
Firm Strategies and Consumer Behavior under Market-Based Sustainability Policies

vorgelegt von
Helene Blanche Naegele, M. A.
geb. in Berlin

von der Fakultät VII – Wirtschaft und Management
der Technischen Universität Berlin
zur Erlangung des akademischen Grades
Doktor der Wirtschaftswissenschaften
– Dr. rer. oec. –

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr. Georg Meran

Gutachter: Prof. Dr. Pio Baake

Dr. habil. Pauline Givord (Centre de recherche en
économie et statistique (CREST), Paris)

Prof. Dr. Christian von Hirschhausen

Tag der wissenschaftlichen Aussprache: 8. Dezember 2017

Berlin 2018

*To my parents,
who made everything possible.*

Abstract

This dissertation studies different approaches to sustainability regulation. The author evaluates the effectiveness of several policy measures in achieving their objective and discusses potential unintended side-effects. Both depend on strategic firm behavior and consumer sensitivity.

Chapter II considers the impact of fuel taxes on the fuel efficiency of new automobiles. The results rely on exhaustive consumer-level data of French automobile purchases; the author estimates the parameters of automobile demand using a nested logit model accounting for heterogeneity of consumer groups. The estimated parameters are used to predict the impact of two hypothetical changes in fuel taxes. The results show that new vehicle purchases react very little to changes in fuel tax and the impact of the examined fuel tax policies is economically small.

Chapter III studies the strategic interactions between two sustainability labels competing for firms to offer their products: it shows why the industry has an interest to introduce its own label. The results rely on a model of oligopolistic firms offering several horizontally (between firms) and vertically (between label qualities) differentiated products. These firms interact with two labeling organizations that pursue different objectives: a for-profit label and an industry standard maximizing joint firm profit. The results show that the industry benefits from introducing an industry standard that reduces competition by segmenting the market; by contrast, a social planner maximizes consumer welfare and total social welfare by maximizing the number of labeled goods. Horizontal differentiation plays a key role for the final market outcome.

Chapter IV measures the magnitude of fixed transaction costs in European emissions trading. Transaction costs are defined in a broad sense as monetary and non-monetary frictions from certificate trading of firms. The results rely on plant-level administrative data from European emissions trading; the author estimates the distribution of fixed transaction costs arising from the use of “normal” European certificates, on the one hand, and international offset certificates, on the other hand. The results show that for most firms, the bulk of transaction costs stems from market participation in general rather than from the use of international certificates. The magnitude of transaction costs is such that a fifth of all firms does not participate in profitable offset trading.

Chapter V studies whether European emissions trading has led to a displacement of European carbon emissions to other parts of the world (“carbon leakage”), both via relocation and via loss of market shares to foreign competitors. A literature survey reveals different approaches to identify carbon leakage empirically. This chapter’s results rely on a combination of sector-level trade data and plant-level data from European emissions trading; using various ways of defining both outcome and stringency of environmental

policy, the author finds no evidence of carbon leakage.

Keywords: automobile demand, carbon dioxide, carbon leakage, climate change, consumer labels, demand estimation, emissions trading, EU ETS, externalities, fuel tax, nested logit model, sustainability, transaction costs.

Zusammenfassung

Diese Dissertation untersucht unterschiedliche Ansätze zur Nachhaltigkeitsregulierung. Die Autorin bewertet die Wirksamkeit von Politikmaßnahmen sowie potenzielle unbeabsichtigte Nebenwirkungen. Beide hängen vom strategischen Verhalten von Unternehmen sowie der Sensitivität der Verbraucher ab.

Kapitel II betrachtet die Auswirkung von Treibstoffsteuern auf die Kraftstoffeffizienz von neuen Automobilen. Die Ergebnisse beruhen auf Verbraucherdaten über französischen Automobilkäufe; die Autorin schätzt die Parameter der Automobilnachfrage mit einem Nested Logit Modell, das die Heterogenität von Verbrauchergruppen berücksichtigt. Die geschätzten Parameter werden verwendet, um die Auswirkung von zwei hypothetischen Änderungen der Treibstoffsteuer zu berechnen. Die Ergebnisse zeigen, dass neue Fahrzeugkäufe sehr wenig auf Änderungen der Treibstoffsteuer reagieren und die Auswirkungen der untersuchten Steuerreformen wirtschaftlich vernachlässigbar sind.

Kapitel III untersucht die strategischen Interaktionen zwischen zwei Nachhaltigkeitslabels im Kaffeemarkt: es zeigt, warum die Röstereibranche ein Interesse daran hat, ihr eigenes Label einzuführen. Die Ergebnisse beruhen auf einem Modell mit Firmen im Oligopol, die mehrere horizontal (zwischen Firmen) und vertikal (zwischen Labelqualitäten) differenzierte Produkte anbieten. Diese Unternehmen interagieren mit zwei Labelorganisationen, die unterschiedliche Ziele verfolgen: ein profitorientierter Lizenzierer und ein Industriestandard, der den gemeinsamen Unternehmensgewinn maximiert. Die Ergebnisse zeigen, dass ein Industriestandard immer versucht, den Wettbewerb durch Segmentierung des Marktes zu reduzieren. Im Gegensatz dazu maximiert ein sozialer Planer die gesamtgesellschaftliche Wohlfahrt durch die Maximierung der Zahl der gelabelten Produkte. Die horizontale Differenzierung spielt eine entscheidende Rolle für das Marktergebnis.

Kapitel IV misst Transaktionskosten im europäischen Emissionshandel. Transaktionskosten werden hier im weiten Sinne als monetäre und nicht monetäre Aufwände von Firmen im Zertifikatshandel definiert. Die Ergebnisse beruhen auf administrativen Daten des europäischen Emissionshandels; die Autorin schätzt die Verteilung der fixen Transaktionskosten, die sich aus der Verwendung von "normalen europäischen Zertifikaten" einerseits und internationalen Offsetzertifikaten andererseits ergeben. Die Ergebnisse zeigen, dass für die meisten Unternehmen der Großteil der Transaktionskosten von der Marktbeteiligung im Allgemeinen und nicht von der Verwendung internationaler Offsetzertifikate stammt. Die Größenordnung der Transaktionskosten ist so, dass ein Fünftel aller Firmen nicht am gewinnbringenden Offsethandel teilnimmt.

Kapitel V untersucht, ob der europäische Emissionshandel zu einer Verschiebung der europäischen CO₂-Emissionen in andere Teile der Welt geführt hat ("carbon leakage"), sowohl durch Verlagerung der Produktion als auch durch Verlust von Marktanteilen an ausländische Wettbewerber. Eine Literaturrecherche zeigt verschiedene Ansätze zur Iden-

tifizierung von carbon leakage in den Daten. Die Ergebnisse dieses Kapitels beruhen auf einer Kombination von Handelsdaten und Daten aus dem europäischen Emissionshandel; mit verschiedenen Definitionen der Ergebnisvariablen sowie der Emissionskosten, findet die Autorin keine Hinweise auf carbon leakage.

Schlüsselwörter: Automobilnachfrage, Emissionshandel, EU ETS, Externalitäten, Klimawandel, Kohlendioxid, Konsumentenlabels, Nachfrageschätzung, Nachhaltigkeit, Nested Logit Modell, Transaktionskosten, Treibstoffsteuer.

Acknowledgements

First and foremost, I would like to thank my supervisors: Pio Baake for giving me the freedom to set my own research agenda, for sharpening my mind with seemingly innocent questions, and for his company and the DIW's best coffee; Christian von Hirschhausen for his passionate policy opinions, for his unconditional support even when my research agenda was distant from his, for being an excellent coach and constantly encouraging me to "start working".

I further thank Pauline Givord for her availability despite all other responsibilities, for her sharp mind and for traveling to Berlin for my defense; Jana Friedrichsen for believing in a better world; Aleksandar Zaklan for his patience and infinite proof-reading; Bernd Fitzenberger for advice on quantile regressions that so few people seem to grasp; Stéphane Caprice and Dennis Rickert for their warm welcome during my research stay at the Toulouse School of Economics; Sylvain Chabé-Ferret for his metaphysical thoughts about standards for good science; Tobias Schmidt for being an inspiring colleague even though he ran away too quickly; and Roland Rathelot for pushing me to pursue a dissertation.

I also thank my friends for supporting me along the way; in particular, Kathrin for constantly challenging me; Frieda for shaping who I am; and Mathis for patiently putting up with what mathematics look like in my head.

Many thanks to Johanna Möllerström and the entire department *Competition and Consumers*. This dissertation would not have been possible without the financial support from the DIW Graduate Center and the Leibniz Gemeinschaft.

Contents

Title page	1
Abstract	iii
Zusammenfassung	v
Contents	vii
List of Tables	xi
List of Figures	xiii
I Introduction	1
1 Sustainability policies	3
1.1 Pollution tax	5
1.2 Cap-and-trade	6
1.3 Private voluntary standards	8
2 Methodology	10
2.1 Nested logit	10
2.2 Binary quantile estimation	11
2.3 Treatment effects	12
3 Contribution	13
3.1 Fuel taxes and automobile fuel efficiency	14
3.2 Competition between sustainability labels	16
3.3 Transaction costs in European emissions trading	17
3.4 Carbon leakage and European emissions trading	17
4 Concluding remarks	18
5 Bibliography	20
II How do fuel taxes impact new car purchases? An evaluation using French consumer-level data	23
1 Introduction	24
2 Choice model	26

3	Data and descriptive evidence	28
3.1	New vehicle registrations	28
3.2	Types of consumers: demographic groups	29
3.3	Diesel and gasoline cars	31
3.4	Cost per kilometer	34
4	Econometric approach	35
4.1	Nested logit estimation	35
4.2	Endogenous variables and instruments	36
5	Empirical results	37
5.1	Aggregate elasticities to fuel price variation	38
5.2	Tax alignment	38
5.3	Carbon tax	40
5.4	Robustness checks	41
6	Conclusion	42
7	Bibliography	45
	Appendices	49
A	Descriptive statistics	49
B	Details on the computation of the elasticities	50
C	Complementary results for the main specification	55
C.1	Raw coefficients	55
C.2	Demand for selected car models	56
D	Testing for weak instruments	60
E	Robustness checks: elasticities	62

III Competition between for-profit and industry labels: the case of social labels in the coffee market 63

1	Introduction	64
1.1	Related literature	65
1.2	Coffee market and fairtrade	66
2	Model	67
2.1	Consumer demand	67
2.2	Firms	70
2.3	Labeling organizations	71
2.4	Game sequence	71
3	Market equilibrium with licensor F only	72
3.1	Consumer prices	72
3.2	Product line decisions	73
3.3	License fee	73
3.4	Label quality	75
3.5	Minimum quality requirement	76
4	Market entry of industry standard I	77
4.1	Consumer prices	77
4.2	Product line decisions	78

4.3	License fee	78
4.4	Label quality of incumbent licensor F	79
4.5	Label quality of new entrant I	81
4.6	Minimum quality requirement	85
5	Conclusion	85
6	Bibliography	87
	Appendices	90
A	Proof of Lemma 1	90
B	Proof of Lemma 2	91
C	Calculations for Proposition 2	92
D	Calculations for Proposition 3	93
E	Calculations for Proposition 4.(iii)	93
 IV Offset credits in the EU ETS: a quantile estimation of firm-level transaction costs		
	costs	95
1	Introduction	96
2	Background	99
2.1	Institutional framework	99
2.2	Why are offset certificates cheaper?	100
2.3	Definition and interpretation of transaction costs	101
3	Model	103
3.1	Emissions trading with offset credits: reference scenario	103
3.2	Trading costs for both certificate markets	105
4	Data and empirical research design	106
4.1	Emissions, allocation and offset entitlement	106
4.2	Price spread and realized savings	107
4.3	Descriptive evidence for transaction costs	109
4.4	Econometric methodology	111
5	Estimation results	112
5.1	Sector-specific results	115
6	Conclusion	116
7	Bibliography	119
	Appendices	122
A	National offset entitlement rules	122
B	Offset use as a fixed horizon problem	123
C	Exogeneity of allocation status	124
D	Parametric estimation results	126
E	Quantile regression fit	128
 V Does the EU ETS cause carbon leakage in European manufacturing?		
1	Introduction	130
2	Trade theory and carbon leakage	133
2.1	(Neo-)classical approach	133

2.2	New trade approach	135
3	Empirical Implementation	136
3.1	Measures of environmental stringency	136
3.2	Identification	138
3.3	Net flows	139
3.4	Bilateral flows	140
4	Data and descriptives	142
4.1	Data	142
4.2	Descriptive statistics	142
5	Results	146
5.1	Net trade flows	146
5.1.1	Main results	146
5.1.2	Sector heterogeneity	147
5.2	Bilateral trade flows	149
6	Summary and conclusions	150
7	Bibliography	153
	Appendices	157
A	Potential bias from using normalized variables	157
B	Additional regression results	159
B.1	Net trade flows	159
B.2	Bilateral trade flows	160

List of Tables

Table	Page
I.1 Overview by chapter: topic, pre-publication and author's contribution	15
II.1 Descriptive statistics: main characteristics of new car registrations 2003-2007 .	29
II.2 Descriptive statistics: average mileage by household characteristics	30
II.3 Elasticities with respect to fuel prices: diesel share, average fleet fuel consumption (L/km) and CO ₂ intensity (g/km)	39
II.4 Percentage impact of a carbon tax and a tax alignment on diesel share, average fleet fuel consumption (L/km) and CO ₂ intensity (g/km)	40
II.5 Robustness checks: percentage impact of carbon tax and tax alignment on diesel share, average fleet fuel consumption (L/km) and CO ₂ intensity (g/km)	41
II.6 Distribution of demographic groups among buyers (%)	49
II.7 Descriptive statistics: car characteristics	50
II.8 Estimates for the coefficient on cost per km β_d	55
II.9 Estimates for the coefficient σ_{1d} (substitutability within model, between engine types)	57
II.10 Estimates for the coefficient σ_{2d} (substitutability within segment, between models)	58
II.11 Estimates for the coefficient on vehicle price γ_d	59
II.12 Demand elasticity for selected models with respect to fuel prices	59
II.13 Conditional F-values of the weak instrument test – instruments for the <i>price</i> . .	60
II.14 Conditional F-values of the weak instrument test – instruments for the <i>market share of the model within its segment</i> $s_{dj s}$	60
II.15 Conditional F-values of the weak instrument test – instruments for the <i>market share of a fuel-type within its model nest</i> $s_{df j}$	61
II.16 Robustness checks: elasticities with respect to fuel prices of diesel share, average fleet fuel consumption (L/km) and CO ₂ intensity (g/km)	62
III.1 Minimum quality requirement \underline{q} set by the social planner as a function of horizontal differentiation μ	85
IV.1 Overview of transaction cost estimates per firm in the EU ETS in the literature	103
IV.2 Descriptive firm statistics	108

IV.3	Estimates from binary quantile estimation and probit regression	114
IV.4	Offset limits from National Allocation Plans	122
IV.5	Parametric mean estimates for transaction costs	127
IV.6	Specification test of binary regression quantile models (predicted and observed probabilities)	128
V.1	Descriptive statistics	143
V.2	Regression results for EU net imports	147
V.3	Regressions of net import flows on environmental cost and its interaction with transport costs and higher order terms of emission cost	148
V.4	Bilateral trade flows in logs of million U.S. dollars and embodied CO ₂ emissions	151
V.5	Regression of net trade flows on emission cost with additional control for scaling variables	158
V.6	Summary overview stating only the coefficient of the ETS stringency variables (for different specifications of net flows)	159
V.7	Summary overview stating only the coefficient of the <i>bilateral ETS stringency variables</i> (for different specifications of bilateral flows)	160
V.8	Summary overview stating only the coefficient of the <i>two separate ETS stringency variables</i> (for different specifications of bilateral flows)	161

List of Figures

Figure	Page
II.1 Nested decision-making structure of the car purchaser	27
II.2 Overview of spatial variation in share of diesel cars and mileage	30
II.3 Diesel fuel prices and market shares in Europe in 2012	32
II.4 Monthly fuel prices and cost per km	32
II.5 Monthly new registrations by fuel-type	33
III.1 Nested decision-making structure of consumers	68
III.2 Product line equilibrium as a function of license fee L_0 and horizontal differ- entiation μ	74
III.3 Equilibrium quality q_0^{F*} as a function of horizontal differentiation μ	76
III.4 Minimum quality requirement \underline{q} as a function of horizontal differentiation μ .	77
III.5 Reaction function of quality q_1^{F*} as a function of industry standard quality q^I .	80
III.6 Preferred product range equilibrium of licenser F as a function of quality q^I and horizontal differentiation μ	81
III.7 Equilibrium qualities q^{I*} and q_1^{F*} as a function of horizontal differentiation μ .	84
III.8 License fee as a function of label quality q_1^F for $\mu = 0.4$ and $q^I = 0.07$	92
III.9 Joint firm profit Π as a function of industry standard q^I given licenser reaction $q_1^{F*}(q^I)$ for $\mu = 0.4$	93
III.10 Comparing license fees from $t = 0$ and $t = 1$	94
IV.1 Stylized illustration of aggregate market equilibrium	101
IV.2 Prices of European certificates (EUA) and offsets (CER)	108
IV.3 Ratio of used offset credits over overall offset entitlement	109
IV.4 Relationship between offset use, offset entitlement and net allocation status . .	110
IV.5 Estimated transaction cost (in €) - quantile plot	115
IV.6 Estimated transaction costs (in €) - distribution function and density	115
IV.7 Observed frequencies and predicted probabilities of quantile method and probit	116
IV.8 Sector-specific quantile estimation results (in €)	117
IV.9 McCrary's test for continuity	125
V.1 Stylized illustration of the pollution haven hypothesis	134

V.2	Correlation of the measures of environmental stringency (scatter and fitted linear trend, sector-year averages across countries, 2007 & 2011)	144
V.3	Distribution of import taxes and transport costs (sector-averages 2004-2011) . .	145
V.4	Imports in value and imports in embodied carbon by bilateral EU ETS treatment status	145

Introduction

Un Chat, nommé Rodilardus,
 Faisait de Rats telle déconfiture
 Que l'on n'en voyait presque plus,
 Tant il en avait mis dedans la sépulture.
 Le peu qu'il en restait, n'osant quitter son trou,
 Ne trouvait à manger que le quart de son soû ;
 Et Rodilard passait, chez la gent misérable,
 Non pour un Chat, mais pour un Diable.
 Or, un jour qu'au haut et au loin
 Le Galand alla chercher femme,
 Pendant tout le sabbat qu'il fit avec sa dame,
 Le demeurant des Rats tint chapitre en un coin
 Sur la nécessité présente.
 Dès l'abord, leur Doyen, personne fort prudente,
 Opina qu'il fallait, et plus tôt que plus tard,
 Attacher un grelot au cou de Rodilard ;
 Qu'ainsi, quand il irait en guerre,
 De sa marche avertis ils s'enfuiraient sous terre ;
 Qu'il n'y savait que ce moyen.
 Chacun fut de l'avis de Monsieur le Doyen ;
 Chose ne leur parut à tous plus salutaire.
 La difficulté fut d'attacher le grelot.
 L'un dit : Je n'y vas point, je ne suis pas si sot ;
 L'autre : Je ne saurais. Si bien que sans rien faire
 On se quitta.
 Jean de la Fontaine (1668)

Old Rodilard, a certain cat,
 Such havoc of the rats had made,
 'Twas difficult to find a rat
 With nature's debt unpaid.
 The few that did remain,
 To leave their holes afraid,
 From usual food abstain,
 Not eating half their fill.
 And wonder no one will
 That one who made of rats his revel,
 With rats pass'd not for cat, but devil.
 Now, on a day, this dread rat-eater,
 Who had a wife, went out to meet her;
 And while he held his caterwauling,
 The unkill'd rats, their chapter calling,
 Discuss'd the point, in grave debate,
 How they might shun impending fate.
 Their dean, a prudent rat,
 Thought best, and better soon than late,
 To bell the fatal cat;
 That, when he took his hunting round,
 The rats, well caution'd by the sound,
 Might hide in safety under ground;
 Indeed he knew no other means.
 And all the rest
 At once confess'd
 Their minds were with the dean's.
 No better plan, they all believed,
 Could possibly have been conceived.
 No doubt the thing would work right well,
 If any one would hang the bell.
 But, one by one, said every rat,
 "I'm not so big a fool as that."
 The plan knock'd up in this respect,
 The council closed without effect.
 Translation by Elizur Wright (1882)

Our society is facing great challenges. Among them, climate change has become more imminent but days where it makes *telle déconfiture*¹ are still impending. Other consequences of the lack of environmental sustainability can be easily felt today. While working on my PhD at the DIW Graduate Center, I traveled to Beijing and coughed my lungs out every day. There was so much air pollution in Teheran that I could not see the base of the TV tower while standing on top of it. I could barely breathe in Mexico City.

In Bogotá, I tried to find a fairtrade farmer, knocking on many doors in vain: it seems the average fairtrade farmer is too poor to have a postal address.² In Amman, my friend

¹French for: total collapse

²I ended up in front of the presidential palace because one farmer wrote down the only postal address

told me her broom closet was originally included in order to house a modern-day slave.

Someone must start taking responsibility, without saying “*Je ne suis pas si sot*”³ or “*Je ne saurais*.”⁴ Some steps are clear, like the need to reduce pollution or to find a *grelot*,⁵ while others much less so, like how to sustain comfortable living standards in the meantime. With commitment and perseverance, I am convinced solutions can be found that move toward a more sustainable economic system: for this, we need to evaluate the possible paths to sustainability, but not just the desirability of the outcome.

1 Sustainability policies

This dissertation assesses, empirically and theoretically, different market-based approaches to sustainability regulation. In particular, I ask whether the regulation alleviates the problem it targets and at what cost. This section introduces the fundamental issues behind sustainability regulation, gives a historic overview of regulatory strategies, and briefly presents the approaches studied in this dissertation. Section 2 then summarizes the methods used in this dissertation and Section 3 provides a summary of each chapter’s contributions; Section 4 concludes.

Sustainability signifies the ability of a system to remain stable and productive indefinitely. This concept contrasts with an economic system that is depleting economic and natural resources and may collapse after depletion. Generally, three aspects of sustainability of production systems are stressed: social well-being, environmental integrity, and economic enterprise. Sustainability regulation often focuses on externalities as the main challenge to environmental sustainability.

An externality arises whenever one party’s actions affect another party, without the first party incurring any cost for this effect on the second: a side-effect of an (economic) activity. A typical example is a production plant polluting a river, such that downstream users suffer from poor water quality without being responsible for, or benefiting from, the polluting activity. In the context of air pollution, Crocker (1966, p. 62) describes the externality phenomenon as the “divorce of emission costs from emission benefits.” Useful concepts to think about negative externalities⁶ include private cost, damage, and social cost. Private cost refers to the polluting party’s cost, for example the cost of production inputs such as labor and capital. Damage refers to the additional cost caused to the party suffering from the externality. Finally, social cost is the sum of damage and private cost.

Typical economic analysis shows that a difference between marginal private cost and marginal social cost is harmful to social welfare: a producer not held responsible for the

he knew.

³French for: “I am not so stupid [to sacrifice].”

⁴French for: “I’m not able.”

⁵French for: bell

⁶As the bulk of environmental economics, this dissertation focuses on negative externalities; for positive externalities, one has to replace damage and cost by benefit.

polluted river water pollutes excessively. Hence the concept of *internalizing* externalities by introducing a correction that aligns the polluting party's private marginal cost with social marginal cost.

One of the major concerns for environmental sustainability in 2017 is climate change induced by anthropogenic greenhouse gas emissions. The link between sustainability and negative externalities is particularly salient here: the consequences of climate change will be felt by future generations, especially in poor countries⁷ – surely not the same agents emitting greenhouse gases today. Without regulatory intervention, emissions exceed welfare-maximizing amounts when one includes future generations' welfare in cost-benefit analysis. Estimates of marginal damage range from \$14/tCO₂e to \$350/tCO₂e (Van den Bergh 2010); other methodologies yield estimates with smaller, but still large, variance (Pindyck 2016). Externalities from greenhouse gases range from increases of heat-related human mortality, over wildfires and floods, to declining water resources, limited food security and destruction of terrestrial and marine ecosystems (IPCC 2014, p. 14).

Moreover, climate change is a global externality: no matter where greenhouse gases are emitted, their impact on the climate is the same (IPCC 2014). Regulating a global externality problem is therefore particularly difficult, as the efforts of some might be undone by free-riding of others. Local pollution, like pollution of lake water, is easier to address because one can rely on existing local institutions.

Although not usually seen as an externality, increasingly international value chains of consumer goods have the unintended side-effect of increasing social inequality. Be it textiles or commodity food products, such as cocoa and coffee, the lack of social sustainability materializes in poverty of farmers, producers and employees. Prices of food commodities are often so low and volatile that producers can barely economically sustain their business. Surprisingly, consumers commonly state that they are in principle willing to pay higher prices to reduce inequality.⁸

In this context, what regulation can lead us toward more sustainability? Such regulation has to contribute to moving toward sustainable production processes (effectiveness); ideally, regulation should be shaped such that it minimizes cost and reduces unintended side-effects (efficiency). This dissertation contributes to the search for effective sustainability policies.

The traditional approach to externalities is to forbid, in a top-down manner, the harmful activity or to prescribe a less harmful alternative, thereby directly regulating quantities of the externality. So-called *command-and-control* policies prescribe emissions standards or force adoption of best available technology. Economists have repeatedly expressed the concern that command-and-control policies lead to inefficiencies: the one-

⁷IPCC (2014, p. 31): "Climate change is a threat to sustainable development[...]. Climate change exacerbates other threats to social and natural systems, placing additional burdens particularly on the poor."

⁸In practice, the efforts however often remain limited. For example, the Fairtrade initiative has very low market shares; the highest Fairtrade market share is in the coffee market, where it amounts to about 2% of globally traded coffee in 2014 (Panhuysen and Pierrot 2014).

technology-fits-all approach does not let any room for shifting the reduction to firms that abate at the lowest cost. In response, market-based approaches have been put forward. “Market-based” policies in this context are defined as policies that regulate either quantities or prices, but leave the other one of the two free to be determined by a market equilibrium. The following subsections briefly presents Pigouvian taxes, emissions trading, and voluntary private standards as alternatives to command-and-control policies. Note that the textbook efficiency of these policies depends on a set of assumptions, in particular most models I allude to in this short introduction are static (not taking into account dynamic considerations, in particular when facing investment decisions) and ignore political economy considerations; in practice, their superiority to command-and-control policies is less clear.

1.1 Pollution tax

The classical economic thinking about externalities is fundamentally shaped by Pigou (1920),⁹ calling for *market-based* approaches using a price mechanism. Pigouvian taxes on negative externalities – or subsidies for positive externalities – have, for a long time, been the economics textbook-solution to externality problems. The tax is a straightforward form of internalizing the externality: the regulator imposes a tax equal to marginal damage to any emitter of a negative externality and gives a subsidy equal to marginal benefit to those causing a positive externality, thereby aligning private and social cost. Once the polluter incurs the cost of his externalities, the market regulates the quantities: higher production costs reduce the amount of externality by reducing the quantity of the polluting good and inciting more emission-efficient production processes. Chapter II studies such a pollution tax.

Examples include fuel taxes, tobacco taxes, deposit-refund systems, and fat taxes. While ideally the amount of the tax reflects the marginal external damage to society, in practice implementation is an “iterative tâtonnement type of planning game” (Weitzman 1974, p. 478) and above mentioned examples can hardly be seen as precise measures of marginal damage.

The first advantage of pollution taxes is that they are relatively simple to implement, because public authorities and firms typically have existing capacities for tax management. Moreover, there is discussion about a “double dividend” of both government revenue and externality reduction when revenue-raising distorting taxes are substituted by more efficient externality taxes (see the critical discussion by Fullerton and Metcalf 1997). In case of technology-dependent externalities, such as CO₂ emissions, a tax has the double effect of inducing immediate action as well as pushing for long term investments in energy-efficiency.

On the negative side, in practice it is rather difficult to determine the marginal damage of externalities, e.g. the cost of a ton of CO₂ emissions: when taxes are too low, damage control is insufficient; when taxes are too high, resources are wasted in excessive damage

⁹ As cited in Baumol and Oates (1971), who coined the term “Pigouvian tax.”

control. Pindyck (2016), for example, discusses the difficulty to measure the social cost of carbon; his survey reveals the large variance of estimates even among experts.

In his seminal paper, Weitzman (1974) compares quotas (quantity regulation) to taxes (prices). His starting point is the paradoxical observation that “the average economist in the Western marginalist tradition has at least a vague preference toward indirect control by prices, just as the typical non-economist leans toward the direct regulation of quantities[...]. Certainly a careful reading of economic theory yields little to support such a universal proposition”(Weitzman 1974, p. 477). To Weitzman’s eyes, prices and quantities are equally difficult to determine optimally: the social cost of a potential error of the policy maker varies case by case, and the choice between both policy tools should depend on this relative “cost of a mistake.”

Moreover, there is some mixed evidence that consumers myopically under-invest in future economies from energy-efficiency (Allcott and Wozny 2013): many other factors also impact the sensitivity of consumers. Thus, the impact of fuel taxes on investment in automobile emission-efficiency is an empirical question, as evaluated in Chapter II.

1.2 Cap-and-trade

Radically challenging conventional wisdom, Coase (1960) provocatively asks whether the problem of externalities has been correctly laid out: before then, externalities were thought of as unilaterally caused by an emitter (the upstream firm polluting the river) and affecting an innocent party (the downstream residents eating poisoned fish). In contrast, Coase insists on the *reciprocal* nature of the problem: if my crop is set on fire by the nearby railway, maybe my field is too close to the rails? Coase famously states that *if* there were no costs to negotiation, granting a transferable ownership right to one of the parties could cost-efficiently solve the externality problem, and more effectively so than a command-and-control approach or a tax;¹⁰ only the redistributive question of initial allocation remains.

However, after stating above-mentioned *if*-statement, Coase adds: “This is of course a very unrealistic assumption.” (Coase 1960, p. 15) and continues on the importance of taking into account economic considerations when making legal decisions, as these do have an impact on final resource use in presence of frictions. Over the course of his life,¹¹ Coase kept insisting on the importance of studying, “the importance of transaction costs, the possibility of merger solutions, the costs associated with state action, and the need for a comparative institutional approach” (Medema 2014, p. 111) and was a co-founder of International Society for New Institutional Economics (ISNIE), while the economic community enthusiastically welcomed his negotiation result as a reason *not* to study institutional economics. Ignoring this fundamental *if* leads to what Demsetz (1969) calls a “nirvana approach”: comparing reality with an optimal and, thus, utterly unrealistic, unattainable world.

¹⁰Transferable ownership right in this context is equivalent to full legal liability for either side.

¹¹Indeed, Coase started much earlier to insist, in another context, on the importance of the “cost of using the price mechanism” (Coase 1937, p. 390).

Notwithstanding these considerations, the “Coase” theorem¹² was fundamentally important for newer implementations of sustainability regulation, in particular for developing emissions trading schemes. The main form of emissions trading schemes is cap-and-trade: such a scheme fixes a maximum quantity (cap) of emissions and lets polluting parties subsequently trade the units within this capped quantity. Thus, cap-and-trade is a quantity regulation *via* the price mechanism; it allows firms to shift emissions between regions, years, and sectors. Getting aggregate quantities “right” seems indeed easier than specifying the right quantity for each sector or even firm, as in traditional command-and-control quota policies. Crocker (1966, p. 81) underlines the “information-providing potential of a price system” to regulate air pollution.¹³ Dales (1968) was the first to lay out the applicability of a cap-and-trade scheme as a solution to water pollution, followed by a more general treatment by Montgomery (1972). Chapters IV and V study the European cap-and-trade scheme.

As Professor Dominique Strauss-Kahn put it in my undergraduate economics class, “emissions trading is the solution [to the externality problem] that is at the same time most cherished by economists and most ignored by policy makers.”¹⁴ Today, this statement is no longer entirely correct. Following tentative provisions in the US Clean Air Act in 1977 and more decisive steps in the US Acid Rain Program in Title IV of the 1990 Clean Air Act, emissions trading is no longer a niche idea. In particular, greenhouse gas emissions now fall under the international system of the Kyoto Protocol, the European Emissions Trading System (EU ETS), the California Cap-and-Trade Program, Regional Greenhouse Gas Initiative (RGGI, several US states) and the Australian Clean Energy Act. Currently China is working on its own national scheme. The US also has a trading system for sulfur dioxide and nitrogen oxides.

The key argument in favor of cap-and-trade schemes is the cost-effectiveness results from Montgomery (1972): given a fixed quantity of emissions, overall abatement cost is minimized as abatement is shifted to where abatement is cheapest. The other much-discussed property of emissions trading derives directly from the “Coase” theorem: the fact that final allocation of productive factors is independent of initial allocation, called the *independence property*. “This property is very important because it allows equity and efficiency concerns to be separated in a relatively straightforward manner” (Hahn and Stavins 2011, p. 267). Based on this property, emission certificates were distributed for free during the first two phases of the EU ETS, leading to a massive redistribution of funds across firms. However, the independence property fundamentally relies on the absence of transaction costs. This condition is studied in Chapter IV.

One disadvantage of emissions trading is that it involves the creation of an artificial

¹²The “Coase” theorem was first called so by George Stigler, and might thus rather be called “Stigler theorem.” I follow a suggestion by McCloskey (1998) to add the “quotation marks around the non-Coasean ‘Coase’ theorem.”

¹³Crocker also points out the potential free-rider problem as no one can be excluded from using the (scarce) resource “free air”; he concludes that this problem makes totally decentralized solutions (without quantity regulation by the control authority) undesirable.

¹⁴Cited from the author’s memory of an undergraduate economics class at Sciences Po Paris in 2006.

market: to non-economists, the underlying logic often appears counterintuitive. As there is no natural market and assets are intangible, the regulator has to incur administrative costs for establishing a cap-and-trade scheme. Firms also have to build capacities for emissions trading, additionally to existing structures for tax management. Given the complexity of aggregating national rules, European emissions trading “abounds in loopholes”¹⁵ is vulnerable to regulatory capture (Gawel et al. 2014, p. 176). The most-discussed disadvantage of emissions trading is one that is common to all unilateral regulation of global externalities: if environmental regulation makes it more expensive to produce in one region, production might simply move to another, unregulated region; be it by loss of market share of regulated firms or relocation of their production facilities. In the US, the debate is structured around the “pollution haven hypothesis” and American environmentalists have called for the establishment of measures to correct the distortion at the border in order to protect regulated industries. In the EU, the issue is termed “carbon leakage” by adversaries of environmental policy, who call for a weakening of environmental stringency. The potential for carbon leakage depends, among other factors, on the relative cost impact of environmental policy, trade barriers, and competition intensit. The empirical importance of carbon leakage is evaluated in Chapter V.

1.3 Private voluntary standards

An alternative that has evolved in parallel to above-mentioned public policy approaches are private voluntary standards. In areas where consumers feel that public regulation is not strong enough, firms have voluntarily committed to higher environmental, social, or safety standards: over the last 100 years, voluntary implementation of higher standards is especially common in the food industry. The first movements concentrated on the environmental impact of food production as well as health effects from the use of chemicals. Voluntary self-regulation started moving into mainstream production with the establishment of voluntary third-party verified labels (ecolabels) in the 1970s, for example in Germany with the Bioland label. Following increasing market shares of these ecolabels, national governments established official requirements and labels for ecological food production that effectively established minimum label standards. These labels are still voluntary and complemented by more demanding private standards such as Demeter (Germany) or Nature&Progrès (France). Chapter III studies the interaction when several labels compete in one market.

The success of ecolabels spurred development of labels in other domains. Such labels allow firms to credibly commit to higher production standards than those prescribed by regulation regarding, for example, social sustainability, safety, absence of genetically modified plants, and user friendliness – anything consumers have come to perceive as a negative externality. Voluntary labels solve the problem in a decentralized way: consumers themselves can take on the responsibility of ensuring that acceptable production

¹⁵A VAT carousel fraud on European emissions allowances caused €5 billion in damage for European taxpayers, see for example <https://www.europol.europa.eu/newsroom/news/further-investigations-vat-fraud-linked-to-carbon-emissions-trading-system>; retrieved on the 05/05/2017.

conditions exist. The need for third-party labels can be explained by the fact that most of certified points are not verifiable before purchase and often not even after. Interestingly, third-party sustainability labels often certify characteristics of the production process, rather than intrinsic characteristics of the good. Economics have provided relatively little evidence on how such labels affect product markets in equilibrium.

Fairtrade labels are a particularly interesting case. When production chains are geographically dispersed, national governments struggle to address externalities with standard policy-tools. An example for an international value chain and by far the largest fairtrade sector is the coffee market: while almost all coffee is drunk in Western industrialized countries, its production remains largely in the Global South. Consequently, voluntary labels have a natural advantage over national regulators in such sectors.

Labels pursue different objectives, as they are backed by different labeling organizations. Some labels are established by nongovernmental organizations (NGOs) or by a public regulator to target some dimension of social welfare. Others are organized as private firms, maximizing their own profits from license fees. Finally, labels are often established as industry standards by firms themselves, maximizing industry profits for example in reaction to the establishment of a private label (Fischer and Lyon 2014). Some examples of ecolabel certifiers are Ecocert (for-profit) and the European Union Ecolabel (public); safety-labels for textiles include BlueSign (for-profit) and Oeko-Tex (NGO); sustainability labels in the wood industry are Forest Stewardship Council (NGO), PEFC (industry standard), and the Sustainable Forestry Initiative (industry standard).

Advocates of private voluntary standards underline that each industry knows its own business field best and, thus, is best able to decide on appropriate regulation; be it by establishing its own standard or by deciding which NGO-backed or public standard to offer. Bringing forward an invisible hand argument, they argue that firms have an interest to serve consumer interests, provided the consumer is willing to pay accordingly. Putting all power into the consumer's hand, third-party labels are sometimes portrayed as an opportunity to "shop for a better world." Moreover, when production is distributed around the globe, no government can directly regulate production standards, environmental protection, and labor conditions, so that public policy cannot be a substitute to voluntary industry self-regulation here.

Opponents claim that private voluntary standards often address the issues too superficially and are mere "green-washing" marketing tools. The independence of third-party labels is often not well established, leading to conflicts of interest and credibility issues. A larger problem is that if the government has the aforementioned difficulties in determining the correct level of damage control, it seems rather unlikely that a decentralized mass of consumers can do a better job. When firms are competing in oligopoly, they can establish private voluntary standards that serve as coordination tools in order to segment the market, as shown in Chapter III.

2 Methodology

The chapters of this dissertation use different modeling and estimations techniques that are briefly introduced in this section.

2.1 Nested logit

Both Chapters II and III model consumers taking a discrete choice decision. Such a decision is commonly modeled using the nested logit model. The nested logit is a generalized form of the multinomial logit model, going back to the seminal article by McFadden (1978).

In the general model, the consumer i faces a limited number of K alternatives (products, most of the time), each of them provides him with utility

$$U_{ik}(p_k, X_k, \xi_k, \epsilon_{ik}; \theta), \quad (1)$$

where p_k is the price of alternative k , while X_k are the other observed and ξ_k the unobserved characteristics of this alternative. ϵ_{ik} is an error term and θ a vector of parameters. The consumer chooses option k if it gives him the highest utility:

$$k = \arg \max_j U_{ij} \text{ with } j = 0, 1, \dots, K. \quad (2)$$

In the simple logit model, the consumer's utility is of the functional form

$$U_{ik} = \alpha + \beta p_k + \gamma X_k + \xi_k + \epsilon_{ik} \quad (3)$$

where the residual ϵ_{ik} is assumed to follow an extreme value type I distribution. Integrating over this distribution allows me to determine each consumer's probabilities to choose a particular alternative k , and given homogeneous consumer preferences this probability equals the market share s_k . For simplicity, I denote $\delta_k = \alpha + \beta p_k + \gamma X_k + \xi_k$ the deterministic part of utility:

$$s_k = \frac{\exp(\delta_k)}{\sum_{j=0}^K \exp(\delta_j)} \quad (4)$$

Note that I need to fix the utility of one of the goods in order to normalize the equation system. Usually, the product 0 is defined to be the outside good, i.e. no purchase, with utility $\delta_0 = 0$, so that I have:

$$\ln(s_k/s_0) = \alpha + \beta p_k + \gamma X_k + \xi_k \quad (5)$$

which is the equation that allows me to estimate parameters α , β and γ , treating the unobserved product characteristic ξ_k as an error term.

This model implies the *independence of irrelevant alternatives* (IIA): the choice between two alternatives is assumed independent of the existence of another third alternative.

This assumption is often considered unrealistic, famously illustrated by the red bus/blue bus problem in McFadden (1980).

The nested logit model attempts to alleviate this problem by adding a nested group structure over alternatives, assuming IIA within nests but allowing substitution to depend on nest structure. Technically, this dependence is achieved by allowing error terms to be correlated within a nest. In the car market application of Chapter II, for example, I assume that sports car drivers substitute more easily to another sports car than to a multi-van

Consumer utility from good $k \in \mathcal{T}_g$ is then defined as

$$U_{ik} = \delta_k + \zeta_{ig} + (1 - \sigma)\epsilon_{ik} \quad (6)$$

where both ϵ_{ik} and $\zeta_{ig} + (1 - \sigma)\epsilon_{ik}$ are assumed to follow an extreme value type I distribution. σ then measures the strength of within-nest correlation; when $\sigma = 0$, the nesting structure is irrelevant and the nested logit collapses into a simple multinomial logit model. By integrating, one can derive expressions for market shares s_g for the market share of nest g and $s_{k|g}$ for the market share of good k *within* nest g :

$$s_{k|g} = \frac{\exp(\delta_k/(1 - \sigma))}{D_g} \text{ with } D_g = \sum_{j \in \mathcal{T}_g} \exp(\delta_j/(1 - \sigma)) ; \quad (7)$$

$$s_g = \frac{D_g^{(1-\sigma)}}{\sum_{h=0}^G D_h^{(1-\sigma)}} ; \quad (8)$$

$$s_k = s_{k|g} \times s_g . \quad (9)$$

When estimating nested logit models, it is important to account for the endogeneity of the price p_k and within market shares $s_{k|g}$ by using instrumental variables.

2.2 Binary quantile estimation

Chapter IV uses a more recently developed way of modeling discrete choice: binary quantile estimation (Kordas 2006). This estimation technique uses binary decisions by firms to infer information about an underlying continuous distribution of transaction costs.

Most econometric methods commonly used concentrate on the mean: certainly the mean is not the only informative parameter, but it has convenient statistical properties. Koenker and Bassett Jr. (1978) famously challenged this practice by introducing the computationally more cumbersome quantile estimation method that allows the econometrician to estimate conditional distributions.

While their original model applied to continuous outcomes, Kordas (2006) built on the (smoothed) maximum score estimator (Manski 1975, Horowitz 1992) to establish a methodology to estimate binary regression quantiles. The model assumes that there is a

latent continuous variable Y^* , of which only a binary indicator Y is observed:

$$Y_i^* = X_i' \beta + \epsilon_i \quad (10)$$

$$Y_i = \mathbb{1}\{Y_i^* > 0\}, \quad (11)$$

where X_i is a vector of covariates for observation i , ϵ_i is a random error term and β is a set of parameters of interest. If Y^* was observable, one could estimate a quantile regression following Koenker and Bassett Jr. (1978) for each quantile $\tau \in (0, 1)$:

$$Q_{Y^*|X}(\tau) := F_{Y^*|X}^{-1}(\tau) = X' \beta \quad (12)$$

where $Q_{Y^*|X}(\cdot)$ and $F_{Y^*|X}(\cdot)$ are the conditional quantile and distribution functions of Y^* . Given however that Y^* is not observable in many settings, one can use the fact that quantile estimates are robust to a monotone transformation of the outcome variable (Koenker and Hallock 2001). An indicator function is a monotone transformation, so the conditional quantile distribution of Y is given by

$$Q_{Y|X}(\tau) = \mathbb{1}\{X' \beta \geq 0\} \quad (13)$$

Although I only observe a binary outcome Y , I can draw conclusions on the conditional distribution $F_{Y^*|X}(\cdot)$ of the latent continuous variable Y^* .

If the error term is assumed to be independent and identically distributed, following a normal distribution, the median quantile of equation (13) can be estimated with a standard probit model. However, this assumption might not be appropriate in many situations, e.g. if the error distribution is skewed. Binary quantile estimation allows the researcher to remain agnostic about the distribution of the error term, making this method particularly robust to outliers.

Just like the probit regression relies on an assumption of a mean zero error term, the binary quantile regression assumes that the conditional *median* error is zero. In practice, the binary quantile estimator at the median ($\tau = .5$) maximizes the number of “correct predictions.” The estimation of this model involves optimization over a complex function. In Chapter IV, I use simulated annealing which has the advantage of being more robust to starting values, local optima, and discrete parts of the objective function. Such methods have only recently become available with the dramatic increase of available computing power.¹⁶

2.3 Treatment effects

While the previous methods use a discrete choice approach, Chapter V is based on continuous outcomes in a treatment effect analysis framework.

The standard way to think about the treatment effect is shaped by Rubin’s (1974) model for causal inference using counterfactual outcomes. The fundamental problem is

¹⁶Nevertheless, the main estimation of Chapter IV needs almost six hours to run.

that the outcome for a particular individual usually is observed with the treatment or without, but rarely both. Nevertheless, one would like to find the treatment effect τ , given by:

$$\tau = E(Y|T = 1, X) - E(Y|T = 0, X) \quad (14)$$

where Y is some outcome, X a set of covariates and $T = \{0, 1\}$ the treatment status. When treatment is randomized across individuals and compliance is full, the estimate for the average treatment effect is simply given by the difference between sample mean of the treated and the sample mean of the untreated. In practice, situations with incomplete compliance – when some assigned to the treatment do not follow – or even observational, non-randomized data are more common; a large literature developed to address these problems, notably by the work of Heckman, Angrist, and Imbens.

Chapter V uses such a framework in a somewhat unusual setting: observations i are in this case sectors and treatment can be defined as either binary or continuous, relating this work to studies with multi-valued treatment (also called treatment intensity or dose-response). In the most simple version, one assumes a constant unit treatment effect $Y_j - Y_{j-1} = \beta$ for all j and all sectors, so that the model can be estimated using linear regression models. I further assume strong unconfoundedness, i.e. conditional on a set of covariates, environmental policy stringency (the treatment) does not depend on import intensity of a sector (the outcome). I thus estimate

$$Y_i = \alpha + \beta\theta_i + \gamma X_i + \epsilon_i \quad (15)$$

where X_i is a set of covariates, ϵ_i a random error term and θ_i a continuous treatment measure. Alternatively, I also test whether my results differ when I use a binary indicator function $\tilde{\theta}_i = \mathbb{1}(\theta_i > 0)$. However, Angrist and Imbens (1995) remind us that collapsing a multi-valued treatment into a binary treatment indicator – above/below some cut-off – generally bias the estimates.

3 Contribution of this dissertation

Each chapter concentrates on one of the above-mentioned policies: Chapter II studies the impact of fuel taxes on new automobile fuel efficiency; Chapter III studies the strategic quality setting of sustainability labels; Chapter IV estimates transaction costs in European emissions trading; and Chapter V searches for evidence of carbon leakage in European emissions trading. Table I.1 provides an overview of chapter titles, co-authors and pre-publications.

Chapters II, IV and V study policies that address environmental sustainability concerns, while Chapter III focuses on social sustainability. Chapter V additionally relates to the economic sustainability of emissions trading. Chapters II, IV and V rely on econometric analysis, while Chapter III develops a theoretical model. Chapters II and III present a hypothetical ex ante evaluation of a policy, while the last two chapters provide an ex post assessment of particular aspects of emissions trading.

While all of the chapters belong to environmental economics, this dissertation is positioned at the crossroads with several other sub-fields of economics. Chapters II, III and IV use approaches from industrial organization. Chapter V is based on literature from trade and international economics, and Chapter IV relies on the concept of transaction costs that is central to new institutional economics.

In practice, the impact of a policy depends very much on the agent's sensitivity to economic incentives, which Chapters II and V attempt to measure. The desirability of voluntary approaches and industry self-regulation depends on strategic firm interactions as studied in Chapter III. Finally, the crucial *if* of the "Coase" theorem determines if one can use a solution relying on negotiation or not, which Chapter IV studies for the European emissions trading scheme. Note that these policies address the problem from very different angles, so that the unit of analysis is sometimes the consumer (Chapter II and III), sometimes the firm (Chapters III and IV) or even sector-level aggregates of firms (Chapter V).

The following subsections go more into detail on each chapter's contribution.

3.1 Fuel taxes and automobile fuel efficiency: the importance of consumer elasticity

Following efforts to address emissions from private road transports and to ensure political independence from oil-producing countries, automobile fuels are among the most heavily taxed goods categories in Europe. The impact of fuel taxes depends on consumer elasticity in two dimensions: in the short run, consumers can adjust their mileage driven with their current car (intensive margin); in the long run, consumers can invest in more efficient cars (extensive margin); additionally, there is potentially an interaction between margins, as consumers who invested in a more efficient might react less in their mileage or even drive more, the so-called "rebound effect."

Chapter II evaluates the impact of a hypothetical fuel tax on the extensive margin, i.e. on new car purchases. The research relies on exhaustive consumer-level data of monthly registration of new cars in France. I use information on the car holder to account for heterogeneous preferences across purchasers and identify demand parameters through the large oil price fluctuations of this period. The results suggest that the sensitivity of short-term demand with respect to fuel prices is generally low and, in particular, for corporate purchases.

Using the estimated parameters of consumer demand to compute elasticities, Chapter II estimates the *ex ante* impact of two different policies. First, a policy equalizing diesel and gasoline taxes would reduce the share of diesel-engines in new car purchases, without substantially changing the average fuel consumption or CO₂ emission levels of new cars. Second, I suggest a revenue-equivalent carbon tax that would be at 51 €/ton of CO₂. Again, this policy has only small effects on average fuel consumption or average CO₂ emission levels of new cars.

