Thesis on Environment and Transport Economics

Lichao Chen

Submitted in fulfillment of the requirements for the degree of

Doctor of Philosophy

University College London (UCL)

Department of Economics

October 2025

Declaration

I, Lichao Chen, confirm that the work presented in my thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

This thesis examines how supply-side carbon pricing and demand-side purchase subsidies reshape market structure and welfare in Europe's two largest transport sectors: aviation and automobiles. Chapter 1 analyses the European airline industry using a two-stage entry-and-pricing model that captures key institutional features such as airport slot constraints and point-to-point business models. Fixed-cost parameters are estimated through a hybrid approach combining moment inequalities with maximum likelihood, ensuring policy simulations remain consistent with observed market structures. The analysis shows that carbon pricing induces asymmetric network adjustments concentrated among low-cost and regional carriers, while full-service groups at hub airports remain relatively resilient. Although industry profits decline, the reallocation of capacity improves allocative efficiency and redistributes welfare unevenly across Europe. Chapter 2 evaluates electric-vehicle purchase subsidies in the UK, France, and Germany (2010–2021) using a random-coefficients logit demand model with micro-moment calibration and a static Bertrand pricing framework. The results show that the expansion of the EV market has been driven mainly by product innovation and model fleet turnover rather than by flat purchase incentives. Subsidy effectiveness varies across countries, with limited impacts in the UK and France but stronger effects in Germany. An income-targeted subsidy design achieves similar emissions reductions at substantially lower fiscal cost and with greater equity. Together, the chapters demonstrate that environmental policies operate through distinct mechanisms—reconfiguring airline networks and influencing car buyers' choices—and that well-designed instruments can achieve decarbonisation with higher efficiency and fairer distributional outcomes.

Impact Statement

This thesis investigates how climate policy reshapes competition, consumer choice, and welfare in Europe's transport sector, treating aviation and road transport as complementary testbeds for the supply- and demand-side instruments at the core of the EU's decarbonisation strategy. The contribution is twofold. Substantively, it provides comparable, policy-relevant evidence on who bears the costs and who benefits when carbon prices and purchase subsidies interact with market structure and household heterogeneity. Methodologically, it advances estimation and counterfactual design so that simulated policies are disciplined by observed networks and by micro-evidence on preferences, improving external validity for ex-ante policy appraisal.

Chapter 1 is, to our knowledge, the first paper to study how carbon regulation alters competition on a European airline route network, jointly modelling endogenous network formation and price competition under binding slot constraints. It estimates a two-stage game in which airlines choose routes and frequencies before pricing, and introduces a hybrid identification strategy: moment inequalities identify linear fixed-cost components, while maximum likelihood recovers the full distribution of unobserved fixed-cost shocks. This permits drawing shock realisations that exactly rationalise the observed network, eliminating baseline drift and anchoring counterfactuals in the data. The counterfactuals—implemented via an iterative equilibrium algorithm under EU-ETS-style carbon pricing—show asymmetric network responses concentrated among low-cost and regional carriers, with full-service groups at congested hubs comparatively resilient. Despite lower industry profits, carbon pricing can raise total welfare by reallocating aircraft toward higher-value services, while redistributing welfare geographically across Europe. Conceptually, the chapter reframes carbon policy from a pure cost shock to a force that reconfigures market structure in a highly congested, mixed-model industry, and methodologically contributes a portable way to combine set- and point-identification for credible network counterfactuals.

Chapter 2 provides a unified, multi-country structural evidence on the drivers of Europe's EV uptake and on the incidence and design of purchase subsidies. Using annual model-level sales and characteristics for the UK, France, and Germany (2010–2021), augmented with micro-moments that link purchasing to income, it estimates a random-coefficients logit with a static Bertrand supply side. Decompositions show that the surge in EV market share is driven primarily by product turnover and attributes—not by flat purchase subsidies—with limited effects in the UK and France and larger effects in Germany. The income-linked micro-moments reveal pronounced heterogeneity: poorer households are more price-sensitive yet have weaker baseline EV tastes, implying uneven welfare gains under uniform subsidies. The chapter then designs income-based subsidies and shows they can deliver comparable or higher EV shares at substantially lower fiscal cost, while reshaping the distribution of benefits toward lower-income households—thus clarifying efficiency–equity trade-offs and offering implementable policy rules for Europe, where income-based EV subsidies are not yet standard practice.

Together, the chapters show that climate policy's real effects emerge from its interaction with industrial organisation on the supply side and with heterogeneous preferences on the demand side. By pairing network-credible airline counterfactuals with income-disciplined EV demand, the thesis offers a coherent template for appraising transport decarbonisation policies that is both empirically grounded and directly usable by policymakers.

Acknowledgements

I thank my main supervisor, Lars Nesheim, for his encouragement and guidance throughout this long journey. When I felt lost in my first year, he generously supported my shift into empirical industrial organisation, which proved to be an inflection point in my PhD. His mentoring combines calm, thoughtful advice with the freedom to explore my own interests rather than prescribing a fixed path. Whenever difficulties arose—whether in data acquisition, high-performance computing, or career planning—he consistently went out of his way to help. I am also grateful to my second supervisor, Joao Granja, for his help in countless situations, and to Aureo de Paula for invaluable suggestions on virtually every aspect of both chapters.

I thank all co-authors and collaborators. In particular, I am profoundly grateful to Ertian Chen, whose brilliance and extensive comments have continually inspired me to push further. I am also grateful to colleagues in the Air Transportation Systems Lab at the UCL Energy Institute, including Andreas Schafer, Lynnette Dray, Khan Doyme, Peggy Li, Brian Pearce, Olivier Dessens, Yagmur Gok, and Boning Yang, for generously providing key European airline data and for their many helpful suggestions from an industry perspective. I will miss our weekly meetings.

I would also like to thank the IO group in the department, especially Javier Boncompte and Giuseppe Forte, for numerous insightful discussions on empirical IO. I am also thankful to participants at seminars and workshops for their constructive comments throughout.

Finally, I am grateful to all data providers who made this dissertation possible, and to Simeon Duckworth and Mateusz Mysliwski for exceptional industry connections and conversations that were crucial for an empirical IO project of this scope.

Introduction

This thesis examines how environmental policy reshapes competition, consumer choice, and welfare in Europe's two largest transport sectors: aviation and automobiles. It combines a supply-side study of carbon regulation in airline markets with a demand-side study of purchase subsidies in car markets to deliver comparable, policy-relevant evidence on the incidence and efficiency of climate instruments. In both chapters, policy counterfactuals are analysed using parameters estimated from the data: in aviation, within a route-network formation framework under binding capacity (slot) constraints; in the automotive sector, through the design and evaluation of income-based electric-vehicle subsidies. Together, these analyses yield an integrated, empirically grounded framework for ex-ante appraisal of decarbonisation policies.

Chapter 1: Carbon regulation and competition in European aviation

Chapter 1 examines how carbon regulation alters competition on a European airline route network while jointly modelling endogenous network formation and price competition under binding airport slot constraints. The analysis uses proprietary itinerary-level data for 2019 with prices, frequencies, travel times, and passenger volumes at the route–airline level, supplemented by metropolitan populations and airport-to-airport distances. The model is a two-stage oligopoly: in the first stage, carriers choose route networks and frequencies subject to aircraft and slot constraints; in the second stage, they set prices conditional on those choices.

Fixed-cost parameters are recovered with a hybrid identification strategy. Moment inequalities identify linear components of entry and frequency costs, and maximum likelihood recovers the distribution of unobserved fixed-cost shocks. This allows the baseline network to be rationalised exactly, which improves the credibility of counterfactuals by avoiding drift from the observed equilibrium. Policy experiments implement an EU-ETS-style carbon price and compute new network equilibria through an iterative best-response procedure in which airlines re-optimise route choices subject to capacity limits.

Four results follow. First, demand and cost differences between full-service and low-cost carriers are pronounced, especially in the valuation of hub access and fixed costs. Second, carbon pricing produces asymmetric network adjustments that are concentrated among low-cost and regional carriers, while full-service groups anchored at congested hubs are comparatively resilient. Third, welfare is redistributed across space as connectivity patterns change, with gains and losses varying across European regions. Finally, despite lower industry profits, overall welfare can rise because capacity is reallocated toward higher-value services. Conceptually, the chapter reframes carbon policy as a force that reconfigures market structure in a slot-constrained, mixed-business-model industry. Methodologically, it contributes a portable way to combine set and point identification for network-credible counterfactuals.

Chapter 2: Purchase subsidies, drivers of EV uptake, and equity in Europe

Chapter 2 provides unified, multi-country structural evidence on the drivers of Europe's electric-vehicle uptake and on the incidence and design of purchase subsidies. It assembles annual model-level sales and characteristics for the UK, France, and Germany from 2010 to 2021 and augments these data with micro moments that link vehicle purchasing to income. The empirical framework is a random-coefficients logit demand system paired with a static Bertrand supply side, with interactions that allow preferences to evolve by time, brand, and fuel type.

Three findings stand out. First, decompositions attribute most of the EV market-share growth to new-model introductions and changes in product attributes rather than to flat purchase subsidies. Second, subsidy effectiveness differs across countries, with relatively modest impacts in the UK and France and larger impacts in Germany. Third, welfare gains are uneven across income groups: lower-income households are more price sensitive but exhibit weaker baseline EV tastes, which limits their realised benefits under uniform subsidies. Motivated by this incidence pattern, the chapter designs income-based subsidies targeted at low- and middle-income households. Counterfactuals show that such designs can deliver comparable or higher EV sales at substantially lower fiscal cost while shifting benefits toward lower-income consumers. Budget-equivalent comparisons clarify efficiency–equity trade-offs and provide implementable rules for European policymakers.

Synthesis and policy relevance

Together, the chapters show that the real effects of climate policy are shaped by product (route) entry and exit on the supply side and by heterogeneous preferences on the demand side. In aviation, carbon pricing interacts with congestion, network choices, and business-model heterogeneity to reshape competitive structure and the spatial distribution of welfare. In the automotive sector, purchase incentives interact with product turnover and income-linked tastes to alter fleet composition and the distribution of benefits across households. By pairing network-credible airline counterfactuals with income-disciplined EV demand, the thesis provides a coherent template for appraising transport decarbonisation policies—externally valid, attentive to distributional outcomes, and directly applicable to policy design.

Contents

Declaration 2

Thesis Abstract 4

Impact Statement 6

Acknowledgements 9

Thesis Introduction 11

1. Carbon Regulation and Competition in European Airline Industry 19

1.1 Introduction 21

1.2 Background and Data 26

1.2.1 The Unique Structure of the European Airline Market 26

1.2.2 Data 32

1.3 Model 45

1.3.1 Model Setting: Demand and Price Competition 47

1.3.2 Model Setting: Entry and Frequency 51

1.3.3 Information Structure and Equilibrium Concept 53

1.4 Identification and Estimation Strategies 56

1.4.1 Construction of Moment Inequalities 59

1.4.2 Feasible Alternative Route Network 61

1.4.3 Estimation of the Linear Fixed Cost Parameters 64

1.4.4 Estimation of the Distribution of Fixed Cost Shocks 66

1.5 Estimation Results 69

1.5.1 Demand Estimates 69

1.5.2 Marginal Cost on Passengers 75

1.5.3 Linear Fixed Cost Estimation 78

1.5.4 Fixed Cost Shocks Estimation 87

1.6 Counterfactual Experiment on Carbon Policy 91

1.6.1 Drawing Fixed Cost Shocks from Estimated Distributions 94

1.6.2 Equilibrium Concept and Iterative Algorithm 96

1.6.3 Counterfactual Results 99

1.7 Conclusion 113

2. Does The Subsidy on Electric Vehicles Reach Its Target? Evidence From The European Car Market 116

2.1 Introduction 118

2.2 Literature 124

2.3 Background and Data 128

2.3.1 Industry Background 128

2.3.2 Policy Background 131

2.3.3 Data (Car Sales) 132

2.3.4 Data (Socio-Demographic) 136

2.4 Empirical Model 139

2.4.1 Set-up 139

2.4.2 Demand 140

2.4.3 Supply 144

2.5 Estimation **146**

2.5.1 Instrumental Variables 146

2.5.2 Identification 147

2.5.3 Estimation Approach 148

2.6 Estimation Results **152**

2.7 Welfare Analysis **156**

2.8 Counterfactual One: Driver of the Market Share Changes **158**

2.9 Counterfactual Two: Income-based Subsidy Policy **163**

2.10 Conclusion **169**

2.11 Tables **171**

2.12 Figures **183**

Chapter 1

Carbon Regulation and Competition in European Airline Industry

with Ertian Chen (UCL), Lars Nesheim (UCL), and Andreas Schafer (UCL)

Abstract

This paper quantifies the effects of carbon regulation on airline competition and endogenous route network formation in Europe. We estimate a two-stage structural model where full-service and low-cost carriers choose route networks before competing on prices. Our framework is the first to jointly analyse these dynamics while accounting for unique European features like binding airport slot constraints. A novel estimation strategy combines moment inequalities with maximum likelihood to ensure credible policy simulations. Counterfactual analysis of a carbon tax reveals asymmetric impacts: network changes are concentrated amongst low-cost and regional airlines, whilst full-service carriers prove resilient. The policy also induces a significant geographic redistribution of welfare, benefiting Central and Eastern Europe at the expense of remote regions. Importantly, whilst airline profits fall, the policy can be total welfare-enhancing by forcing a competitive reallocation of aircraft that improves allocative efficiency.

Keywords: European Aviation, Carbon Policy, Route Entry & Exit

JEL Codes: L52, L62, L90

1.1 Introduction

The aviation industry, a cornerstone of global connectivity, presents one of the most formidable challenges in the transition to a low-carbon economy. In 2022, the sector accounted for nearly 4% of the European Union's total greenhouse gas emissions, a figure projected to grow rapidly as air travel rebounds and expands (European Commission; International Energy Agency). The scale of the problem is striking: a single round trip from Lisbon to New York emits carbon equivalent to heating a typical EU home for an entire year. While technological advancements are promising, their large-scale impact remains a distant prospect. Gains from more efficient new-generation aircraft are incremental and slow to materialise across a global fleet with operational lifespans often exceeding 25 years; historically, these efficiency gains have been outpaced by demand growth. Sustainable Aviation Fuels (SAFs), a key alternative, face prohibitive costs and severe supply constraints, with production capacity at less than 1% of global demand. Truly transformative solutions, such as hydrogen or electric propulsion, are not expected to be commercially viable for most routes until the 2030s at the earliest. Consequently, in the short to medium term, meaningful emissions reductions will likely depend less on innovation and more on regulation.

This paper investigates the competitive and network-level effects of increasingly stringent carbon regulations within the European airline industry—one of the world's largest and most complex aviation market. The European market's structure is unique, shaped by factors that distinguish it sharply from its North American or Asia-Pacific counterparts. It is the world's most mature low-cost carrier (LCC) market, where budget airlines operate a majority of short-haul flights, fostering intense price competition. This is coupled with a network dominated by direct, point-to-point routes, a departure from the hub-and-spoke models prevalent in the U.S. Furthermore, the continent suffers from severe infrastructure constraints, hosting nearly half of the world's most congested, slot-coordinated airports (IATA), and features high aircraft utilisation rates driven by the operational efficiency and rapid turnaround times pioneered by LCCs.

This market structure gives rise to a uniquely intense competitive dynamic, driven by the deep bifurcation of airline business models. Full-service carriers (FSCs) typically operate from major, congested primary airports, leveraging network economies to serve both point-to-point and international connecting traffic. In contrast, the highly prevalent LCCs exploit a point-to-point model, often from smaller, secondary airports within the same metropolitan area to minimise operational costs. These divergent strategies create starkly different cost structures and fare strategies; LCCs leverage their operational efficiencies to offer lower base fares and unbundled services, capturing a price-sensitive market segment of a scale not seen elsewhere. Crucially, the point-to-point model affords LCCs greater strategic flexibility in network expansion. By serving a wider portfolio of cities, LCCs possess a combinatorially larger set of feasible new routes to enter, allowing them to rapidly redeploy aircraft to capture emerging demand in markets that may be too thin or unprofitable for the more rigid hub-and-spoke structure of an FSC. This fundamentally alters the calculus of route entry and profitability across the continent.

It is within this complex competitive environment that Europe is implementing some of the world's most stringent aviation carbon policies, which are poised to significantly affect airline operations. The progressive phasing out of free allowances under the EU Emissions Trading System (EU-ETS) is set to dramatically increase the effective carbon price for airlines by 2026. This is compounded by the ReFuelEU mandate, which requires an increasing blend of Sustainable Aviation Fuels (SAFs)—

a technology that remains several times more expensive than conventional jet fuel and faces significant production shortfalls. These cost shocks will not be neutral; they will disproportionately impact airlines based on their business models, route structures, and margins, making the interaction between regulation and competition a first-order question for the industry's future.

To that end, this paper is the first to jointly analyse the effects of carbon regulation on airline competition and endogenous route network formation in the European context. We ask: How do the distinct business models of FSCs and LCCs shape their strategic responses to rising carbon costs? How does regulation alter market structure through route entry and exit? And what are the ultimate consequences for consumer welfare and its geographic distribution across Europe?

To answer these questions, we develop and estimate a two-stage static game of oligopoly competition. In the first stage, airlines endogenously choose their route networks and flight frequencies. In the second stage, conditional on the established network, they compete on prices. Our primary methodological contribution lies in a novel, hybrid estimation strategy for the fixed costs of route operation. We first use moment inequalities to set-identify the linear parameters of the fixed cost function, following recent advances in the industrial organisation literature (Ho and Pakes [2014]; Pakes et al. [2015]). We then, in a novel second step, estimate the full distribution of the unobserved fixed cost shocks via maximum likelihood. This allows us to draw a specific realisation of these shocks in a way that ensures our counterfactual simulations begin from the *actual observed network*, lending significant credibility to our policy analysis. Demand and marginal cost parameters are estimated using established methods (Berry and Jia [2010]; Bontemps et al. [2023]).

Our counterfactual analysis simulates the impact of a carbon tax, implemented through the EU Emissions Trading System (EU-ETS). The simulation finds a new network equilibrium using an iterative algorithm where airlines sequentially re-optimise their route choices. Our key findings are fourfold. First, our estimates reveal stark differences in the demand and cost structures of FSCs and LCCs, particularly in the valuation of hub airports and the underlying distribution of fixed costs. Second, the impacts of a carbon tax are highly asymmetric: network adjustments are concentrated amongst LCCs and smaller regional carriers, whilst large FSCs with valuable and congested hubs prove remarkably resilient. Third, the policy induces a significant geographic redistribution of welfare. Central and Eastern European countries benefit from intensified competition on shorter routes, while remote regions like Iceland and Norway suffer from reduced connectivity. Finally, we find that despite reducing airline profits, carbon pricing can be *total welfare-enhancing*. The policy not only prices the environmental externality but also forces a competitive reallocation of aircraft that improves allocative efficiency in an imperfectly competitive market, suggesting a potential "double dividend".

Literature and Contribution

This paper contributes to four distinct strands of literature. First, we build on the rich body of work estimating structural models of airline competition. The vast majority of this research, however, focuses on the U.S. market, where hub-and-spoke networks and connecting traffic are central to competition (Bontemps et al. [2023]; Yuan and Barwick [2024]; Aguirregabiria and Ho [2012]). Our focus on the intra-European market is not merely a geographic shift; it necessitates a fundamental change in modelling approach. With direct, point-to-point flights accounting for 94% of passengers in our data, complex network spillovers are of second-order importance. Instead, the critical institutional feature is the prevalence of binding slot constraints at major airports. This motivates one of our key modelling innovations: we frame an airline's choice not as a series of independent entry/exit decisions, but as a problem of redeploying a fixed stock of aircraft—and their associated valuable slots—across a set of feasible routes.

Second, we contribute to the small strand but growing literature on competition in the European airline market. While existing studies have examined specific features such as slot allocation (Marra [2024]), LCC subsidies (Bontemps et al. [2024]), or mergers (Bergantino et al. [2024]), our paper provides the first integrated analysis of how the defining features of European competition—the FSC versus LCC dynamic, and hub airports congestion—interact with environmental regulation to shape market-wide outcomes. We also add to the literature on other slot constraints (Park [2020]; Ciliberto and Williams [2014]; Forbes and Lederman [2009]; Argyres et al. [2024]) by incorporating their effects structurally through aircraft utilisation constraints, reflecting the reality that an airline exiting a route from a congested hub must redeploy its aircraft to retain its valuable slot.

Third, we advance the literature on the economic impacts of carbon regulation. While many studies on carbon pricing focus on environmental efficacy or policy design (Timilsina [2022]), we examine how such policies fundamentally reconfigure an oligopolistic industry. Our approach is related in spirit to Ryan [2012], who studies environmental regulation in the U.S. cement industry using a dynamic oligopoly model. While a full dynamic model is computationally infeasible for the thousands of city-pair markets in our setting, we adapt the core insight: environmental policy is not just a cost shock, but a catalyst for changes in market structure, concentration, and welfare. This paper is the first to apply this lens to the European airline industry, quantifying the competitive fallout of its unique and stringent carbon policies.

Finally, we make a methodological contribution to the estimation of entry games. While two-stage models are common, our hybrid estimation of fixed costs is novel.

The existing literature listed above typically either ignores unobserved fixed cost shocks in counterfactuals or relies solely on moment inequalities, which can only set-identify parameters. By combining moment inequalities for the linear parameters with a maximum likelihood estimation of the full shock distribution, we can point-identify the group-specific means and variances of these shocks. This allows us to then draw realisations of the shocks that perfectly rationalise the observed network as the baseline equilibrium, which is missing in all previous literature. This step is crucial for ensuring the accuracy and credibility of policy counterfactuals, as it eliminates the "baseline drift" that can undermine simulations in complex structural models.

The remainder of this paper is structured as follows. Section 2 reviews the European airline market and our dataset. Section 3 presents the two-stage model. Section 4 discusses estimation and identification. Section 5 reports parameter estimates. Section 6 presents the counterfactual analysis of the EU-ETS. Section 7 concludes.

1.2 Background and Data

1.2.1 The Unique Structure of the European Airline Market

The Rise of European Low-Cost Carriers Following the deregulation of European aviation, low-cost carriers (LCCs) have fundamentally reshaped the continent's competitive landscape. Their market share surged from just 5.3% in 2001 to approximately 35% by 2022. As Figure 1 illustrates, LCCs now consistently account for nearly half of all intra-European passenger traffic. The scale of this transformation is exemplified by Ryanair, which in 2023 carried 182 million passengers—more than any single full-service carrier (FSC) in Europe.[1] The LCC sector itself is heterogeneous, comprising two main archetypes: subsidiaries of legacy FSC groups (such as Vueling, Eurowings, and Transavia) and independent, 'pure-play' LCCs (such as Ryanair, EasyJet, and Wizz Air). It is this latter group, with its distinct business models, that has been the primary driver of market disruption.

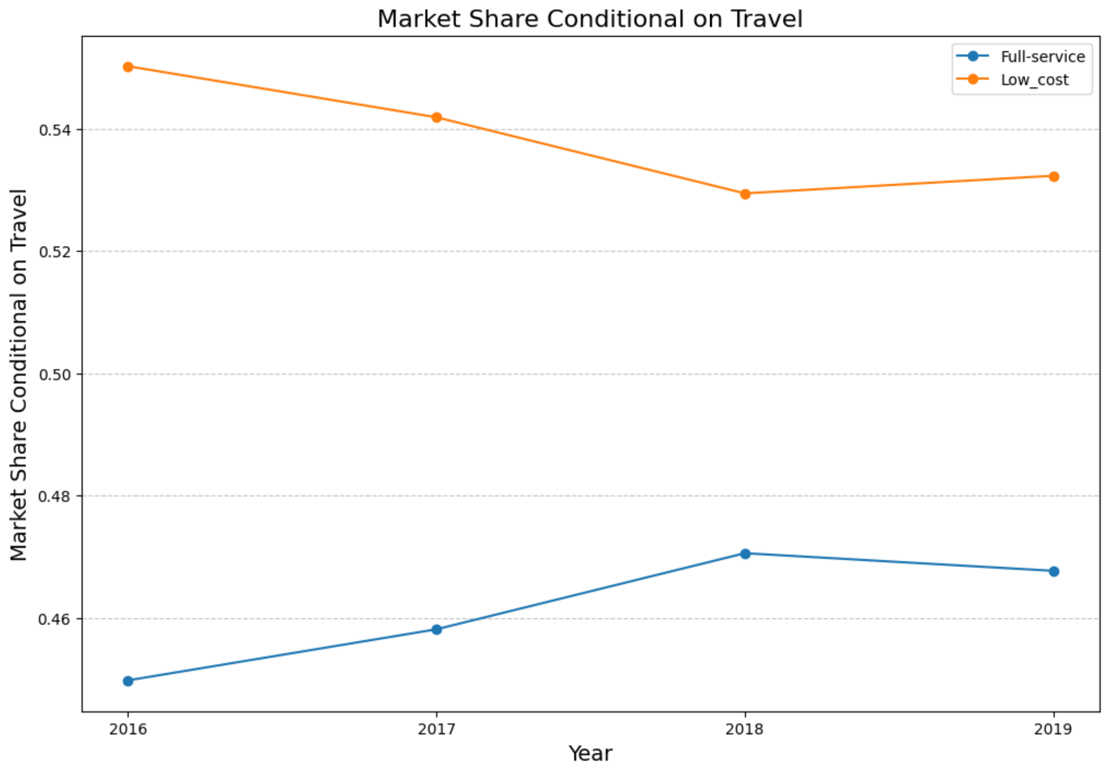

Figure 1: Market Share for Different Types of Airlines

Full-service carriers (FSCs) and low-cost carriers (LCCs) operate according to starkly different business models. First, a primary strategic divergence lies in network architecture. Full-service carriers (FSCs) typically employ a *Hub-and-Spoke* model, where the network is anchored by one or more major hubs—such as London Heathrow for British Airways, Paris Charles de Gaulle for Air France, and Amsterdam Schiphol for KLM. These hubs serve a dual purpose: they achieve economies of

[1] Statista.

scale in ground operations and, critically, aggregate short-haul European traffic to feed lucrative long-haul intercontinental services. In stark contrast, LCCs operate decentralised *Point-to-Point* (P2P) networks, which offer direct flights between a wide array of city pairs, providing greater routing flexibility.

The distinction is visually apparent in Figure 2, which contrasts the hub-centric network of Air France-KLM with the diffuse, web-like structure of Ryanair. While Ryanair maintains large operational bases at airports like London Stansted, these do not function as connecting hubs for transfer passengers; their strategic role is to serve large origin-destination markets, not to facilitate transfers, underscoring the airline's strict adherence to the P2P model.[2]

Second, cost structures differ markedly. FSCs incur higher per-passenger and per-flight costs, driven by operations at expensive hub airports, lower fleet utilisation, and premium offerings like business class and meal services. According to KPMG, the cost per available seat kilometre for LCCs (excluding fuel) is 20%–30% lower than for FSCs,[3] granting them a substantial pricing advantage.

Third, service levels and airport selection strategies diverge. LCCs 'unbundle' their product, earning a significant portion of revenue from ancillary fees for services like baggage handling and seat selection.[4] In contrast, FSCs traditionally offer a more inclusive fare. This strategic bifurcation extends to airport choice, which is particularly notable in Europe's multi-airport metropolitan areas. FSCs typically operate from large international hubs, while LCCs favour smaller, secondary airports. London provides the clearest example across its six airports: Heathrow serves

[2]Ryanair operates de facto hubs at London Stansted and Dublin. However, these are primarily used as operational bases for aircraft and do not function as international hubs in the FSC sense. Their significance lies in serving large local markets rather than facilitating connecting traffic.

[3]KPMG.

[4]While FSCs increasingly adopt similar pricing practices, they are still generally perceived as offering higher service quality. See: Daily Telegraph.

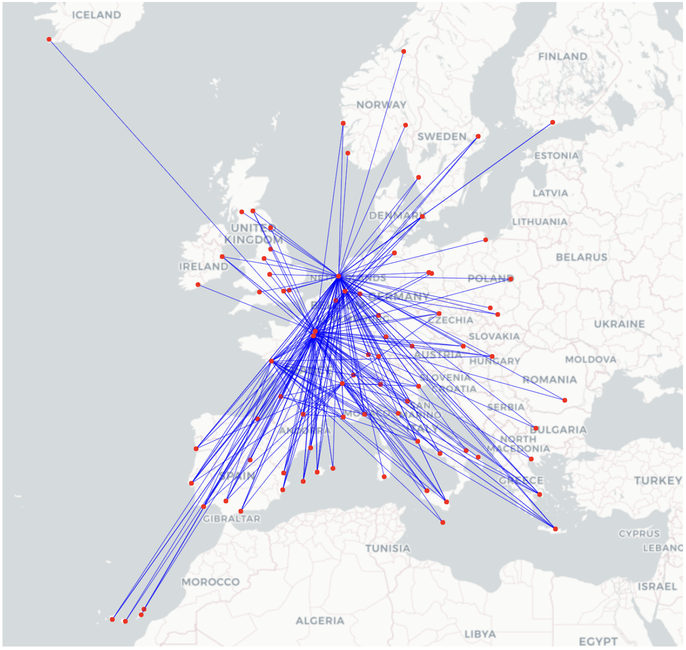

Air France-KLM Group Route Map

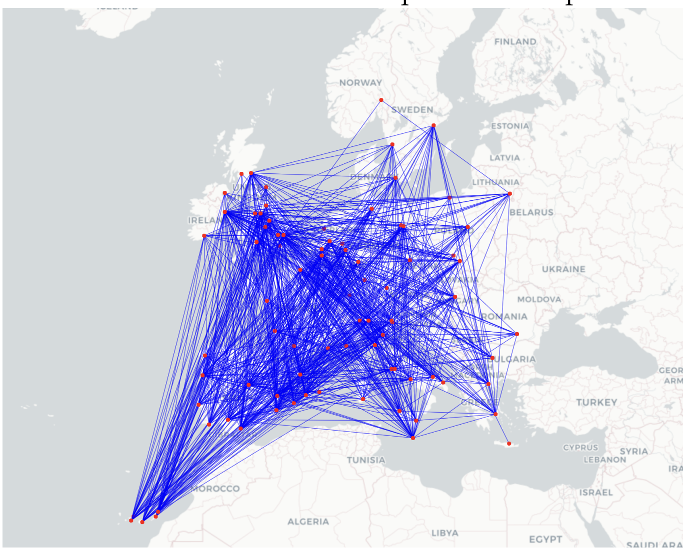

Ryanair Route Map

Figure 2: Comparing Route Map Between Full-service and Low-cost Airlines

almost exclusively FSCs as the principal international hub; Gatwick accommodates both; Stansted and Luton are major LCC bases; and the City and Southend airports cater to specialised segments.[5] Although Heathrow is the most connected, its severe

[5]London City mainly serves business routes (e.g., London–Paris or London–Frankfurt), while

capacity constraints and high airport charges make it economically unattractive to the LCC business model.

To formalise these strategic distinctions, we embed the defining characteristics of FSC and LCC business models directly within our structural framework. Differences in network formation are captured through constraints on feasible route entries, while variations in service quality and airport choice are modelled via airline- and airport-specific fixed effects. Finally, divergent operational efficiencies are reflected in differentiated cost structures. This specification is crucial for identifying the heterogeneous strategic responses of each airline type to a uniform carbon regulation.

Binding Slot Constraints in European Airports Europe is home to some of the world's most congested airports. Several major hubs—most notably London Heathrow—have operated at or near full capacity for decades, yet expanding this infrastructure faces significant barriers, including stringent regulatory constraints, political opposition, and financing challenges (ACI Europe). Consequently, airline operations are governed by a rigid system of slot controls. A slot grants an airline the right to use a runway for takeoff or landing during a specific time window, and its allocation is critical in capacity-constrained environments. The scale of this issue is unique to Europe, which accounts for nearly half of all IATA Level-3 slot-coordinated airports worldwide—those where demand consistently exceeds capacity.[6]

The allocation of these scarce slots is a contentious issue in European aviation policy. The current system, established in 1993, is built on a "grandfathering" principle, where an airline retains its historical slots provided it meets a minimum

Southend is dominated by charter airlines.

[6]IATA classifies airports into three categories: Level-1 airports have no significant congestion; Level-2 airports may require coordination; Level-3 airports consistently face demand that exceeds available capacity. This system is widely used to measure airport congestion.

usage threshold—typically 80% in a given season.[7] This "use it or lose it" rule creates powerful incumbency advantages, particularly for established national carriers who can maintain control over valuable slot portfolios. The immense strategic value of these slots has led to perverse incentives, most notably the operation of near-empty "ghost flights" during periods of low demand, flown solely to satisfy usage requirements and prevent the forfeiture of a prized asset.[8]

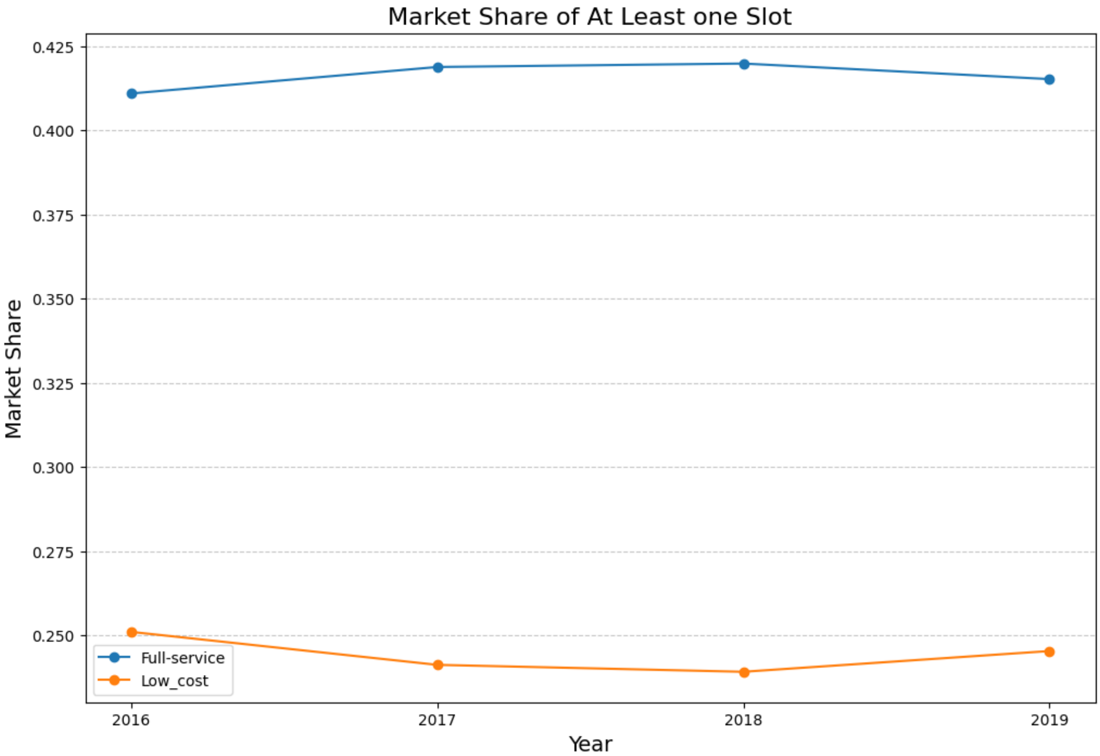

Figure 3: Proportion of Passengers Travelling from/to at Least one Hub

Rather than proposing specific reforms to the EU's slot allocation system (see Marra [2024] for a comprehensive treatment), this paper incorporates the binding nature of slot constraints directly into our modelling framework. We treat slot availability as a key constraint that limits the set of feasible frequency adjustments an

[7]European Union.

[8]This phenomenon was widely reported during the pandemic. See: Forbes.

airline can make when entering new routes. The strategic importance of these airports varies significantly by business model, a pattern illustrated in Figure 3, which shows that passengers flying with full-service carriers are far more likely to travel through slot-controlled airports than their low-cost counterparts. This underscores the fundamental trade-off these airports present: they offer superior connectivity and infrastructure but impose higher operating costs and congestion. We explore precisely how these airport characteristics shape airline revenues, costs, and network expansion strategies for each carrier type in Sections 3 and 4.

1.2.2 Data

A significant challenge in studying the European airline market is the limited availability of public data, in stark contrast to the U.S. market where sources like the DB1B dataset are readily accessible. To overcome this, our analysis is built upon a proprietary dataset from Sabre Market Intelligence,[9] a global distribution system that provides travel reservation and pricing tools for many of Europe's largest airlines, including IAG Group, Air France-KLM Group, Lufthansa Group, EasyJet, and Wizz Air. Because this system is actively used by airlines for fare optimisation, it offers highly accurate, itinerary-level pricing information—a critical component for demand estimation that is typically absent from public administrative datasets.

The raw Sabre data are organised at the itinerary level, defined as a specific airline's service between an origin and destination airport. Each observation includes key characteristics such as average airfare (price), flight frequency, travel time, and passenger volume. We make two key processing decisions. First, given that only 6% of European passengers in our sample travel on connecting flights, we restrict our analysis to the direct flight market. Second, following the methodology of Yuan

[9] Sabre.

and Barwick [2024] and Bontemps et al. [2023], we aggregate directional itineraries (e.g., A to B and B to A by the same airline) into a single non-directional route. This aggregation reflects the operational reality that airlines typically operate return services with nearly identical prices and frequencies, allowing us to sum passenger counts across both directions.

Our sample period comprises the four quarters of 2019. We select this year as our baseline because it represents the most recent period of 'normal' market conditions before the profound structural disruption caused by the COVID-19 pandemic. As of our analysis, data reflecting a full post-pandemic recovery are not yet available,[10] making 2019 the most suitable reference for a stable equilibrium from which to analyse policy-induced changes to route networks.

Finally, we supplement the Sabre data with two external sources: metropolitan population data from Eurostat,[11] which we use to construct our market size variable, and airport-to-airport surface distances obtained via the *Google Distance Matrix API*.

Market Definition A critical modelling choice in our analysis is the definition of a market. We define each market as a city-pair, representing the origin and destination metropolitan areas, observed quarterly in 2019. We use city-pairs rather than airport-pairs for two primary reasons related to competitive dynamics and consumer behaviour.

First, a city-pair definition is essential to capture the true extent of competition between full-service and low-cost carriers. LCCs frequently operate from secondary airports within a major metropolitan area to minimise fees and operational costs—a

[10]Indeed, it is debatable whether the European airline market has fully returned to its pre-pandemic structure and performance.

[11]European Union.

strategy exemplified by their near-total absence from London Heathrow and strong presence at airports like Gatwick, Luton, and Stansted. Defining markets at the airport-pair level would therefore fail to capture the direct substitution that occurs between an FSC at a primary hub and an LCC at a secondary airport serving the same urban catchment area.

Second, consumers exhibit distinct preferences for different airports within the same city. Primary hub airports typically offer superior ground services and public transport connections, but are often more congested. Conversely, non-hub airports may be located farther from the city centre, imposing additional ground travel time. To capture this variation in consumer utility, our demand model incorporates dummy variables for all major hub airports, allowing us to estimate the value consumers place on hub status.

To implement this, we adopt Eurostat's official definition of "metropolitan regions" to delineate city boundaries.[12] This classification provides a robust and economically meaningful definition of a city, accounting for dimensions such as population density, commuting flows, and transport integration. This approach ensures that all airports located within a given pair of metropolitan regions are correctly treated as effective substitutes within the same air travel market.

Our definition captures not only well-known multi-airport systems like Greater London and the Paris metropolitan area, but also less obvious, economically integrated regions such as the Düsseldorf/Dortmund/Cologne corridor in Germany's Ruhr area. It also includes cross-border metropolitan regions with integrated transport systems, such as Copenhagen/Malmö and Vienna/Bratislava. For our analysis, we include all metropolitan areas with populations exceeding 850,000, which collectively account for over 90% of all European air travel.

[12] European Union.

Hubs, Airline Aggregation, and Slot Constraints Hub airports are central to the operations of full-service carriers (FSCs). Table 1 lists the parent airline groups, their associated operating carriers, and designated hub airports, using the industry-standard IATA codes employed throughout this paper.

Parent	**Airlines**	**Hubs**
BA	British Airways, Iberia, Aer Lingus, Vueling	LHR, MAD, DUB, BCN, FCO
AF	Air France, KLM, Transavia	CDG, AMS
LH	Lufthansa, Austria Airline Swiss, Brussels Airline, Eurowings	FRA, MUC, ZRH VIE, BRU
SK	Scandinavian Airlines	CPH, ARN, OSL
AY	Finnair	HEL
A3	Aegean Airlines	ATH
LO	LOT Polish Airlines	WAW

Note: The parent airline's code is assigned to the largest airline within the parent company. Hub airports represent the central hubs for all airlines under the same parent company.

Table 1: List of Parent-Constituted Airlines-Hubs

In this study, we define airlines at the parent company level rather than at the level of individual operating carriers. For instance, 'BA' represents the International Airlines Group (IAG), which encompasses not only British Airways but also other major carriers such as Iberia (Spain's flag carrier), Aer Lingus (Ireland's national airline), and the low-cost subsidiary Vueling. Consequently, the hub network attributed to 'BA' includes the primary hubs of these carriers, such as Madrid (MAD) and Dublin (DUB), in addition to London Heathrow (LHR).

This parent-level aggregation is motivated by two factors. First, it reduces the number of distinct airline entities from over 35 to a computationally tractable 14, simplifying both estimation and empirical identification. Second, it reflects the economic reality that carriers within the same parent group typically coordinate operations through code-sharing and complementary routes, behaving more like a unified strategic actor than as separate competitors.

As previously discussed, Europe is home to many of the world's most congested airports. London Heathrow, for instance, has operated at over 98% of its runway capacity for more than two decades, with expansion severely limited by a combination of local opposition, land constraints, environmental concerns, and the political influence of incumbent airlines seeking to protect their market power.[13]

To incorporate this critical operational reality, we classify all major hub airports used by FSCs—as defined in Table 1—as slot-controlled. This classification is empirically grounded, as all 18 of these airports are designated as Level 3 congested under the IATA system. We then operationalise this constraint in both our estimation and counterfactual exercises by limiting the ability of airlines to alter their existing route networks or increase frequencies, as detailed in the subsequent sections.

