


**Towards sustainable transport:
macroeconomic and
microeconometric perspectives**

Ciarán Mac Domhnaill

A thesis submitted for the degree of PhD
At the
University of St Andrews



2024

Full metadata for this thesis is available in
St Andrews Research Repository
at:

<https://research-repository.st-andrews.ac.uk/>

Identifier to use to cite or link to this thesis:

DOI: <https://doi.org/10.17630/sta/829>

This item is protected by original copyright

This item is licensed under a
Creative Commons Licence

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

ABSTRACT

Progressing towards sustainable transport will be critical to meeting the Paris Climate Change Agreement's target of limiting the global temperature rise to 1.5 degrees Celsius. This thesis explores questions relevant to this transition. Have we reached 'peak car'? Did Brexit cause a shift from road freight transport to the more energy efficient mode of short sea shipping? Can ride hailing platforms advance transport sustainability?

In Chapter 1, using 1950-2010 data on 88 countries, I demonstrate that the number of private cars per GDP in the economy, or 'car intensity', evolves in a hump-shaped pattern during economic development. I develop a general equilibrium model to argue that structural transformation can generate this trend. My calibrated model can account for a quarter of observed variation in car intensity among 54 countries in 2010. Counterfactual exercises show that peak intensity is lower for economies that develop later.

In Chapter 2, I conduct a difference-in-differences analysis of the effect of Brexit on maritime cargo volumes. I examine 2013-2022 Eurostat port-level data and find a 22 per cent decrease in EU-UK roll-on roll-off (Ro-Ro) volumes, and a 54 per cent decrease in Ireland-UK Ro-Ro volumes, due to Brexit. I find a concurrent increase of 147 per cent in Ireland-France Ro-Ro cargo, indicating a diversion from the UK land-bridge route to direct short sea shipping routes. I estimate that emissions would be roughly 60 per cent lower on the direct route.

In Chapter 3, I employ Scottish Household Survey 2012-2019 travel diary data from 16,712 individuals in a difference-in-differences examination of how ride hailing affected the use of other transport modes. Results reveal a small complementary effect on the use of public transport relative to driving a car in Glasgow, which is more pronounced among individuals who are younger, male, and with higher household income.

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my supervisors, Radoslaw Stefanski and David A. Jaeger, for all their expertise and support with this research project.

This endeavour would not have been possible without the scholarship funded by Sir Bob and Mike Reid.

Special thanks also go to my PhD colleagues: Nicolò Bandera, Francesca Chiaradia, Qunli Hu, Mario Lupoli, Arham Nahar, Janie Olver, Richard Sparkes, Jacob Stevens, Kairat Umargaliev, and Lazaros Zografopoulos. It was a genuine privilege to get to know you all and to work alongside you in the Old Burgh School.

I am also grateful to my annual progress reviewers, Irina Merkurieva and Roderick McCrorie, for their guidance and advice.

For Chapter 1, I would like to recognise helpful comments from Alex Trew, Julieta Caunedo, Margaret Leighton, Ozge Senay, Sudharshan Reddy Paramati, Chenchuan Shi and participants at the annual conferences of the Society for Economic Dynamics, International Transport Economics Association, Irish Economic Association and Royal Economic Society PhD Conference.

For Chapter 2, I would like to recognise helpful comments from Seán Lyons, Luc Bridet, Margherita Negri, Lorenzo Neri and David Escamilla-Guerrero.

For Chapter 3, I would like to recognise helpful comments from Seán Lyons, Lorenzo Neri, Anthony Higney, Timothy J. Moore, and participants at the Second Scotland and Northern England Conference in Applied Microeconomics.

Finally, I could not have undertaken this journey without the continuous support of Mum, Dad and my brother Odhrán.

DECLARATION

Candidate's declaration

I, Ciarán Mac Domhnaill, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 43,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in January 2021.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date 15 December 2023

Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date 15 December 2023

Signature of supervisor

PERMISSION FOR PUBLICATION

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Ciarán Mac Domhnaill, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 15 December 2023

Signature of candidate

Date 15 December 2023

Signature of supervisor

UNDERPINNING RESEARCH DATA OR DIGITAL OUTPUTS

Candidate's declaration

I, Ciarán Mac Domhnaill, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 15 December 2023

Signature of candidate

FUNDING

This work was supported by the University of St Andrews (Sir Bob Reid Handsel Scholarship).

