZEV credits, model years 2009–2017*

	Criteria		
Tier	UDDS Range (mi)	Fast Refueling	Credits
Type I	$[50, 75)$	–	2
Type I.5	$[75, 100)$	–	2.5
Type II	≥ 100	–	3
Type III	≥ 200	–	4
Type III	≥ 100	Yes	4
Type IV	> 200	Yes	5
Type V	≥ 300	Yes	7 (9 after 7/2015)

The credit requirement in each year is formulated as fixed percentage of the manufac-

[1] See "Tesla profits could be challenged by Calif. credit-rule change" (Mark Rechtin, Automotive News, 8/5/13).

turer's "production volume" of non-zero-emission passenger cars and light-duty trucks. (Before 2009 only light-duty trucks under 3750 pounds loaded weight were counted; between model years 2009 and 2012 this cutoff was raised to 8500 pounds.) The manufacturer chose in each year whether its production volume was its same-year sales or a function of past sales. In model years 2009 through 2011, the past-sales function was the average of sales in model years 2003–2005; in model years 2012 through 2017, the past-sales function in year t was the average of sales in model years $t-6$ through $t-4$ (13 CCR §1962.1(b)(1)(B)).

The credit requirement percentage for applicable manufacturers is shown in Table A.2 (source: 13 CCR §1962.1(b)(2)).

Table A.2: *Large Volume Manufacturer requirements by model year, 2009–2017*

Model Years	Minimum ZEV
2009–2011	2.475%
2012–2014	0.790%
2015–2017	3.000%

In model years 2009–2011, large manufacturers could opt for a lower requirement of 0.205% if they did not use traded credits to meet it (13 CCR §1962.1(b)(2)(B)).

A.1.1 Travel provision

The travel provision allows manufacturers to count credits from certain vehicles sold in one ZEV state toward requirements in all ZEV states. It does not have to be exercised in the same model year the car was delivered; a credit can be banked or traded and then traveled later.

In model year 2009, credits for ZEVs travel one-for-one between California and the other ZEV states (13 CCR §1962.1(d)(5)(E)). Because model year 2009 ZEVs were not sold commercially, we ignore them.

In model years 2010–2017, credits for ZEVs travel proportionally. Suppose an eligible vehicle by manufacturer m, which earns x credits, is placed into service in a ZEV state in model year t. Then, if the manufacturer chooses to travel the credit, it translates into the following in each state s (including the state where it was originally placed into service):

$$x \cdot \frac{\text{Sales Volume}_{m,t,s}}{\text{Sales Volume}_{m,t,\text{CA}}},$$

where Sales Volume is the same-year sales of non-zero-emission cars and light-duty trucks in the state.