As I refrain from taking any hypothesis on mileage and mileage elasticity of consumers, I cannot identify whether consumers under-invest in fuel efficiency relative to

Table I.1: Overview by chapter: topic, pre-publication and author's contribution

Ch.	Title	Co-Authors	Pre-Publication	Contribution
I	Overview of market-based sustainability policies	Single author	(not published)	
II	How do fuel taxes impact new car purchases? An evaluation using French consumer-level data	Pauline Givord, Céline Grislain-Letrémy	DIW Discussion Paper 1428, 2014; Revise&Resubmit at Energy Economics; Follow-up project of the author's master thesis	Author worked on the initial draft, including in particular the literature review. Final development of the model, computational implementation, interpretation of results and writing was collaborative.
III	Competition between for-profit and industry labels: the case of social labels in the coffee market	Pio Baake	DIW Discussion Paper 1686, 2017; under revision	Research was initiated by the author, responsible for writing the manuscript. Establishment of the model was collaborative.
IV	Offset credits in the EU ETS: a quantile estimation of firm-level transaction costs	Single author	DIW Discussion Paper 1513, 2015; Environmental and Resource Economics, forthcoming. The final publication is available at Springer via https://doi.org/10.1007/s10640-017-0111-1 .	The author is responsible for all parts of the research.
V	Does the EU ETS cause carbon leakage in European manufacturing?	Aleksandar Zaklan	DIW Discussion Paper 1689, 2017; under revision	Research was initiated by the author, responsible for developing the model, performing data analysis and interpreting the results. Writing, as well as data collection and management were collaborative.

potential savings from more efficient cars: the rationality of consumer response is not assessed here. Nevertheless, in order to impact average fuel efficiency by a magnitude not only statistically significant but also economically meaningful, fuel tax would have to be much higher, raising the question of political feasibility.

3.2 Competition between sustainability labels: how an industry standard strategically impacts product lines

Private voluntary standards (labels) offer the consumer the possibility to enforce certain points of production practices beyond legal minimum levels. From a firm perspective, labels are a tool to differentiate products from their competitors and/or their own product lines.

The coffee market is characterized by labels more than any other product category. Coffee is a product with an international value chain, as coffee farming is not possible in most coffee consuming countries, for example in Europe. Moreover, world coffee prices have dramatically fallen since the end of the Cold War (which marked the end of the International Coffee Agreement) so that many coffee farmers live at mere subsistence level. Consumers are worried about such inequality and see the coffee farmers' poverty as a negative side-effect of mainstream conventional coffee production. Different countries show different market constellations of product line rivalry: firms specialize in some countries, while they compete on all market segments in others, similar to the results of my model.

Chapter III examines a market with labels of different quality and objectives. I model how an industry standard interacts with the fairtrade label, facing firms in duopoly that decide which labels to offer. The incumbent fairtrade label maximizes its own profit and is challenged by an industry standard that maximizes joint firm profits. Using a nested logit, the result of this multi-stage game depends crucially on the (exogenous) degree of horizontal differentiation. The industry label always wants to segment the market, if possible, and attempts to distort the number of labeled products downwards. For high levels of horizontal differentiation, the industry cannot coordinate on not competing on all labels; for low levels of horizontal differentiation, the market is segmented; finally, for intermediate levels of horizontal differentiation, the industry sets a strategically low standard that in equilibrium reduces the overlap of product lines from different firms and thereby reduces competition. At all levels of horizontal differentiation, the industry benefits from the introduction of the industry standard.

A social planner would like to prevent such a distortion, as this maximizes both consumer utility and aggregate social welfare. I explore whether there is scope for a policy imposing a minimum label quality and find that a minimum label is only binding in the cases with intermediate horizontal differentiation in which the industry standard strategically reduces the number of available products. For very high or very low levels of horizontal differentiation, the industry standard cannot strategically induce more segmented product lines; in these cases, a minimum label quality is not binding.

3.3 Transaction costs in European emissions trading: the fundamental *if* in the Coase theorem

As underlined by, for example, Rose and Stevens (2001), the cost-effectiveness from cap-and-trade schemes stems largely from the transferability of pollution in time (banking/borrowing), between firms (trade), and across regions (linking of regional emissions trading schemes). Many hopes have in particular been put into the last point, trading between geographical regions; however, critiques argue that industrialized countries should not hamper economic development of the Global South by paying them for *not* producing. When putting into place the EU Emissions Trading System (EU ETS), the EU thus introduced a possibility to substitute *some* foreign emission reduction effort to domestic effort via offset certificates, but also strongly limited the overall amount of such offset certificates.

An offset certificate is created from emission reductions in unregulated regions, certified by the responsible UN Environment Programme (UNEP). Offset certificates have been cheaper than European certificate at all times, although they are substitutes within the EU ETS. Thus, firms had a strong incentive to use offset certificates up to their firm-specific quota. However, a considerable number of firms did not exhaust their offset quota and, by doing so, seemingly forwent profits.

While most literature on emissions trading evaluates the efficiency of regulation in a frictionless world, in practice firms incur costs when complying with regulation. Transaction costs of emissions trading include information gathering, forecasting of allowance prices, finding trading partners, bargaining, contracting, and managing price risk or, instead, the costs of out-sourcing the whole trading process – costs that are contingent on actively buying and selling certificates. In order to assess the relevance of emissions trade related fixed transaction costs, Chapter IV examines the use of international offset credits in the EU ETS. It establishes a model of firm decision under fixed (quantity-invariant) entry costs and estimates the magnitude of trading costs rationalizing firm behavior using semi-parametric binary quantile regressions.

The resulting cost estimates are sizable: they prevent a fifth of all firms, especially small emitters, from using offset certificates. Comparing binary quantile results with probit estimates shows that high average transaction costs result from a strongly right-skewed underlying distribution. For most firms, the bulk of transaction costs stems from certificate trading in general, rather than additional participation in offset trading.

3.4 Carbon leakage and European emissions trading: “much ado about nothing”?

When the EU ETS was introduced, carbon leakage was one of the biggest concerns of policy makers, industry representatives, and academics. The fundamental dilemma is that the regulator, on the one hand, wants the social cost of emissions to be passed through to firms and consumers in order to align private and social costs. On the other hand, the policy maker is concerned about the effect of the policy on product prices, because

this makes foreign products (produced in non-regulated regions) more attractive. This problem stems inherently from the unilateral nature of this effort to reduce a global externality. The debate is particularly salient in Europe, where the EU ETS covers emissions of many traded sectors.

In a first step, Chapter V surveys how carbon leakage and the pollution haven hypothesis have been identified in previous literature with particular attention to the definition of outcome and policy treatment variables.

In a second step, Chapter V uses a panel of trade and input-output data from the Global Trade Analysis Project (GTAP), in order to compute trade flows in value and in embodied carbon, and combine it with administrative plant-level data from the EU ETS. This allows me to account for direct and indirect (through electricity use) environmental cost from cap-and-trade regulation. I consider both bilateral trade flows and net imports. I do not find any evidence in favor of carbon leakage caused by the EU ETS during its first two trading phases.

4 Concluding remarks

The challenges to contain climate change and to reduce global poverty are just some aspects of a global search for sustainability. The relatively new consciousness about our impact on future generation's living conditions gives us a responsibility to address these problems quickly.

This dissertation assesses the effects and side-effects of certain approaches to sustainability regulation. Chapter II finds that households react little to fuel-tax incentives in their vehicle choice and firms react even less. Chapter III finds that strategic interaction between labels reduce the welfare benefit of labels. Chapter IV attempts to put a price tag on inertia and the burden of regulatory complexity of European emissions trading. Chapter V finally finds that alarming scenarios of side-effects of European emissions trading – production relocation and competitiveness loss of European firms – have not materialized.

Economic systems are embedded in social structures. The isolated view of individual policy measures easily allows the economist to compute optimal behavior in stylized models, but empirical evaluation typically finds that agents react much less. A typical example are the contrasting findings of ex ante models predicting dramatic levels of carbon leakage, and ex post econometric studies that fails to find any evidence of carbon leakage (such as Chapter V). Overall, the surrounding social and organizational elements increase the system's inertia. The findings about transaction costs in this dissertation underline the importance of keeping complexity at bay.

Thus, it seems as if society needs to take much more drastic steps if change is to be achieved quickly. More drastic regulation necessarily implies higher costs and must be backed by a strong commitment to sustainability. Emissions trading is an example of lacking commitment: policy makers both want to make emissions more costly for firms and at the same time protect firms from these costs. Solutions to the carbon leakage

problem include border-carbon adjustments (carbon-based tariffs) and output-based allocation; while the former is very likely to be against WTO laws, the latter unfortunately undoes (part of) the incentive effect that is the *raison d'être* of emissions trading. The point that has so far not attracted a great deal of attention is that carbon leakage is not – at current price levels – an actual problem (Chapter V): firms that remain in Europe despite labor costs being 10 to 30 times higher than in emerging nations (e.g. Schröder 2016) do not relocate for emissions costs of €5/tCO₂e.

More research is necessary in order to identify how sustainability policies can have the largest effect at the smallest cost. However, this search for efficiency should not conceal the hard truth that change comes at some cost and that clear priorities and determination are key to move away from today's unsustainable practices.

5 Bibliography

- Allcott H and Wozny N (2013) Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics*, 96(5): 779–795.
- Angrist J. D and Imbens G. W (1995) Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association*, 90(430): 431–442.
- Baumol W. J and Oates W. E (1971) The use of standards and prices for protection of the environment. *The Swedish Journal of Economics*, 73(1): 42–54.
- Coase R (1960) The problem of social cost. *Journal of Law & Economics*, 3(4): 87–137.
- Coase R. H (1937) The nature of the firm. *Economica*, 4(16): 386–405.
- Crocker T. D (1966) The structuring of atmospheric pollution control systems. In H. Wolozin (ed.) *The economics of air pollution*, New York: W. W. Norton & Company: 81–84.
- Dales J. H (1968) Land, water, and ownership. *The Canadian Journal of Economics*, 1(4): 791–804.
- Demsetz H (1969) Information and efficiency: another viewpoint. *The Journal of Law and Economics*, 12(1): 1–22.
- Fischer C and Lyon T. P (2014) Competing environmental labels. *Journal of Economics & Management Strategy*, 23(3): 692–716.
- Fullerton D and Metcalf G. E (1997) Environmental taxes and the double-dividend hypothesis: did you really expect something for nothing? *Chicago-Kent Law Review*, 73: 221–1393.
- Gawel E, Strunz S, and Lehmann P (2014) A public choice view on the climate and energy policy mix in the EU: How do the emissions trading scheme and support for renewable energies interact? *Energy Policy*, 64: 175–182.
- Hahn R. W and Stavins R. N (2011) The effect of allowance allocations on cap-and-trade system performance. *The Journal of Law and Economics*, 54(S4): S267–S294.
- Horowitz J (1992) A smoothed maximum score estimator for the binary response model. *Econometrica*, 60(3): 505–531.
- IPCC (2014) Fifth assessment report (AR5). Synthesis Report, Intergovernmental Panel on Climate Change.
- Koenker R and Bassett Jr. G (1978) Regression quantiles. *Econometrica*, 46(1): 33–50.
- Koenker R and Hallock K (2001) Quantile regression: an introduction. *Journal of Economic Perspectives*, 15(4): 43–56.

- Kordas G (2006) Smoothed binary regression quantiles. *Journal of Applied Econometrics*, 21(3): 387–407.
- La Fontaine J. d (1668, reprinted 2007) *Conseil tenu par les rats*, Chap. II.2: Editions Flammarion.
- (1882, reprinted 2005) *The Council Held By the Rats or Who Will Bell the Cat?*, Chap. II.2: Projekt Gutenberg.
- Manski C (1975) Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics*, 3(3): 205–228.
- McCloskey D (1998) The so-called Coase theorem. *Eastern Economic Journal*, 24(3): 367–371.
- McFadden D (1978) Modeling the choice of residential location. *Transportation Research Record*, 673: 72–77.
- (1980) Econometric models for probabilistic choice among products. *The Journal of Business*, 53(3): S13–S29.
- Medema S. G (2014) Neither misunderstood nor ignored: the early reception of Coase’s wider challenge to the analysis of externalities. *History of Economic Ideas*, 22(1): 111–132.
- Montgomery D (1972) Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3): 395–418.
- Panhuysen S and Pierrot J (2014) Coffee barometer 2014. Technical report, Tropical Commodity Coalition, The Hague.
- Pigou A. C (1920) *The economics of welfare*, London: McMillan&Co.
- Pindyck R. S (2016) The social cost of carbon revisited. Technical Report, National Bureau of Economic Research.
- Rose A and Stevens B (2001) An economic analysis of flexible permit trading in the Kyoto Protocol. *International Environmental Agreements*, 1(2): 219–242.
- Rubin D. B (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5): 688–701.
- Schröder C (2016) Industrielle Arbeitskosten im internationalen Vergleich. *IW-Trends. Vierteljahresschrift zur empirischen Wirtschaftsforschung*, 43(3): 17–26.
- Van den Bergh J (2010) Externality or sustainability economics? *Ecological Economics*, 69(11): 2047–2052.
- Weitzman M. L (1974) Prices vs. quantities. *The Review of Economic Studies*, 41(4): 477–491.

II

How do fuel taxes impact new car purchases? An evaluation using French consumer-level data

Quitte à pleurer, je préfère
pleurer dans une Jaguar.

Françoise Sagan
on the rationality of automobile
purchases.

1 Introduction

In France, road transport produces more than a third of total CO₂ emissions and much higher shares of other greenhouse gases.¹ On the one hand, this problem might be alleviated by a shift to diesel-fueled cars, as diesel is more dense in energy and diesel-engines particularly efficient in using it: typically, a diesel car produces less CO₂ per km than a similarly-sized gasoline-fueled car. On the other hand, diesel cars also produce medically hazardous fine particles (in particular black carbon) and nitrogen oxides (NO_x). Thus, policy makers are facing both a global climate problem as well as a local health issue; shifting toward more diesel-fueled cars might alleviate the global externality, but increase local concerns.

Facing the conundrum between global and local pollution, European policy makers have, for a long time, opted to support diesel, particularly in France (Hivert 2013). This tax advantage for diesel is fueling a renewed debate, sparked by episodes of smog. In December 2016, air quality in France dropped so low that the government heavily restricted driving and Paris authorities have banned the oldest and most polluting vehicles from the city center, pledging “an end to diesel” in Paris by 2020. As stressed, for instance by Mayeres and Proost (2001), environmental benefits of diesel cars have been overestimated: the environmental costs of diesel cars are much higher than those of gasoline cars. While diesel cars emit more fine particles and NO_x that harm human health, new technology decreases the spread between CO₂-emission-efficiency of diesel and gasoline cars.² The production of diesel-models is also more CO₂ intensive because they are heavier. Against this background, in 2015 the French government announced the progressive reduction of the relative tax advantages for diesel fuel.³ This tax alignment adds to a previous “carbon tax” passed in France in 2003 at a modest €15 per tonne of CO₂. Such a tax is proportional to the amount of CO₂ emitted and is, thus, considered a more efficient incentive aligning directly the private cost to the consumer and the externality cost to society.

Emissions from road transport depend heavily on the vehicle fleet in circulation, as cars are durable goods, and regulation affecting the entry of new vehicles impact emissions for a long time. While mandatory standards (command-and-control regulation) were the most prominent regulation until the 1990s, alternative regulations have been tested since, in particular economic incentives such feebates or fuel taxes.^{4,5} Fuel taxes

¹See <http://www.citepa.org/en/air-and-climate/analysis-by-sector/transport>; retrieved on 14/03/2015.

²Miravete et al. (2015) go as far as to argue that diesel-friendly policy in Europe is essentially a non-tariff trade barrier against American manufacturers.

³The difference was reduced from 14.9 cent in 2015 to 11.7 cent in 2016 and 9.4 cent in 2017. The path to full equalization such as described in this chapter has yet to be defined; see <http://www.douane.gouv.fr/articles/a12285-carburants-gazole-super-e10-taux-de-taxe-par-region>; retrieved on 05/09/2017.

⁴Besides recent scandals show that standards seem difficult to enforce effectively.

⁵Feebates, a system combining fees (for more polluting cars) and rebates (for less polluting cars) were implemented in several European countries in the 2010s. This mechanism is expected to shift consumer expenses toward less polluting goods, and to be self-financed as the fees should compensate the rebates.

have the advantage of affecting both the present and future emissions: car owners are immediately encouraged to drive less with their current car when fuel prices rise, while at the same time investment in fuel-efficient cars becomes more attractive. On this latter aspect, some previous results, based mostly on the US market, emphasize an “energy paradox,” meaning that consumers systematically undervalue future economies of energy-efficiency (e.g. Allcott and Wozny 2014); others, like Sallee et al. (2016) or Busse et al. (2013) find no evidence of such consumer myopia. Meta-studies (Helfand et al. 2011, Greene 2010) find that the empirical evidence about the energy paradox is inconclusive.

The effect of fuel taxes on carbon emissions depends on the extent to which car owners react to fuel price, i.e. whether such taxes are able to change the composition of the vehicle fleet toward more fuel efficiency (greenhouse gases) and how the share of diesel cars evolves (local pollution). This chapter estimates the short-term sensitivity of automobile purchases to changes in fuel prices in France. We evaluate the impact of two (hypothetical) fuel tax policies on aggregate characteristics of the vehicle fleet in circulation, leaving aside the question whether consumers adjust their mileage both to changing fuel prices and to changing fuel efficiency of their car (potential rebound effect).⁶ We contribute to the literature by addressing the aggregate impact on the composition of the vehicle fleet in circulation, disregarding whether a low sensitivity to fuel prices is due to elastic mileage or to consumer myopia.

We use French car registration data from 2003 to 2007, which includes exhaustive information about both household and firm automobile purchases. Our main focus lies on the aggregate impact of fuel taxes on fuel consumption, CO₂ emission intensity and the share of diesel purchases. Our dataset links technical car characteristics to information on the car holder. This enables us to define consumer types to account for heterogeneity in preferences across purchasers. In particular, we can separate between private consumers and firms. While the latter represent more than one-third of purchases of new cars in France (over our period), virtually no evidence exists so far on their responsiveness to changes in fuel prices.

As it is common in this literature, we rely on a static discrete choice model assuming that the decision to buy a specific car depends on several car characteristics, including the cost per kilometer. The nested logit specification enables us to model substitution patterns depending on car market segments and on fuel-type versions. Over these five

However, D'Haultfœuille et al. (2014) show that the French experience has led to unexpected results. In absence of previous empirical evidence on consumers elasticities to car prices, the feebate system has resulted in a sharp increase in car sales, but also in CO₂ intensity. This disappointing result is partly explained by a “rebound effect”: with a more fuel-efficient car, the cost per kilometer is lower, which may induce more driving.

⁶There are four components to the reaction of total emissions to fuel taxes: the direct mileage elasticity to fuel prices, the elasticity of the new car's fuel efficiency to fuel prices (analyzed here), the elasticity of mileage to this new fuel efficiency and the elasticity of car lifetime. Frondel and Vance (2014), for example, examine the first point and find that the elasticity of mileage to fuel prices is not significantly different for diesel and gasoline drivers. We examine the second point. Small and Van Dender (2007) examine the third point. Adda and Cooper (2000) work on the fourth point.

years, monthly fuel prices vary considerably. We identify the impact of fuel cost in car choice using time variation in fuel prices and cross-sectional differences in fuel efficiency. We deduce the elasticity of demand for cars with respect to an increase in fuel taxes.

Our results suggest that short-term sensitivity of demand with respect to fuel prices is generally low, but presents significant heterogeneity across purchasers. The difference between private and corporate purchases is particularly salient: firms are much less reactive than households. We use our estimates to simulate the impact of two hypothetical policies, the equalization of diesel and gasoline taxes and a “carbon tax.” Both policies increase taxes relative to the *status quo* but they are calibrated to be revenue-equivalent to each other.⁷ Assuming that consumers react to changes in final consumer prices without distinguishing between taxes or oil price changes, our results suggest that equalizing diesel and gasoline taxes would reduce the market share of diesel cars (from 69% to 65%) in the short-run without notably changing average fleet fuel consumption or CO₂ intensity. The carbon tax leaves the diesel share almost constant and has a similarly small impact on the other two outcomes. Overall, fuel taxes do not seem an effective tool to influence car choices.

This chapter is in line with the literature on the impact of fuel prices on the automobile sector. Most papers focus on American data (Allcott and Wozny 2014, Busse et al. 2013, Klier and Linn 2010) and concentrate on the question of consumer rationality, as reviewed in Greene (2010) and Helfand et al. (2011), while we choose to take a policy maker’s perspective and concentrate on the aggregate vehicle fleet characteristics. Klier and Linn (2013), who evaluate the effect of fuel prices on new vehicle fuel economy in the eight largest European markets (including France), observe strong differences between European and American markets. Most of this existing literature relies on data with little or no information on consumers, while we have individual data matching cars to consumers and can identify corporate purchases. Previous results for France suggest that the elasticity of fuel demand to fuel prices in France is heterogeneous across demographic groups (Clerc and Marcus 2009), depending notably on working status. We only estimate short-run reactions, as we take supply as given: list prices can be adjusted in the medium-term and the set of available cars might change in the long-run.

The chapter is organized as follows. Section 2 explains our assumptions on the decision making process. Section 3 presents the data and some descriptive statistics. The model is presented in Section 4. Section 5 discusses results and robustness tests, and Section 6 concludes.

2 Choice model

To model market shares of new vehicles, we rely on a standard discrete choice model with differentiated products. More specifically, we assume that the purchaser buys one product maximizing his utility that is a linear function of new vehicle characteristics and a vehicle-specific unobserved effect. The individual valuation of these vehicles may vary

⁷As a consequence, our carbon tax scenario is more ambitious than the tax voted in France.

among individuals, like e.g. Allcott and Wozny (2014), tracing back to seminal work by McFadden (1978).

We assume that the consumer decision can be modeled as a hierarchical choice, choosing first a car segment (i.e. SUV, compact, etc; see list in Table II.1), then a model (combination of nameplate and car body style) within this segment, and, finally, one of the two fuel-type versions of this model.⁸ While this structure is largely ad hoc, it seems empirically validated by our parameter estimates (see Appendix C on page 55). The nested logit model yields heterogeneous substitution patterns between products that are more or less similar; for instance a sporty BMW Z3 is more substitutable to a BMW Z4 than to a bulky Renault Kangoo. We also consider an outside option, which is not to buy any new vehicle.⁹ This substitution pattern is represented in the tree diagram of Figure II.1.

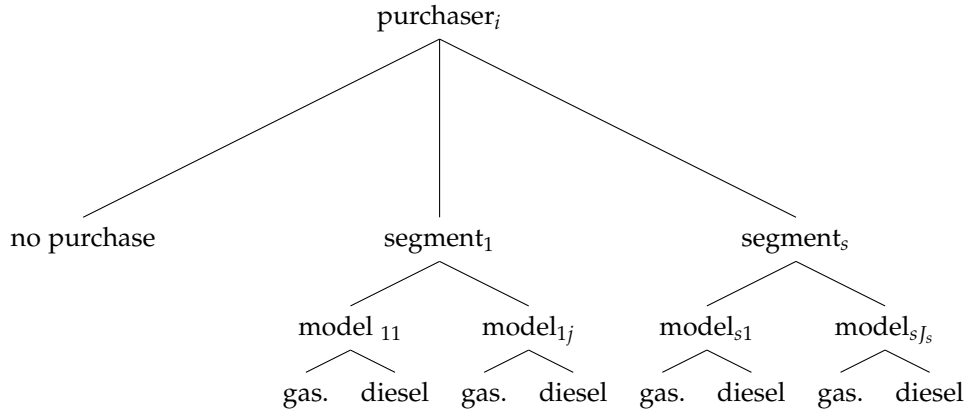


Figure II.1: Nested decision-making structure of the car purchaser

The individual utility of choosing the product with model (combination of nameplate and car body style) j , fuel-type f and segment s , for purchaser i at month t is written:

$$u_{ijft} = \alpha_i + \beta_i p_{jft}^{km} + \gamma_{1i} p_{jft} + \gamma_{2i} X_{jft} + \xi_{ijft} + \epsilon_{ijft}, \quad (1)$$

where p_{jft} denotes the car price and X_{jft} represents the characteristics of new cars. p_{jft}^{km} is the cost at time t for the amount of fuel needed to drive one km with the model j of fuel-type f .¹⁰ ξ_{ijft} measures the unobserved (to the econometrician) preference for product jf . As such, it captures attributes like perceived quality, design and reputation.

We rely on a nested logit specification with two nesting levels to reflect our decision

⁸In order to clarify the vocabulary, *nameplate* refers to the brand name of the car, for instance Corolla, Prius. Within the same nameplate, there are usually several *models* that are defined in this chapter by the intersection of a nameplate and a body style, i.e. Corolla sedan or Corolla station wagon. Each *model* typically exists as two different *products*, i.e. in a diesel- and a gasoline-version.

⁹As we consider monthly sales, the outside option's market share is likely to be much larger than any other option's share. For the sake of comparison, over the period the number of new cars registered a month ranges from 75,000 to 160,000 vehicles, for around 37.5 millions of drivers in France.

¹⁰Another way to look at this would be to multiply the fuel consumption by the number of kilometers expected by the purchaser and using some sort of discounting; this is equivalent to our presentation if β_i is defined to include this expected number of kilometers and discount factor of purchaser i .

process of Figure II.1. This means we assume the error term can be decomposed as:

$$\epsilon_{ijft} = v_{ist} + (1 - \sigma_{2i})(v_{ijt} + (1 - \sigma_{1i})e_{ijft}), \quad (2)$$

where v_{ijt} measures the individual preference for unobserved characteristics of model j common to both fuel versions, for example design, while v_{ist} is the consumer's overall preference for segment s , for example status symbol value of SUVs. The remaining error e_{ijft} is assumed to be independent and identically distributed according to an extreme value distribution. There is a unique distribution for v_{ist} and v_{ijt} such that ϵ_{ijft} follows an extreme value distribution (Cardell 1997). This specification is standard in this literature (see in particular Berry 1994).

The parameters σ_{1i} and σ_{2i} capture the correlation between individual preferences for cars within nests, as defined above. As shown by McFadden (1978), the nested logit model is consistent with random-utility maximization for values of σ_{1i} and σ_{2i} between 0 and 1. $\sigma_{1i} = 0$ means that substitution effects are identical across and within model,¹¹ while a high σ_{1i} , approaching 1, implies a high correlation between preferences for both fuel-versions of the same model. $\sigma_{2i} = 0$ implies that the purchaser is a priori indifferent to substitute between models within and across segments (see for example Verboven 1996 for a more complete discussion of these terms).

3 Data and descriptive evidence

3.1 New vehicle registrations

We use the exhaustive dataset of all new cars registered in France from January 2003 to November 2007, provided by the Association of French Automobile Manufacturers (CCFA, *Comité des Constructeurs Français d'Automobiles*), giving us over 7 million observed registrations. As a feebate scheme was introduced in January 2008, which dramatically changed the demand for fuel economy, we only use data up to the date of its announcement in November 2007.¹²

Our data includes all information necessary for the registration of a new car, i.e. both technical specifications of the car as well as demographic information on the purchaser. The CCFA has further linked this data to list prices of new cars as provided by the car manufacturers.¹³

A product is defined by brand, nameplate (Corolla, Kangoo, etc.), fuel-type (diesel or gasoline),¹⁴ CO₂ intensity class and body style (for instance city-car and sedan).¹⁵ More-

¹¹“Within-model” substitution refers to the substitution between the gasoline-powered and the diesel-powered versions of the same model.

¹²See D’Haultfœuille et al. (2014) for an analysis of this policy and a description of this dataset.

¹³List prices may differ from the actual selling prices, which are unobserved.

¹⁴We exclude electric and hybrid vehicles as they constitute a tiny share of the French market over the examined period.

¹⁵The definition seeks to be detailed enough to avoid the aggregation of heterogeneous products. At the same time, a too narrow definition yields many zero monthly market shares, which have to be dropped by

over, the dataset contains other characteristics like number of doors, horsepower, weight, cylinder capacity. Given the outlined structure of the decision process, we exclude models available with only one fuel-type; this is only the case for exceptional cars which represent overall 7% of sales.¹⁶

Table II.1: Descriptive statistics: main characteristics of new car registrations 2003-2007

	Products	Sales-weighted		Products	Sales-weighted
<i>By type of car-body</i>			<i>By class of CO₂ (g/km)</i>		
City-car	3%	7%	≤100	0%	0%
Compact	14%	34%	101 to 120	4%	18%
Sedan	33%	24%	121 to 140	9%	27%
Minivan	13%	24%	141 to 160	14%	33%
Utilitarian	6%	4%	161 to 200	29%	21%
Sport	20%	3%	201 to 250	26%	6%
All-road/SUV	10%	5%	>250	18%	2%
<i>By horsepower</i>			<i>By type of fuel</i>		
≤60	14%	34%	Gasoline	57%	32%
61 to 100	35%	60%	Diesel	42%	74%
101 to 140	27%	10%			
141 to 180	13%	2%			
>180	10%	1%			
Number of products and observations				2,148	7,828,903

Source: CCFA, authors' calculations.

3.2 Types of consumers: demographic groups

Our administrative registration data match every sale of a new car with information on the new car owner. We can distinguish between private buyers and firms. Fuel price elasticities are likely to be related to consumer characteristics such as income, working status and area of residence. Most of the relevant literature on fuel elasticity relies on aggregate data, but as noted by Bento et al. (2012), this omission might entail erroneous findings about fuel economy valuation.

In order to account for heterogeneous preferences, we split our sample into consumer types based on demographic characteristics: we differentiate three firm sectors and three occupational types of private consumers. We further differentiate types based on geography and income, resulting in 28 distinct consumer types. These categories aim at capturing factors essential to vehicle choice and fuel-price sensitivity: mileage and preference for diesel cars, as well as a comfort-price trade-off. The location additionally captures the extent to which a buyer can substitute with other means of transports (bike,

definition: the logit model does not accommodate zero market shares, conceptually, and we cannot take the log of zero, practically. The definition used here is similar to Allcott and Wozny (2014) and somewhat more detailed than those used in most of the literature (e.g. Goldberg 1995, and Verboven 1996).

¹⁶One of the robustness checks verifies that this assumption is not crucial for the results, cf. Section 5.4.

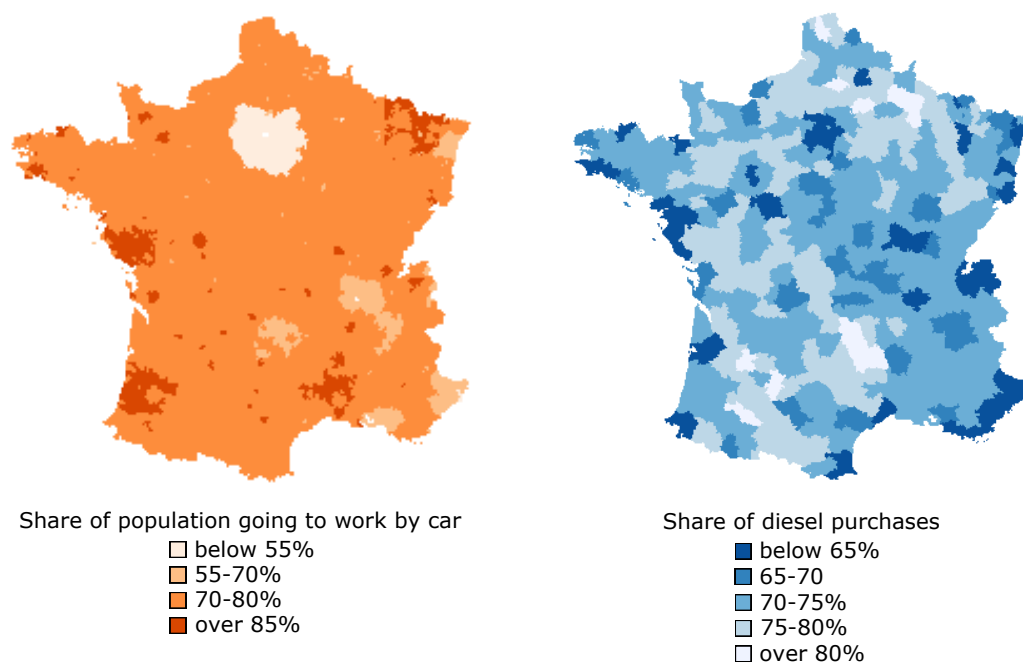


Figure II.2: Overview of spatial variation in share of diesel cars and mileage

Source: CCFA (left graphic) and INSEE National Transport and Travel Survey 2007 (right graphic), maps generated by the authors using R and GEOFLA base maps.

public transport, etc.). The groups are designed in a way to explain as much variation in diesel share, annual mileage¹⁷ and car price as possible.

Table II.2: Average mileage by household characteristics (private consumers only), km/year

Income	Not employed		Employed	
	Low	High	Low	High
Urban	10,850	10,950	14,950	15,600
Suburb./rural	10,750	14,300	16,250	18,850
Paris urban	9,750		14,050	
Paris suburban	11,950		18,350	

Source: INSEE National Transport and Travel Survey 2007, author's calculations.

For both private consumers and firms, we differentiate between types of residence areas. Residence area (rural or urban) accounts for differences in average travel distance and the availability of means of transport other than the car. Residence area is derived from the postal code: we sort areas of residence between urban Paris, the larger Paris metropolitan region,¹⁸ other urban areas and suburban/rural zones. Different types of

¹⁷Information on annual mileage is available for households only and not by age group, computed from INSEE National Transport and Travel Survey 2007, see Table II.2.

¹⁸In the following, we use the term "Paris" or "urban Paris" for Paris and its close and densely populated suburbs (departments Paris (75), Hauts-de-Seine (92), Seine-Saint-Denis (93), Val-de-Marne (94) and some adjoining municipalities) while "Paris metropolitan region" or "suburban Paris" describe the rest of the Île-de-France region.

residence areas have considerably different average travel times and distances (Baccaini et al. 2007). The average yearly mileage is consistently smaller in the Paris region with its dense public transportation network than in other comparable areas.

Activity status is an additional important factor for private owners, as employed consumers have larger mileage across all geographic areas, shown in Table II.2. Indeed, the difference between average yearly mileage ranges from around 10,000 km/year for non-active households living in urban Paris, to almost twice more for working households living in wealthy suburban areas. As shown in Clerc and Marcus (2009), French private consumer elasticity to fuel prices largely depends on whether the consumer uses their car to go to work, as commuting represents the majority of kilometers driven in France. The Paris region is again special to this extent as reflected in Figure II.2, which shows that this region has an exceptionally low share of people using their car to go to work. We consider the three groups: young employed under the age of 30, employed (over 30-year-old), and not employed, with the latter including retirees and unemployed.

We moreover split households according to income. We proxy the buyer income by the median earnings of their age group at the precise municipality (*"commune"*) of each consumer and define two groups corresponding to the upper and lower half of this distribution. As group sizes are smaller in the Paris region, we do not distinguish along income dimensions for this region (see Table II.6 in the Appendix on page 49 for group sizes).

Little is known about the factors of heterogeneity in mileage for firms; thus, we use the same geographic partition as for households as it is partly related to infrastructure facilities. We also differentiate with respect to the business sector that is available in the data: industry and agriculture, rental, and trade/services. Little is known in the literature about the factors influencing firm's fuel-price sensitivity.

3.3 Diesel and gasoline cars

As shown by Hivert (2013), the advantage given to diesel cars in France is particularly salient in international comparison. Figure II.3 illustrates this specific position of France among European countries. Outside Europe, policies are much less favorable for diesel and diesel-engines virtually do not exist: in both Japan and the US, diesel cars make up about 2% of the overall vehicle fleet in circulation (Cames and Helmers 2013).

Within the time frame of the data used in this chapter, from January 2003 through November 2007, gasoline and diesel prices became more variable, with a general upward trend, after some time of relative stability, as shown in Figure II.4. Fuel prices varied considerably between €1.01 per liter and €1.38 per liter of gasoline, and between €0.75 and €1.21 per liter of diesel;¹⁹ this variation is about the same order of magnitude as the

¹⁹Monthly fuel prices are obtained from the French Ministry of Environment; we use sales-weighted national average prices available at <http://www.developpement-durable.gouv.fr/Prix-de-vente-moyens-des,10724.html>; retrieved on 06/12/2016. For diesel prices we use the price of car diesel oil (*"gazole"*), while for gasoline price we use premium unleaded gasoline (*"super sans plomb 95"*). All price indications in this chapter are deflated by the French




Figure II.3: Diesel fuel prices and market shares in Europe in 2012

Source: European Automobile Manufacturer's Association (ACEA). Price advantage of diesel is defined as the price differential (including taxes) between diesel and super unleaded gasoline (95 RON) divided by the latter.




Figure II.4: Monthly consumer fuel prices (incl. taxes) and cost per km (resulting from fuel prices (€) and fuel consumption (L/km) of new car purchases)

Source: French Ministry of Ecology and CCFA, authors' calculations.

policies we consider in this chapter.

Pre-tax prices for gasoline and diesel are highly correlated (correlation over 0.95) and

National Statistical Institute (INSEE) consumer price index, taking January 2008 as reference. Local prices are available only since 2007 and cannot be used here. However, the spatial variation is much lower than the temporal variation: the relative standard deviation is below 2 % for monthly fuel prices measured at the local (French "département") level in 2007, while it is above 10% for national monthly prices over the period 2003-2007.




Figure II.5: Monthly new registrations by fuel-type (in thousands, raw and smoothed series, studied period shaded in blue)

Source: CCFA, authors' calculations.

their difference is small (between -3 and 9 cents), so that we assume price variations of both depend equally on oil prices. The final fuel tax rates result from the combination of a fuel-type specific lump sum tax²⁰ and the proportional VAT of 19.6%. Over the whole examined period, diesel fuel prices are significantly lower than gasoline prices (Figure II.4) because of the lower TICPE tax on diesel fuel: in 2011, the consumption tax on energy products reached €0.61 per liter of gasoline, while it was €0.44 per liter of diesel. Moreover, firms benefit from an 80% rebate on VAT for diesel only, meaning that firms have an even stronger incentive to invest in diesel cars.

Diesel has a higher energy content so it produces more CO₂ per liter than gasoline: one liter of gasoline is transformed to 2.33 kg of CO₂ while one liter of diesel is transformed to 2.63 kg of CO₂.²¹ Besides this important global greenhouse gas, diesel cars also emit local pollutants like NO_x, as well as fine particles (see e.g. Cames and Helmers 2013). As a consequence, the French government has decided to adjust diesel taxation.

In France until 2017, the number of diesel cars sold has consistently been higher than the number of gasoline cars, and this difference has been increasing over the period under study in this chapter (Figure II.5). The overall number of new registrations is strongly seasonal, but is virtually constant over the years. The details of the choice between diesel and gasoline cars is amply discussed by Rouwendal and de Vries (1999).

Beyond fuel taxation, firms face an annual tax related both to the CO₂ class and to the fuel-type. Prior to 2004, the amount of this tax depended on horsepower; since 2004, it depends on CO₂ class, which is closely related to horse power but slightly less favorable

²⁰Consumption tax on energy products, "Taxe intérieure de consommation sur les produits énergétiques" (TICPE).

²¹The differences in CO₂ intensity are due to the differences in density of the fuel-types, see for example Demirel (2012). The mass of CO₂ per liter of fuel that weighs less than a kg might seem surprising; it results of the association of carbon elements from the fuel and ambient oxygen.

to diesel cars.²² As it may impact the preferences of firms toward one or other class, we use dummies for CO₂ classes in our estimations. This also accounts for marketing-based preferences for CO₂ classes (Koo et al. 2012) beyond direct valuation of fuel cost savings.

3.4 Cost per kilometer

Our focus lies on the consumer sensitivity to fuel prices when buying a new vehicle, via the cost of driving. We thus focus on the impact of the expected cost $E(p_{jft}^{km})$ at time t for the amount of fuel f needed to drive one km with the car jf . By definition, it depends on the car's fuel consumption ϕ_{jf} in L/100km, its fuel-type f (diesel or gasoline) and the expectations about fuel prices.

$$E(p_{jft}^{km}) = 1/100 \times \phi_{jf} [\mathbb{1}_{f=diesel} E_t(p_D) + \mathbb{1}_{f=gas} E_t(p_G)],$$

where p_D and p_G denote the fuel prices including tax for one liter of diesel and gasoline, respectively. ϕ_{jf} denotes the car's fuel consumption, measured in L/100km, which is the inverse of fuel efficiency as typically used in the US, measured in miles per gallon (MPG). Note that it is not equal to the total amount of fuel consumed, which results from the product of fuel consumption and mileage.

As a car is a durable good, the decision to buy a given product jf at time t should take into account the discounted utility of the future utilization of this car net of operating cost. We need therefore to take an assumption on how purchasers forecast future gasoline prices: according to Anderson et al. (2013), consumer beliefs regarding future fuel prices are indistinguishable from a no change forecast, consistent also with a random walk. However, given that new cars are rarely sold "off the rack," it usually takes a few months between purchase and the actual delivery and registration, which is our point of data collection. Thus, in our estimates, we do not use the contemporaneous fuel price but rather a three months lag of fuel prices. Alternative approaches in the literature include using moving averages, which are for example consistent with a purchaser belief in mean-reversion of fuel prices. In a model similar to ours, Klier and Linn (2013) use both current fuel prices and moving averages, and find that this assumption has no significant impact on parameter estimates, but standard errors are larger with moving averages.²³

Across different cars in our data, the price of driving one kilometer, i.e. the product of fuel price p_f and fuel consumption ϕ in liters per 100 km, covers a wide range from €2.60 per 100 km up to €30.9 per 100 km depending on the car.

²²The yearly amount of the tax ranges from €750 for the smaller cars to €4,500 for the biggest ones in 2014.

²³In an earlier version of this chapter, we estimated the results using moving averages over 6 months before purchase without finding significantly different results.

4 Econometric approach

4.1 Nested logit estimation

We take advantage of the fact that our data matches consumers and products: we assume that systematic differences in the valuation of the different characteristics are captured by observed purchaser types. We thus use the 28 consumer types as specified in Section 3.2 and estimate our model separately for each demographic group. Our approach is an alternative to two common ways to include demographic variation: random coefficient models *à la* BLP (Berry et al. 1995) and linear specifications as in Goldberg (1998). First, random coefficient models allow preferences to be shaped by aggregate distributions of household demographics, which is useful when only aggregate data is available.²⁴ As relevant heterogeneity is assumed to be observed and captured by the demographic groups here, we can refrain from using such complex models (see also Grigolon and Verboven 2014). Second, our specification is more flexible than the solution of, for instance, Goldberg (1995, 1998) who makes certain parameters linearly dependent on household demographics by including interactions of purchaser and product characteristics.

We thus aggregate individual choices within each consumer type, in order to recover the market shares of each product jf (model j of fuel-type f) up to an identifying normalization. As usual in the literature, identification stems from the normalization of the outside good's value to zero. As an intermediary step, we thus obtain a linear specification for the market share s_{djft} of the product jf at time t among consumer type d relatively to s_{d0t} the market share of the outside good for that same demographic group:

$$\ln(s_{djft}) - \ln(s_{d0t}) = \alpha_d + \beta_d p_{jft}^{km} + \gamma_{1d} p_{jft} + \gamma_{2d} X_{jft} + \sigma_{1d} \ln(s_{df|j}) + \sigma_{2d} \ln(s_{dj|s}) + \xi_{djft}, \quad (3)$$

where $s_{df|j} = \frac{s_{djft}}{s_{dj t}}$ is the relative share of purchases of fuel-type f within purchases of model j in each month t and $s_{dj|s} = \frac{s_{dj t}}{s_{ds t}}$ is the relative share of model j within the sales of segment s .

However, these shares are defined over the entire potential market size, which in our case – as in virtually all cases – is unknown. Indeed, this market size should contain only those who consider buying a car in a given period (and maybe *decide* not to). As detailed information on this market size is unknown, using some approximation is a standard procedure in this literature (for instance the seminal papers by McFadden 1978, Goldberg 1995), using for example most recent estimates of the population size or the number of people holding a driver's license. This number dramatically overstates the actual market with durable goods like cars, because in each given month only a small fraction of consumers considers buying a car. Moreover, when a large portion of new car registrations are made by firms and not by private owners, it is not clear whether

²⁴However, this comes at the cost of high computational complexity. This complexity is also shown to lead to numerical instability in some cases: Knittel and Metaxoglou (2014) find results often depend on starting values and optimization algorithms.

the number of driving license holders is relevant. Huang and Rojas (2014) show both theoretically and practically that coefficients estimated using such a wrong market size may be considerably biased.

To avoid this potential bias, we follow a suggestion by Huang and Rojas and reformulate Equation (3): by using quantities rather than market shares, the market size cancels out on the left-hand side. We are left with the log of the outside good's quantity, which we can move to the right-hand side and estimate it as part of the time-specific constant. Given the highly seasonal fluctuations of the number of purchases in Figure II.5, we allow this constant to vary with year and calendar month. The overall market size and the outside good quantity are not necessary to compute the relative shares $s_{dj|s}$ and $s_{df|j}$. Our main estimation equation is thus:

$$\ln(q_{djft}) = \alpha_d + \beta_d p_{jft}^{km} + \gamma_{1d} p_{jft} + \gamma_{2d} X_{jft} + \sigma_{1d} \ln(s_{df|j}) + \sigma_{2d} \ln(s_{dj|s}) + y_d + m_d + \xi_{djft}, \quad (4)$$

where q_{djft} stands for the number of sales of product jf . The characteristics of the new car, namely horsepower, CO₂ class, number of doors, fuel-type, car body (sedan, sport, compact, etc.) and brand are controlled for. Year and calendar month dummies, y_d and m_d , account for temporal trends as well as seasonality in aggregate new cars purchases.

The main parameter of interest is the parameter β_d measuring sensitivity to fuel prices. We use the parameters of Equation (4) to compute the fuel price elasticity, which takes into account both direct and indirect effects of an increase in fuel prices in the market share of one specific car. This elasticity can be approximated by:²⁵

$$\eta_{dsjf} = \frac{\partial s_{dsjf} / s_{dsjf}}{\partial p^e / p^e},$$

$$\approx (1 + t^{VAT}) p^e \left(\frac{\beta_d}{1 - \sigma_{1d}} \phi_{sjfd} + \left(\frac{\beta_d}{1 - \sigma_{2d}} - \frac{\beta_d}{1 - \sigma_{1d}} \right) \bar{\phi}_{sjd} + \frac{\beta \sigma_{2d}}{1 - \sigma_{2d}} \bar{\phi}_{sd} \right). \quad (5)$$

Equation (4) is estimated using the generalized method of moments separately for each demographic group, assuming these groups homogeneous enough to include only buyers with the same demand parameters.

4.2 Endogenous variables and instruments

Gas prices can be considered as exogenous in the French case, as France represents about 2% of world oil consumption and produces less than 0.1% of the world production.²⁶ French gas prices are defined by the international energy market, on which France has only a limited weight (which may be not the case for the US, see Davis and Kilian (2011) for a discussion).

By contrast, the vehicle price p_{jft} is endogenous, as it is the result of demand and supply which by assumption vary with the unobserved attractiveness ξ_{djft} . As it is usual in the literature, we use a set of instruments based on the characteristics of potential substi-

²⁵Details of elasticity computation are given in the Appendix B on page 50.

²⁶In 2009, see <http://www.eia.gov/countries/country-data.cfm>; retrieved on 14/03/2015.

tutes aiming at capturing market density, and thus beyond production cost, the variation in mark-ups. More specifically, in a multi-product Bertrand competition framework, one can derive a set of instruments based on the sums of each characteristics of other models produced by the *same firm* in the same segment and those of *competing firms* (Berry et al. 1995, henceforth “BLP”). This measure is computed twice; once over all products within the same nest, and another time over all products in all other nests. Importantly, we use yearly list prices and thus assume that purchase prices do not vary with fuel prices. In the short term, this is likely to be true, as list prices are set on a much longer horizon than fuel prices; in the long term, list prices can obviously adapt to fuel price variation.

Armstrong (2016) argues that in markets with a large number of heterogeneous goods, BLP instruments are no longer sufficiently strong. Thus, we add cost-shifters, such as the prices of raw materials, that provide exogenous variations in market prices as they are related to supply but not demand. Thus, we use the price indices of iron (current and lagged value) and indices of export prices of tires as instruments, both weighted by the car’s weight. These cost shifters appear strongly correlated to vehicle prices.

Within segment, the market share $s_{dj|s}$ is endogenous by definition. As for the price, we use BLP-style instruments for this variable and further add the number J_s of offered goods per segment s .

Finally, we instrument the within-model market share $s_{df|j}$ by the difference in characteristics of gasoline and diesel versions, as well as the difference in costs shifters for these two versions, capturing the relative attractiveness of each version.

As pointed out by Bound et al. (1995), using many over-identifying restrictions as we do can lead to misleading results if the instruments are weak. In case of only one endogenous variable, it is now common to test the strength of the instruments by using on the first-stage F-values, as proposed by Stock and Yogo (2005). As shown by Sander-son and Windmeijer (2016), this method can be extended to regressions with multiple endogenous variables: for each endogenous variable, the relevant test statistic is then the first-stage F-value *conditional* on the other two endogenous regressors, that can be compared to the values tabulated by Stock and Yogo (2005). We compute these test statistics for each of our three endogenous variables and for each demographic group. At a 5% significance level, we can reject for most regressions a bias of the 2SLS regression relative to an OLS of more than 5%; in only two cases (out of eighty) we can only reject biases superior to 20% (cf. Tables II.13, II.14 and II.15 in the Appendix on page 60). One case is problematic, as we cannot reject that our instruments are too weak to identify the within-model parameter σ_{1d} for the purchases by car rental companies in the Paris suburban area. This group is small and aggregate results are virtually identical if we drop it. Thus, we are confident that our results are not biased by weak-instrument effects.