Aircraft Choice and Utilisation While aircraft selection is a key strategic lever for airlines, we do not explicitly model this dimension of choice. Our focus on intra-European routes justifies this simplification, as this market is overwhelmingly dominated by two highly comparable aircraft families: the Boeing 737 and the Airbus A320. These single-aisle aircraft exhibit similar characteristics in seating capacity, fuel efficiency, and emissions per kilometre. The use of larger, twin-aisle aircraft is rare and economically unviable for these short-haul operations,[14] while LCCs further homogenise their fleets by operating a single aircraft type to maximise efficiency. Given this relative uniformity, we abstract from aircraft heterogeneity.[15]

A more salient constraint in the European context is aircraft utilisation, which directly limits an airline's ability to adjust flight frequencies. Two features of the

[13]For further discussion, see European Union.

[14]An exception is the occasional use of A350s on routes such as London–Helsinki.

[15]Although variants within aircraft families—such as the A319, A320, and A321—exist on intra-European routes, they usually share the same engine type and have similar fuel consumption. Newer, more fuel-efficient models like the A320NEO and B737MAX were still relatively rare in 2019.

market in our 2019 sample period are critical. First, European carriers operate with high levels of fleet efficiency, meaning most aircraft were already near full operational capacity, leaving little slack to increase total network frequency.[16] Second, the continent's airlines were not undergoing significant fleet expansion during this period. Given the long lead times for aircraft orders—typically three to five years—rapid capacity growth was not feasible, and no large-scale orders were pending delivery.

These realities of a fixed fleet size and high utilisation rates motivate a key feature of our modelling framework: we treat network adjustments not as unconstrained entry and exit decisions, but as a problem of aircraft redeployment. Consequently, if an airline exits a route, we assume the freed aircraft is reallocated to another profitable opportunity within its feasible network.

We formally define an airline's set of feasible redeployment routes based on two conditions: (1) the airline must already maintain a presence in both endpoint cities, and (2) the route must already be served by at least one other airline in the market. The first condition reflects the operational incentive to expand within an existing network, leveraging established infrastructure and personnel. This constraint also captures a key structural difference between business models: FSCs, with their hub-and-spoke networks, serve fewer cities and tend to add spokes from their hubs, whereas LCCs' more diffuse point-to-point presence gives them a combinatorially larger set of feasible new routes. The second condition ensures that we only consider routes with demonstrated underlying demand, avoiding economically implausible connections between remote or low-traffic locations. Together, these constraints define a realistic choice set for aircraft redeployment that shapes the simulated network adjustments and resulting profit outcomes.

[16] See the report from Eurocontrol.

Summary Statistics Table 2 presents the summary statistics for our 2019 sample, which reveals a highly concentrated market structure. The industry is dominated by 14 parent airline groups, with the six largest—the three primary full-service carriers (IAG, Air France-KLM, and Lufthansa Group) and the three primary low-cost carriers (Ryanair, EasyJet, and Wizz Air)—collectively accounting for 87% of all intra-European passenger traffic. Competition is widespread, as only 4,039 of the 11,292 total itineraries operate on monopoly routes, indicating that most markets feature multi-product competition. The central role of hubs is also evident, with nearly 22% of all itineraries involving a flight to or from a designated hub airport. A typical route in our sample has an average fare of approximately $86, a frequency of roughly one flight per day, and a travel distance of about 1,400 kilometres (a flight duration of just under two hours). In total, the routes in our 2019 sample served over 350 million passengers.

Table 3 further quantifies the strategic divergence in network structure by providing key metrics for hub cities. While LCCs do not operate formal hubs in the traditional sense, we identify the two most connected cities in each LCC's network for comparative purposes. Panels (a) and (b) reveal that FSCs maintain far greater connectivity from their hubs and operate at significantly higher frequencies, particularly on dense business routes. For instance, Lufthansa Group (LH) operates approximately 40 daily flights between its hubs in Munich and Düsseldorf, while IAG (BA) operates 35 between Madrid and Barcelona.

This contrast is starkly illustrated in Panels (c) and (d), which measure network concentration. Nearly 70% of Air France–KLM's entire route network touches its hubs in Paris or Amsterdam, a clear empirical signature of a Hub-and-Spoke model. In contrast, LCCs exhibit much lower concentration levels, with their routes more evenly distributed across a wide range of cities, reflecting their decentralised Point-

(a) Sizes:		
Number of firms	14	
Number of products (itineraries)	11292	
Number of markets	7025	
Number of the hub itineraries	2432	
Number of monopoly itineraries	4039	
Number of city pairs	2003	
Number of passengers (1 million)	354.0	
Number of quarters	4	
(b) Market shares by airline (excl. outsude option)		
BA	0.16	
AF	0.09	
LH	0.12	
FR (LCC)	0.25	
U2 (LCC)	0.21	
W6 (LCC)	0.04	
Other	0.13	
(c) Demand and cost variables	**Mean**	**St.Dev**
Fare (100 USD)	0.86	0.57
Frequency (Daily)	0.95	1.74
Distance (1,000 km)	1.38	0.73
Market Size (1 million)	2.82	2.01
Product Shares	1.48e-02	2.40e-02
(d) Market level statistics	**Mean**	**St.Dev**
Number of products	2.07	1.11
Average Fare in Similar Markets	0.86	0.15

Note: This table presents key summary statistics for the sample drawn from the four quarters of 2019. Hub itineraries are defined as those where at least one of the origin or destination airports is classified as a hub airport. Market shares in panel (b) exclude outside options, such as individuals choosing not to travel or opting for alternative modes of transportation. Fares are calculated as the average fare across all tickets for a specific itinerary. Comparable markets used to compute average fares are defined as those with similar distances (±10%).

Table 2: Summary Statistics

to-Point strategy.[17]

[17]Wizz Air shows a relatively high concentration rate, primarily because it operated a much smaller network in 2019 compared to the other airlines. This is also reflected in its smaller market share. Since then, Wizz Air has expanded significantly, and its hub concentration is now closer to that of Ryanair and EasyJet.

Airlines	Top Hub	Hub Index	Freq	Second Hub	Hub Index	Freq
(a) Full service:						
BA	Madrid	60	2.3	London	56	2.6
AF	Amsterdam	73	2.1	Paris	52	1.9
LH	Frankfurt	66	3.0	Munich	64	4.0
(b) Low Cost:						
FR	Dublin	61	0.9	London	56	1.2
U2	London	61	1.8	Geneva	51	0.7
W6	Budapest	37	0.4	Bucharest	27	0.4
Airlines	Hub1	Concentration		Hub2	Concentration	
(c) Full service:						
BA	Madrid	14%		London	25%	
AF	Amsterdam	36%		Paris	37%	
LH	Frankfurt	19%		Munich	18%	
(d) Low Cost:						
FR	Dublin	7%		London	7%	
U2	London	12%		Geneva	9%	
W6	Budapest	18%		Bucharest	14%	

Note: This table presents key summary statistics for each airline's hub cities and their characteristics. The Hub Index represents the total number of cities served by the hub, indicating its level of connectivity. Freq refers to the average frequency of all itineraries to/from a specific hub. Concentration refers to the proportion of itineraries to/from this hub city relative to the total number of itineraries.

Table 3: Summary Statistics of Hub Cities

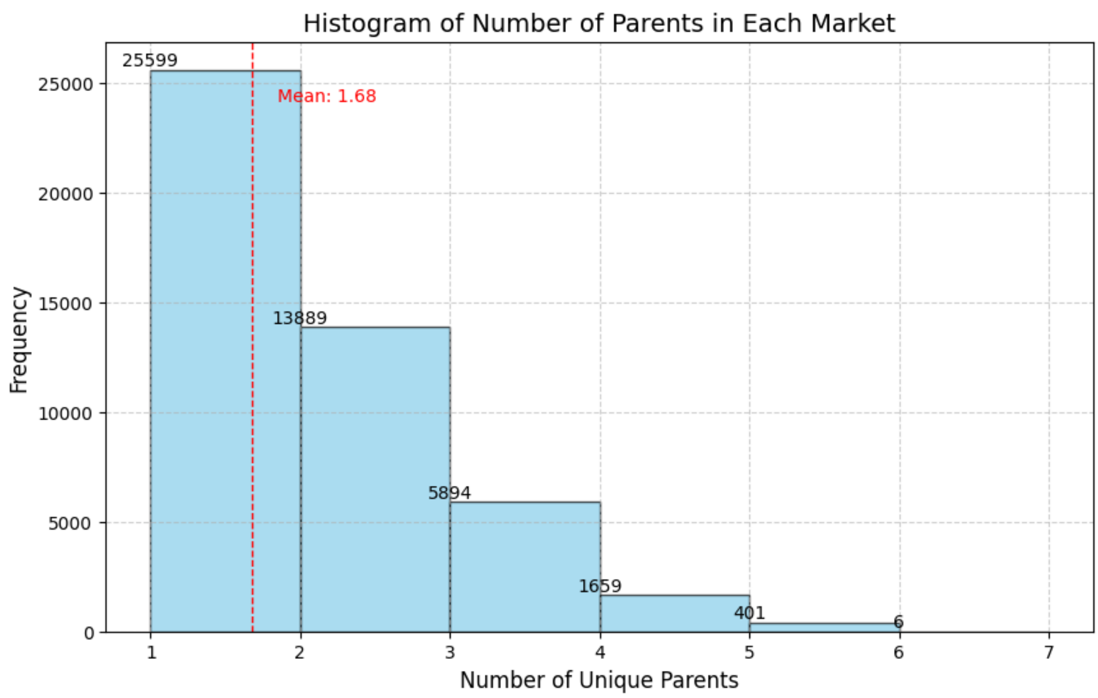

Figure 4: Number of Airlines in Each Market

Several key features of the market's competitive structure are illustrated in Figures 4 through 8. The overall level of competitive intensity is high; as shown in Figure 4, nearly half (46%) of all markets are served by more than one airline group. This has a clear disciplinary effect on pricing (Figure 5), where average fares decrease as the number of competitors increases, consistent with standard oligopoly theory. Fares in monopoly markets are not only higher but also exhibit greater dispersion. The data also reveal stark strategic differences between carrier types. FSCs operate, on average, nearly three times as many routes involving a hub as LCCs (Figure 6). This hub-centric strategy is also reflected in their dynamic network adjustments (Figure 7): FSCs frequently alter their portfolio of hub-related routes in response to seasonal demand, while LCCs rarely do, a pattern likely driven by slot constraints at major airports. Finally, Figure 8 highlights how these different strategies translate into market share. While LCCs like Ryanair (FR) and EasyJet (U2) lead in terms of passenger volume, FSCs such as IAG (BA) and Lufthansa Group (LH) dominate when measured by revenue and frequency, reflecting their focus on premium services and dense schedules.

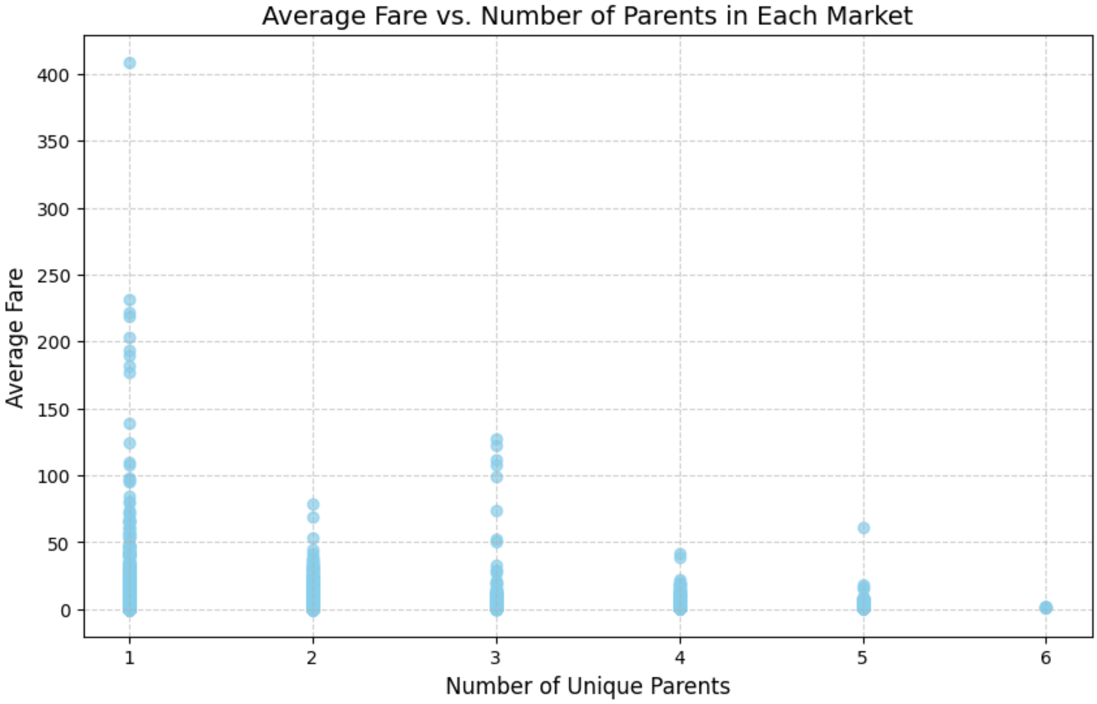

Figure 5: Average Fare vs. Number of Airlines in Each Market

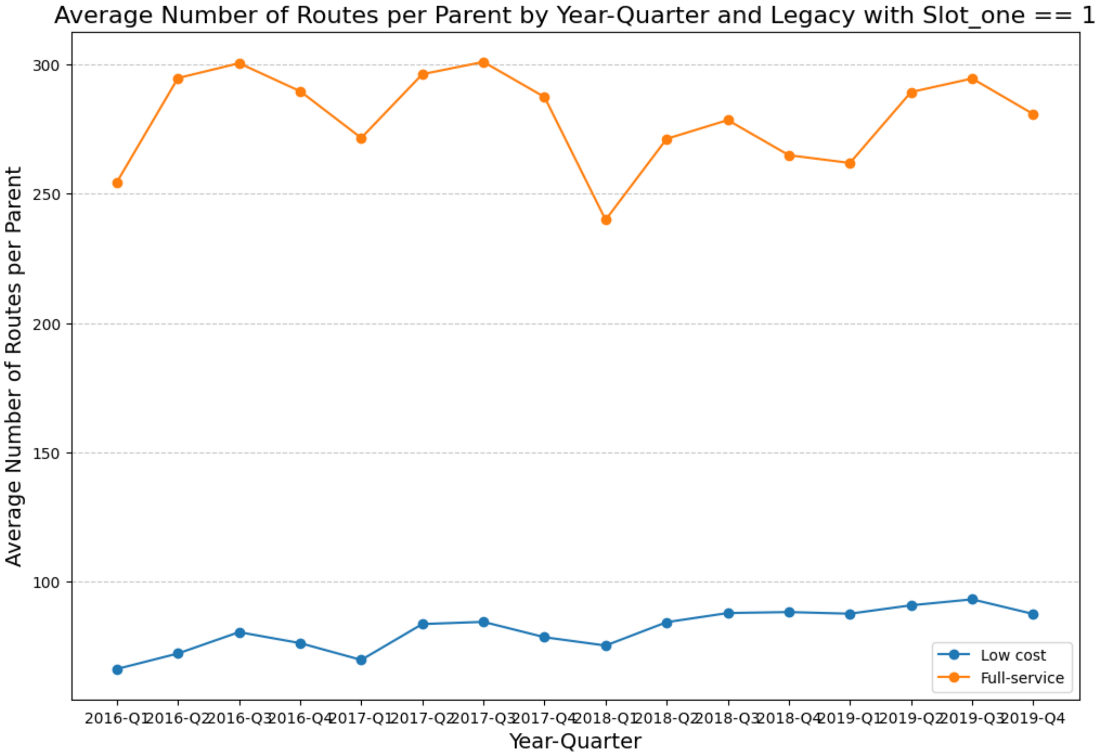

Figure 6: Number of Routes Linking at Least One Hub Airport

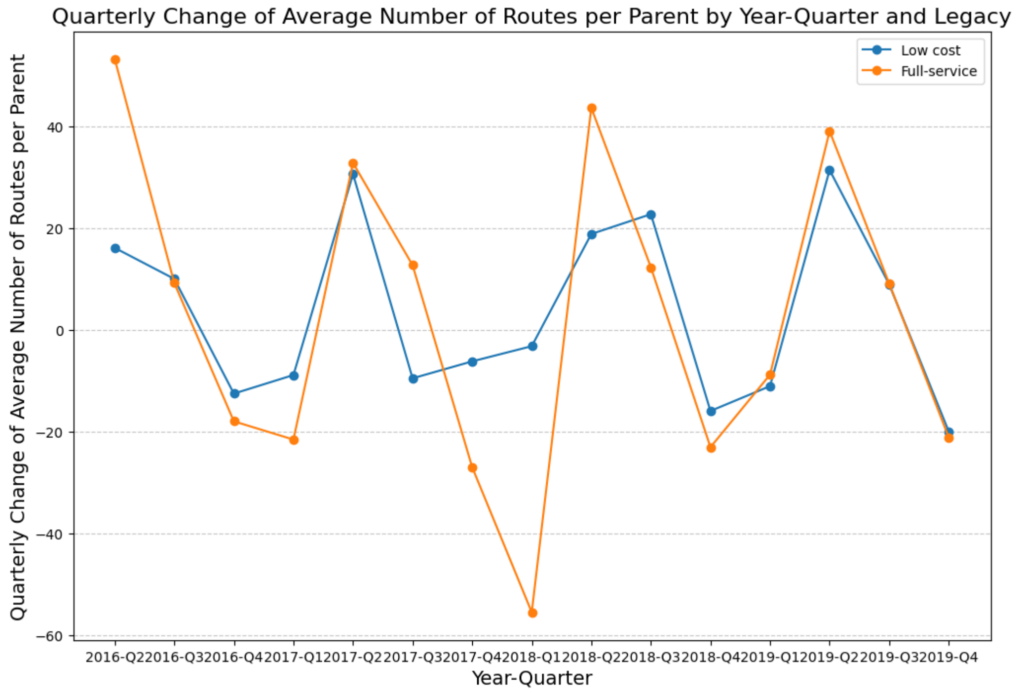

Quarterly Numbers of Routes Change

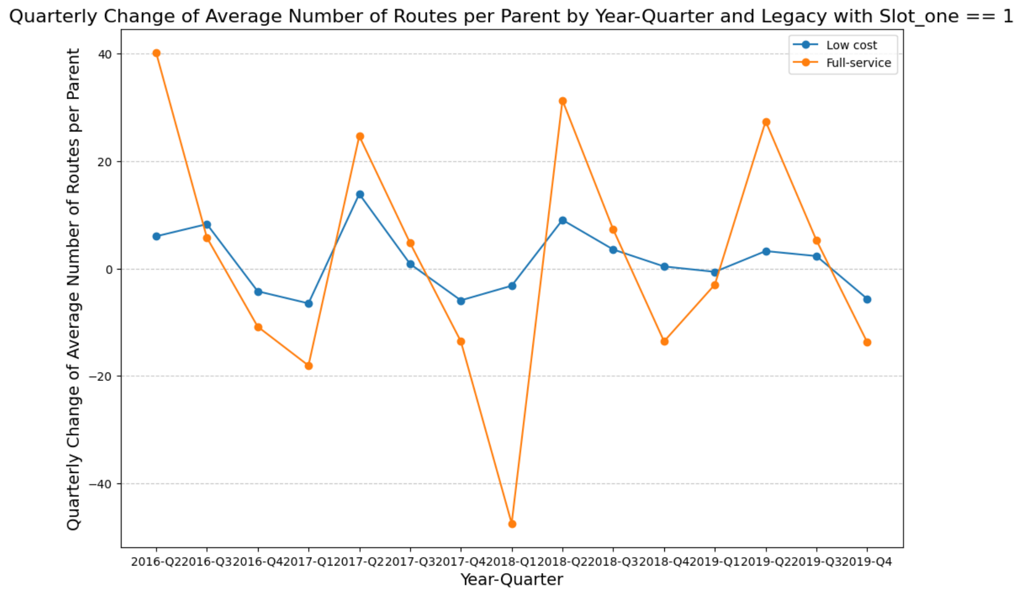

Quarterly Numbers of Routes Change for Hub Airports

Figure 7: Comparing Quarterly Numbers of Routes Change

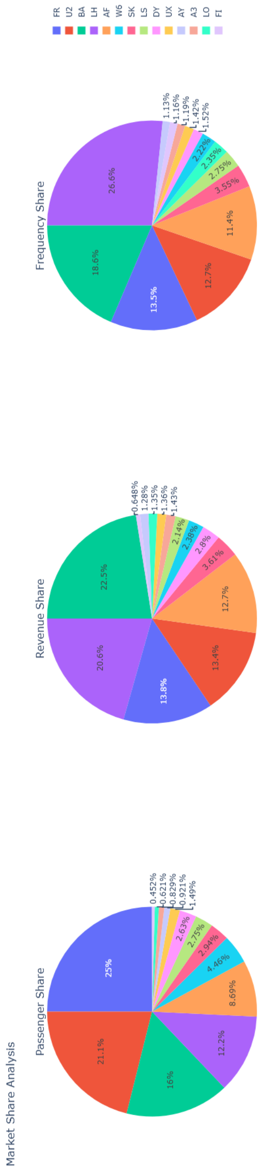

Figure 8: Market Share Analysis

1.3 Model

This section introduces a static model for airlines' entry, flight frequency, and pricing decisions similar to Yuan and Barwick [2024] and Bontemps et al. [2023].

The model has two stages. In the first stage, airlines simultaneously decide routes to enter, thereby shaping the overall flight network. If an airline enters, it also determines the flight frequency. In the second stage, airlines compete on prices to attract customers.

The total cost of operating a route in the airline industry consists of three main components: airport costs, flight costs, and passenger-related costs (Ciliberto et al. [2021]). *Airport costs* include fees for landing, parking, terminal access, and security, as well as ground handling expenses at both origin and destination airports. *Flight costs* refer to per-flight expenditures such as fuel, pilot salaries, and wages for flight attendants. *Passenger costs* cover per-passenger items like ticketing, in-flight catering, insurance, and liability charges. Beyond these recurring operational expenses, airlines also incur a one-time *entry cost* when launching a new route. This includes expenses related to hiring personnel, coordinating logistics, and marketing the new service.

As previously discussed, Europe is home to some of the most congested airports in the world—an issue that carries important implications for both consumer behaviour and airline operations. On the demand side, passengers often favour flights through major hub airports, influencing how airlines design their networks. On the cost side, congestion drives up landing fees and other charges, raising both per-passenger and per-flight operating costs. In addition, regulatory constraints—most notably slot control policies—limit airlines' ability to freely enter or exit congested markets. These constraints do not affect all carriers equally: their impact varies by business model, shaping the strategic responses of full-service and low-cost airlines in distinct ways.

All these factors play a crucial role in understanding airline competition and route entry decisions. We will explore them in greater detail in this section and further elaborate on them in the next section, where we specify the estimation and identification approach.

1.3.1 Model Setting: Demand and Price Competition

In the second stage, the route network and frequencies are fixed. There are $m \in M$ markets in total, each defined by a city-pair, which specifies the origin and destination cities for a given quarter of the year. We adopt the city-pair definition rather than the airport-pair for the reasons discussed earlier.

Each product $j \in J_m$ in market m represents an airline itinerary in a given quarter of 2019, specifying an airline $f \in F$ operating from airport $o \in O$ in city $a \in A$ to airport $d \in D$ in city $b \in B$. We consider itineraries in a one-directional manner, meaning that the origin and destination airports are not distinguished. This approach is commonly adopted in the literature (Ciliberto et al. [2021], Bontemps et al. [2023]). For instance, in the London–Paris market, examples of products include British Airways flying from Heathrow to Charles de Gaulle, Air France from Heathrow to Orly, and EasyJet from Gatwick to Charles de Gaulle. Airlines are defined at the parent company level.

As our study focuses on intra-European[18] flights, we only consider direct flights. Short-haul indirect flights are uncommon in Europe for both full-service and low-cost carriers, unlike in the United States, where previous studies primarily examine network spillover effects from indirect flights across the route network. In each market, airlines face product-specific demand, incur a per-passenger marginal cost, and set prices to maximise profits under complete information and Bertrand competition. The following sections detail the demand and supply settings.

[18]Flights to and from Armenia, Azerbaijan, Georgia, Belarus, Moldova, Serbia, Ukraine, Russia, and Turkey are excluded due to their non-compliance with current European aviation policy, despite their geographical location in Europe.

Demand: Consumers make a discrete choice from all available products in a given market to maximise their indirect utility. We adopt a nested logit model, following Bresnahan [1987], Berry [1994], and Berry and Jia [2010]. In this model, all air travel products form one nest, while all other options—including not travelling or using alternative transport modes such as road, high-speed rail (HSR), and ferry—form another.

The indirect utility received by individual passenger i from choosing product j in market m for air travel (inside goods) is given by:

$$U_{ijm} = \beta \mathbf{X_{jm}} - \alpha p_{jm} + \xi_{jm} + v_{im}(\lambda) + \lambda \epsilon_{ijm} \tag{1}$$

In (1), $\mathbf{X_{jm}}$ is a vector of product characteristics and p_{jm} is a scalar of the average fare for an itinerary. β and α are vector and scalar respectively. ξ_{jm} is the structural unobserved (to researchers) error and it can potentially be correlated with prices. λ is the nesting parameter which defines the extent of the correlation among consumer's idiosyncratic shock within a nest. Finally, ϵ_{ijm} should follow a type I extreme value distribution.

We include in $\mathbf{X_{jm}}$ a range of product characteristics. First, we incorporate the logarithm of flight frequency for each itinerary. Higher frequency offers consumers more travel options and greater flexibility, allowing them to adjust their plans in response to unexpected events.

Second, we include fixed effects for airlines and cities, as is standard in the literature. Given the prevalence of multi-airport systems in European metropolitan areas, passengers may exhibit distinct preferences for different airports within the same city. To capture this, we also include fixed effects for major hub airports.

Third, we include an airline $\times$ airport interaction term for selected major hub airports and their respective national carriers. This fixed effect is important because major European hubs serve not only intra-European passengers but also a substantial volume of intercontinental transfer passengers who fall outside the scope of our analysis[19]. We include this additional fixed effect to control for the influence of intercontinental layover traffic on observed demand patterns at key hub airports.

Finally, we include seasonal fixed effects to capture the variation between the summer peak season and the winter off-peak season.

The utility for the outside good is normalised as: $U_{0m} = v_{im}(\lambda) + \lambda\epsilon_{i0m}$. Then, the market share for product j in market m at nest g is:

$$s_{jm|g}(\mathbf{X_m}, \mathbf{p_m}, \xi_\mathbf{m}, \theta_\mathbf{d}) = \frac{exp((\beta\mathbf{X_{jm}} - \alpha p_{jm} + \xi_{jm})/\lambda_g)}{\sum_{k\in g} exp((\beta\mathbf{X_{km}} - \alpha p_{km} + \xi_{km})/\lambda_g)} \tag{2}$$

And the probability of choosing a nest g is:

$$s_{mg}(\mathbf{X_m}, \mathbf{p_m}, \xi_\mathbf{m}, \theta_\mathbf{d}) = \frac{(\sum_{k\in g} exp((\beta\mathbf{X_{km}} - \alpha p_{km} + \xi_{km})/\lambda_g))^{\lambda_g}}{1 + (\sum_{k\in g} exp((\beta\mathbf{X_{km}} - \alpha p_{km} + \xi_{km})/\lambda_g))^{\lambda_g}} \tag{3}$$

In equation (2) and (3), $\mathbf{X_m} := (\mathbf{X_{jm}} : j \in J_m)$, $\mathbf{p_m} := (p_{jm} : j \in J_m)$, $\xi_\mathbf{m} := (\xi_{jm} : j \in J_m)$, $\theta_\mathbf{d} := (\beta, \alpha)$. J_m is the product set in market m and λ_g is the nesting parameters for nest g. Given that we only have two nests, then nest g just includes all products excluding the outside good. The unconditional choice probability for product j in market m is:

$$s_{jm}(\mathbf{X_m}, \mathbf{p_m}, \xi_\mathbf{m}, \theta_\mathbf{d}) = s_{jm|g} \cdot s_{mg} \tag{4}$$

[19]Specifically, international transfer passengers travelling on a single itinerary with a short layover are not captured in our dataset and are excluded from the demand estimation. However, some travellers choose to extend their stopover to visit the hub city itself. In such cases, the intra-European leg of the journey is included in our demand sample.

The market size in a city-pair is the geometric mean of the two end points (Berry and Jia [2010]):

$$MS_m = exp(\frac{\ln(pop_a) + \ln(pop_b)}{2}) \tag{5}$$

where pop_a and pop_b are the population of two endpoints. The advantage of using the geometric means instead of the absolute mean is that it balance large and small cities in a more moderate way.

Supply Airlines simultaneously set the prices for all their offered products in each market m to maximise the variable profits, under complete information:

$$\pi_f = \sum_{m \in M} \sum_{j \in J_{fm}} (p_{jm} - MC_{jm}) \cdot s_{jm} \cdot MS_m \tag{6}$$

where J_{fm} is the *observed* set of products offered by firm f in market m and MC_{jm} is product j's *per-passenger* marginal cost. Without any constraint, airline f's equilibrium pricing choices are determined by the first order condition of maximising the variable profits for each product $j \in J_{fm}$:

$$p_m - c_m = \Delta_m^{-1} s_m \tag{7}$$

Equation (7) is the vector-matrix form of first order conditions in a specific market m where p_m, c_m, and s_m representing three $J_m \times 1$ vectors of prices, marginal costs, and market shares respectively. Δ_m is a $J_m \times J_m$ intra-firm (negative, transposed) demand derivatives:

$$\Delta_m = \mathcal{H}_m \odot \frac{\partial s'_m}{\partial p_m} \tag{8}$$

Where $\mathcal{H}_m$ is the market level ownership matrix where $\mathcal{H}_{mjk}$ equals one if product j and k belong to the same airline. We further decompose the marginal cost as a function of observed and unobserved characteristics:

$$MC_{jm} = \theta_s \mathbf{W_{jm}} + \omega_{jm} \tag{9}$$

where $\mathbf{W_{jm}}$ is a vector of marginal cost shifters that are observed by the researcher and ω_{jm} represents the marginal cost shifters that are unobersved by the researcher.

As on the demand side, we include various product characteristics in $\mathbf{W_{jm}}$. First, we incorporate distance and its squared term to account for the impact of route length on crew labour costs for serving an additional passenger. Second, we include the log of flight frequency, as itineraries with higher frequency typically have lower load factors, increasing the marginal cost of accommodating an extra passenger. Third, consistent with the demand-side specification, we include airline, city, and airport dummies to capture fixed effects. Finally, we incorporate seasonal fixed effects to account for cost variations between peak and off-peak travel periods.

1.3.2 Model Setting: Entry and Frequency

In the first stage, each airline *simultaneously* determines its route network and chooses the flight frequency for every active route in a given quarter. When an airline operates in a market, it incurs a fixed cost that includes operational expenses and entry-related adjustment costs. Let A_{fj} denote an indicator for whether airline f operates product $j \in G_{ft}$, where G_{ft} is the set of all feasible routes for airline f in quarter t. A route is defined as a unique tuple $(City_x, City_y, Airport_x, Airport_y)$, where $City_x$ and $City_y$ are the endpoint cities and $Airport_x$ and $Airport_y$ the corresponding airports. Routes are treated as unidirectional since directional attributes are nearly identical in the data: airlines typically return aircraft to their origin after a flight, particularly full-service carriers. Thus, London–Paris–LHR–CDG and

Paris–London–CDG–LHR are considered the same route. Finally, $Freq_{fj}$ denotes the frequency of product j operated by airline f.

Entry and exit are defined at the *product level* (Quarter–Airline–City Pair–Airport Pair), not at the *market level* (Quarter–Airline–City Pair). This distinction reflects the multi-airport structure of many European metropolitan areas. As noted earlier, full-service and low-cost carriers often operate from different airports within the same city. Modelling entry and exit only at the market level would obscure this distinction, which is already incorporated in the demand estimation via airport-specific fixed effects.

Airlines incur fixed costs when actively operating a route. The total fixed cost of airline f in quarter t is the sum of fixed costs across all active routes:

$$FC_{ft} = \sum_{j \in G_{ft}} FC_j \cdot A_{fj}. \tag{10}$$

We further decompose FC_j as:

$$FC_j(\mathbf{Z_j}, \theta_{fc}; \kappa_j) = \theta_{fc}\mathbf{Z_j} + \kappa_j, \tag{11}$$

where $\mathbf{Z_j}$ denotes route characteristics affecting fixed costs. First, we include a 'frequency $\times$ distance' term to capture operating costs: longer and more frequent routes incur higher fuel, maintenance, and labour costs. Second, we include market size in $\mathbf{Z_j}$, since entering larger markets typically requires greater investment. Unlike much of the literature, we do not include a constant term in $\mathbf{Z_j}$ to represent the pure entry effect, because our difference-based moment inequalities cannot jointly identify both the lower and upper bounds of such an intercept.

Nevertheless, the pure entry effect remains important. Rather than estimate a constant directly, we model the fixed cost shock κ_j as normally distributed:

$$\kappa_j \sim \mathcal{N}(\mu_g, \sigma_g^2). \tag{12}$$

Here, shocks vary across groups g, each with its own mean and variance. The mean can be interpreted as the pure entry effect. We define groups as follows: the three largest full-service airlines (BA, AF, LH), the three largest low-cost carriers (FR, U2, W6), and a pooled group of smaller airlines. For the three full-service carriers and the pooled group, we further distinguish between routes involving at least one of the carrier's hubs and routes without hubs, reflecting the special role of hubs in generating revenue[20]. In total, this yields 11 groups, each with a distinct mean and variance. These additional distributional assumptions on fixed cost shocks make our model more flexible than approaches in the existing literature, which typically do not impose such structure. We discuss these distributional assumptions in Section 4, where we compare them with prior work and explain how we estimate the group-specific means and variances.

1.3.3 Information Structure and Equilibrium Concept

Information Structure: In the first stage, airlines simultaneously decide whether to enter a route and, if so, determine the corresponding frequency. Each airline has perfect information on product characteristics $\mathbf{Z_j}$, $\mathbf{X_{jm}}$, and $\mathbf{W_{jm}}$, which correspond to fixed costs, demand, and marginal costs, respectively. They also know the parameter values θ_{fc}, θ_d, and θ_s. Consequently, when making entry and frequency decisions in the first stage, airlines are aware of the linear part of values of fixed costs, marginal costs, and demand.

[20]This may include loyal business travellers, international transfer passengers, and economies of scale in labour and maintenance.

The model includes one set of fixed cost shocks κ_j in the first stage, and two sets of shocks in the second stage: demand shocks ξ_{jm} and marginal cost shocks ω_{jm}. In the second stage, it is standard to assume that airlines observe the realised values of ξ_{jm} and ω_{jm} and then set prices optimally. However, we assume that airlines only know the distributions of demand and marginal cost shocks, not their realisations. This timing assumption follows prior work, including Aguirregabiria and Ho [2012], Sweeting [2013], Eizenberg [2014], and Yuan and Barwick [2024], where firms learn the true values of demand and marginal cost shocks only after committing to entry.

In addition, airlines are assumed to have complete information on the set of fixed cost shocks κ_j for all products j in a given quarter. This assumption is consistent with the strand of literature using moment inequalities to estimate fixed cost parameters[21].

Equilibrium Concept: The equilibrium of this two-stage game is a Subgame Perfect Nash Equilibrium (SPNE). We now describe the optimisation problem for each airline in quarter t:

- *First stage:* Given the optimal route networks of competitors $\mathbf{A}^*_{-\mathbf{ft}} := (A_{-fj} : j \in G_{-ft})$, the frequencies of competitors' products $\mathbf{Freq}^*_{-\mathbf{ft}} := (Freq_{-fj} : j \in G_{-ft})$, the complete set of fixed cost shocks $\kappa := (\kappa_j : j \in G_{ft}, \forall f)$, and the optimal second-stage pricing profile $\mathbf{p}^*_{\mathbf{t}} := (p^*_{jm} : j \in J_{ft}, m \in M, \forall f)$, airline f chooses its route network $\mathbf{A}^*_{\mathbf{ft}}$ and frequencies $\mathbf{Freq}^*_{\mathbf{ft}}$ to maximise expected profit:

[21] Sabal [2025] assumes that fixed cost shocks are private information across firms, which implies conditional independence. However, his context is the automobile industry, where the fixed costs of developing new models are likely private and independent across products. By contrast, in the airline industry, fixed cost shocks are widely recognised as common knowledge among competitors and are typically correlated rather than conditionally independent.

$$(\mathbf{A^*_{ft}}, \mathbf{Freq^*_{ft}}) = \underset{\mathbf{A_{ft}}, \mathbf{Freq_{ft}}}{\arg\max} \Big\{ \mathbf{E}_{\xi,\omega} \left[\pi_f(\mathbf{A_{ft}}, \mathbf{Freq_{ft}}, \mathbf{p^*_{ft}}(\xi,\omega); \theta_d, \theta_s \mid \mathbf{A^*_{-ft}}, \mathbf{Freq^*_{-ft}}, \mathbf{p^*_{-ft}}(\xi,\omega)) \right]$$
$$- FC_{ft}(\mathbf{A_{ft}}, \mathbf{Freq_{ft}}, \kappa; \theta_{fc}) \Big\}. \tag{13}$$

In equation (13), we omit the arguments $\mathbf{Freq_{ft}}$ and $\mathbf{Freq^*_{-ft}}$ in $\mathbf{p^*_t}$ to simplify notation. The choices $\mathbf{A_{ft}}$ and $\mathbf{Freq_{ft}}$ also determine the product attributes $\mathbf{Z_j}$, which enter the fixed cost function. In equilibrium, condition (13) must hold for all airlines. Expectations are taken over the unobserved shocks ξ and ω, which are not realised until after entry. In practice, we independently draw 36 pairs of (ξ, ω) for each Quarter–Airline–Route from the empirical distribution of recovered shocks on observed routes. These draws are stored and reused in counterfactual simulations to ensure consistency between estimation and counterfactual analysis.

- *Second stage:* Airlines compete on fares in each market following Bertrand-style competition. In market m, given the active airlines' frequencies from the first stage $\mathbf{Freq^*_m} := (Freq_{fm} : f \in F)$ and competitors' pricing strategies $\mathbf{p^*_{-fm}}(\xi_\mathbf{m}, \omega_\mathbf{m})$, where $\xi_\mathbf{m}$ and $\omega_\mathbf{m}$ are the realised demand and cost shocks for all products in m, airline f's optimal pricing strategy solves:

$$\mathbf{p^*_{fm}}(\xi_\mathbf{m}, \omega_\mathbf{m}) = \underset{\mathbf{p_{fm}}(\xi_\mathbf{m}, \omega_\mathbf{m})}{\arg\max} \; \pi_{fm}(\mathbf{Freq^*_m}, \mathbf{p_{fm}}(\xi_\mathbf{m}, \omega_\mathbf{m}); \theta_d, \theta_s \mid \mathbf{p^*_{-fm}}(\xi_\mathbf{m}, \omega_\mathbf{m})). \tag{14}$$

Airlines may operate multiple products in the same market due to multi-airport structures. In equation (14), π_{fm} denotes total profits from all products of airline f in market m. Unlike in (13), no expectation is taken since shocks are realised in the second stage. Equation (14) must hold for all airlines in all markets in equilibrium.

The existence and uniqueness of equilibrium in the second-stage pricing game are established by Nocke and Schutz [2018] for multi-product nested logit models. By contrast, equilibrium in the first-stage entry game is not guaranteed to exist or be unique. As noted by Bontemps et al. [2023] and Yuan and Barwick [2024], multiple equilibria $\mathbf{A}^*$ may arise due to spillovers in demand, marginal costs, and fixed costs. Following standard practice, we assume the existence of at least one equilibrium. Crucially, our moment inequality approach to estimating fixed costs does not require solving for the industry equilibrium, allowing us to sidestep the issue of multiplicity. We revisit the equilibrium concept when discussing counterfactual simulations in later sections.

1.4 Identification and Estimation Strategies

Endogenous Entry: Route network choices and flight frequencies directly affect the indirect utility of itinerary selection in the second stage. This relationship creates a fundamental challenge in the product entry literature: firms *endogenously* select markets based on anticipated private gains. To address this issue, Ciliberto et al. [2021], following Ciliberto and Tamer [2009], simultaneously estimate entry and pricing decisions whilst allowing for correlation between cost and demand shocks. These authors assume a joint normal distribution for all shocks—fixed cost, marginal cost, and demand—and estimate the joint variance-covariance matrix using a probit approach. Alternatively, Aguirregabiria and Ho [2012] and Yuan and Barwick [2024] assume firms realise demand and marginal cost shocks only after market entry. Rather than imposing specific distributional assumptions on fixed cost shocks, they exploit revealed preference arguments to construct moment inequalities for *set* estimation of fixed cost parameters.[22] In practice, instruments are introduced to eliminate fixed cost shocks when constructing moment inequalities.

Our estimation approach synthesises both methodologies. We employ IV-moment inequalities to set-estimate the linear parameters θ_{fc} for frequency $\times$ distance and market size. However, given that our deviations involve redeploying aircraft to existing routes, we can only estimate terms that *vary* across alternative routes. The constant term cannot be estimated as it cancels on both sides of the inequality. We therefore impose distributional assumptions on fixed cost shocks (described in Section 3) and estimate their mean and variance directly from the likelihood of revealed preference inequalities. Notably, airlines maintain complete information on all fixed cost shocks, preserving the arguments for handling multiple equilibria in moment inequality estimation. Whilst our estimation approach to fixed cost shock's distribution resembles Ciliberto et al. [2021], our estimation focuses on fitting residuals to the most appropriate distribution rather than solving endogenous entry problems under incomplete information.