CONTENTS

Abstract	iii
Acknowledgements	v
Declaration	vii
Permissions	ix
Underpinning Research Data or Digital Outputs	xi
Funding	xiii
List of Figures	xix
List of Tables	xxiv
Introduction	1
1 Driving over the hill?	7
1.1 Introduction	7
1.1.1 Peak car	8
1.1.2 Structural transformation	14
1.1.3 Environmental Kuznets Curve	17
1.1.4 This study	18
1.2 The facts	19
1.2.1 Car use per GDP	19
1.2.2 Car intensity	23
1.2.3 Carbon emissions intensity	26
1.2.4 Explaining these facts	27
1.2.5 Structural transformation	30
1.3 The model	36
1.3.1 Consumers	37
1.3.2 Firms	37
1.3.3 Carbon emissions	40
1.3.4 Competitive equilibrium	41

CONTENTS

1.3.5	<i>Characterisation</i>	41
1.3.6	<i>Model discussion</i>	42
1.4	<i>Model calibration</i>	45
1.5	<i>Baseline simulation</i>	50
1.5.1	<i>Calibrated period: 1995-2016</i>	51
1.5.2	<i>External validity: 1970-2016</i>	53
1.5.3	<i>Cross-sectional model fit</i>	55
1.5.4	<i>Case studies: France, South Korea and China</i>	57
1.6	<i>Role of structural transformation</i>	61
1.6.1	<i>Counterfactual 1: No structural transformation</i>	62
1.6.2	<i>Counterfactual 2: Later structural transformation</i>	63
1.7	<i>Empirical analysis</i>	64
1.7.1	<i>Method</i>	65
1.7.2	<i>Data</i>	68
1.7.3	<i>Results</i>	69
1.8	<i>Discussion</i>	73
1.8.1	<i>Policy implications</i>	75
1.8.2	<i>Limitations and strengths</i>	76
2	<i>All at sea?</i>	79
2.1	<i>Introduction</i>	79
2.1.1	<i>How has Brexit affected trade?</i>	80
2.1.2	<i>How has Brexit affected maritime transport?</i>	84
2.1.3	<i>This study</i>	85
2.2	<i>Methods</i>	87
2.2.1	<i>Theoretical framework</i>	87
2.2.2	<i>Study design</i>	89
2.2.3	<i>Alternative specifications</i>	94
2.2.4	<i>Data</i>	98
2.2.5	<i>Outcome variable</i>	101
2.3	<i>Results</i>	106
2.3.1	<i>Cargo volumes</i>	106
2.3.2	<i>Robustness</i>	108
2.4	<i>Energy consumption and emissions</i>	112
2.5	<i>Implicit tariff on land-bridge</i>	116
2.6	<i>Discussion</i>	118
2.6.1	<i>Policy implications</i>	123
2.6.2	<i>Limitations and strengths</i>	124

3	All hail?	127
3.1	<i>Introduction</i>	127
3.1.1	<i>What is ride hailing?</i>	128
3.1.2	<i>Who uses ride hailing?</i>	130
3.1.3	<i>How does ride-hailing affect other forms of transport?</i>	131
3.1.4	<i>How has regulation responded to ride hailing?</i>	135
3.1.5	<i>This study</i>	136
3.2	<i>Methods</i>	137
3.2.1	<i>Theoretical framework</i>	137
3.2.2	<i>Study design</i>	141
3.2.3	<i>Alternative specifications</i>	147
3.2.4	<i>Data</i>	151
3.2.5	<i>Outcome variables</i>	156
3.2.6	<i>Difference-in-differences variables</i>	159
3.3	<i>Results</i>	160
3.3.1	<i>Main transport mode</i>	160
3.3.2	<i>Average journey speed</i>	163
3.3.3	<i>Robustness</i>	164
3.3.4	<i>Mechanism</i>	166
3.4	<i>Discussion</i>	168
3.4.1	<i>Policy implications</i>	172
3.4.2	<i>Limitations and strengths</i>	173
	Conclusions	175
	Appendix A Appendices to Chapter 1	179
A.1	<i>Data</i>	179
A.1.1	<i>Main country panel: 54 countries</i>	179
A.1.2	<i>Extended country panel: 88 countries</i>	188
A.1.3	<i>Great Britain</i>	189
A.1.4	<i>Empirical analysis</i>	190
A.2	<i>Disaggregating non-agriculture</i>	191
A.3	<i>Model theory</i>	195
A.3.1	<i>Consumers</i>	195
A.3.2	<i>Firms</i>	196
A.3.3	<i>Counterfactual 1: No structural transformation</i>	198
	Appendix B Appendices to Chapter 2	201
B.1	<i>Model theory</i>	201
B.2	<i>Additional tables and figures</i>	203