5 Empirical results

Our aggregate outcomes of interest are: the share of diesel cars (local pollution), average fleet fuel consumption (international fuel dependency) and average CO₂ intensity (global

pollution).

The presentation of the empirical results proceeds in three steps: first, we present the aggregate elasticities of market shares, diesel share, fuel consumption, and CO₂ emission intensity.²⁷ Then, these elasticities are used to compute ex ante estimates of the impact of two policies, one equalizing tax on diesel and gasoline; the other taxing carbon directly. The two policy scenarios are calibrated such that they are revenue-equivalent for the implementing government in absence of consumer reaction. The raw coefficients cannot be interpreted directly, but we discuss them in the Appendix C on page 55, where we also compute the demand elasticities for some popular car models.

5.1 Aggregate elasticities to fuel price variation

We model the aggregate elasticities to a change in fuel prices (both gasoline and diesel) through an international oil price shock. As diesel engines tend to be more efficient with an average fleet fuel consumption of 5.6L/100km versus 6.8L/100km for gasoline engines (Table II.7 in the Appendix on page 50), an increase of fuel prices raises the share of diesel cars among new purchases π^D (see elasticity η_D in Table II.3).²⁸ Consequently, the average fleet fuel consumption decreases as well as average CO₂ intensity. However, all these effects have a small magnitude.

These results can be compared to some previous estimates obtained in the literature. Using aggregated data on several European car markets, Klier and Linn (2011) estimate that a 1\$ increase in fuel prices per gallon would increase the average miles-per-gallon (MPG) efficiency in France by 0.21, implying an average fuel consumption elasticity η_ϕ of -0.017.²⁹ This value is similar to our estimate and much lower than the value they find for the US: there, 1\$ decreases the average MPG by 1.03, implying an average fuel consumption elasticity of -0.042. Our estimate is smaller than the estimates by Clerides and Zachariadis (2008), who find a short term elasticity of average fleet fuel consumption to fuel prices equal to -0.08 for the EU, using aggregate data. Klier and Linn (2011) also estimate that a hypothetical policy equalizing diesel and gasoline prices reduces the diesel market share in France by 1.4 percentage points only; much less than suggested by our estimate of around 4 percentage points.

5.2 Tax alignment

These estimates allow us to simulate the impact of a policy that aligns diesel and gasoline taxes. Leaving gasoline taxes unchanged, this policy raises diesel taxes by almost a third, from 43 cent/liter to 60 cent/liter. Furthermore, this policy abandons the VAT advantage for corporate diesel cars, increasing it to the standard rate of 19.6%.

²⁷See the Appendix B on page 50 for details on the computation of these elasticities.

²⁸ π^D is the market share of diesel cars *among purchased cars* whereas the market shares s_j , s_s etc. are defined on the whole market, including the outside good.

²⁹Brons et al. (2008) analyze more in detail the aggregate elasticity of fuel demand, resulting of the elasticities of mileage, fuel consumption and car ownership; their meta-study also finds this elasticity to be empirically small.

Table II.3: Elasticities with respect to fuel prices: diesel share, average fleet fuel consumption (L/km) and CO₂ intensity (g/km)

	Diesel share	Fuel cons.	CO ₂
	η_D	η_ϕ	η_{CO_2}
Households	0.026*** (0.003)	-0.013*** (0.001)	-0.015*** (0.001)
Firms	0.017*** (0.003)	-0.004*** (0.001)	-0.006*** (0.001)
Total	0.029*** (0.003)	-0.010*** (0.001)	-0.012*** (0.001)

Source: CCFA, authors calculations. Estimates rely on the parameters of Equation (4) estimated by GMM separately for each type of consumers. Standard errors in parentheses are estimated by bootstrap (500 replications).

As expected, the induced variation in diesel share is negative and strong: since taxes only increase for diesel, they would push many purchasers to substitute for a gasoline-fueled car. We find that such a policy would reduce the aggregate share of diesel cars in overall sales by 5.9%, that is from 69% to 65% (Table II.4). This decrease in diesel sales comes mostly from households who substitute much more easily away from diesel engines, rather than from firms (7.4% and 3.6% reduction, respectively).

This result can be compared to the one in Klier and Linn (2011) who also evaluate a hypothetical policy of equalizing diesel and gasoline prices. At the European level, their estimates suggest that the impact of such a policy on the market share of diesel cars would be negligible (less than 1%). Two elements explain this difference. First, our analysis is focused on France, where the gap between gasoline and diesel taxes is the highest of all countries they consider: the hypothetical policy change is strong which is not the case for other countries.³⁰ Second, as they emphasize, Klier and Linn (2011) cannot distinguish in their data company cars from privately owned cars. According to our estimates, firms are much less sensitive to fuel prices (Table II.3).

Gasoline cars consume more liters of fuel per km but produce 13% less CO₂ per liter of fuel – gasoline is a less energy-rich combustible. The effect of a demand shift toward gasoline cars on CO₂ is thus a priori ambiguous. According to our estimations, substitutions between gasoline and diesel cars have only a marginal impact on both fuel consumption of the new vehicle fleet and CO₂ intensity. It *increases* fuel consumption (Table II.4) and *reduces* the average CO₂ intensity of newly purchased cars. Both effects are significant but small: in spite of the large jump in diesel tax, average fleet fuel consumption increases only by 0.44% and average CO₂ intensity decreases by 0.12%. The absolute magnitudes of these changes are small: fuel consumption increases by 26 mL/100km from the average of 6L/km and CO₂ intensity is reduced by 180mg/km from the average of 152g/km.

³⁰Estimates detailed by countries are available in a previous working paper (Klier and Linn 2011). They obtain that the diesel market share in France would decrease by 1.4 percentage points. This reduction is higher than the effect in most other countries they examine.

Table II.4: Percentage impact of a carbon tax and a tax alignment on diesel share, average fleet fuel consumption (L/km) and CO₂ intensity (g/km)

	Tax alignment			Carbon tax		
	Diesel share	Fuel cons.	CO ₂	Diesel share	Fuel cons.	CO ₂
	$\Delta^{tD} \eta_D$	$\Delta^{tD} \eta_\phi$	$\Delta^{tD} \eta_{CO_2}$	$\Delta^{tC} \eta_D$	$\Delta^{tC} \eta_\phi$	$\Delta^{tC} \eta_{CO_2}$
Households	-7.43*** (0.36)	0.50*** (0.03)	-0.13*** (0.01)	0.15** (0.07)	-0.43*** (0.02)	-0.43*** (0.02)
Firms	-3.55*** (0.46)	0.28*** (0.09)	-0.11* (0.06)	0.65*** (0.12)	-0.21*** (0.03)	-0.15*** (0.03)
Total	-5.94*** (0.32)	0.44*** (0.04)	-0.12*** (0.02)	0.59*** (0.07)	-0.37*** (0.02)	-0.33*** (0.02)

Source: CCFA, authors calculations. Estimates rely on the parameters of Equation (4) estimated by GMM separately for each type of consumers. Instrumental variables for prices are the price indices of iron (current and lagged value) and indices of export prices of tires, interacted with the car model's weight. Standard errors in parentheses are estimated by bootstrap (500 replications).

5.3 Carbon tax

We also predict the impact of a carbon tax, i.e. a tax increase that is proportional to the carbon emissions of each fuel-type. The amounts are calibrated such that the government revenue is equal to the previous tax alignment policy, yielding a price of €51 per tonne of CO₂. This results in an increase of 11.9 cent/liter of gasoline and 13.4 cent/liter of diesel, representing around 9% of the average end-user price.³¹ A very similar but less ambitious policy has been voted in France in 2014, leading to a progressive increase in fuel taxes up to €30.5/tCO₂ in 2017.³²

The impact $\Delta^{tC} \eta_D$ of this carbon tax policy on the share of diesel engines sold is positive, but very small: it increases the diesel share by 0.6% (Table II.4). This is the result of two contrasting effects: on the one hand, the carbon tax is higher on diesel than on gasoline, but on the other hand, diesel cars are more fuel-efficient. The incentive for purchasers to buy more fuel-efficient cars seems to dominate. The carbon tax reduces average fleet fuel consumption as well as average CO₂ intensity (Table II.4). The impacts are significant but again very small. The fuel consumption decreases by 0.37%, which is however only around 22 mL/100km from the average of 6L/km; CO₂ emission intensity shift by 0.33% which is 500mg/km from the average of 152g/km.

The impact of both policies on fuel consumption and CO₂ intensity is economically small. The main difference is that leveling out the diesel tax advantage induces a noticeable shift away from diesel engines, thus reducing local pollution. Moreover, the carbon tax achieves a larger reduction in CO₂ intensity and furthermore reduces fuel consumption, thus leading – on its modest level – to a lower dependency on foreign petrol imports.

³¹This scenario maintains the VAT rebate for diesel cars of corporate consumers.

³²See the website of the French ministry of environment:

<https://www.ecologique-solaire.gouv.fr/fiscalite-carbone>; retrieved on 09/09/2017.

Table II.5: Robustness checks: percentage impact of carbon tax and tax alignment on diesel share, average fleet fuel consumption (L/km) and CO₂ intensity (g/km)

	Tax alignment			Carbon tax		
	Diesel share	Fuel cons.	CO ₂	Diesel share	Fuel cons.	CO ₂
	$\Delta^{t_D} \eta_D$	$\Delta^{t_D} \eta_\phi$	$\Delta^{t_D} \eta_{CO_2}$	$\Delta^{t_C} \eta_D$	$\Delta^{t_C} \eta_\phi$	$\Delta^{t_C} \eta_{CO_2}$
<i>Main specification - including degenerate nests (gas- or diesel-only models)</i>						
Households	-9.37*** (0.35)	0.80*** (0.03)	0.02*** (0.01)	0.55*** (0.08)	-0.50*** (0.02)	-0.47*** (0.02)
Firms	-3.85*** (0.41)	0.24*** (0.08)	-0.19*** (0.05)	0.70*** (0.11)	-0.23*** (0.02)	-0.17*** (0.03)
Total	-7.15*** (0.32)	0.62*** (0.04)	-0.06*** (0.02)	0.99*** (0.06)	-0.44*** (0.02)	-0.36*** (0.02)
<i>Alternative specification - Nests (segment > model)</i>						
Households	-9.17*** (0.38)	0.79*** (0.03)	0.02*** (0.01)	0.49*** (0.08)	-0.48*** (0.03)	-0.45*** (0.02)
Firms	-3.51*** (0.54)	0.22** (0.09)	-0.17*** (0.06)	0.59*** (0.14)	-0.22*** (0.03)	-0.16*** (0.03)
Total	-6.84*** (0.40)	0.59*** (0.05)	-0.05** (0.02)	0.97*** (0.08)	-0.42*** (0.02)	-0.35*** (0.02)
<i>Main specification - BLP-instruments only</i>						
Households	-8.67*** (0.39)	0.60*** (0.04)	-0.13*** (0.01)	0.28*** (0.08)	-0.51*** (0.02)	-0.50*** (0.02)
Firms	-3.12*** (0.55)	0.29*** (0.10)	-0.06 (0.06)	0.65*** (0.15)	-0.16*** (0.03)	-0.10*** (0.03)
Total	-6.27*** (0.42)	0.49*** (0.05)	-0.10*** (0.02)	0.84*** (0.07)	-0.42*** (0.02)	-0.36*** (0.02)
<i>Main specification - without purchaser heterogeneity</i>						
Total	-7.45*** (0.77)	0.45*** (0.07)	-0.26*** (0.02)	0.26*** (0.05)	-0.61*** (0.06)	-0.60*** (0.05)

Source: CCFA, authors calculations. Estimates rely on the parameters of Equation (4) estimated by GMM separately for each type of consumers. Instrumental variables for prices are the price indices of iron (current and lagged value) and indices of export prices of tires, interacted with the car model's weight. Standard errors in parentheses are estimated by bootstrap (500 replications).

5.4 Robustness checks

We estimate several alternative specifications to check that results are not driven by our main specification choice, but also to emphasize the impact of individual hypothesis underlying this main specification. On the whole, the estimated impact of our policy scenarios remains at a similar order of magnitude across specifications.

Our first test includes all models, i.e. including those that are available only with either gasoline or diesel motor. In our main specification, we drop these models as they lead to “degenerate” nests at the end of the decision tree, where a model-branch only includes one product. While the aggregate elasticities (Table II.16 in the Appendix on page 62) appear similar to our main specification, the policy simulation shows that this model slightly over-estimates the policy impact while leading broadly to the same conclusions.

In the same spirit, our second test uses a more commonly used model accounting only for two levels: purchasers choose a segment and then a product within that segment. The two fuel-type versions of a model then count as independent products, which is the same as constraining all σ_{1d} coefficients to zero. The elasticities are similar to the previous test

(Table II.5) and just slightly stronger than our main specification. Although the changes are small, we still reject this more constrained model as in our main estimation σ_{1d} was significantly different from zero for almost all demographic groups (Table II.9 in the Appendix on page 57).

Our third test drops the cost-shifter instruments and includes only the BLP-style instruments. Again, the elasticities are very similar and the policy impacts give the same intuition, but overstate the impact of a carbon tax on the diesel share.

As a last test, we estimate the model jointly for all demographic groups, which means we do not account for consumer heterogeneity. Bento et al. (2012) suggest that unaccounted heterogeneity biases estimated elasticity downwards, which we do not find here (Table II.16 in the Appendix on page 62). Quite the contrary, elasticities and estimated policy impacts overstate the consumer reaction in our case (Table II.5).

Our main specification still seems most appropriate, but these alternative specifications do not dramatically change the implications of this chapter.

6 Conclusion

This chapter estimates the short-term impact of fuel prices on new automobile purchases of both households and firms. These estimates allow us to compute elasticities which we aggregate to estimate ex ante the impact of two tax reforms. Using a nested logit specification, we control for hedonic valuation of a large range of car characteristics. We also account for taste heterogeneity between consumer groups, in particular between private and corporate purchases.

Our aggregate outcomes of interest are: the share of diesel cars (local pollution), average fleet fuel consumption (international fuel dependency), and average CO₂ intensity (global pollution). We use our estimates to examine a (hypothetical) policy equalizing tax levels on gasoline and diesel. We find that this policy decreases the share of diesel cars in sales from 69% to 65%. As purchasers would substitute to (less efficient) gasoline cars, the average fuel consumption would rise in response to this policy, while at the same time average CO₂ intensity would slightly *decrease* as gasoline cars emit less CO₂ per liter of used fuel. The examined carbon tax – which implements a much higher carbon price than the recently voted French policy – is expected to slightly increase the share of diesel cars among new purchases. It decreases both fuel consumption and CO₂ intensity significantly, but the overall amounts stay low.

All in all, the estimated effects of these two tax policies are significant but economically small in the short-run, i.e. holding supply constant. This is even more noteworthy, as one might argue that our policy scenarios are somewhat overly ambitious and might not be politically feasible. Overall, fuel taxes do not appear to be a strong policy tool for influencing car choices in the short-run.

An important advantage is provided by our individual registration data, as we can account for purchaser heterogeneity and our estimates are thus less prone to omitted sorting bias. Indeed, purchaser types react differently to fuel tax changes. A large part

of aggregate market reaction comes from households, and particularly from urban and non working consumers. To our knowledge, this important distinction between household and firm purchases is not accounted for in earlier related literature, although firm purchases constitute about a third of the market in our sample. Corporate purchases are particular important for the diesel share, as firms buy a lot more diesel-powered cars and are less likely to substitute away from them.

A limitation of this chapter is that our simple demand model does not take into account long-run shifts on the supply side. While one can be confident that the monthly fuel price variation used for identification in this article does not impact the characteristics of available cars instantaneously, it is likely that producers react more to long-term shifts: if fuel efficiency becomes more valuable, they might in the medium-run adjust their list prices and in the long-run adjust the products developed and offered. For Klier and Linn (2011) this means that these short-run results underestimate the true impact on fuel efficiency and emissions, which would be enhanced by the producer's reactions. However, as shown by Verboven (2002), producer price reaction should counteract purchaser reaction to changes in differential fuel taxation. However, one could argue like Goldberg (1998) that a short-term consumer reaction as small as suggested by our estimates is unlikely to shift supply, so that the long-run effect should be small as well.

The aim of environmental policy is ultimately not to increase fuel efficiency, but to decrease CO₂ emissions which result from the interaction of fuel consumption and mileage. Additional research is needed to clarify the impact of fuel efficiency on car mileage. Previous research suggests that rebound effects might reduce any impact on fuel consumption (see for example Austin and Dinan 2005, Frondel et al. 2012), so that our (already small) estimated effects become even less economically and environmentally significant. Nevertheless, the change in the composition of the vehicle fleet impacts fuel efficiency in the long run as cars circulate on average for 13 years in France (Bilot et al. 2013).

We do not use any data on mileage nor assume anything on car lifetime and discounting, so that we remain agnostic on the actual profit a consumer realizes with fuel efficiency. As a consequence, we cannot evaluate welfare effects of the policy such as Bento et al. (2009) or Bureau (2011) or the rationality (or myopia) of consumers such as reviewed in Greene (2010) and Helfand et al. (2011). To our knowledge, there is no study that includes mileage elasticity to fuel prices and to fuel efficiency, as well as potentially elastic lifetime, so that computations usually remain back-of-the-envelope sketches (e.g. Grigolon et al. 2014, Allcott and Wozny 2014, Busse et al. 2013³³). Nevertheless, our estimated consumer reactions are too small to fully account for the change in operating cost if utilization does not change. In this light, it may seem surprising that corporate purchases are even less reactive to fuel price changes than household purchases. However, similar results have been obtained on the market for airline tickets. Firms can deduce total fuel cost from taxes and may be able to pass costs through to consumers. These

³³These papers account for mileage at a detailed car- or consumer-level but assume zero elasticity; they can thus not account for well documented phenomena such as the "rebound effect" (Small and Dender 2007).

factors may explain why they react less to fuel prices than households. Further research is needed to clarify whether this is due to differences in mileage or whether there are behavioral and organizational factors at play.

7 Bibliography

- Adda J and Cooper R (2000) Balladurette and juppette: a discrete analysis of scrapping subsidies. *Journal of Political Economy*, 108(4): 778–806.
- Allcott H and Wozny N (2014) Gasoline prices, fuel economy, and the energy paradox. *Review of Economics and Statistics*, 96(5): 779–795.
- Anderson S. T, Kellogg R, and Sallee J. M (2013) What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management*, 66(3): 383–403.
- Armstrong T. B (2016) Large market asymptotics for differentiated product demand estimators with economic models of supply. *Econometrica*, 84(5): 1961–1980.
- Austin D and Dinan T (2005) Clearing the air: the costs and consequences of higher CAFE standards and increased gasoline taxes. *Journal of Environmental Economics and Management*, 50(3): 562–582.
- Baccaini B, Sémécurbe F, and Thomas G (2007) Les déplacements domicile-travail amplifiés par la périurbanisation. *INSEE Première*, 1129: 1–4.
- Bento A. M, Goulder L. H, Jacobsen M. R, and Von Haefen R. H (2009) Distributional and efficiency impacts of increased US gasoline taxes. *American Economic Review*, 99(3): 667–699.
- Bento A. M, Li S, and Roth K (2012) Is there an energy paradox in fuel economy? A note on the role of consumer heterogeneity and sorting bias. *Economics Letters*, 115(1): 44–48.
- Berry S, Levinsohn J, and Pakes A (1995) Automobile prices in market equilibrium. *Econometrica*, 63(4): 841–890.
- Berry S. T (1994) Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2): 242–262.
- Bilot H, Breteau V, and Weber S (2013) Quels effets d’un changement de taxation des carburants sur la diesélisation du parc automobile et les émissions de polluants ? *La revue du CGDD*, Juin 2013: 49–57.
- Bound J, Jaeger D, and Baker R (1995) Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American Statistical Association*, 90(430): 443–450.
- Brons M, Nijkamp P, Pels E, and Rietveld P (2008) A meta-analysis of the price elasticity of gasoline demand. A SUR approach. *Energy Economics*, 30(5): 2105 – 2122.
- Bureau B (2011) Distributional effects of a carbon tax on car fuels in France. *Energy Economics*, 33(1): 121 – 130.

- Busse M. R, Knittel C. R, and Zettelmeyer F (2013) Are consumers myopic? Evidence from new and used car purchases. *American Economic Review*, 103(1): 220–56.
- Cames M and Helmers E (2013) Critical evaluation of the European diesel car boom-global comparison, environmental effects and various national strategies. *Environmental Sciences Europe*, 25(1): 1–22.
- Cardell N. S (1997) Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, 13(2): 185–213.
- Clerc M and Marcus V (2009) Élasticités-prix des consommations énergétiques des ménages. Working Papers of the DESE 8, INSEE.
- Clerides S and Zachariadis T (2008) The effect of standards and fuel prices on automobile fuel economy: an international analysis. *Energy Economics*, 30(5): 2657–2672.
- Davis L. W and Kilian L (2011) Estimating the effect of a gasoline tax on carbon emissions. *Journal of Applied Econometrics*, 26(7): 1187–1214.
- Demirel Y (2012) *Energy: Production, Conversion, Storage, Conservation, and Coupling*: Springer.
- D'Haultfœuille X, Givord P, and Boutin X (2014) The environmental effect of green taxation: the case of the French bonus/malus. *The Economic Journal*, 124(578): F444–F480.
- Fronzel M, Ritter N, and Vance C (2012) Heterogeneity in the rebound effect: further evidence for Germany. *Energy Economics*, 34(2): 461 – 467.
- Fronzel M and Vance C (2014) More pain at the diesel pump? An econometric comparison of diesel and petrol price elasticities. *Journal of Transport Economics and Policy (JTEP)*, 48(3): 449–463.
- Goldberg P. K (1995) Product differentiation and oligopoly in international markets: the case of the U.S. automobile industry. *Econometrica*, 63(4): 891–951.
- (1998) The effects of the corporate average fuel efficiency standards in the US. *The Journal of Industrial Economics*, 46(1): 1–33.
- Greene D. L (2010) How consumers value fuel economy: a literature review. Technical Report, Environmental Protection Agency.
- Grigolon L, Reynaert M, and Verboven F (2014) Consumer valuation of fuel costs and the effectiveness of tax policy: evidence from the European car market. Discussion Paper 10301, CEPR.
- Grigolon L and Verboven F (2014) 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.

- Helfand G, Wolverton A et al. (2011) Evaluating the consumer response to fuel economy: a review of the literature. *International Review of Environmental and Resource Economics*, 5(2): 103–146.
- Hivert L (2013) Short-term break in the French love for diesel? *Energy Policy*, 54: 11–22.
- Huang D and Rojas C (2014) Eliminating the outside good bias in logit models of demand with aggregate data. *Review of Marketing Science*, 12(1): 1–36.
- Klier T and Linn J (2010) The price of gasoline and new vehicle fuel economy: evidence from monthly sales data. *American Economic Journal: Economic Policy*, 2(3): 134–153.
- (2011) Fuel prices and new vehicle fuel economy in Europe. Working Papers 1117, Massachusetts Institute of Technology, Center for Energy and Environmental Policy Research.
- (2013) Fuel prices and new vehicle fuel economy: comparing the United States and Western Europe. *Journal of Environmental Economics and Management*, 66(2): 280–300.
- Knittel C. R and Metaxoglou K (2014) Estimation of random-coefficient demand models: two empiricists' perspective. *Review of Economics and Statistics*, 96(1): 34–59.
- Koo Y, Kim C. S, Hong J, Choi I.-J, and Lee J (2012) Consumer preferences for automobile energy-efficiency grades. *Energy Economics*, 34(2): 446 – 451.
- Mayeres I and Proost S (2001) Should diesel cars in Europe be discouraged? *Regional Science and Urban Economics*, 31(4): 453 – 470.
- McFadden D (1978) Modeling the choice of residential location. *Transportation Research Record*, 673: 72–77.
- Miravete E. J, Moral M. J, and Thurk J (2015) Innovation, emissions policy, and competitive advantage in the diffusion of European diesel automobiles. Discussion Paper 10783, CEPR.
- Rouwendal J and de Vries F (1999) The taxation of drivers and the choice of car fuel type. *Energy Economics*, 21(1): 17 – 35.
- Sallee J. M, West S. E, and Fan W (2016) Do consumers recognize the value of fuel economy? Evidence from used car prices and gasoline price fluctuations. *Journal of Public Economics*, 135: 61–73.
- Sanderson E and Windmeijer F (2016) A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics*, 190(2): 212–221.
- Small K. A and Dender a, K. V (2007) Fuel efficiency and motor vehicle travel: the declining rebound effect. *Energy Journal*, 28(1): 25–51.

- Small K. A and Van Dender K (2007) Fuel efficiency and motor vehicle travel: the declining rebound effect. *The Energy Journal*: 25–51.
- Stock J. H and Yogo M (2005) Testing for weak instruments in linear IV regression. In D. W. K. Andrews and J. H. Stock (eds.) *Identification and Inference for Econometric Models*: Cambridge University Press: 80–108, Cambridge Books Online.
- Verboven F (1996) International price discrimination in the European car market. *The RAND Journal of Economics*, 27(2): 240–268.
- (2002) Quality-based price discrimination and tax incidence: evidence from gasoline and diesel cars. *The RAND Journal of Economics*, 33(2): 275–297.

Appendices

A Descriptive statistics

Table II.6: Distribution of demographic groups among buyers (%)

	<i>Private consumers</i>						
	Not employed		Young employed (<30)		Employed (≥30)		
Income	Low	High	Low	High	Low	High	Total
Urban	150,214	82,692	389,903	192,957	679,981	646,949	2,142,696
	5.0%	2.5%	8.7%	8.3%	1.7%	1.5%	27.6%
Suburban/rural	136,187	116,348	246,876	331,066	450,728	564,686	1,845,891
	1.7%	1.5%	3.2%	4.2%	5.8%	7.2%	23.6%
Paris urban	40,298		186,758		486,700		713,756
	0.5%		2.4%		6.2%		9.1%
Paris suburban	11,069		45,160		81,893		138,122
	0.1%		0.6%		1.0%		1.8%
Total	536,808		1,392,720		2,910,937		4,840,465
	11.3%		27.3%		23.5%		62.1%
	<i>Firm purchases</i>						
	Industry & agriculture		Car rental & repairing		Trade & services		Total
Urban	307,871		1,261,364		374,754		1,567,383
	3.9%		16.1%		4.8%		24.8%
Suburban/rural	113,947		66,416		137,182		383,855
	1.5%		0.8%		1.8%		4.1%
Paris urban	203,606		313,880		172,532		565,762
	2.6%		4.0%		2.2%		8.8%
Paris suburban	7,674		4,083		25,129		47,902
	0.1%		0.1%		0.3%		0.5%
Total	633,098		1,645,743		709,597		2,564,902
	8.1%		21.0%		9.1%		38.2%

Source: CCFA, authors' calculations.

Table II.7: Descriptive statistics of car characteristics

	Mean	Coefficient of variation (%)	Percentiles		
			25%	Median	75%
Gasoline (N= 2,376,527)					
Car price (€)	16,606	69.4	11,738	13,975	18,800
Cost of driving 100 km (€)	8.4	22.7	7.3	8.1	9.1
Horse power (kW)	70	48.8	54	60	80
Fuel consumption (L/100km)	6.8	21.7	6.0	6.5	7.4
CO ₂ intensity (g/km)	159.3	21.7	139.0	152.0	172.0
Diesel (N= 5,452,376)					
Car price (€)	22,968	41.0	16,783	21,875	26,236
Cost of driving 100 km (€)	5.7	27.1	4.8	5.4	6.3
Horse power (kW)	78	34.6	63	78	88
Fuel consumption (L/100km)	5.6	24.5	4.7	5.4	6.0
CO ₂ intensity (g/km)	147.0	24.5	124.0	141.0	157.0

Note: The coefficient of variation, or unitized risk, is the ratio of the standard error to the mean.

Source: CCFA, authors' calculations.

B Details on the computation of the elasticities

For the computation of elasticities it is useful to introduce a decomposition of the fuel price per km

$$p_j^{km} = \phi_j(1 + t^{VAT}) [p^e + 1_{j=diesel}(t_D) + 1_{j=gas}(t_G)] , \quad (6)$$

where ϕ_j denotes the characteristic fuel-consumption (inverse of fuel-efficiency, thus in liter of fuel per kilometer), t_D and t_G the consumption tax rates for energy products for one liter of diesel and gasoline, respectively, and t^{VAT} the VAT rate. The fuel prices excluding taxes for diesel and gasoline are very similar and strongly correlated as they are driven by oil prices; for the sake of simplicity, we thus assume that the fuel prices per liter excluding taxes are the same for diesel and gasoline, denoted p^e .

The demand elasticity η_{sjf} for a given product with respect to oil price p^e exclusive of tax at a given point in time can be computed using parameters corresponding to the demand model. Fuel prices affect *all* products proportionally to their fuel consumption: both the nominator and the denominator of the market shares are impacted. In order to find this elasticity, let us differentiate Equation (3) for the model j in segment s and of fuel-type f , using the definition of the cost per kilometer:³⁴

$$\frac{\partial s_{sjf}}{s_{sjf}} - \frac{\partial s_0}{s_0} = \beta \partial p^e (1 + t^{VAT}) \phi_{sjf} + \sigma_1 \left(\frac{\partial s_{sjf}}{s_{sjf}} - \frac{\partial s_j}{s_j} \right) + \sigma_2 \left(\frac{\partial s_j}{s_j} - \frac{\partial s_s}{s_s} \right) \quad (7)$$

³⁴For the sake of readability, we omit the index for demographic groups and do not state the obvious aggregation over these groups for all equations in this section.

or slightly rearranged:

$$\partial s_{sjf} - \frac{\partial s_0}{s_0} s_{sjf} = \beta \partial p^e (1 + t^{VAT}) \phi_{sjf} s_{sjf} + \sigma_1 (\partial s_{sjf} - s_{sjf} \frac{\partial s_j}{s_j}) + \sigma_2 s_{sjf} (\frac{\partial s_j}{s_j} - \frac{\partial s_s}{s_s}). \quad (8)$$

We then aggregate this last equation over both fuel-type versions of the same model, to obtain the change in the market share of one model j in one segment s :

$$\begin{aligned} \partial s_j - \frac{\partial s_0}{s_0} s_j &= \sum_{f \in j} (\partial s_{sjf} - \frac{\partial s_0}{s_0} s_{sjf}) \\ &= \beta \partial p^e (1 + t^{VAT}) \underbrace{\sum_{f \in j} \phi_{sjf} s_{sjf}}_{\bar{\phi}_j s_j} \\ &\quad + \sigma_1 \left(\underbrace{\sum_{f \in j} \partial s_{fjs}}_{\partial s_j} - \frac{\partial s_j}{s_j} \underbrace{\sum_{f \in j} s_{sjf}}_{\partial s_j} \right) \\ &\quad + \sigma_2 \left(\frac{\partial s_j}{s_j} - \frac{\partial s_s}{s_s} \right) \underbrace{\sum_{f \in j} s_{sjf}}_{s_j} \end{aligned}$$

We define $\bar{\phi}_j$ as the sales-weighted average fuel consumption of both fuel-type versions of the same model j . Thus we obtain that

$$(1 - \sigma_2) \frac{\partial s_j}{s_j} = \beta \partial p^e (1 + t^{VAT}) \bar{\phi}_j - \sigma_2 \frac{\partial s_s}{s_s} + \frac{\partial s_0}{s_0}. \quad (9)$$

Aggregating further, we can also recover the relative variation in the market share of segment s ($\frac{\partial s_s}{s_s}$) or of the outside good ($\frac{\partial s_0}{s_0}$) by summing on respectively all cars in the same segment, and all new cars. For segment s , we obtain that:

$$\frac{\partial s_s}{s_s} = \beta \partial p^e (1 + t^{VAT}) \bar{\phi}_s + \frac{\partial s_0}{s_0},$$

while for the overall number of sold cars we get:

$$\frac{\partial s_0}{s_0} = -\beta \partial p^e (1 + t^{VAT}) \bar{\phi} (1 - s_0).$$

Combining these expressions in 7 we finally can compute the elasticity η_{sjf} as:

$$\begin{aligned} \eta_{sjf} &= \frac{\partial s_{sjf} / s_{sjf}}{\partial p^e / p^e} \\ &= \beta (1 + t^{VAT}) p^e (\rho_1 \phi_{sjf} + (\rho_2 - \rho_1) \bar{\phi}_j - (\rho_2 - 1) \bar{\phi}_s) - \beta (1 + t^{VAT}) p^e \bar{\phi} (1 - s_0) \\ &\approx \beta (1 + t^{VAT}) p^e (\rho_1 (\phi_{sjf} - \bar{\phi}_j) + \rho_2 (\bar{\phi}_j - \bar{\phi}_s) + \bar{\phi}_s), \end{aligned} \quad (5)$$

where $\rho_i = \frac{1}{1 - \sigma_i} \in [1, +\infty]$. The demand elasticity depends on the parameter β measuring

sensitivity to fuel prices, the VAT rate t^{VAT} ,³⁵ as well as on the current price of fuel and the car's fuel consumption ϕ_{sjf} relative to the average fuel economy of its substitutes (within the same model $\bar{\phi}_j$, within its segment $\bar{\phi}_s$ and among all sales $\bar{\phi}$). The share of the outside good s_0 is very close to 1, as a monthly frequency is high compared to vehicle lifetime: most people do not buy a car in any given month and monthly sales are small compared to the market size. Thus, the second term involving $\bar{\phi}(1 - s_0)$ is negligible.

The easier purchasers substitute between fuel-type versions of the same model, resp. between models within a segment, the higher is σ_1 , resp. σ_2 , and, thus, the higher is ρ_1 , resp. ρ_2 . Intuitively speaking, a higher correlation of preference for similar products (same nests) leads to a relatively higher weight put onto the comparison with these similar products.

Obviously, diesel taxes affect cars differently depending on their fuel-type. Using our main model defined in Equation (4), the elasticity $\eta_{sjf}^{t_D}$ of demand for a given car sjf with respect to an increase in diesel tax (holding gasoline tax constant) can be computed as:

$$\begin{aligned} \eta_{sjf}^{t_D} &= \frac{\partial s_{sjf} / s_{sjf}}{\partial t_D / t_D} \\ &= \beta(1 + t^{VAT}) t_D \left(\rho_1 (\mathbb{1}_{f=diesel} \phi_{sjf} + (\rho_2 - \rho_1) \pi_j^D \bar{\phi}_j - (\rho_2 - 1) \pi_s^D \bar{\phi}_s) - \right. \\ &\quad \left. \beta(1 + t^{VAT}) t_D \bar{\phi}^D \pi^D (1 - s_0) \right) \\ &\approx \beta(1 + t^{VAT}) t_D \left(\rho_1 (\mathbb{1}_{f=diesel} \phi_{sjf} - \pi_j^D \bar{\phi}_j) + \rho_2 (\pi_j^D \bar{\phi}_j - \pi_s^D \bar{\phi}_s) + \pi_s^D \bar{\phi}_s \right). \end{aligned} \quad (10)$$

where the indicator $\mathbb{1}_{f=diesel}$ takes the value 1 if the vehicle sjf is running on a diesel engine, π_j^D is the share of diesel in sales of model j , π_s^D is the share of diesel in sales of segment s , and π^D is the overall market share of new diesel cars (among purchases). $\bar{\phi}^D$ is the mean fuel consumption of new diesel cars (sales-weighted average). Again, $(1 - s_0)$ is very close to zero and this elasticity can be closely approximated by the first part of the equation.

Intuitively, an increase in the diesel tax rate has a direct negative impact for all diesel cars. However, this effect may be reduced if its substitutes are also impacted by this increase. The effect for gasoline cars of a diesel tax is expected to be positive.

On a more aggregate level, we examine the impact of an increase in fuel prices on the composition of the automobile fleet, with a particular focus on the amount of diesel cars purchased. More specifically, we evaluate the elasticity of the share of diesel cars among new purchases π^D . Assuming again that an international oil price shift equally affects both gasoline and diesel pre-tax prices, such a price shift would change the share

³⁵This is specific to the French form of petrol tax: as the fuel-type specific taxes are of a lump-sum form, they do not play a role here. The t^{VAT} is the same for both fuel-types.

of diesel cars by η_D . In the simple logit demand, this change can be computed as:

$$\begin{aligned}
 \eta_D &= \frac{\partial \pi^D / \pi^D}{\partial p^e / p^e} \\
 &= \frac{\sum_{s,j,f} \mathbb{1}_{f=diesel} s_{sjf} \eta_{sjf}}{\sum_{s,j,f} \mathbb{1}_{f=diesel} s_{sjf}} - \frac{\partial(1-s_0)}{\partial p^e} \frac{p^e}{1-s_0} \\
 &= \beta(1+t^{VAT}) p_e \left(\rho_1 (\bar{\phi}^D - \bar{\phi}_j) + \rho_2 (\bar{\phi}_j - \bar{\phi}_s) + \bar{\phi}_s - \bar{\phi} \right), \\
 &= \frac{\beta(1+t^{VAT}) p_e}{\pi^D(1-s_0)} \sum_{s,j} s_j \left(\underbrace{\rho_1 \pi_j^D (\phi_j^D - \bar{\phi}_j)}_{S_1} + \underbrace{\rho_2 (\pi_j^D - \pi_s^D) \bar{\phi}_j}_{S_2} + \underbrace{(\pi_s^D - \pi^D) \bar{\phi}_s}_{S_3} \right), \quad (11)
 \end{aligned}$$

which involves weighted averages of fuel consumption, where the weights are given by the share of diesel sales.³⁶ $\bar{\phi}_j = \sum_{s,j} \frac{\pi_j^D s_j}{\pi^D(1-s_0)} \bar{\phi}_j$ is the average fuel consumption weighted by the share of diesel per model, whereas $\bar{\phi}_s = \sum_s \frac{\pi_s^D s_s}{\pi^D(1-s_0)} \bar{\phi}_s$ is the average weighted by the diesel share per segment. ϕ_j^D is the fuel consumption of the diesel version of model j . π_j^D , resp. π_s^D , is the share of diesel among purchases of model j , resp. of segment s .

The interpretation of this equation is not straightforward. In the simplest logit case ($\sigma_1 = \sigma_2 = 0$), $\eta_D = \beta(1+t^{VAT}) p_e (\bar{\phi}^D - \bar{\phi})$. Naturally, η_D depends on the average fuel consumption of diesel cars relative to the overall average fuel consumption. $\bar{\phi}^D - \bar{\phi}$ is always negative because diesel cars are more fuel-efficient. β is negative as well, so that η_D is positive: if fuel prices increase, purchasers substitute to more fuel-efficient diesel cars and their share among purchases increases.

In a nested setup, the effect is less straightforward, but we still expect a positive sign. Indeed, the first term S_1 in Equation (11) involves the difference between diesel fuel consumption and average fuel consumption; again, this change is expected to be negative as diesel engines tend to be more fuel-efficient. However, we do not have such an unambiguous relation for the two other terms S_2 and S_3 .³⁷ Both ρ_1 and ρ_2 are positive and larger than one. In practice ρ_2 is smaller than ρ_1 , so that η_D is most strongly impacted by the first element of the parenthesis, which is likely to be positive.

Similarly, the elasticity of the share of diesel cars π^D to a change in fuel taxes (holding gasoline taxes constant) $\eta_D^{t_D}$ may be written:

$$\begin{aligned}
 \eta_D^{t_D} &= \frac{\partial \pi^D / \pi^D}{\partial t_D / t_D} \\
 &= \beta(1+t^{VAT}) p_e \left(\rho_1 (\bar{\phi}^D - \bar{\pi}_j^D \bar{\phi}_j) + \rho_2 (\bar{\pi}_j^D \bar{\phi}_j - \bar{\pi}_s^D \bar{\phi}_s) + \bar{\pi}_s^D \bar{\phi}_s - \bar{\phi} \right). \quad (12)
 \end{aligned}$$

³⁶With any variable A we denote $\tilde{A} = \sum_{s,j,f} \frac{s_{sjf}}{\pi^D(1-s_0)} A_{sjf} \mathbb{1}_{f=diesel}$ this variable weighted by the share of the diesel version amongst all diesel cars (for example, $\tilde{\phi}_{sjf}$ corresponds to the average fuel consumption of diesel cars $\bar{\phi}^D$).

³⁷The last term for example does not have a well defined sign. For example in the case of only two segments in proportion s_1 and $(1-s_1)$, this term is proportional to $s_1(1-s_1)(\pi_{s_1}^D - \pi_{s_2}^D)(\bar{\phi}_{s_1} - \bar{\phi}_{s_2})$. One cannot exclude that this term is positive, for example if cars have a much higher fuel consumption on average in the segment with the higher share of diesel cars.

This elasticity η_D^{tD} depends only on the fuel consumption of diesel cars and on their relative share among purchases: the lower their fuel consumption, the smaller the impact of a diesel tax increase.

Finally, we can also compute the elasticity η_ϕ (respectively η_{CO_2}) of the average fuel consumption (respectively of average CO₂ intensity) of new cars with respect to fuel prices p^e and to fuel taxes.

$$\begin{aligned}\eta_\phi &= \frac{\partial \bar{\phi} / \bar{\phi}}{\partial p^e / p^e} \\ &= \beta(1 + t^{VAT}) \frac{p^e}{(1 - s_0)\bar{\phi}} \sum_{j,s,f} (\phi_{sjf} s_{sjf} (\rho_1(\phi_{sjf} - \bar{\phi}_j) + \rho_2(\bar{\phi}_j - \bar{\phi}_s) + \bar{\phi}_s - \bar{\phi})) .\end{aligned}\quad (13)$$

For example, in the simple logit demand model, η_ϕ simplifies to:

$$\eta_\phi = \beta(1 + t^{VAT}) p^e \left(\frac{\overline{\phi^2} - \bar{\phi}^2}{\bar{\phi}} \right), \quad (14)$$

with $\overline{\phi^2}$ is the mean of squared fuel consumption of new vehicles. The impact of an oil price shock on average fuel consumption depends thus on the ratio of the variance and the mean of fuel consumption. Both the variance and the mean of ϕ are always positive, so that η_ϕ is always negative in the simple logit case: when fuel prices increase, we expect to find that average fuel consumption is reduced. In the more realistic nested logit demand model, the conclusion is less straightforward. Again, we have some intuition for the first term of Equation (13) which is of first order in the sum: it can be simplified rewritten as $\beta \rho_1 \sum_{s,j} \pi_j^D (1 - \pi_j^D) s_j (\phi_j^D - \phi_j^G)^2$ and is thus expected to be negative.

The elasticity of average fuel consumption η_ϕ^{tD} (respectively $\eta_{CO_2}^{tD}$) to a change in diesel tax (holding gasoline tax constant) can be written in case of a simple logit demand model:

$$\begin{aligned}\eta_\phi^{tD} &= \frac{\partial \bar{\phi} / \bar{\phi}}{\partial t_D / t_D} \\ &= \beta t_D (1 + t^{VAT}) \underbrace{\frac{\beta \pi^D}{\bar{\phi}}}_{<0} \left(\underbrace{\overline{\phi_D^2} - \bar{\phi}_D^2}_{>0} + (1 - \pi^D) \bar{\phi}_D \underbrace{(\bar{\phi}_D - \bar{\phi}_G)}_{<0} \right).\end{aligned}\quad (15)$$

This elasticity depends on the fuel consumption of diesel cars and on their relative share among purchases compared with the average fuel consumption. The sign is not clear-cut. An increase in the diesel tax can reduce the share of diesel cars, which are more fuel-efficient. The higher the gap between the average fuel consumption of gasoline and diesel cars, the higher the increase in the average fuel emissions of new cars. This effect may be partially offset by the dispersion in fuel emissions of diesel cars, as we expect that an increase in diesel prices has more impact on less fuel-efficient cars. Overall, we expect that a rise in diesel tax increases the average fuel emissions of new cars if diesel cars are much more fuel-efficient than gasoline cars and that the diesel share is not too high.

C Complementary results for the main specification

C.1 Raw coefficients

Tables for estimated coefficients are not directly interpretable. This is why the body of this chapter concentrates on elasticities and counterfactual policy impacts. The coefficients β_d measure each demographic group's direct sensitivity to fuel prices. As expected, β_d is statistically significant for most demographic groups and is always negative when significantly different from zero: as fuel prices increase, the utility from any given car decreases (Table II.8).

We find substantial heterogeneity in the relative magnitude of β_d across purchaser types. The heterogeneity in this sensitivity parameter depends on three main factors: first, the flexibility of the consumer's car usage (if he can adjust his car mileage, the fuel efficiency becomes less important for his purchasing decision); second, whether the consumer buys fuel-efficient cars no matter what (there might not be much of a margin to react on for some consumers); and finally, the consumer's income and preferences for other characteristics of the car.

Table II.8: Estimates for the coefficient on cost per km β_d

<i>Private consumers</i>						
Income	Not employed		Young professional		Employed (>30)	
	Low	High	Low	High	Low	High
Urban	−0.11*** (0.02)	−0.08*** (0.02)	−0.15*** (0.02)	−0.13*** (0.02)	−0.13*** (0.02)	−0.14*** (0.01)
Suburb./rural	−0.08*** (0.02)	−0.11*** (0.02)	−0.10*** (0.02)	−0.15*** (0.02)	−0.10*** (0.02)	−0.15*** (0.01)
Paris urban		−0.10*** (0.02)		−0.09*** (0.02)		−0.10*** (0.01)
Paris suburban		−0.03 (0.02)		−0.08*** (0.02)		−0.10*** (0.01)
<i>Firm purchases</i>						
Sector	Agriculture & industry	Car rental		Trade & services		
Suburban/rural	−0.01 (0.01)	−0.03 (0.04)		−0.06*** (0.01)		
Urban	−0.09*** (0.02)	−0.16*** (0.03)		−0.10*** (0.01)		
Paris urban	−0.07*** (0.02)	0.08*** (0.02)		−0.01 (0.01)		
Paris suburban	−0.01 (0.02)	0.01 (−)		−0.04 (0.02)		

Source: CCFA, authors' calculations. Bootstrap standard errors in parentheses. Equation (4) is estimated by GMM separately for each type of purchasers. Other controlling variables include horsepower, brand fixed effects, segment fixed effects, class of CO₂, month-year effects, and price. Instrumental variables for prices are the price indices of iron (current and lagged value) and indices of export prices of tires (both interacted with the car's weight), BLP-style instruments and differences of characteristics between gasoline and diesel versions. The estimation of car rental purchases in the Paris suburban area appears to have a problem of weak instruments (see Section 4.2) and does not converge for all bootstrap draws, so that we give no bootstrap error term for it.

Among private consumers, the effect of fuel price increases is stronger for employed consumers (Table II.8). Working people have to drive more and travel distances cannot

be easily reduced; they are thus expected to be the more responsive to fuel price changes. This effect is less strong in the Paris region, where more public transport alternatives are available.

Generally, firms react less strongly to fuel prices than private consumers. Among other factors this may be due to firms' ability to pass through fuel costs to the consumer and to smaller absolute fuel price variations when VAT refund is taken into account. Within firms, we see considerable heterogeneity (Table II.8). The most responsive firms are in urban areas except Paris. In the Paris metropolitan region, sensitivity is particularly low and almost never significant.

However, because of the nested logit specification, the magnitude of the parameters is not directly informative on the actual fuel prices elasticities. One has to consider indirect effects due to the correlation (and thus higher potential substitution) between gasoline and diesel versions of the same model captured by σ_{1d} , as well as substitution within segment σ_{2d} . The estimates for these parameters are as expected all between 0 and 1. σ_{1d} is on average 0.5 implying a relatively high correlation between the two fuel-type versions of the same model (Table II.9, while σ_{2d} is relatively low, on average 0.2, implying a relatively low correlation within segments (Table II.10). If the purchaser has a preference for a particular model, he substitutes easily between gas and diesel versions when fuel prices change, rather than switching to a different model and only reluctantly switches segment. Intensity of substitution between the gasoline and diesel versions of the same model appears to be higher in urban areas (including Paris urban and metropolitan areas) than in rural areas. Indeed, while diesel cars yield savings in running costs for long journeys, this advantage is not clear cut for city driving.

The signs of other variables' coefficients are as expected; in particular, the vehicle price impacts utility negatively (Table II.11).

C.2 Demand for selected car models

For a given product, the demand elasticity to fuel prices depends on the car's fuel consumption (relative to competing products) and on the preferences of the consumer types that buy this car (Table II.12). For the sake of illustration, we compute different elasticities η_{jf} implied by the previously presented parameters for some selected cars, as well as the shifts in demand $\Delta^{t_c}\eta_{jf}$ and $\Delta^{t_D}\eta_{jf}$ corresponding to the equalization of diesel and gasoline taxes (t_D) and the carbon tax (t_c), respectively.

An increase in fuel prices (both gasoline and diesel) reduces demand for all cars ($\eta_{jf} < 0$), but the magnitude varies: Table II.12 gives only a sample of the most popular cars in our data, where the Peugeot 307 gasoline model had an elasticity with respect to fuel price of -0.17, while the Citroen C3 gasoline model had an elasticity of -0.34. An increase in diesel fuel tax strongly lowers the demand for diesel cars ($\Delta^{t_D}\eta_{jf} < 0$); for example the sales of the Audi A6 with diesel engine would decrease by 18.2% (Table II.12). At the same time, such a policy has a small but significantly positive effect on the demand for gasoline cars, reflecting a substitution effect.