There are several reasons to impose a distributional assumption on the fixed cost shocks. First, because $\mathbf{Z_j}$ contains no constant term, the mean of the shocks is not expected to be zero. This mean has a clear economic interpretation: it captures the average pure entry cost together with any unexplained components of airline net profits. Second, we need to draw fixed cost shocks from the estimated distribution to ensure that all revealed-preference inequalities hold, which is particularly important for counterfactual analysis. Ignoring the fixed cost shocks, as in some studies, would generate a large number of Quarter–Airline–Route deviations under the estimated θ_{fc} from moment inequalities. This would further distort the coun-

[22] See also Pakes et al. [2015], Ho and Pakes [2014], and Houde et al. [2023].

terfactual results because the benchmark network would deviate substantially from the observed network. We provide further details in Section 6.

Choice Set: Our approach differs fundamentally from the existing entry literature. In most studies, airlines enter a route if the expected profit is positive and exit if it is negative. This reduces the entry problem to *independent* binary choices across all available routes, allowing standard logit or probit models to be applied. We do not believe this binary setting fits the European airline market. As discussed in earlier sections, the key resource constraints are aircraft availability and slot capacity at congested airports. This means airlines cannot freely choose flight frequencies because they may lack sufficient aircraft or secured slots.

The existing literature either abstracts entirely from frequency decisions or treats frequency decision as an unconstrained problem to maximise expected profit, both of which are unrealistic in the European context. Instead, for every observed Quarter–Airline–Route combination, we assume the choice set consists of redeploying all aircraft used on that route (i.e., the total frequency) to any route in the airline's feasible set[23], plus an exit option. Hence, rather than making independent binary entry choices across all routes, the airline selects the *best* available route conditional on a fixed number of aircraft.

This setup creates two key differences when forming the estimation inequalities. First, even if the expected profit of an alternative route is positive, the airline will not enter if a more profitable redeployment exists. Second, we require only that the expected profit of the observed route be positive, since it represents the best option under revealed preferences. We do not require alternative routes in the choice set to have negative profits, as long as their profits do not exceed those of the observed

[23]The feasible set requires that the airline already serves both endpoint cities and that the market (city pair) is served by at least one other airline.

route. In the counterfactual analysis, an airline exits only if all feasible routes yield negative profits.

Section 4.1 introduces the instrument-based construction of moment inequalities, Section 4.2 sets the criteria for feasible alternative route networks, Section 4.3 presents the estimation strategy for the linear parameters θ_{fc}, and Section 4.4 details the maximum-likelihood estimation of the fixed cost shock distribution.

1.4.1 Construction of Moment Inequalities

Let $\Pi_f(\mathbf{A_f^*}, \mathbf{Freq_f^*}, \kappa; \theta_{fc}|\mathbf{A_{-f}^*}, \mathbf{Freq_{-f}^*})$ and $\Pi_f(\mathbf{A_f^a}, \mathbf{Freq_f^a}, \kappa; \theta_{fc}|\mathbf{A_{-f}^*}, \mathbf{Freq_{-f}^*})$ denote airline f's expected profit under the observed and alternative route networks $\mathbf{A_f^*}$ and $\mathbf{A_f^a}$, given competitors' observed optimal network $\mathbf{A_{-f}^*}$. This Π_f corresponds to the value on the right-hand side of equation (13). For notational simplicity, we omit the time index t, the optimal second-stage prices $\mathbf{p_t^*}$, the demand and marginal cost parameters θ_d and θ_s, and the demand and marginal cost shocks ξ and ω. By revealed preference, the observed network must yield at least as high an expected profit as any alternative network:

$$\Pi_f(\mathbf{A_f^a}, \mathbf{Freq_f^a}, \kappa; \theta_{fc}) - \Pi_f(\mathbf{A_f^*}, \mathbf{Freq_f^*}, \kappa; \theta_{fc}) = \Delta\Pi_f(\mathbf{A_f^*}, \mathbf{Freq_f^*}, \mathbf{A_f^a}, \mathbf{Freq_f^a}; \theta_{fc}) + \tau_f^a(\kappa) \leq 0. \tag{15}$$

The term $\Delta\Pi_f$ captures the profit difference between the observed (optimal) network and a counterfactual alternative, excluding the fixed cost shocks. The term $\tau_f^a((\kappa)$ represents the difference in fixed cost shocks between the observed and alternative networks. We further decompose $\Delta\Pi_f$ into the difference in second-stage variable profits and the difference in the linear component of fixed costs:

$$\Delta\Pi_f = \big[\mathbf{E}_{\xi,\omega}[\pi_f(\mathbf{A_f^a}, \mathbf{Freq_f^a})] - \mathbf{E}_{\xi,\omega}[\pi_f(\mathbf{A_f^*}, \mathbf{Freq_f^*})]\big] - \sum_{j \in G_f} \big[\theta_{fc}\mathbf{Z_j^a} \cdot A^a_{fj} - \theta_{fc}\mathbf{Z_j^*} \cdot A^*_{fj}\big]. \tag{16}$$

Here, $\mathbf{E}_{\xi,\omega}[\pi_f(\mathbf{A_f^a}, \mathbf{Freq_f^a})]$ and $\mathbf{E}_{\xi,\omega}[\pi_f(\mathbf{A_f^*}, \mathbf{Freq_f^*})]$ denote the second-stage expected profits under the alternative and observed networks, evaluated at the optimal prices for each draw of ξ and ω. To speed computation, we use a computing cluster with 36 cores per node to parallelise the calculation of expected second-stage profits.

Unlike the U.S. airline industry, where a change on one direct route can have knock-on effects on all connecting routes in a Hub-and-Spoke network, the intra-European airline industry is dominated by direct Point-to-Point networks. Because we evaluate only direct flights, as long as the product set in a market remains unchanged, prices and demand in that market are unaffected by changes in other markets[24]. In the second part of equation (16), $\mathbf{Z_j^a}$ and $\mathbf{Z_j^*}$ denote the vectors of market characteristics for each product under the alternative and observed networks.

The difference between the observed and alternative profit residuals $\tau^a_f(\kappa)$ in equation (15) is interpreted as measurement errors following Houde et al. [2023]. We use a vector of non-negative instruments H to cancel the measurement errors:

$$\mathbf{E}[H \cdot \Delta\Pi_f(\mathbf{A_f^*}, \mathbf{Freq_f^*}, \mathbf{A_f^a}, \mathbf{Freq_f^a}; \theta_{fc})] + \underbrace{\mathbf{E}[H \cdot \tau^a_f(\kappa)]}_{=0} \leq 0 \tag{17}$$

We use exogenous factors such as market size and distance as instruments. Then we construct the sample moment inequalities to estimate the fixed cost coefficients θ_{fc} following Pakes et al. [2015]:

[24]Strictly speaking, markets are not completely independent in our setting because deviations involve redeploying existing aircraft to alternative routes. However, once the route network and frequencies are fixed, conditions in one market do not directly affect another.

$$\frac{1}{N^a} \sum_{A_f^*, Freq_f^*, A_f^a, Freq_f^a} H_h \cdot [\Delta \Pi_f(\mathbf{A_f^*}, \mathbf{Freq_f^*}, \mathbf{A_f^a}, \mathbf{Freq_f^a}; \theta_{fc}] = \overline{m}_h(\theta_{fc}) \leq 0 \tag{18}$$

Equation (18) holds for all $h = 1, \dots, |H|$, where $|H|$ is the total number of instruments and thus the number of moment inequalities. N^a denotes the number of feasible alternative route networks for a given observed Quarter–Airline–Route. We now explain how the feasible alternative route network is defined.

1.4.2 Feasible Alternative Route Network

Exploring all possible alternative route networks is computationally infeasible because the number of combinations grows exponentially (2^M) and flight frequencies vary continuously, so we instead select a carefully constructed subset of alternatives that captures the most relevant and informative deviations, drawing on previous literature and the operational realities of the European airline market and guided by the following principles:

First, we consider only **single–market deviations**, following Yuan and Barwick [2024] and Bontemps et al. [2023]. This assumption is well suited to the European market, where the dominance of point–to–point operations limits interdependence across routes. It is less appropriate for the U.S. airline market tho, where Hub–and–Spoke networks dominate and a change in one "spoke" can generate extensive knock–on effects throughout the network via the hubs. In Europe, such spillovers arise only at a handful of large hubs and mainly affect indirect connections that combine an international long–haul flight with a "direct" intra–European short–haul leg. For these airports, the frequency constraints below ensures that the model remains consistent with operational realities.

Second, when an airline seeks an alternative route to replace an existing one, it must **redeploy the same aircraft** from the exited route to the new route. This implies that the alternative route must be operated at the **same frequency** as the original route. This contrasts with the U.S. literature, where entry and exit decisions are typically modelled independently. The rationale is that aircraft are high–value assets that generate revenue only when in service, and European carriers—especially low–cost airlines—maintain high utilisation rates. Rather than asking whether entering a new market is profitable, we ask which feasible route yields the highest profit conditional on the available aircraft (i.e., the fixed frequency). This approach better reflects actual airline behaviour and sharpens identification of the fixed–cost parameters.

Third, a route is deemed feasible for an airline if two conditions hold: (1) the airline already operates in both endpoint cities of the route, as in Berry [1992]; and (2) the route is currently served by at least one airline. The first condition captures the practical requirement for existing infrastructure and resources at both endpoints and highlights differences between full–service and low–cost carriers. Full–service airlines, which typically follow hub–based strategies, serve fewer cities and expand primarily from their hubs, whereas low–cost carriers serve many more cities and therefore face a larger feasible choice set. The second condition rules out routes with negligible or unobserved demand, such as those connecting economically isolated regions.

These constraints substantially reduce the number of feasible alternative route networks and, more importantly, focus the analysis on realistic and economically meaningful deviations, thereby avoiding distortions from extreme or implausible scenarios. Figure 9 summarises the network characteristics across airlines. In the *top–left* panel, network size refers to the total number of City–Airport pairs served by

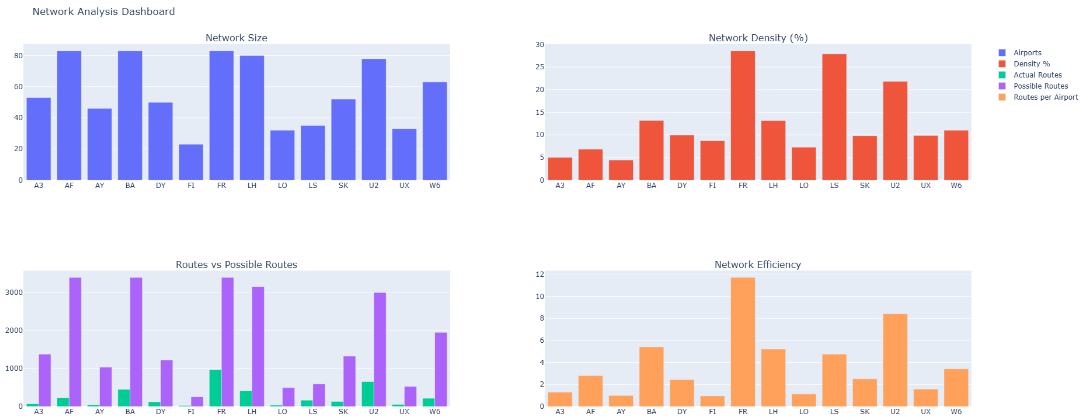

Figure 9: Network Analysis

each airline; a larger network size implies a broader set of feasible route alternatives. The three largest full–service airlines (AF, BA, LH) and the three largest low–cost carriers (FR, U2, W6) exhibit the largest network sizes. The *bottom–left* panel compares the number of observed routes with the number of potential deviation routes. A larger gap indicates weaker direct connectivity between city pairs served by an airline. By this measure, full–service carriers show lower connectivity across their served cities, reflecting their hub–and–spoke business model. This pattern is reinforced in the *top–right* panel, where low–cost carriers display higher network density percentages. Finally, the *bottom–right* panel reports network efficiency, defined as the average number of unique routes per City–Airport, which measures how intensively each served City–Airport is used. Low–cost carriers again score higher on this metric: for example, Ryanair (FR) operates on average more than ten unique routes per City–Airport it serves, whereas Air France (AF) averages fewer than three.

1.4.3 Estimation of the Linear Fixed Cost Parameters

We use two sets of moment inequalities, as shown in equation (15), to estimate the linear fixed–cost parameters θ_{fc}. First, the observed route must generate the highest expected profit among all feasible alternative routes operated at the **same frequency** (the redeployment argument). Second, the observed route must yield a non–negative expected profit (the entry argument).

Because we focus on single–market deviations, any deviation affects profits only in two markets: the market of the exited route and the market of the new entry route. Profits in all other markets remain unchanged. Under this structure, the first bracket of $\Delta\Pi_f$ in equation (16) for the redeployment inequality can be rewritten as

$$\big[\mathbf{E}_{\xi,\omega}[\pi_f(\mathbf{A}_f^a, \mathbf{Freq}_f^a)] - \mathbf{E}_{\xi,\omega}[\pi_f(\mathbf{A}_f^*, \mathbf{Freq}_f^*)]\big] = \mathbf{E}_{\xi,\omega}(\pi_{fj'}^a) - \mathbf{E}_{\xi,\omega}(\pi_{fj^*}^o). \tag{19}$$

Here $\mathbf{E}_{\xi,\omega}(\pi_{fj'}^a)$ is the expected profit from the alternative <u>single</u> new route (product) j' to which airline f redeploys its aircraft. $\mathbf{E}_{\xi,\omega}(\pi_{fj^*}^o)$ is the expected profit from the original observed route (product) j^*. Equation (19) therefore shows that, under a single–market deviation, the total profit difference reduces to the profit gap between the deviated route and the substituted route[25].

The second bracket of equation (16) can be simplified in the same way:

$$\sum_{j \in G_f} \left[\theta_{fc}\mathbf{Z}_j^a A_{fj}^a - \theta_{fc}\mathbf{Z}_j^* A_{fj}^*\right] = \theta_{fc}\big(\mathbf{Z}_{j'}^a - \mathbf{Z}_{j^*}^o\big). \tag{20}$$

Here $\mathbf{Z}_{j'}^a$ and $\mathbf{Z}_{j^*}^o$ are the vectors of two market attributes—frequency×distance

[25] In rare cases where an airline operates multiple products within the same market—only in large cities such as London or Paris—we assume that the airline exits all such products and reallocates the combined aircraft to the alternative route.

and market size—for the alternative and observed routes. Consequently, equation (18) becomes

$$\frac{1}{N^a}\sum_{j',j^*} H_h\left[\big(\mathbf{E}_{\xi,\omega}(\pi^a_{fj'}) - \mathbf{E}_{\xi,\omega}(\pi^o_{fj^*})\big) - \theta_{fc}\big(\mathbf{Z}^a_{j'} - \mathbf{Z}^o_{j^*}\big)\right] \leq 0. \tag{21}$$

For the instruments H_h, we follow standard practice and define them as dummy variables indicating whether a market's exogenous characteristics—such as market size or population—fall within the hth bin. Bins are created by evenly dividing the range of each characteristic into discrete intervals. As in prior studies, we vary the number of bins to test the robustness of the results.

Identification of θ_{fc} relies on the <u>differences</u> in attributes between observed and alternative routes. Any attribute that does not vary across routes—such as a constant term or a full–service airline dummy—cannot be identified through these inequalities. The second set of moment inequalities, imposing non–negative profit for each observed route, is defined similarly:

$$\frac{1}{N^o}\sum_{j^*_c} H^o_h\left[\mathbf{E}_{\xi,\omega}(\pi^o_{fj^*}) - \theta_{fc}\mathbf{Z}^o_{j^*}\right] \geq 0. \tag{22}$$

Here N^o is the number of unique observed routes, and the instruments H^o_h are constructed in the same way, although the number of dummies may differ.

At first glance, the inequalities in equation (23) seem to allow identification of terms that do not vary between observed and alternative routes, such as a constant or a full–service dummy. However, the redeployment inequalities in equation (22) identify only an <u>upper bound</u> for such terms, not a lower bound. Unlike the classic entry literature—where moment inequalities are based on a binary entry/exit choice—our inequalities cover only the "entry" side because of the frequency constraint. In practice, we find that even the upper bound is very loose regardless of the number of instruments. We therefore interpret these constant terms as part of the mean of the fixed cost shock κ, which is estimated separately in the next subsection.

The key difference between our moment inequalities and those in the existing literature lies in how the bounds of the identification set are determined. In our framework, the bounds are driven by the sign of the differences in market characteristics between observed and deviated routes. By contrast, previous studies typically derive these bounds from firms' entry and exit decisions. Identification in our setting therefore depends heavily on variation in characteristics between observed and alternative routes. This motivates the use of the interaction term $\underline{\text{frequency} \times \text{distance}}$ rather than frequency alone. Because frequency remains constant between the deviated and original routes, its stand–alone effect cannot be separately identified.

To implement these inequalities, we compute roughly 12.6 million expected profit values. These calculations are parallelised across a computing cluster to ensure feasible computation times.

1.4.4 Estimation of the Distribution of Fixed Cost Shocks

While the previous subsection provides estimates of the linear component of the fixed–cost parameters θ_{fc}, we still need to estimate the means and variances of the fixed–cost shocks as defined in equation (12). This step is essential because we must draw the realised shocks for each observed and alternative route in order to satisfy the revealed–preference inequalities in equation (15) for every observed–alternative pair.

The current IV–based moment inequalities in equations (21) and (22) only guarantee that the *weighted average* of the inequalities holds. This is sufficient for estimation but problematic for counterfactual analysis. In practice, we find that if the fixed–cost shocks are ignored, more than 50% of the revealed–preference inequalities are violated under the estimated $\hat{\theta}_{fc}$. When this occurs, the benchmark route network—defined as the optimal network under the estimated parameters following Yuan and Barwick [2024]—deviates substantially from the observed network. Such a gap undermines the interpretation of the counterfactual results because the benchmark for comparison is such different compared to the observed route network.

Our goal is therefore to estimate the distribution of fixed–cost shocks so that all inequalities implied by revealed preference hold, conditional on the estimated linear fixed–cost parameters. These conditions include

$$\mathbf{E}_{\xi,\omega}(\pi^a_{fj'}) - \hat{\theta}_{fc}\mathbf{Z}^a_{j'} + \kappa_{j'} \leq \mathbf{E}_{\xi,\omega}(\pi^o_{fj^*}) - \hat{\theta}_{fc}\mathbf{Z}^o_{j^*} + \kappa_{j^*}, \tag{23}$$

And the requirement that all observed routes earn non–negative expected profits:

$$\mathbf{E}_{\xi,\omega}(\pi^o_{fj^*}) - \hat{\theta}_{fc}\mathbf{Z}^o_{j^*} + \kappa_{j^*} \geq 0. \tag{24}$$

Inequality (23) must hold for every combination of observed Quarter–Airline–Route j^* and alternative Quarter–Airline–Route j', and inequality (24) must hold for every unique observed route. Let $\hat{\Pi}_{fj} = \mathbf{E}_{\xi,\omega}(\pi_{fj}) - \hat{\theta}_{fc}\mathbf{Z}_j$ denote the known part of the estimated route's expected profit. With this notation, inequality (23) can be rewritten as

$$\kappa_{j'} \leq \hat{\Pi}^o_{fj^*} - \hat{\Pi}^a_{fj'} + \kappa_{j^*} \tag{25}$$

And inequality (24) as

$$\kappa_{j^*} \geq -\hat{\Pi}^o_{fj^*}. \tag{26}$$

Each observed route j^* is associated with many alternative routes j'. Fix attention on a particular observed route j^* and its set of alternatives, the joint probability that all inequalities in equation (25) hold is

$$\prod_{j'} \Phi\left(\frac{\hat{\Pi}^o_{fj^*} - \hat{\Pi}^a_{fj'} + \kappa_{j^*} - \mu_{gj'}}{\sigma_{gj'}}\right). \tag{27}$$

Given j^*, the inequalities for different alternatives j' are independent. The joint probability is therefore the product of standard normal CDFs with mean $\mu_{gj'}$ and standard deviation $\sigma_{gj'}$, where g indexes the group of the alternative route j'. Groups are defined by both the airline and whether the route connects a hub. Thus, even alternatives of the same airline may have different means and variances depending on whether the alternative route link hubs.

Combining equations (25), (26), and (27), the joint probability for an observed route j^* is

$$\int_{-\hat{\Pi}^o_{fj^*}}^{\infty} \prod_{j'} \Phi\left(\frac{\hat{\Pi}^o_{fj^*} - \hat{\Pi}^a_{fj'} + \kappa_{j^*} - \mu_{gj'}}{\sigma_{gj'}}\right) \phi\left(\frac{\kappa_{j^*}}{\sigma_{gj^*}}\right) \frac{1}{\sigma_{gj^*}} d\kappa_{j^*}. \tag{28}$$

The integral in equation (28) is taken over a *truncated* normal distribution of κ_{j^*}, where the lower bound is set by inequality (26). The total log–likelihood for all observed routes of all airlines is then

$$\sum_{j^*} \ln \int_{-\hat{\Pi}^o_{fj^*}}^{\infty} \prod_{j'} \Phi\left(\frac{\hat{\Pi}^o_{fj^*} - \hat{\Pi}^a_{fj'} + \kappa_{j^*} - \mu_{gj'}}{\sigma_{gj'}}\right) \phi\left(\frac{\kappa_{j^*}}{\sigma_{gj^*}}\right) \frac{1}{\sigma_{gj^*}} d\kappa_{j^*}. \tag{29}$$

This structure resembles a standard probit model but with key differences. First, in a conventional binary entry model, the entry and exit shocks are independent because they never appear in the same inequality. Here, all alternative–route shocks associated with a given observed route are correlated through equation (25). We must therefore evaluate each observed route and all of its alternatives jointly. Second, the classic probit model does not involve a truncated normal distribution, whereas truncation arises here because the entry condition applies only to observed routes and not to alternatives.

Compared with the moment–inequality literature, which typically includes an intercept in the linear component, our maximum–likelihood approach offers several advantages. First, it can point identify both the mean and the variance of the fixed–cost shocks, whereas moment inequalities provide only set identification. Second, it allows us to estimate a rich set of group–specific mean terms. As discussed earlier, we model eleven different groups defined by airline type and hub status, and we estimate a distinct mean for each group. Such richness would be practically infeasible in a pure moment–inequality framework, because increasing the number of parameters causes the identified sets to interact and greatly complicates estimation.

1.5 Estimation Results

1.5.1 Demand Estimates

Table 4 reports the demand–estimation results for the core parameters and selected fixed effects. We use two types of instruments: (i) the number of products offered in each market, and (ii) the average fare of routes with similar distances[26]. Following the standard BLP literature, we report the first–stage F–statistic for price as an indicator of instrument relevance. The first–stage F–statistic is 172.55, and the heteroskedasticity–robust F–statistic is 91.16. Both far exceed the conventional threshold of 10, indicating that the instruments are strongly correlated with price and are unlikely to suffer from weak–instrument problems.

[26] Routes with distances between 99% and 101% of the current route, taking advantage of the dense European airline network.

For clearer economic interpretation, we translate the estimated coefficients into willingness–to–pay (WTP) measures.[27] On average, consumers are willing to pay about \$22.4 for a one–unit increase in the log of daily flight frequency, reflecting the high value passengers place on schedule convenience.

The distance coefficients reveal a more nuanced trade-off. The linear term is positive, and the squared term is also positive but smaller, producing a convex relationship. Evaluated at the sample mean distance of roughly 1,000 km, consumers are willing to pay approximately

$$\frac{0.325 + 2 \times 0.145 \times 1}{|-5.426|} \times 100 \approx \$8.6$$

for an additional 1,000 km of travel. This WTP increases with distance because the marginal utility of distance becomes more positive on longer routes. Economically, passengers are less price-sensitive on longer flights where practical travel alternatives are limited. At shorter distances the incremental value is smaller—especially in Europe, where rail and car travel provide strong substitutes—and could even turn negative for very short segments if the squared term dominates.

Seasonal preferences are also evident. Relative to the baseline quarter (Q1), consumers value spring (Q2) flights about \$9.8 more and summer (Q3) flights about \$3.3 more, while winter (Q4) shows no significant difference. Carrier and hub effects are also obvious. Full-service airlines receive sizable premia: consumers are willing to pay roughly \$32.7 more to fly with the Air France–KLM Group than with Ryanair on the same route. Air France–KLM also enjoys strong hub advantages: itineraries involving CDG or AMS are valued about \$36 higher than competing services on

[27] WTP is calculated as the ratio of the coefficient of interest to the absolute value of the price coefficient. It measures how much more consumers are willing to pay for a one–unit change in a product characteristic, holding utility constant.

identical markets.

The estimated nesting parameter is 0.885 and highly significant. In the nested-logit framework, this parameter captures the correlation in unobserved utility among products within the same nest. Here all airline itineraries form one nest and the outside option forms the other. A value close to one indicates strong substitution among airline products and substantial correlation in their unobserved components (for example, common shocks such as weather or macro-demand factors). It also confirms that the nested-logit specification is appropriate, lying comfortably between the simple logit case ($\lambda = 0$) and the degenerate case of perfectly correlated products ($\lambda = 1$).

Overall, the WTP estimates highlight the key drivers of consumer choice in European short-haul aviation: a high value on frequency, a non-linear premium for longer distances, pronounced seasonal patterns, and significant brand and hub advantages.

Figure 10 shows the distribution of the recovered demand shocks ξ for one full–service airline (BA) and one low–cost airline (FR). In both cases, the mean is close to zero and the spread between the 5th and 95th percentiles is roughly one, indicating a shape that is close to a standard normal distribution. This pattern is consistent with the BLP framework, where ξ captures unobserved product–market characteristics that are orthogonal to the instruments after estimation. A distribution centred near zero with unit–like dispersion suggests that the instruments successfully purge price endogeneity and that the structural error behaves like a well–specified mean-zero disturbance. Such a pattern is typically viewed as evidence of a good model fit: if the model were badly misspecified, the recovered ξ would display large systematic biases or heavy tails rather than a near-Gaussian shape. Similar distributions are found for other airlines, reinforcing the conclusion that the

Table 4: Demand Estimation Results: Core Parameters and Selected Fixed Effects

Variable	**Coefficient**	**Std. Error**
Core Demand Parameters		
Price (in 100 USD)	-5.426***	0.515
Log Frequency	1.217***	0.037
Distance (in 1,000 KMs)	0.325**	0.163
Distance^2	0.145***	0.033
Nesting Parameter	0.885***	0.054
Q2	0.533***	0.084
Q3	0.181***	0.065
Q4	-0.018	0.063
Airline Fixed Effects		
Airline_BA	3.449***	0.401
Airline_AF	1.663***	0.236
Airline_LH	3.585***	0.452
Airline_FR	-0.094*	0.053
Airline_W6	-0.263***	0.089
Airport Fixed Effects		
Airport_AMS	-0.580***	0.082
Airport_FRA	-0.621***	0.105
Airport_MAD	-1.554***	0.118
Airport_BCN	-1.701***	0.103
Airport_VIE	-0.742***	0.058
City Fixed Effects		
London/Southend/Cambridge	-1.100***	0.160
Paris/Pontoise	-1.309***	0.145
Amsterdam/Rotterdam	-0.580***	0.082
Dusseldorf/Dortmund/Cologne	-0.608***	0.194
Rome	-1.412***	0.107
Madrid	-1.554***	0.118
Airline–Airport Fixed Effects		
AF_CDG	1.987***	0.323
AF_AMS	1.951***	0.358
BA_LHR	0.633	0.417
LH_FRA	-0.037	0.259

Notes: Stars denote significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Seasonal dummies (Q2–Q4) control for quarterly variation. Airline–airport interactions capture route-specific advantages.

demand system captures the main drivers of consumer choice. The estimation is based on a total of 11,062 observed airline–market combinations.

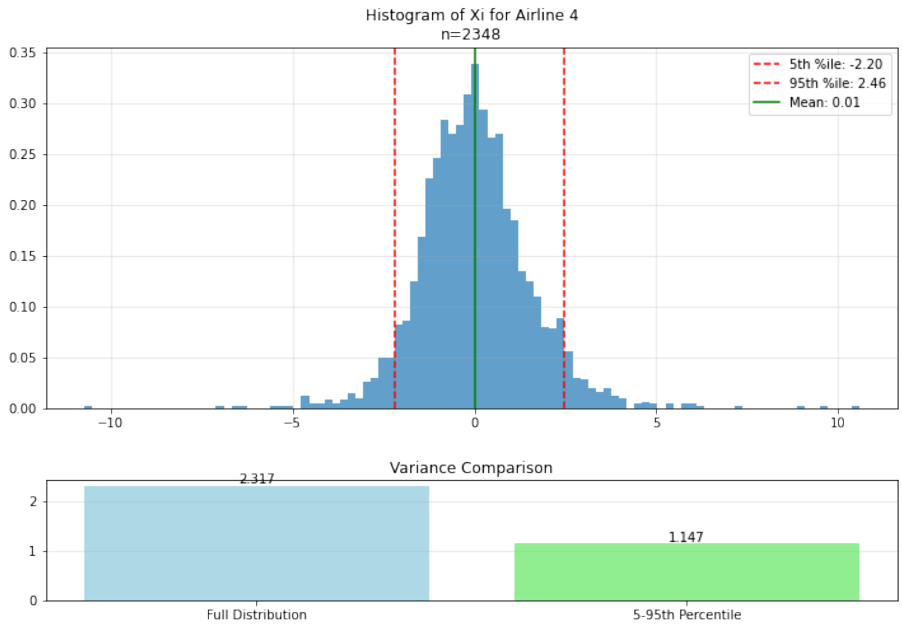

ξ Plot for BA (IAG Group)

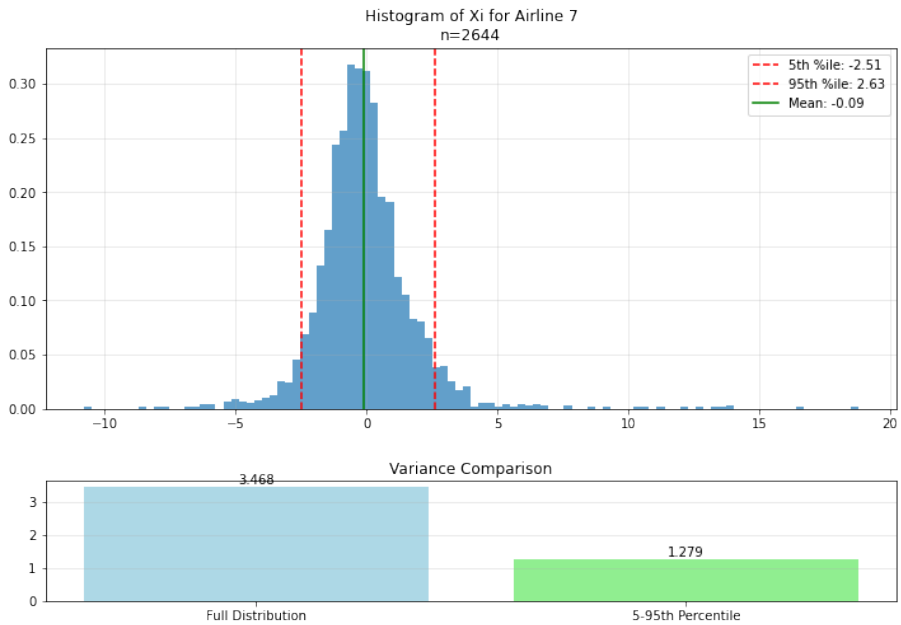

ξ Plot for FR (Ryanair)

Figure 10: ξ Plot for Two Representative Airlines

Figure 11 shows the distribution of own- and cross-price elasticities implied by

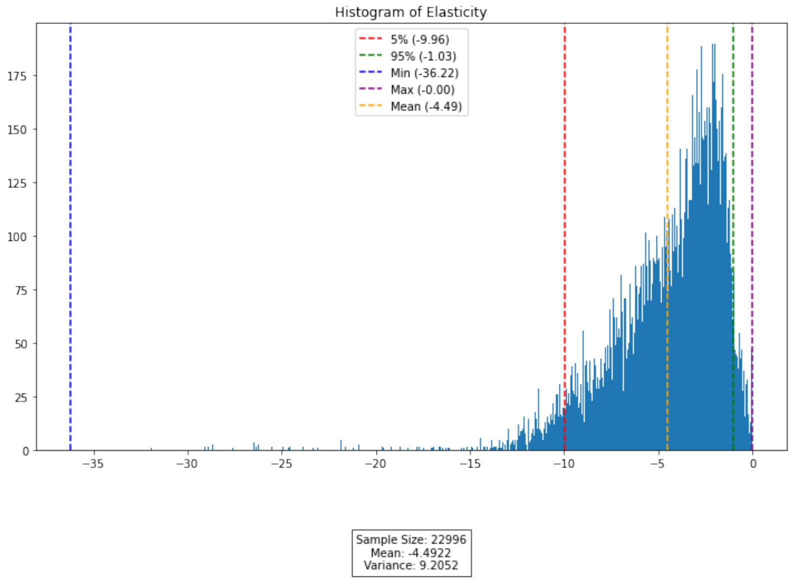

Own-Price Elasticity Plot

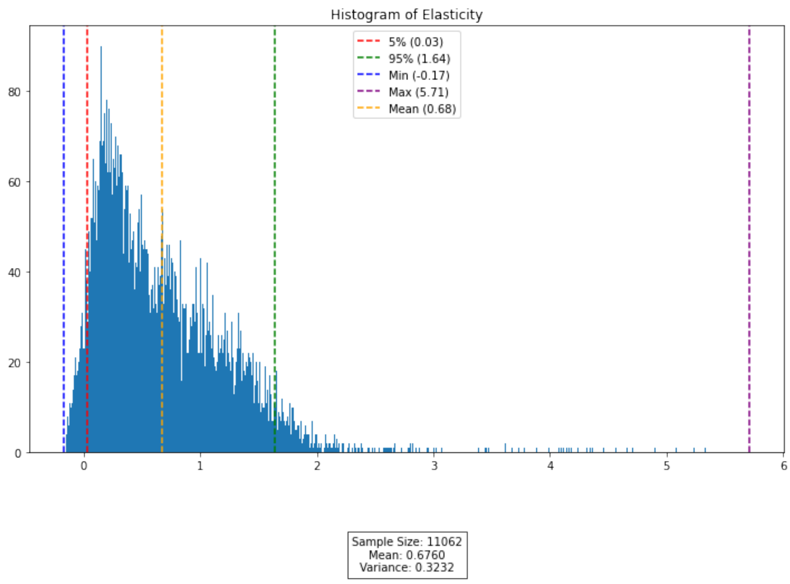

Cross-Price Elasticity Plot

Figure 11: Price Elasticity Plot

the BLP estimates. The average own-price elasticity is approximately –4.49, which is close to the estimate reported in Bontemps et al. [2023] (–3.78) and notably more elastic than the values from the two-consumer-type model in Berry and Jia [2010].

Elasticities of this magnitude are consistent with evidence from the airline industry, where empirical studies of European short-haul markets typically find own-price elasticities ranging between –3 and –5 for leisure-dominated routes. Such values indicate that passengers are quite sensitive to fare changes: a 1% increase in price leads, on average, to roughly a 4.5% decrease in demand. This high responsiveness reflects the availability of close substitutes—both between airlines on the same city pair and across alternative modes of transport. The cross-price elasticities, while smaller in absolute value, confirm significant substitution across carriers operating in the same market, reinforcing the interpretation of a highly competitive environment.

Several factors help explain why our estimate is more elastic than the typical values reported in the U.S. literature. European air travellers are generally more price sensitive, as documented in both empirical studies and industry reports. This heightened sensitivity reflects the greater presence of low–cost carriers, denser and more competitive point-to-point networks, and, on average, lower income levels across Europe.[28] In contrast, the estimated cross–price elasticities are mostly positive, consistent with standard substitution patterns among competing airline products and confirming that passengers readily switch to rival carriers when relative fares change.

1.5.2 Marginal Cost on Passengers

Marginal costs are recovered from the first–order condition in Equation (7). Table 5 reports the estimates. The average marginal cost per passenger is \$67.6, roughly 30% lower than comparable U.S. estimates such as those in Yuan and Barwick [2024], who report mean marginal costs around \$95 per passenger for similar short–haul markets.

[28] See, for example, IATA's report: IATA.

Table 5: Marginal Cost Estimation Results: Core Parameters and Selected Fixed Effects

Variable	**Coefficient**	**Std. Error**
Core Cost Parameters		
Distance (in 1,000 KMs)	0.098***	0.020
Distance^2	0.024***	0.006
Frequency	0.064***	0.005
Q2	0.090***	0.010
Q3	0.021**	0.010
Q4	-0.006	0.010
Airline Fixed Effects		
Airline_AF	0.608***	0.019
Airline_BA	0.744***	0.016
Airline_LH	0.831***	0.018
Airline_FR	0.020	0.013
Airline_W6	-0.044**	0.022
Airport Fixed Effects		
Airport_FRA	0.099***	0.011
Airport_CDG	0.217***	0.026
Airport_LHR	0.175***	0.027
Airport_AMS	0.117***	0.010
City Fixed Effects		
London/Southend/Cambridge	0.150***	0.017
Paris/Pontoise	0.041**	0.022
Amsterdam/Rotterdam (Randstad)	0.117***	0.01
Frankfurt/Mannheim	0.099***	0.011
Dusseldorf/Dortmund/Cologne	0.243***	0.022
Statistics	*Value*	
Average Marginal Cost	\$67.6	
Average Markup	\$17.9	
Average Percentage Markup	35.2%	
Average Profit	560,586	

Notes: Stars denote significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Seasonal dummies (Q2–Q4) control for quarterly variation. Selected fixed effects are displayed for airlines, cities, and airports with high statistical significance.

This difference is in line with broad industry evidence. European carriers consistently report lower unit operating costs than their U.S. counterparts. For example, IATA cost benchmarking shows that European short–haul airlines have cost per available seat kilometre (CASK) roughly 20–35% below that of major U.S. legacy carriers over the past decade, largely because of a higher share of low–cost carriers, denser route networks, and more efficient aircraft utilisation.[29] Low–cost carriers such as Ryanair and Wizz Air routinely report CASK levels less than half of those of U.S. full–service carriers, and their presence drives average European unit costs downward even for network airlines.

The average route distance in our sample is 1,407 kilometres (about 875 miles), which implies a unit cost of roughly \$0.05 per kilometre or \$0.08 per mile. These figures closely match international benchmarks: Berry and Jia [2010] report about \$0.06 per mile for U.S. domestic flights, while Yuan and Barwick [2024] find around \$0.08 per mile. IATA cost data for European short–haul operations similarly cluster in the \$0.05–\$0.09 per mile range once adjusted for fuel prices and exchange rates, reinforcing the plausibility of our estimates.

The implied markup is also sizeable. The average markup is \$17.9, corresponding to an average percentage markup of 35.2% and an average per–route profit of roughly \$0.56 million. These figures are broadly consistent with European airline financial statements and with the 25–35% margin estimates commonly reported for competitive U.S. domestic routes. Higher airport charges and slot constraints in Europe may also sustain slightly higher margins even in markets served by multiple carriers.

The cost coefficients reveal clear economic patterns. Both distance and distance^2

[29]See IATA Annual Review (various years) and InterVISTAS (2015) Estimating Air Travel Demand Elasticities, which report CASK figures for major world regions.

are positive and significant, implying that marginal cost rises at an increasing rate with route length. This convexity reflects the growing cost of fuel, crew time, and maintenance over longer legs and is consistent with engineering cost studies for narrow–body fleets. The coefficient on frequency is also positive, in contrast to many U.S. studies where frequency often lowers marginal cost by spreading fixed expenses across more departures. In Europe, two factors likely drive this difference. First, European carriers operate with consistently high load factors—often above 85%—leaving little unused capacity to absorb additional flights. Second, high–frequency services are typically short–haul "city–hopper" routes (e.g., London–Amsterdam or Madrid–Barcelona) where airlines deploy smaller regional jets with higher per–seat operating costs.[30]

As expected, full–service carriers face higher marginal costs than low–cost airlines, and operating from large hub airports (e.g., FRA, CDG, LHR) is also associated with higher costs. These patterns mirror industry evidence on cost heterogeneity: full-service airlines incur higher labour and service costs, while congested hubs impose higher landing fees and turnaround expenses. The recovered marginal–cost distribution, together with realistic markups and distance–cost relationships, supports the internal consistency of our model and aligns well with both academic estimates and industry cost benchmarks for European short–haul aviation.

1.5.3 Linear Fixed Cost Estimation

Table 6 presents the attributes and expected profits of observed and alternative routes by airline. For each metric, Diff (%) is computed as (Observed − Alt.)/Observed $\times$ 100. This difference is the key variable used in the moment in-

[30]For example, British Airways frequently operates Embraer 190s from London City Airport to destinations such as Dublin and Amsterdam, which raises per–passenger marginal costs relative to larger narrow–body aircraft.

Table 6: Route's attributes and Expected Profits for Observed and Alternative Network

Airline	Frequency × Distance			Market Size (millions)			Expected Variable Profit ($000s)			Obs. (000s)
	Observed	Alt.	Diff (%)	Observed	Alt.	Diff (%)	Observed	Alt.	Diff (%)	
Full-Service Airlines										
BA (British Airways)	1,483	2,042	-37.7	3.258	2.622	+19.5	901.2	813.1	+9.8	2,051
AF (Air France)	1,504	2,315	-53.9	3.489	2.587	+25.9	1,090.7	337.3	+69.1	1,386
LH (Lufthansa)	1,613	3,020	-87.3	3.329	2.636	+20.8	603.4	778.5	-29.0	2,127
Low-Cost Airlines										
FR (Ryanair)	611	617	-1.1	2.620	2.944	-12.4	474.7	431.4	+9.1	2,559
U2 (easyJet)	749	1,028	-37.2	2.641	2.748	-4.0	761.2	809.7	-6.4	2,487
W6 (Wizz Air)	479	442	+7.6	2.252	3.160	-40.3	375.0	363.7	+3.0	763
Regional Airlines										
A3 (Aegean)	765	571	+25.4	3.020	3.480	-15.2	213.2	244.6	-14.7	242
AY (Finnair)	1,314	1,099	+16.3	2.026	3.486	-72.0	410.3	686.2	-67.2	136
SK (SAS)	842	1,209	-43.7	2.284	3.456	-51.3	518.5	704.6	-35.9	331
LO (LOT Polish)	1,649	1,737	-5.3	3.230	4.088	-26.6	287.2	403.5	-40.5	60
UX (Air Europa)	971	1,231	-26.8	3.601	4.004	-11.2	506.7	626.8	-23.7	51
DY (Norwegian)	438	661	-50.8	1.798	3.050	-69.7	529.3	668.0	-26.2	237
FI (Icelandair)	3,706	1,595	+57.0	0.937	4.501	-380.2	421.1	684.1	-62.5	21
LS (Jet2)	1,320	827	+37.3	1.956	3.325	-70.0	408.9	417.0	-2.0	124
Market Average	1,072	1,546	-44.2	2.896	2.822	+2.6	679.5	624.5	+8.1	12,578

Notes: Frequency × Distance measured in daily flights × km. Market size represents quarterly geometric mean of the population at two endpoints in millions. Alt. denotes alternative routes. Diff (%) shows percentage change from observed to alternative values. Obs. denotes the count of pairs between one observed route and one alternative route in thousands. Full-service airlines comprise legacy carriers, while regional airlines include smaller carriers serving specific markets.

equalities to estimate θ_{fc}.