CONTENTS

B.2.1	<i>Descriptive statistics</i>	203
B.2.2	<i>Additional results</i>	204
B.2.3	<i>Untransformed PPML coefficients</i>	211
Appendix C	<i>Appendices to Chapter 3</i>	217
C.1	<i>Data</i>	217
C.2	<i>Additional figures and tables</i>	220
C.2.1	<i>Respondent characteristics</i>	220
C.2.2	<i>Additional results</i>	221
C.2.3	<i>Untransformed multinomial logistic regression coefficients</i>	230
C.3	<i>Transport in Scotland</i>	231
C.3.1	<i>City transport profiles</i>	232
C.3.2	<i>Changes in transport infrastructure</i>	234
C.3.3	<i>Changes in transport prices</i>	235
C.3.4	<i>Changes in aggregate transport use</i>	237
References		251

LIST OF FIGURES

1.1	Car passenger kilometres per GDP and real GDP per capita, 29 countries 1970-2019. Axes transformed to \log_e scales. Points show decile averages. Sources: Author's analysis; OECD 2017; Feenstra, Inklaar, and Timmer 2015.	20
1.2	Car VKM intensity and private car intensity, Great Britain 1950-2019. Sources: Author's analysis; UK Department for Transport 2021b, 2021a; Feenstra, Inklaar, and Timmer 2015.	23
1.3	Private car intensity and real GDP per capita, extended panel. Axes transformed to \log_e scales. Sources: Author's analysis; Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.	24
1.4	Carbon emissions intensity and real GDP per capita, extended panel 1971-2019. Axes transformed to \log_e scales. Points show decile averages. Sources: Author's analysis; IEA 2022c; Feenstra, Inklaar, and Timmer 2015.	27
1.5	Peak level of private cars per capita and real GDP per capita, extended panel 1950-2010. X-axis transformed to \log_e scale. Sample limited to countries with GDP per capita >20,000 PPP USD and where annual growth levels of vehicles per capita were below 2 per cent in each of the final 3 years of the sample period. Source: Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.	28
1.6	PKM per GDP across all transport modes and real GDP per capita, Great Britain 1952-2019. Axes transformed to \log_e scales. Source: Author's analysis; Feenstra, Inklaar, and Timmer 2015; UK Department for Transport 2013.	29
1.7	Private cars per GDP, road fuel prices and real GDP per capita, Great Britain 1955-2010. Y-axis transformed to \log_e scale. Road fuel prices are typical January retail prices in current pence per litre, indexed relative to 1955. Petrol price is for lead replacement petrol 1955-1988, and premium unleaded petrol 1989-2010. Source: Author's analysis; UK Department for Energy Security and Net Zero 2023b; Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.	30
1.8	Changes in sectoral motor vehicle intensity over structural transformation, main panel. Sources: Author's analysis; OECD 2021; ILO 2020; Feenstra, Inklaar, and Timmer 2015.	35
1.9	Structural transformation, baseline model simulation (blue dashed line) and data (red solid line), 1995-2016 calibration period.	52