Table II.9: Estimates for the coefficient σ_{1d} (substitutability within model, between engine types)

<i>Private consumers</i>						
Income	Not employed		Young professional		Employed (>30)	
	Low	High	Low	High	Low	High
Urban	0.41 *** (0.04)	0.48 *** (0.04)	0.51 *** (0.03)	0.51 *** (0.03)	0.55 *** (0.02)	0.59 *** (0.02)
Suburb./rural	0.45 *** (0.04)	0.41 *** (0.03)	0.38 *** (0.03)	0.41 *** (0.03)	0.55 *** (0.02)	0.52 *** (0.02)
Paris urban		0.30 *** (0.04)		0.62 *** (0.03)		0.62 *** (0.02)
Paris suburban		0.10 (0.06)		0.34 *** (0.04)		0.57 *** (0.03)
<i>Firm purchases</i>						
Sector	Agriculture & industry		Car rental		Trade & services	
Suburban/rural	0.29 *** (0.03)		0.26 *** (0.08)		0.24 *** (0.03)	
Urban	0.33 *** (0.03)		0.18 *** (0.04)		0.23 *** (0.03)	
Paris urban	0.17 *** (0.04)		−0.16 *** (0.04)		0.18 *** (0.03)	
Paris suburban	0.77 *** (0.05)		0.42 (−)		0.60 *** (0.05)	

Source: CCFA, authors' calculations. Bootstrap standard errors in parentheses. Equation (4) is estimated by GMM separately for each type of purchasers. Other controlling variables include horsepower, brand fixed effects, segment fixed effects, class of CO₂, month-year effects, and price. Instrumental variables for prices are the price indices of iron (current and lagged value) and indices of export prices of tires (both interacted with the car's weight), BLP-style instruments and differences of characteristics between gasoline and diesel versions. The estimation of car rental purchases in the Paris suburban area appears to have a problem of weak instruments (see Section 4.2) and does not converge for all bootstrap draws, so that we give no bootstrap error term for it.

Table II.10: Estimates for the coefficient σ_{2d} (substitutability within segment, between models)

<i>Private consumers</i>						
Income	Not employed		Young professional		Employed (>30)	
	Low	High	Low	High	Low	High
Urban	0.11 *** (0.02)	0.13 *** (0.02)	0.22 *** (0.02)	0.19 *** (0.02)	0.32 *** (0.01)	0.39 *** (0.01)
Suburb./rural	0.14 *** (0.02)	0.16 *** (0.02)	0.23 *** (0.01)	0.21 *** (0.01)	0.28 *** (0.02)	0.34 *** (0.01)
Paris urban		0.17 *** (0.02)		0.26 *** (0.02)		0.37 *** (0.02)
Paris suburban		0.21 *** (0.02)		0.20 *** (0.02)		0.30 *** (0.02)
<i>Firm purchases</i>						
Sector	Agriculture & industry		Car rental		Trade & services	
Suburban/rural	0.08 *** (0.02)		0.16 *** (0.03)		0.01 (0.02)	
Urban	0.07 *** (0.02)		0.08 *** (0.03)		0.16 *** (0.02)	
Paris urban	0.12 *** (0.03)		0.10 *** (0.02)		0.24 *** (0.02)	
Paris suburban	0.28 *** (0.03)		0.22 (-)		0.32 *** (0.03)	

Source: CCFA, authors' calculations. Bootstrap standard errors in parentheses. Equation (4) is estimated by GMM separately for each type of purchasers. Other controlling variables include horsepower, brand fixed effects, segment fixed effects, class of CO₂, month-year effects, and price. Instrumental variables for prices are the price indices of iron (current and lagged value) and indices of export prices of tires (both interacted with the car's weight), BLP-style instruments and differences of characteristics between gasoline and diesel versions. The estimation of car rental purchases in the Paris suburban area appears to have a problem of weak instruments (see Section 4.2) and does not converge for all bootstrap draws, so that we give no bootstrap error term for it.

Table II.11: Estimates for the coefficient on vehicle price γ_d

<i>Private consumers</i>						
Income	Not employed		Young professional		Employed (>30)	
	Low	High	Low	High	Low	High
Urban	−0.63*** (0.05)	−0.57*** (0.05)	−0.30*** (0.04)	−0.31*** (0.04)	−0.21*** (0.03)	−0.12*** (0.03)
Suburb./rural	−0.65*** (0.05)	−0.66*** (0.05)	−0.42*** (0.04)	−0.30*** (0.04)	−0.36*** (0.03)	−0.15*** (0.03)
Paris urban		−0.36*** (0.05)		−0.32*** (0.04)		−0.21*** (0.03)
Paris suburban		−0.20*** (0.05)		−0.25*** (0.04)		−0.14*** (0.03)
<i>Firm purchases</i>						
Sector	Agriculture & industry	Car rental		Trade & services		
Suburban/rural	−0.22*** (0.03)	−0.29*** (0.08)		−0.10*** (0.03)		
Urban	−0.01 (0.03)	0.14*** (0.05)		−0.00 (0.03)		
Paris urban	−0.01 (0.03)	−0.03 (0.04)		−0.09*** (0.03)		
Paris suburban	−0.14*** (0.03)	−0.28 (−)		−0.27*** (0.05)		

Source: CCFA, authors' calculations. Bootstrap standard errors in parentheses. Equation (4) is estimated by GMM separately for each type of purchasers. Other controlling variables include horsepower, brand fixed effects, segment fixed effects, class of CO₂, month-year effects, and price. Instrumental variables for prices are the price indices of iron (current and lagged value) and indices of export prices of tires (both interacted with the car's weight), BLP-style instruments and differences of characteristics between gasoline and diesel versions. The estimation of car rental purchases in the Paris suburban area appears to have a problem of weak instruments (see Section 4.2) and does not converge for all bootstrap draws, so that we give no bootstrap error term for it.

Table II.12: Demand elasticity for selected models with respect to fuel prices

model (segment)	fuel	CO ₂ (g/km)	fuel cons. (L/km)	η_{jf}	$\Delta^{td} \eta_{jf}$ (%)	$\Delta^{tc} \eta_{jf}$ (%)
Audi A6 (sedan)	gasoline	236.9	10.2	−0.22*** (0.03)	1.17*** (0.22)	−6.73*** (0.89)
Audi A6 (sedan)	diesel	200.1	7.6	−0.29*** (0.02)	−18.20*** (1.55)	−9.39*** (0.60)
Citroen C3	gasoline	147.8	6.4	−0.34*** (0.02)	2.46*** (0.23)	−10.62*** (0.51)
Citroen C3	diesel	112.8	4.3	−0.19*** (0.01)	−13.48*** (0.69)	−6.55*** (0.32)
Peugeot 307 (sport)	gasoline	192.7	8.3	−0.17*** (0.01)	1.57*** (0.08)	−4.29*** (0.21)
Peugeot 307 (sport)	diesel	159.0	6.0	−0.32*** (0.01)	−18.62*** (0.87)	−9.41*** (0.43)
Renault Twingo (compact)	gasoline	137.0	5.9	−0.32*** (0.01)	0.86*** (0.03)	−9.78*** (0.44)
Renault Twingo (compact)	diesel	113.0	4.3	−0.25*** (0.01)	−15.62*** (0.93)	−7.30*** (0.37)

Source: CCFA, authors' calculations. Equation (4) is estimated by GMM separately for each type of consumers. Standard errors in parentheses are estimated by bootstrap (500 replications).

D Testing for weak instruments

Table II.13: Conditional F-values of the weak instrument test – instruments for the *price*

	<i>Private consumers</i>					
	Not employed		Young employed (<30)		Employed (>30)	
Income	Low	High	Low	High	Low	High
Urban	35.1***	31.9***	51.8***	51.7***	47.8***	49.0***
Suburban/rural	31.7***	37.6***	64.8***	70.9***	51.3***	51.5***
Paris urban	20.4***		42.2***		44.2***	
Paris suburban	16.3**		36.6***		39.4***	
	<i>Firm purchases</i>					
	Industry & Agriculture		Car rental		Trade & services	
Urban	42.2***		11.5**		39.7***	
Suburban/rural	45.9***		34.9***		37.2***	
Paris urban	52.4***		36.3***		34.6***	
Paris suburban	14.2**		14.1**		15.4**	

Note: Stars denote conditional F-values beyond the critical value (at 5% significance level) for different levels of maximal bias of the IV estimator relative to OLS; *** stands for a maximal bias of 5%, ** for 10%, * for 20%.

Table II.14: Conditional F-values of the weak instrument test – instruments for the *market share of the model within its segment* $s_{dj|s}$

	<i>Private consumers</i>					
	Not employed		Young employed (<30)		Employed (>30)	
Income	Low	High	Low	High	Low	High
Urban	68.1***	71.3***	61.4***	62.4***	60.7***	53.6***
Suburban/rural	71.5***	73.3***	71.1***	64.0***	58.9***	55.7***
Paris urban	53.4***		58.3***		49.3***	
Paris suburban	36.5***		53.5***		54.8***	
	<i>Firm purchases</i>					
	Industry & Agriculture		Car rental		Trade & services	
Urban	45.9***		34.9***		37.2***	
Suburban/rural	42.2***		11.5**		39.7***	
Paris urban	52.4***		36.3***		34.6***	
Paris suburban	14.2**		14.1**		15.4**	

Note: Stars denote conditional F-values beyond the critical value (at 5% significance level) for different levels of maximal bias of the IV estimator relative to OLS; *** stands for a maximal bias of 5%, ** for 10%, * for 20%.

Table II.15: Conditional F-values of the weak instrument test – instruments for the *market share of a fuel-type within its model nest* $s_{df|j}$

	<i>Private consumers</i>					
	Not employed		Young employed (<30)		Employed (>30)	
Income	Low	High	Low	High	Low	High
Urban	26.5***	22.6***	32.6***	32.3***	41.9***	44.7***
Suburban/rural	23.6***	27.9***	31.7***	38.1***	43.7***	44.3***
Paris urban	15.0**		27.3***		31.7***	
Paris suburban	16.7**		22.2***		26.5***	
	<i>Firm purchases</i>					
	Industry & Agriculture		Car rental		Trade & services	
Urban	24.2***		21.6***		25.8***	
Suburban/rural	32.4***		6.3*		28.6***	
Paris urban	15.2**		21.0***		20.8***	
Paris suburban	11.7**		2.9		10.4*	

Note: Stars denote conditional F-values beyond the critical value (at 5% significance level) for different levels of maximal bias of the IV estimator relative to OLS; *** stands for a maximal bias of 5%, ** for 10%, * for 20%.

E Robustness checks: elasticities

Table II.16: Robustness checks: elasticities with respect to fuel prices of diesel share, average fleet fuel consumption (L/km) and CO₂ intensity (g/km)

	Diesel share	Fuel cons.	CO ₂
	η_D	η_ϕ	η_{CO_2}
<i>Main specification - including degenerate nests (gas- or diesel-only models)</i>			
Households	0.044*** (0.003)	-0.015*** (0.001)	-0.018*** (0.001)
Firms	0.017*** (0.003)	-0.004*** (0.001)	-0.006*** (0.001)
Total	0.045*** (0.002)	-0.011*** (0.001)	-0.015*** (0.001)
<i>Alternative specification - Nests (segment > model)</i>			
Households	0.042*** (0.003)	-0.014*** (0.001)	-0.017*** (0.001)
Firms	0.015*** (0.004)	-0.004*** (0.001)	-0.006*** (0.001)
Total	0.044*** (0.003)	-0.010*** (0.001)	-0.014*** (0.001)
<i>Main specification - BLP-instruments only</i>			
Households	0.033*** (0.003)	-0.015*** (0.001)	-0.017*** (0.001)
Firms	0.017*** (0.004)	-0.003*** (0.001)	-0.004*** (0.001)
Total	0.039*** (0.003)	-0.011*** (0.001)	-0.014*** (0.001)
<i>Main specification - without purchaser heterogeneity</i>			
Total	0.039*** (0.004)	-0.028*** (0.003)	-0.025*** (0.002)

Source: CCFA, authors calculations. Estimates rely on the parameters of Equation (4) estimated by GMM separately for each type of consumers. Standard errors in parentheses are estimated by bootstrap (500 replications).

III

Competition between for-profit and industry labels: the case of social labels in the coffee market

I pity the man who wants a coat so cheap that the man or woman who produces the cloth will starve in the process.

Benjamin Harrison
on the importance of fair trade.

1 Introduction

Over the past decades, consumers have become more and more interested in the social and environmental impact of their consumption. However, most sustainability aspects of a product are difficult for consumers to verify, even after purchase, meaning that the promise of a responsible production process is essentially a *credence attribute* that cannot be verified either before or after purchase. Firms increasingly use voluntary third-party labels to solve their credibility problem.

The coffee market has a particularly large number of well-established sustainability labels; the most important being Fairtrade, Rainforest Alliance, and UTZ Certified. These target the well-being of farmers and the environmental impact of production. The stringency of the labels varies: Fairtrade, for example, guarantees a price premium for farmers, while the price premia established by UTZ Certified and Rainforest Alliance are lower and not guaranteed.¹ When it comes to social sustainability labels, higher farmgate prices are seen by consumers as higher quality and justify higher prices.

When each firm can offer several differentiated products, various constellations of product lines can arise. In the coffee example, an international comparison illustrates this multitude of possible product line constellations: in Germany, most roasters² offer a range of products including conventional, i.e. not labeled, and labeled coffee of several labels (head-to-head competition). In other countries, such as Finland,³ coffee roasters have specialized so that each label is only offered by one roaster (market segmentation).

This paper establishes a model of label competition, between a for-profit label and an industry standard. To start with, we model the firms' choice of a third-party label offered by a for-profit licensor in the first period. We are interested in the interaction between the licensor, which sets a license fee and a label quality, and firms, which decide on their product line and their prices. Each firm can offer several goods that are differentiated both horizontally between firms and vertically through quality. In a second period, we allow firms to establish its own labeling organization – an industry standard – that maximizes joint firm profit. We then analyze how the industry standard sets its quality and what product lines are offered in equilibrium.

In both periods, we find that equilibrium product lines depend crucially on the degree of (exogenous) horizontal differentiation: the market is segmented if horizontal

¹Fairtrade Labeling Organizations International (FLO) guarantee a price premium at the farmgate of \$0.20/lb (since 2011) over the stock market price. UTZ Certified in 2012 reported sales prices that result in an average premium of \$0.04/lb over the price index of the International Coffee Organization; the prices for Rainforest Alliance are not known but they reported a premium of \$0.11/lb in 2009 (Potts et al. 2014). In 2012, Fairtrade and Rainforest Alliance had similar market shares of 2-3% worldwide while UTZ had almost twice as much, with much larger market shares in countries like the United States, Germany, and Great Britain.

²The German coffee market is dominated by JDE/Mondelez, Aldi, Tchibo, Melitta and Dallmayr; together they hold 90% of the market (Villas-Boas 2007, adjusted for the merger of JDE/Mondelez in 2015).

³In Finland, per capita coffee consumption is the highest in the world. The average Finn consumes 9-10 kg of roasted coffee annually; approximately four cups per day (Valkila et al. 2010). There are just two major companies on the Finnish coffee market: Meira and Paulig.

differentiation is weak, i.e. each label is offered by one firm only. In contrast, firms are in head-to-head competition when horizontal differentiation is strong, that is both firms offer all available labels. When there are two labels and horizontal differentiation is intermediate, the industry standard strategically distorts its quality downwards in order to induce a segmented product line.

Overall, we illustrate why an industry facing a third-party label has an interest in establishing its own industry standard: the presence of a second label reduces the fees set by the for-profit licensor, and an additional vertically differentiated good increases product lines, thereby increasing overall demand. Moreover, for intermediate levels of horizontal differentiation, the industry standard strategically reduces competition by reducing product line overlap, thereby increasing mark-ups.

We further ask whether regulation in form of a minimum quality requirement for labels, such as established in organic farming, can increase welfare. In the first period with one label, a minimum quality requirement increases the label's standard, thereby increasing welfare. Welfare increases if firms are in head-to-head competition, but the minimum quality requirement cannot affect the equilibrium product line. In the second period with two labels, the social planner can set its minimum quality requirement such that it prevents the industry standard's strategic downward distortion, thereby maximizing the number of labeled products. Whenever the industry standard does not strategically distort its quality downwards, the social planner aims at setting *lower* qualities than the industry standard. In these cases, a minimum quality requirement does not bind and does not impact welfare: the duopoly firms in equilibrium differentiate *too much* from conventional market and too little from the higher label.

In the remainder of this paper, we begin by discussing the relevant literature and explain the context of the coffee market and fairtrade research. We then explain the model and each player's objectives in Section 2. We first solve the first period with only the for-profit licensor in Section 3. Then, we solve the model in the second period upon entry of an industry standard in Section 4. For each period, we explore whether there is scope for a government-imposed minimum quality requirement. Finally, we conclude in Section 5.

1.1 Related literature

Our model features both vertical differentiation between labels and horizontal differentiation between firms. Methodologically, this study relies on a large literature using the nested logit model established by McFadden (1978). In particular, the version of Anderson and De Palma (1992) with multi-product firms allows us to explicitly model the endogenous substitution elasticity between labels depending on label differentiation. Gallego and Wang (2014) use such a nested logit to account for horizontal and vertical differentiation.

Von Schlippenbach and Teichmann (2012) and Yu and Bouamra-Mechemache (2016) model how standards are used by different agents (retailers, resp. manufacturers) to strengthen their bargaining power within the vertical supply chain. The choice of firms

in duopoly adopting a labeled product line also relates to product line rivalry (e.g. Avenel and Caprice 2006). Cheng and Peng (2012) show the importance of strategic effects in quality setting when a firm can offer more than one vertically differentiated product.

A growing literature is studying voluntary third-party certification, for a review see Bonroy and Constantatos (2015). In particular, newer papers study the interaction between several labeling organizations and firms, focusing on endogenous quality levels. Fischer and Lyon (2014) model the rivalry between an ecolabel set by an NGO and an industry-standard in the forestry sector and find that the industry-standard lowers environmental benefits even if consumers are perfectly informed. Poret (2016) models the competition between two NGOs setting labels with different objectives. Similarly to this study, Bottega et al. (2009) study the interaction between a regulator, an industry standard and a for-profit licensor. However, all these studies consider simple market constellations (monopolist/single-good duopoly), following in particular the model by Heyes and Maxwell (2004). Finally, a strand of literature explores the effect of consumer confusion when several labels coexist or monitoring is imperfect (Harbaugh et al. 2011, Mahenc 2010, Mason 2011), whereas we assume that consumers observe label quality perfectly.

1.2 Coffee market and fairtrade

In our model, the incumbent labeling organization maximizes its profit. Previous theoretical research on fairtrade has modeled an NGO label maximizing farmer welfare (Podhorsky 2015, Richardson and Stähler 2014, Chambolle and Poret 2013). However, it is difficult to argue that the FLO price policy is aimed at maximizing farmer welfare. A concise theoretical model by Janvry et al. (2015) shows how farmer rents are eroded by unlimited entry of farmers, such that in equilibrium all the price premium goes to the licensor in form of the farmer annual fee. Crucially, fairtrade guarantees prices, but not sales, such that fairtrade-labeled farmers typically sell large proportions of their production as conventional coffee, i.e. without the label at world-market prices (e.g. Valkila and Nygren 2009, Panhuysen and Pierrot 2014).⁴ Moreover, annual license fees are high for both roasting companies and, in particular, for farmers, which contrasts with the idea that an NGO maximizes label participation.⁵

This paper concentrates on the impact of labels in the consumer country, excluding the farmer from the picture: we interpret fairtrade as a quality label. Fairtrade coffee is an amply available commodity and farmers have no market power. Johannessen and Wilhite (2010) estimate that about 75% of value added in fairtrade coffee remain in the consumer

⁴Panhuysen and Pierrot (2014) show that about a quarter of certified coffee production is sold with a label.

⁵Under standard assumptions, an NGO label maximizes access to its label and sets its fee as low as possible, that is equal to the cost of monitoring (cf. Bottega and De Freitas 2009), which is normalized to zero in our model. If the cost of the label is zero, then our model predicts that it is always an equilibrium for both firms to offer the label. Only in markets with very weak horizontal differentiation, market segmentation might be an additional equilibrium. However, this does not reflect the reality of coffee markets.

country. Empirical evidence suggests that farmers receive a higher price for fairtrade coffee than for conventional coffee (Beuchelt and Zeller 2011, Dragusanu and Nunn 2014, Arnould et al. 2009), but the impact on income is small at best when controlling for selection into the labeling scheme (Ruben and Fort 2012, Saenz Segura and Zuniga-Arias 2008, Beuchelt and Zeller 2011). Dragusanu et al. (2014) review this literature in more detail.

Nevertheless, marketing and experimental research has consistently shown that consumers have a positive willingness-to-pay for fairtrade products (e.g. Basu and Hicks 2008, Pelsmacker et al. 2005, Loureiro and Lotade 2005). A rational consumer understands that it is welfare-enhancing for a farmer to sell more fairtrade coffee, once he has incurred the fixed entry costs of labeling. Moreover, Friedrichsen and Engelmann (2017) and Teyssier et al. (2014) show that social image concerns play a role, so that consumers enjoy being seen buying fairtrade products. Another possible explanation of the widespread support of the fairtrade system is that consumers are not aware of the dynamic effects of the fairtrade system leading to an excessively large number of certified farmers. We assume that consumers derive a homogeneous positive utility from higher coffee prices at the farmer level, leaving aside the debate whether these preferences are due to social image, *warm glow* (Andreoni 1989), or pure altruism.

2 Model

We analyze a game with two periods, each consisting of several stages. The game involves two labeling organizations $s = F, I$, two horizontally differentiated firms $i = 1, 2$, and homogeneous consumers which value quality positively. Firms can offer several vertically differentiated products: they always supply a product of conventional market quality q^C and can additionally opt for one or both labels. We assume that firms cannot credibly offer qualities higher than conventional market quality $q^C = 0$ without getting labeled by a labeling organization.⁶ The labeling organizations decide on qualities q^F and q^I , guaranteed by their respective label. The for-profit licensor moreover sets a license fee L . Subsection 2.4 provides a detailed overview of the game sequence.

2.1 Consumer demand

To capture both horizontal and vertical product differentiation, we specify consumer demand using a nested logit model (cf. McFadden 1978, Anderson and De Palma 1992). In our model, products become closer substitutes when their qualities become more similar. This section derives the demand equations in the case where both firms offer both labels. The firms' market shares and demand functions for other product line constellations can be derived analogously.

⁶The certification and labeling process is assumed to be credible and to guarantee that labeled products fulfill the quality requirements defined by the licensors. We further assume that consumers are perfectly informed about the qualities chosen by the licensors.

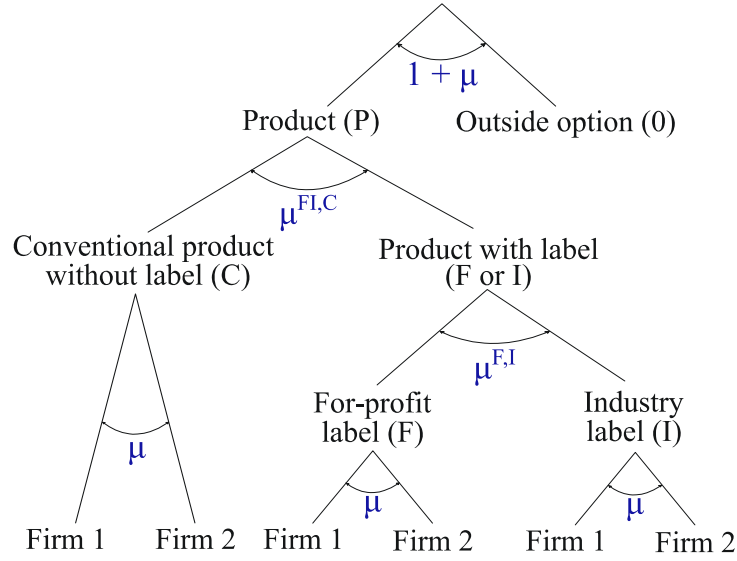


Figure III.1: Nested decision-making structure of consumers

Assume that each firm offers three products with qualities q^F , q^I , and q^C , then Figure III.1 shows the decision structure of consumers. Each of the homogeneous consumers buys one unit or opts for the outside good. Consumers decide *if* they want to buy any product (decision between nest P , for product, and nest 0, for outside option). If consumers choose nest P , they decide between products with and without labels (decision between nests FI and C). Within nest FI consumers choose between labels (between nests F and I). Finally, within each nest s with $s = F, I, C$ consumers decide from which firm they buy. Figure III.1 illustrates this decision structure, where the substitution parameters $\mu^{FI,C}$, $\mu^{F,I}$ and μ are explained below.

Proceeding backwards, consider first consumers' decision within nest s ($s = F, I, C$) between both firms' goods. Each consumer chooses the firm i that maximizes his indirect utility

$$u_i^s = \bar{u} + v(q^s) - p_i^s + \mu \epsilon_i^s, \quad (1)$$

where \bar{u} is the consumer's direct utility of the product, $v(q^s)$ denotes the additional utility from consuming quality q^s and p_i^s the price of firm i 's product with quality q^s . ϵ_i^s is an error term that is distributed with the extreme value distribution. In the example of fairtrade coffee, quality is defined by the farmgate prices guaranteed by the labeling organization. The parameter $\mu > 0$ measures the degree of horizontal differentiation between the two firms such that μ approaching zero translates into perfect competition within the final market. Consumers have a homogeneous valuation of quality $v(q^s)$ which is strictly increasing and strictly concave in q^s :

$$v^s = v(q^s) = \sqrt{\frac{q^s}{1 + q^s}}. \quad (2)$$

Integrating equation (1) over the distribution of the stochastic term ϵ_i^s , as it is standard in

nested logit models,⁷ we obtain firm i 's within-nest market shares $\mathbb{P}_{i|s}$ for nest s

$$\mathbb{P}_{i|s} = \frac{\exp((\bar{u} + v^s - p_i^s)/\mu)}{\exp((\bar{u} + v^s - p_i^s)/\mu) + \exp((\bar{u} + v^s - p_j^s)/\mu)} = \frac{\exp((\bar{u} + v^s - p_i^s)/\mu)}{\exp(A^s/\mu)} \quad (3)$$

$$\text{with } A^s = \mu \ln \left[\exp((\bar{u} + v^s - p_i^s)/\mu) + \exp((\bar{u} + v^s - p_j^s)/\mu) \right] \quad (4)$$

A^s measures the expected utility of nest s (given previous choices at higher nest levels), which is called the inclusive value in nested logit models.

Consider next the choice between F -labeled goods and I -labeled goods. The utility u^s of the nest s for $s = F, I$ is then defined as

$$u^s = A^s + \mu^{F,I} \epsilon^s \text{ for } s = F, I \quad (5)$$

where ϵ^s is a nest-specific error term that is distributed extreme value and substitution between nests F and I is given by

$$\mu^{F,I} = \mu + \frac{v^F - v^I}{1 + v^F - v^I}. \quad (6)$$

The specification of $\mu^{F,I}$ implies that $\mu^{F,I}$ approaches μ if the labels become more similar, i.e. if q^F approaches q^I . As before, integrating over the stochastic term's distribution, we obtain the market shares for nest F (and analogously for nest I):

$$\mathbb{P}_{F|FI} = \frac{\exp(A^F/\mu^{F,I})}{\exp(A^I/\mu^{F,I}) + \exp(A^F/\mu^{F,I})} = \frac{\exp(A^F/\mu^{F,I})}{\exp(A^{FI}/\mu^{F,I})} \quad (7)$$

$$\text{with } A^{FI} = \mu^{F,I} \ln \left[\exp(A^I/\mu^{F,I}) + \exp(A^F/\mu^{F,I}) \right]. \quad (8)$$

Moving upwards, consider now the choice between choosing a labeled product or choosing conventional quality. The utility of the nest FI and of the nest C are defined as

$$u^{FI} = A^{FI} + \mu^{FI,C} \epsilon^{FI} \text{ and } u^C = A^C + \mu^{FI,C} \epsilon^C \quad (9)$$

where $\epsilon^{FI}, \epsilon^C$ is a nest-specific error term distributed with the extreme value distribution and $\mu^{FI,C}$ characterizes the substitution between FI and C . In analogy to equation (6) we use the following functional form

$$\mu^{FI,C} = \mu + \frac{v^F + v^I (v^F - v^I)}{1 + v^F + v^I (v^F - v^I)} \quad (10)$$

Integrating gives the market share of the conventional products, given the consumer

⁷See econometrics textbooks, e.g. Train (2009), for more details on the derivation of market shares in the standard nested logit.

buys any product (nest P)

$$\mathbb{P}_{C|P} = \frac{\exp(A^C/\mu^{FI,C})}{\exp(A^C/\mu^{FI,C}) + \exp(A^{FI}/\mu^{FI,C})} = \frac{\exp(A^C/\mu^{FI,C})}{\exp(A^P/\mu^{FI,C})} \quad (11)$$

$$\text{with } A^P = \mu^{FI,C} \ln \left[\exp(A^C/\mu^{FI,C}) + \exp(A^{FI}/\mu^{FI,C}) \right]. \quad (12)$$

Finally, consider the choice between buying any of the considered goods or the outside good, i.e. a substitute from another product category or nothing. Again, the utility of the nest P is defined as

$$u^P = A^P + \gamma \epsilon^P \quad (13)$$

where ϵ^P is a nest-specific error term distributed with the extreme value distribution and the substitution between the firms' products and an outside good is defined as

$$\gamma = 1 + \mu. \quad (14)$$

Normalizing the outside good's utility to zero, we obtain the probability to buy any product, i.e. the aggregated market share of both firms \mathbb{P}_P :

$$\mathbb{P}_P = \frac{\exp(A^P/\gamma)}{\exp(A^P/\gamma) + 1} \quad (15)$$

$$\text{with } A = \gamma \ln \left[\exp(A^P/\gamma) + 1 \right]. \quad (16)$$

Note that the definitions of the substitution parameters ensure that we always have $0 \leq \mu \leq \mu^{FI,I} \leq \mu^{FI,C} \leq \gamma$ such that goods within a nest are equally or more similar than goods from different nests.⁸ Furthermore, consumers' preferences exhibit love of variety as the inclusive values in all nests increase in the number of products offered.

Summarizing and normalizing the total mass of consumers to 1, demand for firm i 's products can be written as

$$D_i^F = \mathbb{P}_P \mathbb{P}_{FI|P} \mathbb{P}_{F|FI} \mathbb{P}_{i|F}, \quad (17)$$

$$D_i^I = \mathbb{P}_P \mathbb{P}_{FI|P} \mathbb{P}_{I|FI} \mathbb{P}_{i|I}, \quad (18)$$

$$D_i^C = \mathbb{P}_P \mathbb{P}_{C|P} \mathbb{P}_{i|C} \quad (19)$$

2.2 Firms

Firms decide, first, which label to acquire and, second, how to set product prices. Conventional quality q^C can be offered without any certification. Hence, we assume without loss of generality that firms always offer q^C and choose the profit maximizing price for

⁸We adopt the notation from Anderson and De Palma (1992), with substitution parameters at each nest level, which is formally equivalent to the notation more common in econometrics (e.g. Train 2009), where the highest parameter γ is normalized to 1 and substitution parameters σ_k of lower nest levels are defined as $\mu^{FI,C}/\gamma$ and $\mu^{FI,I}/\mu^{FI,C}$ and $\mu^S/\mu^{FI,I}$. Therefore, our restriction on parameters ($0 \leq \mu \leq \mu^{FI,I} \leq \mu^{FI,C} \leq \gamma$) is equivalent to the restriction $\sigma_k \in (0, 1)$ in econometric work.

this quality.⁹

We assume that marginal production costs $c(q^s)$ are equal for both firms, as well as constant and linearly increasing in q^s :

$$c(q^s) = q^s. \quad (20)$$

We define mark-up as difference of price p_i^s and marginal cost q^s . Firm profits are then sum of the demand D_i^s multiplied by the mark-up for each of its products, minus a license fee L if the firm offers label F . As an example, if both firms offer F , I and C , the firm i 's profits are given by $\bar{\Pi}_{i:FIC|FIC}$:

$$\bar{\Pi}_{i:FIC|FIC} = D_i^F(p_i^F - q^F) + D_i^I(p_i^I - q^I) + D_i^C p_i^C - L \quad (21)$$

$$= \Pi_{i:FIC|FIC} - L, \quad (22)$$

where for readability, $\Pi_{i:FIC|FIC}$ is the firm's gross profit before payment of the fee to the licensor.

2.3 Labeling organizations

Both labeling organizations do not face any costs. The first labeling organization is licensor F that maximizes its profit. The licensor's profit Γ is given by the number of firms offering an F -labeled good multiplied by its license fee L . The second labeling organization is an industry standard I that maximizes joint profit of both firms; it does not charge any fees and has no own profit.

Both labeling organizations strategically set the quality of their respective label q^I and q^F . We assume that qualities chosen by the labeling organizations as well as the license fee are public information, without any room for private negotiation.

Licensor F is the established label and is challenged by the industry standard I . In the first period, we model the situation with only for-profit licensor F . The second period is modeled as a Stackelberg game: industry standard I enters and sets q^I taking into account the strategic adjustment of the licensor F 's quality q^F and license fee L .

2.4 Game sequence

We analyze a game with two periods. In the first period $t = 0$, there is only one label, offered by the for-profit licensor. Licensor and firms play the following four stage game with perfect information:

Stage 0.1: Licensor F sets its license fee L_0 and its quality q_0^F ;

Stage 0.2: Firms $i = 1, 2$ choose which label to offer, i.e. decide on their product line;

Stage 0.3: Firms set the consumer prices p_i^s for $s = F, C$;

⁹Stated differently, a firm's decision not to offer quality q^C is equivalent to charging an infinitely high price for this quality, which is never optimal for a firm.

Stage 0.4: Consumers choose their favorite product and buy 1 unit or opt for the outside good.

In the next period $t = 1$, the industry standard I enters the market, so that there is an additional stage 0, followed again by the previous four stages:

Stage 1.0: Industry standard I sets quality q^I ;

Stage 1.1: Licensor F sets license fee L_1 and its quality q_1^F ; the licensor cannot undercut his previous quality: $q_1^F > q_0^F$;

Stage 1.2: Firms $i = 1, 2$ choose which label(s) to offer, i.e. decide on their product line;

Stage 1.3: Firms set the consumer prices p_i^s for $s = F, I, C$;

Stage 1.4: Consumers choose their favorite product and buy 1 unit or opt for the outside good.

Note that we assume that the incumbent for-profit licensor in $t = 1$ cannot decrease its quality q_1^F below its equilibrium monopoly value q_0^{F*} from $t = 0$, without seriously harming its brand image. For simplicity, we further assume that the licensor in the first period does not anticipate the entry of the industry standard in the second period. In the following, we solve the game by backward induction.

3 Market equilibrium with licensor F only

We first look at the first period $t = 0$ before entry of the industry standard, i.e. with only a for-profit licensor F . The game starts with licensor F setting its quality q_0^F and fee L_0 . Both firms can decide to offer an F -labeled product, there are thus three possible market constellations: both firms offer F or one firm offers F or no firm offers F . Since conventional quality $q^C = 0$ can be offered without any certification, we can restrict the analysis to the cases where the product line offered by each firm comprises at least C . Additionally, there can be no equilibrium in which licensor F sells no license; as consumers value quality and variety positively and licensors have no cost, choosing some $q^F > q^C$ and an arbitrarily small but positive license fee L , licensor F can always earn a positive profit.

3.1 Consumer prices

We first compute product price equilibria in stage 0.3 for all product lines and label qualities. Let $y, z \in \{FC, C\}$, we use the notation $\Pi_{i:y|z}$ for a firm i 's profit when it plays y and the other firm plays z . Maximizing firm profit¹⁰ $\Pi_{y|z}$ with respect to prices, we find:

¹⁰We omit the firm index i if no confusion is possible.

Lemma 1 *For all possible product line constellations, there are unique equilibrium prices p_i^s with $s = F, C$ and $i = 1, 2$ in stage 0.3; moreover*

- (i) *when the product lines are symmetric (both firms offer the same qualities), prices are symmetric;*
- (ii) *when firms compete head-to-head $\{FC, FC\}$ (both firms offer all qualities), the symmetric prices are given by the marginal production costs $c(q^s)$ plus a constant mark-up.¹¹*

Proof. See Appendix on page 90. ■

We let $\Pi_{y|z}^*$ denote a firm's reduced profit when it plays y and the other firm plays z , given optimal price setting by both firms.

3.2 Product line decisions

Turning to stage 0.2 of the game and analyzing the firms' product line decisions, we compute the firms' best responses in choosing whether to offer an F -labeled good. Assume firm 1 offers F and C , then firm 2's best response is given by

$$\max\{\Pi_{FC|FC}^* - L_0, \Pi_{C|FC}^*\} \quad (23)$$

Solving the respective maximization problem if firm 1 does not offer F and using symmetry allows us to numerically compute the equilibrium in stage 0.2 of the game.

Figure III.2 illustrates product line equilibria for different values of license fee L_0 and horizontal differentiation μ .¹² The lower the horizontal product differentiation μ , the less profitable it is for both firms to offer the labeled good simultaneously ($\{FC, FC\}$), as fiercer competition reduces their mark-ups. If the license fee L_0 is too high, neither of the firms offers the labeled good ($\{C, C\}$).

3.3 License fee

When deciding on its license fee L_0 , the licensor F has two options: it can aim at selling licenses for its label to both firms or it can decide to sell just one license. Selling to both firms requires a low license fee, whereas selling to only one firm allows for a higher license fee. Maximizing its profits, the licensor sets the fee such that firms are just indifferent, i.e. at the edge of an area in Figure III.2, either from $\{FC, FC\}$ to $\{FC, C\}$, or from $\{FC, C\}$ to $\{C, C\}$.

Assume that the licensor aims at selling its label license to both firms, inducing symmetric, head-to-head competition. The licensor then sets its fee such that both firms prefer offering FC rather than offering C ; the fee equals the deviation profit given the

¹¹Considering a different nest structure Anderson and De Palma (1992) also obtain that equal mark-ups are optimal.

¹²Horizontal differentiation μ is by definition between zero and infinity. However, our figures show only the range until $\mu = 1$ as the results do not change qualitatively for higher values of μ .




Figure III.2: Product line equilibrium in stage 0.2 as a function of license fee L_0 and horizontal differentiation μ (for $q_0^F = 0.23$)

other firm also offers FC :

$$L_0^{sym} = \Pi_{FC|FC}^* - \Pi_{C|FC}^* . \quad (24)$$

We verify numerically, that this license fee L_0^{sym} indeed ensures that both firms want to offer label F :

$$\Pi_{FC|FC}^* - L_0^{sym} = \max\{\Pi_{FC|FC}^* - L_0^{sym}, \Pi_{C|FC}^*\} \text{ for all } \mu \text{ and } q_0^F \quad (25)$$

Assume in contrast that the licensor intends to establish market segmentation $\{FC, C\}$ as an equilibrium in stage 0.2 of the game. The licensor then sets its fee such that the firm offering FC does not want to deviate to offering C only; the fee equals the deviation profit of the firm offering FC given the other firm offers only C :

$$L_0^{seg} = \Pi_{FC|C}^* - \Pi_{C|C}^* . \quad (26)$$

When the licensor sets this fee L_0^{seg} , the other firm could also start offering F , leading again to the symmetric head-to-head equilibrium, but we verify numerically that the potential entrant would always be worse off.¹³ The license fee L_0^{seg} thus ensures that firms play the market segmentation equilibrium in stage 0.2 for all μ and q_0^F .

Summarizing, the licensor effectively chooses the equilibrium played in stage 0.2 by setting its fee. The licensor's preference between both outcomes depends both on (exogenous) horizontal differentiation μ and on (endogenous) vertical differentiation from label quality q_0^F . The licensor induces the product line equilibrium that gives him the highest profit Γ_0 :

$$\Gamma_0 = \max\{2L_0^{sym}, L_0^{seg}\} . \quad (27)$$

Numerically, we find that for strong market differentiation with $\mu > 0.48$, the licensor

¹³We always have $\Pi_{FC|FC}^* - \Pi_{C|FC}^* < \Pi_{FC|C}^* - \Pi_{C|C}^*$: offering F is always more profitable when the other firm does not offer F .

prefers the head-to-head equilibrium for all q_0^F ; for weak market differentiation with $\mu < 0.43$, the licensor always prefers market segmentation. In the relatively small range between these values, the comparison depends on label quality q_0^F .

3.4 Label quality

The licensor profit Γ_0 depends on label quality q_0^F , and each product line equilibrium has different first-order conditions. Using the envelope theorem, we compute the licensor's first-order conditions for optimal label quality q_0^{F*} :

$$\frac{\partial \Gamma_0}{\partial q_0^F} = \begin{cases} \left[\frac{\partial \Pi_{i:FC|FC}}{\partial q_0^F} + \sum_s^{F,C} \frac{\partial \Pi_{i:FC|FC}}{\partial p_j^s} \frac{\partial p_j^s}{\partial q_0^F} \right] - \left[\frac{\partial \Pi_{i:C|FC}}{\partial q_0^F} + \sum_s^{F,C} \frac{\partial \Pi_{i:C|FC}}{\partial p_j^s} \frac{\partial p_j^s}{\partial q_0^F} \right] = 0 \text{ if } \Gamma_0 = 2L_0^{sym} \\ \frac{\partial \Pi_{i:FC|C}}{\partial q_0^F} + \frac{\partial \Pi_{i:FC|C}}{\partial p_j^C} \frac{\partial p_j^C}{\partial q_0^F} = 0 \text{ if } \Gamma_0 = L_0^{seg}. \end{cases} \quad (28)$$

In the first line of equation (28), the licensor maximizes the deviation profit, that is the difference between the equilibrium played and the most profitable alternative, taking into account cross-price effects. The interests of licensor and industry are not aligned: a quality q_0^F that maximizes only the first element $\Pi_{FC|FC}^*$ would maximize joint licensor and industry profits, while the licensor also wants to make the firm's best alternative (second bracket) less profitable by reducing quality q_0^F in equilibrium. In the second line of equation (28), the licensor maximizes its customer's profit.

The licensor has to trade off selling two cheaper licenses for its label versus selling one more expensive license. Let L_0^{sym*} , resp. L_0^{seg*} , denote the license fees with optimal quality q_0^{F*} maximizing the license fee in the head-to-head, resp. segmented, case. The licensor wants to play the symmetric head-to-head equilibrium $\{FC, FC\}$ if $2L_0^{sym*} > L_0^{seg*}$.¹⁴ The trade-off crucially depends on horizontal differentiation μ : the higher μ , i.e. the lower the intensity of competition between the firms, the more profitable it is for a firm to offer a label that is also offered by the other firm; and higher surplus for the firm directly translates into higher license fees.

We numerically solve the first-order conditions of equation (28) for all values of horizontal differentiation μ , compare the resulting licensor profits for each equilibrium and find that there is a single threshold:

Proposition 1 ($2L_0^{sym*} - L_0^{seg*}$) *increases monotonically with horizontal differentiation μ : above $\mu = 0.46$, the licensor prefers symmetric head-to-head competition, selling two licenses; below this threshold, it prefers market segmentation, selling just one license.*

Figure III.3 shows the optimal quality chosen by licensor F : for $\mu < 0.46$, it is more profitable for the licensor to set a high quality and a high fee, attracting only one firm in

¹⁴This is related to, but not equal to the comparison $2[\Pi_{FC|FC}^* - \Pi_{C|FC}^*]$ versus $[\Pi_{FC|C}^* - \Pi_{C|C}^*]$, as the licensor sets different optimal qualities in each case.




Figure III.3: Equilibrium quality q_0^{F*} in stage 0.1 as a function of horizontal differentiation μ

the market segmentation equilibrium. For $\mu > 0.46$, licenser F chooses a low quality and fee but sells its label to both firms, leading to the head-to-head equilibrium in stage 0.2. Within a given product line equilibrium, the optimal quality q_0^{F*} generally increases with horizontal differentiation μ (for $\mu > 0.05$).

3.5 Minimum quality requirement

In order to evaluate the scope for regulatory intervention, we define social welfare. As we compute the label quality given the duopoly's pricing game, these are second-best values. Following the example of organic certification, there is potentially scope for a government-imposed minimum quality requirement for fairtrade labels. As in organic certification, this standard would leave the conventional market unchanged but raise the label's quality to a regulated minimum level \underline{q} . We assume that the social planner cannot force labeling organizations to adjust downwards.

In nested logit models, expected consumer surplus S is the inclusive value of the highest nest level; here, it is thus the inclusive value at the decision level to buy the product or the outside good from equation (16). Social welfare W is the sum of the consumer surplus, the firms' profits and the licenser fee:

$$S = \gamma \ln \left(\exp(A^P/\gamma) + 1 \right) \quad (29)$$

$$W = S + \sum_{j=1,2} \Pi_{j:y|z} \quad (30)$$

where $\Pi_{j:y|z}$ again denotes the profits of firm j when it plays y and the other firm plays z , with $y, z \in \{FC, C\}$.

We find that the social planner wants to maximize the number of available products. Thus, the social planner always wants both firms to offer F -labeled goods. However, for horizontal differentiation below $\mu < 0.46$ the social planner would have to *decrease* the label's quality to induce the head-to-head equilibrium, which he cannot do by assump-




Figure III.4: Minimum quality requirement \underline{q} as a function of horizontal differentiation μ (unregulated equilibrium quality q_0^{F*} from stage 0.1 in gray)

tion. For weak horizontal differentiation μ , the social planner sets the optimal quality for the segmented market constellation $\{FC, C\}$. As shown in Figure III.4, the social planner always sets a minimum quality requirement above the equilibrium quality of licenser F .

4 Market entry of industry standard I

In the beginning of the second period $t = 1$, an industry standard I enters the market and announces its quality. The for-profit licenser F can then adjust its quality and license fee. Firms can now decide whether to offer one or both labels. This increases the number of possible market constellations, but we can again restrict the analysis to cases where each label is offered at least by one firm.

We assume that the for-profit licenser cannot undercut its quality q_0^{F*} from the previous period. This is motivated by the observed qualities of labels in the coffee market, where the incumbent licenser has never adjusted downwards and new entrants have always established less stringent standards than the incumbent.

Our results show that the industry benefits from introducing an industry standard, because an additional good increases overall demand (less people opt for the outside good) and reduces the license fee. Moreover, for intermediate horizontal differentiation, one firm stops offering an F -labeled good, thereby reducing competition in that nest and payments to the licenser.

4.1 Consumer prices

We find that Lemma 1 can be generalized to a situation with two labeling organizations:

Lemma 2 *For all possible product line constellations, there are unique equilibrium prices p_i^s with $s = F, I, C$ and $i = 1, 2$ in stage 1.3; moreover*

- (i) *when the product lines are symmetric (both firms offer the same qualities), prices are symmetric;*
- (ii) *when firms compete head-to-head $\{FIC, FIC\}$ (both firms offer all qualities), prices are given by the marginal production costs $c(q^s)$ plus a constant and symmetric mark-up;*
- (iii) *in partial market segmentation $\{FIC, IC\}$, the mark-up on the F-labeled product is higher than the mark-up on the I-labeled product and the difference in mark-ups decreases when the label qualities become more similar.*

Proof. See Appendix on page 91. ■

With $y, z \in \{FIC, FC, IC, C\}$, we let $\Pi_{i;y|z}^*$ denote firm i 's reduced profit with unique profit-maximizing prices when it plays y and the other firm plays z .

4.2 Product line decisions

Turning to stage 1.2 of the game and analyzing the firms' decision to offer one or both labels, we compute the firms' best responses to each other's product line. Assume firm 1 offers FIC . Then, firm 2's best response is given by

$$\max \left\{ \Pi_{FIC|FIC}^* - L_1, \Pi_{FC|FIC}^* - L_1, \Pi_{IC|FIC}^*, \Pi_{C|FIC}^* \right\} \quad (31)$$

Solving the respective maximization problem for all other strategies of firm 1 and using symmetry allows us to fully characterize the equilibrium in stage 1.2 of the game. The equilibrium played in stage 1.2 of the game depends on μ , as well as on qualities q^I , q_1^F and fee L_1 .

4.3 License fee

As in the first period, when deciding on its license fee L_1 , the licensor has two options: it can aim to sell its label to both firms or it can decide to sell it to just one firm. In stage 1.1, this trade-off depends on μ as before and the quality q^I previously set by the Stackelberg leader industry standard I . The relevant cases are symmetric head-to-head competition $\{FIC, FIC\}$ and full market segmentation $\{FC, IC\}$, as before, plus additionally partial market segmentation $\{FIC, IC\}$. We also compute equilibrium qualities and prices for all other possible cases, but this section concentrates on the relevant cases, i.e. cases that are equilibria under certain conditions.

Assume first that the licensor aims at inducing the symmetric head-to-head equilibrium $\{FIC, FIC\}$, i.e. firms compete on all labels. In this case, the licensor sets its fee L_1 such that neither of the two firms offering FIC wants to deviate to offering IC only; the fee equals their deviation profit, given the other firm also offers FIC :

$$L_1^{sym} = \Pi_{FIC|FIC}^* - \Pi_{IC|FIC}^* \quad (32)$$

Numerically, we verify that firms indeed play the head-to-head equilibrium in stage 1.2 of the game when the licensor sets its fee at L_1^{sym} , as we have for all μ , q_1^F and q^I :

$$\Pi_{FIC|FIC}^* - L_1^{sym} = \max \left\{ \Pi_{FIC|FIC}^* - L_1^{sym}, \Pi_{FC|FIC}^* - L_1^{sym}, \Pi_{IC|FIC}^*, \Pi_{C|FIC}^* \right\}$$

Secondly, assume that the licensor aims to establish partial segmentation – as we call the product line constellation $\{FIC, IC\}$ – as an equilibrium in stage 1.2 of the game. Then, the licensor sets its fee L_1 such that the firm offering FIC has no interest to deviate to offering IC ; the fee equals the deviation profit of this firm, given the other firm offers IC .¹⁵

$$L_1^{pscg} = \Pi_{FIC|IC}^* - \Pi_{IC|IC}^*. \quad (33)$$

Third, assume that the licensor aims to establish full market segmentation $\{FC, IC\}$ as an equilibrium in stage 1.2 of the game. The licensor sets its fee L_1 such that the firm offering FC has no interest in deviating to offer IC ; the fee equals the deviation profit of this firm, given the other firm plays IC :

$$L_1^{seg} = \Pi_{FC|IC}^* - \Pi_{IC|IC}^*. \quad (34)$$

The second element of L_1^{pscg} and L_1^{seg} is identical, so that the licensor's preference between full market segmentation and partial market segmentation is determined by the first element. If the licensor chooses the higher of these two fees with a license fee defined as $\max\{L_1^{pscg}, L_1^{seg}\}$, we numerically verify that both firms have no interest in deviating from the chosen constellation for all μ , q_1^F and q^I .¹⁶ In both cases, the licensor sells just one license fee.