For the Frequency×Distance measure, full–service airlines (FSCs) consistently operate longer routes than low–cost carriers (LCCs) in both observed and alternative networks. This pattern reflects the distinct business models of the two groups. Hub–and–spoke networks naturally link longer city–pairs to one or more hubs, whereas point–to–point strategies favour shorter sectors to maximise daily aircraft utilisation—a hallmark of the European low–cost model.[31]

The difference term for Frequency×Distance shows that all FSCs would, if free to redeploy aircraft, shift toward alternative routes with longer distances than those currently operated. Because frequencies are held constant, the choice sets of alternative network for FSCs contain routes of greater length than their current networks. Among LCCs, the picture is more heterogeneous. easyJet, often described as a "hybrid" or semi–full–service carrier, displays a pattern similar to the FSCs, consistent with its strategy of operating both dense leisure city–pairs and key primary airports. Ryanair's Frequency×Distance shows almost no difference between observed and alternative routes. This is intuitive because Ryanair already operates the most extensive network in Europe, serving nearly every major city–pair of economic relevance, so potential alternatives offer similar distances and therefore limited scope for reallocation. Wizz Air, headquartered in Budapest and heavily focused on Eastern Europe, shows a positive difference (alternative routes shorter on average). Many of its feasible redeployments link medium–sized cities in Central and Eastern Europe, because the largest Eastern European markets are already present in its current network. Several regional carriers, such as Finnair and Icelandair, exhibit very high Frequency×Distance values, reflecting the remote geographic position of their hubs

[31] EUROCONTROL's Data Snapshots document the higher average stage length of full-service carriers and the shorter, more numerous sectors flown by European LCCs.

(Helsinki and Reykjavik) and the long sectors required to connect them to the rest of Europe.

Full–service airlines also tend to serve markets with larger populations than low–cost carriers. The difference terms show that FSCs' alternative routes generally connect smaller endpoint populations than their observed routes. This is natural because most large European city–pairs are already covered in their current networks, so remaining alternatives involve thinner markets. In contrast, all LCCs have alternative routes with larger market size than the observed routes. This reflects their tactical avoidance of certain large markets in the observed network—often to avoid higher landing fees, congestion charges, or labour costs at primary airports—and their focus on secondary airports around major metropolitan areas.[32]

Expected variable profits from second–stage price competition also differ sharply by business model. FSCs earn higher expected profits than LCCs on both observed and alternative routes, reflecting their ability to command price premia through brand reputation, business–class demand, and hub connectivity. The differences between observed and alternative profits are larger for FSCs than for LCCs, a result of strong hub effects. Most of the profitable hub routes for FSCs are already included in their current networks; alternative routes are therefore more likely to be non–hub markets where network economies are weaker and price competition is stronger. This large gap supports our decision to allow hub and non–hub routes within the same full–service airline to follow different distributions of fixed–cost shocks. Among LCCs, profit differences are modest, consistent with already optimised point–to–point schedules and intense fare competition on thick leisure markets. All regional carriers show higher expected profits for alternative routes. Their

[32]For example, Ryanair often uses airports such as Charleroi for Brussels and Beauvais for Paris, allowing it to tap large catchment areas while avoiding the high costs of main hubs.

continued operation of current networks is likely sustained by substantial subsidies and public–service obligations, which reduce the private incentive to redeploy capacity even when profitable alternatives exist.[33]

Large full-service and low-cost carriers have far more observations than regional airlines. On the one hand, this richer data generates greater variation for the moment-inequality estimation. On the other hand, it allows us to estimate the fixed-cost distribution separately for each of the six largest European airlines—a level of flexibility that is rarely achievable in the existing literature.

Taken together, the counterfactual reallocations push full–service airlines toward longer routes but smaller markets, while low–cost airlines tilt toward larger markets with limited changes in route distance—patterns consistent with their point–to–point, high–utilisation model. These findings are in line with European industry norms: high LCC penetration and sustained load factors above 85%[34] leave little spare capacity for large schedule reoptimisations, while hub–and–spoke legacy carriers face an inherent trade–off between maintaining network connectivity and chasing pure market size.[35]

We use standard exogenous instruments—such as market population and route distance—following the moment–inequality literature. Table 7 reports the identified sets for the linear fixed–cost parameters when varying the number of instruments used in the difference and observed–profit inequalities. When the number of instruments is small, the identified sets are wide: for example, with 20 difference inequal-

[33]Regional carriers in Europe frequently receive national or EU subsidies, particularly on thin peripheral routes; see European Commission reports on Public Service Obligation (PSO) routes. It is worth noting, however, that large legacy carriers such as Air France also receive state support, especially during crises.

[34]IATA reports European carriers' 2023 load factor at about 83.8%; Ryanair's investor reports indicate load factors above 93% in recent years.

[35]See EUROCONTROL on LCC market share and network structure, and IATA monthly and annual market analyses for European load factors and capacity trends.

Table 7: Set-Identified Fixed Cost Bounds Under Different Instrument Combinations (in \$100s)

IV Count		**Frequency × Distance**		**Market Size**	
Diff. Ineq.	**Obs. Ineq.**	**Lower**	**Upper**	**Lower**	**Upper**
	10	770	1,743	445	834
20	20	770	1,657	445	834
	30	770	1,657	445	834
	10	1,086	1,417	433	654
30	20	1,086	1,417	433	654
	30	1,086	1,387	433	654
	10	Empty Set		Empty Set	
40	20	Empty Set		Empty Set	
	30	Empty Set		Empty Set	

Notes: Each coefficient is set-identified using moment inequalities. "Diff. Ineq." refers to difference-based moment inequalities requiring the observed route to have the highest expected profit among alternatives. "Obs. Ineq." refers to moment inequalities requiring observed routes to have non-negative profits. Both sets use Market Size and Distance as instruments, with counts shown in the first two columns. Frequency × Distance measured in daily flights × thousands of kilometers. Market size is the geometric mean of endpoint populations in millions. Empty sets indicate over-identification where no parameter values satisfy all moment conditions.

ities and only 10 observation inequalities, the bounds for the Frequency×Distance coefficient range from 770 to 1,743 (in \$100s). As the number of instruments increases, the bounds tighten considerably, reflecting stronger informational content and more stringent moment restrictions. When the instrument count becomes too large (e.g., 40 difference inequalities), the feasible set collapses to an <u>empty set</u>, indicating that the additional moments over–identify the model and no parameter vector can satisfy all conditions simultaneously. This pattern is consistent with the findings of Yuan and Barwick [2024] and other applications of set–identified moment inequalities, where the tension between sampling variation and strong instrument sets can lead to empty identified sets.

Compared with existing studies, our identified sets are substantially tighter. This improvement reflects the richer variation in our European route–level data and the greater economic content of our moment inequalities, which exploit both deviation profits and non–negative observed profits across a large network of markets. The high number of deviations further enhances the precision of the moments, allowing for sharper bounds even with relatively few instruments.

Interestingly, the bounds for the Frequency×Distance coefficient change very little when we vary the number of instruments for the observation inequalities while holding the difference–inequality instruments fixed. For example, with 20 difference inequalities, increasing the observation instruments from 10 to 30 leaves the lower and upper bounds essentially unchanged. This suggests that identification of deviation–sensitive parameters is primarily driven by the first set of moments (the difference inequalities), where most of the empirical variation occurs when airlines consider redeployment. In contrast, the second set of moments (non–negative profit conditions) adds little incremental identifying power once the core deviation structure is captured.

From an economic perspective, these results imply that the key identifying information for fixed costs comes from the behaviour of airlines when they face realistic redeployment choices rather than from simple entry profitability conditions. This is consistent with industry intuition: in congested European markets, airlines continually re–optimise aircraft allocation across feasible routes, and the profit differentials between observed and alternative routes reveal more about underlying fixed costs than the mere fact that a route remains active.

Given the estimation results and the sample averages (Distance = 1,350 km; Market size = 2.896 million), the implied per-flight linear component of the fixed cost (distance term + market-size term) evaluates to the midpoint value[36] of approximately \$3,602.89.

The estimated figure above represents the linear contribution to per-flight fixed cost that is attributable to the observable covariates (distance and market size) evaluated at sample means. It does not include the stochastic fixed-cost shock κ. Equivalently, this amount is the per-flight value implied by the estimated linear parameters alone and therefore constitutes a partial yet comprehensive measure of per-flight operating cost.

We clarify the economic interpretation of the \$3.6k fixed-cost component through three key considerations, contextualised within European airline industry benchmarks.

First, our framework separately identifies per-passenger marginal costs during

[36] Computation (units: USD per flight; coefficients reported in the table are in \$100s):

$$\text{Distance part} = \frac{(1086 + 1387)}{2} \times 100 \div 90 \times 1.350 = 1236.5 \times 100 \div 90 \times 1.350 \approx 1{,}854.80,$$

$$\text{Market size part} = \frac{(433 + 654)}{2} \times 100 \div 90 \times 2.896 = 543.5 \times 100 \div 90 \times 2.896 \approx 1{,}748.09,$$

$$\text{Total (Distance + Market)} \approx 1{,}854.80 + 1{,}748.09 \approx \$\mathbf{3{,}602.89}.$$

demand estimation. Industry-standard cost metrics typically reported on a per-flight basis—such as fuel allocated by available seat-kilometres (ASK), variable handling fees, or aircraft turnaround costs—are captured both within our marginal-cost and fixed cost estimates. Consequently, the \$3.6k figure should not be interpreted as the total "per-flight" operating cost reported in airline financial statements. Instead, it represents the specific component of fixed costs that scales linearly with distance and market size within our econometric specification.

Second, industry benchmarking employs standardised metrics such as CASK (cost per available seat-kilometre). European carriers exhibit substantial variation: ultra-low-cost carriers report CASK values of 3.2-4.2 US cents, whilst full-service carriers exceed 10 US cents (CAPA, 2025; Wizz Air H1 FY24). Since CASK declines systematically with stage length, converting to per-flight equivalents requires aircraft-specific adjustments. For narrowbody aircraft typical of European short-haul operations, industry sources report operating costs of \$2,900-\$3,200 per block hour for A320/B737 aircraft (OPShots, 2015; Simple Flying, 2024), suggesting a 1.5-hour flight at 1,300km incurs approximately \$4,350-\$4,800 in total costs. Within this context, our \$3.6k estimate represents a plausible fixed-cost component, accounting for roughly 75-80% of total per-flight costs.

Third, whilst many empirical studies normalise fixed-cost shocks to have zero mean, we do not impose this restriction. A non-zero mean of κ implies that a portion of the true per-flight fixed cost could be partially absorbed into the intercept term. Specifically, total per-flight fixed cost comprises three components: (i) the linear term computed here (\$3.6k), (ii) the intercept (mean of κ), and (iii) the route-specific idiosyncratic shock. Therefore, actual per-flight fixed costs (linear component + intercept + shock) will exceed \$3.6k in many cases. Industry evidence confirms substantial variation in total costs: European low-cost carriers report per-passenger costs ranging from €40 for Ryanair to €79 for easyJet (excluding fuel), whilst legacy carriers like IAG and Lufthansa operate at €159-164 per passenger (The Flight Club, 2025). With typical load factors of 85-90% on 150-180 seat aircraft, this translates to per-flight costs varying from approximately \$5,100 to \$25,000 across different business models (EUROCONTROL, 2024), supporting our framework where fixed costs include both the linear component and additional stochastic elements. We will estimate the distribution of fixed cost shocks κ in the next subsection.

1.5.4 Fixed Cost Shocks Estimation

We estimate the means and variances of the fixed-cost shock distributions by maximizing the likelihood in equation (29). Before applying the gradient–based Newton method, we evaluate the likelihood over a grid of mean and variance values to confirm the presence of a unique global maximum.[37] Table 8 reports the estimation results. We estimate separate distributions for the three largest full-service carriers and the three largest low-cost carriers, as well as a pooled distribution for all regional airlines. For the full-service carriers and the pooled group, we further distinguish hub and non-hub routes. In total, this yields 11 distinct fixed-cost shock distributions.

The results demonstrate remarkable consistency with the chosen linear fixed cost parameters, indicating minimal unexplained variation in our previous moment inequality estimation, which has been effectively absorbed into the fixed cost shocks. The stability of these parameters across different specifications suggests robust identification of the underlying cost structure.

The most significant and striking finding emerges from the systematic differences between hub routes and non-hub routes. This pattern reinforces the critical impor-

[37] We use a 50×50 grid, with ranges spanning from the minimum to the maximum of the expected variable profit from price competition, net of the linear component of fixed cost.

Table 8: Distribution for Fixed Cost Shocks by Airlines and Hub Status

		Lower Bound		Middle Point		Upper Bound	
Airline	Route Type	μ	σ	μ	σ	μ	σ
Full-Service Airlines							
AF	Non-hub	-3.59	3.82	-3.69	3.87	-3.94	3.99
	Hub	3.61	4.17	3.72	4.19	3.97	4.26
BA	Non-hub	-5.08	4.99	-5.08	5.00	-5.08	5.04
	Hub	2.97	4.22	2.97	4.27	2.97	4.36
LH	Non-hub	-5.53	4.22	-5.53	4.24	-1.66	4.39
	Hub	3.16	4.99	3.16	5.03	7.47	5.08
Low-Cost Airlines							
FR	All routes	-2.87	3.61	-2.87	3.62	-2.87	3.67
U2	All routes	-8.18	7.20	-8.18	7.17	-8.18	7.13
W6	All routes	-5.07	2.61	-5.07	2.61	-5.07	2.64
Pooled (Other Airlines)							
Pooled	Non-hub	-5.64	3.67	-6.11	3.72	-5.61	3.70
	Hub	10.22	2.12	10.05	2.15	10.53	2.19

Notes: The reported parameters are the mean (μ) and standard deviation (σ) of these distributions, expressed in units of 10^5 USD. The terms Lower Bound, Middle Point, and Upper Bound refer to the set estimates of the linear fixed-cost parameters.

tance of hub operations for full-service airlines' strategic positioning and network economics. Specifically, hub routes for all full-service airlines exhibit positive mean fixed cost shocks, whilst all non-hub routes and routes operated by low-cost carriers demonstrate consistently negative means. The mean of fixed cost shocks can be explained as the pure "entry cost" plus any unexplained effects from the structural model and data that render the observed route optimal amongst all available alternatives. The positive mean for hub routes indicates that, all else equal, there exists a strong unexplained positive effect favouring observed hub routes, making them preferable despite potentially higher underlying costs—otherwise, they would not represent the optimal route choice.

The positive hub route's mean can be explained by established industry knowledge and practices. First, the fundamental purpose of many hub routes operated by major full-service airlines in Europe is to facilitate international transfer passengers connecting to or from long-haul flights.[38] Evaluating these routes in isolation, considering only point-to-point demand, would rarely demonstrate profitability. The network effects and connecting passenger flows create substantial value that is not captured in simple route-level analysis.

Second, European full-service airlines frequently benefit from extensive government incentive schemes and subsidies, which fundamentally differs from the competitive landscape faced by major US carriers. Many European airlines remain government-owned or receive substantial state support. Research by Transport & Environment reveals that the aviation sector receives €26.4 billion of indirect subsidies annually through tax breaks on VAT (€13.6 billion) and fuel tax exemptions (€10.7 billion).[39] This contrasts sharply with the operational environment of the major US carriers, which operate as privately-owned entities with limited government support. The European model of state involvement creates implicit incentives for maintaining certain routes that serve broader economic or political objectives beyond pure commercial viability.

The estimated means of non-hub routes and low-cost airlines' routes are considerably more likely to resemble the pure entry cost estimates, given the substantially fewer non-economic reasons to operate such routes. This analytical distinction provides valuable insights into the underlying cost structures across different airline business models. On this front, we find that full-service airlines generally exhibit

[38] Centreforaviation.com

[39] Transport & Environment

significantly higher pure entry costs compared to their low-cost counterparts, reflecting fundamental differences in operational complexity, service standards, and market positioning strategies.

Ryanair (FR) demonstrates the lowest pure entry cost amongst all airlines in our sample, which aligns perfectly with their well-documented ultra-low-cost operational model. In contrast, easyJet exhibits the highest entry costs amongst low-cost carriers, a finding that reflects their fundamentally different strategic positioning within the European aviation market. Unlike Ryanair's exclusive focus on secondary airports, easyJet deliberately operates from primary airports at considerably higher costs, positioning itself to compete directly with full-service airlines at their traditional hub airportsIndustry Report. This strategic choice necessitates substantially higher entry costs as easyJet must overcome the established advantages of incumbent full-service carriers whilst operating in more expensive airport environments. The airline's emphasis on higher frequency services and premium airport access creates natural barriers to entry that require substantial initial investment to establish competitive viability.

Regarding the reasonableness of our estimated pure entry costs, the figures demonstrate remarkably strong consistency with both established academic literature and comprehensive industry benchmarks. For instance, the estimated entry cost for British Airways on an average non-hub route approximates $0.5 million per quarter, translating to nearly $2 million annually. This magnitude aligns exceptionally well with industry cost structures documented in recent literature, including Yuan and Barwick [2024], and demonstrates robust consistency with comprehensive industry cost analyses conducted across multiple aviation markets.

Industry reports indicate that route establishment costs typically range between

\$1.5-3 million annually for full-service carriers on medium-haul European routes[40]. These estimates encompass not only direct operational costs but also substantial fixed investments required for viable route operations, including airport slot acquisition (exceeding \$500,000 for premium European airports), ground handling arrangements, marketing expenditure, and regulatory compliance costs. The convergence of our econometric estimates with these industry benchmarks provides compelling validation of our structural modelling approach.

1.6 Counterfactual Experiment on Carbon Policy

Carbon taxation has emerged as a pivotal policy instrument in the European aviation industry, with significant implications for airline operational costs and route economics. The current regulatory framework centres on the EU Emissions Trading System (EU ETS), which has experienced substantial price volatility and structural reforms in recent years. According to the International Emissions Trading Association, the average EU ETS carbon price is expected to rise from €84.4 per tonne during 2022-2025 to almost €100 per tonne during 2026-2030 (Statista). Critically, the system's application to aviation has been significantly strengthened, with 25% fewer free allowances allocated to aircraft operators in 2024, and complete removal of free allocation scheduled for 2026 (European Commission). This regulatory tightening ensures that airlines will face substantially higher carbon costs in the immediate future.

International organisations project even more dramatic carbon price escalations over the coming decades. Advanced modelling by Enerdata indicates that EU ETS prices will progressively increase after 2030, reaching around €130/tCO2 in 2040,

[40] IATA

before rapidly escalating to exceed €500/tCO2 by 2044 (Enerdata). These projections, spanning from approximately \$100 to \$500 per tonne over the next two decades, translate to substantial operational cost increases for airlines. For typical narrow-body aircraft operating intra-European routes, these carbon prices correspond to additional costs ranging from approximately \$1 to \$5 per kilometre flown, depending on fuel efficiency and carbon content assumptions.

Concurrent with carbon pricing pressures, the aviation industry faces mounting fuel cost challenges through two primary mechanisms. First, conventional aviation fuel supplies are increasingly constrained by environmental regulations and policy frameworks designed to reduce fossil fuel dependency. Second, mandatory sustainable aviation fuel (SAF) adoption requirements impose substantial cost premiums on airlines. Current market data indicates that SAF costs between two to seven times more than traditional jet fuel, whilst EASA's 2024 assessment shows conventional aviation fuel priced at €734 per tonne compared to aviation biofuels at €2,085 per tonne. Industry projections suggest that SAF prices will remain two to three times higher than conventional jet fuel until 2030(World Economic Forum), creating persistent upward pressure on airline fuel costs beyond carbon taxation effects.

Given these converging cost pressures from both carbon pricing mechanisms and fuel supply constraints, we implement five counterfactual scenarios that increase the Frequency $\times$ Distance coefficient by 1,000, 2,000, 3,000, 4,000, and 5,000 respectively. This parametric approach captures the combined effects of escalating carbon taxation and higher fuel prices within a realistic range of \$1-5 per additional kilometre flown. The lower bound reflects current EU ETS price levels with modest SAF adoption, whilst the upper bound corresponds to high carbon price scenarios with extensive SAF mandates. This specification provides a comprehensive framework for analysing how European airlines might adapt their route networks and pricing strategies under increasingly stringent environmental policies, thereby offering valuable insights for both industry stakeholders and policymakers navigating the transition to sustainable aviation.

What are the anticipated outcomes of route network deviation arising from increased Frequency $\times$ Distance coefficients? In the absence of competitive effects, airlines should exhibit a marginal preference for shorter routes over longer alternatives. This occurs because, holding optimal fares and expected variable profits constant, the increase in the linear component of fixed costs from larger coefficients imposes more severe penalties on longer routes compared to shorter ones. Given that frequency remains constant across all routes within the same choice set, the additional fixed cost can be expressed as: Extra coefficients $\times$ constant frequency $\times$ distance. In essence, maintaining frequency at constant levels, shorter routes generate fewer emissions and consequently incur lower additional emission-related costs.

However, airline competition also plays a crucial role in route selection decisions, as network deviations affect all airlines' expected variable profits from price competition across both origin and destination markets. The competitive interdependencies mean that when one airline adjusts its route portfolio in response to carbon pricing, this creates ripple effects throughout the whole market structure for the market it left and the market it entered, potentially altering the competitive dynamics on remaining routes and influencing rivals' strategic responses.

The empirical results presented subsequently demonstrate that whilst the absolute number of route changes remains relatively modest compared to the entire network, the economic implications are substantially more pronounced. Specifically, we observe significant effects on fare structures, passenger volumes, and expected profitability that extend far beyond the directly affected routes, highlighting the interconnected nature of airline network competition and the amplified welfare effects of environmental policy interventions in aviation markets.

Subsection 6.1 describes the methodology for drawing fixed cost shocks from the truncated distribution based on revealed preference constraints and our estimated distributional parameters. Subsection 6.2 introduces the equilibrium concept employed in the counterfactual analysis and details the iterative algorithm used to achieve such equilibrium. Subsection 6.3 presents the counterfactual results, examining how route networks evolve under different carbon tax scenarios, the resulting changes in fares, passenger volumes, and airline profits, and the overall welfare implications of the proposed carbon taxation policy.

1.6.1 Drawing Fixed Cost Shocks from Estimated Distributions

As detailed in the previous section, we draw realisations of the fixed cost shocks from their respective distributions. This represents a crucial methodological innovation that distinguishes our approach from all other research in the literature, which typically ignores fixed cost shocks in counterfactual simulations. The realisation of fixed cost shocks critically determines whether observed routes generate the highest net profit amongst all available alternatives. Ignoring these shocks would render the optimal network under our estimates substantially divergent from the actual observed network, thereby complicating the interpretation of counterfactual results and potentially undermining their policy relevance.

We draw fixed cost shocks whilst ensuring that all revealed preference constraints are satisfied. It is important to note that each observed route is associated with numerous alternative routes, and different observed routes may share common alternative options. We draw realisations of fixed cost shocks for each Quarter-Airline-Route tuple and apply the same realisation for identical tuples across different choice sets, encompassing both observed and alternative routes. In essence, we assume that fixed cost shock realisations remain constant across different decision-making contexts, provided the tuple specification is identical. Our dataset comprises 64,661 unique Quarter-Airline-Route tuples, which is substantially smaller than the total number of observed-alternative route pairs detailed in Table 6.

Operationally, we first draw realisations of all unique alternative Quarter-Airline-Route shocks $\hat{\kappa}_{j'}$ freely from their dedicated distributions. Subsequently, given all $\hat{\kappa}_{j'}$ values, we draw shocks for unique observed Quarter-Airline-Routes $\hat{\kappa}_{j^*}$ from truncated dedicated distributions where the lower bound satisfies:

$$\hat{\kappa}_{j^*} \geq \max\left\{\hat{\Pi}^{a}_{fj'} - \hat{\Pi}^{o}_{fj^*} + \hat{\kappa}_{j'}, -\hat{\Pi}^{o}_{fj^*}\right\} \tag{30}$$

Inequality (30) ensures that all revealed preference constraints hold under the realised shocks. This methodology guarantees that when we insert the original estimated coefficients into the counterfactual analysis, we recover precisely the observed route network, thereby rendering comparisons across different counterfactual scenarios both meaningful and empirically grounded.

We draw shock realisations separately for the lower bound, midpoint, and upper bound of the linear parameter estimates, providing robustness checks across the identified parameter space whilst maintaining consistency with our estimation framework.

1.6.2 Equilibrium Concept and Iterative Algorithm

As in the estimation section, the counterfactual experiment encounters similar computational difficulties as the number of possible deviations explodes with the number of markets. It is computationally infeasible to compute all $2^M - 1$ route networks with M representing the number of City-Airport combinations.

Unlike the equilibrium concept employed in the theoretical model, which requires optimality for all possible deviations, we impose a more restrictive equilibrium concept consistent with the literature by allowing only single-market deviations (Yuan and Barwick [2024], Jackson and Wolinsky [1996]). Although this restrictive equilibrium concept differs from the theoretical model, it remains consistent with the empirical approach used to estimate the parameters, thereby ensuring that counterfactual simulation results are directly comparable to the observed route network.

When calculating expected variable profits from price competition, we employ the same 36 draws of ξ and ω in counterfactual simulations as in the estimation stages to minimise numerical fluctuation from different draws of demand and marginal cost shocks.

Once we focus exclusively on single-market deviations, the order of evaluating Quarter-Airline-Route combinations becomes crucial for the final converged network. In practice, we first define a sequence of airline moves in each quarter and each airline's best response within each move. The network is updated following each airline move, with airlines moving sequentially and responding to their best response when facing the most up-to-date network in the sequence. The sequence is defined as follows.

We begin with the observed route network and calculate revenue for each Airline-Route in a specific quarter, where revenue equals optimal airfare multiplied by passenger numbers. We then calculate total revenue for all Airline-Routes in a given market (city pairs) and rank them from the highest-revenue (largest) market to the lowest-revenue (smallest) market. Airlines move first in larger markets, then progress to smaller markets. This ordering mimics real-world behaviour, as airlines typically prioritise larger markets over smaller ones and are likely to make decisions for larger markets first. Within each market, the airline with the highest combined revenue (largest) for all routes it operates in that market moves first, followed by smaller airlines within that market. Larger airlines in a market are likely established incumbents, whilst other smaller airlines are likely followers.

The route network changes continuously as airlines deviate to new routes. This affects both total revenue for a market and revenues for airlines still operating in that market during subsequent visits, potentially altering the revenue-based ordering for both markets and airlines within markets. In practice, we first rank markets by revenues from the starting network in each iteration. This order remains unchanged throughout the iteration, even if deviations alter revenue. In the first iteration, the starting network is the observed network. In subsequent iterations, the starting network represents the most up-to-date network after visiting all markets sequentially in the previous iteration. In other words, there is no *reversal* of market orders within one iteration. However, airline ordering within a specific market should be continuously updated by previous deviations. If, during visits to previous (larger) markets, airlines deviate to a (smaller) market, once we evaluate the smaller market subsequently, we must consider the previous deviation and rank airline revenue using the current market structure following previous deviations.

Airlines will exit the market completely if the best expected net profit in the current evaluation remains negative. In the original case where no additional carbon-related costs are imposed, there are no exits compared to the benchmark observed route network. At the other extreme, if carbon-related costs are infinite, airlines would exit all routes, regardless of sequential orders.

Deviations may also alter the set of feasible routes faced by each airline given our previous definition. The set of feasible routes is defined as airlines already serving both endpoint cities, with routes already served by at least one other airline. When airlines deviate, the total universe of served cities may change. Additionally, a route currently served by at least one airline may become empty if that airline deviates from or exits that route. We do not account for changes in the set of feasible routes to avoid further complicating our counterfactual exercise. Instead, we employ fixed sets derived from the observed route network. This approach is justified because if a city is currently served by an airline, or if a market is currently served by at least one airline, it is likely that this airline maintains some presence in this city, or that demand for the market is not negligible, even if both feasible route criteria fail in later counterfactual iterations. In essence, we always provide airlines with the largest possible feasible route network available when making choices.

Finally, since we employ exactly the same 36 draws of ξ and ω, we use pre-computed profits for both observed and alternative routes from the fixed cost estimation, provided the market structure[41] remains unchanged. This significantly accelerates computation, particularly in early iterations where deviations have yet to impact the broader network.

We perform counterfactual analysis separately for 4 quarters, 5 different carbon prices, and 3 sets of linear parameters θ_{fc}. In all cases, the algorithm converged, with convergence occurring on average around 5 iterations.

[41]This means that all products in one market (city pair), defined by unique Airline-Routes and their frequency, remain exactly the same for the current counterfactual network compared to the observed network in this market.

1.6.3 Counterfactual Results

Table 9 presents a comprehensive analysis of network changes under different counterfactual scenarios across quarters. We disaggregate the results by airline type, as they exhibit markedly different behavioural patterns in response to carbon pricing.

Total Routes denotes the total number of unique routes in our sample. Unsurprisingly, the summer peak season (Q2, Q3) features substantially more routes than the winter off-peak season (Q1, Q4). Low-cost carriers contribute the majority of unique routes, followed by full-service airlines and regional carriers. This distribution reflects the European aviation market structure, where low-cost carriers have significantly expanded their route networks following liberalisation.

Total Deviations measures the reallocation of aircraft to different routes by exiting current markets and entering new ones. The number of deviations increases monotonically with carbon cost intensity—from 40 under low carbon costs to 222 under ultra-high costs in peak season. Whilst deviations represent approximately 3–7% of total routes, their impact extends far beyond this percentage, as discussed below. Crucially, the burden falls disproportionately on regional carriers, who can account for over 120 deviations compared to just 8–54 for larger airlines. This fragility stems from regional airlines' longer average route lengths and thinner profit margins. In contrast, full-service carriers benefit from hub economies of scale and established market positions, whilst low-cost carriers operate higher-frequency, shorter-haul routes that are less carbon-intensive per passenger-kilometre.

Exit Number captures routes where airlines withdraw completely rather than redeploying aircraft. Pure exits increase substantially with carbon costs, rising from as few as 7 to as many as 51 under extreme scenarios. Remarkably, exits concentrate

Table 9: Network Analysis Results - All Carbon Scenarios

Metric	Low				Medium				High				VHigh				UHigh			
	Q1	**Q2**	**Q3**	**Q4**	**Q1**	**Q2**	**Q3**	**Q4**	**Q1**	**Q2**	**Q3**	**Q4**	**Q1**	**Q2**	**Q3**	**Q4**	**Q1**	**Q2**	**Q3**	**Q4**
Total Routes																				
All	2406	2868	2978	2523	2406	2868	2978	2523	2406	2868	2978	2523	2406	2868	2978	2523	2406	2868	2978	2523
Full Service	730	863	909	771	730	863	909	771	730	863	909	771	730	863	909	771	730	863	909	771
Low Cost	1276	1481	1520	1336	1276	1481	1520	1336	1276	1481	1520	1336	1276	1481	1520	1336	1276	1481	1520	1336
Regional	400	524	549	416	400	524	549	416	400	524	549	416	400	524	549	416	400	524	549	416
Total Deviations																				
All	86	100	40	84	120	150	60	100	136	176	62	116	166	194	86	144	176	222	110	188
Full Service	26	26	8	18	32	38	12	22	36	38	12	24	42	40	16	36	44	54	28	46
Low Cost	16	20	8	24	24	28	10	26	30	32	12	30	34	38	18	28	40	46	20	38
Regional	44	54	24	42	64	84	38	52	70	106	38	62	90	116	52	80	92	122	62	104
Exit Number																				
All	28	7	15	16	32	10	15	17	33	10	15	18	37	29	15	20	51	38	17	28
Full Service	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2
Low Cost	28	7	15	16	30	7	15	16	29	7	15	16	27	7	15	16	29	7	15	16
Regional	0	0	0	0	2	3	0	1	4	3	0	2	10	22	0	4	22	31	2	10
Total Markets																				
All	1548	1811	1867	1615	1548	1811	1867	1615	1548	1811	1867	1615	1548	1811	1867	1615	1548	1811	1867	1615
Full Service	659	783	828	702	660	787	829	704	661	786	829	706	662	787	829	708	663	790	830	708
Low Cost	1084	1258	1279	1126	1085	1263	1279	1126	1086	1264	1279	1128	1088	1265	1281	1128	1089	1266	1281	1132
Regional	363	466	484	381	370	473	488	385	373	478	491	389	377	490	496	396	376	491	497	403
Changed Markets																				
All	91	80	45	69	121	115	59	90	136	137	65	103	162	168	81	125	178	194	95	163
Full Service	23	24	8	15	29	35	11	19	34	35	11	22	39	35	14	34	42	49	25	44
Low Cost	43	26	22	37	51	34	24	41	54	37	26	45	55	41	32	42	62	47	34	52
Regional	39	47	18	36	59	72	30	46	67	90	34	58	90	121	47	74	101	134	55	104
Affected Routes																				
All	213	183	95	154	277	268	125	208	312	323	140	236	375	372	175	282	401	433	202	365
Full Service	29	28	6	15	37	35	10	16	39	40	10	17	47	38	13	34	52	54	24	43
Low Cost	49	27	23	36	60	34	25	44	64	39	27	47	62	41	36	43	69	47	39	52
Regional	37	42	17	32	51	63	26	41	58	81	29	51	85	99	41	61	96	113	49	86
Total Iterations																				
Counts	3	3	2	3	5	4	3	3	4	3	3	3	4	13	4	3	4	6	4	3

almost exclusively amongst low-cost carriers, with only scattered exits by regional carriers and virtually none by full-service airlines across all scenarios. This pattern reflects the thin operating margins inherent to the low-cost business model, as evidenced by the lower expected variable profits shown in Table 6. The near-complete absence of full-service carrier exits underscores the powerful role of hub networks and sunk investments in maintaining route viability even under extreme carbon pricing—a finding consistent with evidence that full-service airlines exhibit greater route persistence due to network effects and slot constraints at major airports.

Total Markets refers to unique city pairs served in our sample. *Changed Markets* denotes city pairs experiencing altered market structure between counterfactual and observed networks, where market structure encompasses all product offerings governing competitive dynamics. Regional airlines again show the greatest sensitivity: under ultra-high carbon costs, approximately 25–27% of regional airline markets experience structural changes, compared to less than 10% for full-service and low-cost carriers. This disparity reflects regional carriers' focus on thinner routes with fewer competitors, where the exit or entry of even a single airline fundamentally alters market conditions.

Affected Routes quantifies all routes in the observed network experiencing market structure changes. Revenue and profit or those routes change despite unchanged exogenous attributes through competitive spillovers. Notably, the affected routes can reach 10–20% of the total network—substantially exceeding the direct impact measured by total deviations alone. This multiplier effect demonstrates that carbon pricing's competitive consequences extend well beyond the routes directly restructured. Unlike previous metrics, affected routes distribute more evenly across airline types relative to their network sizes, suggesting that competitive interdependencies propagate throughout the network regardless of the type of airlines.

In summary, carbon pricing triggers cascading effects throughout the European aviation network, with regional carriers bearing disproportionate adjustment costs whilst full-service carriers demonstrate remarkable resilience. Crucially, the true economic impact extends far beyond the routes directly restructured: competitive spillovers affect a substantial portion of the network, with up to 433 routes experiencing altered profit conditions in peak season under ultra-high carbon costs. This multiplier effect—where market structure changes propagate throughout interconnected city-pair markets—highlights the critical importance of general equilibrium considerations in transport policy evaluation. Analyses focusing solely on direct route adjustments would severely underestimate the policy's full economic consequences.

How do ticket prices, passenger numbers, and airlines' net profits adjust when an additional carbon cost is imposed and the route network is re-optimised? Addressing these questions clarifies which types of routes shift for different airline groups and sets up the welfare analysis that follows. Table 10 reports the results for Q2, which we highlight because the Q2 network is the largest and therefore most informative; the other quarters display very similar patterns.

We report average fares separately for routes that are common to both the baseline and counterfactual networks and for routes that appear only in one network (i.e., non-common routes that are either newly added or dropped). Average fares on common routes remain stable across scenarios because the underlying market structure on those links—such as the set of active competitors and their relative positions—changes little, so second-stage pricing incentives are largely preserved. In contrast, for non-common routes, the pattern depends on airline type. For large full-service and low-cost carriers, the routes that are discontinued or replaced tend to have below-average fares, which suggests that re-optimisation targets links where

Table 10: Q2 Combined Analysis - Fares, Passengers, and Net Profits

Metric	Airline	Baseline	Low			Medium			High			VHigh			UHigh		
		Total	Common	Only BL	Only CF	Common	Only BL	Only CF	Common	Only BL	Only CF	Common	Only BL	Only CF	Common	Only BL	Only CF
Fares	All Airlines	0.92	0.92	1.10	1.03	0.92	1.11	1.03	0.91	1.12	0.99	0.91	1.13	0.99	0.91	1.14	1.00
	Full Service	1.35	1.35	1.31	1.24	1.35	1.30	1.25	1.35	1.31	1.19	1.35	1.31	1.26	1.35	1.32	1.23
	Low Cost	0.59	0.59	0.48	0.59	0.59	0.52	0.54	0.59	0.50	0.52	0.59	0.56	0.53	0.59	0.59	0.56
	Regional	1.15	1.15	1.36	1.10	1.15	1.29	1.09	1.14	1.30	1.07	1.14	1.27	1.05	1.14	1.25	1.05
		BL	**Total**	**Comm**	**Non**	**Total**	**Comm**	**Non**	**Total**	**Comm**	**Non**	**Total**	**Comm**	**Non**	**Total**	**Comm**	**Non**
Passengers	All Airlines	100.8	106.1	3.7	1.6	108.2	3.7	3.8	109.1	3.8	4.5	108.7	3.6	4.3	109.2	3.4	5.1
	% Change		**(+5.3%)**			**(+7.4%)**			**(+8.2%)**			**(+7.9%)**			**(+8.4%)**		
	Full Service	38.8	40.2	3.1	-1.7	40.2	3.1	-1.6	40.3	3.0	-1.5	40.1	2.9	-1.7	40.8	2.7	-0.7
	% Change		**(+3.6%)**			**(+3.7%)**			**(+3.8%)**			**(+3.3%)**			**(+5.1%)**		
	Low Cost	49.0	50.3	0.5	0.9	51.1	0.5	1.7	51.2	0.6	1.6	51.4	0.5	2.0	51.3	0.5	1.8
	% Change		**(+2.7%)**			**(+4.4%)**			**(+4.5%)**			**(+5.0%)**			**(+4.8%)**		
	Regional	13.0	15.6	0.2	2.4	16.9	0.2	3.7	17.6	0.2	4.4	17.2	0.2	4.0	17.1	0.1	4.0
	% Change		**(+19.9%)**			**(+30.1%)**			**(+35.5%)**			**(+32.2%)**			**(+31.5%)**		
Net Profit	All Airlines	88.9	85.8	-2.0	-1.1	82.7	-4.6	-1.6	79.8	-7.1	-2.0	77.0	-9.6	-2.3	73.8	-11.9	-3.2
	% Change		**(-3.5%)**			**(-7.0%)**			**(-10.3%)**			**(-13.4%)**			**(-17.0%)**		
	Full Service	33.2	32.0	-0.8	-0.5	30.6	-2.0	-0.7	29.2	-3.3	-0.7	27.9	-4.6	-0.7	26.4	-5.6	-1.2
	% Change		**(-3.8%)**			**(-8.1%)**			**(-12.1%)**			**(-15.9%)**			**(-20.4%)**		
	Low Cost	44.5	43.3	-0.9	-0.3	42.3	-1.8	-0.4	41.3	-2.7	-0.5	40.3	-3.6	-0.5	39.3	-4.5	-0.7
	% Change		**(-2.6%)**			**(-5.0%)**			**(-7.2%)**			**(-9.4%)**			**(-11.8%)**		
	Regional	11.2	10.5	-0.3	-0.3	9.9	-0.7	-0.6	9.3	-1.1	-0.8	8.8	-1.4	-1.1	8.1	-1.7	-1.4
	% Change		**(-5.9%)**			**(-11.8%)**			**(-17.3%)**			**(-21.6%)**			**(-27.4%)**		

Notes: Fares: Common = Routes in both baseline and counterfactual; Only BL/CF = Routes only in baseline/counterfactual. Passengers (millions) and Net Profit (10^8 USD: BL = Baseline; Comm = Common route difference; Non = Non-common route difference. % = change from baseline.

competition is stronger and price premia are limited—often because the routes do not connect hubs or major cities and hence cannot sustain higher markups. For small regional carriers, however, the deviated (dropped) routes typically have higher average fares; these links are frequently (near-)monopoly services connecting remote, long-distance pairs that are relatively costly to operate under higher carbon prices and that face thinner demand.

The mechanism linking carbon costs to fares operates through the network rather than directly through prices in our two-stage framework. In the second stage, pricing depends on the contemporaneous competitive structure of the realised network, not directly on costs. Consequently, higher carbon costs influence fares by altering which routes are profitable to operate, which then changes competitive intensity on the resulting network. This implies that higher carbon costs do not mechanically translate into higher pass-through to prices or lower total passenger numbers. Empirically, across counterfactuals, newly chosen routes exhibit higher average fares than common routes, but this difference does not necessarily grow monotonically with the level of the carbon cost. As carbon costs rise, some previously optimal links become unprofitable and exit; this selection margin widens the gap between the average fares of old and new routes, which is precisely what we observe.