LIST OF FIGURES

1.10	Intensity, baseline model simulation (blue dashed line) and data (red solid line), 1995-2016 calibration period.	53
1.11	Sectoral motor vehicle intensity in current prices, baseline model simulation for agriculture (red solid line) and non-agriculture (blue dashed line), 1995-2016.	54
1.12	Structural transformation, baseline model simulation (blue dashed line) and data (red solid line), 1970-2016. Grey shaded area shows model calibration period.	55
1.13	Motor vehicle intensity in constant prices, baseline model simulation (blue dashed line) and data (red solid line), 1970-2016. Grey shaded area shows model calibration period.	55
1.14	Simulated and observed motor vehicle intensity in constant prices, main panel 2010 (Great Britain = 1). Red dashed line shows 45-degree line.	56
1.15	Agriculture employment shares, 1970-2010, France (red solid line), South Korea (blue dashed line) and China (light blue dot-dashed line). Source: Author's analysis; ILO 2020.	58
1.16	Baseline simulations of motor vehicle intensity for France, South Korea and China. Sources: Author's analysis; Palgrave Macmillan Ltd (2013b) 60	
1.17	Simulations 1970-2120 using baseline model (red solid line), one-sector model (dark blue dot-dashed line), and later structural transformation model (light blue dashed line).	65
1.18	Estimated country fixed effects and 95 per cent confidence intervals from regression model of car intensity, main panel 1970-2010. Standard errors clustered at country level. Sources: Author's analysis; Palgrave Macmillan Ltd 2013b.	70
1.19	LOWESS, car intensity and structural transformation, main panel 1970-2010. LOWESS carried out using running-line least-squares smoothing. Points show country-year observations. Red solid line shows smoother with bandwidth = 0.8. Blue dashed line shows bandwidth = 0.6. Orange dot-dashed line shows bandwidth = 0.4. Sources: Author's analysis; ILO 2020; Palgrave Macmillan Ltd 2013b.	71
2.1	Cargo volumes by cargo type, 2013-2022. Sources: Author's analysis; Eurostat 2022a.	103
2.2	Total cargo volumes by partner country, Irish main ports 2018. Projection: WGS84. Sources: Author's analysis; Eurostat 2022a; Esri 2022.	104
2.3	Total cargo volumes by partner entity, 2013-2022. Sources: Author's analysis; Eurostat 2022a.	104
2.4	Ro-Ro cargo volumes by partner country, Irish main ports. Projection: WGS84. Sources: Author's analysis; Eurostat 2022a; Esri 2022.	105
2.5	Regression difference-in-differences estimates and 95 per cent confidence intervals for Brexit effect on quarterly cargo volumes, generalised models 2013-2022. Exponentiated coefficients. Robust standard errors clustered at partner country level and transformed using delta method. Sources: Author's analysis; Eurostat 2022a.	111

2.6	Typical routes for Ireland-France Ro-Ro cargo. Land-bridge route involves Ro-Ro vessel from Dublin to Holyhead, road truck from Holyhead to Dover, and Ro-Ro vessel from Dover to Calais. Direct route involves Ro-Ro vessel from Rosslare to Cherbourg. Projection: WGS84. Sources: Author's analysis; Eurostat 2022a.	114
2.7	Intensity factors of road freight (truck with gross vehicle weight of 10-20 tonnes towing a medium load on a motorway) and short sea shipping (Ro-Ro vessel with deadweight tonnage of 5,000-9,999 and medium to heavy load). Emissions figures are for emissions arising from fuel combustion during vehicle or vessel use. Sources: Author's analysis; CE Delft 2021; International Maritime Organization 2021.	115
2.8	Implicit tariff on UK land-bridge route due to Brexit depending on elasticity of substitution between land-bridge and direct route. Source: Author's analysis.	119
3.1	Frequency of Google searches for 'uber' search term, 2012-2019. Source: Author's analysis; Google Trends 2022.	129
3.2	Population by city, 2012 and 2019. Sources: Author's analysis; Office for National Statistics 2020.	146
3.3	Journeys starting or ending in Edinburgh, Glasgow, Dundee or Aberdeen, 2012-2019. Line weight is defined by the number of journeys between two local authorities. Point size is defined by the number of journeys within a city. Source: Author's analysis; SHS 2020.	153
3.4	Journey main mode choice by city, 2012-2019. Car as passenger includes taxi. Other includes bicycle, motorcycle/moped, ferry, air, horse-riding and tram. Sources: Author's analysis; SHS 2020.	157
3.5	Journey main mode choice and mean of calculated average journey speeds, 2012-2019. Control group comprised of Dundee and Aberdeen. Ride hailing became available in Glasgow and Edinburgh between 2015 and 2016 surveys (grey dashed line). Sources: Author's analysis; SHS 2020.	158
3.6	Percentage of journeys by assigned city, before and after availability of ride hailing. Sources: Author's analysis; SHS 2020.	159
3.7	Regression difference-in-differences estimates coefficients and 95 per cent confidence intervals for effect of ride hailing effect on public transport, generalised model 2012-2019. Exponentiated coefficients. Robust standard errors clustered at individual level and transformed using delta method. Uber operating in Glasgow and Edinburgh from 2016 onwards (red dashed line). Sources: Author's analysis; SHS 2020.	166
3.8	Regression difference-in-differences estimates coefficients and 95 per cent confidence intervals for effect of ride hailing effect on average journey speed of road-based journeys, generalised model 2012-2019. Robust standard errors clustered at individual level. Uber operating in Glasgow and Edinburgh from 2016 onwards (red dashed line). Sources: Author's analysis; SHS 2020.	167