Summarizing this section, the licensor's profit Γ_1 can be written as:

$$\Gamma_1 = \max\{2L_1^{sym}, L_1^{pscg}, L_1^{seg}\} \quad (35)$$

4.4 Label quality of incumbent licensor F

As in the case with only one label (period $t = 0$), the optimal label quality q_1^{F*} maximizes the license fee. We assume that the incumbent for-profit licensor cannot decrease its quality below its monopoly value, q_0^{F*} , without seriously harming its brand image. For simplicity, we further assume that the licensor in the first period does not anticipate the entry of the industry standard in the second period. Using the envelope theorem, we can

¹⁵Theoretically, the possible alternative profits are $\Pi_{IC|IC}$ and $\Pi_{C|IC}$. However, we numerically have $\Pi_{IC|IC} > \Pi_{C|IC}$ for all q^I , q_1^F and μ .

¹⁶We numerically compute the equilibria for all L , μ , q_1^F and q^I : for many parameter constellations, the licensor cannot induce partial or full segmentation, but he can always induce the one that gives him the higher pay-off.




Figure III.5: Reaction function of quality q_1^{F*} as a function of industry standard quality q^I for $\mu = 0.5$ (equilibrium quality q_0^{F*} from the previous period in gray)

again write down the corresponding first-order conditions:

$$\frac{\partial \Gamma_1}{\partial q_1^F} = \begin{cases} \left[\frac{\partial \Pi_{i:FIC|FIC}}{\partial q_1^F} + \sum_s^{F,I,C} \frac{\partial \Pi_{i:FIC|FIC}}{\partial p_j^s} \frac{\partial p_j^s}{\partial q_1^F} \right] - \left[\frac{\partial \Pi_{i:IC|FIC}}{\partial q_1^F} + \sum_s^{F,I,C} \frac{\partial \Pi_{i:IC|FIC}}{\partial p_j^s} \frac{\partial p_j^s}{\partial q_1^F} \right] = 0 \text{ if } \Gamma_1 = 2L_1^{sym} \\ \left[\frac{\partial \Pi_{i:FIC|IC}}{\partial q_1^F} + \frac{\partial \Pi_{i:FIC|IC}}{\partial p_j^I} \frac{\partial p_j^I}{\partial q_1^F} + \frac{\partial \Pi_{i:FIC|IC}}{\partial p_j^C} \frac{\partial p_j^C}{\partial q_1^F} \right] = 0 \text{ if } \Gamma_1 = L_1^{pseg} \\ \left[\frac{\partial \Pi_{i:FC|IC}}{\partial q_1^F} + \frac{\partial \Pi_{i:FC|IC}}{\partial p_j^I} \frac{\partial p_j^I}{\partial q_1^F} + \frac{\partial \Pi_{i:FC|IC}}{\partial p_j^C} \frac{\partial p_j^C}{\partial q_1^F} \right] = 0 \text{ if } \Gamma_1 = L_1^{seg} \end{cases} \quad (36)$$

In the first line of equation (36), the licensor sets its quality q_1^F combining the effect on the firm's profits against the effect on the firm's best alternative. Both $\Pi_{FIC|FIC}^*$ and $\Pi_{IC|FIC}^*$ increase in q_1^F as it increases the differentiation between nests F and I , and decrease in q^I as it decreases differentiation between nests F and I . The two qualities are strategic complements: the higher the quality q^I of the industry standard, the higher the optimal quality q_1^{F*} of the licensor, allowing him to set a higher fee L_1^{sym} . In the two latter cases of equation (36), there is no such strategic element and the licensor set its quality q_1^F maximizing the profits of the firm offering F .

As an example, Figure III.5 plots the reaction function of the licensor quality q_1^{F*} to industry standard quality q^I for horizontal differentiation $\mu = 0.5$. For small q^I , the licensor induces partial market segmentation $\{FIC, IC\}$; for large q^I , the licensor induces head-to-head competition $\{FIC, FIC\}$. In the head-to-head equilibrium, the licensor distorts its quality downwards to increase its license fee by reducing the deviation profit, which explains the discontinuity in Figure III.5. Within a product line equilibrium, q_1^{F*} is increasing in q^I .




Figure III.6: Preferred product range equilibrium of licensor F as a function of quality q^I and horizontal differentiation μ

Let L_1^{sym*} , resp. L_1^{seg*} and L_1^{pseg*} , denote the license fees with optimal quality, i.e. quality q_1^{F*} maximizing the license fee in the head-to-head, resp. fully and partially segmented, case. We numerically compute the optimal qualities for all μ and q^I and then compare $2L_1^{sym*}$, L_1^{pseg*} , and L_1^{seg*} . Figure III.6 plots the resulting preferred product line of licensor F .

Proposition 2 *For strong horizontal differentiation $\mu > 0.61$, the licensor F induces head-to-head competition $\{FIC, FIC\}$ independently of industry standard quality q^I . For weak horizontal differentiation, $\mu < 0.05$, the licensor induces full market segmentation $\{FC, IC\}$. For intermediate values of μ , the product line equilibrium depends on industry standard quality q^I (Figure III.6).*

Details on numerical calculations: See Appendix on page 92. ■

The licensor prefers partial market segmentation over head-to-head competition if $L_1^{pseg*} > 2L_1^{sym*}$. Intuitively, low quality q^I decreases vertical competition between nests F and I , while low horizontal differentiation μ increases competition within nests. More in detail, lower horizontal differentiation μ increases the benefit of being the only firm offering an F -labeled good ($FIC|IC$ versus $IC|IC$) and increases the potential fee L_1^{pseg*} . At the same time, a lower μ decreases the mark-ups on the F -labeled product when both firms offer FIC ($FIC|FIC$ versus $IC|FIC$) and decreases the potential fee L_1^{sym*} .

The licensor prefers full market segmentation over partial market segmentation when $\Pi_{FIC|IC}^* < \Pi_{FC|IC}^*$. For low values of μ and q^I , the competition within nests is so strong that competing within a label market is not profitable: offering FC is better than offering FIC , given the other firm offers IC .

4.5 Label quality of new entrant I

In stage 1.0 of the second period, industry standard I sets its quality q^I , anticipating the equilibria in the following stages of the game, in particular the reaction of licensor F . The industry standard can influence the licensor by strategically setting its quality q^I . Propo-

sition 2 and Figure III.6 showed the levels of horizontal differentiation for which the industry standard can set its quality such that the licensor plays a segmentation equilibrium. We first compute the joint firm profit in the three cases mentioned before: head-to-head competition, partial segmentation, and full segmentation, subsequently comparing these three cases. Generally, firms want to segment the market as much as possible: the less product lines overlap, the higher joint firm profit.

The industry standard I maximizes joint profit of both firms. When the licensor induces head-to-head competition, we can use the expression for licensor fee L_1^{sym} from equation (32) to get an expression for joint firm profit:

$$\begin{aligned}\Pi^{sym} &= 2(\Pi_{FIC|FIC}^* - L_1^{sym}) \\ &= 2\Pi_{IC|FIC}^*\end{aligned}\quad (37)$$

In a partially segmented setting, where both firms offer an I -labeled good, but only one of them offers an F -labeled product, we can use the expression for licensor fee L_1^{pseg} from equation (33) to get an expression for joint firm profit:

$$\begin{aligned}\Pi^{pseg} &= (\Pi_{FIC|IC}^* - L_1^{pseg}) + \Pi_{IC|FIC}^* \\ &= \Pi_{IC|IC}^* + \Pi_{IC|FIC}^*\end{aligned}\quad (38)$$

For the full market segmentation equilibrium, we can use the expression for licensor fee L_1^{seg} from equation (34) to get an expression for joint firm profit:

$$\begin{aligned}\Pi^{seg} &= (\Pi_{FC|IC}^* - L_1^{seg}) + \Pi_{IC|FC}^* \\ &= \Pi_{IC|IC}^* + \Pi_{IC|FC}^*\end{aligned}\quad (39)$$

Comparing joint firm profits Π^{sym} and Π^{pseg} , the industry prefers partial market segmentation $\{FIC, IC\}$ over head-to-head competition if $\Pi_{IC|IC}^* > \Pi_{IC|FIC}^*$, i.e. if offering IC is more profitable when the other firm offers IC than if the other firm offers FIC . Numerically, this is almost always the case, because a firm offering FIC obtains a higher overall market share than a firm offering IC . Only for weak horizontal differentiation μ and exceptionally large vertical differentiation (low q^I and high q^F), the industry prefers head-to-head competition and this extreme region is never an equilibrium.

Comparing joint firm profits Π^{pseg} and Π^{seg} , the industry prefers full market segmentation over partial market segmentation if $\Pi_{IC|FC}^* > \Pi_{IC|FIC}^*$. Numerically, we verify that offering IC is always more profitable if the other firm offers FC than if the other firm offers FIC , because a firm benefits from being the only firm offering I -labeled goods. Firms thus always want to segment the market passing from $\{FIC, IC\}$ to $\{FC, IC\}$. However, we have seen in the previous section that the licensor does not play this equilibrium unless horizontal differentiation μ is weak.

Let us summarize the comparisons between the relevant cases both for the licensor and the industry: the licensor wants to play the head-to-head equilibrium when μ and q^I are high; the partially segmented equilibrium when μ is intermediate and q^I is low; and

the fully segmented equilibrium when μ and q^I are low (see Figure III.6). The industry always wants market segmentation.

Combining this finding about the industry's preferred market outcome with the licensor's reaction in Figure III.6 allows us to determine the equilibrium market constellations that are determined by the industry standard's quality q^I .

Proposition 3 *Depending on the degree of horizontal differentiation μ , the industry standard sets its quality q^{I*} following*

μ	equilibrium	q^{I*}	q_1^{F*}
> 0.61	head-to-head $\{FIC, FIC\}$	$\arg \max \{\Pi^{sym}\}$	$\arg \max \{L_1^{sym} q^{I*}\}$
$[0.16, 0.61]$	partially segmented $\{FIC, IC\}$	$\max\{q^I L_1^{pseg*} \geq 2L_1^{sym*}\}$	$\arg \max \{L_1^{pseg} q^{I*}\}$
$[0.13, 0.16]$	fully segmented $\{FC, IC\}$	$\max\{q^I L_1^{seg*} \geq L_1^{pseg*}\}$	$\arg \max \{L_1^{seg} q^{I*}\}$
$(0, 0.13]$	fully segmented $\{FC, IC\}$	$\arg \max \{\Pi^{seg}\}$	$\arg \max \{L_1^{seg} q^{I*}\}$

Details on numerical calculations: See Appendix on page 93. ■

When horizontal differentiation is strong ($\mu > 0.61$), the industry standard maximizes its profit in the symmetric head-to-head constellation from equation (37) by setting quality q^I under following first-order condition:

$$\frac{\partial \Pi^{sym}}{\partial q^I} = \frac{\partial \Pi_{IC|FIC}^*}{\partial q^I} + \frac{\partial \Pi_{IC|FIC}^*}{\partial q_1^F} \frac{\partial q_1^F}{\partial q^I} = 0 \quad (40)$$

The industry standard I maximizes the firms' surplus from offering the label taking into account that a higher q^I also induces a higher q_1^F (see discussion in Subsection 4.4). This strategic effect increases q^I , relative to the solution maximizing only the direct effect on $\Pi_{IC|FIC}^*$.

If the industry standard can induce partial market segmentation $\{FIC, IC\}$ with a positive quality q^I (i.e. when horizontal differentiation μ is intermediate with $\mu \in [0.16, 0.61]$), then the industry prefers this outcome over head-to-head competition. For intermediate horizontal differentiation μ , the industry standard sets its quality low enough to make the licensor just indifferent between playing the head-to-head equilibrium $\{FIC, FIC\}$ and partial market segmentation $\{FIC, IC\}$:

$$q^{I*} = \max\{q^I | L_1^{pseg*} \geq 2L_1^{sym*}\} \quad (41)$$

Thus, the equilibrium quality q^{I*} is lower than the quality that solves the first-order condition $\partial \Pi^{pseg} / \partial q^I = 0$, but the gain of playing an equilibrium with fewer products is high enough to compensate for the distortion in quality q^I . Graphically, the quality q^{I*} in Figure III.7 can be deduced from Figure III.6, as it is on the border between the area inducing $\{FIC, FIC\}$ and the area inducing $\{FIC, IC\}$.

Similarly, if the industry standard can induce full market segmentation $\{FC, IC\}$ with a positive quality q^I (i.e. when horizontal differentiation μ is sufficiently small with $\mu \in$




Figure III.7: Equilibrium qualities q^{I*} and q_1^{F*} as a function of horizontal differentiation μ (q_0^{F*} from the previous period with one label in gray)

$[0.13, 0.16]$), then the industry prefers this outcome over partial market segmentation. The optimal quality q^{I*} makes the licensor just indifferent between $\{FIC, IC\}$ and $\{FC, IC\}$:

$$q^{I*} = \max\{q^I | L_1^{seg*} \geq L_1^{pseg*}\} \quad (42)$$

Again, the quality q^{I*} in Figure III.7 is graphically on the border between the area inducing $\{FIC, IC\}$ and the area inducing $\{FC, IC\}$ in Figure III.6.

For weak horizontal differentiation μ ($\mu \in (0, 0.13]$), the industry standard can play an interior solution to its first-order condition in the fully segmented constellation, maximizing profits from equation (39). Analogously to the head-to-head case, the first-order condition in case of full market segmentation is

$$\frac{\partial \Pi^{seg}}{\partial q^I} = \frac{\partial \Pi_{IC|IC}^*}{\partial q^I} + \frac{\partial \Pi_{IC|FC}^*}{\partial q^I} + \frac{\partial \Pi_{IC|FC}^*}{\partial q_1^F} \frac{\partial q_1^F}{\partial q^I} = 0 \quad (43)$$

Figure III.7 represents the equilibrium quality q^{I*} for different values of μ , as detailed in Proposition 3. Comparing Proposition 3 with the results in the case with only one labeling organization in Proposition 1, we understand that the industry standard effectively reduces the offer of F -labeled products for horizontal differentiation $\mu \in [0.46, 0.61]$ (shaded area in Figure III.7).

Proposition 4 *The industry benefits from introducing the industry standard I, because*

- (i) *offering another vertically differentiated product increases total demand;*
- (ii) *for $\mu \in [0.46, 0.61]$, the introduction of the industry standard induces one firm to stop offering an F -labeled good, thereby reducing competition and payments to the licensor;*

- (iii) at any given horizontal differentiation μ , the introduction of the industry standard lowers the license fee.

Details on numerical calculations: See Appendix on page 93. ■

4.6 Minimum quality requirement

We use the same definitions of consumer surplus and social welfare as in equations (29) and (30). As before, we find that welfare increases in the number of products offered. The social planner wants to counteract the industry standard's effort to restrict product lines and reduce overlap. However, a minimum quality requirement is only binding for labels that are in equilibrium below this minimum standard \underline{q} . If the lower label is raised to the minimum standard, then the licensor strategically adjusts the higher label.

Table III.1 shows how the social planner determines the optimal minimum standard \underline{q} . In the two polar cases – for very large and very small μ – where the industry standard plays an interior solution, the minimum quality \underline{q} is not binding because the industry standard is already too high, leading to over-differentiation from the conventional market C relative to welfare optimizing values. In these markets, a minimum quality requirement cannot impact the status quo. A minimum quality requirement can only have a welfare-enhancing effect in the markets where the industry strategically distorts its quality to induce market segmentation. In these cases, the social planner solves the same equation as the industry standard, albeit the industry standard wants to be marginally below the solution inducing partial segmentation (resp. full segmentation) while the social planner wants to be marginally above inducing head-to-head product competition (resp. partial segmentation).

Table III.1: Minimum quality requirement \underline{q} set by the social planner as a function of horizontal differentiation μ

μ	equilibrium played	\underline{q}
> 0.61	head-to-head $\{FIC, FIC\}$	not binding
$[0.50, 0.61]$	head-to-head $\{FIC, FIC\}$	$\arg \max \{W^{sym}\}$
$[0.35, 0.50]$	head-to-head $\{FIC, FIC\}$	$\min\{q^I L_1^{pseg*} \leq 2L_1^{sym*}\}$
$[0.14, 0.35]$	partially segmented $\{FIC, IC\}$	not binding
$[0.09, 0.14]$	partially segmented $\{FIC, IC\}$	$\min\{q^I L_1^{seg*} \leq L_1^{pseg*}\}$
$[0.01, 0.09]$	fully segmented $\{FC, IC\}$	not binding

5 Conclusion

Our model describes the interaction between two firms and two labeling organizations of different quality; one of the labeling organizations is a for-profit licensor, the other one is an industry standard. We first model how a for-profit licensor sets its fee and quality

when it is the only labeling organization on a market with two firms. We then allow for the entry of an industry standard and model the competition between two labels. In order to model sensible substitution patterns, we develop a discrete choice model with both horizontal differentiation (exogenously given) and vertical differentiation (from endogenous product quality) using a nested logit.

Our results show that the equilibrium product line depends on horizontal differentiation: the market is segmented if horizontal differentiation is weak, while firms are in head-to-head competition when horizontal differentiation is strong. In summary, firms seek vertical differentiation when horizontal differentiation is low.

We further find that the industry benefits from reducing overlap in the firms' product lines; against this background, the industry standard can serve as a coordination tool to induce market segmentation and increase profits. Interestingly, there are cases where firms play the fully segmented equilibrium where not all firms offer *I*-labeled goods, even though the industry standard charges no license fee.

Social welfare always benefits from head-to-head competition in our setting, reflecting a fundamental love of variety of consumers as well as a benefit from stronger competition. This leads to a conflict between industry and consumers, where the former want to reduce product lines such that they do not overlap and the latter want to maximize product diversity. A minimum standard set by the regulator can improve the situation in some cases. In other cases, however, the industry standard, set as an interior solution, is too high relative to the welfare-maximizing minimum standard: firms in duopoly benefit from differentiating more than the welfare-maximizing level.

Our results shed some light on product line decisions in complex markets like the one for coffee: as we noted in the beginning, the product line equilibria in different national coffee markets are very different, with some featuring head-to-head competition (Germany) and others market segmentation (Finland), consistent with our theoretical analysis. Moreover, our model explains why the coffee industry collectively has an interest to introduce an industry standard. In practice, industry-related labels like UTZ and Rainforest Alliance have gained popularity in recent years. As the marginal production costs are lower, the global quantities of coffee sold under these industry-related labels are three times higher than the quantity sold under the Fairtrade label (Panhuysen and Pierrot 2014). It remains an open question however, whether industry standards are strategically distorted downwards in order to decrease competition. Overall, there remains considerable scope for further research: for example, our model is limited to the strategic interactions within one country, whereas in practice labeling organizations set their license fees on a global scale for many heterogeneous countries.

6 Bibliography

- Anderson S. P and De Palma A (1992) Multiproduct firms: a nested logit approach. *The Journal of Industrial Economics*, 40(3): 261–276.
- Andreoni J (1989) Giving with impure altruism: applications to charity and ricardian equivalence. *Journal of Political Economy*, 97(6): 1447–1458.
- Arnould E. J, Plastina A, and Ball D (2009) Does fair trade deliver on its core value proposition? Effects on income, educational attainment, and health in three countries. *Journal of Public Policy & Marketing*, 28(2): 186–201.
- Avenel E and Caprice S (2006) Upstream market power and product line differentiation in retailing. *International Journal of Industrial Organization*, 24(2): 319–334.
- Basu A. K and Hicks R. L (2008) Label performance and the willingness to pay for Fair Trade coffee: a cross-national perspective. *International Journal of Consumer Studies*, 32(5): 470–478.
- Beuchelt T. D and Zeller M (2011) Profits and poverty: certification’s troubled link for Nicaragua’s organic and fairtrade coffee producers. *Ecological Economics*, 70(7): 1316–1324.
- Bonroy O and Constantatos C (2015) On the economics of labels: how their introduction affects the functioning of markets and the welfare of all participants. *American Journal of Agricultural Economics*, 97(1): 239–259.
- Bottega L and De Freitas J (2009) Public, private and nonprofit regulation for environmental quality. *Journal of Economics & Management Strategy*, 18(1): 105–123.
- Bottega L, Delacote P, and Ibanez L (2009) Labeling policies and market behavior: quality standard and voluntary label adoption. *Journal of Agricultural & Food Industrial Organization*, 7(2): 1–15.
- Chambolle C and Poret S (2013) When fairtrade contracts for some are profitable for others. *European Review of Agricultural Economics*, 40(5): 835–871.
- Cheng Y.-L and Peng S.-K (2012) Quality and quantity competition in a multiproduct duopoly. *Southern Economic Journal*, 79(1): 180–202.
- Dragusanu R, Giovannucci D, and Nunn N (2014) The economics of Fair Trade. *Journal of Economic Perspectives*, 28(3): 217–36.
- Dragusanu R and Nunn N (2014) The impacts of Fair Trade certification: evidence from coffee producers in Costa Rica. Working Paper, Harvard University.
- Fischer C and Lyon T. P (2014) Competing environmental labels. *Journal of Economics & Management Strategy*, 23(3): 692–716.

- Friedrichsen J and Engelmann D (2017) Who cares about social image? Discussion Paper 1634, Deutsches Institut für Wirtschaftsforschung (DIW).
- Gallego G and Wang R (2014) Multiproduct price optimization and competition under the nested logit model with product-differentiated price sensitivities. *Operations Research*, 62(2): 450–461.
- Harbaugh R, Maxwell J. W, and Roussillon B (2011) Label confusion: the Groucho effect of uncertain standards. *Management Science*, 57(9): 1512–1527.
- Heyes A. G and Maxwell J. W (2004) Private vs. public regulation: political economy of the international environment. *Journal of Environmental Economics and Management*, 48(2): 978–996.
- de Janvry A, McIntosh C, and Sadoulet E (2015) Fair trade and free entry: can a disequilibrium market serve as a development tool? *Review of Economics and Statistics*, 97(3): 567–573.
- Johannessen S and Wilhite H (2010) Who really benefits from fairtrade? An analysis of value distribution in fairtrade coffee. *Globalizations*, 7(4): 525–544.
- Loureiro M. L and Lotade J (2005) Do fair trade and eco-labels in coffee wake up the consumer conscience? *Ecological Economics*, 53(1): 129–138.
- Mahenc P (2010) Wasteful labeling. *Organization*, 7(2): 6.
- Mason C. F (2011) Eco-labeling and market equilibria with noisy certification tests. *Environmental and Resource Economics*, 48(4): 537–560.
- McFadden D (1978) Modeling the choice of residential location. *Transportation Research Record*, 673: 72–77.
- Panhuysen S and Pierrot J (2014) Coffee barometer 2014. Technical report, Tropical Commodity Coalition, The Hague.
- Pelsmacker P. D, Driesen L, and Rayp G (2005) Do consumers care about ethics? Willingness to pay for fair-trade coffee. *The Journal of Consumer Affairs*, 39(2): 363–385.
- Podhorsky A (2015) A positive analysis of fairtrade certification. *Journal of Development Economics*, 116: 169–185.
- Poret S (2016) Label battles: competition among NGOs as standard setters. Working Papers 2016-01, Aliss.
- Potts J, Lynch M, Wilkings A, Huppé G, Cunningham M, and Voora V (2014) The state of sustainability initiatives review 2014: standards and the green economy. Technical report, International Institute for Sustainable Development (IISD) and the International Institute for Environment and Development (IIED).
- Richardson M and Stähler F (2014) Fair Trade. *Economic Record*, 90(291): 447–461.

- Ruben R and Fort R (2012) The impact of fair trade certification for coffee farmers in Peru. *World Development*, 40(3): 570–582.
- Saenz Segura F and Zuniga-Arias G (2008) *The impact of Fair Trade*: Wageningen Academic Publishers.
- Teyssier S, Etilé F, and Combris P (2014) Social-and self-image concerns in fair-trade consumption. *European Review of Agricultural Economics*, 42(4): 579–606.
- Train K (2009) *Discrete choice methods with simulation*: Cambridge University Press.
- Valkila J, Haaparanta P, and Niemi N (2010) Empowering coffee traders? The coffee value chain from Nicaraguan Fair Trade farmers to Finnish consumers. *Journal of Business Ethics*, 97(2): 257–270.
- Valkila J and Nygren A (2009) Impacts of fair trade certification on coffee farmers, cooperatives, and laborers in Nicaragua. *Agriculture and Human Values*, 27(3): 321–333.
- Villas-Boas S. B (2007) Vertical relationships between manufacturers and retailers: inference with limited data. *The Review of Economic Studies*, 74(2): 625–652.
- Von Schlippenbach V and Teichmann I (2012) The strategic use of private quality standards in food supply chains. *American Journal of Agricultural Economics*, 94(5): 1189–1201.
- Yu J and Bouamra-Mechemache Z (2016) Production standards, competition and vertical relationship. *European Review of Agricultural Economics*, 43(1): 79–111.

Appendices

A Proof of Lemma 1

In the symmetric head-to-head case $\{FC, FC\}$, profits are:

$$\Pi_{i:FC|FC} = D_i^F(p_i^F - q^F) + D_i^C p_i^C \quad (44)$$

Analyzing the first-order conditions, we find that there is a unique mark-up δ such that

$$p_i^{F*} = p_j^{F*} = p_i^{C*} + \delta = p_j^{C*} + \delta \quad (45)$$

with δ implicitly given by

$$\delta = 2\mu \left(1 - \frac{\mu}{\mu + \gamma \left[1 + (\exp(A^C / \mu^{F,C}) + \exp(A^F / \mu^{F,C}))^{\mu^{F,C} / \gamma} \right]} \right) \quad (46)$$

Simple calculations show that with $p^s = \delta + q^s$ the right hand side of the last equation is decreasing in δ , which establishes uniqueness. Furthermore, numerical calculations show that the second order conditions are satisfied at $p_i^s - q^s$. The same strategy applies for the symmetric equilibrium $\{C, C\}$.

In the asymmetric segmented case $\{FC, C\}$, firm i playing FC has the first-order conditions:

$$\frac{p_i^F - q^F}{p_i^C} = \frac{D_i^C \partial D_i^C / \partial p_i^F - D_i^F \partial D_i^C / \partial p_i^C}{D_i^F \partial D_i^F / \partial p_i^C - D_i^C \partial D_i^F / \partial p_i^F} \quad (47)$$

Substituting the demand functions and the respective derivatives leads to

$$\frac{p_i^F - q^F}{p_i^C} = 1 + \frac{(\mu^{F,C} - \mu) \exp(p_i^C / \mu)}{\mu \exp(A^C / \mu)} \quad (48)$$

The mark-up on the C -labeled good is identical to the mark-up on the F -labeled good when their qualities are equal, i.e. $\mu^{F,C} = \mu$. If the labels are vertically differentiated, then the mark-up on the F -labeled product is higher, as this is the market where the firm offering FC is in monopoly.

Furthermore, differentiating both sides of the equation with respect to p_i^F shows that the left-hand side is increasing in p_i^F while the right-hand side is decreasing in p_i^F . Additionally, using the solution of this equation numerical calculations show that

$$\frac{\partial \Pi_{i:FC|C}}{\partial p_i^C} = (p_i^F - q^F) \frac{\partial D_i^F}{\partial p_i^C} + D_i^C + p_i^C \frac{\partial D_i^C}{\partial p_i^C} = 0 \quad (49)$$

has exactly one solution in p_i^C .

Applying the same procedure for firm j we obtain

$$\frac{\partial \Pi_{j:C|FC}}{\partial p_j^C} = D_j^C + p_j^C \frac{\partial D_j^C}{\partial p_j^C} = 0 \quad (50)$$

has exactly one solution in p_j^C .

B Proof of Lemma 2

In the symmetric head-to-head case $\{FIC, FIC\}$, profits are:

$$\Pi_{i:FIC|FIC} = D_i^F(p_i^F - q^F) + D_i^I(p_i^I - q^I) + D_i^C p_i^C \quad (51)$$

Analyzing the corresponding first-order conditions, we find again, as in Lemma 1 that there is a unique mark-up δ :

$p^{s*} - q^s = \delta$ with δ implicitly given by

$$\delta = 2\mu \left(1 - \frac{\mu}{\mu + \gamma \left[1 + (\exp(A^C/\mu^{FI,C}) + \exp(A^{FI}/\mu^{FI,C}))^{\mu^{FI,C}/\gamma} \right]} \right)$$

Simple calculations show that with $p^s = \delta + q^s$ the right hand side of the last equation is decreasing in δ , which establishes uniqueness. Furthermore, numerical calculations show that the second order conditions are satisfied at $p_i^s - q^s$. As mentioned in the proof of Lemma 1, an analogous result holds for $\{FC, FC\}$ and $\{C, C\}$.

In the $\{FC, IC\}$ case, firm i playing FC has the first-order conditions:

$$\frac{p_i^F - q^F}{p_i^C} = \frac{D_i^C \partial D_i^C / \partial p_i^F - D_i^F \partial D_i^C / \partial p_i^C}{D_i^F \partial D_i^F / \partial p_i^C - D_i^C \partial D_i^F / \partial p_i^F} \quad (52)$$

Substituting the demand functions and the respective derivatives leads to

$$\begin{aligned} \frac{p_i^F - q^F}{p_i^C} &= \Psi \frac{\mu \exp(p_j^C/\mu) + \mu^{FI,C} \exp(p_i^C/\mu)}{\mu^{F,I} \exp[(v^F - p_i^F)/\mu^{F,I}] + \mu^{FI,C} \exp[(v^I - p_j^I)/\mu^{F,I}]} \\ \text{with : } \Psi &= \frac{\mu^{F,I} \exp(A^{FI}/\mu^{F,I})}{\mu \exp(A^C/\mu)} \end{aligned}$$

Furthermore, differentiating both sides of the equation with respect to p_i^F shows that the left-hand side is increasing in p_i^F while the right-hand side is decreasing in p_i^F . Additionally, using the solution of this equation numerical calculations show that

$$\frac{\partial \Pi_{i:FC|IC}}{\partial p_i^C} = (p_i^F - q^F) \frac{\partial D_i^F}{\partial p_i^C} + D_i^C + p_i^C \frac{\partial D_i^C}{\partial p_i^C} = 0 \quad (53)$$

has exactly one solution in p_i^C . Applying the same procedure for firm j we obtain

$$\frac{p_j^I - q^I}{p_j^C} = \Psi \frac{\mu \exp(p_i^C/\mu) + \mu^{FI,C} \exp(p_j^C/\mu)}{\mu^{F,I} \exp[(v^I - p_j^I)/\mu^{F,I}] + \mu^{FI,C} \exp[(v^F - p_i^F)/\mu^{F,I}]} \quad (54)$$

Again, while the left-hand side is increasing in p_j^I , the right-hand side is decreasing in p_j^C




Figure III.8: License fee as a function of label quality q_1^F for $\mu = 0.4$ and $q^I = 0.07$ (with $q_0^{F*} = 0.28$)

and

$$\frac{\partial \Pi_{j:IC|FC}}{\partial p_j^C} = (p_j^I - q^I) \frac{\partial D_j^I}{\partial p_j^C} + D_j^C + p_j^C \frac{\partial D_j^C}{\partial p_j^C} = 0 \quad (55)$$

has exactly one solution in p_j^C .

In the $\{FIC, IC\}$ case, we also compute the first-order conditions for the firm i playing FIC . Substituting the demand functions and the respective derivatives leads to

$$\frac{p_i^F - q^F}{p_i^I - q^I} = 1 + \frac{(\mu^{FIC} - \mu) \exp(p_i^I / \mu)}{\mu \exp(A^I / \mu)} \quad (56)$$

If the labels are vertically differentiated, then the mark-up on the F -labeled product is higher, as this is the market where the firm offering FIC is in monopoly.

The proof for uniqueness of equilibrium prices works identically to the previously shown full market segmentation $\{FC, IC\}$ case.

C Calculations for Proposition 2

For determining the equilibrium in stage 1.1 where the licensor sets its fee and quality, we numerically compute for each value of industry standard q^I and horizontal differentiation μ the optimal licensor quality q^{F*} for each of the three fees L^{sym} , L^{pseg} and L^{seg} . We then compare the reduced licensor profit with optimal quality $2L^{sym*}$, L^{pseg*} and L^{seg*} and keep the case that maximizes licensor profits. This gives us the reaction function of the licensor $q_1^{F*}(q^I)$ shown in Figure III.5.

As an illustration, Figure III.8 plots the license fee as a function of label quality q_1^F for horizontal differentiation $\mu = 0.4$. At this level of horizontal differentiation, the licensor chooses the highest fee between L_1^{sym} and L_1^{pseg} . Moreover, it cannot undercut the label quality from the previous period with only one label q_0^{F*} drawn as a gray line. For $q^I = 0.07$, the maximum is such that the licensor chooses the partially segmented market constellation. When q^I increases, the symmetric equilibrium becomes more attractive and the distance between the maxima of the two curves explains the jump on Figure III.5.




Figure III.9: Joint firm profit Π as a function of industry standard q^I given licenser reaction $q_1^{F*}(q^I)$ for $\mu = 0.4$


D Calculations for Proposition 3

In order to determine the equilibrium in stage 1.0 where the industry decides on its standard, we first compute the licenser reaction in q_1^F and L_1 for each level of horizontal differentiation μ and each industry standard q^I . We then determine for each horizontal differentiation μ , the q^{I*} that maximizes joint firm profit Π .


As an example, Figure III.9 shows the joint firm profit Π for different values of industry standard q^I , holding horizontal differentiation μ fixed at 0.4. There is a jump in the curve, because the licenser F switches from partial segmentation to head-to-head competition when q^I increases above 0.085. The joint firm profit is maximized by the corner solution ensuring partial segmentation.

E Calculations for Proposition 4.(iii)

Figure III.10 shows that the license fees are systematically lower upon entry of the industry standard. In the first graphic, with $\mu = 0.6$, we compare head-to-head competition license fees L_0^{sym} and L_1^{sym} for different values of licenser quality q^F and industry standard q^I ; L_1^{sym} is always smaller. In the second graphic, with $\mu = 0.1$, we compare segmented (resp. partially segmented) market license fees L_0^{seg} and $\max\{L_1^{pseg}, L_1^{seg}\}$ for different values of licenser quality q^F and industry standard q^I ; $\max\{L_1^{pseg}, L_1^{seg}\}$ is always smaller.



(a) License fee for head-to-head competition for possible values of q^I and q_t^F with $\mu = 0.6$



(b) License fee for segmented equilibria for possible values of q^I and q_t^F with $\mu = 0.1$

Figure III.10: Comparing license fees from $t = 0$ and $t = 1$

IV

Offset credits in the EU ETS: a quantile estimation of firm-level transaction costs

This is, of course, a very unrealistic assumption.

Coase (1960)
on the Coase theorem.

1 Introduction

The EU Emissions Trading System (EU ETS) aims at achieving the EU's carbon emission goals at minimum cost. Instead of imposing a tax, the policy determines an emission cap and lets the market determine the equilibrium emissions price. Ideally, all firms incur the same price for emissions and abatement is realized where it is cheapest, such that the aggregate abatement cost is minimized.

However, abatement and certificate costs are not the only costs arising from an emissions trading scheme: just like any other regulation, this policy has to be implemented by firms, causing a wide range of administrative, managerial, and information-related transaction costs. Typically, such frictions are unobserved by the econometrician. Presumably, many firms themselves do not track the value of their employees' time and resources spent in the course of EU ETS compliance and optimization. This chapter considers such unobserved trading cost, i.e. transaction costs that are conditional on trading.

This chapter focuses on the possibility for firms to use not only European certificates but also international offset credits. The EU ETS is linked to the international certificate market of the Kyoto Protocol. On aggregate, these additional foreign certificates increase the cap for European polluters and decrease their compliance cost. Offset credits were cheaper than European credits (European Union Allowances, EAUs) throughout Phase II of the EU ETS (2008-2012). However, the EU limited the quantity of offset credits by a firm-specific offset quota (*entitlement*). For the firms, offset usage was an unambiguous way to reduce compliance cost. Nevertheless, over twenty percent of regulated firms did not use any offsets.

This chapter uses firm-level data on offset usage to estimate the distribution of fixed trading costs, both for general entry into certificate trade and for offset use in particular. It brings together elements, first, from theoretical literature on transaction costs in emissions trading; second, from empirical literature on transaction costs in European emissions trading; and, third, from the small literature on the use of offset certificates in the EU ETS. Methodologically, this research uses binary quantile methodology.

While the abatement incentives of cap-and-trade schemes are amply discussed, most of the literature does not consider transaction costs. However, emissions trading – just like any other market transaction – is unlikely to be completely free of frictions. In his seminal article, Coase (1960) underlines that the irrelevance of initial property allocation for final resource allocation holds only if “costs to use the price mechanism” are negligible. The theoretical model of Stavins (1995) focuses on variable (quantity-dependent) trading costs, i.e. transaction costs arising from each certificate traded. Singh and Weninger (2016) build on this seminal work and show what distinguishes the impacts of variable and fixed (quantity-independent) trading costs. Fixed trading costs, as analyzed in this chapter, suppress some of the potential trades and lead to capacity- and certificate-underutilization; they also make initial allocation non-neutral, as firms only trade if their emissions and initial allocation are far away from each other.

Empirical evidence on transaction costs in environmental policy is scarce, as McCann et al. (2005) note in their literature review. Literature suggests that transaction costs and

other market imperfections have hampered the impact of US environmental trading programs (Tietenberg 2006, Hahn and Hester 1989). For example, Atkinson and Tietenberg (1991) argue that trading is too scarce to reach a cost-effective outcome; they claim that this inefficiency stems from the bilateral, sequential nature of trades leading to frictions and thus transaction costs in a broad sense.

Concerning the EU ETS, the literature generally finds that small firms behave more “passively” and that many firms lack the inherent institutional capacity for active trading, for instance Sandoff and Schaad (2009) on a sample of Swedish firms. Many German small and medium enterprises trade only at the end of the year and *only if* the grandfathered allocation does not suffice (Löschel et al. 2011). Schleich and Betz (2004) state that for small firms, transaction costs likely exceed certificate cost. Zaklan (2013) shows that most transactions take place between plants belonging to the same firm, which might be a way to reduce trading cost. Surveys show that large emitters face smaller per-tonne transaction costs (Heindl 2017, Jaraitė et al. 2010, Löschel et al. 2010, 2011). For example, Jaraitė et al. (2010) estimate that in Ireland *per tonne* transaction costs of the largest firms were €0.05 per tonne of emissions, while they were up to €2 per tonne for small firms. This suggests that transaction costs are mostly composed of fixed (quantity-independent) costs, potentially combined with smaller variable (per unit) costs. However, different authors use different definitions of transaction costs, making literature comparison difficult. Some studies include monitoring, reporting and validation (MRV) costs that occur for *all* regulated firms, while others concentrate on transaction costs that occur conditionally on trading.

Virtually all empirical work on trading costs in the EU ETS relies on survey-data, except Jaraitė-Kažukauskė and Kažukauskas (2015) who use transaction data from Phase I (2005-2007). They find that trading costs were a substantial factor inhibiting firms from actively trading European certificates, but they do not directly estimate their magnitude.

While the previously cited literature examines trading schemes with only one type of certificate, few articles deal with linked schemes with two certificate types. Trotignon (2012) shows that firms initially used few offsets until 2011, when there was a sharp increase in offset usage. He estimates the cumulated savings of firms at €1.5 billion. Ellerman et al. (2016) provide an aggregate description through the end of Phase II in 2012.

Binary choice methods are an established way to identify latent variables that shape behavior around some cut-off. In particular, one can identify unobserved costs from observed participation behavior to some cost-saving or profit-yielding activity. Anderson et al. (2011) use this approach on the marginal costs of regulating fuel-standards by observing to what extent car producers use a regulatory loophole of known costs to avoid the regulation on corporate fuel efficiency standards. Attanasio and Paiella (2011) similarly identify fixed household costs of financial market activity from household’s participation choice in the market. Conceptually, this resembles the present chapter, which identifies fixed costs by measuring the returns that firms forwent by avoiding trade. Quantile models are developed by Koenker and Bassett Jr. (1978), and applied to binary

choice by Kordas (2006). Belluzzo Jr (2004) uses them to estimate the distribution of willingness-to-pay for a public good, analogous to the present chapter: I measure transaction costs here from the observed “unwillingness-to-benefit” of firms. Going beyond usual estimation of the mean, this quantile methodology allows me to estimate the median as well as (a discrete approximation to) the whole distribution of transaction costs across 19 quantiles.

This chapter provides both an analytical and empirical contribution to the literature. First, it describes the observed offset usage behavior. Among the firms not using offsets, there are mostly small firms and, more particularly, firms with generous free allocations of European certificates. Across all firms, forgone revenue from unused offsets adds up to around €1.37 billion.

In a second step, I argue that firms’ reluctance to trade can be interpreted as transaction costs. Without such unobserved transaction costs, the offset entitlement would be an unequivocal “free lunch” opportunity. The share of firms incurring this opportunity cost can only be explained by the interference of some unobserved frictions: trading costs, as defined in this chapter, can include employees’ time/salaries, training and consultancy costs. Trading costs are assumed fixed (quantity-independent) and payable whenever a firm first decides to purchase emissions certificates in general or offset credits in particular; therefore, they might also be called entry costs.

The theoretical section lays out how trading costs change the firms’ optimization problem. Building on the standard model, I introduce a second type of certificate and fixed transaction costs. Such costs make the firms’ free allocation of certificates non-neutral, as firms with allocations larger than their emission do not *need* to engage in emissions trading: they can avoid transaction costs of active trading, such that they are less likely to use their offset entitlement. The model establishes a link between, on one hand, the decision to trade on the offset market and, on the other hand, both the initial net allocation status and offset entitlement. This relies on the fundamental assumption that a firm enters offset trading if and only if (observed) trading benefits exceed (unobserved) trading costs.

The empirical section uses this insight to estimate the latent transaction costs rationalizing a firm’s decision to not to enter the offset market. To the best of my knowledge, this is the first study to estimate costs using binary quantile regression. I identify the distribution of two transaction cost components: general trading cost and offset-specific cost.¹ The empirical results show that trading cost to the offset market is low for most firms, with a median of €905. The general trading cost is much higher with a median cost of €7,770. However, the estimated distribution of these costs is highly skewed, such that the means are much higher than the medians (€21,519 for mean general entry and €83,675 for offset market entry), resulting from some large outliers. Thus, a probit regression of the conditional mean is misleading about the costs faced by the majority of firms. Although these transaction costs are often small compared to other production

¹Note that I only consider fixed transaction costs that are conditional on trading *any* amount. The terms transaction cost and trading cost thus apply interchangeably.

factors, they make the use of offsets unprofitable for 21% of the firms. For bigger firms, investment in offset certificates mostly remains profitable.

The remainder of this chapter is organized as follows. After introducing the institutional and legal framework of international offset certificates (Section 2.1), I briefly explain the aggregate impact of offset trading in the EU ETS (Section 2.2) and the definition of transaction costs in this context (Section 2.3). I then set up a model of firm-behavior in the reference case, i.e. without any transaction/entry costs (Section 3.1), which I extend by adding entry costs (Section 3.2). Finally, I present the data and some stylized facts, explain the econometric methodology (Section 4) and present the estimated distribution of transaction costs (Section 5).

2 Background

The EU ETS and the international offset credits are based on a complex regulatory framework. This section briefly explains the key elements of this regulation. It further sketches out the aggregate mechanics of introducing a second type of certificate into an emissions trading system. Finally, this section explains in detail the specific transaction costs examined in this chapter.

2.1 Institutional framework

Each year, the European Union issues EU emissions allowances (EUAs) that, in total, equal the overall EU ETS emission cap. In Phase II – the period under study here – virtually all these certificates were distributed free of charge to the regulated firms, according to their historical emission levels (*grandfathered allocation*). At the end of each year, firms have to report their emissions and *surrender* certificates equaling their emissions: one for each tonne of CO₂. Other greenhouse gases are included as well, for instance methane (CH₄) and nitrous oxide (N₂O). Emissions of these other gases are converted with specific factors to CO₂ equivalent masses; hence the use of *tonnes of CO₂ equivalent* (tCO₂e) as a unit measuring quantities of certificates. Used certificates disappear, while unused certificates are *banked*, as they remain valid in subsequent years.

In order to coordinate international emission reduction efforts and to lower abatement cost for EU-based companies, the EU linked its ETS to the international framework established by the United Nations Framework Convention on Climate Change (UNFCCC, 1992) and the Kyoto Protocol. According to these international conventions, suitable projects that save emissions in unregulated parts of the world² can be validated and certified by UNEP. This procedure then generates Certified Emission Reductions (CERs, from Clean Development Mechanism) or Emission Reduction Units (ERUs, from Joint Implementation) that can be used to cover emissions in regulated parts of the world. CERs and ERUs are commonly called *international offset certificates*.³ The EU does not dis-

²Kyoto “non-Annex I” countries, in practice mostly China, Ukraine and India.

³CERs and ERUs can be used interchangeably under this legislation. I only use the term “offsets” from here, as everything applies equally to CERs and ERUs.

tribute offset certificates, meaning that firms can only use them after actively acquiring them, either by conducting projects generating offsets or by buying them on the market.

Within their obligations under the EU ETS, firms could substitute a limited number of European certificates with offset certificates. Such a substitution is attractive because offset certificates are cheaper than European certificates. However, to ensure that the bulk of emission reduction was achieved domestically, the EU restricted the quantity of offsets usable by each firm. The exact definition of this quota depends on the national government, but most countries computed it as a percentage share of the grandfathered allocation, cf. Table IV.4 on page 122. This yields a firm-specific *offset entitlement*, as a product of firm-specific allocation and country-/sector-specific percentage share. While European certificate allocations were distributed each year, the total offset entitlement was determined only in 2008; once fixed entitlements could then be used at any point in time over Phase II.

Offset entitlements were set in advance for the entire Phase II. In the middle of Phase II (April 2009), EU Directive 2009/29/EC announced that the usage limits of certain offsets would be transferable (*bankable*) into Phase III (2013-2020);⁴ however it was unclear what amounts and which types of certificates were involved. It was clear that “industrial gas” certificates, which constituted the bulk of offsets traded (Ellerman et al. 2016), would no longer be valid. Due to institutional obstacles, the final regulation ensuring the bankability and its conditions only appeared in November 2013,⁵ i.e. *after* the original claims for Phase II expired. From the perspective of a firm acting during Phase II, the end of Phase II had therefore to be considered as the temporal limit when planning the use of its offset entitlement.⁶

An alternative explanation for limited offset use would be that offset use was limited by supply side constraints. However, the data shows that offsets were always amply available: the central registry of the UNEP shows that the number of offsets generated at the end of 2012 was much higher than aggregate offset usage rights within the EU.⁷ Offset prices collapsed to virtually zero after the end of Phase II, which shows that the EU ETS demand was the driving force behind offset valuation.

2.2 Why are offset certificates cheaper?

Before looking at the impact of transaction costs, it is useful to consider the impact of offset certificates in general (without transaction costs) and, in particular, to show why they

⁴Phase III mainly extended the provisions of Phase II, in particular emission certificates from Phase II remained valid in Phase III. Important changes included new allocation rules, a reduction of free allocation combined with an increase in certificate auctioning, and the inclusion of the air transport sector.

⁵Commission Regulation (EU) No 1123/2013

⁶See Appendix B on page 123 for more detail.

⁷Theoretically, in addition to EU firm-level demand (analyzed in this chapter) there was scope for additional demand coming from the state-level; however, at the state-level of the Kyoto framework, offsets were perfect substitutes for Assigned Amount Units (AAUs). Given the large AAU overallocation to ex-Soviet Union states (so-called “hot air”), the evidence suggests that AAUs are usually sold far below the price of EUAs, CERs, and ERUs (Aldrich and Koerner 2012).




Figure IV.1: Stylized illustration of aggregate market equilibrium with two alternative offset supply levels

have been cheaper than European certificates. Transaction costs are added in Section 3.2. International offset credits cover emissions from geographic regions not previously included in the scope of EU ETS. As such, they are a *spatial flexibility mechanism* (Stevens and Rose 2002) allowing firms to abate where it is cheapest and have the abatement credited via the creation of offset credits. The introduction of offsets increases the overall cap imposed by the EU ETS. Potentially, the cap could increase by an amount equal to the sum of all firms' offset quotas (*entitlements*).⁸

Figure IV.1 illustrates the resulting market equilibrium: in an unregulated situation, emissions have no cost and firms emit $e_{\text{unreg'd}}^*$. In an ETS without offset credits, the standard result for emissions trading holds: the market clears at the regulated maximum emission level \bar{e} at price p^e , equal to the marginal abatement cost at \bar{e} (Trotignon 2012). When offsets are introduced, they are perfect substitutes for European certificates up to the quota. When offsets are costly to produce (supply q_{low}^o), their availability increases the overall cap, lowers the price and moves the equilibrium to e_{low}^* , where prices are set at the level for which offset supply clears. This equalizes European certificate and offset prices $p_{\text{low}}^e = p_{\text{low}}^o$. When offset creation is cheap (supply q_{hi}^o), firms would like to buy more offset certificates than allowed and emit up to e_{hi}^* . The aggregate offset quota \bar{q}^o binds in that case. The resulting constrained equilibrium at $\bar{e}' = \bar{e} + \bar{q}^o$, no longer ensures equal prices: European certificates trade at marginal abatement cost p_{hi}^e at \bar{e}' . The over-supply of offset certificates drives their price down to p_{hi}^o . The price differential $\Delta_p = p^e - p^o$ is always positive or zero; its magnitude depends on the difficulty to generate offsets and on the stringency of the offset quota.