Passenger volumes can increase relative to the baseline even when carbon costs are higher, because fares on common routes are nearly unchanged while the re-optimised network may attract additional demand on newly added links. In our results, total passenger numbers rise in all scenarios. For large carriers, the aggregate change is modest, reflecting the fact that their core, high-capacity networks remain largely intact. For regional airlines, the proportional increase is more pronounced, because their newly selected routes feature substantially lower average fares relative to the routes they discontinue, which stimulates demand. A decomposition confirms that most of the passenger growth for regional carriers originates from non-common (i.e., newly operated) routes.

Higher passenger numbers and nearly stable fares do not guarantee higher net profits, because fixed and carbon-related costs also increase. In the baseline, we estimate total net profits across all airlines at approximately $8.89 billion (USD). This magnitude aligns extremely well with industry evidence for 2019, which places European airlines' net profits at roughly €7 billion (EUROCONTROL) and is consistent with IATA's regional benchmarks (IATA).[42] Across counterfactuals, net profits decline for all airline groups and the losses increase with the carbon cost, as expected. Full-service and regional carriers experience the largest reductions, which is consistent with their longer average stage lengths—implying higher carbon exposure per link—and with the relatively lower average fares on the routes they select after re-optimisation.

How do the results in Table 10 translate into the welfare analysis? In Table 11, we report the change in airlines' net profit (producer surplus), the change in consumer surplus (measured as equivalent variation), and the carbon-related revenues paid by airlines—either to government in the form of a carbon tax or to fuel suppliers via higher prices for sustainable fuels. We then present the combined change in producer and consumer surplus, followed by the total welfare change obtained by adding carbon revenues to these surplus components. For completeness, we also report the change in daily flown distance, which provides a transparent proxy for daily carbon savings: flying fewer total kilometres implies lower emissions, holding aircraft technology and load factors fixed.

The results display several robust patterns. Net profit losses become more se-

[42] Differences stem from currency units (USD vs. EUR), data vintages, coverage definitions, and the fact that our aggregates are model-based.

Table 11: Q2 Welfare Analysis - Net Profit Changes, Consumer Surplus, Carbon Revenue, and Flight Distance

Metric	Baseline	Low			Medium			High			VHigh			UHigh		
	Total	Total	Comm	Non	Total	Comm	Non	Total	Comm	Non	Total	Comm	Non	Total	Comm	Non
Net Profit	88.9	-3.1	-2.0	-1.1	-6.2	-4.6	-1.6	-9.2	-7.1	-2.0	-11.9	-9.6	-2.3	-15.1	-11.9	-3.2
% Change		**(-3.5%)**			**(-7.0%)**			**(-10.3%)**			**(-13.4%)**			**(-17.0%)**		
Consumer Surplus	19.4	+0.97	+1.12	-0.15	+1.44	+1.61	-0.17	+1.60	+1.63	-0.03	+1.53	+1.61	-0.09	+1.74	+1.94	-0.19
% Change		**(+5.0%)**			**(+7.4%)**			**(+8.2%)**			**(+7.9%)**			**(+9.0%)**		
Surplus Change	–	-2.13	-0.88	-1.25	-4.76	-2.99	-1.77	-7.60	-5.47	-2.03	-10.37	-7.99	-2.39	-13.36	-9.96	-3.39
Carbon Revenue	–	3.11	–	–	6.18	–	–	9.21	–	–	12.08	–	–	14.84	–	–
Total Welfare Change	–	+0.98	–	–	+1.42	–	–	+1.61	–	–	+1.71	–	–	+1.48	–	–
Daily Flight Distance Δ	2,863.0	-29.0	–	–	-65.8	–	–	-91.8	–	–	-140.5	–	–	-188.7	–	–
% Change		**(-1.0%)**			**(-2.3%)**			**(-3.2%)**			**(-4.9%)**			**(-6.6%)**		

Notes: Net Profit and Consumer Surplus in 10^8 USD; Flight Distance in 10^3 km. Comm = Common route changes; Non = Non-common route changes (entry/exit effects). Surplus Change = Net Profit change + Consumer Surplus change. Carbon Revenue = [(1.25 + increment) × FD_counterfactual - 1.25 × FD_baseline] × 100/90 / 10^3, where increment is 1, 2, 3, 4, 5 for Low, Medium, High, VHigh, UHigh scenarios. Total Welfare Change = Surplus Change + Carbon Revenue. Flight Distance Δ shows FD_delta (counterfactual minus baseline). % = percentage change from baseline.

vere as the carbon cost rises, reflecting higher per-kilometre operating costs and re-optimisation away from previously profitable links. By contrast, consumer surplus increases in all scenarios, because newly entered routes tend to exhibit lower fares and attract additional passengers, consistent with the demand and pricing movements documented in Table 10. On balance, the combined consumer–producer surplus becomes more negative at higher carbon cost levels, largely because growing profit shortfalls outweigh consumer gains.

Carbon revenues are constructed from the change in total frequency–distance cost in each counterfactual relative to the baseline. Two forces operate simultaneously. First, the per-kilometre charge increases with the carbon price (or the renewable-fuel premium). Second, total frequency–distance adjusts as airlines enter and exit routes in response to profitability. As the per-kilometre cost rises, the revenue component increases mechanically, while network adjustments can either amplify or partially offset this depending on how total operated distance responds.

The most consequential finding concerns total welfare, defined as the sum of the surplus changes and carbon revenues. Once carbon revenues are included, the net welfare effect becomes positive rather than negative. Intuitively, the carbon charge reduces distortions on two margins. It prices the externality directly (the Pigouvian channel) and, through network re-optimisation, can temper mark-ups on links with pronounced market power—improving allocative efficiency even before counting the environmental benefits of lower emissions. This mechanism is consistent with established results on corrective taxation and the "double-dividend" discussion in environmental economics, where revenue recycling and competitive reallocation can yield welfare gains in already distorted markets. In the European airline context—where many routes are effectively monopolies or duopolies—this channel is particularly salient. The policy implication is that, provided the raised revenues are used productively (for example, to reduce other distortionary charges or to support efficiency-enhancing infrastructure), carbon pricing can deliver broad social gains in addition to its primary environmental objectives.

Finally, the overall impact of the carbon policy is distributed unevenly across European countries. Figure 12, Figure 13, and Figure 14 report the percentage changes in consumer surplus, airlines' net profits, and total welfare by country in Q2 under the UHigh counterfactual. To attribute route-level changes to countries, we weight each route's contribution by the population shares of its origin and destination cities and then aggregate to the country level. For cities that straddle national boundaries (e.g., Copenhagen/Malmö or Vienna/Bratislava), we split each measure evenly across the two affected countries to avoid double counting.[43]

[43] This allocation preserves country aggregates while remaining neutral with respect to cross-border functional city regions.

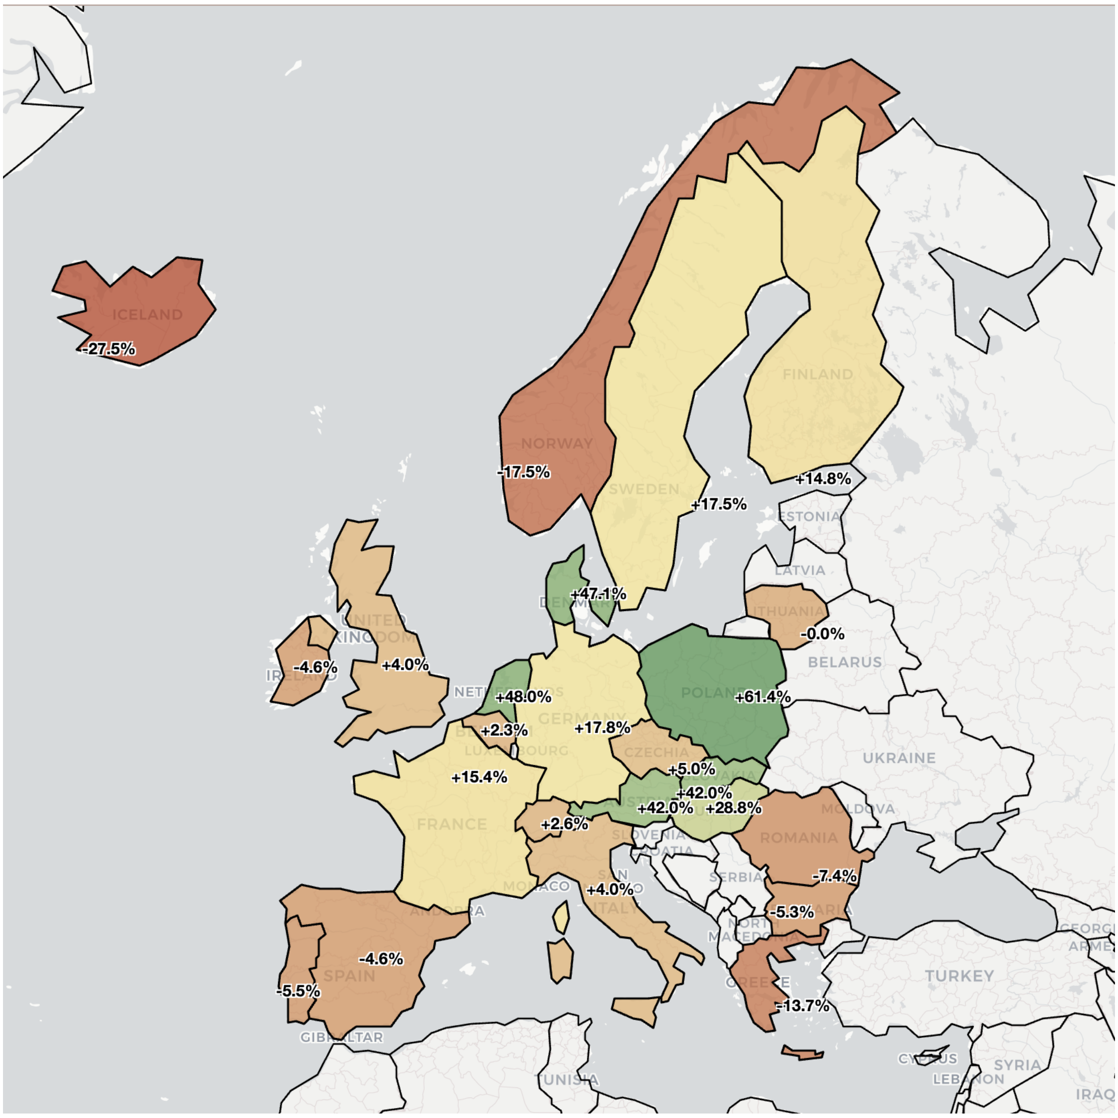

Figure 12: Change of Consumer Surplus by Countries (Q2, UHigh)

Consumer surplus increases most strongly, in percentage terms, in several Central and Eastern European countries such as Poland, Slovakia, and Hungary. Two forces account for this pattern. First, many of the newly entered routes under the carbon constraint are relatively short-haul links within the region, which typically sustain lower average fares and attract higher passenger volumes once the network is re-optimised. Second, these countries offer dense catchment areas with multiple viable secondary airports, expanding entry options and intensifying competition on newly operated links. By contrast, peripheral and island geographies such as Iceland, Norway, Greece, and Portugal experience declines in consumer surplus. Longer stage lengths in these regions raise carbon-related costs per flight, and the re-optimised network is more likely to cancel or deviate from thin, long-haul leisure routes; both the reduction in available links and the higher average fares on surviving routes depress passenger volumes. These directional effects are consistent with industry evidence that carbon exposure scales with distance and that periphery markets rely disproportionately on long sectors with limited substitution options.[44]

[44] For background, see industry discussions of distance-related carbon cost exposure and the resilience of short-haul, multi-airport networks in Europe in 2019–2023 reporting (e.g., EUROCONTROL network and market monitors; IATA regional outlooks).

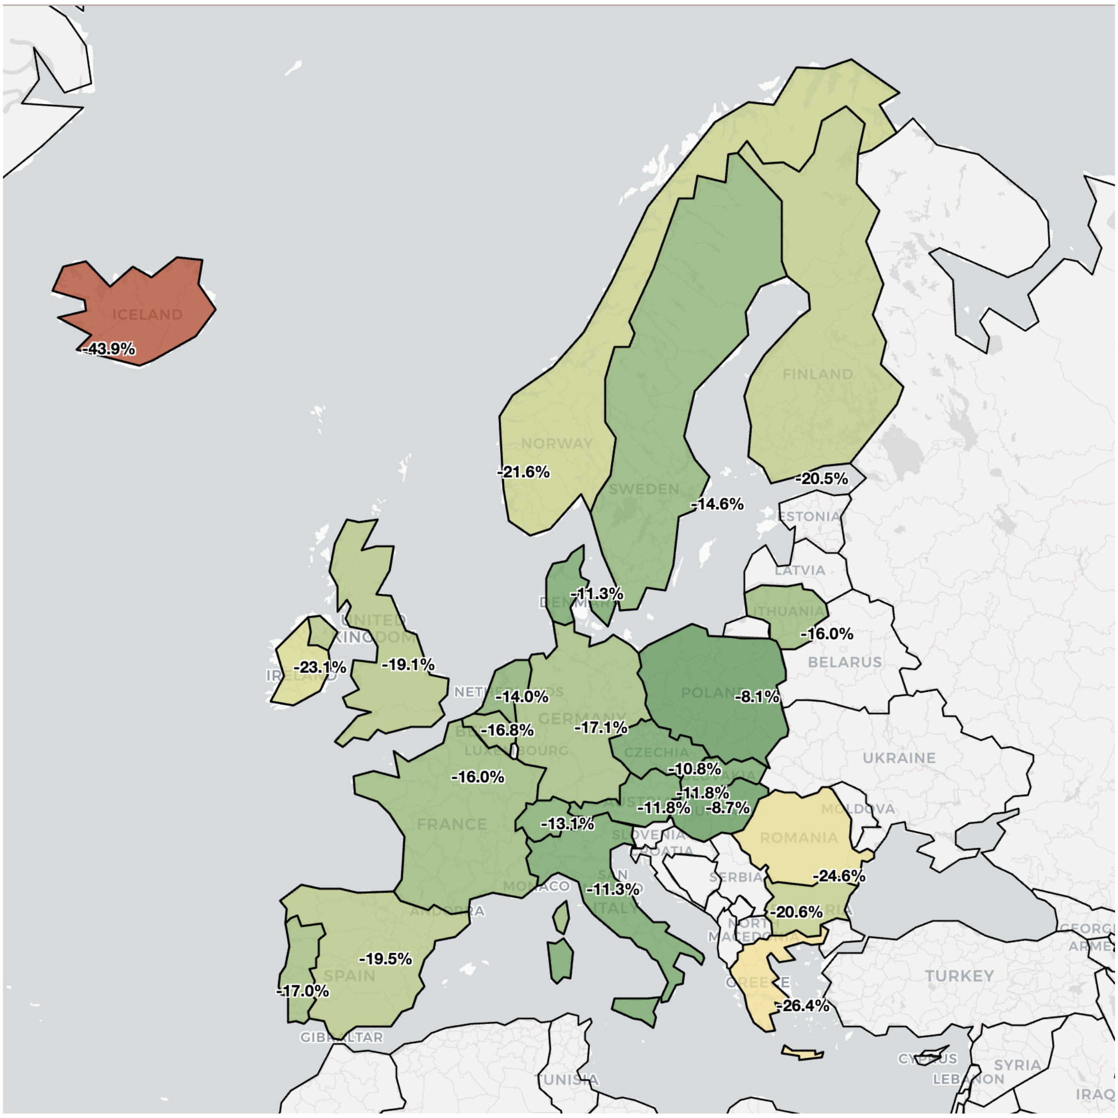

Figure 13: Change of Net Profit by Countries (Q2, UHigh)

Net profit changes in Figure 13 follow a similar geography. Peripheral countries see the largest percentage drops, reflecting both higher incremental carbon costs on longer average stage lengths and a greater incidence of route exits, which remove positive-contribution links from the portfolio. In Central and Eastern Europe, the declines are more muted because newly entered, demand-rich short-haul routes can offset part of the cost increase, and the shorter sectors imply a smaller per-flight carbon cost uplift. This asymmetry aligns with pre-existing carrier network strategies: ultra- and low-cost carriers have concentrated growth in Central/Eastern Europe using short-haul, high-frequency networks and secondary airports, while full-service and regional operators disproportionately serve longer or thinner markets where the fixed and carbon-related cost burden is harder to dilute.

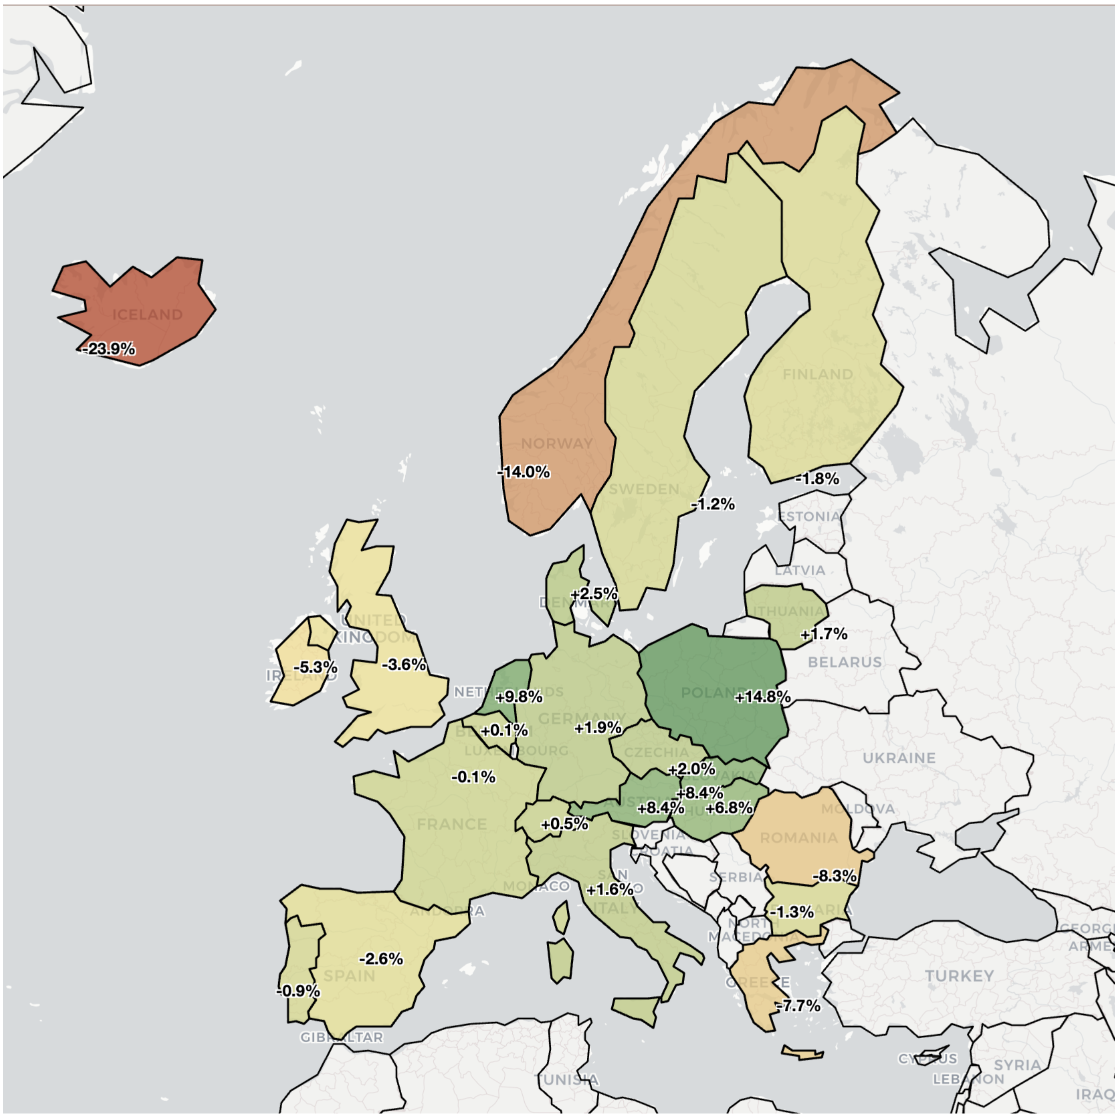

Figure 14: Change of Total Welfare by Countries (Q2, UHigh)

Total welfare changes mirror the joint behaviour of consumer surplus and producer surplus. Central and Eastern European countries emerge as net beneficiaries under the UHigh policy; for example, Poland records a projected total welfare gain of about 14.8%. In contrast, remote areas such as Iceland and Norway experience the largest welfare losses, driven by reduced network connectivity and higher average travel costs on retained routes. For countries hosting large hub airports—such as the UK, France, and Germany—aggregate welfare changes are comparatively small. Hub networks tend to preserve core trunk routes even under higher operating costs, which stabilises both prices and volumes on the common network segment; as a result, the national aggregates move little despite rebalancing at the margin across thinner spokes. Overall, these patterns are consistent with the notion that carbon pricing reshapes the extensive margin of route choice more forcefully in peripheral, long-haul-dominated markets than in central, short-haul-dense systems.

1.7 Conclusion

This paper develops a static, two-stage oligopoly model tailored to European short-haul aviation. In the first stage, airlines choose route networks and frequencies subject to capacity constraints at congested airports; in the second stage, they compete in multiproduct Bertrand pricing within each city-pair market. Demand is specified as a nested logit with airline, airport, and city fixed effects, and the estimated nesting and price coefficients imply own-price elasticities consistent with strong within-market substitution. The cost side features distance- and frequency-dependent marginal costs and fixed costs with a linear component and an unobserved shock that varies by business model and hub status. Linear fixed-cost parameters are set-identified using moment inequalities constructed from revealed-preference comparisons over feasible alternatives, while the full distribution of fixed-cost shocks is point-identified by maximum likelihood. This hybrid strategy allows us to draw shock realisations that rationalise the observed baseline network and to discipline policy simulations with the estimated primitives.

Counterfactuals impose an EU-ETS–style carbon cost proportional to operated distance (or fuel burn) and recompute network equilibria. To preserve a credible baseline, we draw route- and period-specific fixed-cost shocks from the estimated distributions and ensure that the observed network is an equilibrium under the estimated parameters. We solve for counterfactual networks using a sequential best-response algorithm in which airlines make single-market entry, exit, and frequency adjustments until no profitable deviation remains. Under higher carbon costs, aircraft are reallocated toward shorter, higher-density links; route exits are concentrated among low-cost and regional carriers, while full-service groups anchored at hub airports are comparatively resilient. Consumer surplus tends to rise where intensified competition lowers fares on newly operated or expanded links, while producer surplus falls; once carbon revenues are included, total welfare effects can be positive. Geographically, welfare changes are uneven: central and better-connected regions gain more, whereas peripheral and island geographies face larger risks from connectivity losses.

Substantively, the paper shows that carbon pricing is not merely a uniform cost shifter; by interacting with congestion, network choice, and heterogeneous business models, it reshapes competitive structure and the distribution of welfare. Methodologically, combining moment inequalities for linear fixed costs with likelihood-based recovery of shock distributions enables network-credible counterfactuals that limit baseline drift. Several extensions would further strengthen external validity if additional data—such as itineraries for international transfer passengers and airport-level slot allocations—were available. First, incorporating transfer passengers would better capture the economics of full-service hub operations, where connecting traffic is a major revenue driver. Second, replacing reduced-form capacity measures with slot microdata would allow explicit slot constraints at each hub. Third, allowing endogenous fleet dynamics would permit capacity to grow or contract in response to policy. Finally, expanding the feasible choice set beyond the current observed network to include currently unserved city pairs would test the robustness of entry margins under alternative expansion paths. These extensions are feasible within the present framework and, given appropriate data, would sharpen welfare and incidence conclusions for European aviation policy.

Chapter 2

Does The Subsidy on Electric Vehicles Reach Its Target? Evidence From The European Car Market

Abstract

This paper examines the key drivers behind the adoption of electric vehicles in the European car market. By combining annual car sales data with vehicle characteristics from the UK, France, and Germany between 2010 and 2021, along with detailed micro survey data from national transport surveys, I estimate a random coefficient logit model. The model incorporates interactions between time dummies, brand dummies, and fuel-type dummies to capture shifts in consumer preferences towards different fuel types over time. Counterfactual analysis reveals that subsidies have minimal impact on the change in EV market shares in the UK and France. Instead, the primary factor driving EV adoption is the evolution of the product lineup and vehicle attributes. Welfare analysis shows that consumers across different income levels are affected heterogeneously by these market changes. Finally, I propose an alternative income-based subsidy policy that achieves the same level of EV sales with less than half the current expenditure.

Keywords: Electric Vehicles, Purchasing Subsidy, Income-based Policy

JEL Codes: L52, L62, L90

2.1 Introduction

The automobile industry has undergone significant changes in the past decade, most notably with the sharp increase in market share for electric vehicles (EVs). According to the International Energy Agency (IEA), 14% of all new cars sold in 2022 were electric[45], up from around 9% in 2021 and less than 5% in 2020 (Global EV Outlook 2023, IEA). Three markets have contributed most to this surge in global sales. China accounted for nearly 60% of global electric car sales, with one in four new cars on the road now being electric. Europe is the second-largest market for EVs, with more than one in five new cars being electric. Given the recent spike in oil prices, largely due to the conflict in Ukraine, the global EV market is expected to continue its rapid growth.

Despite similar market shares, the specifics vary widely across regions. In Europe, EVs tend to be much more expensive compared to internal combustion engine (ICE) vehicles, such as petrol and diesel cars. Since 2015, the average price of EVs in Europe and the U.S. has risen significantly, whereas in China, EV prices have dropped by nearly half. According to JATO Dynamics, a reputable automotive industry research firm, the average EV in Europe cost just under €49,000 in 2015. By the first half of 2022, this had increased to almost €56,000, a rise of nearly 14%. As a result, an EV in Europe is now 27% more expensive than a petrol car. In contrast, the average EV in China cost nearly €67,000 in 2015, but this dropped to less than €32,000—a 52% reduction.

Another report by JATO Dynamics shows that in May 2021, EVs were on average 52% more expensive than ICE cars in the UK and 54% more expensive in the Netherlands. The relatively high EV prices in Europe are primarily due to produc-

[45] Electric cars include both battery electric vehicles (BEVs) and plug-in hybrid vehicles (PHEVs).

tion costs, as many of the raw materials needed for key EV components, such as batteries, must be imported from East Asia. Additionally, European car manufacturers are reluctant to fully transition from ICE vehicles to EVs, as they are hesitant to give up their competitive advantage in the ICE market. Even though the higher purchase price of an EV can be partly offset by lower running costs compared to ICE vehicles, owning an EV in Europe remains more expensive than in other parts of the world.

In reality, EVs in Europe have slowly become a luxury item, affordable mainly to wealthier consumers. Unless more affordable models are introduced soon, achieving net-zero emissions will be challenging, as the majority of vehicle usage comes from middle- to lower-income consumers.

Driven primarily by environmental concerns, governments around the world have implemented various policies to encourage the adoption of electric vehicles (EVs). Two of the most influential policies are subsidies for EV purchases and the planned bans on the development and sale of new internal combustion engine (ICE) vehicles in the near to medium future. The two most common forms of EV subsidies are attribute-based subsidies and flat subsidies. For example, China and Japan use attribute-based subsidies, where the subsidy amount increases with the driving range, though in different ways[46]. In contrast, most European countries offer a flat subsidy to EV buyers[47].

The rationale behind the decision of European countries to apply a flat subsidy across all EV models and to all potential buyers is not entirely clear[48]. Nevertheless,

[46]In China, this is a step-wise function, while in Japan it is a linear function.

[47]It is important to note that the subsidy amount is not entirely flat in Europe. For instance, there are specific criteria that must be met to qualify for the subsidy (usually based on emission levels or the retail price). However, the variation in subsidy amounts across different models is generally minimal.

[48]This may be due to capacity issues, where the cost of implementing a more flexible policy is prohibitively high.

the uniform subsidy has a varied impact on EV models with different attributes. More importantly, it creates disparate effects on consumers with differing preferences, often influenced by demographic factors such as income level and geographic location.

On the other hand, governments worldwide have committed to phasing out the sale of new internal combustion engine (ICE) cars in the near future. Norway, for example, aims to achieve 100% electric vehicle sales by 2025, setting one of the most ambitious goals globally. In 2020, the UK government made a historic step towards net-zero emissions by advancing the end date for new petrol and diesel car sales to 2030, a decade earlier than its previous target (GOV.UK)[49]. The European Union has also agreed to ban all new ICE car sales by 2035, despite initial resistance from Germany and some Eastern European countries (Fit for 55: zero CO2 emissions for new cars and vans by 2035, 14 February 2023).

Car manufacturers have had to align with these policies and significantly shift their development focus towards electric vehicles. For instance, Mercedes has announced that starting in 2025, all new vehicle platforms will be EV-only, while Volvo has committed to producing only EVs by 2030. This shift has been particularly challenging for European manufacturers, who previously promoted diesel cars as a cleaner alternative to petrol and invested heavily in that technology. However, the so-called "Dieselgate" scandal on September 18, 2015, revealed that Volkswagen Group's diesel vehicles were cheating on emissions tests. This shocking discovery highlighted that diesel vehicles were far more polluting than official figures suggested, even more so than petrol cars (BBC, 4 November 2015)[50]. As a result, public opin-

[49] Although this was revised in September 2023 to reset the date to 2035, the overarching goal remains unchanged. See https://www.bbc.co.uk/news/uk-politics-66871457.

[50] This could partly explain the high EV prices in Europe, as development efforts were focused on another technology for a considerable time.

ion on diesel cars has likely been influenced by the scandal, pushing consumers to seek alternative fuel options, whether petrol or electric.

Both subsidies and the ban on ICE cars could, in principle, facilitate the sale of electric vehicles (EVs). However, the key question remains: which factor contributes most to the increase in EV market share? Understanding the primary driver behind EV adoption is crucial because it affects how consumer surplus varies across different groups. If the ban on ICE cars is the main driver, the reduced availability of ICE options could force some consumers to purchase expensive EVs, especially if they need a car. This is likely to negatively impact consumer welfare, particularly for poorer households, who often hold unfavourable views of electric vehicles. According to the Society of Motor Manufacturers and Traders (SMMT), more than half of EV owners were in the top 20% of earners, while those in the lowest two income brackets accounted for just 4% of EV owners in the first nine months of 2020.

On the other hand, if monetary incentives are the main driver, consumer surplus is likely to increase, albeit at the expense of the social planner's surplus. However, the heterogeneous preferences among consumer groups mean that the welfare gains from the flat subsidies used in Europe are unevenly distributed. Wealthier consumers, who might have purchased EVs even without subsidies, benefit more, while lower-income consumers may still find EVs too expensive, even with subsidies. This creates a twofold issue: first, higher-income consumers gain disproportionately under the current policy; second, subsidies are effectively wasted on wealthier consumers who would have bought EVs regardless, diverting resources that could be better used elsewhere. An income-based subsidy policy might be a more efficient way to achieve the same targets while better preserving consumer welfare and reducing government spending.

This paper aims to address three key research questions: 1.What is the main driving force behind the observed uptake of EVs? 2.How is consumer welfare affected differently across various income groups? 3.Can we design a more efficient policy that achieves the same goals while better protecting consumer welfare?

A carbon-free car industry is unsustainable if EVs remain the domain of affluent consumers while poorer consumers are forced to switch due to the elimination of alternatives. Wealthier consumers, typically living in urban areas, use their cars less frequently than lower-income consumers in rural areas, who have fewer public transportation options[51]. Achieving a carbon-free car market is impossible without the participation of lower-income consumers, who should not be compelled to switch solely through restrictive measures.

The discrete choice model is commonly used to estimate demand parameters in this type of research. Three specific models are prevalent in the literature: the nested logit model (Berry [1994]), the random coefficient logit model (Berry et al. [1993], hereafter BLP), and the random coefficient nested logit model (Brenkers and Verboven [2006]; Grigolon and Verboven [2014]). As Grigolon and Verboven [2014] point out, both the nested logit model and the random coefficient logit model are special cases of the random coefficient nested logit model. The differences between these models stem from how they account for consumer heterogeneity. In the nested logit model, heterogeneity is defined by nests, where residuals of products within the same nest are correlated. In the random coefficient logit model, heterogeneity is captured by random coefficients on product characteristics. The random coefficient nested logit model combines these two approaches. The distinction between the nested logit model and the random coefficient logit model lies in their different distributional assumptions. As discussed by Cardell [1997], the nested logit model

[51]This context applies to Western Europe, with its high population density and strong public transportation networks. The situation differs in other regions, such as North America.

can be viewed as a special case of the random coefficient model, where the random coefficient applies to group dummy variables forming the nests, under a specific distributional assumption to ensure the overall residual follows an extreme value distribution.

In this paper, we use the random coefficient logit model, while incorporating potential nesting variables into the non-linear terms. This approach captures much of the heterogeneity that the nested logit model would account for, through the use of random coefficients, as highlighted by Grigolon and Verboven [2014]. It is well known that models of this type can encounter endogeneity issues, as price and other characteristics may be correlated with unobserved (to the econometrician) structural errors. To address this, we employ instrumental variables (IVs) from Gandhi and Houde [2019], building on the IVs introduced by Berry [1994][52].

On the supply side, profit-maximising firms are assumed to engage in static Bertrand competition, setting prices for each product. Marginal costs are assumed to be correlated with product characteristics and are jointly estimated with the demand side, following standard practice in the literature. I have chosen not to introduce additional dynamic elements, such as product entry and exit, endogenous product attributes, or supply-side investments, as this paper primarily focuses on demand-side policy design and welfare.

The main finding shows that growth in EV market share is driven primarily by changes in the product portfolio, notably the introduction of new models and improvements in attributes, rather than by uniform purchase subsidies. Subsidies play a limited role in the UK and France and a larger one in Germany. Income also shapes preferences for EVs. Lower-income households are more price sensitive and,

[52] Other potential IVs are also used in the literature. For example, Berto Villas-Boas [2007] uses input prices interacted with product dummy variables as instruments, while Hausman [1975] and Nevo [2001] rely on prices of the same product in other markets.

on average, place a lower intrinsic value on EVs than higher-income households. Counterfactual exercises indicate that an income-targeted subsidy can match or exceed the sales achieved under a flat scheme at substantially lower fiscal cost, while shifting benefits toward lower-income households and improving the overall distribution of welfare.

2.2 Literature

This study contributes to three key areas of literature. First, it adds to the research on public policy for promoting electric vehicle (EV) adoption, as reviewed by Rapson and Muehlegger [2023]. Existing studies, such as those by Linn [2022] and Xing et al. [2021], examine income-targeted subsidy designs in the US, showing that these policies are more cost-effective. However, there are notable differences between the US and European markets that influence policy outcomes. In the US, credit systems for zero-emission vehicles (ZEV) and vehicle greenhouse gas (GHG) emissions play a significant role, whereas these systems are absent in Europe. As a result, subsidies in the US act more like a consumption tax credit, whereas in Europe, they tend to reduce EV prices by lowering marginal costs[53]. Additionally, European countries have set firm deadlines to phase out internal combustion engine (ICE) cars, which has led to the discontinuation of new ICE models—a major factor driving the market shift towards EVs in Europe, a shift that is less pronounced in the US.

This paper extends the analysis of how subsidy policies impact different consumer groups, particularly in terms of income. For instance, Muehlegger and Rapson [2022] investigate the California Enhanced Fleet Modernisation Programme

[53] For example, Linn [2022] suggests that subsidies in the US may increase EV prices as the credit price falls. In Europe, however, subsidies drive down EV prices by reducing marginal costs.

(EFMP), which provides retire-and-replace subsidies for EV purchases, focusing on low- and middle-income households. However, their analysis is limited to the EV market, whereas this paper considers the entire new vehicle market in Europe. This paper also explores the interaction between demand-side policies, such as subsidies, and supply-side dynamics. Armitage and Pinter [2021] study the effects of supply-side policies, like zero-emission mandates, alongside demand-side incentives in the US. While their work compares these different policy approaches, this paper focuses specifically on optimising demand-side policies, given the existing flat subsidy structure in Europe. Finally, Remmy [2022] addresses the indirect network effects between EV sales and charging infrastructure. Although Remmy compares the impact of purchase subsidies and charging station subsidies in Germany, this paper delves deeper into how purchase subsidies affect consumer heterogeneity, particularly in terms of income, across multiple European countries[54]. To my knowledge, most existing research on EV adoption focuses on the US, Chinese, or Scandinavian markets, with the exception of Remmy [2022], who studies Germany. This paper, however, broadens the scope by comparing policy effects across the UK, France, and Germany, providing a more comprehensive understanding of EV market dynamics in Europe.

Second, this paper contributes to the broader research on environmental regulations in the automobile industry. Reynaert [2021] examines European emission standards for cars and finds that the policy fell short of its goals, as firms resorted to technological adaptations and manipulation of emission tests to reduce emissions, ultimately failing to improve overall welfare. This paper extends this analysis by evaluating the changes in vehicle fleets following the ban on ICE cars, which can

[54]While Remmy examines EV range choices on the supply side, this paper focuses on income heterogeneity on the demand side, which presents a different challenge in equilibrium analysis.

be seen as the most extreme form of emission regulation. Other relevant studies include Ale-Chilet et al. [2021], who explore collusion among firms in response to imperfectly monitored environmental regulations, as well as research by Anderson and Sallee [2011], Anderson and Sallee [2016], Holland et al. [2009], Klier and Linn [2012], Whitefoot et al. [2017], Whitefoot and Skerlos [2012], and Ito and Sallee [2018].

Third, this paper also contributes to the literature on the impact of the 'dieselgate' scandal. Bachmann et al. [2019] treat the scandal as a natural experiment, finding that it damaged the collective reputation of both diesel cars and German manufacturing, even for products that passed emission tests after the scandal. This indicates a shift in consumer preferences away from diesel cars. This paper delves further into this topic, demonstrating that the negative effects on diesel cars have primarily shifted consumer demand towards petrol cars rather than electric vehicles[55]. Other studies, such as those by Strittmatter and Lechner [2020], Che et al. [2018], and Ater and Yoseph [2022], examine the scandal's impact on Volkswagen vehicles in the used car market. Additionally, Griffin and Lont [2018] and Barth et al. [2022] investigate the financial market repercussions for VW and other major automakers. However, this research focuses exclusively on the new car market.

The empirical analysis yields several key findings. First, the estimation results show that the correlation between income level and preference for EVs varies across regions. Specifically, French consumers generally prefer non-BEVs over BEVs, though this preference gap narrows as income increases. In contrast, German consumers show a stronger preference for BEVs, with this disparity widening at higher income levels. UK consumers, on the other hand, tend to favor EVs, but there is no significant variation across income groups. These results are derived while

[55]Though the impact varies across different consumer groups.

controlling for all other characteristics and prices. Borenstein and Davis [2016]) document that high-income households were significantly more likely to be early adopters of EVs, which aligns with the findings for France and Germany. Several factors, such as maintenance costs, the second-hand car market, and financial loans, could explain this heterogeneous preference. However, the presence of such preference heterogeneity highlights the inequality in EV access, underscoring the need for alternative policies that efficiently target specific consumer subsets.

Second, the counterfactual analysis indicates that the subsidy policy is most effective in Germany, contributing up to an 11% increase in EV market share. In contrast, the subsidy only results in a 2% increase in EV market share in the UK and France[56]. In the UK, observed and unobserved characteristic changes are the primary drivers of EV uptake, while in France, fleet changes play the dominant role. In Germany, EV adoption is evenly driven by all changing factors. Across all three countries, however, the decline in diesel car sales is primarily due to fleet changes rather than changes in characteristics, preferences, or subsidy policies, indicating that the 'dieselgate' scandal had only a minor effect on diesel car sales compared to the ban on ICE cars.

Third, the welfare analysis reveals that overall consumer surplus increased in all three countries under the subsidy scheme. However, these benefits are not evenly distributed across different income groups. In France, consumers in the lowest two income groups actually experienced a decrease in surplus, while in the UK, the per-consumer surplus increase for the highest income group is nearly six times greater than that for the lowest income group. This raises important questions about how to design a more equitable policy that maintains the trend of EV uptake while

[56] A similar result is found in Archsmith et al. [2022], which concludes that by 2035, EV market share will be determined more by non-monetary factors than by subsidies.

addressing income disparities. I propose an alternative income-based subsidy policy to better target the consumers who most need EVs.

Fourth, I designed and evaluated an alternative income-based subsidy policy. Making EVs more affordable for lower-income consumers is crucial for achieving net-zero targets. In the counterfactual scenario, I demonstrate that an income-based policy could achieve higher EV market shares and greater consumer surplus using the same budget. Additionally, the total subsidy spending required to reach the same EV market share is substantially lower under an income-based subsidy, suggesting that such a policy would benefit both consumers and the government.

The remainder of the paper is organised as follows: Section 2 provides a detailed description of the industry and data. Section 3 presents the model setup, estimation methods, and results. Section 4 discusses the counterfactual results. Section 5 compares the current flat subsidy with a flexible subsidy based on income levels. Section 6 concludes.

2.3 Background and Data

2.3.1 Industry Background

The global sales of electric vehicles (EVs) have surged over the past decade. Before 2010, EVs were marginal products, often used as prototypes, with a negligible global market share. However, thanks to technological advancements, several popular EV models were introduced around 2010, including the Nissan Leaf (2010), Vauxhall Ampera (2011), Renault Zoe (2012), and Tesla Model S (2012). Additionally, growing environmental concerns and the ripple effects of the 'dieselgate' scandal in 2015 likely accelerated adoption. Governments worldwide have also implemented generous subsidies to boost EV sales. As a result, 10.2 million EV units were sold in 2022, accounting for 14% of global vehicle sales. However, electrification remains highly concentrated in China, Europe, and the United States, as EV supply requires an electric distribution grid and sufficient electric generation capabilities. China represents the largest market share, with 57.8% of global EV sales in 2022, driven by tax incentives and exemptions from local vehicle sales quotas for EVs. Europe follows as the second-largest market with 25.4% of sales, although there are significant variations across countries. Norway, rich in electricity, boasts the highest EV sales share globally, with 72% of all new vehicle sales being electric. In comparison, the figure is around 25% for Germany and 15% for France and the United Kingdom. Understanding the reasons behind the increase in EV market share is vital for designing better policies to accelerate EV uptake without forcing consumers to adopt EVs by eliminating ICE options, especially in less affluent countries.