LIST OF FIGURES

3.9	Use of car as passenger as journey stage among public transport journeys, before and after availability of ride hailing. Car as passenger includes taxi. Sources: Author’s analysis; SHS 2020.	168
A.1	Using OECD Input-Output Table data on motor vehicle inputs to calculate a proxy measure of sectoral car intensity.	185
A.2	Ternary plot of sectoral labour shares and aggregate car intensity, main panel 1970-2019. Triangle sides show percentages of total employment in agriculture (agri.), industry (ind.) and services (ser.) sectors. Points, representing country-year pairs, are coloured by aggregate car intensity in private motor vehicles per million USD of real PPP GDP. Sources: Author’s analysis; ILO 2020; Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.	192
A.3	Changes in sectoral motor vehicle intensity over structural transformation as regression lines, main panel 1995-2016. X-axis extended over full domain and reversed for illustration. Sources: Author’s analysis; OECD 2021; ILO 2020; Feenstra, Inklaar, and Timmer 2015.	194
A.4	Market system in my economic model.	195
A.5	Consumer’s utility maximisation problem.	196
A.6	Firm’s production function.	197
B.1	Total cargo volumes by partner country, Irish main ports 2018. Line width determined by cargo volume. Projection: WGS84. Sources: Author’s analysis; Eurostat 2022a; Esri 2022.	203
B.2	Cargo volumes by partner entity, EU main ports 2013-2022. Sources: Author’s analysis; Eurostat 2022a.	205
B.3	Cargo volumes by partner entity, Irish main ports 2013-2022. Sources: Author’s analysis; Eurostat 2022a.	206
B.4	Main ports in Ireland. Projection: TM65/Irish Grid. Sources: Author’s analysis; Marine Institute 2023; Eurostat 2022a.	207
B.5	Regression difference-in-differences estimates and 95 per cent confidence intervals for Brexit effect on quarterly cargo volumes, generalised models 2013-2022. Robust standard errors clustered at partner country level. Sources: Author’s analysis; Eurostat 2022a.	215
C.1	Flow chart of study sample size from Scottish Household Survey 2012-2019. Sources: Author’s analysis; SHS 2020.	218
C.2	Random adult characteristics, control and treatment groups 2012-2019. Control group comprised of Dundee and Aberdeen. Axis shows percentage of random adults. Sources: Author’s analysis; SHS 2020.	224
C.3	Age and total household net annual income distributions of random adults, control and treatment groups 2012-2019. Control group comprised of Dundee and Aberdeen. Sources: Author’s analysis; SHS 2020.	225
C.4	Socio-demographic characteristics of random adults, 2012-2019. Sources: Author’s analysis; SHS 2020.	226

C.5	Regression difference-in-differences estimates coefficients and 95 per cent confidence intervals for effect of ride hailing effect on public transport, generalised model 2012-2019. Robust standard errors clustered at individual level. Uber operating in Glasgow and Edinburgh from 2016 onwards (red dashed line). Sources: Author’s analysis; SHS 2020. . . .	230
C.6	Roads and public transport stops, 2020. Source: Author’s analysis; UK Department for Transport 2020a; Ordnance Survey 2020.	231
C.7	Population density and density of transport infrastructure by city. Source: Author’s analysis; Office for National Statistics 2020; Ordnance Survey 2020; UK Department for Transport 2020a, 2020b.	234
C.8	Transport infrastructure in Scotland, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author’s analysis; Transport Scotland 2020; UK Department for Transport 2020b.	236
C.9	Transport prices in the UK, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author’s analysis; UK Department for Energy Security and Net Zero 2023c; UK Department for Transport 2023; Office of Rail and Road 2022.	238
C.10	Aggregate measures of car use by city or region, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author’s analysis; UK Department for Transport 2022; Transport Scotland 2020.	240
C.11	Aggregate measures of public transport use by city or region, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author’s analysis; Transport Scotland 2020.	244