2.3 Definition and interpretation of transaction costs

The EU ETS causes direct costs through abatement and certificate prices. Besides this direct (and intended) cost, the EU ETS causes a number of (unintended) information-, administration- and management-related frictions, which in this chapter are broadly

⁸See Mansanet-Bataller et al. (2011), Nazifi (2013) for more details from a finance perspective.

understood under the term transaction costs.

I separate these costs into two parts according to their contingency: the first are “administrative costs” due to mandatory actions, such as costs for monitoring, reporting and validating emissions (MRV) as well as the EU registry service charges. These administrative costs are unavoidable and thus cannot explain firms’ (non-)entry to the offset market. The second, generally known as trading or entry costs, are the consequence of voluntary trading choices, such as information gathering, forecasting of certificate prices, finding trading partners, bargaining, contracting, managing price risk, or finally simply the costs of out-sourcing the whole trading process. This chapter concentrates on the latter, i.e. trading costs, which are defined as all frictions that are important for a firm’s decision to actively enter the certificate market. While some firms have to purchase certificates, others have allocations large enough to avoid any active involvement in the certificate market.

This definition is narrower than in other works which consider the overall cost of establishing, managing, monitoring and enforcing a policy (Krutilla and Krause 2010, Joas and Flachsland 2016).⁹ However it is also broader than the definition used in some of the literature, as it includes *all frictions* preventing firms from entering the certificate market, in particular it includes outsourcing costs and purely psychological factors that discourage managers from devoting resources to certificate trading.

Heindl (2012) finds that information-procurement alone – the biggest upfront cost – costs firms about 17 employee-workdays. He also finds that information and trading costs do not depend on firm size. While this indicates fixed costs, most surveys present their results on a per-tonne basis, i.e. interpreting them as variable rather than fixed costs, cf. Table IV.1. None of them asks about offset-related costs. The brokerage fees of an individual transaction are low,¹⁰ while there are upfront entry costs. Just as an example, setting up a trading account at the ICE (the biggest exchange, clearing about 90% of emission certificate trade in Europe) costs €2,500 in direct fees,¹¹ while an individual transaction thereafter costs only cents.¹²

A multitude of news and data providers (Point Carbon), consulting firms (ICIS/Tschach), and financial transaction services (brokerage like TFS Green, exchange platforms like ICE) have emerged. The fact that firms use such costly services indicates a lack of cost-free information. Moreover, descriptive management literature highlights

⁹In particular, this chapter concentrates on costs borne by firms and does not take into account what Joas and Flachsland (2016) call “public-sector costs” borne by the regulatory authority.

¹⁰Convery and Redmond (2007) establish a list of direct transaction fees: brokers have large minimum trade sizes and take between 1 and 5 cent fee per certificate (tCO₂e). Exchanges take smaller trades and charge between 0.5 and 3 cent per certificate.

¹¹As indicated on <https://www.theice.com/fees>; retrieved on 01/03/2015.

¹²Internationally operating firms could decide to create offset certificates in their own plants abroad, rather than purchasing the certificates on a market place. This chapter assumes that the large majority of firms bought their certificates, which matches anecdotal evidence about offsets. However, this claim cannot be proven due to data restrictions. If this claim is not true, the estimations in this chapter remain valid, but their interpretation changes from trading costs to transaction costs in the generation of offsets.

Table IV.1: Overview of transaction cost estimates per firm in the EU ETS in the literature

	Average transaction costs	Cost structure	Scope	Time
Heindl (2012)	€4,193 information €4,659 trading €12,223 MRV	fixed + variable	Germany	2009 and 2010 (yearly)
Jaraitė et al. (2010)	€71,860 early implementa- tion €74,180 MRV	variable	Ireland	Phase I
Löschel et al. (2010)	€1.79/tCO ₂ e if emissions < 25,000t €0.36/tCO ₂ e if emissions ≥ 25,000t	variable	Germany	2009
Löschel et al. (2011)	€11,136 MRV and infor- mation €2,654 trading	fixed + variable	Germany	2010
Jaraitė- Kažukauskė and Kažukauskas (2015)	show significance, no magnitude		EU	Phase I

Source: Cited studies and author's computation from estimates stated therein.

the discrepancy between actual and intended market practice: firms use simple heuristics instead of fully optimizing their behavior (e.g. Veal and Mouzas 2012). These anecdotal elements support the idea of transaction costs, even though firms may rarely account for them as such explicitly.

3 Model

First, a static model describes firm's optimization problem in presence of two types of emission certificates without transaction costs. In a second step, I examine how behavior changes in the presence of fixed transaction costs. Simply put, firms always want to use offset credits, unless transaction costs are higher than potential returns from using the cheaper offset credits. Given the institutional background, the model is static with just one period corresponding to Phase II of the EU ETS.

3.1 Emissions trading with offset credits: reference scenario without trading costs

For the purpose of this chapter, it is useful to look at firms' optimization problem aggregated over Phase II. As a reference case, this Subsection extends the standard emissions trading model with a second type of certificate and *without* adding trading costs. Firms can separate the decision of emission levels and produced quantities from the partitioning between European and offset certificates.

In the absence of offsets, marginal abatement cost is constant across firms and equal to the European certificate price p^e in equilibrium (e.g. Montgomery 1972). Each firm i

jointly produces some quantity y and emissions e , maximizing profits:

$$\max_{y_i, e_i, q_i^e, q_i^o} \pi = y_i - C(y_i, e_i) - p^e(q_i^e - q_i^{e0}) - p^o q_i^o, \quad (1)$$

$$\text{subject to } e_i = q_i^o + q_i^e, \quad (2)$$

$$q_i^o \leq \bar{q}_i^o, \quad (3)$$

where π is profit and $C(y_i, e_i)$ production cost, which depends on emissions e_i and output y_i sold at a price normalized to 1. $C(y_i, e_i)$ is assumed continuous and twice differentiable. I assume that reducing emissions at a given production level increases cost, $C_e < 0$.¹³ q_i^o is the amount of offsets and q_i^e the amount of European certificates used. At the beginning of Phase II, firms are given a free allocation of European certificates q_i^{e0} and a firm-specific offset entitlement \bar{q}_i^o . They can buy and sell European certificates at market price p^e and offsets at price p^o .

The firm must simultaneously solve three problems: decide on the produced quantity y_i^* , determine the emission level e_i^* , and split compliance (i.e. an amount of certificates equal to e_i) between the international offset and European certificates. To satisfy the first-order condition, emissions e_i^* have to be such that marginal abatement cost is equal to the marginal certificate price, and production such that y_i^* that marginal production cost (including compliance cost) is equal to 1 (price normalization).

The compliance cost is composed of the cost of buying the certificate quantities q_i^e and q_i^o necessary to cover the emission level e_i^* , abatement cost and the forgone revenue of adjusting production relative to a production level that would be optimal at zero emission cost. The marginal cost is either p^e or p^o depending on which type of certificate is used to cover the last (marginal) emission. Offsets are perfect substitutes for European certificates up to the quota; their price difference is thus zero or positive: $p^e - p^o =: \Delta_p \geq 0$.¹⁴ The result is straightforward: as a perfect substitute at a lower price, offset credits are unambiguously preferable to European certificates, up to the regulated entitlement \bar{q}_i^o . Only if emissions are above \bar{q}_i^o , the firm covers the remaining emissions by using the more expensive European certificates. Compared to a system with only European certificates, the firm saves an amount equal to $\bar{q}_i^o \Delta_p$. The optimization problem can be simplified as

¹³ C_y and C_e denote the partial derivatives with respect to y and e , respectively. The production cost function includes abatement cost, as the marginal cost of reducing emissions by a tonne at same output equals $-C_e$ (see Singh and Weninger 2016, for further details).

¹⁴For the purpose of this chapter, I only consider situations in which offset certificates are strictly cheaper than European certificates, as the alternative where both prices are equal is qualitatively not different from a system without offsets. Moreover, the data reveals that in practice there has always been a clear price discount for offset certificates.

follows:¹⁵

$$\max_{e_i} \pi(y^*(e_i), e_i) = \begin{cases} y^*(e_i) - C(y^*(e_i), e_i) - p^o e_i, & \text{if } 0 < e_i \leq \bar{q}_i^o \\ y^*(e_i) - C(y^*(e_i), e_i) - p^e e_i + \bar{q}_i^o \Delta_p, & \text{if } \bar{q}_i^o < e_i \end{cases} \quad (4)$$

In the EU ETS, the offset entitlement \bar{q}_i^o is, in practice, small compared to emissions. Virtually all firms need to use European certificates in addition to offsets, meaning that the constraint in equation (3) is binding. The usual result that marginal abatement cost are equalized across firms at the price level p^e remains valid.

3.2 Trading costs for both certificate markets

I now assume that firms face some general entry trading cost to enter any certificate market, i.e. the cost of setting up a trading department no matter the type of certificates. Only once they have such a trading department, they actively enter certificate trading and can incur an additional cost contingent on entering the offset market. They can avoid both costs if they only use their freely allocated European certificates. Firms with emissions greater than their allocation have to buy certificates and cannot avoid the general component of trading cost. Profit equation (1) has now two additional fixed cost terms:

$$\begin{aligned} \pi &= y^*(e_i) - C(y^*(e_i), e_i) - p^e q_i^e - \mathbb{1}^e \kappa^e - p^o q_i^o - \mathbb{1}^o \kappa^o, \\ &= y^*(e_i) - C(y^*(e_i), e_i) - p^e e_i - \mathbb{1}^e (\kappa^e + \mathbb{1}^o (\kappa^o - \Delta_p q_i^o)), \end{aligned} \quad (5)$$

$$\text{where } \mathbb{1}^o = 1 \text{ iff } q_i^o > 0 \quad (6)$$

$$\mathbb{1}^e = 1 \text{ iff } q_i^o > 0 \vee q_i^e - q_i^{e0} > 0 \quad (7)$$

where a firm incurs general entry trading costs κ^e if it buys any certificates, but also needs to pay additional information costs, κ^o , to enter the less well-known offset market. Firms that are “long” in equilibrium, i.e. which received more free allocations than needed for their emissions ($q_i^{e0} > e_i^*$), are not obliged to purchase certificates. “Short” firms cannot behave “autarkic” (Jong and Zeitzberger 2014): they must enter the market to buy some certificates and, thus, consider the general trading cost κ^e sunk when deciding about offset usage. The impact of transaction costs on offset usage and incurred total cost depends on the relative magnitudes of κ^o , $\kappa^o + \kappa^e$ and $\bar{q}_i^o \Delta_p$.

As usual with fixed entry costs, firms enter trading if, and only if, profits are higher with entry relative to non-entry. Given the specific cost structure assumed here, short firms enter the offset market if $\kappa^o < \bar{q}_i^o \Delta_p$, while long firms enter if $\kappa^o + \kappa^e < \bar{q}_i^o \Delta_p$. Thus, entry to the offset market is a binary choice, yielding “all-or-nothing” behavior.¹⁶ In this situation, grandfathered allocations create a discontinuity that impacts firm behavior.

¹⁵The allocation term $p^e q_i^{e0}$ in equation (1) is a choice-independent lump-sum transfer and can be dropped from the maximization problem.

¹⁶This part assumes that firms have emissions greater than their offset entitlement, which is the case for over 98% of the firms.

This assumes firms take their allocation status as given when deciding about their entry to the offset market. The fixed cost at emission level $e_i = q_i^{e0}$, i.e. the switching point between short and long, could cause firms to restrict their emissions to q_i^{e0} . Appendix C on page 124 formalizes this condition and tests whether there is any empirical evidence for such behavior, i.e. bunching of firms at the threshold. While theoretically possible, there is no empirical evidence for such an adjustment. Trading costs do not impact the marginal cost-benefit trade-off: both above and below q_i^{e0} firms face a certificate price of p^e , such that the main mechanism of the ETS is independent of fixed transaction costs.¹⁷

Let firm “net allocation status” $\mathbb{1}_i^{long}$ be a dummy variable indicating that allocation q_i^{e0} is larger than emissions e_i^* ,¹⁸ and $\mathbb{1}_i^o$ is again the dummy indicating the use of offset certificates.

$$\mathbb{1}_i^o = \begin{cases} 1 & \text{if } \bar{q}_i^o \Delta_p > \kappa^o + \mathbb{1}_i^{long} \kappa^e, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

4 Data and empirical research design

I use administrative data from the EU ETS. Descriptive data analysis reveals four stylized facts that my empirical analysis relies on: (a) offset certificates are cheaper than European certificates; (b) virtually all firms have emissions greater than their offset entitlement; (c) a non-negligible number of firms (22%) does not use their offset entitlements; and (d) the distributions of firms’ emissions and entitlements are highly dispersed.

4.1 Emissions, allocation and offset entitlement

This chapter mainly relies on compliance data of the European ETS Registry (European Union Transaction Log, EUTL), which combines all member states’ national registries of Phase II (2008-2012). This comprehensive administrative data comprises the allocated European certificates, verified emissions, and surrendered certificates (EUAs, CERs and ERUs) for *all* 13,590 plants subject to the ETS.

I aggregate the data over Phase II, because offset quotas were defined over the whole period and could be used at any point during the phase, without any yearly constraint, so that the decision whether to use offsets was ultimately only revealed once, on the last day of Phase II. The data does not contain transactions *per se*, but all firms using offsets must have acquired them previously. Firms had no interest to stockpile offsets beyond

¹⁷An underlying assumption is that firms take prices as given: every individual firm is too small to consider its own impact on the price level, i.e. it has no market power on the certificate market. On the aggregate, p^e depends on the number of firms using offset certificates. To the extent that transaction costs reduce access to the offset market, they are neither neutral for p^e nor, consequently, for y^* and e^* : second-order effects decrease the offset price p^o and increases the European certificate price p^e . While these price effects are essential for a general equilibrium and welfare assessment, they are not informative on transaction costs and are beyond the scope of this chapter.

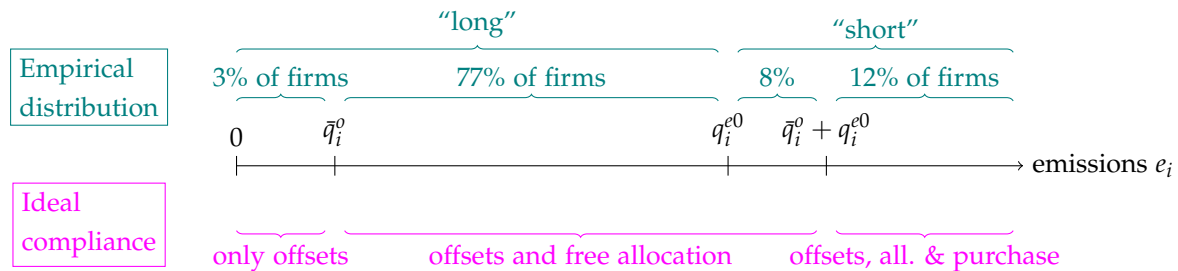
¹⁸The dummy variable is defined at the firm level, thus allowing for cost-free within-firm trade. Moreover, it includes dynamic considerations: given firms could bank certificates, $\mathbb{1}_i^{long} = 1$ if the *cumulative* sum of emissions does not exceed the *cumulative* sum of allocation in any year of Phase II.

the end of Phase II if they could also use them for compliance: in this chapter, offset usage is thus equated with offset acquisition. Moreover, all firms which were “short” in allocation, i.e. had emissions larger than their free allocation, had to buy certificates, either European or offset.¹⁹

A matching with Bureau van Dijk’s Orbis company database reveals ownership structures that link many of these individual plants.²⁰ This matching matters as the relevant decision likely happens at the firm level, even though regulation, allocation, and offset entitlements are defined at plant level. After some data cleaning,²¹ around 9,000 plants belonging to 4,578 firms remain. Over half of the plants belong to firms that own just one plant.

The plant-specific offset quota (*entitlement*) \bar{q}_i^o is the product of a country-specific offset percentage multiplied by the plant’s free allocations q_i^{e0} over Phase II. For the purpose of this chapter, the entitlement has been computed using this rule and verified using the International Credit Entitlement tables published by the EUTL in 2014.

Allocations have been generous, such that 80% of the firms could cover all of their emissions using only grandfathered allocations; these firms are called the “long” firms in the remainder of this chapter. Offset entitlement \bar{q}_i^o is so small that only 2.8% of firms are able to comply by using offsets only. Table IV.2 shows that free allocation has, on average, been just above emissions. Firms have a wide variety of sizes, with some firms owning up to 158 plants and being active in 11 sectors or 17 countries.



4.2 Price spread and realized savings

Daily price data for offsets (CERs) and European certificates (EUAs) is available from Intercontinental Exchange. Offsets are expected to trade at a lower price compared to

¹⁹There are certainly *some* firms which entered the market without being legally obliged by being short. If many firms fall into this case, the ratio between offset cost and general cost is biased toward general cost, while the overall distribution still holds. In presence of transaction costs however, only short firms have an interest to buy additional European certificates.

²⁰For more information on this extensive matching to the “global ultimate owner” level, see Jaraitė et al. (2013); or their website <http://fsr.eui.eu/CPRU/EUTLTransactionData.aspx>; retrieved on 06/09/2016.


²¹Plants from countries that do not participate in the standard way, as described in Section 2.1 (Estonia, Iceland, Lithuania, Liechtenstein, Malta and Norway; 220 plants), and firms that have offset-use beyond the legal limit (most likely because of merger and acquisition transactions that are unobserved in this data set; 94 plants) are excluded. Also excluded are about 4,000 plants that never registered any emissions, ceased existing in 2011/12, or have their first emissions after 2009.

Table IV.2: Descriptive firm statistics

	Mean	Median	SD	Min	Max
Number of countries active	1.13	1	.728	1	17
Number of plants	1.88	1	5.03	1	158
Number of sectors active (NACE definition)	1.12	1	.566	1	11
Free allocated EUAs (ktCO ₂ e)	1,975	112	13,831	.015	380,586
Emissions (ktCO ₂)	1,919	78.5	16,148	.003	563,608
International credit entitlement (ktCO ₂ e)	272	12	2,335	.001	91,537
Used offset credits (ktCO ₂ e)	208	8.34	1,494	0	55,536
Savings from offset use (k €)	799	31.2	5,836	0	217,412
Unexploited profits from offsets (k €)	627	22	7,370	.00465	200,316
Firms using all offset entitlement (in %)	50.5				
Firms using no offsets (in %)	22				
N					4,578

Source: EUTL and author's computations.

European certificates. Indeed, offsets have always traded at a positive discount from European certificates. Figure IV.2 shows that the price differential was rather small in the beginning. After few months, the spread increased and offsets have been up to €7 cheaper than European certificates, with a mean price difference of €3.60.



(a) Prices of EUAs and CERs on the secondary market

(b) EUA-CER price spread

Figure IV.2: Prices of EU certificates and offsets

Source: www.theice.com

This price spread allowed firms to achieve considerable savings,²² reaching €217.4 million for the largest firm. 78% of all firms have used offsets; together, they saved €3.6 billion compared to using only European certificates.²³ The 22% of firms that did

²²Savings are approximated by multiplying the annual average price spread with the amount of offset certificates used in that year, because the actual transaction prices are not observed.

²³These numbers take prices as given, so they cannot be interpreted as the general-equilibrium savings from offset usage: as seen in Section 2.2, the counterfactual EUA price in absence of offset credits would have been higher than the observed prices. the estimates used in Table IV.2 are, thus, a lower bound for the

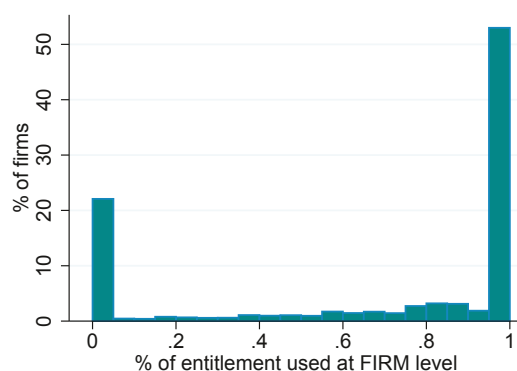


Figure IV.3: Ratio of used offset credits over overall offset entitlement

Source: EUTL and author's computations.

not use offsets could have used another 288 million tCO₂e certificates and generated €1.37 billion at 2012 prices. Among firms that used offsets, firms have saved on average €799,000, while the median is only €31,200.

4.3 Descriptive evidence for transaction costs

Many firms did not use their offset entitlements. Given the large supply of offset certificates and their low price, this is surprising. Factors that prevented firm entry are interpreted as transaction costs by this chapter, such as the costs of information procurement and other frictions.

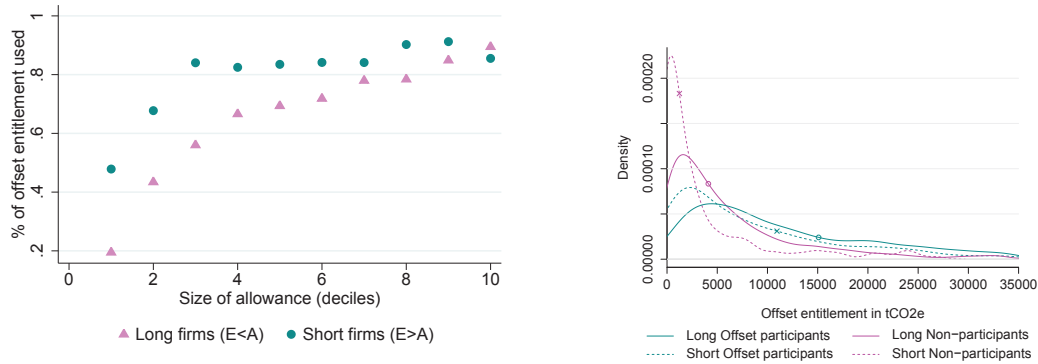
The stylized facts supporting the idea of fixed (rather than quantity-dependent) costs are (a) a largely binary behavior between using either the maximum allowed or no offsets at all; (b) the non-neutrality of European certificate net allocation status for entering the offset market; and (c) an increasing likelihood of entry to the offset market as offset entitlement increases.

The offset usage of firms mostly followed a binary “all-or-nothing” pattern, suggesting the presence of fixed trading cost. Figure IV.3 shows used offsets as a percentage of the total offset entitlement: over half of the firms used *all* their offset entitlements and almost a quarter of the firms used *none*. While per-unit costs could lead to intermediate usage rates, fixed entry costs for market entry can explain such binary behavior. Most multi-plant firms with intermediate usage are composed of plants that exhibit an all-or-nothing behavior: this results hint at coordination problems within firms.²⁴

Transaction costs depend on initial allocation. Short firms are legally bound to trade, meaning that they should consider general trading costs as sunk, whereas offset-specific cost applies to both long and short firms. Moreover, with fixed costs, firms with large off-

de facto achieved savings from offset usage. Stephan et al. (2014) estimate demand elasticity as being high, such that actual firms' savings may be as high as €20 billion, as offset availability decreased the overall stringency of the cap. Moreover, it does not account for the incurred transaction costs.

²⁴My personal interviews (not representative) revealed that large firms sometimes have sub-firms that are not well integrated, for example for recently acquired sub-firms, so that emissions trading departments do not take them into account.



(a) Offset usage by allocation and according to deciles of offset entitlement, i.e. potential profit
 (b) Distribution of offset entitlement in the different groups of firms

Figure IV.4: Relationship between offset use, offset entitlement and net allocation status
 Source: EUTL and author's computations. Density estimation using Gaussian kernel from `density()` in R, with smoothing bandwidths calculated by Silverman's rule of thumb; for readability, the graph is cut at 50 ktCO₂e, although both densities continue beyond. Crosses and circles indicate median values.

set entitlements are more likely to trade, as the potential gain becomes larger compared to entry costs. Figure IV.4a shows the interaction between size and allocation status: at lower size deciles, firms use offsets rarely, with a large difference between long and short firms. As size increases, firms become more likely to use offsets, while at the same time the difference between long and short firms becomes less marked. At the tenth size decile, virtually all firms trade and there is no significant difference between long and short firms' behavior.

Assuming that firms make rational decisions, plants which do not trade must estimate their trading costs to be higher than their potential profit, such that the mean offset entitlement multiplied by the mean price spread gives us a lower bound of the magnitude of these transaction costs (similar to the reasoning in Attanasio and Paiella 2011). At the same time, the opposite is true for firms that do enter the offset market. These two distributions largely overlap, but Figure IV.4b shows that the means and medians are strongly different. In general firms that do not use offsets tend to be smaller, with half of firms below 3,600 tCO₂e of offset entitlements (while the median is 16,600 tCO₂e for firms using offsets). Nevertheless, both distributions stretch out above 50,000 tCO₂e, showing that the separation is not clear cut. The largest firm without offset use has a 262,000 tCO₂e entitlement; 9% of the firms have larger entitlements and they *all* enter the offset market. Among firms with offset use, the size distribution of long and short firms is similar. On the opposite, small short firms are overrepresented in the group that does not use offsets.

The size distribution of firms' offset entitlements in Figure IV.4b is highly dispersed; similar levels of inequality are found for emissions, number of plants and grandfathered allocations. The empirical methods used need to be chosen such that they are robust to rare but extremely large outlier firms.

4.4 Econometric methodology

The model links binary firm behavior, i.e. using any offset credits or not, and the magnitudes of unknown entry costs κ^o and κ^e , to the known quantities q_i^{e0} , e_i and $\bar{q}_i^o \Delta_p$. We want to measure the latent fixed transaction cost κ_i^* , while observing only the binary outcome $\mathbb{1}_i^o$ equal to 1 if κ_i is smaller than opportunity cost $\bar{q}_i^o \Delta_p$:

$$\begin{aligned} \mathbb{1}_i^o &= \mathbb{1}\{\bar{q}_i^o \Delta_p > \kappa_i^*\} \\ &= \mathbb{1}\left\{ \underbrace{\bar{q}_i^o \Delta_p}_{\text{potential profit}} > \underbrace{\kappa^o + \kappa^e \mathbb{1}_i^{long}}_{\text{trading cost}} + \epsilon_i \right\} \end{aligned} \quad (9)$$

In this binary choice setup, $\bar{q}_i^o \Delta_p$ is the firm-specific cut-off value relevant for the decision to trade. Other than in most binary choice settings with a common cut-off at zero, for instance standard probit, a firm-specific cut-off allows us to identify an intercept as it fixes a scale for the two estimated parameters κ^o and κ^e in terms of units of $\bar{q}_i^o \Delta_p$ (i.e. euros).²⁵

This method relates to binary methods to measure “willingness-to-pay” (WTP). Here, rather than estimating WTP, I identify transaction costs by interpreting any forgone profits $\bar{q}_i^o \Delta_p$ as “unwillingness-to-benefit” or, in other words, opportunity costs. If the error term is assumed to be independent and identically distributed, following a normal distribution, equation (9) would describe a standard probit model in which coefficients are normalized such that the coefficient of the potential profit equals 1. The other coefficients then measure transaction costs in euros, as when willingness-to-pay is estimated by normalizing the utility of income to 1.²⁶ However, the stylized facts presented in Section 4 strongly suggest that this homoskedastic normality assumption does not hold; consequently, probit is not an appropriate model. If the distribution of transaction costs is skewed, an estimation of the mean cost is not the most representative summary statistic as it might be driven by large outliers.

Following empirical work by Kordas (2006) and Belluzzo Jr (2004), I estimate a range of binary quantile regressions to analyze the conditional *distribution* of transaction costs rather than just the conditional mean. This semi-parametric method is more robust to non-symmetric error distributions and outliers. For all quantiles $\tau \in (0, 1)$, I define the conditional quantile $Q_{\kappa^*}(\tau)$ as the τ^{th} quantile of the transaction cost distribution F_{κ^*} :

$$Q_{\kappa^*}(\tau | \mathbb{1}_i^{long}) := F_{\kappa^*}^{-1}(\tau) = \kappa_t^o + \kappa_t^e \mathbb{1}_i^{long} \quad (10)$$

These quantiles are identified using the observed offset-market entry $\mathbb{1}_i^o$ and the

²⁵ \bar{q}_i^o is measured in tCO2e of offset entitlement and Δ_p is the mean price spread measured in €/tCO2e.

²⁶ The standard normalization of a probit sets the standard deviation σ to 1; in contrast, the standard deviation is a free parameter here (see Train 2009).

monotone transformation of equation (9). Then $Q_{1_i^o}(\tau)$ may be written as:²⁷

$$Q_{1_i^o}(\tau | 1_i^{long}, \bar{q}_i^o \Delta_p) = \mathbb{1}\{\bar{q}_i^o \Delta_p \geq \kappa_\tau^o + \kappa_\tau^e 1_i^{long}\} \quad (11)$$

The probit regression draws its identification from the conditional mean assumption $E(\epsilon_i | x) = 0$ and the normality assumption, while the following methodology estimates the median and draws its identification from the assumption that the conditional *median* error is zero. The earliest estimator using this semi-parametric assumption is the maximum score estimator by Manski (1975). At the median with $\tau = .5$, this estimator maximizes the number of “correct predictions” using an indicator function:

$$\max_{\kappa_\tau^o, \kappa_\tau^e} S_{n\tau}(\kappa_\tau^o, \kappa_\tau^e; \bar{q}_i^o \Delta_p) = n^{-1} \sum_{i=1}^n [\mathbb{1}_i^o - (1 - \tau)] \mathbb{1}\{\bar{q}_i^o \Delta_p - \kappa_\tau^o - \kappa_\tau^e 1_i^{long} \geq 0\} \quad (12)$$

Similar to the median, we can estimate other conditional quantiles. While intuitive, this estimator is not continuous, which makes it difficult to optimize and determine standard errors. To resolve this issue, Horowitz (1992) formulates a smoothed maximum score estimator using a kernel function to obtain a continuous function of the estimated parameters, which Kordas (2006) extends to quantiles other than the median. The smoothed binary quantile estimator at quantile $\tau \in (0, 1)$ solves the following problem:

$$\max_{\kappa_\tau^o, \kappa_\tau^e} S_{n\tau}^*(\kappa_\tau^o, \kappa_\tau^e; h_n, \bar{q}_i^o \Delta_p) = n^{-1} \sum_{i=1}^n [\mathbb{1}_i^o - (1 - \tau)] \Phi\left((\bar{q}_i^o \Delta_p - \kappa_\tau^o - \kappa_\tau^e 1_i^{long})/h_n\right) \quad (13)$$

where $\Phi(\cdot)$ is a continuous, differentiable kernel function and h_n an appropriate bandwidth that tends to zero as the sample size increases.

The estimation of this model involves optimization over a complex function, in particular when using the discrete version of equation (12). I use R to implement Kordas’ S-Plus/Fortran code to perform simulated annealing following the algorithm of Goffe et al. (1994). Simulated annealing has the advantage of being more robust to starting values, local optima and discrete parts of the objective function; although computationally more demanding, the full code including bootstrapping runs in less than six hours. With a large sample, such as the one used in this chapter, the results of Manski’s discrete quantile maximum estimator and of Horowitz’ smoothed estimator turn out to be virtually identical. Standard errors are calculated by bootstrap methods.

5 Estimation results

According to my results, transaction costs are around €100,000 on average. Their distribution is skewed: many firms face small transaction costs, while a few firms have high costs. In particular, the offset-specific cost is much smaller than general entry cost for

²⁷We observe a transformation of the latent variable by an indicator function that is a monotone transformation. See Koenker and Hallock (2001) on the equivariance of quantile estimates to monotone transformations.

most firms. This section illustrates how quantile regressions can add valuable information if the underlying distribution is asymmetric.

The binary quantile regression estimates the distribution of transaction costs from which each firm draws its transaction cost. As this distribution is not assumed to follow a known functional form, it is described here by estimating 19 quantiles, from the 5th to the 95th percentile in steps of 5 percentage points. For better readability, Table IV.3 shows only selected quantiles, while Figure IV.5 shows the full estimation for all quantiles (19 separate estimations).

The transaction cost components are measured in units of potential profit, i.e. in euros. The median offset-specific cost κ^o is estimated around €905, which means that a short firm with enough offset entitlement to generate €905 of offset revenue has a 50% chance of participating. While transaction costs are low, at around €500 for the lower quarter of the transaction cost distribution, their values are high at the upper end with €201,919 for the highest quantile ($\tau = .95$). The distribution for κ^e indicates that long firms (with generous initial allocations) are much more reluctant to trade. At the median, their behavior is consistent with an *additional* cost equivalent to €7,770. This goes up to the higher quantile estimates around €41,900 for $\tau = 0.95$. A long firm thus needs potential profits of €7,770+ €905= €8,675 to have a 50% probability to use offsets.

The quantile analysis reveals that the transaction cost distribution spans a large range and is strongly skewed: while the difference between the median quantile and lower quantiles is small, there are large outliers driving the estimates of the highest quantiles. Consequently, the means (bottom of Table IV.3)²⁸ can be misleading about the transaction cost distribution. Of a similar order of magnitude, the probit estimates of the conditional mean are also much higher than the median.²⁹ Figure IV.5 plots probit estimates with a cross and adds the distribution of the normal error to represent the distribution implied by probit assumptions.³⁰ Despite the similar means, quantile and probit estimates are significantly different for most quantiles and yield different perspectives on the transaction cost distribution. For virtually all quantiles, the impact of net allocation status exceeds the offset-specific cost: the bulk of transaction costs stems from the general cost component κ^e . Firms thus refrained from using their offset entitlement not because of offset-specific trading costs, but rather to avoid certificate trading altogether. However, the means, both from probit and from quantile regression (bottom of Table IV.3), obscure this finding and suggest that transaction cost for offset are *on average* larger than the ones for general trading. There are some large outliers in the distribution of κ^o .

These results are more intuitive if we switch the axis of the standard quantile plot Figure IV.5, such that we obtain the estimated cumulated density function of firm's trans-

²⁸Means from the quantile regression are computed with the following steps: (a) estimate quantile parameters in 5% steps from the 5th percentile to the 95th; (b) predict market entry probability depending on firm characteristics (see Appendix E on page 128); (c) impute transaction cost from τ equal to predicted probability; and (d) take average across all observed firms.

²⁹More detail on these parametric estimations can be found in Appendix D on page 126.

³⁰Due to the renormalization, the error term does not follow a standard normal distribution, instead having a larger standard deviation.

Table IV.3: Estimates from binary quantile estimation and probit regression

τ	All firms		Manufacturing		Electricity	
	$\hat{\tau}^o$	$\hat{\tau}^e$	$\hat{\tau}^o$	$\hat{\tau}^e$	$\hat{\tau}^o$	$\hat{\tau}^e$
	offset-sp.	general	offset-sp.	general	offset-sp.	general
0.05	35.0*** [25; 152]	1.0 [-94; 587]	950.6*** [345; 1,308]	12.5 [-89; 2,335]	35.6*** [21; 143]	-13.2 [-87; 1,140]
0.1	35.0*** [30; 344]	1.0 [-96; 1,824]	1,013.8*** [354; 1,359]	936.0 [-64; 2,919]	32.7*** [25; 284]	-7 [-93; 1,373]
0.25	472.9*** [35; 587]	2,817.5*** [1,444; 4,675]	965.0*** [330; 1,378]	2,732.8*** [797; 4,429]	338.5*** [32; 906]	4,198.3*** [718; 7,867]
0.5	904.7*** [378; 2,753]	7,769.5*** [3,976; 10,616]	1,045.3*** [340; 1,538]	5,417.8*** [4,015; 10,696]	917.0*** [393; 5,169]	7,695.8*** [2,880; 15,417]
0.75	9,352.6*** [2,746; 12,741]	17,876.2*** [9,995; 30,478]	1,295.6*** [393; 11,331]	21,376.0*** [11,276; 36,002]	12,587.2*** [3,970; 26,390]	15,291.6** [1,466; 29,466]
0.9	28,392.9*** [17,596; 99,858]	57,135.0** [1,712; 165,116]	21,426.0*** [11,018; 52,336]	63,250.2*** [32,879; 132,068]	88,307.2*** [29,228; 170,252]	108,950.3* [1,223; 141,695]
0.95	201,919.4*** [79,334; 304,069]	7,184.6*** [264; 476,038]	301,294.8*** [23,545; 309,215]	13,588.2* [102; 486,145]	165,532.4*** [65,021; 236,666]	31,900.4* [5,274; 388,442]
Mean	83,675	21,519	123,133	64,269	65,322	62,542
Probit	109,557.2***	44,302.5***	173,656.8***	98,911.7***	48,632.4**	4,059.1
N	4,578		2,938		1,640	

Note: Function optimized by simulated annealing, significance and point-wise 95% confidence intervals are determined by bootstrap (500 replications). Columns 1 and 2 show the result of the binary quantile regression, dependent variable is the offset use dummy equal to 1 if the firm used any offsets, regressors are forgone profits, i.e. offset entitlement multiplied by price spread (coefficient normalized to one), "long" allocation dummy equal to 1 if the firm could cover all emissions with its allocation and a constant. Columns 3 to 6 show the result of the same regression with additional dummies for sector affiliation (and their interaction with the allocation dummy). Manufacturing includes cement, pulp and paper, glass, ceramics, metals, oil refining and "other."

action costs as shown in Figure IV.6a. One can infer a probability density function from this cumulated density function by using standard kernel density methods (Figure IV.6b). Again, these figures show how some large outliers drive the high mean of κ^o : the tail of the probability density function of the offset-specific transaction cost shows a bump that is driven by the only four non-participating firms with potential profits above €200,000. The mean and thus the results of a probit regression may be considered to be a misleading statistic in such a case. Figure IV.7 compares the estimated probability of entering the offset market from the probit and quantile model to the observed frequencies at different entitlement magnitudes. Particularly for smaller emitters, the quantile method predicts entry probability much better than the probit. Analogously, the fit of the quantile estimation is strong if evaluated with the method outlined by Kordas (2006), i.e. checking whether predicted and observed probabilities coincide (cf. Appendix E on page 128).³¹

³¹The better fit does not come as a surprise: the quantile model fits 38 free parameters, while the probit only fits three free parameters.

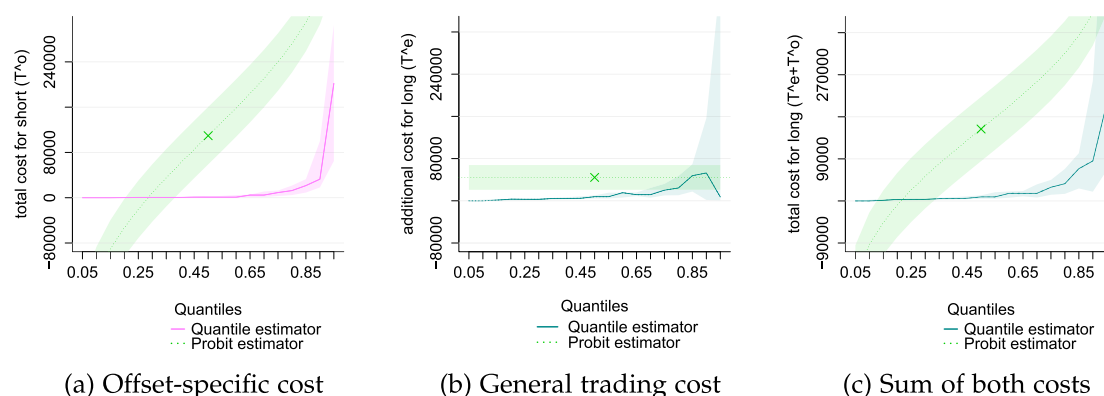


Figure IV.5: Estimated transaction cost (in €) - quantile plot

Note: Quantile estimates for all 5th percentiles from 5% to 95%. The dotted green line shows the mean estimate from probit, the shaded bands represent the point-wise 95% confidence intervals.

5.1 Sector-specific results

While the data set is too small for a full sector-specific analysis, the right-hand side of Table IV.3 shows the result separately for manufacturing and electricity-generating firms. Electricity and heat generation account for one-third of all firms, but half of total emissions. Electricity firms are known to have active and sophisticated compliance and trading behavior, likely because of the experience from electricity trading (Heindl 2017, Jong and Zeitzberger 2014).

Results (Table IV.3 and Figure IV.8) show that the sectors explain some of the observed transaction cost heterogeneity: while costs are similar around the median for manufacturing and electricity firms, I do not find any large outliers in the electricity sector, meaning that this sector's means are considerably lower than that of manufacturing. Thus, a handful of manufacturing firms is driving the high result at the 95th percentile of the pooled estimation.

For electricity and heat generation firms, the estimates for several quantiles of the

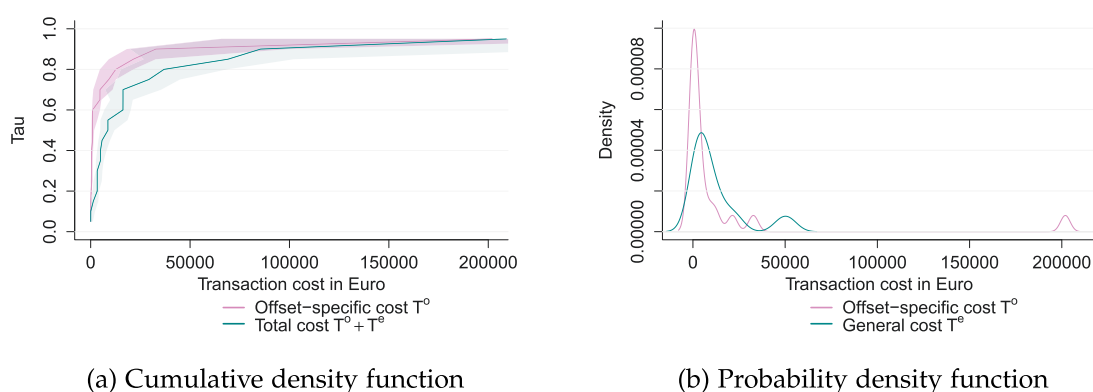


Figure IV.6: Estimated transaction costs - distribution function and density

Note: Probability density function in Figure IV.6b estimated from cumulative density function in Figure IV.6a using kernel density in R.

general cost κ^e are not significant. Moreover, the probit estimate is not significant. As virtually all large electricity firms trade emission certificates and most are short on certificates, this general component is difficult to identify: as noted in Jong and Zeitlberger (2014) (on Phase I), the “energy sector was ‘forced’ to purchase certificates as it was the only ‘under-allocated’ sector”, although this refers to Phase I. The offset-specific cost remains significant and similar to the main estimation.

For manufacturing, both estimates are similar to the ones for the general case: means are much higher than medians, offset-specific costs are less relevant than general costs for most of the distribution and, nevertheless, the means are higher for κ^o . Carbon cost is a less important cost factor for manufacturing firms and they own, on average, fewer plants with smaller emissions than electricity and heat generation firms.

6 Conclusion

Within their obligations from the EU ETS, firms had the opportunity to reduce expenses by using their right to substitute European certificates with international offset certificates: a priori, it is profitable to use cheaper offset certificates. However, many firms did not use their offset entitlement. After briefly explaining the aggregate mechanics of offsets in the EU ETS, this chapter shows the impact of fixed transaction costs on offset usage and estimates a (discrete approximation to the) distribution of fixed transaction costs rationalizing firms’ entry into the offset market.

Prior work mostly uses survey data to show that compliance with the EU ETS generates transaction costs. To the best of my knowledge, this is the first study to establish a framework to assess the magnitude of transaction costs empirically through the use of compliance data. Moreover, this is the first study to use binary quantile methodology, which allows for comparing quantiles (in particular the median) of the cost distribution, thus revealing its skewness. Entry costs are estimated to be at the median €7,770




Figure IV.7: Observed frequencies and predicted probabilities of quantile method and probit (cut at 40,000 tCO₂e for better readability)

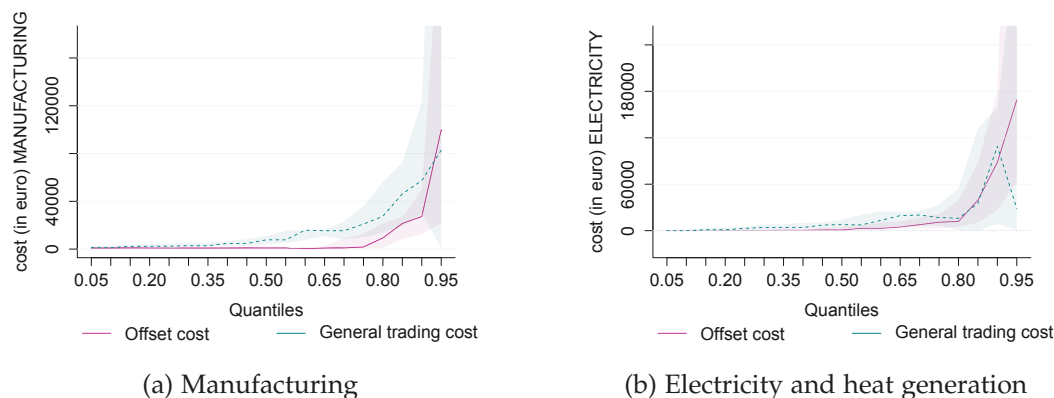


Figure IV.8: Sector-specific quantile estimation results (in €)

(mean €21,519) for general entry to the certificate market (any type of certificate) plus €905 (mean €83,675) for entry to the offset market: for the majority of firms, transaction costs are mostly due to trading in general, rather than the offset-specific entry costs. Not surprisingly, trading costs are more relevant for manufacturing firms than for electricity firms. Virtually all electricity firms use offsets, which can be explained by their large emissions (fixed costs are then more easily overcome) and the large cost share of emissions in electricity production (such that electricity firms usually have specialized emissions trading departments). The quantile estimation shows, thus suggesting that these means are largely driven by few large outliers. Consequently, this chapter illustrates the advantage of using binary quantile methods in addition to the typical parametric approaches.

Environmental policy aims at reducing ecological harm at minimum cost to society. Nevertheless, most academic and policy-related work only accounts for the compliance and abatement costs of the EU ETS. However – just like any regulation – the EU ETS causes administrative and management-related transaction costs. My estimates suggest that trading costs are relevant in practice: firms significantly deviate from the scenario without transaction costs. Designing policy remains “an empirical matter” (Montero 1998). Usually, regulation aims at giving the optimal incentive structure, while this chapter argues that regulatory complexity also creates costs. As the objective of a regulation becomes more complicated, there appears to be a trade-off between incentive perfection and a need to keep complexity for the regulated firms at bay – incentives only work as intended if they are understood and implemented at low cost. From this perspective, this chapter aims at contributing to the practical debate about the shape of environmental policy. Empirical evidence for transaction costs calls for simpler policy designs, rather than more sophisticated (but complicated) policy designs.

Note that with fixed trading costs, only large firms benefit from the cost reduction of offset certificates, meaning that small firms are disadvantaged. On this point, some action has been taken with the possibility for small emitters to opt-out of the EU ETS.³²

³²The possibility for such an opt-out for firms with emissions below 25,000 tonnes was created with

Alternatively, Heindl (2017) and Joas and Flachsland (2016) suggest moving to more upstream regulation.

This chapter addresses only part of the actually arising transaction costs: all administrative costs that are not contingent on trading – cannot be influenced by firm behavior – are unable to be captured using my methodology, for instance monitoring and reporting costs and registry fees.³³ My estimates are thus only one part of the costs that should be included in the policy discussion. Importantly also, these transaction costs are not synonymous with overall efficiency loss: while effort spent in information gathering certainly does not improve welfare, a real welfare effect analysis needs to look at the bigger picture of the general equilibrium. It would be interesting to estimate the impact of offset certificates on European certificate prices, as well as to examine more closely the price distortions (both on European certificates and offsets) caused by trading costs.

The estimated residual transaction cost is essentially a black box measuring *all* the frictions preventing firms from investing in offsets. It remains to be analyzed in detail what these costs include and how they could be reduced to implement a less distortionary policy. In fact, this chapter encompasses both “hard” financial costs and more “soft” behavioral factors, such as inattention, salience, or risk aversion, etc. Importantly, firms’ aversion to use offsets could also be due to reputation considerations, as offsets have received bad press in most countries. However, we are talking about the behavior of firms, therefore psychological factors should play less of a role than they do for consumer decisions.

Directive 2009/29/EC in June 2009 (Art. 27). However, not all member states have implemented this rule and in Germany, for example, only a handful of installations have used this option.

³³Registry fees in Phase II ranged from €100 for the period to €15,000 per year, depending on the country and (for some countries) emission size, cf. EUTL website.