Regarding EV sales in Europe by manufacturer, the Volkswagen Group holds the largest market share (19.6%) as of April 2023, with popular models like the ID.3 and ID.4. Stellantis follows with a 13.9% share, thanks to its compact models such as the Fiat 500e and Peugeot 208 EV. Tesla ranks third with 12.6%, the only non-European manufacturer in the top five. The Tesla Model Y is currently the best-selling EV in Europe, with over 10,000 units sold.

According to the European Alternative Fuels Observatory (EAFO), the combined EV sales in the UK, France, and Germany account for over 50% of all EV sales in Europe[57][58]. In terms of market share, EVs constitute 2.94% of the total passenger cars on the road in the UK, 2.58% in France, and 3.78% in Germany. Although

[57]Europe refers to the EU27 plus the UK, Norway, Iceland, Switzerland, Turkey, and Liechtenstein.

[58]The total numbers of electric passenger cars in the UK, France, and Germany are 1,094,821, 1,174,187, and 1,967,099, respectively.

these figures might seem modest, the new car registration data for 2022 reveals that 19.24% of new car sales in the UK were electric. The corresponding figures for France and Germany are 20.48% and 18.62%, respectively.

Meanwhile, the variety of car models with different powertrains has changed dramatically over the last decade. According to the International Energy Agency (IEA) Global EV Outlook, the number of available EV models reached 500 in 2022, up from below 450 in 2021 and more than double the figure from 2018-2019. At the same time, the number of ICE car models has steadily decreased, with a compound annual growth rate of minus 2% over the 2016-2022 period, bringing the total to about 1,300 models in 2022. In Europe, the number of available ICE options was 8% lower in 2022 than in 2016, largely due to the introduction of more EV models and the phasing out of ICE models. More details will be provided in the data description subsection.

EVs are increasingly likely to be SUVs, with nearly 40% of all BEV models being SUVs[59], which is 10% higher than the average percentage of SUVs across all models. Moreover, EVs are much more expensive in Europe than in other major markets like China and the US. In China, the best-selling electric cars in 2022 were the Wuling Mini BEV, a small model priced at under USD 6,500, and BYD's Dolphin, another small model priced below USD 16,000. Together, these two models accounted for nearly 15% of Chinese BEV passenger car sales, illustrating the demand for smaller models. In contrast, the best-selling small BEVs in France, Germany, and the United Kingdom – Fiat's 500, Peugeot's e-208, and Renault's Zoe – were all priced above USD 35,000 (Global EV Outlook 2023, IEA). The correlations between fuel type, body type, brand, and prices, which are evident in the dataset, create an endogeneity problem in estimation. To address this issue, several dummy variables

[59] From an engineering perspective, it is also easier to develop an EV as an SUV.

and their interactions are included in the regression to prevent omitted variable bias, which is typically overlooked in non-EV-focused automobile literature (for example, Berry, Levinsohn, and Pakes (1995)).

2.3.2 Policy Background

Governments in many countries have implemented policies to boost EV sales. Some of the earlier measures include subsidies for new EV buyers, grants for scrapping old vehicles, reductions in vehicle registration fees and taxes, public investments in charging infrastructure, and subsidies for home chargers. Additional measures include the establishment of clean city zones. More recently, policies have shifted towards the supply side of the industry, with investments in battery technology, fast charging infrastructure, and more[60]. In more developed markets, such as China and several European countries, governments are progressively reducing or phasing out incentive schemes for electric cars, shifting focus towards other sectors like heavy transport and charging infrastructure[61]. This raises questions about the effectiveness of subsidies and provides useful insights into when and how subsidies should be implemented in emerging EV markets.

Figure 15 presents the direct consumer subsidies for electric vehicles in the UK, France, and Germany from 2010 to 2022. The UK and France were early adopters of subsidies, introducing them as far back as 2010. In contrast, Germany began subsidising EV purchases much later, in 2016[62]. However, the amount of subsidy in UK and France has decreasing after 2016. Both countries put a price cap to

[60]Industry-wide initiatives include the Inflation Reduction Act (IRA) in the United States and the Green Deal Industrial Plan in the European Union.

[61]For example, the UK government phased out the plug-in grant for electric cars on 14 June 2022 (GOV.UK)

[62]This delay may be attributed to the influence of powerful German car manufacturers and their lobbying efforts, which were less prominent in the UK and France.

qualify the subsidy and ruled out the subsidy for PHEV completely over the years[63]. Although being late on early subsidising, Germany actually increased their subsidy amount recently. Also, by the end of 2022, Germany is the only country which still subsidising the PHEV[64]. As a result, I find that Germany is the only country where the subsidy on EV actually makes a significant impact on the market share after the 'dieselgate' scandal. The amount of the subsidy for the UK and France is quite similar. However, UK decreased the amount of the subsidy at a much radical pace.

2.3.3 Data (Car Sales)

The data used in this paper are drawn from several sources. I obtained monthly new car registration data from IHS Markit, covering the period from January 2010 to September 2022. This dataset includes key car characteristics such as manufacturer, MSRP (in local currency and USD), horsepower (HP), body type, and fuel type. To make the estimation more tractable, I processed the raw data as follows:

First, I aggregated trim-level observations[65] to model-level observations using sales-weighted averages, as is common in the literature (e.g., Grigolon and Verboven (2014); Grieco, Murry, and Yurukoglu (2021)). Focusing on the model level simplifies the logit estimation by reducing the number of products in the analysis. I also aggregated the monthly sales data to an annual level.

Second, I removed car models with very high prices (greater than $150,000) and very low sales (fewer than 200 units per month). Extremely expensive models are likely to be luxury goods, which can distort price elasticity estimates, potentially

[63]This is not shown in the plot but the price cap drops over the years.

[64]Part of the reason of the increased subsidy comes from the effect of the pandemic. See Die Bundesregierung (The Federal Government) (2020-09-22). "Climate-friendly transport: Promoting the conversion to electric mobility". Cabinet of Germany. Retrieved 2021-01-25.

[65]In the automobile industry, "trims" refer to different versions of a model, such as those with leather seats or premium sound systems.

resulting in positive elasticity. Similarly, low-sales models can create unusual odds ratios, making logistic regression results less reliable.

Third, I grouped car manufacturers by their parent companies (e.g., Volkswagen Group) rather than individual brands (e.g., Audi), as profit maximisation often occurs at the parent company level. Models from brands within the same parent company often share key components, such as chassis and engines. For example, the Audi A3 and Volkswagen Golf share the same platform (MQB) and engine choices (EA888). Parent companies typically make joint decisions for their subsidiaries. A few mergers and acquisitions occurred during the sample period: PSA acquired Vauxhall and Opel from General Motors in 2017, and PSA and FCA merged to form the Stellantis group in 2022. I assigned brands to their parent companies based on the time of the sample observation to minimise distortion in competition[66]. After this processing, the dataset includes approximately 20 firms.

Fourth, I created a new dummy variable to indicate whether a brand is considered luxurious. Brands such as Audi, BMW, Mercedes, Subaru, Volvo, Tesla, Jaguar, and Land Rover are classified as luxury brands, while others are considered affordable[67]. This classification follows industry norms and, more importantly, the sales-weighted average prices of luxury brands are much closer to each other compared to those of affordable brands, suggesting that luxury brands are more likely to be close competitors. Introducing the luxury dummy is crucial for several reasons. First, it addresses important endogeneity issues because it correlates with prices, body types, and fuel types, as shown in the summary statistics later. Second, the luxury dummy correlates with income levels, reflecting consumers' varying

[66]I treated PSA and FCA as separate entities for the entire sample period, as the merger was not fully completed by the end of 2022.

[67]For example, Audi is classified as a luxury brand, whereas other brands within the Volkswagen Group, such as Skoda, are classified as affordable.

preferences for brand premiums.

Figure 16 shows the evolution of market share by fuel type in the UK, France, and Germany, based on the disaggregated monthly data. Overall, the market share for diesel cars has decreased sharply. Before the 'dieselgate' scandal in September 2015, the market share for diesel cars was relatively stable in all three countries. However, the scandal damaged the reputation of diesel vehicles and raised public awareness about the environmental and health impacts of diesel emissions. Consequently, the market share for diesel cars declined significantly after 2015. In the UK, the market share fell from 45% in 2015 to less than 10% in 2022. In France, it dropped from 50% to less than 15% in 2022. Germany experienced a relatively moderate decline from around 40% to 20%. Such a drastic change in market structure within less than a decade is rare in any industry, raising an important question: what are the underlying reasons for this shift?

It is unclear whether the decline in diesel car sales is attributable to changes in subsidies, shifts in consumer preferences, alterations in the product list and car characteristics, or a combination of these factors. In the early years following the 'dieselgate' scandal, there was no significant change in the market share of electric vehicles (EVs), which remained just above 3% by the end of 2019 in all three countries. However, from 2020 onwards, the market share of EVs surged, reaching approximately 35% of new car sales in 2022. These observations prompt several intriguing questions. First, the impact of subsidies is contentious, as both the UK and France significantly reduced their subsidy amounts after 2019. Second, it is uncertain whether the reduction in diesel car shares is translating into increased EV sales, given that the shifts in market shares for diesel cars and EVs do not occur simultaneously. This suggests that different factors may be influencing the market changes for diesel cars and EVs.

Table 12 and table 13 present the summary statistics for different subsets of the disaggregated data across all three countries. Specifically, I examine the key characteristics across the entire sample, as well as within SUV, luxury, and EV sub-samples. There are strong correlations between SUVs, luxury brand models, and EVs in all three countries. SUVs and luxury brand models each account for approximately 30% of all models, while EV models make up 20% of the total models in the full sample. However, the proportion of EV models is 5 to 10 percentage points higher in SUV sub-samples and 10 to 20 percentage points higher in luxury sub-samples compared to the full sample. Similar patterns are observed with luxury models within the SUV and EV sub-samples. Given these strong correlations, it is crucial to include all three dummy variables in the model to avoid omitted variable bias. Additionally, SUVs, luxury brand models, and EVs generally have higher prices and horsepower on average, which addresses another potential endogeneity issue related to price.

Figure 17 illustrates the change in the total number of models by fuel type. It is evident that the number of diesel car models has significantly declined since 2016. In the UK, the number of diesel car models peaked at 120 in 2016 and decreased to just 30 models by 2022, representing a reduction of nearly 75% over this period. Similar trends are observed in France and Germany. Conversely, the number of EV models has increased from fewer than 50 in 2016 to 150 in 2022. By around 2020, the number of EV models surpassed the number of diesel car models, with the gap widening over time. The number of petrol car models has remained relatively stable throughout the sample period. In a logit model framework, the reduction in diesel car options has forced consumers to switch to petrol cars, EVs, or other alternatives, even if some of these options are less desirable. The significant change in the product list may have a substantial impact on consumer welfare, though whether it will result in an increase or decrease remains uncertain.

In conclusion, the sales data shows a sharp increase in the EV market share and a decline in diesel car sales over the sample period. The summary statistics also reveal strong correlations among SUVs, luxury brand models, and EVs, underscoring the need to include dummy variables in the analysis. Furthermore, the significant growth in the number of EV models, coupled with the shrinking number of diesel car models, suggests additional factors driving market changes beyond subsidy policies, changes in characteristics, and shifts in consumer preferences.

2.3.4 Data (Socio-Demographic)

I use micro-survey data on individual purchasing decisions to highlight the importance of consumer heterogeneity. The survey data for the UK is sourced from the National Travel Survey (NTS) conducted by the Department for Transport. This survey is conducted annually and links purchasing behaviour with income levels[68]. Each year, approximately 3,700 individuals from 7,000 households in England participate in the survey. Since the survey focuses on existing car ownership rather than new car sales, I use the 2021 survey results to represent the most current market accumulation. In contrast, public sources on individual purchasing decisions in France and Germany are less accessible. One notable source is the European Alternative Fuels Observatory (EAFO) Consumer Monitor & Survey (European Commission, 2022), which was conducted across the EU27 countries.

The survey, launched in September 2022, provides up-to-date information on accumulated stock. In France, the survey received 1,703 valid responses, including 58 BEV drivers and 1,645 non-BEV drivers. In Germany, the survey included 1,648

[68]In the survey, consumers are categorised into five income quintiles. Individual income levels are not directly observed.

observations, with 57 BEV owners and 1,591 non-BEV owners. A typical French BEV driver is characterised as a male aged 35 or younger, living in a detached house, with a monthly income between €2,000 and €3,999, and holding a university degree or higher. A typical German BEV driver is a male aged 33 to 55, residing in a detached house, with a monthly income between €2,000 and €3,999, and also holding a university degree or higher (EAFO Consumer Monitor 2022).

Table 14 presents the summary statistics from the survey data used for micro moments in estimation. Due to the survey's structure, models in the UK are classified as either EV or ICE cars, while in France and Germany, they are divided into BEVs (a subset of EVs) and non-BEVs. However, this different classification does not affect the main argument or the creation of micro moments. It is evident that there is a negative correlation between income levels and the ownership of electric vehicles. Wealthier consumers are more likely to own an EV (or BEV) compared to poorer consumers, who are more likely to own an ICE car (or non-BEV)[69]. This observation highlights the heterogeneous preferences for EVs, which will be incorporated into the random coefficients of the demand estimation.

It is essential to create a micro sample for demand estimation, with the goal of making it as representative as possible of the unobserved survey samples. To achieve this, I randomly draw individual income observations from the national income distribution. In the UK, data on income distribution is obtained from the Survey of Personal Incomes (HMRC). For France and Germany, the data is sourced from the EU-SILC and ECHP surveys (Eurostat). Figure 18 illustrates the evolution of disposable incomes (after tax) for the 10th, 30th, 50th, 70th, and 90th percentiles[70].

[69]The proportion of very wealthy consumers owning an EV in the UK appears higher compared to the proportion of the top income group owning a BEV in France and Germany. This discrepancy arises from different definitions of income groups, with the very wealthy in the UK representing the top 20% of earners, a broader category than the top income groups in France and Germany.

[70]The currency used in the UK is pounds sterling (GBP), while the currency adopted in France

France shows the lowest income for the bottom 10% of the population, with minimal changes over the past decade. However, France has the highest income levels for the low to medium percentiles (30th and 50th) compared to the UK and Germany. The top 10% of the population in the UK are generally wealthier than their counterparts on the continent. All three countries have experienced income growth over the last decade. However, disposable incomes have increased at a higher rate in the UK and Germany compared to France, where growth has been nearly stagnant. Wealthy individuals have seen a more significant increase in income than poorer consumers. For example, in the UK, the income threshold for the top 10% rose by £10,200, whereas the median income increased by only £6,700. In Germany, the income threshold for the top 10% increased by €12,008, nearly doubling the increase in the median income.

As previously mentioned, the random draw from the income distribution should closely replicate the survey data. To create a pseudo-survey sample, I first derive unconditional probabilities for each income group based on the conditional probabilities provided in Table 3, alongside the number of drivers for different fuel types. Next, I generate a pseudo-survey sample with 200 observations for each market (or time period) and calculate the number of consumers in various income groups using the derived unconditional probabilities. Finally, I randomly draw income observations for each income group from the interpolated income distribution, according to the number of consumers in each group.

Income distribution changes over time, and it is unlikely that consumer preferences remain independent of income levels across different periods. However, it is probable that individuals in the same relative income position (for example, the median) share similar preferences. Therefore, I adjust the cutoffs for different income

and Germany is euros (EUR).

groups based on their current relative income position while maintaining the fixed proportion of each income group.

2.4 Empirical Model

2.4.1 Set-up

I require a model that realistically captures substitution patterns between electric and combustion engine cars on the demand side, along with a competition model on the supply side. To address this, I use the random coefficients logit model (RC) from Berry et al. [1993], paired with Bertrand competition on the supply side. Alternative models include the nested logit model (NL) from Berry [1994] and the random coefficient nested logit model (RCNL) from Brenkers and Verboven [2006]. As noted by Cardell [1997] and Grigolon and Verboven [2014], the nested logit model involves a specific type of random coefficients on group dummy variables[71]. Generally, the NL model is more tractable but less flexible due to its distributional assumptions, while the RC model offers greater flexibility but is computationally demanding. The RCNL model combines the features of both the NL and RC models but often yields insignificant results, as the heterogeneity is estimated twice, both through random coefficients and nesting parameters. In practice, I find that the RC model effectively captures heterogeneity compared to the NL and RCNL models, consistent with findings in the automobile industry literature[72]. This paper does not aim to offer methodological innovation but to utilise existing models to understand recent

[71]Specifically, these (somewhat) random coefficients must follow a particular distribution to comply with the assumption that the error term is distributed according to an extreme value distribution.

[72]In practice, I attempted to separate luxury and EV models into distinct nests. The NL model performed poorly relative to the RC models when luxury and EV were included as dummy variables. The RCNL model provided similar estimates, but the price elasticity was excessively high in absolute value, suggesting model mis-specification.

developments in the Western European automobile industry.

Consumers choose products to maximise their indirect utility, exhibiting heterogeneity in response to prices and product characteristics on the demand side. On the supply side, firms engage in Bertrand competition, simultaneously setting prices to maximise overall profit across all products in a repeated game. Consumers select products from the available options in a given market. This model represents a repeated static Bertrand competition. Although a dynamic model could provide additional insights into firm decision-making, it would also significantly enlarge the state space, potentially rendering it infeasible to solve[73][74]. Since this paper primarily focuses on the impact of consumer subsidies, a repeated static Bertrand competition model is a reasonable simplification for the supply side. On the demand side, carmakers typically update their models every 7-8 years, and consumers usually keep a vehicle for 5-6 years. Given the significance of cars for many consumers, especially those in rural areas, it is unlikely that they would delay purchases in anticipation of future events; rather, they are more likely to choose alternative options (Remmy [2022]). Hence, it is also reasonable to assume that consumers' buying decisions are static.

2.4.2 Demand

I treat the UK, France, and Germany as isolated markets and compare the estimation and counterfactual results. There are $t = 1, 2, 3, \ldots, T$ markets across all three countries. Following Berry et al. [1993], markets are defined by time. In this

[73] Consider a simple product entry and exit game: in each market, firms must decide the number of new models to introduce and old models to phase out, as well as the characteristics of each new model. They must do this based on information about all available models, not just their own but also those of competitors. Such problems are typically too complex to solve using standard methods.

[74] However, I am working on a follow-up paper that examines the dynamic effects of introducing new EV models using value function approximation and neural networks.

paper, I use annual sales data from 2010 to 2021 to form a total of $T = 12$ markets. Each market contains $j = 1, 2, 3, \ldots, J_t$ products, produced by $f = 1, 2, 3, \ldots, F_t$ firms. The number of products may vary between markets. Each market contains $i = 1, 2, 3, \ldots, I_t$ consumers with different income levels, drawn from the income distribution. Consumers choose from existing products or an outside option, $j = 0$. The market size is captured by the total population for the UK, France, and Germany. Hence, consumers who choose the outside option either buy a used car or do not purchase a car at all.

Observed demand-side product characteristics form the $N \times K_1$ matrix of linear characteristics, X_1, and the $N \times K_2$ matrix of nonlinear characteristics, X_2, which is typically a subset of X_1. The structural error term, ξ, represents unobserved characteristics for all product-market bundles, with dimension $N \times 1$. In market t, demographics d form an $I_t \times D$ matrix for the agent's observed characteristics. The agent's unobserved characteristics are represented by the $I_t \times K_2$ matrix v, where the elements are drawn independently from a known distribution (e.g., normal distribution).

The indirect utility of agent i in market t from purchasing product j is:

$$U_{ijt} = \delta_{jt} + \mu_{ijt} + \epsilon_{ijt} \tag{31}$$

I use $V_{ijt} = \delta_{jt} + \mu_{ijt}$ to denote the model predicted utility excluded the idiosyncratic shocks. Vector δ is so called the mean utility in literature and the functional form is given below:

$$\delta_{jt} = \beta^{hp} \cdot Hpwt_{jt} + \beta^{SUV} \cdot SUV_{jt} + Pre16 \cdot \beta^{Pre16} \cdot x_{jt} + After16 \cdot \beta^{After16} \cdot x_{jt} + \xi_{jt} \tag{32}$$

Here, $Hpwt_{jt}$ represents the horsepower, and SUV_{jt} is a dummy variable indicating whether the body type is an SUV. The vector x_{jt} includes all interactions between a brand dummy (for luxury or affordable brands) and a fuel-type dummy (for Petrol, Diesel, and EV). Thus, x_{jt} is a 1×6 vector. As mentioned earlier, including x_{jt} is crucial for mitigating the omitted variable problem, given the strong correlation between SUVs, luxury brands, and EVs.

EV-specific characteristics, such as range and charging time, would ideally be included. Unfortunately, I do not have access to this information in the dataset. Since firms' decisions do not involve choosing range and charging time, as in other studies (Remmy [2022]; Armitage and Pinter [2021]), the EV dummy should adequately capture the effects of EV-specific characteristics. I also allow consumer preferences for x_{jt} to vary before and after 2016, in an attempt to account for any potential changes in preferences due to the dieselgate scandal and other factors, should they exist.[75]

There are $K_1 = 15$ linear characteristics in total. To further control for unobserved characteristics, ξ_{jt}, I decompose it as follows:

$$\xi_{jt} = \xi_f + \xi_{year} + \tilde{\xi}_{jt} \tag{33}$$

Here ξ_f is the fixed effect for firm f and ξ_{year} is the annual fixed effect to capture the industry wide shock in a given year (for example, the 2020 covid pandemic). Finally, $\tilde{\xi}_{jt}$ is the pure unobserved characteristic to form the moment function for identification.

[75] An alternative approach would be to create a linear time-trend variable, rather than using collapsed time dummies as in Remmy [2022]. However, most estimators become insignificant under the time-trend setting, suggesting that any changes in preference are minor, as will be seen later in the estimation results.

The agent-specific utility is divided into following parts:

$$\mu_{ijt} = -\frac{\alpha}{Y_{it}} \cdot P_{jt} + \gamma^{ICE} \cdot Incomegroup_{it} \cdot ICE_{jt} + \gamma^{EV} \cdot Incomegroup_{it} \cdot EV_{jt} \quad (34)$$

Here, P_{jt} represents the final price of product j at time t, after deducting any subsidies from the listed price. Y_{it} denotes the income level of agent i at time t. Wealthier consumers are less sensitive to price changes compared to poorer consumers. Each consumer is categorised into one of five income groups: very poor, poor, medium, rich, and very rich, aligning with the micro-moments observed in the survey data. Consumer preferences for internal combustion engine (ICE) vehicles and electric vehicles (EVs)[76] are allowed to vary across different income groups. This captures important heterogeneity in vehicle preference, which would otherwise be overlooked in a standard setting. Lastly, consumer preferences for luxury brands are subject to a random shock, σv_{it}, where σ is the standard deviation and v_{it} is randomly drawn from a standard normal distribution.

Consumer i in market t chooses alternative products to maximise her utility. Given that the idiosyncratic shock ϵ_{ijt} follows a type-I extreme value distribution, the probability of consumer i to choose product j in market t is:

$$s_{ijt} = \frac{expV_{ijt}}{1 + \sum_{k \in J_t} expV_{ikt}} \quad (35)$$

Aggregate market shares are obtained by integrating over the distribution of individual heterogeneity, which includes both the agent's observed and unobserved characteristics. The integration is typically solved using numerical approximations. I utilise the Halton sequence for numerical integration, which functions similarly to

[76]Given the nature of the survey data, in France and Germany, ICE vehicles include all non-BEV models (Petrol, Diesel, and PHEV), while EVs only include BEV models. In the UK, ICE vehicles include Petrol and Diesel models, while EVs include both BEV and PHEV models.

a quasi-Monte Carlo simulation.[77] The Halton sequence generates an $I_t \times K_2$ matrix of integration nodes, v, and an $I_t \times 1$ vector of integration weights, w, such that:

$$s_{jt} \approx \sum_{i \in I_t} w_{it} s_{ijt} \tag{36}$$

2.4.3 Supply

I model the profit-maximisation price decisions for F multi-product carmakers in each market t. I assume that all product characteristics are fixed so as not to further complicate the model. The observed product characteristics form the $N \times K_3$ matrix, X_3, into which any non-price characteristics can be included. An additional $N \times 1$ vector of unobserved supply-side characteristics, ω, is introduced, acting similarly to the vector ξ on the demand side.

Firm f sets prices for each product in its portfolio $J_{ft} \subset J_t$ to maximise total profit at time t:

$$\pi_{ft} = \sum_{j \in J_{ft}} (p_{jt} + \lambda_{jt} - c_{jt}) s_{jt} \tag{37}$$

Here λ_{jt} denotes the subsidy for product j in time t. There are different ways to model the asymmetrical price between supply-side and demand-side where the listed price is the choice variable. The current setting is used because I use the final price on the demand side. The first order conditions are, in vector-matrix form:

$$p + \lambda - c = \underbrace{\Delta^{(-1)} s}_{\eta} \tag{38}$$

Here η denotes the markups and it depends on Δ, a $J_t \times J_t$ matrix of intra-firm

[77] Other options include Monte Carlo, Gauss-Hermite quadrature, sparse grids, and Latin hypercube sampling.

(negative, transposed) demand derivatives:

$$\Delta = -\mathcal{H} \odot \frac{\partial s'}{\partial p} \tag{39}$$

Here $\mathcal{H}$ denotes the market-level ownership matrix with dimension $J_t \times J_t$ where element $H_{jk} = 1$ if the same firm produces product j and k, and 0 otherwise.

I pull the subsidy and the marginal cost term together to create a standard notation as in Berry, Levinsohn, and Pakes (1995) where:

$$c' = c - \lambda \tag{40}$$

By doing this, it becomes easier to further decompose the *de facto* marginal cost c' and estimate the relevant coefficients. Additionally, this creates a simplified notation for counterfactual scenarios without subsidies, by adding the subsidy to the estimated marginal cost. The explanatory variables for the marginal cost are assumed to be exactly the same as the linear characteristics on the demand side:

$$log(c'_{jt}) = \gamma^{hp} log(Hpwt_{jt}) + \gamma^{SUV} SUV_{jt} + \gamma^{EV} EV_{jt} + \gamma^{Diesel} Diesel_{jt} + \omega_{jt} \tag{41}$$

I take the logarithmic form of the marginal cost. A simplified set of variables is used here compared to the demand side, as employing the same set would lead to insignificant results. This suggests that marginal costs do not change significantly before and after 2016. Insignificant results can create substantial bias in the counterfactual analysis; hence, I have limited the variables to those that remain relevant. Environmental policies in other sectors may affect the marginal cost of producing one unit of an existing model, particularly if the cost of raw materials increases. Technological advancements could also alter marginal costs, especially for electric vehicles (EVs). Hsieh et al. [2019] show a clear downward trend in the price of lithium-ion cells, a key component in EV batteries. As on the demand side, I also include firm fixed effects and annual fixed effects:

$$\omega_{jt} = \omega_f + \omega_{year} + \tilde{\omega}_{jt} \tag{42}$$

2.5 Estimation

2.5.1 Instrumental Variables

Demand-side estimation in the random coefficient logit model faces the well-known endogeneity issue, where the price may be correlated with unobserved characteristics. Even if additional variables correlated with the price are included, the unobserved characteristics may still contain important factors that correlate with the price. One notable example is fuel efficiency, often measured in miles per gallon (MPG), which unfortunately is not observed in my dataset. Suitable instruments not only mitigate the endogeneity issue but also help to identify the random coefficients, thus serving a dual purpose.

Recent literature has highlighted that the widely used BLP instruments, which involve summing the product characteristics, tend to perform rather poorly (Reynaert and Verboven [2014]; Gandhi and Houde [2019]). The main reason is that the BLP model is not 'optimal' in the sense defined by Chamberlain [1987] for optimal instruments. While the sum of characteristics contains relevant information about the price, it also introduces a significant amount of noise. To address this, I combine three sets of instruments to filter the most relevant information regarding prices.

I use differentiation instruments, as introduced by Gandhi and Houde [2019]. The idea is to select close substitutes based on their relative positions in the characteristic space. I employ the quadratic variant, which sums the squared differences between product characteristics across the *entire* sample:

$$\begin{aligned} Z_{jtk}^{Quad,Other}(X) &= \sum_{k \in J_{ft} \setminus \{j\}} d_{jtkl}^2 \\ Z_{jtk}^{Quad,Rival}(X) &= \sum_{k \notin J_{ft}} d_{jtkl}^2 \end{aligned} \tag{43}$$

Here $d_{jtkl} = x_{ktl} - x_{jtl}$ is the difference between products j and k in terms of characteristic l. $Z_{jtk}^{Quad,Other}$ is the instrument regarding firm's internal distance and $Z_{jtk}^{Quad,Rival}$ is the instrument for the product from rival firms. Finally, I also include the government subsidy as a instrument as in Armitage and Pinter [2021].

2.5.2 Identification

The aforementioned instruments aid in identifying the linear parameters on both the demand and supply sides. Market shares vary across different products and over time, whereas the explanatory variables (characteristics) primarily vary over time. It is well known that the BLP-style model is only identified when the outside option is specified and the indirect utility for the outside option is normalised to $U_{i0t} = \epsilon_{i0t}$. The firm's first-order conditions, combined with optimal prices and the log form of marginal cost relative to characteristics, assist in identifying the supply parameters.

On the heterogeneous side, market shares, non-linear characteristics, consumer demographics, and micro-moments help identify the demographic coefficients Π and variance σ. This is because consumer income levels vary over time. All theoretical identification results can be found in Berry and Haile [2014], Berry and Haile [2024], Salanié and Wolak [2022], and Conlon and Gortmaker [2023]. More details on constructing the moment functions will be discussed in the next section.

2.5.3 Estimation Approach

On the demand side, there are three sets of parameters to estimate: the linear parameter β, the demographic heterogeneity matrix Π, and the random heterogeneity Σ. Including the supply side, there is also γ. All non-concentrated-out parameters are denoted by θ. The estimation process follows several steps. First, the mean utility $\delta(\theta)$ must be computed for each market. Second, the structural error terms $\hat{\xi}(\theta)$ and $\hat{\omega}(\theta)$ are estimated and used to form the moment functions. Third, additional moment functions based on micro-moments from survey data are calculated. Finally, the GMM objective function is constructed using the moment functions and the optimal weighting matrix. I will now describe each step in detail.

To compute the mean utilities $\delta(\theta)$, I use the SQUAREM accelerated iteration method from Varadhan and Roland [2008], and Reynaerts et al. [2012]. This method is typically three to six times faster than the standard iteration of Berry et al. [1993], as it employs a first-order squared non-monotone extrapolation scheme. The iteration is defined as:

$$\delta_{jt}^{h+1} \longleftarrow \delta_{jt}^{h} - 2\alpha^{h} r^{h} + (\alpha^{h})^{2} v^{h} \tag{44}$$

Where

$$\alpha^{h} = \frac{(v^{h})' r^{h}}{(v^{h})' v^{h}}, r^{h} = f(\delta_{jt}^{h}) - \delta_{jt}^{h}, v^{h} = f(f(\delta_{jt}^{h})) - 2f(\delta_{jt}^{h}) + \delta_{jt}^{h} \tag{45}$$

Here $f(\delta_{jt}^{h}) = \delta_{jt}^{h} + (1-\rho)(\log s_{jt} - \log s_{jt}(\delta^{h}, \theta))$ denotes the original contraction. Once the mean utility is calculated, marginal cost are then computed according to firm's first order conditions:

$$c_{jt}(\theta) = p_{jt} - \eta_{jt}(\theta) \tag{46}$$

Concentrated out parameters are computed with linear IV-GMM where:

$$\begin{bmatrix} \hat{\beta}^{ex} \\ \hat{\gamma} \end{bmatrix} = (X'ZWZ'X)^{-1}X'ZWZ'Y(\theta) \tag{47}$$

Where

$$X = \begin{bmatrix} X_1^{ex} & 0 \\ 0 & X_3 \end{bmatrix}, Z = \begin{bmatrix} Z_D & 0 \\ 0 & Z_S \end{bmatrix}, Y(\theta) = \begin{bmatrix} \delta(\theta) - X_1^{en}\hat{\alpha} \\ \tilde{c}(\theta) \end{bmatrix} \tag{48}$$

Here I use the same notation defined in Conlon and Gortmaker [2020] where the first row in $Y(\theta)$ includes the endogenous part. This is different from the standard notation in Berry et al. [1993] where the endogenous coefficients has to be jointly estimated with the supply side. The reason of doing so is to increase the accuracy of the estimators on endogenous variables. Once the linear parameters are concentrated out, the (estimated) structural errors are:

$$\begin{bmatrix} \hat{\xi}(\theta) \\ \hat{\omega}(\theta) \end{bmatrix} = \begin{bmatrix} \hat{\delta}(\theta) - X_1\hat{\beta} \\ \tilde{c}(\theta) - X_3\hat{\gamma} \end{bmatrix} \tag{49}$$

Estimated structural errors are then used to construct the moment conditions $E[g_{D,jt}] = E[g_{S,jt}] = 0$ where:

$$\begin{bmatrix} g_{D,jt} & g_{S,jt} \end{bmatrix} = \begin{bmatrix} \hat{\xi}_{jt}Z_{D,it} & \hat{\omega}_{jt}Z_{S,jt} \end{bmatrix} \tag{50}$$

Here $Z_{D,jt}$ and $Z_{S,jt}$ are $N \times M_D$ and $N \times M_S$ matrices of the demand and supply instruments. To bring down the dimensionality of the moment condition, sample analogues $\overline{g}$ for each instruments are used to construct the objective function for the GMM estimation where:

$$\overline{g} = \begin{bmatrix} \overline{g}_D \\ \overline{g}_S \end{bmatrix} = \frac{1}{N} \begin{bmatrix} \sum_{j,t} Z'_{D,jt}\hat{\xi}_{jt} \\ \sum_{j,t} Z'_{S,jt}\hat{\omega}_{jt} \end{bmatrix} \tag{51}$$

More detailed micro data on individual choices can be used to supplement the standard demand- and supply-side moments $\overline{g}_D$ and $\overline{g}_S$ above with an additional $m = 1, 2, ..., M_M$ micro moments, $\overline{g}_M$, for a total of $M = M_D + M_S + M_M$ moments:

$$\overline{g} = \begin{bmatrix} \overline{g}_D \\ \overline{g}_S \\ \overline{g}_M \end{bmatrix} \tag{52}$$

Conlon and Gortmaker [2023] provides a standardised framework for incorporating micro moments into BLP estimation. The base idea is to create distance moments based between the statistics observed in the survey data and model predictions:

$$\overline{g}_{M,m} = f_m(\overline{v}) - f_m(v) \tag{53}$$

Here $f_m(\cdot)$ is a function that maps micro moment parts $\overline{v}$ or v into a micro statistic. This is typically a conditional expectation. Like in the demand- and supply-moments, I use the sample analogue version. More details can be checked in Conlon and Gortmaker (2023). Finally, the GMM problem becomes:

$$\min_{\theta} q(\theta) = \overline{g}(\theta)' W \overline{g}(\theta) \tag{54}$$

Here $q(\theta)$ is the GMM objective on the non-linear parameters θ. W is the GMM optimal weighting matrix with dimension $M \times M$. Traditionally, the 2SLS weighting matrix is used in the first stage:

$$W = \begin{bmatrix} (Z_D' Z_D/N)^{-1} & 0 \\ 0 & (Z_S' Z_S/N)^{-1} \end{bmatrix} \tag{55}$$

With two-step GMM, which is usually the case, W is updated before the second stage according to:

$$W = S^{-1} \tag{56}$$

Matrix S could take different forms depending on the optimisation context. I use the clustered weighting matrices with $c = 1, 2, ..., C$ clusters:

$$S = \frac{1}{N} \sum_{c=1}^{C} g_c g_c' \tag{57}$$

Where, letting the set $J_{ct} \subset J_t$ denote products in cluster c and market t,

$$g_c = \sum_{t \in T} \sum_{j \in J_{ct}} g_{jt} \tag{58}$$

The robust standard errors are adjusted for each cluster. For the (re)calculation of equilibrium prices in the counterfactual analysis and data simulation, I use the reformulation of the FOC of the profit maximisation problem by Morrow and Skerlos [2011], which enhances performance in terms of speed and convergence compared to the conventional Newton's method. For the software, I utilise the PyBLP package developed by Conlon and Gortmaker [2020] and Conlon and Gortmaker [2023] to perform the estimation. Several micro-moments are constructed to match the observations in the survey data. I use the 'BFGS' algorithm for optimisation, setting the tolerance to $1e^{-4}$. The 'SQUAREM' iteration method is employed, with the convergence tolerance set to $1e^{-14}$.

It is well-known that the GMM objective function is generally non-convex. Additionally, numerical issues arise when random coefficients for dummy variables become too large. Given the scope of the GMM problem, I impose a bound on random coefficients, drawing from the extensive results in the automobile literature. I also experiment with different starting points to find the global minimum within the bound, and use both first-order and second-order gradients to ensure the optimiser does not get stuck at a saddle point.

2.6 Estimation Results

Table 15 presents the estimated coefficients for the linear parameters. Columns two to four represent the demand-side coefficients, while columns five to seven display the supply-side coefficients. The top panel shows the universal variables that remain constant over time, whereas the bottom two panels illustrate the coefficients of interacted dummies for pre-2016 and post-2016 levels on the demand side. All else being equal on the demand side, consumers in all three countries prefer vehicles with higher horsepower. German consumers place a higher value on each additional unit of horsepower compared to consumers in the UK and France. Consumers also show a preference for SUVs over other body types, with UK consumers deriving the least utility premium from SUVs.

In the bottom two panels, changes in preferences exhibit distinct patterns across the three countries. French consumers' preference for all brand $\times$ fuel-type intersections declines, suggesting a generally weaker automobile market after 2016. In contrast, UK consumers' preference for all vehicle types increased after 2016. In Germany, the preference for models from affordable brands increased, while the preference for luxury brand models decreased after 2016, indicating a weaker market for luxury brand sales overall.

On the supply side, producing models with higher horsepower is more costly. SUVs have a significant impact on marginal costs in the UK and France, but not in Germany. In the lower panel, EVs are more expensive to produce in the UK, cheaper in Germany, and roughly the same cost in France when compared to diesel cars. I also attempted to include the luxury dummy on the supply side. However, the marginal cost for luxury brands proved insignificant in all three countries, indicating that, all else being equal, there is no substantial difference in marginal production costs between luxury and affordable brands.

Table 16 presents the results from the non-linear estimates, which include interaction terms between income levels and car characteristics. Additionally, Table 16 shows the price elasticity for different fuel-type vehicles. Recall that Table 14 demonstrates a strong correlation between income levels and the likelihood of being an EV (BEV) driver. Two factors could explain this observed survey result. First, wealthy consumers may be less sensitive to price changes compared to poorer consumers, increasing their likelihood of purchasing more expensive EVs. This effect is confirmed in all three countries, as the coefficients for price/income are consistently negative. Second, affluent consumers may have a stronger preference for EVs than poorer consumers. This reflects a combination of *between* and *within* effects. Wealthier consumers may have a greater preference margin between EVs and ICE cars (*within*), compared to poorer consumers (*between*).

In the UK, while consumers in all income groups prefer ICE vehicles over EVs, the preference gap narrows as income increases. In France, the preference gap between BEVs and non-BEVs remains relatively stable, regardless of changes in consumers' income. Germany is the only country where some income groups prefer BEVs over non-BEVs. Except for the poorest consumers, all other income groups showed a preference for BEVs, with the preference margin increasing as consumers become wealthier. All estimates are statistically significant.

Table 16 also presents the price elasticity for all three countries in the total sample, as well as for each fuel-type sub-sample. The overall elasticity is consistent with estimates in the existing literature on the automobile industry (Berry et al. [1993]; Petrin [2002]; Brenkers and Verboven [2006]; Grigolon and Verboven [2014]; Grieco et al. [2024]; Linn [2022]; Remmy [2022]; Armitage and Pinter [2021]). Germany exhibits the highest price elasticity across the three countries.

In comparing the elasticity across the fuel-type sub-samples, consumers in the UK and France are more price-sensitive towards EVs than ICE cars. However, in Germany, EVs show the lowest price elasticity among all fuel types. This result aligns with the previous finding that, except for lower-income consumers, German consumers exhibit a strong preference for EVs.

Figure 19 presents the estimated marginal costs across different fuel types and brand models. In general, diesel cars are more expensive to produce than petrol cars, and EVs are even costlier to manufacture than diesel cars. Additionally, luxury models are generally more expensive to produce than affordable models. It is important to note that the estimated marginal costs are derived after accounting for the purchase subsidy. The *actual* marginal costs should be the estimated values plus the current subsidy amount. In the UK and Germany, the real marginal costs for EVs initially declined and then stabilised, whereas in France, the marginal costs for EVs continued to rise. On the brand side, subsidies affect both luxury and affordable brands in a similar manner. Overall, marginal costs have increased over the past decade.

In general, diesel cars are more expensive to produce than petrol cars, and EVs are even costlier to manufacture than diesel vehicles. Additionally, luxury models are generally more expensive to produce than affordable models. It is important to note that the estimated marginal costs were derived by factoring in the purchase subsidy. The *actual* marginal costs should therefore be the estimated values plus the current subsidy amount. In the UK and Germany, the real marginal costs for EVs initially decreased before stabilising, whereas in France, the marginal costs for EVs continued to rise. On the brand side, subsidies affect both luxury and affordable brands in a similar way. Overall, marginal costs have increased over the past decade.

Figure 20 represents the average markups by fuel and brand types. The average markups in the UK and France decreased throughout the 2010s. This is similar to the findings of Grieco et al. [2024], where they observed a similar decline in markups in the US automobile market. Potential reasons for this decrease include increasing competition and rising raw material prices. However, markups remained relatively constant in Germany, which aligns more closely with the results found in macroeconomic literature (for example, De Loecker et al. [2020]). Germany also has the lowest average markups among the three countries. This can be partially explained by the fact that, unlike the UK and France, Germany has a large and highly competitive car manufacturing industry, leading to greater competition and lower markups.