LIST OF TABLES

1.1	Private motor vehicles per GDP and real GDP per capita in extended panel 1950-2010	25
1.2	Changes in sectoral motor vehicle inputs over process of structural transformation for main panel 1995-2016	34
1.3	Estimation of sectoral motor vehicle-labour elasticities of substitution in main panel 2016	47
1.4	Calibrated parameter values: main panel aggregates 1995-2016	49
1.5	Cross-sectional model fit for main panel 2010	57
1.6	Regression model of car intensity across main panel 1970-2010	72
2.1	Descriptive statistics for total trade cargo weights of EU-27 main ports 2013-2022	101
2.2	Descriptive statistics for total trade cargo weights of Irish main ports 2013-2022	101
2.3	Regression difference-in-differences estimates of Brexit effect on cargo in EU-27 main ports 2013-2022	107
2.4	Regression difference-in-differences estimates of Brexit effect on cargo in Irish main ports 2013-2022	109
2.5	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022	109
3.1	Descriptive statistics of included journeys, 2012-2019	152
3.2	Descriptive statistics of included random adults, 2012-2019	154
3.3	Descriptive statistics for continuous variables of included journeys 2012-2019	157
3.4	Regression difference-in-differences estimates of ride hailing effect on choice of main mode	161
3.5	Regression difference-in-differences estimates of ride hailing effect on choice of main mode	163
3.6	Regression difference-in-differences estimates of ride hailing effect on journey duration, distance and speed	164
A.1	Mapping International Standard Industrial Classification of All Economic Activities (Revision 4) sections to broad sectors	181
A.2	Mapping International Comparison Programme 2005 categories to broad sectors	183
A.3	Descriptive statistics of selected variables for main panel 1970-2010	191
A.4	Correlation coefficients of selected variables for main panel 1970-2010	191

A.5	Changes in sectoral motor vehicle inputs over process of structural transformation for main panel 1995-2016	193
A.6	Re-calibrated one-sector parameter values: main panel aggregates 1995-2016	200
B.1	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022	207
B.2	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022	208
B.3	Regression difference-in-differences estimates of Brexit effect on cargo type shares in EU-27 main ports 2013-2022	209
B.4	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022	210
B.5	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022	211
B.6	Regression difference-in-differences estimates of Brexit effect on cargo in EU-27 main ports 2013-2022	212
B.7	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022	212
B.8	Regression difference-in-differences estimates of Brexit effect on cargo type shares in EU-27 main ports 2013-2022	213
B.9	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022	213
B.10	Regression difference-in-differences estimates of Brexit effect on cargo in Irish main ports 2013-2022	214
B.11	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022	214
B.12	Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022	215
C.1	Collapsing categorical variable of mode choice	219
C.2	Socio-demographic characteristics of Glasgow treatment group and control group 2012-2015	221
C.3	Socio-demographic characteristics of Edinburgh treatment group and control group 2012-2015	222
C.4	Socio-demographic characteristics of Glasgow treatment group in 2012 and 2019	222
C.5	Socio-demographic characteristics of Edinburgh treatment group in 2012 and 2019	223
C.6	Socio-demographic characteristics of control group in 2012 and 2019	223
C.7	Regression difference-in-differences estimates of ride hailing effect on choice of main mode	227
C.8	Regression difference-in-differences estimates of ride hailing effect on choice of public transport as main mode in Glasgow	228
C.9	Regression difference-in-differences estimates of ride hailing effect on choice of main mode	241

LIST OF TABLES

C.10 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	242
C.11 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	243
C.12 Post-treatment coefficient for use of car as passenger as part of journey	244
C.13 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	245
C.14 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	246
C.15 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	247
C.16 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	248
C.17 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	249
C.18 Regression difference-in-differences estimates of ride hailing effect on choice of main mode	250

INTRODUCTION

The United Nations (UN) identified sustainable transport as being at the heart of sustainable development and a central enabler of a range of the Sustainable Development Goals (UN 2021c). In the United Kingdom (UK), transport accounted for 34 per cent of all carbon emissions in 2022, with the majority of this stemming from road transport (UK Department for Energy Security and Net Zero 2023a). Therefore, progressing towards sustainable transport will be critical to meeting the Paris Climate Change Agreement's target of limiting the global temperature rise to 1.5 degrees Celsius (Masson-Delmotte et al. 2018).