7 Bibliography

- Aldrich E and Koerner C (2012) Unveiling Assigned Amount Unit (AAU) trades: current market impacts and prospects for the future. *Atmosphere*, 3(1): 229–245.
- Anderson B. J, Convery F, and Di Maria C (2011) Technological change and the eu ets: The case of ireland. Working Paper 43, IEFE Bocconi.
- Atkinson S and Tietenberg T (1991) Market failure in incentive-based regulation: the case of emissions trading. *Journal of Environmental Economics and Management*, 21(1): 17–31.
- Attanasio O and Paiella M (2011) Intertemporal consumption choices, transaction costs and limited participation in financial markets: reconciling data and theory. *Journal of Applied Econometrics*, 26(2): 322–343.
- Belluzzo Jr W (2004) Semiparametric approaches to welfare evaluations in binary response models. *Journal of Business & Economic Statistics*, 22(3): 322–330.
- Coase R (1960) The problem of social cost. *Journal of Law & Economics*, 3(4): 87–137.
- Convery F and Redmond L (2007) Market and price developments in the European Union emissions trading scheme. *Review of Environmental Economics and Policy*, 1(1): 88–111.
- Ellerman A. D, Marcantonini C, and Zaklan A (2016) The European Union Emissions Trading System: ten years and counting. *Review of Environmental Economics and Policy*, 10(1): 89–107.
- Elsworth R, Worthington B, and Morris D (2012) *Help or hindrance? Offsetting in the EU ETS*: Sandbag.
- Goffe W, Ferrier G, and Rogers J (1994) Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*, 60(1): 65–99.
- Hahn R and Hester G (1989) Where did all the markets go – an analysis of EPA’s emissions trading program. *Yale Journal on Regulation*, 6: 109–153.
- Heindl P (2012) Transaction costs and tradable permits: Empirical evidence from the EU emissions trading scheme. Discussion Paper 12-021, Zentrum für Europäische Wirtschaftsforschung (ZEW).
- (2017) The impact of administrative transaction costs in the EU emissions trading system. *Climate Policy*, 17(3): 314–329.
- Horowitz J (1992) A smoothed maximum score estimator for the binary response model. *Econometrica*, 60(3): 505–531.
- Jaraitè J, Convery F, and Di Maria C (2010) Transaction costs for firms in the EU ETS: lessons from Ireland. *Climate Policy*, 10(2): 190–215.

- Jaraitė J, Jong T, Kažukauskas A, Zaklan A, and Zeitzberger A (2013) Matching EU ETS accounts to historical parent companies. A technical note, European University Institute, Florence.
- Jaraitė-Kažukauskė J and Kažukauskas A (2015) Do transaction costs influence firm trading behaviour in the European emissions trading system? *Environmental and Resource Economics*, 62(3): 583–613.
- Joas F and Flachsland C (2016) The (ir)relevance of transaction costs in climate policy instrument choice: an analysis of the EU and the US. *Climate Policy*, 16(1): 26–49.
- Jong T and Zeitzberger A (2014) EU emissions trading by energy firms. Research Paper 99, Robert Schuman Centre for Advanced Studies.
- Koenker R and Bassett Jr. G (1978) Regression quantiles. *Econometrica*, 46(1): 33–50.
- Koenker R and Hallock K (2001) Quantile regression: an introduction. *Journal of Economic Perspectives*, 15(4): 43–56.
- Kordas G (2006) Smoothed binary regression quantiles. *Journal of Applied Econometrics*, 21(3): 387–407.
- Krutilla K and Krause R (2010) Transaction costs and environmental policy: an assessment framework and literature review. *International Review of Environmental and Resource Economics*, 4(3-4): 261–354.
- Löschel A, Brockmann K, Heindl P, Lutz B, and Schumacher J (2011) Hoher Anpassungsbedarf im EU-Emissionshandel ab 2013 – deutliche Defizite bei der Vorbereitung in den Unternehmen. CO2 Barometer 2011, KfW/ZEW.
- Löschel A, Kiehl K, Heindl P, Lo V, Koschel H, and Koesler S (2010) Effizienzpotenziale des Emissionshandels noch nicht ausgeschöpft – Strategien und Management deutscher Unternehmen. CO2 Barometer 2010, KfW/ZEW.
- Mansanet-Bataller M, Chevallier J, Herve-Mignucci M, and Alberola E (2011) EUA and sCER phase II price drivers: unveiling the reasons for the existence of the EUA - sCER spread. *Energy Policy*, 39(3): 1056–1069.
- Manski C (1975) Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics*, 3(3): 205–228.
- McCann L, Colby B, Easter W, Kasterine A, and Kuperan K (2005) Transaction cost measurement for evaluating environmental policies. *Ecological economics*, 52(4): 527–542.
- McCrary J (2008) Manipulation of the running variable in the regression discontinuity design: a density test. *Journal of Econometrics*, 142(2): 698–714.
- Montero J.-P (1998) Marketable pollution permits with uncertainty and transaction costs. *Resource and Energy Economics*, 20(1): 27–50.

- Montgomery D (1972) Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3): 395–418.
- Nazifi F (2013) Modelling the price spread between EUA and CER carbon prices. *Energy Policy*, 56: 434–445.
- Sandoff A and Schaad G (2009) Does EU ETS lead to emission reductions through trade? The case of the Swedish emissions trading sector participants. *Energy Policy*, 37(10): 3967–3977.
- Schleich J and Betz R (2004) EU emissions trading and transaction costs for small and medium sized companies. *Intereconomics*, 39(3): 121–123.
- Singh R and Weninger Q (2016) Cap-and-trade under transactions costs and factor irreversibility. *Economic Theory*: forthcoming.
- Stavins R (1995) Transaction costs and tradeable permits. *Journal of Environmental Economics and Management*, 29(2): 133–148.
- Stephan N, Bellassen V, and Alberola E (2014) Use of Kyoto credits by European industrial installations: from an efficient market to a burst bubble. Climate report 43, CDC Climat.
- Stevens B and Rose A (2002) A dynamic analysis of the marketable permits approach to global warming policy: a comparison of spatial and temporal flexibility. *Journal of Environmental Economics and Management*, 44(1): 45–69.
- Tietenberg T (2006) *Emissions trading: principles and practice*: Resources for the Future.
- Train K (2009) *Discrete choice methods with simulation*: Cambridge University Press.
- Trotignon R (2012) In search of the carbon price. The European CO₂ emission trading scheme: from ex ante and ex post analysis to the protection in 2020. Ph.D. dissertation, Paris-Dauphine University.
- Veal G and Mouzas S (2012) Market-based responses to climate change: CO₂ market design versus operation. *Organization Studies*, 33(11): 1589–1616.
- Williams R (2010) Fitting heterogeneous choice models with oglm. *Stata Journal*, 10(4): 540–567.
- Zaklan A (2013) Why do emitters trade carbon permits? Firm-level evidence from the European emission trading scheme. Discussion Paper 1275, Deutsches Institut für Wirtschaftsforschung (DIW).

Appendices

A National offset entitlement rules

Table IV.4: Offset limits from National Allocation Plans

	Annual Cap Ph.II (MMt CO ₂ e)	Offset limit (%)	Annual offset limit (MMt CO ₂ e)	Banking/ Borrow- ing	Industry	Energy	Other sector differ- entia- tion	Included in this chapter
Austria	30.7	10	3.1	Yes/yes				
Belgium	58.5	8.4	4.9	-	Flanders 24% Walloon 4%	Flanders 7% Walloon 8%		
Bulgaria	42.3	12.6	5.3	Yes/yes				
Cyprus	5.48	10	0.5	Yes/yes				
Czech Rep.	86.8	10	8.7	Yes/yes				
Denmark	24.5	17	4.2	Yes/yes	6.50%	28.70%		
Estonia	12.72	10	1.3	No/no	(started only in 2011)			No
Finland	37.6	10	3.8	Yes/Yes	8 / 8.5%	8.5 / 9.5 /23.9%		
France	132.8	13.5	17.9	Yes/Yes				
Germany	453.1	22	99.7	Yes/Yes				
Greece	69.1	9	6.2	Yes/Yes				
Hungary	26.9	10	2.7	-				
Ireland	22.3	10	2.2	Yes/Yes	5%	11%	Cement 11%	
Italy	195.8	15	29.4	Yes/no	7.2%	Electricity 19.3%	"Other" com- bustion 7.2%	
					Ferrous metal 16.7%	Refineries 13.2%		
Latvia	3.43	10	0.3	Yes/Yes				
Lithuania	8.8	20	1.8	No/no				No
Luxembourg	2.5	10	0.3	Yes/Yes				
Malta	2.1	10	0.2	Yes/Yes				No
Netherlands	85.8	10	8.6	Yes/Yes				
Norway		13		Yes/No	13% of actual emissions (rather than allocation)			No
Poland	208.5	10	20.9	Yes/No				
Portugal	34.8	10	3.5	Yes/Yes				
Romania	75.9	10	7.6	Yes/Yes				
Slovakia	30.9	7	2.2	Yes/Yes				
Slovenia	8.3	15.8	1.3	Yes/Yes				
Spain	152.3	20.6	31.4	Yes/No	7.90%	42%		
Sweden	22.8	10	2.3	Yes/Yes				
UK	246.2	8	19.7	Yes/No	8%	9.30%		

Source: Elsworth et al. (2012)

B Offset use as a fixed horizon problem

Offset use was introduced for Phase II of the EU ETS; however, later it was extended to Phase III. This chapter assumes that firms operated under a fixed horizon conjecture, i.e. believed that their offset entitlement ended with the end of Phase II. While my assumption is an approximation, I rely on regulatory evolution and empirical elements to justify this assumption.

The ETS is based on the Kyoto Protocol (to the United Nations Framework Convention on Climate Change) that was signed in 1997, ratified by the EU in May 2002, and became effective in February 2005. It established the possibility to “transfer emission reduction” across geographical regions using flexibility mechanisms, mainly the named offset credits (CER and ERU). Thus, the Kyoto Protocol is the legal basis for the creation and validity of offsets.

The EU ETS was established with the “Emissions Trading Directive”(Directive 2003/87/EC in October 2003). In 2004, the EU ratified the “Linking Directive” (Directive 2004/101/EC in October 2004) as a basis for Phase II (2008-2012). It allowed the use of offset certificates, but left it to Member States to regulate the details. Concrete provisions are introduced “until 31 December 2012”, while nothing points out how regulation will change after the end of Phase II.

After the introduction of offset certificates to the EU ETS by all Member States, Directive 2009/29/EC (June 2009) set out to harmonize offset use across Member States. However, the Kyoto Protocol’s first commitment period ended in 2012 without a follow-up treaty being effective until the time of writing of this chapter in 2017. The 2009 Directive’s preamble clearly states the uncertainty of international climate negotiations:

Once there is an international agreement on climate change, additional use of CERs and ERUs should be provided for[...] In the absence of such an agreement, providing for further use of CERs and ERUs would undermine this incentive.

Moreover, problems with the environmental value-added of certain project-types have become clear and “measures may be applied to restrict the use of specific credits from project types”(Art. 11a(9) of amended Emissions Trading Directive). The 2009 Directive remains elusive however on exactly which offsets will not be usable anymore in Phase III. Commission Regulation (EU) No 550/2011 (June 2011) prohibited so-called industrial gas projects. As it was not clear to certificate holders which certificates fell under this definition, incentives have been particularly strong to submit any purchased offsets before the end of Phase II. As the end of Phase II approached, stakeholders began to worry about the legal foundation of offset use after May 2013. The EU Commission’s publication “Questions and answers on use of international credits in the third trading phase of the EU ETS”³⁴ attempts to reduce uncertainty, but many answers start with “details in this regard will be determined in a forthcoming amendment.” Legal advisory pages

³⁴Published on 14/11/2011 under http://ec.europa.eu/clima/policies/ets/markets/docs/q_a_20111114_en.pdf; retrieved on 05/08/2016.

published cautious warnings about the lack of legal base for offsets beyond the end of the Kyoto Protocol first commitment period (which equaled the end of Phase II).³⁵

Finally, it was not until Commission Regulation (EU) No 1123/2013 in November 2013 that the bankability of offset entitlements was finally confirmed and specified in detail. This regulation was established *after* the end of Phase II; no regulation regarding offset prolongation was effective at the moment when firms surrendered the last certificates for Phase II. Except for new entrants, no new entitlements were created, but operators can use up remaining entitlements using certificates from specific project-types in LDCs. Until then, the bulk of offset certificates had been based on industrial-gas projects from India, China, and Ukraine (Ellerman et al. 2016) and, thus, were suddenly useless.

Anecdotaly, in the compliance data, all of the largest 9% of firms (in terms of emissions) used up their offset entitlement at the end of Phase II. These are the firms have the largest stakes and that were particularly well informed and/or influential, so it would be unlikely that the small firms driving to a large extent the identification of this chapter were better informed about the actual bankability of offset entitlements.

C Exogeneity of allocation status

In Section 3.2, I claim that firms do not strategically constrain their emissions to be just below allocation level, even though firms face a cost jump when emissions increase beyond this level. This assumption is important, as I use the fact that *short* firms, with emissions above allocations, are constrained to trade while *long* firms can choose whether to incur trading costs. This methodology is flawed if *transaction costs lead firms to manipulate their net allocation status*. This section argues that this case is unlikely to be relevant in actual practice. Theoretically, firms choose their production and emissions given production cost and certificate prices; the additional transaction cost is likely to be smaller than the cost of adjusting emissions *and production*. Empirically, there is no significant discontinuity around the net allocation status threshold.

First, note that the firm faces the same marginal cost p^e for emissions both below and above the allocation level, such that marginal abatement cost does not play a role. Thus, the firm compares two situations: one situation where it reduces emissions to allocation level q_i^{e0} , producing $y^*(q_i^{e0})$ *without* incurring entry costs, and another situation where its optimal emission level e_i^* equalizes CO₂ price, the firm buys additional certificates and incurs trading cost. The firm reduces its emissions to q_i^{e0} if the change in profit $\Delta\pi$

³⁵For example <http://www.emissions-euets.com/cers-erus-market-as-from-2013> or <http://ieta.org/the-consequences-of-the-durban-cop-for-the-carbon-market-and-climate-finance>; both retrieved on 05/08/2016.


resulting from this reduction is positive:

$$\begin{aligned} \Delta\pi &= \pi(y^*(q_i^{e0}), q_i^{e0}) - \pi(y^*(e_i^*), e_i^*) \\ &= \underbrace{y^*(q_i^{e0}) - y^*(e_i^*)}_{\Delta_{y^*} < 0} - \underbrace{C(y^*(q_i^{e0}), q_i^{e0}) + C(y^*(e_i^*), e_i^*)}_{\Delta_C \geq 0} - \underbrace{p^e (q_i^{e0} - e_i^*)}_{\Delta_{qe} < 0} + \underbrace{\kappa^e + \min\{\kappa^0 - \bar{q}_i^0 \Delta_p, 0\}}_{> 0: \text{transaction costs}} \end{aligned} \quad (14)$$

$$\underbrace{\hspace{10em}}_{< 0: \text{deviation from (least-cost) optimum}} \quad (15)$$

By assumption, we are looking at cases where optimal emissions $e_i^* > q_i^{e0}$ and, thus, $y^*(e_i^*) > y^*(q_i^{e0})$. By definition of the optimal emission level e_i^* , the first part of $\Delta\pi$ is negative while the transaction cost terms of equation (15) are positive.

A priori this assumption cannot be verified in practice; information on prices, quantities y and production costs are not available, thus neither the cost function $C(y, e)$ nor the profit change $\Delta\pi$ can be estimated. Instead, one way of gathering (descriptive) evidence on this point is to check whether we observe any crowding or “bunching” of emissions just below $e_i = q_i^{e0}$. If firms were manipulating their net allocation status, the distribution of this ratio would be somewhat discontinuous around $e_i/q_i^{e0} = 1$. Figure IV.9 implements McCrary’s test for continuity (McCrary 2008). The estimated densities on the left and on the right of the cut-off where $q_i^{e0} = e_i$ seem smooth on Figure IV.9a: at a discontinuity magnitude of .0116 (in logs) and a standard error of .1133, we cannot reject the hypothesis that there is no bunching around the threshold, or put differently, that the ratio’s density function is continuous around this point. Moreover, restraining emissions to become *long* should be particularly relevant for firms that do not use offset certificates. Therefore Figure IV.9b shows the McCrary test only for the firms that do not trade in the offset market: there is still no significant bunching at $e_i = q_i^{e0}$ (discontinuity estimate at -.3910 with standard error of .2766).



(a) All firms;
discontinuity estimate (in log) .0117, se .1133

(b) Only firms not using offsets;
discontinuity estimate (in log) -.3910, se .2766

Figure IV.9: McCrary’s test for continuity of the running variable (ratio emissions/allocations)

Note: Estimated using Stata DCdensity command by Kovak and McCrary, available under <http://eml.berkeley.edu/~jmccrary/DCdensity/>

Anecdotal and survey evidence (Löschel et al. 2010, 2011) suggests that firms do not have precise and continuous control over their emissions, or rather that there are considerable transaction costs to obtain such control. Only large companies regularly track their emissions throughout the year. The trading scheme's incentives to reduce emissions do not work on a short-term "accurate to the tonne" level, but rather on a long-term technology-inducing level.

Most technologies are such that in the short term the actual technological margin to reduce emissions without a complete corresponding reduction of output is limited; reducing emissions by a certain share is thus equivalent to reducing production by the same share. After all, emissions are just one production cost factor among many others and the short-run flexibility of the cost function is usually low. Emission reductions are mostly accomplished in the long term through technical change, whereas this chapter examines short term behavior. Even for a small difference between e_i^* and q_i^{e0} it is likely that Δ_π is negative.

A notable exception might be emission savings by electricity generating plants, as some firms have scope for fuel-switching across different plants and emission costs are a more important cost factor in this industry (Jong and Zeitlberger 2014). However, the McCrary test also does not show a significant bunching if we are only looking at electricity firms.

While theoretically not fully sound, the assumption of exogenous allocation status thus seems empirically valid and in line with anecdotal evidence.

D Parametric estimation results

In the standard way to estimate the parameters of equation (9), one assumes that error term ϵ_i follows a standard normal distribution. The model then becomes a standard probit model: opportunity cost $\bar{q}_i^o \Delta_p$ is included as a regressor and coefficients are normalized so that the coefficient on \bar{q}_i^o equals -1. The estimation equation reads:

$$\mathbb{1}_i^o = \mathbb{1} \left\{ \beta_0 + \beta_1 \mathbb{1}_i^{long} + \beta_2 \bar{q}_i^o \Delta_p + \epsilon_i > 0 \right\} \quad (16)$$

Standard statistical packages normalize the standard deviation σ to 1. A re-normalization then yields the parameters of interest:³⁶

$$\hat{\kappa}^o = -\frac{\hat{\beta}_0}{\hat{\beta}_2}; \quad \hat{\kappa}^e = -\frac{\hat{\beta}_1}{\hat{\beta}_2}; \quad \hat{\sigma} = \frac{1}{\hat{\beta}_2} \quad (17)$$

The stylized facts presented in Section 4 strongly suggest that this homoskedastic normality assumption does not hold. As shown before, the distribution of offset entitlements is highly skewed with some firms more than 500 times bigger than the median. Some firms with high \bar{q}_i^o still do not exploit their offset entitlement, such that the distribution of ϵ_i from the transaction cost equation (9) is likely to have some large outliers. The (condi-

³⁶Standard errors for the re-arranged parameters are computed using Stata's nlcom command, based on the delta method.

Table IV.5: Parametric mean estimates for transaction costs

	Probit	Heterosk. probit	Probit with sectors	Heterosk. probit with sectors
\hat{T}^o (intercept)	109,557*** (4.24)	102,660*** (4.36)		
\hat{T}^e (1^{long})	44,302*** (3.70)	42,798*** (3.79)		
\hat{T}^o Manufacturing			171,436*** (4.48)	161,416*** (4.63)
\hat{T}^e Manufacturing			96,475*** (4.42)	92,138*** (4.53)
\hat{T}^o Electricity			48,383** (2.58)	278,077*** (4.54)
\hat{T}^e Electricity			4,169 (0.25)	5,065 (0.32)
σ	192,950*** (5.77)	182,835*** (6.04)	192,434*** (5.82)	182,472*** (6.09)
γ		6.96e-08*** (18.15)		6.95e-08*** (18.24)
R^2	.1274	.128	.1372	.1378
Completely determined	371	.	369	.
N	4,578	4,578	4,578	4,578

Note: t statistics in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

tional) mean is a statistic much more sensitive to outliers than the (conditional) median; differently put, the normal distribution assumption has light tails that, consequently, give large weight to outliers.

A slightly more flexible functional form relaxing the homoskedasticity assumption, would be the mixed probit: error terms are still assumed to have a normal distribution, but the variance scales with the size (here \bar{q}_i^o) of the firm. In such a location-scale model, the variance of each ϵ_i depends on some scaling variable and a parameter γ (to be estimated):

$$\epsilon_i \sim \mathcal{N}(0, \sigma_i^2), \text{ where } \sigma_i = \exp(\bar{q}_i^o \gamma) \quad (18)$$

This section shows the results for both assumptions, while claiming that they are not an appropriate description of the data. The results the probit estimation in both the homoskedastic and (linearly) heteroskedastic versions are shown in Table IV.5.³⁷ The costs indicated are measured in euros, as they are normalized by the cut-off value's $\bar{q}_i^o \Delta_p$ coefficient. The estimate for κ^o , the transaction cost for offset usage, exceeds the estimate for κ^e , while both are significant. When I include the sectors, the estimates for transaction costs in the manufacturing sector are much higher than in the electricity and heat generation sector. In particular, the general trading cost, κ^e , seems not relevant for electricity and heat generating firms.

³⁷Estimated using Stata ogln command by Williams (2010).

E Quantile regression fit

Kordas (2006) suggests verifying the fit of the quantile regressions by predicting probability intervals for each observation and verifying that each interval group has an entry rate close to the predicted probability. In order to predict probabilities from the binary quantile regression, one needs to find the smallest quantile $\hat{\tau}_i$ such that the profit-net-of-transaction costs is positive:

$$\hat{\tau}_i = \operatorname{argmin}\{\tau : \bar{q}_i^o \Delta_p - \kappa_\tau^o - \mathbb{1}_i^{\text{long}} \kappa_\tau^e \geq 0\} \quad (19)$$

Then this gives us an interval for the conditional entry probability:

$$\hat{P}_i \in [1 - \hat{\tau}_i, 1 - \hat{\tau}_{i-1}] \quad (20)$$

where $\hat{\tau}_{i-1}$ is the quantile immediately preceding $\hat{\tau}_i$.

For the data used in this chapter, this provides the predicted and observed probabilities displayed in Table IV.6. Except for the lowest quantile, the models seem to fit the data reasonably well. On the opposite, the probit model predicts for *all* firms an entry probability above 50%: one could say that *all* non-participating firms are unpredicted outliers (false-negatives) with the probit model.

Table IV.6: Specification test of binary regression quantile models (predicted and observed probabilities)

Predicted probability	<15%	[15-25%]	[25-35%]	[35-45%]	[45-55%]	[55-65%]	[65-75%]	[75-85%]	>85%
Number of observations	85	130	65	49	153	414	613	971	2,098
Observed frequency	11%	17%	32%	43%	46%	58%	72%	81%	94%

V

Does the EU ETS cause carbon leakage in European manufacturing?

I got 99 problems but [this] ain't one.

after *Jay-Z*

1 Introduction

Climate change caused by CO₂ emissions is a global problem, but efforts to reduce CO₂ emissions are mostly regional. In Europe and some states in the U.S., for example, policy initiatives exist, but no binding international agreement is in force.¹ Unilateral, geographically limited policies increase production costs for domestic producers who compete internationally with producers from unregulated regions. This asymmetry raises the fear of carbon leakage – a shift of CO₂ emissions from a region with emission constraints to an unconstrained area, via a change in relative competitiveness in an open global economy. Carbon leakage is a concern both in the academic debate and in policy circles (Ellerman et al. 2016). As climate change depends on aggregate global emissions, carbon leakage threatens to undo the effects of unilateral policy efforts.² If carbon leakage occurs, the region implementing the policy suffers from a decrease in output and a consequent loss in employment and welfare, additionally to an ineffective environmental policy. The issue is particularly salient when manufacturing sectors are affected by an emissions policy, as they often produce goods that are both carbon intensive and heavily traded.

In this paper, we ask whether the EU ETS, the most important unilateral emissions policy to date, has caused carbon leakage in European manufacturing sectors. More specifically, we test if (parts of) the evolution of sectoral trade intensities can be explained by the stringency of environmental policy. Our empirical analysis in this paper does not provide any evidence of carbon leakage.

In theory, carbon leakage occurs between a domestic region featuring an emissions policy and a foreign region with no policy or a less stringent policy. It results from the combination of two effects: (i) relocation, when domestic firms shift their production to foreign countries to evade the increased production cost imposed by the environmental policy; and (ii) changes in market shares, when domestic firms lose market share to unregulated foreign competitors, who become more competitive as they do not have to bear the additional cost burden.³ Both effects translate directly into trade flows: for a given level of domestic consumption of a carbon-intensive product, carbon leakage

¹The 2015 Paris agreement is neither binding nor does it involve symmetric compliance costs.

²Carbon leakage is a case of the pollution haven effect – which has hitherto mainly been considered in the context of local pollutants (Ederington et al. 2005, Levinson and Taylor 2008) – applied to the global pollutant carbon dioxide. The pollution haven hypothesis states that polluting industries relocate to where pollution is cheap. With local pollutants and a pollution haven effect, the pollution is at least being *displaced*, i.e. the region implementing an environmental policy benefits from less local pollution in exchange for a loss in industrial production. With a global pollutant, carbon leakage *undoes* either part or all of the policy's mitigation effect, depending on the rate of leakage.

³Additionally, carbon leakage can also occur through a drop in the price of emission intensive commodities, usually fossil fuels whose prices are formed globally, due to a fall in global demand for these commodities as a result of the domestic environmental policy. Lower global energy prices may lead to an increase in the demand for fuels in the foreign region, leading to increased energy use there and, thus, carbon leakage (Harstad 2012, Jensen et al. 2015). We do not address this energy price channel of carbon leakage in this paper. However, we believe that to date it is of minor relevance in the case of the EU, the focus of our empirical analysis, as neither its share in global energy demand nor the stringency of its emissions policy are significant enough to materially affect prices in global energy markets.

leads to a higher share of imports in total consumption of the home region and to lower exports.

In practice, the case for carbon leakage is not clear cut. First, the difference in emission cost between Europe and emerging economies has so far been moderate, in particular relative to differences in labor cost. Labor unit cost in Europe is about 10 to 30 times higher than in emerging countries (Schröder 2016). Even though the emission cost is typically zero in other parts of the world, our data show that the emission cost imposed by the EU ETS is below 0.65% of material cost for 95 percent of European manufacturing sectors. Thus, the additional cost introduced by European emissions policy is comparatively small. Second, firms relocating production to a foreign region must pay fixed relocation costs. Relocation also has opportunity costs in the home market, such as a weaker market position and less influence in bargaining with policy makers. Third, emissions policies often combine costs and subsidies. For example, European manufacturing firms received large amounts of free emissions allowances (“free allocation”), which may be sufficient to counter the leakage risk (EU 2014, Schmidt and Heitzig 2014).⁴ Our data reveal that most sectors received a net subsidy from emissions trading, once free allocation is taken into account. Fourth, the business literature predicts an inverse effect of environmental regulation (Porter hypothesis): the negative competitiveness effects of unilateral environmental policy may be offset by successful incentives to innovate in lower-carbon products, spurring a broader productivity increase for firms affected by environmental policies (Porter and Van der Linde 1995). Innovation may be incentivized through the emission price signal (Calel and Dechezleprêtre 2016) or by providing explicit R&D subsidies in parallel (Acemoglu et al. 2012, Aghion et al. 2016).

Our empirical analysis is based on the argument that leakage can be measured through *changes* in trade flows, as they include both leakage channels: production relocation away from Europe and loss of European firms’ market shares. We create a dataset of global trade flows, emission costs and control variables by combining data from the Global Trade Analysis Project (GTAP) with data from the EU’s Transaction Log (EUTL), the administration’s repository of data on emissions, allocations of allowances and transactions in the EU ETS. While GTAP is frequently used for research on computable general equilibrium models, it has recently also been used for empirical research on international trade (Caron et al. 2014). Brunnermeier and Levinson (2004) stress the importance of using panel data, as we do in this study, to account for unobserved heterogeneity of sectors and trading partners. We estimate the effect of four potential measures of the EU ETS’s stringency on trade flows in European manufacturing. Our measures of policy stringency account for both direct and indirect emission costs. Indirect emission cost arises from electricity use: industrial consumers of electricity pay at least part of the costs of embodied emissions, as power producers pass through their emission cost to

⁴Free allocation of emissions allowances based on historical emission levels is an expensive measure to counter carbon leakage: in Phase II of the EU ETS (2008-2012), each year the regulator distributed close to 2 billion tCO₂e allowances for free, which at the average 2012 price of 10.42€/tCO₂e amounts to a yearly opportunity cost of €20.84 billion for free allocation compared to full auctioning.

wholesale prices of electricity (Fabra and Reguant 2014, Hintermann 2016). We use two measures of trade: first, we compute CO₂ emissions embodied in the traded goods, and second we use trade value in U.S. dollars. Flows in embodied CO₂ are computed from input-output tables and measure the emissions necessary to produce the traded goods. Trade flows in embodied CO₂ emissions are often not available, but they capture carbon leakage better than trade flows in value. In our analysis we follow two approaches suggested by the literature: a traditional approach focusing on net imports (Ederington et al. 2005, Levinson and Taylor 2008) and an approach in the spirit of New trade theory where we evaluate bilateral (two-way) trade flows (Aichele and Felbermayr 2012, 2015).

We find no evidence for carbon leakage in European manufacturing sectors. This result contrasts with predictions from ex ante modeling exercises, but is largely in line with findings from existing empirical research on the carbon leakage hypothesis in the context of the EU ETS.

Given the policy relevance of the leakage issue, a sizable literature, mostly based on ex ante computable general equilibrium (CGE) models, has attempted to predict the extent of carbon leakage from existing policy initiatives and potential modifications (as reviewed by Branger and Quirion 2014, Carbone and Rivers 2017). These ex ante approaches predict strong carbon leakage with leakage rates between 10% and 30% (Carbone and Rivers 2017, IPCC 2007).⁵ However, the predictions of ex ante approaches depend on model assumptions, e.g. whether the model includes relocation costs, and the implementation details of the considered emissions policy. Demailly and Quirion (2006) show that introducing output-based allocation in the EU ETS would eliminate leakage, at the cost of decreasing the incentive for producers to abate emissions. Gerlagh and Kuik (2014) show that allowing for technology spill-overs may even lead to carbon leakage from foreign countries into the EU.

Empirical ex post evidence on carbon leakage is limited. Much of the existing empirical literature considers the pollution haven effect in the U.S., i.e. the effect of increasing the stringency of local pollution regulation on trade flows. These contributions typically test for a link between net trade flows and the stringency of pollution control measures, as captured by the Pollution Abatement Cost (PAC) using survey data of U.S. manufacturers.⁶ The evidence in this literature is mixed. Jaffe et al. (1995) review the early contributions, and conclude that there is little evidence that environmental policy has affected trade flows; like other authors, they point to the relatively small magnitudes of environmental expenditures as an explanation. Dechezleprêtre and Sato (2017) review the more recent literature and conclude that there is some evidence in favor of the pollution haven hypothesis, even if the cost burden is small. In particular, Ederington et al. (2005) and Levinson and Taylor (2008) regress U.S. net imports on PAC and find that

⁵Carbon leakage is usually quantified as the ratio of foreign emission increase over domestic emission reduction. If all domestic emission reduction from environmental policy is shifted abroad, carbon leakage is said to be 100%.

⁶As the PAC survey encompasses a wide mix of environmental policies, this literature cannot attribute effects to a specific policy measure.

environmental policy did impact U.S. trade flows. Aichele and Felbermayr (2015) find a carbon leakage effect of the Kyoto protocol. Based on a “gravity model for carbon” they find that the carbon content of sector-level bilateral trade was significantly impacted by a country’s ratification of the Kyoto Protocol. However, it remains unclear through which channel the Kyoto protocol has induced this effect.

To our knowledge, the carbon leakage hypothesis in the EU ETS has so far not been comprehensively evaluated empirically. Some research addresses the relocation channel: Dechezleprêtre et al. (2014) use a survey of multinational firms and find no evidence that the EU ETS induced the relocation of emission-intensive processes *within* multinational firms. Other research addresses the investment channel: using firm-level data on foreign direct investment (FDI) by German multinational companies, Koch and Basse Mama (2016) find no evidence that the EU ETS has contributed to relocation through an increase in outbound FDI. Martin et al. (2014) conduct a survey of managers; they find that relocation risk is limited and that the current EU ETS rules largely over-compensate many sectors given the small risk of relocation. Finally, one strand of literature examines trade flows in specific sectors: Sartor (2013) finds that the EU ETS has not caused carbon leakage in the aluminum sector, while Branger et al. (2016) find no leakage in the cement and steel sector.

We contribute to the literature in several ways: first, we assess both the relocation and the competitiveness impact of the EU ETS by using global sector-level trade data. This approach complements studies focusing on relocation using firm-level data, e.g. Martin et al. (2014) or Dechezleprêtre et al. (2014). Second, using a broader dataset and focusing on a particular policy initiative whose cost can be captured explicitly, we complement previous work on carbon leakage effects of unilateral climate policy. Third, the input-output information in our data allows us to consider all embodied emissions in our outcome variable (trade flows) and our policy variable (emission policy), i.e. both direct and indirect emissions from electricity use.

In the following, we first review the relevant trade theory in Section 2 and then present our empirical implementation in Section 3. This is followed by a description of the data in Section 4 and presentation of results in Section 5. We summarize and conclude in Section 6.

2 Trade theory and carbon leakage

2.1 (Neo-)classical approach

Classical and neo-classical models rely on Ricardo’s theory of comparative advantage, formalized in the Heckscher-Ohlin-Vanek/Samuelson (HOV) model of international trade. In this view of the world, countries are characterized by their unequal endowment of relatively immobile production factors (land, labor), while sectors differ in their factor-intensities and exhibit constant or decreasing returns to scale. A country has a comparative advantage and will specialize in those sectors that are intensive in its relatively abundant factor. Trade in goods essentially amounts to trading bundles of factor




Figure V.1: Stylized illustration of the pollution haven hypothesis

inputs, such that trade equalizes factor prices across countries.

Pethig (1976) establishes the link between a classical HOV model and the pollution haven hypothesis: emissions can be seen as a production factor, and countries with loose emission regulation are more abundant in this factor.⁷ This allows us to directly apply the findings of the general HOV model to the effect of an emissions policy: countries with high emission costs specialize in low-emission sectors and trade leads to equal pollution cost across countries in equilibrium (similar results can be found in Copeland and Taylor 2004, Motta and Thisse 1994, McGuire 1982). Antweiler et al. (2001) decompose the effect of trade liberalization on pollution into composition, scale and technique effects. Copeland and Taylor (2005) show in a three-region model, that the Kyoto protocol may either increase or decrease world pollution, depending on the model setup.

In order to better understand the concept of carbon leakage, it is useful to consider a stylized illustration of the pollution haven problem, assuming a one-sector economy. Figure V.1 illustrates the case of a homogeneous good, immobile production factors and a large country in a neo-classical model. Without an environmental policy, the country produces Y units and consumes C ; the difference between Y and C is imported. When emissions become costly, e.g. through the introduction of an emissions tax t , the supply curve shifts upwards by Δt and the new level of domestic production is Y' . Consumption does not change, while imports increase. If production is equally emission-intensive everywhere in the world, then the total domestic emission reduction is entirely replaced by an increase in foreign emissions and the total effect for global emission mitigation is zero, i.e. carbon leakage is 100%.⁸

In classical models, unilateral environmental policy unambiguously decreases quantities in the regulated country. However, the marginal effect on revenues (prices times

⁷Pethig (1976) assumes emissions enter through a Cobb-Douglas production function. While a Cobb-Douglas production function per se is a restrictive assumption, Levinson and Taylor (2008) show that this is equivalent to a situation where (a.) firms abate optimally given stringency of environmental policy and (b.) pollution abatement cost can be measured as a fraction of total factor use.

⁸In a more nuanced model, substitution between domestic and foreign products is not perfect, there exist trading costs, and technology is not fixed.

quantities) is not always determined: McGuire (1982) shows that environmental policy drives the regulated country out of production of emission-intensive goods entirely, if factors are mobile (unambiguously reducing exports/increasing imports; similarly to Pethig 1976); whereas it merely breaks factor price equalization and changes world commodity prices, if factors are immobile. In the latter case, revenues decrease if the country is small, but revenues can increase or decrease (not determined) if the country is large. This is why we mainly rely on regressions using trade flows in emissions (quantities) and only add trade flows in values (revenues) for completeness. Measures of value are commonly used in the empirical trade literature, even though it is impossible to disentangle price and quantity effects (De Loecker et al. 2016).

(Neo-)classical models are criticized because they fail to explain that countries simultaneously import and export the same commodity with the same trading partner (called intra-industry trade or two-way trade), which empirically accounts for a sizable share of total trade flows.⁹ As a consequence, empirical work based on (neo-)classical trade models focuses on net trade flows, i.e. on the difference between imports and exports for each trading partner. We follow this approach in the first part of our empirical analysis.

2.2 New trade approach

More recently, the literature on trade theory has turned to New trade theory (Dixit and Stiglitz 1977, Krugman 1980), and “New” new trade theory focusing on heterogeneous firms (Melitz 2003). New trade models typically assume increasing returns to scale, providing a reason for specialization beyond initial factor endowments. Models typically assume a CES utility function, monopolistic competition and consider trade in intermediaries. Equilibrium flows of bilateral trade then depend on the market capacity of the importer and supply capacity of the exporter, as well as sectoral demand elasticities and trade costs. Dixit-Krugman-style models are used to derive a theoretical foundation for the gravity equation (Head and Mayer 2014). By assuming a “love of variety”, New trade theory helps explaining the existence of intra-industry trade. A central result is that the representative consumer spreads consumption evenly over differentiated goods within a sector.

The (neo-)classical argument behind the pollution haven hypothesis can be seen as a competitiveness effect, arising from the full cost pass-through of firms in perfect competition. The impact of environmental regulation on trade flows in New trade approaches is more complex: on the supply side, the cost of emissions enters through higher input prices and reduces quantities and the equilibrium number of firms (and thus produced varieties) in the regulated region, therefore increasing imports and decreasing exports. On the demand side, the policy can impact the domestic price index, which makes the regulated region relatively poorer and dampens the increase in imports by reducing overall consumed quantities. In an application, Aichele and Felbermayr (2015) use a Dixit-Krugman-style model to analyze the impact of the ratification of the Kyoto Protocol on

⁹Moreover, HOV models typically predict more trade than what is found empirically.

CO₂ embodied¹⁰ in trade flows.¹¹

Within a New trade model based on monopolistic competition, the effects of emissions policy both on quantity and on revenue are negative.¹² Empirical applications of New trade models use bilateral trade data, i.e. imports and exports are separate observations. We pursue an analogous approach in the second part of our empirical analysis. Classical models typically feature neither horizontal nor vertical differentiation. New trade models account for horizontal but not vertical differentiation. Throughout this study, we assume that there is only horizontal differentiation within sectors.

3 Empirical Implementation

3.1 Measures of environmental stringency

Following Jaffe et al. (1995) and Brunel and Levinson (2016), we note that there are many possibly ways to measure environmental stringency. Depending on the policy implemented, compliance costs are the sum of costs of abating emissions and cost of remaining emissions. In the case of command-and-control policies, affected firms only pay abatement costs, whereas with a carbon pricing scheme they bear the costs of both abating emissions and paying for remaining emissions. The compliance cost of any policy can be offset through direct transfers to the affected firms.

Much of the empirical literature on the pollution haven effect in the U.S. uses data on the Pollution Abatement Cost (PAC) (e.g. Tobey 1990, Grossman and Krueger 1991, Ederington et al. 2005, Levinson and Taylor 2008). PAC is a summary measure of firms' expenditures on the abatement of local pollutants across a range of policies, based on survey data.¹³ Abatement cost is a reasonable measure for total compliance cost when

¹⁰Embodied CO₂ is computed from input-output tables and measures the emissions necessary to produce the traded goods.

¹¹However, New trade models tend to quickly get intractable. In order to apply the model empirically, Aichele and Felbermayr (2015) simplify by accounting only for trade between two regions, regulated and unregulated, and by dropping trade in intermediary goods.

¹²A simple example of a firm maximizing its profit π shows that revenue (product of price and quantity) decreases with an increase in environmental regulation (i.e. a reduction in emissions e). With quantity $q(p)$ and cost function $c(q(p), e)$ both continuous and twice differentiable, and partial derivatives $q_p < 0$, $c_q > 0$, $c_e < 0$ and $c_{qe} < 0$, we have:

$$\begin{aligned}\pi &= pq(p) - c(q(p), e) \\ \text{firm's FOC } p^*(e) : \frac{\partial \pi}{\partial p} &= pq_p + q - c_q q_p = 0 \\ \text{thus : sign } p_e^* &= \text{sign } c_{qe} < 0 \\ \frac{\partial p^* q(p^*)}{\partial e} &= p_e^* q(p^*) + p^* q_p(p^*) p_e^* = 0 \\ &= \underbrace{p_e^*}_{<0} \underbrace{(q(p^*) + p^*(e) q_p(p^*))}_{<0 \text{ from FOC } p^*(e)} > 0\end{aligned}$$

¹³PAC data have only been collected for the U.S. and the data series was discontinued after 2005.

studying command-and-control policies.

We argue that emission costs are a more appropriate measure of environmental policy stringency when studying an emissions trading scheme; not only because the other element of total compliance cost, that is abatement cost, remains unobservable to the econometrician in the absence of a survey. The available literature finds that abatement in manufacturing sectors due to the EU ETS during the period covered by this paper was modest (Martin et al. 2016), so that the emission cost constitutes the main share of compliance cost.¹⁴ In practice, the emission cost imposed on sectors by the EU ETS is likely to be more precisely measured than PAC: it is based on administrative data reflecting the entire population of production plants regulated under the EU ETS, avoiding potential selection bias and response biases from a voluntary firm survey. Moreover, dealing with one policy only instead of a summary measure as in the case of PAC facilitates the attribution of causal effects.

We suggest several measures of the stringency of environmental policy: a binary treatment indicator θ^1 , and continuous measures of the components of emission cost (direct θ^d , indirect θ^i , and allocation θ^a).

- $\theta_{ist}^1 = 1$ if the sector's activity is explicitly regulated under the EU ETS, and 0 otherwise.¹⁵ The dummy variable $\theta_{ist}^1 = 1$ indicates that producers in sector s of country i are required to participate in the EU ETS's compliance mechanism in year t . In addition to greater policy stringency, the binary indicator might capture transaction costs from being included in the scheme more broadly, such as annual verification of emissions and surrender of allowances.
- $\theta_{ist}^d = P_t^e e_{ist}$, where P_t^e is the allowance price and e_{ist} are the sector's direct emissions covered by the EU ETS. θ_{ist}^d captures the *direct* emission cost imposed by the EU ETS on sector s .¹⁶
- $\theta_{ist}^i = P_t^e elec_{ist}$, where $elec_{ist}$ is the amount of emissions embodied in the sector's consumption of electricity, calculated from input-output data. θ_{ist}^i captures the *indirect* emission cost, as allowance prices are passed through to prices of electricity, so that manufacturers ultimately pay for CO₂ emitted in electricity production (Fabra and Reguant 2014, Hintermann 2016).¹⁷

¹⁴Moreover, the abatement cost – if there is any abatement – should be highly correlated to emission stringency and, thus, emission costs, such that our measure is at least a good proxy for environmental stringency.

¹⁵The targeted sectors are: cement; chemicals, rubber, plastic prods; iron and steel; metal products; paper products; petroleum and coal products; other metals; other minerals (which includes glass and ceramics).

¹⁶In addition to the sector activities included explicitly, secondary activities are included in all sectors, usually in-house electricity generation through combustion installations. θ^d captures all these costs, while $\theta^1 = 0$ in many sectors.

¹⁷Fabra and Reguant (2014) and Hintermann (2016) find that power producers pass through their emission costs to electricity wholesale prices fully. In case pass-through is less than complete, our measure of indirect emission cost constitutes an upper bound.

- $\theta_{ist}^a = P_t^a a_{ist}$, where a_{ist} is the amount of allowances freely allocated to the sector s . θ_{ist}^a captures the lump-sum *subsidy* that is part of the EU ETS; it is not a cost, but a benefit.
- $\theta_{ist}^{tot} = \theta_{ist}^d + \theta_{ist}^i - \theta_{ist}^a$, the total net cost of the EU ETS.

Following a suggestion by Ederington et al. (2005), we normalize these emission cost measures by the sector-level material cost,¹⁸ to account for environmental stringency while eliminating absolute magnitude effects (cf. Appendix A on page 157).

3.2 Identification

When regressing trade flows on environmental stringency, it is important to consider endogeneity concerns and potential omitted variable bias. We also discuss in this section what assumptions are necessary about unobserved foreign emission costs.

We control for unobserved sector heterogeneity by including industry-country and time fixed effects. In the following, we go through the elements of our definition of environmental stringency: the dummy, the emission levels, the allowance price, and the allocation, in order to consider whether remaining variation causes endogeneity of θ .

First, the binary treatment indicator θ^1 indicates that the EU explicitly targets a sector for its primary activity. Did the regulator select sectors for inclusion under the EU ETS based on their leakage risk or trade intensity? Our data indicate that the covered sectors are those with the largest historical emissions, which are determined by their production technology, not by their leakage risk.

Second, a similar argument applies to sectoral emission intensities as included in the continuous stringency measures. Emission levels depend on produced quantities, but we normalize by material cost to obtain emission intensities. We assume that emission intensity results from sector-specific technology, which is fixed in the short term and independent of import intensity. If we did not normalize by dividing through material cost, the common correlation of imports, exports and emissions with produced quantities would lead to spurious correlation. In Appendix A on page 157, we verify that we do not induce a bias by using normalized variables.

Third, allowance prices cannot depend on trade flows at the sector level. This is unlikely to be the case, as none of the manufacturing sectors had emissions large enough to substantially influence the price of CO₂ allowances. In fact, the majority of demand for CO₂ allowances comes from the electricity sector, with over 60% of total emissions in the EUTL in Phase II.¹⁹

Fourth, the definition of θ_{ist}^a and of θ_{ist}^{tot} includes free allocation of emission allowances, which the regulator has explicitly introduced to mitigate the risk of carbon leakage. However, the EU distributed free allocations to *all* sectors in our sample, as they were

¹⁸Alternatively, one could normalize by output, but the correlation between output and material cost is close to one, so that this choice is not relevant in practice.

¹⁹The electricity sector is not included in our analysis directly, as electricity is not traded globally.

all deemed to be at risk of carbon leakage. The level of free allocation is proportional to historical emissions (EU 2014), thus exogenous once we account for industry fixed effects.

Consequently, the risk of endogeneity seems limited. However, there may be omitted variables that drive both trade flows and environmental stringency (energy prices) or that modulate the strength of carbon leakage (transport costs).

Energy input prices are linked to both right-hand and left-hand variables of our regression. In one direction, causality seems excluded: energy prices are determined in the global market, and the impact of the EU ETS on global petrol, coal and gas prices is negligible. In the other direction, increasing energy prices will decrease both trade flows in CO₂-intensive goods, as the rising input costs make them more expensive, and CO₂ allowance prices, as with declining production producers of CO₂-intensive goods demand fewer allowances. This may bias our estimate of carbon leakage upwards. As the central result of our paper is that we do not find any significant carbon leakage effect, this actually strengthens our conclusions.

Trade costs, in particular tariffs and transportation costs, affect how easily a product is traded and, in equilibrium, influence the “home bias.” Consequently, sectors with high transport costs are naturally sheltered from foreign competition, reducing the risk of carbon leakage. An identification problem arises if, as argued by Ederington et al. (2005), there is a positive correlation between transport costs and carbon intensity. If transport costs are particularly high for emission-intensive sectors that also have high emission cost, this would bias our estimate towards zero. To control for this effect, we explicitly include transport cost in all our regressions and perform a robustness test using the interaction of our measures of environmental stringency and transport costs (cf. Section 5.1.2).

Finally, we do not include data on emission policies other than the EU ETS. Therefore, our estimates relate the change in European emission policy to changes in trade flows, taking all other emission policy as *given*. To our knowledge, the only major emission policy during the period 2004–2011 is the Kyoto Protocol: Kyoto signatory countries pledged to reduce emissions or otherwise purchase Kyoto allowances at the *country* level. However, producers from Kyoto signatory countries outside the EU did not face emission costs at the *sector level*.²⁰ In some regions, emissions control policies similar to the EU ETS were introduced after 2011, the final year in our sample, e.g. in California, Quebec and at the provincial level in China.

3.3 Net flows

Following the literature, we examine the data from two angles. First, we consider net trade flows as in the classical approach, i.e. the difference of imports and exports at sector-country-year level (this subsection). Then, we analyze bilateral trade flows, including (two-way) intra-industry trade, at the sector-source-destination-year level in the spirit of

²⁰These country-level emission reductions were easily achieved in most cases, either because of generous targets, e.g. in Russia, or due to emission reductions caused by lower production during the economic crisis that started at the end of the last decade.

New trade theory (Subsection 3.4).

In the vein of Ederington et al. (2005), we estimate the following equation on net trade flows:

$$y_{xst} = \alpha\theta_{st} + \beta\tau_{st} + \gamma F_{st} + \delta t_{st} + \nu_t + \nu_{xs} + \epsilon_{xst} \quad (1)$$

where y_{xst} are the net imports – in value or in embodied carbon – of the EU from sector s and country x in year t . θ_{st} is either the ETS dummy variable θ_{st}^1 , the total net ETS cost θ_{ist}^{tot} , or the vector of emission cost components $[\theta_{st}^d, \theta_{st}^i, \theta_{st}^a]$. τ_{st} is the EU's average import tariff for goods from sector s . ν_t are year fixed effects, ν_{xs} are sector-source country fixed effects. F_{st} is a vector of sector-level factor payments to unskilled labor, skilled labor and capital in percentage of total value added; the factor payment to skilled labor is omitted as the three add up to 1.²¹ t_{st} are transportation costs between source and destination countries, normalized by the free-on-board (FOB) value of trade flows. ϵ_{xst} is an error term. Following another suggestion by Ederington et al. (2005), we normalize trade flows by a sector's total output, in order to compare outcomes of similar magnitude.