Regarding markups for different fuel types, the three countries exhibit distinct patterns. In the UK and France, EV markups have been the lowest among all fuel types since 2016, once the EV market reached a size significant enough to draw meaningful conclusions. However, in Germany, EV markups consistently exceed those of their ICE counterparts. When combined with the findings from Figure 5, the main factor is that EVs are significantly more costly to produce in the UK and France compared to Germany.

Figure 20 also shows that markups for luxury brands are generally lower than those for affordable brands. This is a striking result, considering that luxury brands have higher marginal costs, as shown in Figure 19. The high prices of luxury brand models do not stem from market power but rather from the high marginal costs. Therefore, a more targeted subsidy policy is needed to lower EV prices by reducing the marginal costs for luxury brands, which dominate the new models in the EV market.

To summarise, the estimation results indicate that different income groups exhibit varying price sensitivities and preferences for EVs. On the production side, EVs are generally more costly to produce, and average markups either decrease or remain constant in the UK, France, and Germany.

2.7 Welfare Analysis

In this section, I calculate the consumer surplus for each agent within their respective income group. Additionally, I compute the firm's profits at equilibrium. Table 17 presents the average change in consumer surplus for each income group from 2015 to 2021.

The average consumer surplus across the entire sample decreases in the UK and France, while it increases in Germany. Specifically, consumers in the UK experience the greatest decline, whereas those in Germany see the largest gains. This outcome is partly explained by the fact that German consumers received the most generous subsidies in 2021, while UK consumers benefited from significantly higher subsidies in 2015 compared to 2021. However, the distribution of these benefits is highly unequal across income groups. For example, very wealthy consumers gain substantially more than poorer consumers. A typical consumer in the top 20% income group experiences welfare improvements in all three countries, whereas consumers in the bottom two income groups face welfare losses.

Another noteworthy finding is that consumers in the middle-income group are the most negatively affected by changes in the automobile market. Unlike lower-income consumers, who may opt for outside alternatives, middle-income consumers generally require a car for daily commuting. However, shifts in the market have forced them to switch to more expensive EVs, resulting in significant welfare losses. This striking result highlights the unequal impact of market changes on different consumers. In particular, wealthier consumers benefit more from the rise of EVs, as they receive the same amount of government subsidies but have a stronger preference for EVs compared to poorer consumers. Meanwhile, middle-income consumers are the hardest hit, as EVs are too expensive for them, and the availability of ICE vehicles is simultaneously diminishing.

Figure 21 presents the total profits by fuel and brand types. Profits for EVs have been increasing, while profits for diesel cars have been declining in all three countries since 2016. Despite the generally higher markup for diesel cars compared to EVs (as shown in Figure 20), the sharp decline in diesel car market share has led to decreasing total profits for diesel vehicles. Finally, the total profits for all firms have risen from 2015 to 2021.

In summary, changes in consumer surplus vary significantly across income groups, with wealthier consumers benefiting much more than poorer consumers. The total profits from diesel cars have declined, while profits from EVs have increased. As both overall consumer surplus and firms' profits have increased over time, it appears that overall market efficiency has improved, driven either by advanced technologies like EVs or increased competition. However, government subsidies also distort the market to some extent. Therefore, it is essential to identify the effect of these subsidies on market share changes and overall welfare. This will be discussed in the next chapter.

2.8 Counterfactual One: Driver of the Market Share Changes

Figure 16 shows a drastic increase in EV sales and a significant drop in the market share of diesel cars, but the underlying reasons for this shift in market structure remain unclear. In this section, I use the estimation results to explore the detailed causes of this change through various counterfactual designs. The findings will provide important insights into the effectiveness of current policies and inform recommendations for better policy design in the future.

Table 18 presents the overall change in market shares from 2015 to 2021. Several factors, as predicted by the random coefficient logit model, may have contributed to this shift.

First, purchase subsidies affect the perceived prices and, consequently, the sales of models with different fuel types. Second, the available product list and its characteristics (both observed and unobserved) may shift market shares. Third, consumer preferences for various fuel and brand types may evolve over time, influencing their indirect utility when selecting different models.

It is challenging to analyse these four factors simultaneously. Instead, I examine each factor individually. For instance, to explore the effect of subsidies, I adjust only the marginal costs of subsidised models while keeping all other factors constant. By doing this, I decompose the overall changes into four distinct sections, each representing a driving factor. I denote these factors as F_1, F_2, and F_3, representing subsidies, preferences, characteristics, and product lists, respectively. The overall change in market shares can be decomposed as follows:

$$
\begin{aligned}
s_{jt}^{21} - s_{jt}^{15} &= s_{jt}(F_1^{21}, F_2^{21}, F_3^{21}) - s_{jt}(F_1^{15}, F_2^{15}, F_3^{15}) \\
&= s_{jt}(F_1^{\mathbf{21}}, F_2^{21}, F_3^{21}) - s_{jt}(F_1^{\mathbf{15}}, F_2^{21}, F_3^{21}) \\
&+ s_{jt}(F_1^{15}, F_2^{\mathbf{21}}, F_3^{21}) - s_{jt}(F_1^{15}, F_2^{\mathbf{15}}, F_3^{21}) \\
&+ s_{jt}(F_1^{15}, F_2^{15}, F_3^{\mathbf{21}}) - s_{jt}(F_1^{15}, F_2^{15}, F_3^{\mathbf{15}})
\end{aligned}
\tag{59}
$$

It is important to note that when considering the subsidy effect, the second line of equation (29) represents the *difference* in the subsidy effect between 2015 and 2022, rather than the absolute impact of subsidies compared to a no-subsidy scenario. The overall number of available products has decreased over time, meaning that most models in 2021 had counterparts in 2015.[78] Meanwhile, a significant portion of the models available in 2015 had disappeared by 2021.

In principle, three simulated equilibria are required to fully decompose the overall change in the manner described by equation (29). However, since the total change in market shares is already known, two simulations are sufficient to derive the complete result. In this case, I run counterfactual analyses for subsidies and preferences, allowing the effects of the product list and characteristics to be determined automatically. I also changed the order of the three factors in equation (29) to test the stability of the results. All outcomes were consistent, confirming that the order of the factors does not affect the conclusions. Table 19 presents the simulation results as a percentage of the total absolute changes. This approach is necessary because some of the absolute changes in market shares, particularly for the subsidy effect, are very small.

[78]For example, the C7 version of the Audi A6 was introduced in 2011, while the C8 version was produced starting in 2018. In fact, over 85% of models in 2021 have a counterpart in 2015.

Table 19 indicates that the subsidy effect is negligible in the UK. In France, the subsidy effect accounts for only 16% of the total change in market shares. In contrast, the subsidy effect explains nearly half of the total increase in EV market share in Germany. This is because the subsidy amount in the UK and France was much more generous in 2015 compared to 2021. However, since the EV market share was very small in 2015, the subsidy had little impact in these countries. By 2021, as the EV market matured, the subsidy amount had sharply reduced. In Germany, however, the government started subsidising EVs at a later stage, when the EV market was more established. This made the subsidy policy more effective, contributing to nearly half of the overall changes in EV market shares. These results suggest an important implication: purchase incentives are more effective when the market for new technology is more mature.

The preference effect aims to capture the impact of the dieselgate scandal, as many previous studies have claimed it shifted consumer preferences away from diesel cars in the US, where the scandal originated. However, the simulation results do not fully support the same argument for Western Europe. The preference effect for petrol cars is quite volatile in the UK and France, where preference changes actually led to significant shifts in the opposite direction of the overall petrol car sales trends. One possible reason for this is that the total change in petrol car sales in the UK and France is relatively small compared to Germany.

For diesel cars, while preference shifts account for more than 100% of the drop in diesel car sales in France, they again move in the opposite direction in the UK, indicating that UK consumers actually preferred diesel cars after 2016. German consumer preferences for diesel cars remained relatively unchanged after 2016. One explanation for why the dieselgate scandal had less of an impact in Europe compared to the US is that most of the manufacturers involved in the scandal are European, making them harder to replace in Europe, where they have much stronger regional market power. On the EV side, preferences explain much of the share changes in the UK, have the opposite effect in France, and have only a negligible impact in Germany.

Finally, changes in the product list and characteristics also show interesting results across the three countries. Changes in fleet and characteristics explain more than the actual changes for petrol cars. However, they only explain the changes in diesel cars in the UK and Germany. For EVs, fleet and characteristics changes have no impact on EV sales in the UK, explain half of the changes in Germany, and account for more than the observed increase in France.

The effectiveness of the subsidy depends on the cost pass-through to the price (i.e., the ratio of change in equilibrium price to the change in marginal cost). Previous literature has outlined the pass-through ratio in Bertrand competition, including Zimmerman and Carlson [2010] for linear demand, and Verboven and Dijk [2009]) for logit demand. However, analysing the random coefficient logit demand is more complex. One general finding is that the cost pass-through is larger when the demand curve becomes more convex. Weyl and Fabinger [2013] also confirmed that the pass-through rate can exceed 100%, resulting in an overshifting situation. The authors argue that this phenomenon is more likely to occur in Bertrand competition than in Cournot competition.

Figure 22 presents the average own pass-through ratio for subsidised vehicles in the UK, France, and Germany. It is evident that the demand for subsidised vehicles (BEVs and PHEVs) is sufficiently convex, leading to overshifting in all cases. EV producers in France and Germany have greater market power than those in the UK, allowing them to charge even higher prices with the same increase in marginal costs. If the actual MSRP (prices used in the estimation plus the subsidy) is considered, the model predicts that MSRP would rise in the absence of subsidies. The overall pass-through ratio is decreasing, suggesting a more competitive EV market over time.

Table 20 shows the welfare impact of the three different effects. Many of the results reflect the findings in Table 8. Preference effects negatively impact consumer welfare in the UK and Germany, while French consumers benefit significantly from the shift in preferences. Product and characteristics effects move in the opposite direction compared to preference effects. Overall, UK and French consumers experience a welfare loss, while German consumers enjoy a welfare gain. This outcome is largely driven by the fact that Germany implemented a more generous subsidy policy at a much later stage compared to the UK and France.

In summary, the counterfactual results indicate that the subsidy effect plays a minor role in the change in market shares. However, the primary reasons for these changes vary across countries. In the UK, changes are primarily driven by characteristic and preference effects, while in France and Germany, characteristic and product effects account for most of the changes in market shares. The ineffectiveness of the subsidy policy stems from the fact that the EV market had not yet matured when subsidies were substantial, and the demand was sufficiently convex, allowing EV producers to wield significant market power and distort the subsidy's impact. The key policy implication is that it is more effective to subsidise EVs when the market is more mature (as seen in Germany) and more competitive. Characteristic and product effects dominate the changes in consumer surplus, with changes in characteristics decreasing surplus, while changes in the product list increase surplus.

2.9 Counterfactual Two: Income-based Subsidy Policy

In this section, I construct counterfactuals based on proposed alternative income-based policies. Specifically, two policy strands are considered. The first keeps overall spending constant for a given time period, allowing the subsidy amount to vary by income level. The goal of this policy is to maximise EV sales. The second strand aims to keep the EV market share constant while minimising the total spending required to achieve a specific level of EV market share. I use the differential evolution method to find the optimal alternative subsidy allocations.[79] Given the heavy computational burden of this optimisation, it is infeasible to find the optimal policies for each year from 2010 to 2021. Instead, I select two representative years—2017 and 2021—to compute optimal income-based subsidies.

The year 2017 captures a scenario where the EV market was immature and subsidies were relatively generous in some countries. This result is particularly relevant for emerging EV markets worldwide. The year 2021 reflects a more mature EV market with stricter subsidy policies, which is indicative of markets that already have established EV sectors but aim to reduce financial support. Finally, I compare the consumer surplus differences under no subsidy, uniform subsidy, and income-based subsidy for consumers from various income groups.

Implementing an income-based subsidy policy may seem challenging due to the additional resources required to assess each individual's income. However, such policies are already in place in several US states. For example, the California Clean

[79]Remmy (2023) used the grid search method to solve a different policy redesign. However, given the vast candidate space in my problem and the significant computational time required for each new equilibrium, I employ the differential evolution method, which is designed for this type of optimisation.

Vehicle Assistance Program (CVRP) offers income-based incentives for the purchase or lease of new or used electric vehicles, targeting California residents living in disadvantaged communities (DAC).[80] Similarly, the New York State Drive Clean Rebate provides higher rebate amounts for low- and moderate-income individuals purchasing or leasing electric vehicles.[81] However, I have not observed income-based subsidy policies being implemented in Europe, either in the past or in the near future. This creates an ideal scenario to study the effects of alternative policies, offering implications for both Europe and emerging EV markets globally.

Income-dependent subsidies result in varying outcomes. Since lower-income consumers are more price-sensitive in all three countries, it is "cheaper" to stimulate additional EV sales by subsidising low- and medium-income consumers. However, this does not always lead to higher sales, as lower-income consumers generally exhibit a stronger preference for ICE vehicles over EVs. Table 16 confirms this argument, showing that the gap between preferences narrows as consumer income increases. Additionally, targeting specific income groups for subsidies shifts consumer surplus, as different income groups receive varying amounts of financial support.

Before presenting the simulation results for alternative income-based policies, it is crucial to introduce some comparable benchmarks for evaluating subsidy effectiveness beyond total spending and EV market share. Two new concepts are introduced here. First, I use the subsidy per extra unit of EV sales (*subsidy-per-unit*) to capture the pure monetary incentives. This measure is commonly used in the literature, including by Hughes and Podolefsky [2015], Li [2019], Springel [2021], and Linn [2022]. The difference between the policy subsidy and the *subsidy-per-unit* captures the number of consumers who would not have purchased an EV without

[80] https://cleanvehiclegrants.org/
[81] https://www.nyserda.ny.gov/All-Programs/Drive-Clean-Rebate-For-Electric-Cars-Program

the subsidy. As the *subsidy-per-unit* increases, the policy's effectiveness diminishes. Figure 23 illustrates the annual average *subsidy-per-unit* in the three countries.

Germany has the highest overall *subsidy-per-unit* compared to the UK and France, particularly after 2020. This result suggests that most EV buyers in Germany would still purchase EVs even in the absence of a subsidy policy. The UK, however, has the most effective subsidy policy in terms of triggering additional EV sales. In all three countries, the *subsidy-per-unit* is at least three times the current subsidy, meaning that the majority of EV buyers would continue to purchase EVs even without the subsidy.

The second benchmark is the welfare change per unit of sales (*welfare-per-unit*), which is the sum of the change in consumer welfare, the change in carmaker profits, and the subsidy expenditure, divided by the change in EV sales. This measure, used by Pless and Van Benthem [2019]), captures the social benefits or costs, excluding environmental externalities. Figure 24 illustrates the annual average *welfare-per-unit*.

The result in Figure 24 is quite striking because the welfare changes in all three countries are positive. This contrasts with the findings of Linn [2022] for the US market, where welfare changes were negative. One key difference between this paper and Linn [2022] is that this study does not account for interactions between different policies, which may drive the results in varying directions.

From the components of the *welfare-per-unit*, two main reasons explain the increase. First, the improvement in consumer surplus under the subsidy policy exceeds the subsidy amount itself. This is primarily due to the pass-through rate being greater than 1, as shown in Figure 8. The reduction in marginal costs due to subsidies leads to an even larger reduction in price, thereby boosting consumer surplus. Second, carmakers' profits increase with the subsidy because more consumers are willing to purchase EVs as they become cheaper. In the US, spending on other policies, such as ZEV and GHG credits, potentially drives firms' profits into the negative, which results in a negative *welfare-per-unit*. However, the positive *welfare-per-unit* in this study is not distributed evenly across income groups, as will be discussed later in this section.

The goal is to determine five optimal subsidy amounts for each income group in different scenarios. The following steps are used to conduct the analysis. First, I fix four subsidy amounts and use a root-finding algorithm to pin down the remaining one. The root function depends on the counterfactual scenario. If the goal is to maximise EV shares given a fixed total expenditure, the root objective function ensures that the equilibrium-induced total expenditure matches the fixed amount. If the goal is to minimise spending while maintaining a certain EV share, the root-finding function ensures that the equilibrium-induced EV shares equal the benchmark.

Second, I allow the four subsidy amounts to vary and calculate the simulated EV shares or subsidy amounts derived from the remaining optimal subsidy, as found by the root-finding method. Third, I use the differential evolution method to find the optimal subsidy allocation for either share maximisation or spending minimisation scenarios.

Table 21 presents the simulation results for 2021, when the EV market is more *mature*. It includes the optimal subsidy amounts, original and simulated objectives (including EV shares and total subsidy amounts), as well as *subsidy-per-unit* and *welfare-per-unit* for both the original and simulated values. It is more effective to subsidise low-income consumers, regardless of the objective, as they are more price-sensitive. For high-income consumers, many EV buyers would still purchase EVs even without subsidies. Therefore, it is more cost-efficient to allocate subsidies to low-income rather than high-income consumers in all three countries.

Regarding the simulated objectives, the results are mixed. The EV shares simulated under the income-based policy are not drastically different from those under the uniform policy, with the largest improvement in Germany showing only a 25% increase in EV shares. However, the share-oriented policy does reduce the *subsidy-per-unit* by around 35% to 45%. On the other hand, the spending minimisation policy significantly impacts the outcomes, yielding more than 50% in cost savings while achieving the same EV shares. Moreover, the *subsidy-per-unit* tends to be lower compared to the share maximisation policy.

However, the *welfare-per-unit* decreases under the income-based policy compared to the uniform policy, indicating that the income-based policy is actually more costly in terms of welfare. Further analysis reveals that both average consumer welfare and firms' profits decrease under the new policy, and the sum of these negative effects outweighs the cost savings from the subsidy. This important result suggests that while the income-based policy is more cost-efficient, it comes at the expense of reduced average consumer surplus and firms' profits. It is crucial to balance welfare gains and losses when setting policy from the government's perspective.

Table 22 presents the results for 2017, reflecting an *immature* market situation. The key implications of the results—such as allocating more subsidies to poorer consumers than to wealthier ones, and the minor improvement in EV shares relative to the savings in subsidies—remain consistent with the 2021 findings. The only exception is the subsidy-driven policy in the UK, where most of the subsidy is allocated to medium-income consumers, resulting in increased total subsidy spending.

In the 2017 simulation, the improvements in EV shares are slightly higher in percentage terms compared to the 2021 case, with both France and Germany experiencing a 30% increase in EV shares compared to 20% in 2021. In the bottom panel, the *subsidy-per-unit* and *welfare-per-unit* results also follow similar patterns to those seen in 2021, where both values decrease under the income-based policy.

The results from Table 21 and Table 22 suggest that there is no significant difference in the relative improvement of income-based policies between *mature* and *immature* markets. While the income-based policy makes it cheaper to stimulate an additional EV sale, this comes at the cost of overall welfare, affecting consumers, carmakers, and the government.

Although the subsidy policy undoubtedly increases the average consumer surplus, it is not evenly distributed across all income groups. Table 23 shows the difference in average consumer surplus among the five defined income groups, with and without the subsidy policy (both uniform and income-based) in 2021. Under the uniform subsidy, wealthier consumers gain significantly more welfare compared to poorer consumers, as expected, since poorer consumers are more price-sensitive. However, under the income-based subsidy, the situation is reversed for the very-poor and very-rich consumers, as the optimal subsidy allocation is heavily skewed toward poorer consumers.

The income-based policy results in a U-shaped change in surplus as consumers become wealthier. This is because medium-income consumers are both too price-sensitive compared to the very-rich and receive far less subsidy than the very-poor. When comparing share-oriented and subsidy-oriented optimal policies, it appears that the distortions under the subsidy-driven policy are more moderate.

In summary, I propose an alternative income-based subsidy policy that shows both higher induced EV shares and lower total subsidy spending compared to the uniform subsidy policy. The *subsidy-per-unit* measurement decreases under the income-based policy, indicating it is cheaper to promote an additional EV sale under this approach. However, this comes at the cost of total welfare for consumers, firms, and the government, as the *welfare-per-unit* also declines. Under the uniform policy, wealthier consumers gain more surplus than poorer consumers, while under the income-based policy, welfare gains follow a U-shape as consumer income rises.

2.10 Conclusion

This paper addresses three major questions. First, what drives the recent sharp increase in EV sales? To answer this, I use annual new car sales data from the UK, France, and Germany to estimate consumer preferences via a random coefficient logit model. I then run various simulations allowing key factors affecting EV market shares—subsidies, preferences, and product/characteristics—to change. The results show that the subsidy effect plays a minor role in explaining the rise in EV shares compared to the influence of preferences and product/characteristics changes. This suggests that the primary drivers of EV adoption stem from changes in consumer preferences following the 2015 'dieselgate' scandal and product list/characteristics changes after governments announced plans to ban ICE cars within the next 15 to 20 years.

Second, this paper explores heterogeneous preferences among consumers of different income levels. Using data from national travel surveys, I generate micro-moments to improve the accuracy of the estimation. The random coefficients estimates illustrate that poorer consumers are more price-sensitive than wealthier consumers, and higher-income consumers tend to have a stronger preference for EVs over ICE cars. This heterogeneity in EV preferences means that different income groups are affected unevenly by market changes. Welfare analysis reveals that very-poor consumers experienced a welfare loss, while very-rich consumers saw a welfare gain from 2015 to 2021 in all three countries. This indicates that poorer consumers were more negatively impacted by the sharp rise in EV market shares compared to wealthier consumers.

Third, the paper proposes an alternative income-based subsidy policy and compares its effectiveness to the current uniform subsidy across various dimensions. The optimal income-based subsidy generally allocates more funds to poorer consumers, who are more price-sensitive and therefore more likely to be incentivised to purchase an EV. A significant number of wealthier consumers would buy EVs even without a subsidy. The income-based subsidy policy is up to 60% more cost-effective in boosting EV sales compared to the uniform subsidy. However, this comes at the cost of total welfare, as the combined consumer surplus and firm profits decline by more than the government savings. Nevertheless, unlike in the US market, the welfare of each additional EV sale remains positive, leaving some room to balance environmental externalities.

A significant gap in this paper is the lack of analysis on how car manufacturers decide whether to introduce new EV models or phase out existing ICE models. Both the effects of 'dieselgate' and the impending ban on ICE car sales make introducing new ICE models more costly, and consumer preferences for EVs make it even harder to achieve profits with ICE models. Addressing product entry/exit decisions would require a dynamic discrete/continuous choice model. However, given the large state and choice spaces, traditional methodologies may not be sufficient to solve the model.

On the policy side, instead of providing subsidies to consumers to lower EV purchase costs, governments could also offer subsidies to carmakers to reduce EV production costs. It would be interesting to explore under what conditions one policy outperforms the other and examine the differing welfare impacts of purchase subsidies versus production subsidies. These research questions are left for future study.

2.11 Tables

Table 12: Summary Statistics (UK and France)

UK	Mean	Std	Min	Max		Mean	Std	Min	Max
All cars					**SUVs**				
Price (1e4)	3.990	2.259	0.926	15.00	Price (1e4)	4.960	2.347	1.312	15.00
Horsepower	0.38	0.18	0.12	1.68	Horsepower	0.45	0.19	0.18	1.68
SUV	0.35	0.48	-	-	SUV	-	-	-	-
Luxury	0.33	0.47	-	-	Luxury	0.38	0.49	-	-
EV	0.19	0.39	-	-	EV	0.23	0.42	-	-
Observation	31757				Observation	11004			
Luxury					**EV**				
Price (1e4)	5.687	2.472	1.122	15.00	Price (1e4)	6.220	2.850	1.608	14.98
Horsepower	0.52	0.19	0.12	1.68	Horsepower	0.51	0.25	0.12	1.68
SUV	0.40	0.49	-	-	SUV	0.42	0.49	-	-
Luxury	-	-	-	-	Luxury	0.48	0.50	-	-
EV	0.27	0.45	-	-	EV	-	-	-	-
Observation	10535				Observation	5985			
FR	Mean	Std	Min	Max		Mean	Std	Min	Max
All cars					**SUVs**				
Price (1e4)	3.731	2.244	0.755	15.00	Price	4.788	2.412	1.557	15.00
Horsepower	0.34	0.17	0.12	1.59	Horsepower	0.42	0.19	0.18	1.53
SUV	0.32	0.46	-	-	SUV	-	-	-	-
Luxury	0.24	0.43	-	-	Luxury	0.31	0.46	-	-
EV	0.24	0.43	-	-	EV	0.32	0.47	-	-
Observation	27552				Observation	8709			
Luxury					**EV**				
Price (1e4)	5.946	2.722	1.344	15.00	Price (1e4)	6.067	3.014	1.661	15.00
Horsepower	0.49	0.23	0.12	1.59	Horsepower	0.49	0.26	0.12	1.59
SUV	0.41	0.49	-	-	SUV	0.42	0.49	-	-
Luxury	-	-	-	-	Luxury	0.45	0.50	-	-
EV	0.45	0.50	-	-	EV	-	-	-	-
Observation	6524				Observation	6617			

Table 13: Summary Statistics (Germany)

DE	Mean	Std	Min	Max		Mean	Std	Min	Max
All cars					**SUVs**				
Price (1e4)	3.929	2.217	0.788	15.00	Price (1e4)	4.675	2.213	1.445	15.00
Horsepower	0.39	0.18	0.12	1.68	Horsepower	0.46	0.18	0.18	1.68
SUV	0.30	0.46	-	-	SUV	-	-	-	-
Luxury	0.31	0.46	-	-	Luxury	0.36	0.48	-	-
EV	0.17	0.37	-	-	EV	0.22	0.41	-	-
Observation	44230				Observation	13333			
Luxury					**EV**				
Price (1e4)	5.588	2.380	1.365	15.00	Price (1e4)	5.721	2.978	1.308	15.00
Horsepower	0.52	0.20	0.12	1.68	Horsepower	0.49	0.25	0.12	1.68
SUV	0.35	0.48	-	-	SUV	0.39	0.49	-	-
Luxury	-	-	-	-	Luxury	0.46	0.50	-	-
EV	0.25	0.43	-	-	EV	-	-	-	-
Observation	13833				Observation	7393			

Notes: an observation is a maker-model-month. Prices is the list price (MSRP) minus the subsidy in a given time in US dollars to represent the final price faced by consumers. The horsepower is also normalised by 400 to make the mean closer to 1.

Table 14: Income Level vs Fuel-types

Net income / Fuel-types	EV driver	ICE driver
UK, N = 3,626		
Very poor (1st quintile)	9%	15%
Poor (2nd quintile)	8%	17%
Medium (3rd quintile)	21%	22%
Rich (4th quintile)	25%	23%
Very rich (5th quintile)	37%	23%
Net income / Fuel-types	BEV driver	Non-BEV driver
FR, N = 1,703		
< 800€	3%	4%
800 - 1,999€	10%	32%
2,000 - 3,999€	45%	44%
4,000 - 5,999€	28%	15%
>= 6,000€	14%	5%
DE, N = 1,648		
< 800€	2%	4%
800 - 1,999€	5%	26%
2,000 - 3,999€	36%	45%
4,000 - 5,999€	25%	19%
>= 6,000€	32%	6%

Notes: Net income represents the monthly disposable after-tax income.

Table 15: Estimation Results (Linear Parameters)

Country	UK	France	Germany		UK	France	Germany
Side	demand	demand	demand	**Side**	supply	supply	supply
Variables 1:							
hpwt	4.07	3.87	5.05	log(hpwt)	1.14	1.34	1.21
	(0.31)	(1.50)	(1.14)		(0.03)	(0.16)	(0.05)
SUV	0.52	0.50	0.35	SUV	0.03	0.04	0.01
	(0.09)	(0.09)	(0.09)		(0.01)	(0.02)	(0.02)
Pre 2016:				**All time**			
Afford×Petrol	1.01	5.05	9.09	EV	0.34	0.24	0.16
	(0.32)	(0.94)	(1.03)		(0.03)	(0.04)	(0.04)
Luxury×Petrol	1.16	5.42	9.23	Diesel	0.22	0.22	0.24
	(0.31)	(1.01)	(1.06)		(0.01)	(0.04)	(0.01)
Afford×Diesel	1.14	5.76	9.49				
	(0.31)	(1.07)	(1.14)				
Luxury×Diesel	1.94	6.60	10.31				
	(0.30)	(1.14)	(1.18)				
Afford×EV	1.45	5.79	9.05				
	(0.74)	(1.25)	(1.38)				
Luxury×EV	4.70	5.44	9.30				
	(1.21)	(1.20)	(1.36)				
After 2016:							
Afford×Petrol	4.41	4.03	9.56				
	(0.28)	(2.00)	(1.01)				
Luxury×Petrol	4.57	4.23	9.12				
	(0.34)	(2.01)	(1.02)				
Afford×Diesel	4.06	4.13	9.58				
	(0.30)	(2.03)	(1.13)				
Luxury×Diesel	5.02	5.09	10.04				
	(0.34)	(2.06)	(1.16)				
Afford×EV	5.55	4.57	9.42				
	(0.59)	(2.06)	(1.22)				
Luxury×EV	7.37	4.72	9.16				
	(0.66)	(2.05)	(1.17)				

Table 16: Estimation Results (Non-linear Parameters)

UK

Π	1/incomeUS	poor	medium	rich	veryrich
prices	-3.22 (0.18)				
ICE		-1.10 (0.12)	-2.06 (0.14)	-2.93 (0.17)	-4.27 (0.24)
EV		-1.73 (0.80)	-1.96 (0.66)	-2.88 (0.70)	-4.27 (0.78)

France

Π	1/incomeUS	poor	medium	rich	veryrich
prices	-2.74 (0.74)				
Non-BEV		-3.00 (0.65)	-4.62 (1.02)	-5.80 (1.31)	-6.21 (1.45)
BEV		-4.65 (1.50)	-5.56 (1.20)	-6.41 (1.54)	-6.50 (1.78)

Germany

Π	1/incomeUS	poor	medium	rich	veryrich
prices	-3.10 (0.53)				
Non-BEV		-4.53 (0.57)	-6.89 (0.90)	-8.43 (1.12)	-9.26 (1.27)
BEV		-6.08 (1.24)	-6.71 (1.23)	-7.49 (1.42)	-6.70 (1.53)

	UK	**France**	**Germany**
Elasticity (Overall)	-3.36	-2.77	-3.58
Elasticity (Petrol)	-3.13	-2.58	-3.55
Elasticity (Diesel)	-3.41	-2.88	-3.69
Elasticity (EV)	-4.74	-2.95	-3.42

Income/Country	**UK**	**France**	**Germany**
Very-poor	-52.98	-32.48	-88.57
Poor	-93.07	-25.76	-38.45
Medium	-113.48	-113.53	-35.25
Rich	-183.25	18.38	82.31
Very-rich	36.80	132.32	670.49
Overall	-83.37	-49.90	31.72

Table 17: Change of Consumer Surplus 2016-2021 (Population Averaged)

Fuel/Country	**UK**	**France**	**Germany**
Petrol	-0.398%	0.184%	-0.499%
Diesel	-1.608%	-1.096%	-1.146%
EV	0.417%	0.394%	0.776%

Table 18: Market Share Change 2016-2021

Factors/Fuels	**Petrol**	**Diesel**	**EV**
Subsidy (UK)	0.000%	0.029%	2.887%
Subsidy (France)	-1.129%	0.012%	16.427%
Subsidy (Germany)	1.474%	0.309%	44.880%
Preference (UK)	-393.594%	-20.297%	96.482%
Preference (France)	-1175.581%	184.889%	-148.814%
Preference (Germany)	-79.162%	5.535%	4.553%
Product and Characteristics (UK)	493.595%	120.326%	0.632%
Product and Characteristics (France)	1276.710%	-84.901%	232.387%
Product and Characteristics (Germany)	177.688%	94.156%	50.568%

Table 19: Decomposition of Market Share Changes

Factors/Countries	**UK**	**France**	**Germany**
Subsidy	2.218	10.135	53.699
Preference	-355.602	747.199	-43.104
Product and Characteristics	270.011	-827.234	21.125
Overall	-83.373	-49.900	31.720

Table 20: Decomposition of Consumer Surplus Change

2021	**UK**	**UK**	**France**	**France**	**Germany**	**Germany**
objectives	shares	subsidy	shares	subsidy	shares	subsidy
Very-poor	7065.6	5468.5	11468.2	9607.2	13212.4	11717.6
Poor	4048.1	1709.3	4239.3	3610.9	10407.8	9547.9
Medium	42.8	0.1	130.5	2.4	414.1	266.3
Rich	0.3	0.7	91.7	26.8	50.6	16.9
Very-rich	3.2	0.9	91.2	32.1	137.9	24.5
objectives (old)	0.458%	828.3	0.424%	2368.5	0.804%	11461.1
objectives (new)	0.482%	404.8	0.499%	912.2	1.008%	6133.1
subsidy-per-unit (old)	11670.8	11670.8	14567.4	14567.4	16449.5	16449.5
subsidy-per-unit (new)	6940.6	5703.6	7580.8	5678.3	10362.6	8802.6
welfare-per-unit (old)	10191.2	10191.2	11673.9	11673.9	9931.5	9931.5
welfare-per-unit (new)	8543.8	9633.4	3268.6	6422.7	83.8	3068.9

Table 21: Optimal Income-based Subsidy 2021

2017	**UK**	**UK**	**France**	**France**	**Germany**	**Germany**
objectives	shares	subsidy	shares	subsidy	shares	subsidy
Very-poor	6588.6	256.6	9791.7	8415.1	7856.0	6470.5
Poor	2984.7	6136.5	3819.9	321.2	4872.9	3510.0
Medium	998.9	9018.5	18.2	4.9	114.8	2.6
Rich	24.7	7620.8	103.8	1.1	33.8	0.0
Very-rich	255.9	135.7	85.0	2.0	31.6	0.3
objectives (old)	0.071%	475.6	0.053%	542.1	0.064%	430.6
objectives (new)	0.078%	790.1	0.070%	210.5	0.084%	137.7
subsidy-per-unit (old)	10430.8	10430.8	12436.8	12436.8	12361.1	12361.1
subsidy-per-unit (new)	7951.2	17327.4	7006.7	4828.1	5799.0	3954.0
welfare-per-unit (old)	6546.2	6545.2	8275.1	8275.1	10490.4	10490.4
welfare-per-unit (new)	4243.5	5576.1	2327.9	4237.1	1805.0	5106.0

Table 22: Optimal Income-based Subsidy 2017

2021	**UK**	**UK**	**France**	**France**	**Germany**	**Germany**
objectives	shares	subsidy	shares	subsidy	shares	subsidy
Uniform subsidy/no subsidy						
Very-poor	3.66	—	3.87	—	16.55	—
Poor	2.16	—	3.96	—	17.83	—
Medium	5.03	—	13.05	—	46.86	—
Rich	5.31	—	22.29	—	72.44	—
Very-rich	4.49	—	33.36	—	210.28	—
Income-based subsidy/no subsidy						
Very-poor	15.80	9.96	81.29	34.15	262.24	121.16
Poor	4.84	2.05	6.62	4.39	46.50	37.04
Medium	3.77	2.07	8.77	5.66	19.43	15.45
Rich	3.48	2.39	13.39	8.90	27.40	21.92
Very-rich	3.33	2.29	19.18	12.84	74.60	56.21

Table 23: CS Changes for Income-based and Uniform Policy (2021)