In classical theory, the effect of emission policy on net imports y_{xst} in embodied carbon is unambiguous, but it not always clear for net imports in value. If the EU ETS caused carbon leakage, the coefficient of environmental policy stringency θ_{st} is positive: more stringent policies, i.e. a higher emission cost, decrease carbon exports and/or increase carbon imports, which both translate into higher net imports of embodied carbon.

Year fixed effects control for general business cycles that are not sector-specific and sector-country fixed effects control for partner country size, sectoral specialization and distance to the EU. Our parameter of interest α is identified from the correlation of environmental stringency to *within* sector-country changes beyond the overall business cycle (difference-in-differences). The hypothesis is that increases in net imports should correlate with the stringency of environmental policy; in particular, for some sectors environmental stringency is negligible, so that there is no reason for carbon leakage in these sectors.²²

3.4 Bilateral flows

Relying on a New trade model, Aichele and Felbermayr (2015) use bilateral flow data in traded value and in embodied CO₂ emissions to test for carbon leakage. In this spirit, we

²¹Value added is distributed to unskilled labor, skilled labor and capital. We include factor payments in order to replicate the methodology in Ederington et al. (2005): they argue that including factor payments is not a valid test of the HOV model, but that they are still valid industry control variables. For robustness, we include the same regression without factor payments in the Appendix B.1 on page 159.

²²Indeed, no sector is completely protected from emission costs, as all sectors use at least some electricity. However, sectors like electronic equipment or wearing apparel have measured environmental stringency close to zero (total cost impacts of less than 0.04% of material cost). Our method only identifies sector-specific variation, i.e. if there is a leakage component common to all sectors, it will be filtered out by our fixed effects.

estimate the following equation:

$$y_{xmst} = \alpha^m \theta_{mst} + \alpha^x \theta_{xst} + \beta \tau_{mst} + \gamma F_{mst} + \delta t_{mst} + \nu_{mt} + \nu_{xt} + \nu_{st} + \nu_{mxs} + \epsilon_{mxst} \quad (2)$$

where y_{xmst} is the trade flow – in value or in embodied carbon – from country x to country m in sector s and year t . θ_{mst} is either the ETS dummy variable θ_{mst}^1 , the total net ETS cost θ_{mst}^{tot} , or the vector of emission cost components $[\theta_{mst}^d, \theta_{mst}^i, \theta_{mst}^a]$, of the importer m , and θ_{xst} is the analogously defined variable for the exporter x . τ_{mst} is the destination country's average import tariff for goods of sector s . ν_{mt} and ν_{xt} are country-year fixed effects capturing business cycles at the national level. ν_{st} are sector-year fixed effects capturing global shocks at the sector level. ν_{mxs} are sector-country pair fixed effects capturing sector-specific differences in trade intensity between two trading partners.²³ t_{mst} are transportation costs between source and destination countries, normalized by the FOB value of trade flows. ϵ_{mxst} is an error term.

If the EU ETS caused carbon leakage, the effect of emission policy stringency θ_{mst} (importer) on y_{xmst} is positive and/or the effect of θ_{xst} (exporter) is negative: if a sector underlies more stringent environmental policy and suffers from leakage, then its exports decrease and imports increase. In New trade theory, the effect is unambiguous both for trade flows y_{xmst} in value and in embodied carbon.

Note that Aichele and Felbermayr (2015) define their treatment variable as the difference between ratification status with respect to Kyoto in the importing and in the exporting country: $\theta_{mxst}^1 = (\theta_{mst}^1 - \theta_{xst}^1)$. This is equivalent to constraining the parameters in equation (2) such that $\alpha^m = -\alpha^x$. In addition to the model of equation (2), we also include this specification.

The sectoral business cycle is captured by the sector-year fixed effects.²⁴ As typical in gravity-type estimations, the country-year fixed effects account for country size in the sense of supply capacity and market size that might fluctuate beyond global business cycles. Destination-source-sector fixed effects finally capture national specializations, institutional trade proximity and distance between both countries, i.e. factors that are pair-specific but do not fluctuate. Our parameters of interest α^m and α^x are then identified from the within sector-country-pair *changes* in trade flows beyond general trends and their correlation with changes in environmental stringency.

²³We do not include factor payments in the main regression, as this does not fit with New trade models. For robustness, we include the same regression with factor payments in the Appendix B.2 on page 160; this leads us to the same conclusions as our main specification.

²⁴Note that for each sector-year, we have over 4,000 observations of which around 1,500 are trade flows coming from EU countries.

4 Data and descriptives

4.1 Data

We use two main sources of data, the Global Trade Analysis Project (GTAP) version 9.2,²⁵ and the EU Transaction Log (EUTL).²⁶ We draw data on trade flows, CO₂ emissions, factor payments, transport costs, output, and material costs for the years 2004, 2007, and 2011 from the GTAP database. The EU ETS was introduced in 2005, so that we have one period prior to the policy introduction and two periods after. GTAP 9.2 data are divided into 57 sectors and 140 countries. We aggregate smaller economies into regions, resulting in a dataset of 66 regions. We only keep the manufacturing sectors (25 out of 57 sectors), which are at the heart of the carbon leakage debate.

The major benefit of GTAP is that it offers consistent data at the global level and includes input-output (I-O) information. This allows us to fully account for emissions from both electricity and fossil fuel inputs.²⁷ The I-O data also allow us to compute emissions embodied in the electricity consumed by each sector, i.e. indirect emissions and their cost.

We use data from the EU Transaction Log (EUTL) to compute our measures of policy stringency. The EUTL is an administrative dataset containing official yearly compliance data for all production plants regulated under the EU ETS, starting in 2005. We extract data on emissions and allocations from the EUTL and map them to the 4-digit NACE 2 code using a plant-to-NACE matching provided by the European Commission and compiled as part of the Ownership Links and Enhanced EUTL Project.²⁸ We combine the EUTL data with GTAP data via the International Standard Industrial Classification (ISIC, a UN nomenclature); for this, we match the GTAP classification to ISIC following Huff et al. (2000) and the NACE level EUTL data to ISIC using correspondence tables from Eurostat.²⁹ Finally, we add allowance prices for EU ETS emissions allowances (EUAs) from the European Energy Exchange (EEX).

4.2 Descriptive statistics

In 2004, the year before the EU ETS was introduced, no firms based in Europe were liable for CO₂ emissions, so the cost of embodied carbon was zero for all sectors. θ^1 indicates that 32% of the EU's manufacturing sectors (8 out of 25) were directly targeted by the EU ETS (Table V.1).

Carbon leakage is a medium to long-term phenomenon, so we choose average al-

²⁵See <https://www.gtap.agecon.purdue.edu/databases/v9/> and Aguiar et al. (2016) for further details.

²⁶<http://ec.europa.eu/environment/ets>; retrieved on 02/05/2017.

²⁷Our aim is to capture total emissions, both from fossil fuels and process emissions. As the GTAP data only contain information on emissions from the use of fossil fuels, we correct for process emissions in sectors featuring a significant share of process emissions, i.e. iron and steel, cement and chemicals, using data from the UNFCCC (<http://unfccc.int>).

²⁸<http://fsr.eui.eu/climate/ownership-links-enhanced-eutl-dataset-project/>; retrieved on 05/08/2016.

²⁹<http://ec.europa.eu/eurostat/ramon/index.cfm>; retrieved on 06/09/2016.

Table V.1: Descriptive statistics

	Mean	Median	SD	Min	Max	N
Measures of environmental stringency (2007&2011, only EU countries)						
ETS dummy θ_{ist}^1	31.9%	0	46.6%	0	1	1,247
Direct ETS cost θ_{ist}^d	0.17%	0.00%	0.629%	0	6.65%	1,247
Indirect ETS cost θ_{ist}^i	0.16%	0.07%	0.332%	0	4.37%	1,247
Allocation benefit θ_{ist}^a	0.22%	0.00%	0.845%	0	10.06%	1,247
Total net ETS cost θ_{ist}^{tot}	0.11%	0.05%	0.442%	0	6.81%	1,247
EU net import flows						
Net imports (Mil. US\$)	-146.4	-38.0	3,912.9	-36,892.5	84,150.2	3,075
Net embodied CO ₂ imports (Mt)	0.3	0.0	1.5	-4.7	27.0	3,075
Net imports/output	0.14%	-0.02%	1.5%	-8.2%	33.7%	3,075
Net embodied CO ₂ imports/total emissions	1.10%	0.00%	6.6%	-7.0%	161.9%	3,075
Transport cost/FOB	5.54%	4.35%	7.1%	0.0%	95.8%	3,075
Bilateral flows						
<i>Outcomes</i>						
Trade flow (Mil. US\$)	97.56	1.09	882.00	0	132,123.3	321,360
Embodied CO ₂ flow (Mt)	0.02	0.00	0.22	0	21.65	321,360
Trade flow/output	0.55%	0.03%	4.14%	0	936.33%	321,360
Embodied CO ₂ flow/total emissions	0.56%	0.03%	4.23%	0	959.31%	321,360
Carbon intensity	0.04%	0.01%	0.14%	0	4.46%	320,035
<i>Covariates</i>						
Zero tariffs	46.9%		49.9%	0	1	321,360
Tariff (if not zero)	10.5%	5.4%	28.1%	0	2475.9%	170,689
$F_{unskilled}$ /value added	19.5%	17.1%	11.4%	0	68.6%	321,360
$F_{skilled}$ /value added	24.4%	22.3%	14.2%	0	100.0%	321,360
$F_{capital}$ /value added	56.1%	56.9%	19.5%	0	98.97%	321,360
Transport cost/FOB	4.7%	3.2%	6.5%	0	147.2%	320,035

Note: Measures of environmental stringency are computed for 2007 and 2011, and show the emission cost as a share of sectoral material cost. θ^1 is 1 for sectors explicitly targeted under the EU ETS, while θ^d also contains the direct emission costs of secondary activities. Additionally, θ^i captures indirect emission costs from the use of electricity. Carbon intensity is the ratio of a trade flow's embodied carbon over its value. We drop the observations from the Slovenian refinery sector which is an unrealistic outlier with over 22% of ETS cost. The number of observations is explained as follows: 1,247=2 years*25 sectors*25 countries (- Slovenia refinery); 3,075=3 years*25 sectors*41 partner regions; 321,360=3 years*25 sectors*66 source regions*(66-1) destination regions (- Slovenia refinery); some observations are dropped for carbon intensity and trade costs, as they have zero trade flows.

allowance prices over the EU ETS compliance Phase I (2005-2007) to compute θ^d , θ^i and θ^a in 2007, and the average price for Phase II up to 2011 (2008-2011) for 2011. This leads to allowance prices of €10.45 per metric ton of CO₂ in 2007 and €14.53 in 2011.³⁰

The level of direct emission cost θ^d has an average of 0.17% and is below 1% of material cost for the large majority of sectors. Only the iron & steel, cement, petroleum & coal products, non-ferrous metals (incl. aluminum) and other minerals (incl. glass and ceramics) sectors exceed this threshold.³¹ Free allocation θ^a is on average 0.22%, with allocations up to 10% of material costs; the resulting net direct ETS cost ($\theta^d - \theta^a$) is a net subsidy for the large majority of sectors. In general, allocations over-compensate direct emission cost, as evidenced by the slope of more than 1 between θ^d and θ^a (Figure V.2, left-hand panel).

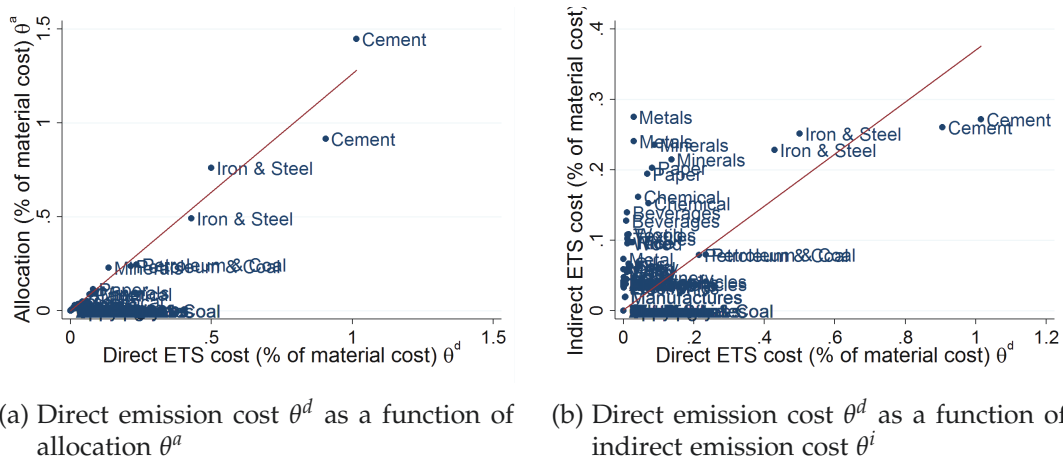


Figure V.2: Correlation of the measures of environmental stringency (scatter and fitted linear trend, sector-year averages across countries, 2007 & 2011)

Source: EUTL and authors' computations.

Sectors also incur indirect ETS costs θ^i from their electricity use. Indirect emission costs account on average for 0.16% of material costs. The largest indirect emission costs occur in the non-ferrous metals and iron & steel sectors, with up to 4.4% of material costs. For most sectors, indirect emission costs from electricity use are higher than direct costs, except for some emission-intensive sectors like cement and iron & steel (Figure V.2, right-hand panel). Adding up direct cost, indirect cost and subtracting the value of free allocation, EU manufacturing sectors were facing a net total emission cost of 0.11% of material cost on average over 2007 and 2011.

Our outcome variables are net trade flows and bilateral trade flows. We measure trade flows in value (U.S. dollars) and in embodied carbon, that is the sum of CO₂ emissions from all combustibles that served as an input to the traded goods (including emissions from electricity generation). In order to account for size effects, we scale net imports with

³⁰Our results are robust to using prices from each year only, instead of multi-year averages.

³¹We excluded the petroleum and coal products sector in Slovenia, which is an outlier with a value of 22.5% of material costs in 2007. Our results are robust to using the full dataset, and to excluding observations at the largest and smallest percentile.

output value and net CO₂ imports with total sectoral carbon emissions. The highest net imports both in value and as a share of output are electronic equipment from China. The highest embodied carbon net imports are in cement, also from China. Overall, Europe is a net importer embodied emissions via manufactured goods.




Figure V.3: Distribution of import taxes and transport costs (sector-averages 2004-2011)
Source: GTAP 9.2 and authors' computations.

Other important determinants of trade flows are import tariffs and transport costs, which may blur the relationship between emission costs and trade flows. Almost half of all trade flows are not subject to import tariffs. Figure V.3 (left-hand panel) shows the average tariffs of the remaining sectors that were 10.5% on average, ranging from almost zero to 62% (for the beverage and tobacco industry). Transport costs on bilateral trade flows amount on average to 4.7% of the FOB value of trade flows. Figure V.3 (right-hand panel) shows that sector averages vary between almost zero and 84%. The sectors with the highest transport costs relative to value are minerals and cement.




Figure V.4: Imports in value and imports in embodied carbon by bilateral EU ETS treatment status

Average difference between pre- and post-treatment country pair-sector averages of the logarithm. 1 are trade flows from an untreated source country to a treated destination country; -1 are trade flows from a treated source to an untreated destination; 0 are trade flows between countries with same treatment status and within untreated sectors.

Source: GTAP 9.2 and authors' computations.

Figure V.4 provides some descriptive evidence on the impact of the EU ETS on bilateral trade flows, similar to a Figure in Aichele and Felbermayr (2015), albeit applied

to the EU ETS instead of the Kyoto protocol. We define a bilateral treatment variable as $\theta_{mxt}^1 = (\theta_{mst}^1 - \theta_{xst}^1) \in \{-1, 0, 1\}$. θ_{mxt}^1 is equal to 1 if the trade flow goes from an untreated source country x to a treated destination country m (within a treated sector); -1 for trade flows from a treated source to an untreated destination (within a treated sector); 0 for trade flows between countries with same treatment status or for trade flows of untreated sectors.

Figure V.4 shows that bilateral trade in value has increased for all values of θ_{mxt}^1 , and the magnitudes broadly match those in Aichele and Felbermayr (2015). Trade in embodied carbon decreased both for trade to and from treated countries, while it increased in non-treated countries and sectors. Carbon leakage would translate into larger imports to and smaller imports from treated sectors, i.e. an increasing slope in both panels of Figure V.4, which does not appear in our data. The shift in trade in embodied carbon found by Aichele and Felbermayr (2015) must have occurred either prior to the introduction of the EU ETS or among non-EU countries who ratified Kyoto. Overall, the descriptive evidence does not suggest that imports in embodied carbon were affected by the introduction of the EU ETS.

5 Results

5.1 Net trade flows

5.1.1 Main results

As a first step, we implement the method suggested by Ederington et al. (2005)³²: using net trade data, we regress net imports in value and in embodied carbon on θ^1 , as well as on the vector $[\theta^d, \theta^i, \theta^a]$. In all regressions, we control for European import tariffs, transport costs (as a percentage of import value), and factor payment shares (as a percentage of value added), as well as for year and sector fixed effects.

The results in Table V.2 show no evidence of carbon leakage. None of the coefficients from regressions of net imports in embodied carbon on emission costs are significant. The (not significant) coefficient on the ETS dummy in Table V.2 column (1), and its 95% confidence interval of $[-.621, .327]$ are consistent with a maximum increase of 0.327 percentage points in net carbon imports in the treated sectors relative to untreated sectors. The confidence intervals in columns (2) to (4) are wider, but the magnitude of the estimates is still small given that the standard deviation of direct ETS cost is 0.6.

The only coefficients that are (weakly) significant appear in the regression of net imports in value on the individual components of emission cost from the ETS: column (7) of Table V.2 shows that net imports increased with direct ETS cost and decreased with allocation. The net effect in column (8) is about zero and not significant. This effect would be compatible with the hypothesis that there the carbon leakage effect has been

³²Ederington et al. (2005) aggregate over all partner countries of the U.S., while we use one observation per year-sector-partner country. Results on aggregate data yield the same result and are available on request, but the sample size shrinks to $N=75$.

Table V.2: Regression results for EU net imports (by partner country and sector)

<i>Emission cost</i>	Net embodied CO ₂ /total CO ₂				Net imports/output			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETS dummy θ_{st}^1	-0.147 (0.242)				0.001 (0.043)			
Direct ETS cost θ_{st}^d		-0.096 (0.272)	0.782 (0.657)			0.034 (0.129)	0.310* (0.129)	
Indirect ETS cost θ_{st}^i			-1.987 (1.983)				0.033 (0.295)	
Allocation benefit θ_{st}^a			-0.495 (0.576)				-0.256* (0.117)	
Total net ETS cost θ_{st}^{tot}				-0.791 (0.892)				0.010 (0.165)
R ²	0.89	0.89	0.89	0.89	0.92	0.92	0.92	0.92
N	3,075 (all columns)							

Notes: OLS regression of outcome on different definitions of ETS cost. Data is a sector-country-level panel for 2004, 2007, and 2011; all regressions include year and sector-country fixed effects and control for factor payment shares, tariffs, and transport costs. Robust standard errors clustered at the sector-partner country level in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

alleviated by free allocation. However, as we cannot find the same effect in trade flows in embodied carbon, we conclude that it must be an artifact of price fluctuations. Indeed, we will see that the significance of this result survives in none of our robustness checks. This result underlines, in our view, the importance of using embodied carbon flows, or at least trade flows in quantities rather than in value when doing an analysis of carbon leakage.

As a robustness test, we replicate the methodology of Levinson and Taylor (2008) by doing the same regression individually for each important trading partner country of the EU. We do not find any significant impact for any country.³³

As a further robustness test, we do the same regressions without controlling for factor payment shares (see Table V.6 in the Appendix on page 159), which exposes the fragility of the previously discussed significance result of column (7) in Table V.2. We provide a compact overview of further regressions, including robustness checks, in Table V.6 in the Appendix, where we only display coefficients of emission cost and their standard errors for all main specifications and robustness tests.

Additionally, we confirm that that our results are not affected by our normalization of the outcome variables (Appendix A on page 157).

5.1.2 Sector heterogeneity

Ederington et al. (2005) hypothesize that transport costs play an important role for carbon leakage, as some sectors are more *footloose* than others: if transport costs are high, industries are relatively more protected from foreign competition, such that environmen-

³³Results are available on request.

tal stringency has different effects for different industries. Our estimate may be biased if transport costs are correlated with environmental stringency: a typical example is the cement industry. In this case, both the measure of environmental stringency and its interaction with transport costs should have a significant negative coefficient.

In our data, we observe a low, but significant, positive correlation (of 0.06) between our measure of direct emission cost θ^d and transport costs, as well as a low, but significant, negative correlation (of -0.05) between our total net cost measure (θ^{tot}) and transport costs. As suggested by Ederington et al. (2005), we correct for this correlation by interacting our measures of policy stringency with transport costs. The results in Table V.3 show no significant effect (and are mostly of a sign not compatible with the carbon leakage hypothesis).³⁴ Thus, we conclude that we do not find evidence that sectors transport costs played a role in mitigating carbon leakage.

Table V.3: Regressions of net import flows on environmental cost and its interaction with transport costs and higher order terms of emission cost

<i>Emission cost</i>	Net embodied CO ₂ /total CO ₂			Net imports/output		
	(1)	(2)	(3)	(4)	(5)	(6)
ETS dummy θ_{mst}^1	-0.258 (0.498)			-0.028 (0.065)		
ETS dummy × Transport cost	-0.002 (0.024)			0.004 (0.005)		
Total net ETS cost θ_{mst}^{tot}		-0.634 (1.396)	-0.444 (1.300)		0.034 (0.176)	0.206 (0.191)
Total net ETS cost × Transport cost		-0.022 (0.077)			0.016 (0.019)	
Total net ETS cost <i>squared</i>			-12.612 (15.791)			1.466 (2.452)
Total net ETS cost <i>cubed</i>			30.657 (42.866)			-4.851 (6.122)
R ²	0.89	0.89	0.89	0.92	0.92	0.92
N	3,075 (all columns)					

Notes: OLS regression of outcome on different definitions of ETS cost. Data is a sector-level for aggregated and sector-country-level panel for "by country", each for 2004, 2007, and 2011; all regressions include year and sector-country fixed effects and control for factor payment shares, tariffs, and transport costs. Robust standard errors clustered at the sector-partner country level in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Our main regression is a linear approximation of the effect of environmental policy on trade. For robustness, we also fit a cubic polynomial (including two higher order terms of θ) to control for heterogeneous effects of the emission cost depending on its level. Table V.3 shows that the higher order terms are never significant, so that we find

³⁴Results on the vector of components of emission costs yield the same result, but are not represented here for compactness; results available on request.

no evidence of nonlinear effects.

5.2 Bilateral trade flows

Bilateral trade flow data provide a richer picture of international trade, accounting for (two-way) intra-industry trade. Our sample size increases dramatically to over 300,000 observations, between 66 source and destination countries for three years and 25 sectors. With these data, we implement the identification strategy of Aichele and Felbermayr (2015): they define a bilateral treatment variable that is the difference of treatment status of destination and source country $\theta_{mxst} = (\theta_{mst} - \theta_{xst}) \in \{-1, 0, 1\}$. This restricts the coefficients on treatment to be of opposite sign and identical magnitude: $\alpha^m = -\alpha^x$. In a second step, we relax this restriction and use separate emission cost measures for source and destination country.

Our regression results are shown in Table V.4. Odd-numbered columns contain specifications using bilateral definitions of our emission cost variables, while even-numbered columns present the results with separate emission cost measures for source and destination country. Carbon leakage is consistent with significant positive coefficients of destination emission cost, and significant negative coefficients of source emission cost, and, thus, a positive effect of the bilateral variables.

Columns (1) and (7) show the specifications corresponding to Aichele and Felbermayr (2015), and both are not significant. Our confidence interval in column (1) is compatible with a maximum increase in carbon imports of 0.037%; and column (7) allows for a maximum increase in imports in value of 0.376%. In contrast, Aichele and Felbermayr (2015) find that Kyoto ratification increases imports by 5% and raises the carbon content of trade (what we call “trade in embodied carbon”) by almost 8%. We conclude that the carbon leakage found by Aichele and Felbermayr (2015) has occurred outside the EU or before the introduction of the EU ETS.

Again, the regressions of embodied carbon are not significant except for the coefficients of indirect emission cost in column (4). However, these coefficients are both positive: both carbon imports *and* exports increased with indirect emission cost.³⁵ This is cannot be interpreted as carbon leakage; it thus must capture some other mechanism making electricity-intensive sectors in the EU more trade-intensive in *both* directions.

For trade flows in value, we have some coefficients that are significantly different from zero, but it is again not consistent with a carbon leakage explanation. The coefficient on the ETS dummy in column (8) has the “wrong” sign for the carbon leakage hypothesis. The signs of the coefficients in columns (9) and (10) have the “right” sign, but are not significant.

Table V.7 in the Appendix on page 160 provides additional results using bilateral variable definitions and alternative outcome variables, analogously to Table V.6 for net flows. Table V.8 in the Appendix on page 161 provides additional results on alternative

³⁵The signs of direct cost and allocation in column (4) are not significant but also identical for source and destination variables. This pattern is surprisingly robust to changes in specification, see Table V.8.

outcomes and specifications. In the majority of cases, the coefficients are estimated with a negative sign and are mostly statistically not significant.

We also explore the influence of fixed effects: our estimation relies on the difference-in-differences between sectors, countries and time. If we do not control for sectoral business cycles (ν_{st}), our results still hold. If we do not control for country-specific business cycles (ν_{mt} and ν_{xt}), the conclusions change and even more so, if we do not control for country-pair effects (ν_{mxs} capturing among other factors distance, essential for gravity estimations).³⁶ However, we believe that controlling for ν_{mt} , ν_{xt} and ν_{mxs} is essential to identification. Controlling for ν_{st} may be optional, and either doing so or not does not change the main results.

Overall, based on our analysis of bilateral trade and emission data, we conclude that the EU ETS did not have a systematic impact on flows of trade or embodied CO₂ emissions. Moreover, there is evidence against the hypothesis that $\alpha^m = -\alpha^x$.

6 Summary and conclusions

This paper considers whether the compliance cost imposed by the EU ETS on producers in European manufacturing sectors has caused carbon leakage. Carbon leakage, a special case of the pollution haven phenomenon, is an important topic in the context of unilateral environmental policy. A unilateral policy intervention changes the relative competitiveness of domestic producers vis-à-vis their global competitors. In the extreme case, carbon leakage undoes the contribution of the unilateral policy to mitigate aggregate global emissions, while the region implementing the policy suffers losses in output, employment, and welfare. This loss in competitiveness due to the EU ETS can occur directly, as producers must abate or pay for the cost of their own emissions, and indirectly through the consumption of electricity, when electricity producers pass through emission costs to power prices. In the EU ETS, the direct emission cost was largely defrayed by free allocation during the period under study, such that the majority of sectors enjoy a net subsidy when net direct costs are considered. Moreover, overall emission costs have so far been small compared to other material costs. In addition to low carbon prices and free allocation, there are further obstacles to leakage: relocation is costly and risky, as the new host region may also introduce corresponding policies in the future. Finally, the EU ETS may also have beneficial effects, such as incentivizing green innovation by producers, which help them become more competitive internationally.

Our empirical analysis is based on the hypothesis that leakage can be measured through changes in trade flows, particularly flows in embodied carbon. This hypothesis can be derived from classical trade theory or from New trade theory. Combining data from GTAP, a global trade dataset with input-output information, and administrative data from the EU ETS, we estimate the effect of various potential measures of the stringency of the EU ETS on trade flows in manufactured goods. Our measures of policy stringency account both for the direct emission cost and the indirect emission cost

³⁶Results available on request.

Table V.4: Bilateral trade flows in logs of million U.S. dollars and embodied CO₂ emissions

Emission cost	ln(embodied carbon)			ln(trade flow)		
	(1)	(2)	(3)	(4)	(5)	(6)
ETS dummy						
- bil. $\theta_{mst}^1 - \theta_{xst}^1$	-0.002 (0.018)					
- dest. θ_{mst}^1		-0.017 (0.025)				
- source θ_{xst}^1		-0.020 (0.025)				
Direct ETS cost						
- bil. $\theta_{mst}^d - \theta_{xst}^d$			0.012 (0.021)			
- dest. θ_{mst}^d				0.033 (0.028)		
- source θ_{xst}^d				0.009 (0.030)		
Indirect ETS cost						
- bil. $\theta_{mst}^i - \theta_{xst}^i$			-0.012 (0.011)			
- dest. θ_{mst}^i				0.049** (0.015)		
- source θ_{xst}^i				0.076*** (0.016)		
Allocation benefit						
- bil. $\theta_{mst}^a - \theta_{xst}^a$			0.005 (0.014)			
- dest. θ_{mst}^a				-0.023 (0.019)		
- source θ_{xst}^a				-0.033 (0.020)		
Total net ETS cost θ^{tot}						
- bil.					-0.007 (0.015)	
- dest.						0.007 (0.018)
- source						0.000 (0.00)
R ²	0.959	0.959	0.959	0.959	0.959	0.962
N	320, 035	320, 035	320, 035	320, 035	320, 035	320, 035

Notes: OLS regression of outcome on different definitions of ETS cost. Data is a country pair-sector-level panel for 2004, 2007, and 2011; all regressions include year-source country, year-destination country and sector-country pair fixed effects and control for factor payment shares and import tariffs. Robust standard errors clustered at the sector-country pair level in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

from electricity use. Our empirical analysis follows two traditions in the trade literature: first, we consider the effect of EU ETS stringency on net trade flows, as suggested by the neoclassical trade literature (Ederington et al. 2005), where we also consider sector heterogeneity. In particular, we test for a potentially stronger effect of EU ETS stringency on *footloose* industries and for nonlinearity of the effects of EU ETS stringency. Second, we follow the New trade literature by analyzing the effect of policy stringency on bilateral trade flows (Aichele and Felbermayr 2015).

We find no evidence that the EU ETS has induced carbon leakage in European manufacturing sectors. This result is in line with existing empirical ex post research on carbon leakage due to the EU ETS, but contrasts with predictions from ex ante modeling exercises. Our results relate to existing work on other environmental policies, like Aichele and Felbermayr (2015) who show that ratification of the Kyoto Protocol has caused carbon leakage. Our results suggest that the leakage found by Aichele and Felbermayr (2015) must have occurred in Kyoto signatory countries who were not part of the EU ETS. Our results suggest Branger et al. (2016) were right to call the debate about carbon leakage “much ado about nothing.”

The absence of trade effects suggests that the barriers preventing leakage are greater than emission costs inducing leakage. Current allowance prices in the EU ETS are low and firms may have some market power. Tariffs and transportation costs are typically higher than CO₂-related costs and contribute to firms’ ability to pass-through at least some of their emission cost to the final consumer without losing significant market share. Additionally, more diffuse factors, such as political risk, exchange rate concerns, and considerations about the availability of qualified labor may limit leakage. Further research will help identify factors mainly responsible for the absence of leakage or find a level of the emission cost for which carbon leakage is a real concern.

The absence of carbon leakage is good news for the political feasibility of unilateral CO₂ policies such as the EU ETS even in a context of globally asymmetric climate policy, at least at current allowance prices. If they do not hamper domestic competitiveness and economic growth, environmental policies are more likely to be implemented.

7 Bibliography

- Acemoglu D, Aghion P, Bursztyn L, and Hémous D (2012) The environment and directed technical change. *The American Economic Review*, 102(1): 131–166.
- Aghion P, Dechezleprêtre A, Hémous D, Martin R, and Van Reenen J (2016) Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1): 1–51.
- Aguiar A, Narayanan B, and McDougall R (2016) An overview of the GTAP 9 data base. *Journal of Global Economic Analysis*, 1(1): 181–208.
- Aichele R and Felbermayr G (2012) Kyoto and the carbon footprint of nations. *Journal of Environmental Economics and Management*, 63(3): 336–354.
- (2015) Kyoto and carbon leakage: An empirical analysis of the carbon content of bilateral trade. *Review of Economics and Statistics*, 97(1): 104–115.
- Antweiler W, Copeland B. R, and Taylor M. S (2001) Is free trade good for the environment? *American Economic Review*, 91(4): 877–908.
- Branger F and Quirion P (2014) Climate policy and the ‘carbon haven’ effect. *Wiley Interdisciplinary Reviews: Climate Change*, 5(1): 53–71.
- Branger F, Quirion P, and Chevallier J (2016) Carbon leakage and competitiveness of cement and steel industries under the eu ets: Much ado about nothing. *The Energy Journal*, 37(3).
- Brunel C and Levinson A (2016) Measuring the stringency of environmental regulations. *Review of Environmental Economics and Policy*, 10(1): 47–67.
- Brunnermeier S. B and Levinson A (2004) Examining the evidence on environmental regulations and industry location. *The Journal of Environment & Development*, 13(1): 6–41.
- Calel R and Dechezleprêtre A (2016) Environmental policy and directed technological change: Evidence from the European carbon market. *Review of Economics and Statistics*, 98(1): 173–191.
- Carbone J. C and Rivers N (2017) The impacts of unilateral climate policy on competitiveness: Evidence from computable general equilibrium models. *Review of Environmental Economics and Policy*, 11(1): 24–42.
- Caron J, Fally T, and Markusen J. R (2014) International trade puzzles: A solution linking production and preferences. *The Quarterly Journal of Economics*, 129(3): 1501–1552.
- Copeland B. R and Taylor M. S (2004) Trade, growth, and the environment. *Journal of Economic literature*, 42(1): 7–71.

- (2005) Free trade and global warming: A trade theory view of the Kyoto protocol. *Journal of Environmental Economics and Management*, 49(2): 205–234.
- De Loecker J, Goldberg P. K, Khandelwal A. K, and Pavcnik N (2016) Prices, markups, and trade reform. *Econometrica*, 84(2): 445–510.
- Dechezleprêtre A, Gennaioli C, Martin R, and Muûls M (2014) Searching for carbon leaks in multinational companies. Working Paper 165, Grantham Research Institute on Climate Change and the Environment.
- Dechezleprêtre A and Sato M (2017) The impacts of environmental regulations on competitiveness. *Review of Environmental Economics and Policy*, 11(2): 183–206.
- Demay D and Quirion P (2006) CO₂ abatement, competitiveness and leakage in the European cement industry under the EU ETS: grandfathering versus output-based allocation. *Climate Policy*, 6(1): 93–113.
- Dixit A. K and Stiglitz J. E (1977) Monopolistic competition and optimum product diversity. *The American Economic Review*, 67(3): 297–308.
- Ederington J, Levinson A, and Minier J (2005) Footloose and pollution-free. *Review of Economics and Statistics*, 87(1): 92–99.
- Ellerman A. D, Marcantonini C, and Zaklan A (2016) The European Union Emissions Trading System: ten years and counting. *Review of Environmental Economics and Policy*, 10(1): 89–107.
- EU (2014) Directive 2003/87/EC of the European Parliament and of the Council of 13 October 2003 Establishing a Scheme for Greenhouse Gas Emission Allowance Trading within the Community and Amending Council Directive 96/61/EC. 2003L0087 - EN - 30.04.2014 - 007.001 - 1.
- Fabra N and Reguant M (2014) Pass-through of emissions costs in electricity markets. *The American Economic Review*, 104(9): 2872–2899.
- Gerlagh R and Kuik O (2014) Spill or leak? Carbon leakage with international technology spillovers: a CGE analysis. *Energy Economics*, 45: 381–388.
- Grossman G. M and Krueger A. B (1991) Environmental impacts of a North American free trade agreement. Working Paper 3914, National Bureau of Economic Research.
- Harstad B (2012) Buy coal! A case for supply-side environmental policy. *Journal of Political Economy*, 120(1): 77–115.
- Head K and Mayer T (2014) Gravity equations: workhorse, toolkit, and cookbook. In G. Gopinath, E. Helpman, and K. Rogoff (eds.) *Handbook of International Economics*, 4: 131–196.

- Hintermann B (2016) Pass-through of CO₂ emission costs to hourly electricity prices in Germany. *Journal of the Association of Environmental and Resource Economists*, 3(4): 857–891.
- Huff K, McDougall R, and Walmsley T (2000) Contributing input-output tables to the GTAP data base. Technical Paper 1, Center for Global Trade Analysis, Purdue University.
- IPCC (2007) Fourth assessment report (AR4). Synthesis Report, Intergovernmental Panel on Climate Change.
- Jaffe A. B, Peterson S. R, Portney P. R, and Stavins R. N (1995) Environmental regulation and the competitiveness of U.S. manufacturing: What does the evidence tell us? *Journal of Economic Literature*, 33(1): 132–163.
- Jensen S, Mohlin K, Pittel K, and Sterner T (2015) An introduction to the green paradox: the unintended consequences of climate policies. *Review of Environmental Economics and Policy*, 9(2): 246–265.
- Koch N and Basse Mama H (2016) European climate policy and industrial relocation: evidence from German multinational firms. SSRN Working Paper.
- Kronmal R. A (1993) Spurious correlation and the fallacy of the ratio standard revisited. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 156(3): 379–392.
- Krugman P (1980) Scale economies, product differentiation, and the pattern of trade. *American Economic Review*, 70(5): 950–959.
- Kuh E and Meyer J. R (1955) Correlation and regression estimates when the data are ratios. *Econometrica*, 23(4): 400–416.
- Levinson A and Taylor M. S (2008) Unmasking the pollution haven effect. *International Economic Review*, 49(1): 223–254.
- Martin R, Muûls M, de Preux L. B, and Wagner U (2014) Industry compensation under relocation risk: a firm-level analysis of the EU emissions trading scheme. *American Economic Review*, 104(8): 2482–2508.
- Martin R, Muûls M, and Wagner U. J (2016) The impact of the European Union Emissions Trading Scheme on regulated firms: What is the evidence after ten years? *Review of Environmental Economics and Policy*, 10(1): 129–148.
- McGuire M. C (1982) Regulation, factor rewards, and international trade. *Journal of public economics*, 17(3): 335–354.
- Melitz M. J (2003) The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6): 1695–1725.

- Motta M and Thisse J.-F (1994) Does environmental dumping lead to delocation? *European Economic Review*, 38(3-4): 563–576.
- Pethig R (1976) Pollution, welfare, and environmental policy in the theory of comparative advantage. *Journal of Environmental Economics and Management*, 2(3): 160–169.
- Porter M. E and Van der Linde C (1995) Toward a new conception of the environment-competitiveness relationship. *The Journal of Economic Perspectives*, 9(4): 97–118.
- Sartor O (2013) Carbon leakage in the primary aluminium sector: What evidence after 6.5 years of the EU ETS? Working Paper 2012-12, CDC Climat.
- Schmidt R. C and Heitzig J (2014) Carbon leakage: grandfathering as an incentive device to avert firm relocation. *Journal of Environmental Economics and Management*, 67(2): 209–223.
- Schröder C (2016) Industrielle Arbeitskosten im internationalen Vergleich. *IW-Trends. Vierteljahresschrift zur empirischen Wirtschaftsforschung*, 43(3): 17–26.
- Tobey J. A (1990) The effects of domestic environmental policies on patterns of world trade: an empirical test. *Kyklos*, 43(2): 191–209.

Appendices

A Potential bias from using normalized variables

The specification of equations (1) and (2) use trade intensity (normalized trade flows) as a dependent variable and environmental policy intensity (normalized emission cost) as main covariate variable. The additional variables transport cost and tariffs are also measured per value unit. We believe this specification represents the relevant magnitudes. Indeed both total emission cost and exports depend on an underlying “sectoral market size” parameter, which would create spurious correlation if not accounted for. Moreover, we compare countries of very different magnitude, where we would face an outlier problem and heteroskedasticity concerns if we were not normalizing. Our normalizations are based on a suggestion of Ederington et al. (2005).

However, the use of ratios is discussed extensively in the statistical literature: when two variables have zero correlation, positive (spurious) correlation might appear in a regression if both left-hand and right-hand side variables are normalized by a common denominator. The bias is even stronger if the variables are correlated with each other and with the common denominator (Kronmal 1993). Note that ratios are generally found to bias the absolute magnitude of estimates *upwards* (e.g. Kuh and Meyer 1955); as we find no significant impact of policy stringency on trade flows, an upward bias would not change our conclusion and in fact strengthens our results. In this Section, we follow the recommendations of Kronmal (1993) to check that results are not an artifact of normalization.

Let Y be an $n \times 1$ vector, Z a diagonal $n \times n$ matrix and X an $n \times p$ matrix, centered such that the mean of each column is zero (e.g. demeaned). Assume that the true model is:

$$Y = \mathbb{1}_n \beta_0 + X \beta_X + Z \mathbb{1}_n \beta_Z + \epsilon \quad (3)$$

where β_0 and β_Z are scalars and β_X a $p \times 1$ vector. $\mathbb{1}_n$ is a $n \times 1$ vector of ones. Our main specification can then be written:

$$Z^{-1}Y = \mathbb{1}_n \alpha_0 + Z^{-1}X \alpha_X + \epsilon \quad (4)$$

Kronmal (1993) shows that estimate $\hat{\alpha}_X$ from least squares of equation (4) is in general a biased estimator of β_X . Indeed, dividing both sides of equation (3) by Z yields

$$Z^{-1}Y = Z^{-1}\mathbb{1}_n \beta_0 + Z^{-1}X \beta_X + \mathbb{1}_n \beta_Z + Z^{-1}\epsilon \quad (5)$$

The least squares estimates of equation (5) are unbiased estimates of the parameters of equation (3). Empirically, this corresponds to estimating equation (4) and adding the scaling variable as an additional right-hand side variable. Dividing the error term by Z results in heteroskedasticity such that OLS is no longer the efficient estimator.

The result in Table V.5 control for the scaling variable, which is total domestic sectoral product for imports in value and total domestic sectoral emissions for imports in embod-

Table V.5: Regression of net trade flows on emission cost with additional control for scaling variables (compare to Table V.2)

<i>Emission cost</i>	Net embodied CO ₂ /total CO ₂				Net imports/output			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ETS dummy θ_{st}^1	−0.265 (0.351)				0.005 (0.053)			
Direct ETS cost θ_{st}^d		−0.258 (0.286)	0.733 (0.655)			−0.007 (0.070)	0.311* (0.129)	
Indirect ETS cost θ_{st}^i			−1.837 (1.937)				0.027 (0.294)	
Allocation benefit θ_{st}^a			−0.470 (0.573)				−0.252* (0.118)	
Total net ETS cost θ_{st}^{tot}				−0.704 (1.151)				0.129 (0.162)
Scaling variable	0.944 (0.73)	0.954 (0.73)	0.870 (0.69)	0.936 (0.73)	−722.0 (580.4)	−711.1 (562.7)	−697.3 (559.5)	−693.5 (552.8)
R ²	0.89	0.89	0.89	0.89	0.92	0.92	0.92	0.92
N	3,075 (all columns)							

Notes: OLS regression of outcome on different definitions of ETS cost, controlling for scaling variables. Data is a sector-country-level panel for 2004, 2007, and 2011; all regressions include year and sector-country fixed effects and control for factor payment shares, tariffs, and transport costs. Robust standard errors clustered at the sector-partner country level in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively.

ied carbon. Including this value is necessary if the true model is given by equation (3) rather than our model as in equation (4). Table V.5 shows that our results are robust to this modification, as magnitudes and significance remain virtually unchanged

For bilateral data, our main specifications in Table V.4 use logarithms rather than normalizing the variables, following Aichele and Felbermayr (2015). The regressions using raw (not normalized or in logs) variables suggest a significant correlation of surprising sign; however, this effect vanishes when using normalized variables or logs. We again test if the normalization for bilateral data is problematic and find that the coefficients change little when including the scaling variable.³⁷

³⁷Results available on request.

B Additional regression results

B.1 Net trade flows

Table V.6: Summary overview stating only the coefficient of the ETS stringency variables (for different specifications of net flows)

	ETS dummy	Total ETS cost	Direct ETS cost	Indirect ETS cost	Allocation
<i>Outcomes</i>	(1)	(2)	(3a)	(3b)	(3c)
Net carbon imports (MtCO ₂)	0.012 (0.083)	0.022 (0.144)	1.102 (1.047)	−0.145 (0.440)	−0.796 (0.743)
Net imports (Mil. US \$)	180.0 (175.7)	267.5 (196.8)	686.8 (600.2)	745.0 (762.7)	−453.1 (381.3)
Net carbon/total carbon (%)	−0.273 (0.354)	0.684 (0.552)	0.782 (0.657)	−1.987 (1.983)	−0.495 (0.576)
-w/o factor payment	−0.147 (0.242)	−0.257 (0.274)	0.957 (0.768)	−1.902 (1.638)	−0.494 (0.482)
Net imports/output value (%)	0.001 (0.052)	0.174 (0.142)	0.310* (0.129)	0.033 (0.295)	−0.256* (0.117)
-w/o factor payment	0.001 (0.043)	0.017 (0.061)	0.221 (0.119)	−0.081 (0.292)	−0.133 (0.101)

Notes: Summary table of regressions of different outcome variables (rows) on different ETS stringency variables (columns). In column 1 and 2, each coefficient comes from a separate regression. In columns 3a to 3c, each row is a regression of the outcome on direct cost, indirect cost and allocation. All regressions include fixed effects and controls mentioned in our main results. Robust standard errors clustered at the sector level in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively.

B.2 Bilateral trade flows

Table V.7: Summary overview stating only the coefficient of the *bilateral ETS stringency variables* (for different specifications of bilateral flows)

	ETS dummy	Total ETS cost	Direct ETS cost	Indirect ETS cost	Allocation
<i>Outcomes</i>	(1)	(2)	(3a)	(3b)	(3c)
Carbon imports (MtCO ₂)	0.000 (0.00)	0.000 (0.00)	0.001 (0.00)	0.000 (0.00)	−0.001 (0.00)
Imports (Mil. US \$)	3.179 (2.61)	0.537 (0.72)	1.063 (1.19)	6.180** (2.17)	0.039 (0.92)
Carbon/total carbon (%)	−0.005 (0.04)	−0.004 (0.01)	0.001 (0.02)	0.020 (0.02)	0.006 (0.02)
Imports/output value (%)	−0.004 (0.04)	−0.002 (0.01)	0.004 (0.02)	0.020 (0.02)	0.004 (0.02)
log(carbon)	−0.002 (0.02)	−0.007 (0.01)	0.012 (0.02)	−0.012 (0.01)	0.005 (0.01)
-w/ factor payment	−0.001 (0.02)	0.009 (0.02)	0.012 (0.02)	−0.012 (0.01)	0.005 (0.01)
log(imports)	0.023 (0.02)	0.027 (0.01)	0.058 (0.03)	0.035* (0.01)	−0.018 (0.02)
-w/ factor payment	0.023 (0.02)	0.054 (0.04)	0.058 (0.03)	0.035* (0.01)	−0.018 (0.02)

Note: Summary table of regressions of different outcome variables (rows) on different ETS stringency variables (columns); bilateral treatment indicator is defined as treatment for the destination country minus treatment for the source country. In column 1 and 2, each coefficient comes from a separate regression. In columns 3a to 3c, each row is a regression of the outcome on direct cost, indirect cost and allocation. All regressions include year-sector and sector-country-pair fixed effects, as mentioned in our main results. Robust standard errors clustered at the sector level in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

Table V.8: Summary overview stating only the coefficient of the *two separate ETS stringency variables* (for different specifications of bilateral flows)

<i>Outcomes</i>	ETS dummy		Total ETS cost		Direct ETS cost		Indirect ETS cost		Allocation	
	source (1a)	destination (1b)	source (2a)	destination (2b)	source (3a)	destination (3b)	source (3c)	destination (3d)	source (3e)	destination (3f)
Carbon imports (MtCO ₂)	-0.010*** (0.00)	-0.010*** (0.00)	0.000 (0.00)	0.001 (0.00)	-0.002 (0.00)	0.001 (0.00)	0.002 (0.01)	0.001 (0.00)	0.000 (0.00)	-0.002 (0.00)
Imports (Mil. US \$)	-9.894 (5.94)	-3.553 (5.93)	-3.890** (1.21)	-2.840* (1.32)	-6.729*** (1.84)	-4.659* (1.99)	-13.620*** (3.13)	-0.930 (2.91)	5.483*** (1.39)	5.595*** (1.45)
Carbon/total carbon (%)	-0.067 (0.05)	-0.077 (0.06)	-0.015 (0.02)	-0.023 (0.02)	-0.007 (0.02)	-0.005 (0.04)	-0.017 (0.02)	0.023 (0.03)	0.001 (0.01)	0.013 (0.03)
Imports/output value (%)	-0.067 (0.05)	-0.076 (0.06)	-0.016 (0.02)	-0.021 (0.02)	-0.011 (0.02)	-0.004 (0.04)	-0.017 (0.02)	0.023 (0.02)	0.004 (0.01)	0.012 (0.03)
log(carbon)	-0.017 (0.02)	-0.020 (0.02)	0.021 (0.02)	0.007 (0.02)	0.009 (0.03)	0.033 (0.03)	0.076*** (0.02)	0.049** (0.02)	-0.033 (0.02)	-0.023 (0.02)
log(imports)	-0.065** (0.02)	-0.020 (0.02)	-0.046* (0.02)	0.007 (0.02)	-0.083 (0.06)	0.032 (0.03)	-0.021 (0.03)	0.050* (0.00)	0.013 (0.03)	-0.022 (0.02)

Notes: Summary table of regressions of different outcome variables (rows) on different ETS stringency variables (columns); each regression includes treatment for the source country and treatment for the destination country. There is one regression for each column number and line. All regressions include year-sector and sector-country-pair fixed effects, as mentioned in our main results. Robust standard errors clustered at the sector level in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% level, respectively.