2.12 Figures

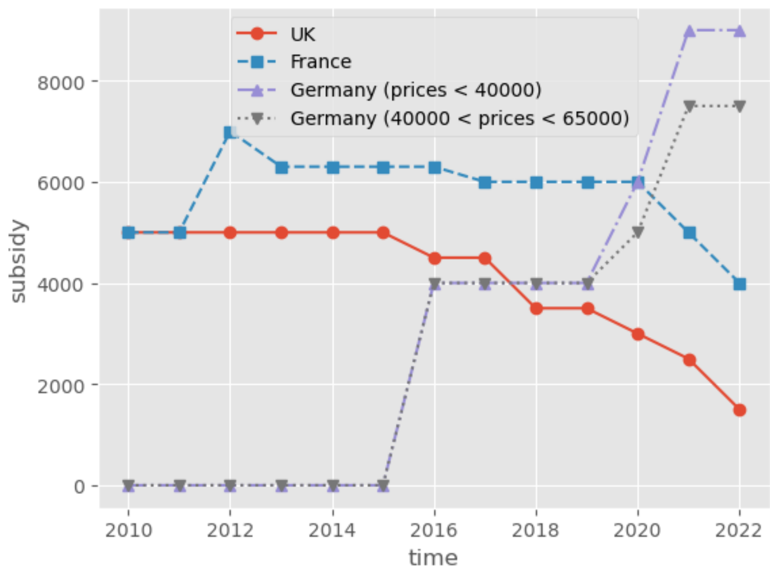

(a) BEV subsidy

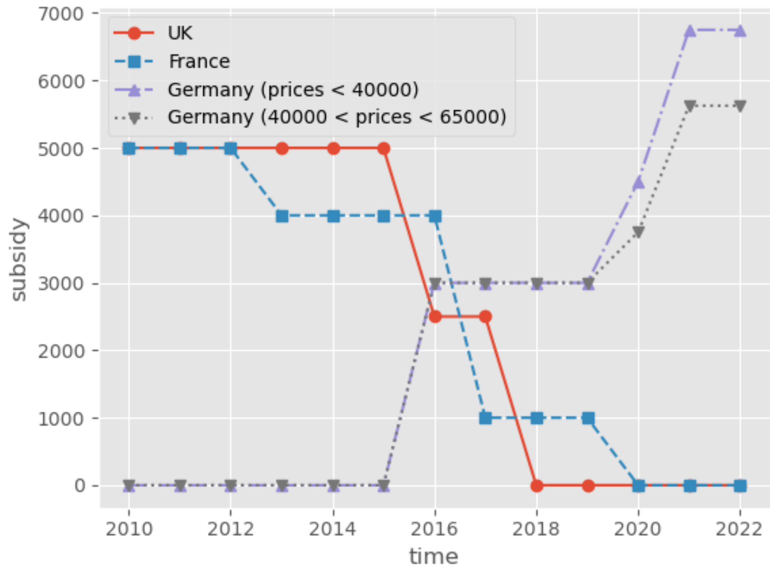

(b) PHEV subsidy

Figure 15: Subsidy Evolution 2010-2022

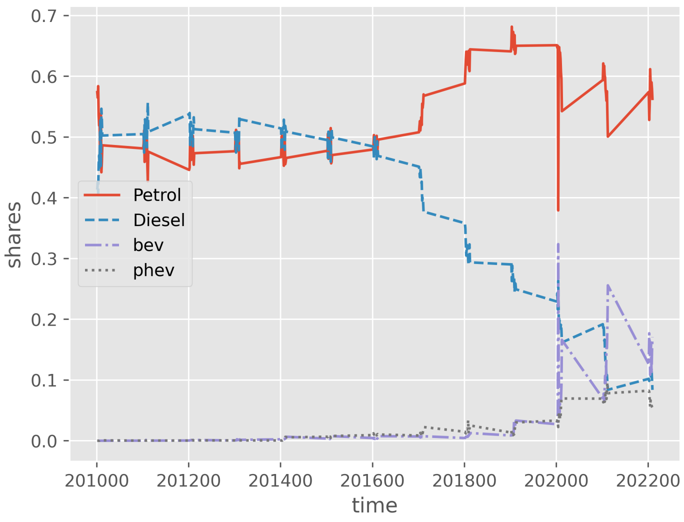

(a) UK Fuel-type shares

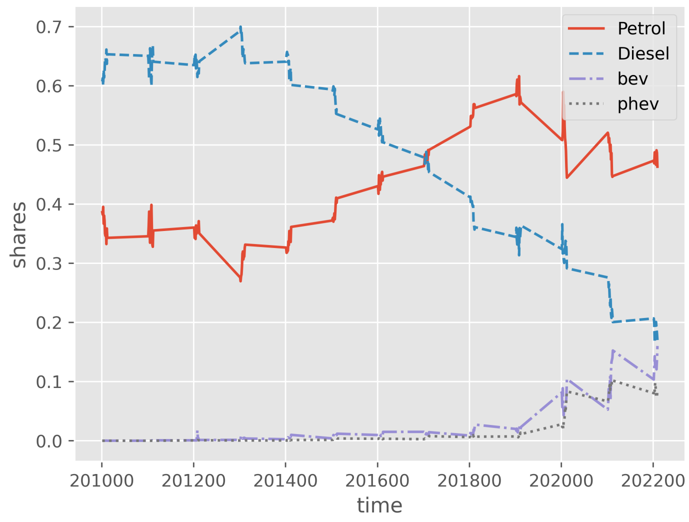

(b) FR Fuel-type shares

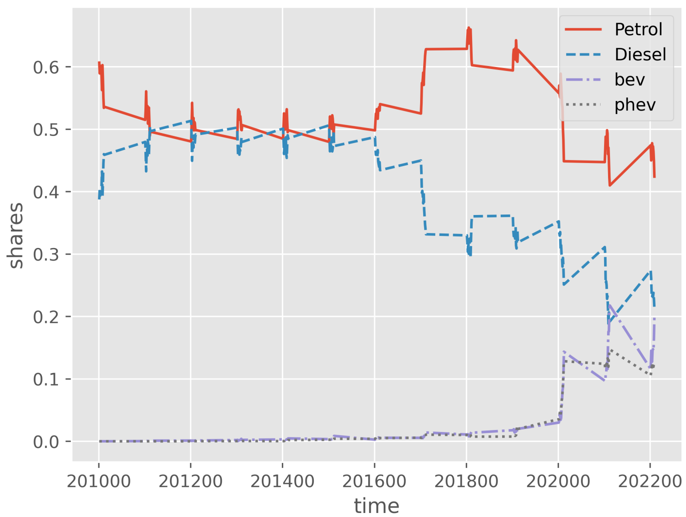

(c) DE Fuel-type shares

Figure 16: Market share by Fuel Types

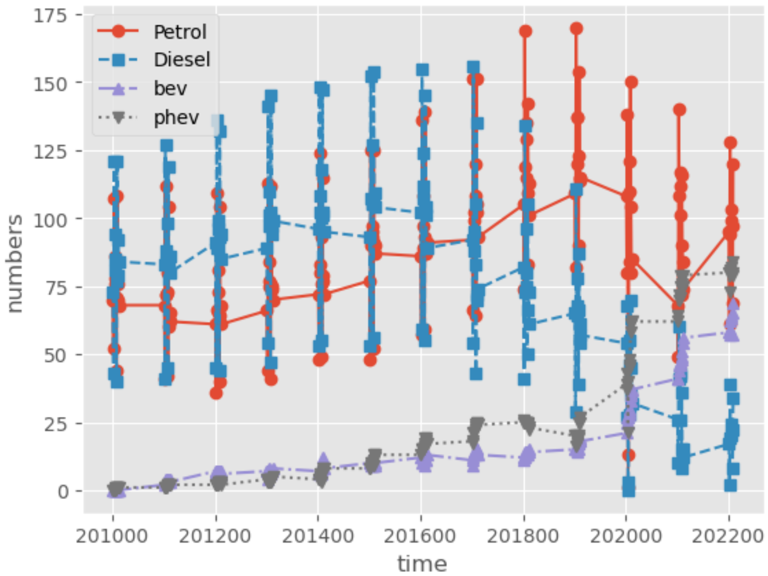

(a) UK Model Numbers

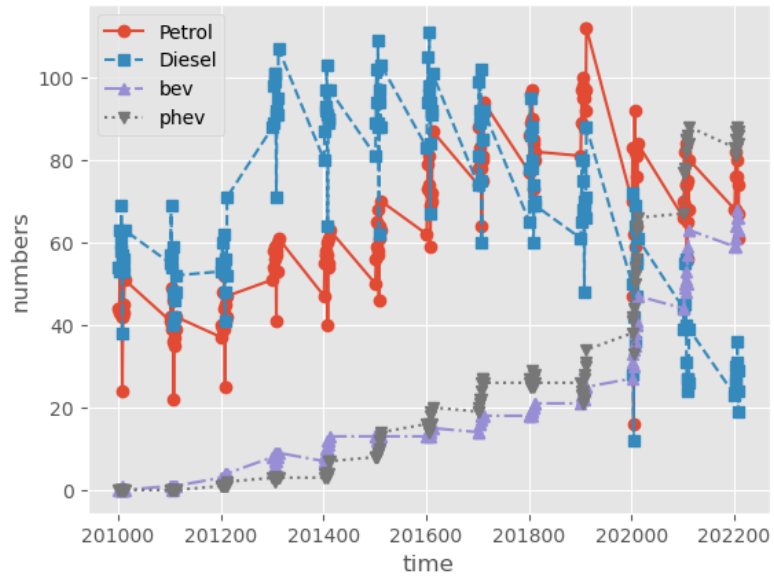

(b) FR Model Numbers

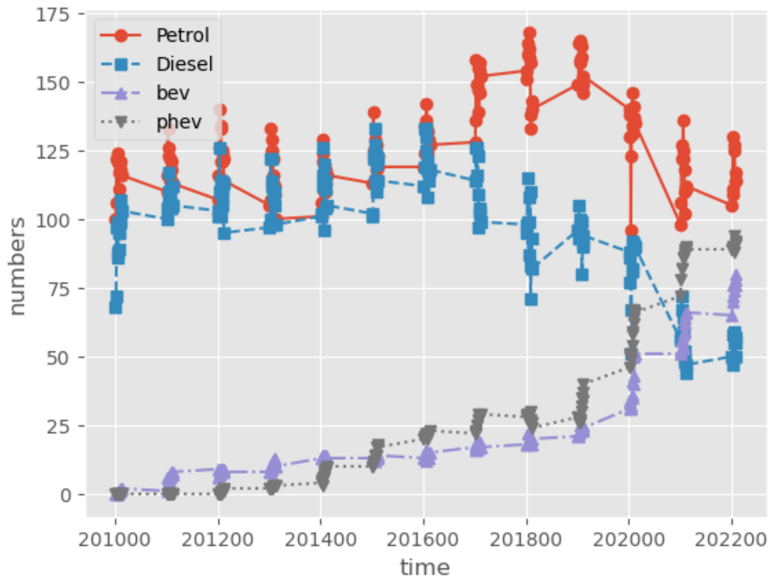

(c) DE Model Numbers

Figure 17: Model Numbers by Fuel Types

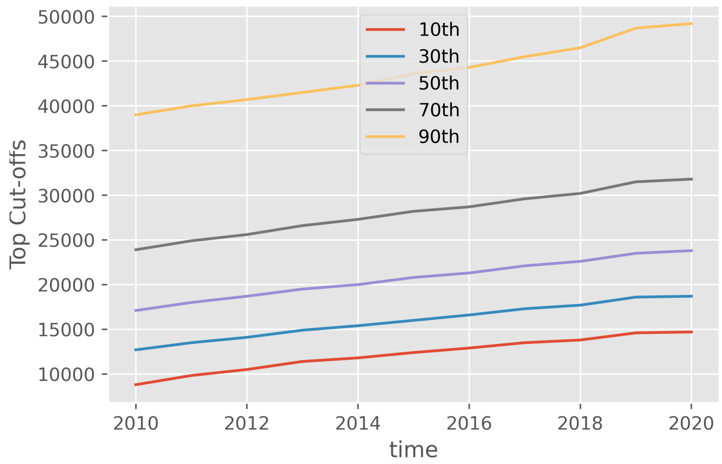

(a) UK income

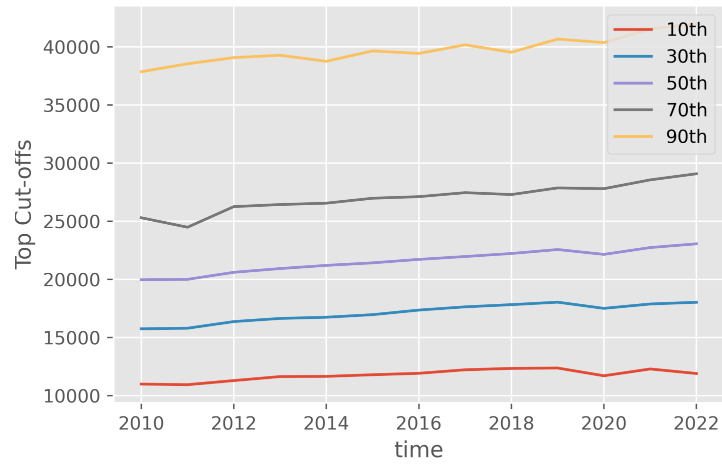

(b) FR income

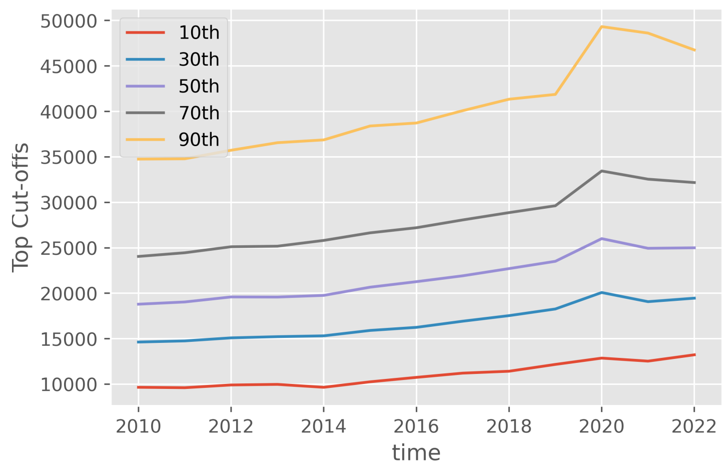

(c) DE income

Figure 18: Evolution of Income distribution 2010-2022

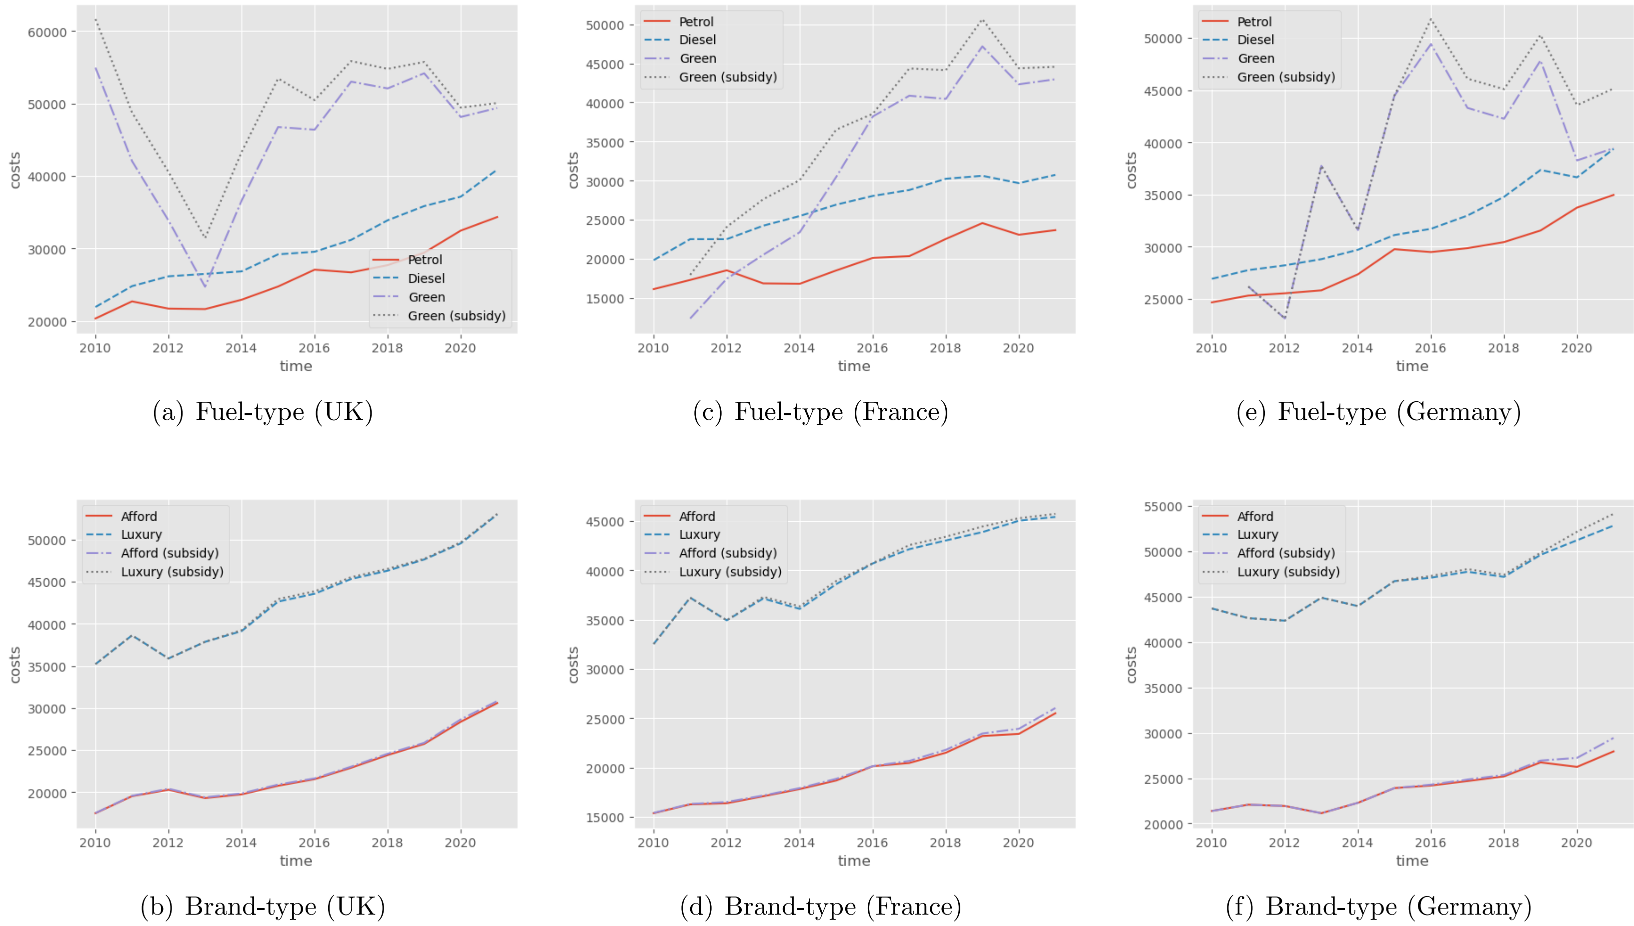

(a) Fuel-type (UK)

(c) Fuel-type (France)

(e) Fuel-type (Germany)

(b) Brand-type (UK)

(d) Brand-type (France)

(f) Brand-type (Germany)

Figure 19: Estimated Marginal Costs

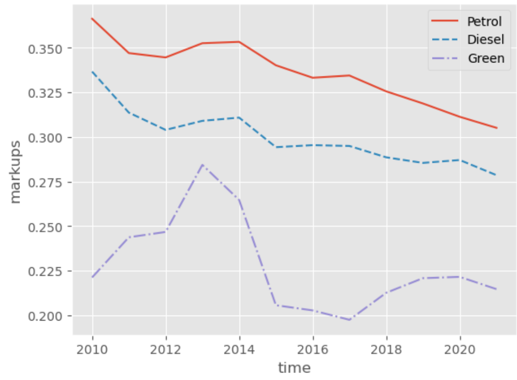

(a) Fuel-type (UK)

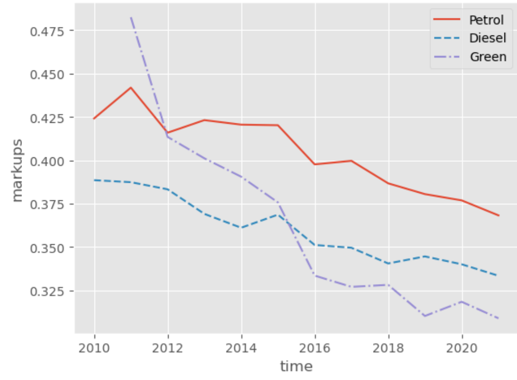

(c) Fuel-type (France)

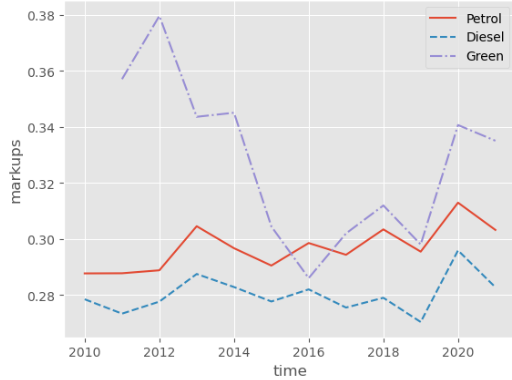

(e) Fuel-type (Germany)

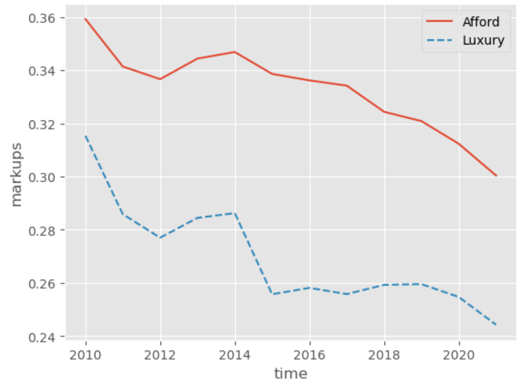

(b) Brand-type (UK)

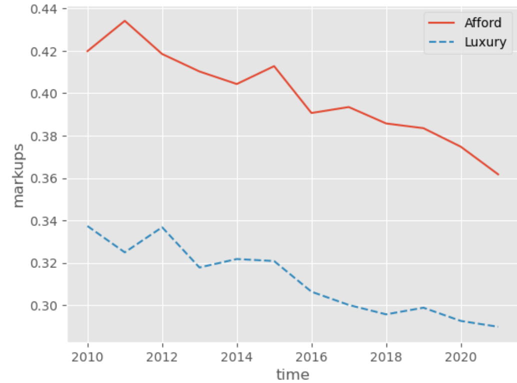

(d) Brand-type (France)

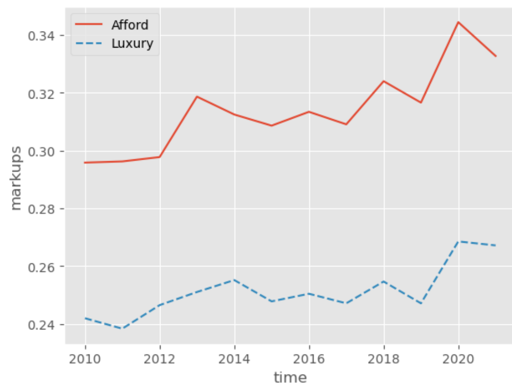

(f) Brand-type (Germany)

Figure 20: Estimated Markups

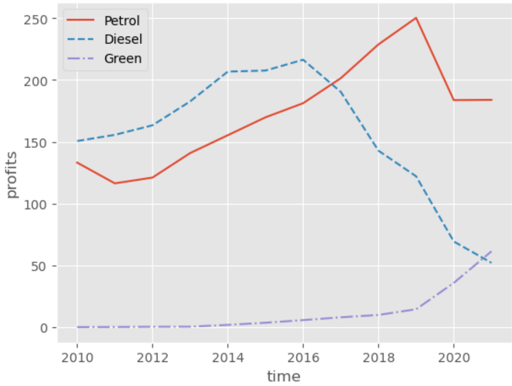

(a) Fuel-type (UK)

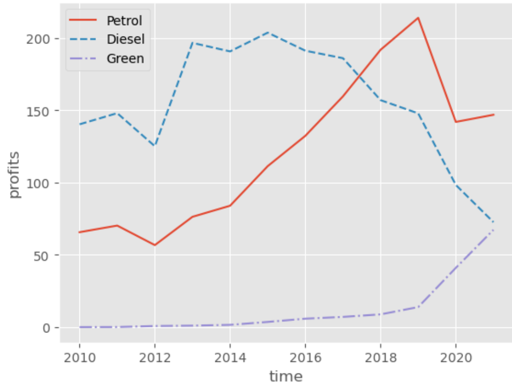

(c) Fuel-type (France)

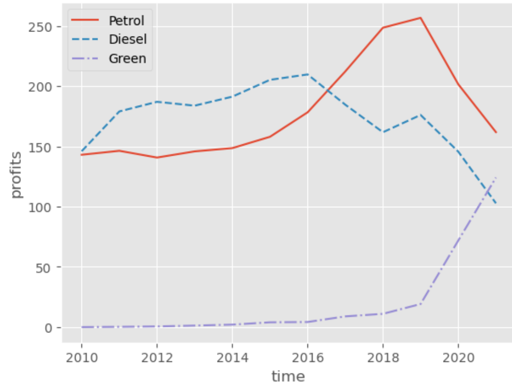

(e) Fuel-type (Germany)

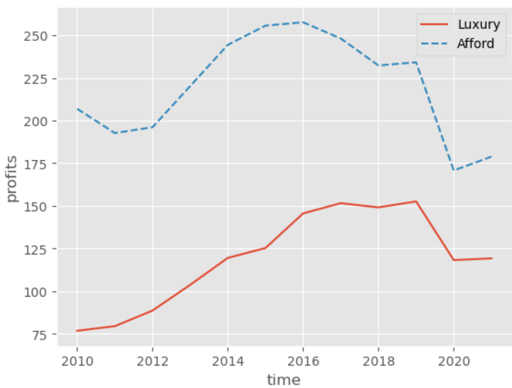

(b) Brand-type (UK)

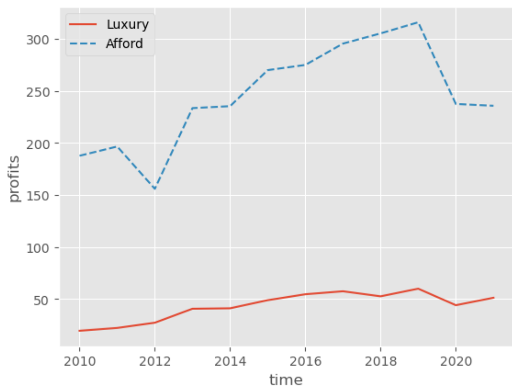

(d) Brand-type (France)

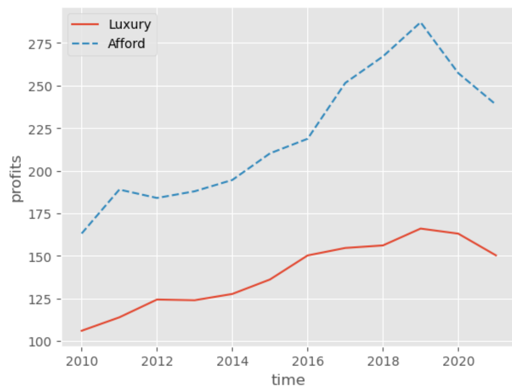

(f) Brand-type (Germany)

Figure 21: Computed Profits

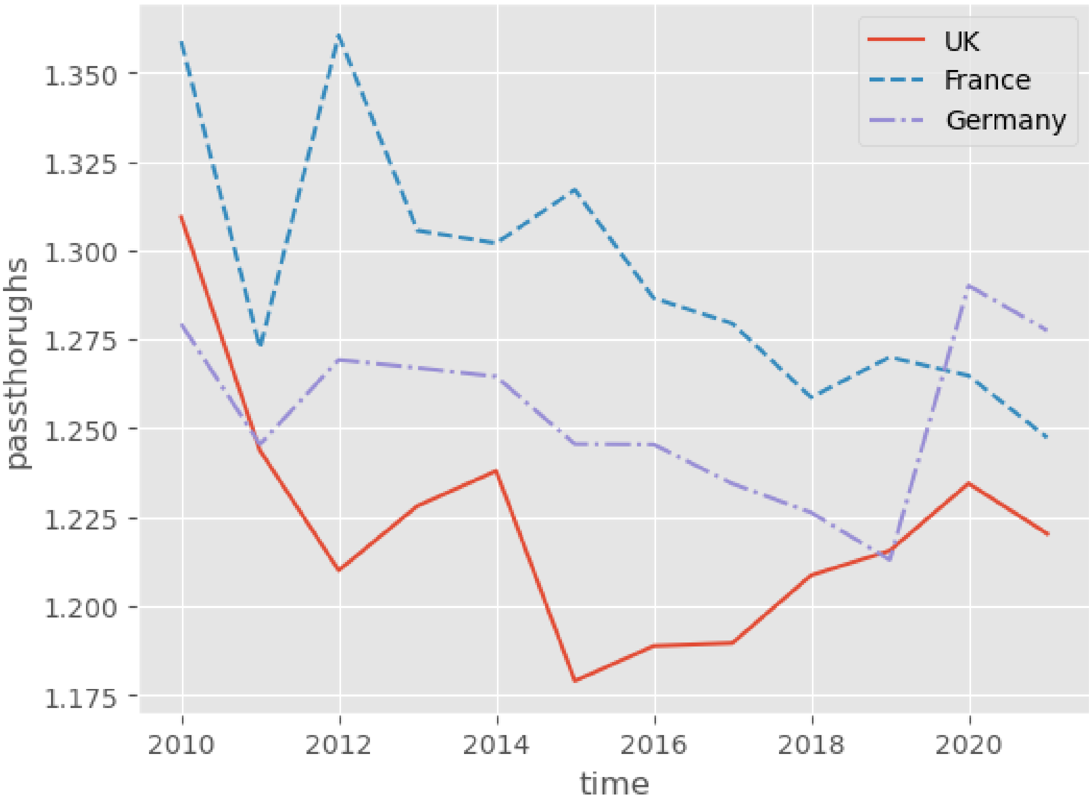

Figure 22: Average Pass-through

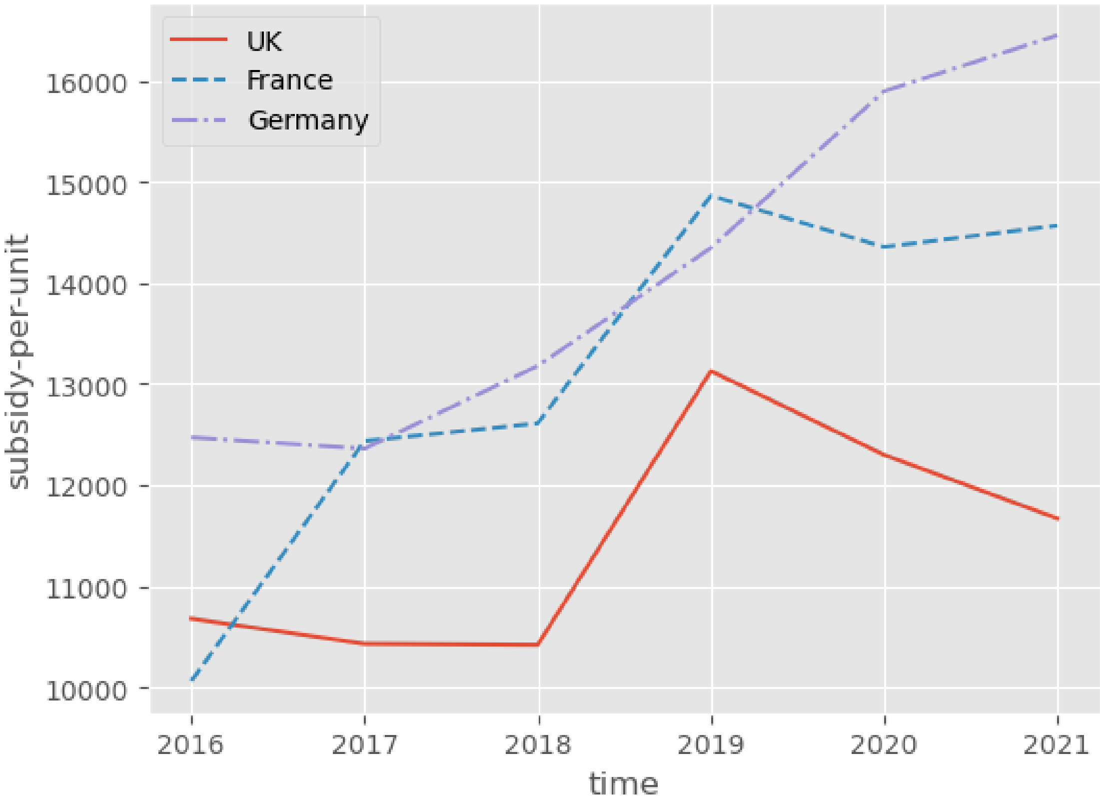

Figure 23: Average *subsidy-per-unit* 2016-2021

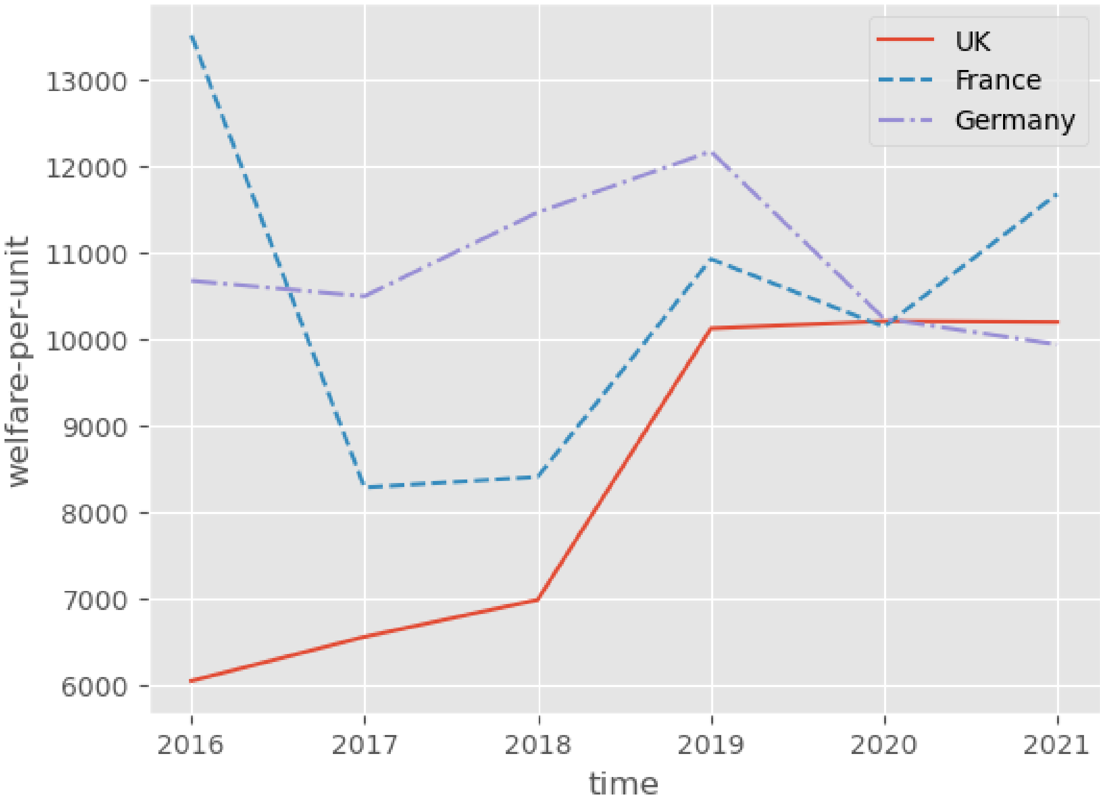

Figure 24: Average *welfare-per-unit* 2016-2021

References

[1] Victor Aguirregabiria and Chun-Yu Ho. A dynamic oligopoly game of the us airline industry: Estimation and policy experiments. Journal of Econometrics, 168(1):156–173, 2012.

[2] Jorge Ale-Chilet, Cuicui Chen, Jing Li, and Mathias Reynaert. Colluding against environmental regulation. Technical report, Toulouse School of Economics (TSE), 2021.

[3] Soren T Anderson and James M Sallee. Using loopholes to reveal the marginal cost of regulation: The case of fuel-economy standards. American Economic Review, 101(4):1375–1409, 2011.

[4] Soren T Anderson and James M Sallee. Designing policies to make cars greener. Annual Review of Resource Economics, 8(1):157–180, 2016.

[5] James Archsmith, Erich Muehlegger, and David S Rapson. Future paths of electric vehicle adoption in the united states: predictable determinants, obstacles, and opportunities. Environmental and Energy Policy and the Economy, 3 (1):71–110, 2022.

[6] Nicholas Argyres, Ricard Gil, and Giorgio Zanarone. Does supply concentration encourage cooperation? evidence from airlines. Management Science, 2024.

[7] Sarah Armitage and Frank Pinter. Regulatory mandates and electric vehicle product variety. Technical report, Working paper, 2021.

[8] Itai Ater and Nir S Yoseph. The impact of environmental fraud on the used car market: Evidence from dieselgate. The Journal of Industrial Economics, 70(2): 463–491, 2022.

[9] Ruediger Bachmann, Gabriel Ehrlich, Ying Fan, Dimitrije Ruzic, and Benjamin Leard. Firms and collective reputation: a study of the volkswagen emissions

scandal. Technical report, National Bureau of Economic Research, 2019.

[10] Florian Barth, Christian Eckert, Nadine Gatzert, and Hendrik Scholz. Spillover effects from the volkswagen emissions scandal: An analysis of stock and corporate bond markets. Schmalenbach Journal of Business Research, 74(1):37–76, 2022.

[11] Angela Stefania Bergantino, Christian Bontemps, Mario Intini, and Ada Spiru. Toward a consolidation of the european airline sector: the potential merger between ita and lufthansa. In Air Transport Research Society World Conference, 2024.

[12] Steven Berry. Estimating discrete-choice models of product differentiation. The RAND Journal of Economics, pages 242–262, 1994.

[13] Steven Berry and Panle Jia. Tracing the woes: An empirical analysis of the airline industry. American Economic Journal: Microeconomics, 2(3):1–43, 2010.

[14] Steven T Berry. Estimation of a model of entry in the airline industry. Econometrica: Journal of the Econometric Society, pages 889–917, 1992.

[15] Steven T Berry. Estimating discrete-choice models of product differentiation. The RAND Journal of Economics, pages 242–262, 1994.

[16] Steven T Berry and Philip A Haile. Identification in differentiated products markets using market level data. Econometrica, 82(5):1749–1797, 2014.

[17] Steven T Berry and Philip A Haile. Nonparametric identification of differentiated products demand using micro data. Econometrica, 92(4):1135–1162, 2024.

[18] Steven T Berry, James A Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium: Part i and ii, 1993.

[19] Sofia Berto Villas-Boas. Vertical relationships between manufacturers and retailers: Inference with limited data. The Review of Economic Studies, 74(2):

625–652, 2007.

[20] Christian Bontemps, Cristina Gualdani, and Kevin Remmy. Price competition and endogenous product choice in networks: Evidence from the us airline industry. Technical report, Working Paper, 2023.

[21] Christian Bontemps, Gianmaria Martini, and Flavio Porta. The effects of lccs subsidies on the tourism industry. TSE Working Paper, 2024.

[22] Severin Borenstein and Lucas W Davis. The distributional effects of us clean energy tax credits. Tax Policy and the Economy, 30(1):191–234, 2016.

[23] Randy Brenkers and Frank Verboven. Liberalizing a distribution system: the european car market. Journal of the European Economic Association, 4(1): 216–251, 2006.

[24] Timothy F Bresnahan. Competition and collusion in the american automobile industry: The 1955 price war. Journal of Industrial Economics, pages 457–482, 1987.

[25] N Scott Cardell. Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. Econometric Theory, 13(2):185–213, 1997.

[26] Gary Chamberlain. Asymptotic efficiency in estimation with conditional moment restrictions. Journal of econometrics, 34(3):305–334, 1987.

[27] Xiaogang Che, Hajime Katayama, and Peter Lee. Willingness to pay for brand reputation: Lessons from the volkswagen emissions scandal. Technical report, Working Paper, Durham University, 2018.

[28] Federico Ciliberto and Elie Tamer. Market structure and multiple equilibria in airline markets. Econometrica, 77(6):1791–1828, 2009.

[29] Federico Ciliberto and Jonathan W Williams. Does multimarket contact facilitate tacit collusion? inference on conduct parameters in the airline industry.

The RAND journal of economics, 45(4):764–791, 2014.

[30] Federico Ciliberto, Charles Murry, and Elie Tamer. Market structure and competition in airline markets. Journal of Political Economy 129.11, pages 2995–3038, 2021.

[31] Christopher Conlon and Jeff Gortmaker. Best practices for differentiated products demand estimation with pyblp. The RAND Journal of Economics, 51(4): 1108–1161, 2020.

[32] Christopher Conlon and Jeff Gortmaker. Incorporating micro data into differentiated products demand estimation with pyblp. Technical report, National Bureau of Economic Research, 2023.

[33] Jan De Loecker, Jan Eeckhout, and Gabriel Unger. The rise of market power and the macroeconomic implications. The Quarterly Journal of Economics, 135 (2):561–644, 2020.

[34] Alon Eizenberg. Upstream innovation and product variety in the us home pc market. Review of Economic Studies, 81(3):1003–1045, 2014.

[35] Silke Januszewski Forbes and Mara Lederman. Adaptation and vertical integration in the airline industry. American Economic Review, 99(5):1831–1849, 2009.

[36] Amit Gandhi and Jean-François Houde. Measuring substitution patterns in differentiated-products industries. NBER Working paper, (w26375), 2019.

[37] Paul LE Grieco, Charles Murry, and Ali Yurukoglu. The evolution of market power in the us automobile industry. The Quarterly Journal of Economics, 139 (2):1201–1253, 2024.

[38] Paul A Griffin and David H Lont. Game changer? the impact of the vw emission-cheating scandal on the interrelation between large automakers' equity and credit markets. Journal of Contemporary Accounting & Economics, 14(2):

179–196, 2018.

[39] Laura Grigolon and Frank Verboven. Nested logit or random coefficients logit? a comparison of alternative discrete choice models of product differentiation. Review of Economics and Statistics, 96(5):916–935, 2014.

[40] Jerry A Hausman. An instrumental variable approach to full information estimators for linear and certain nonlinear econometric models. Econometrica: Journal of the Econometric Society, pages 727–738, 1975.

[41] Kate Ho and Ariel Pakes. Hospital choices, hospital prices, and financial incentives to physicians. American Economic Review, 104(12):3841–3884, 2014.

[42] Stephen P Holland, Jonathan E Hughes, and Christopher R Knittel. Greenhouse gas reductions under low carbon fuel standards? American Economic Journal: Economic Policy, 1(1):106–146, 2009.

[43] Jean-François Houde, Peter Newberry, and Katja Seim. Nexus tax laws and economies of density in e-commerce: A study of amazon's fulfillment center network. Econometrica, 91(1):147–190, 2023.

[44] Cheng-Ta Hsieh, Sheng-Chi Lin, Chih-Hung Lee, Ching-Fang Liu, and Chi-Chang Hu. Designing multifunctional polyethylene-polyimide composite separators for rechargeable lithium-ion batteries. Journal of The Electrochemical Society, 166(14):A3132, 2019.

[45] Jonathan E Hughes and Molly Podolefsky. Getting green with solar subsidies: evidence from the california solar initiative. Journal of the Association of Environmental and Resource Economists, 2(2):235–275, 2015.

[46] Koichiro Ito and James M Sallee. The economics of attribute-based regulation: Theory and evidence from fuel economy standards. Review of Economics and Statistics, 100(2):319–336, 2018.

[47] Matthew O Jackson and Asher Wolinsky. A strategic model of social and

economic networks. Journal of economic theory, 71(1):44–74, 1996.

[48] Thomas Klier and Joshua Linn. New-vehicle characteristics and the cost of the corporate average fuel economy standard. The RAND Journal of Economics, 43(1):186–213, 2012.

[49] Jing Li. Compatibility and investment in the us electric vehicle market. Technical report, Berkerly Haas, 2019.

[50] Joshua Linn. Balancing equity and effectiveness for electric vehicle subsidies. Technical report, Resources for the Future, 2022.

[51] Marleen Marra. A market for airport slots. Available at SSRN, 2024.

[52] W Ross Morrow and Steven J Skerlos. Fixed-point approaches to computing bertrand-nash equilibrium prices under mixed-logit demand. Operations research, 59(2):328–345, 2011.

[53] Erich Muehlegger and David S Rapson. Subsidizing low-and middle-income adoption of electric vehicles: Quasi-experimental evidence from california. Journal of Public Economics, 216:104752, 2022.

[54] Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. Econometrica, 69(2):307–342, 2001.

[55] Volker Nocke and Nicolas Schutz. Multiproduct-firm oligopoly: An aggregative games approach. Econometrica, 86(2):523–557, 2018.

[56] Ariel Pakes, Jack Porter, Kate Ho, and Joy Ishii. Moment inequalities and their application. Econometrica, 83(1):315–334, 2015.

[57] Yongjoon Park. Structural remedies in network industries: An assessment of slot divestitures in the american airlines/us airways merger. US Airways Merger (October 12, 2020), 2020.

[58] Amil Petrin. Quantifying the benefits of new products: The case of the minivan. Journal of political Economy, 110(4):705–729, 2002.

[59] Jacquelyn Pless and Arthur A Van Benthem. Pass-through as a test for market power: An application to solar subsidies. American Economic Journal: Applied Economics, 11(4):367–401, 2019.

[60] David Rapson and Erich Muehlegger. Global transportation decarbonization. Journal of Economic Perspectives, 37(3):163–188, 2023.

[61] Kevin Remmy. Adjustable product attributes, indirect network effects, and subsidy design: The case of electric vehicles. Technical report, University of Bonn and University of Mannheim, Germany, 2022.

[62] Mathias Reynaert. Abatement strategies and the cost of environmental regulation: Emission standards on the european car market. The Review of Economic Studies, 88(1):454–488, 2021.

[63] Mathias Reynaert and Frank Verboven. Improving the performance of random coefficients demand models: The role of optimal instruments. Journal of econometrics, 179(1):83–98, 2014.

[64] Jo Reynaerts, R Varadha, and JC Nash. Enhencing the convergence properties of the blp (1995) contraction mapping. Technical report, KU Leuven, Faculty of Economics and Business (FEB), VIVES-Research Centre . . . , 2012.

[65] Stephen P Ryan. The costs of environmental regulation in a concentrated industry. Econometrica, 80(3):1019–1061, 2012.

[66] Alejandro Sabal. Product entry in the global automobile industry. Technical report, National Bureau of Economic Research, 2025.

[67] Bernard Salanié and Frank A Wolak. Fast, detail-free, and approximately correct: Estimating mixed demand systems. Technical report, Working paper, 2022.

[68] Katalin Springel. Network externality and subsidy structure in two-sided markets: Evidence from electric vehicle incentives. American Economic Journal:

Economic Policy, 13(4):393–432, 2021.

[69] Anthony Strittmatter and Michael Lechner. Sorting in the used-car market after the volkswagen emission scandal. Journal of Environmental Economics and Management, 101:102305, 2020.

[70] Andrew Sweeting. Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. Econometrica, 81(5):1763–1803, 2013.

[71] Govinda R Timilsina. Carbon taxes. Journal of Economic Literature, 60(4): 1456–1502, 2022.

[72] Ravi Varadhan and Christophe Roland. Simple and globally convergent methods for accelerating the convergence of any em algorithm. Scandinavian Journal of Statistics, 35(2):335–353, 2008.

[73] Frank Verboven and Theon van Dijk. Cartel damages claims and the passing-on defense. The Journal of Industrial Economics, 57(3):457–491, 2009.

[74] E Glen Weyl and Michal Fabinger. Pass-through as an economic tool: Principles of incidence under imperfect competition. Journal of political economy, 121(3): 528–583, 2013.

[75] Kate S Whitefoot and Steven J Skerlos. Design incentives to increase vehicle size created from the us footprint-based fuel economy standards. Energy Policy, 41:402–411, 2012.

[76] Kate S Whitefoot, Meredith L Fowlie, and Steven J Skerlos. Compliance by design: Influence of acceleration trade-offs on co2 emissions and costs of fuel economy and greenhouse gas regulations. Environmental science & technology, 51(18):10307–10315, 2017.

[77] Jianwei Xing, Benjamin Leard, and Shanjun Li. What does an electric vehicle replace? Journal of Environmental Economics and Management, 107:102432,

2021.

[78] Zhe Yuan and Panle Jia Barwick. Network competition in the airline industry: An empirical framework. Technical report, National Bureau of Economic Research, 2024.

[79] Paul Zimmerman and Julie A Carlson. Competition and cost pass-through in differentiated oligopolies. Technical report, University Library of Munich, Germany, 2010.