Have we reached 'peak car'? Did Brexit inadvertently cause a shift from road freight transport to the more energy efficient mode of short sea shipping? Can ride hailing platforms advance the sustainability of transport? These questions are all explored in this thesis.

Over the past century, private cars have become central to the way we live across much of the world. But how does an economy's dependency on car use evolve as it gets richer? What role does structural transformation, the reallocation of economic activity between the broad agriculture, manufacturing and services sectors (Herrendorf, Rogerson, and Valentinyi 2014), play in this? I investigate these questions from a macroeconomic perspective in Chapter 1, '**Driving over the hill: Car intensity during structural transformation**'. This is motivated by the debate over whether we have entered peak car, with evidence that private

car use per capita has now peaked in developed economies (see, for example, Bastian, Börjesson, and Eliasson 2016; Millard-Ball and Schipper 2011). Whether we have arrived at peak car is important, as a decrease in car use will be required alongside advances in fuel efficiency and electric vehicle technology to reduce carbon emissions. Rather than studying car use on a per capita basis, I look at the number of licensed private cars required in the economy to produce one unit of gross domestic product (GDP), or ‘car intensity’, as an intrinsic component of car use that may be underlying observed peak car trends.

First, using data collated for a panel of 88 countries from 1950 to 2010, I demonstrate that car intensity evolves in a hump-shaped pattern over the course of economic development. In other words, as GDP per capita increases, car intensity initially rises sharply before reaching a peak and later decreasing. This dynamic, in turn, produces a hump-shaped relationship between the level of car use per GDP and GDP per capita.

Structural transformation has been highlighted as a key feature of modern economic growth (Kuznets 1973). Could this global phenomenon be involved in generating the hump shape in car intensity? Second, therefore, I develop a multi-sector general equilibrium model to argue that structural transformation can have a fundamental role in producing this relationship between car intensity and GDP per capita. The model is similar to the models of structural transformation proposed by Duarte and Restuccia (2010) and Stefanski (2014), with a shift in labour away from agriculture engendered by non-homothetic preferences for the agriculture good. I incorporate motor vehicles in the model as an input to production alongside labour, with a sector-specific elasticity of substitution between motor vehicle and labour inputs measuring how easy it is to switch between these factors. Due to an endogenous evolution in the aggregate elasticity of substitution during structural transformation, my calibrated model reflects the

observed hump-shaped relationship between car intensity and GDP per capita and can account for just under a quarter of observed variation in car intensity among a cross section of 54 countries in 2010.

Third, I conduct counterfactual exercises using this model to demonstrate the role of structural transformation in inducing a hump shape in intensity as the economy develops. Unlike my baseline model, which is characterised by structural transformation from an agriculture sector to a non-agriculture sector, a one-sector model with no structural transformation fails to produce the pattern. In addition, I find economies that develop later reach a peak level of car intensity at a later stage, but that this peak occurs at a lower level than that experienced by countries that shifted away from agriculture earlier. This finding is relevant to developing economies that have not yet arrived at a peak level of car intensity. Finally, I also carry out a semi-parametric regression analysis to empirically test the relationship between car intensity and structural transformation, again detecting a hump-shaped association.

Over half of road transport emissions in the UK is accounted for by passenger cars. However, another significant source of transport emissions is road freight, with over 15 per cent of UK road transport emissions arising from heavy goods vehicles in 2019 (UK Department for Business, Energy and Industrial Strategy 2022). In Europe, one approach to reducing transport emissions is a policy aim to replace road freight transport with more efficient modes, including ‘short sea shipping’, where possible (European Commission 2011).

Transporting goods through the UK has traditionally been a popular route for trade between Ireland and continental Europe, known as the UK ‘land-bridge’, due to its superior transit time and flexibility relative to the direct short sea shipping routes (Breen et al. 2018; Vega, Feo-Valero, and Espino-Espino 2018; Lawless

and Morgenroth 2017). However, since the start of 2021, Brexit imposed non-tariff barriers to trade such as customs checks on these cargo volumes. How did this affect cargo volumes on the land-bridge route? In Chapter 2, '**All at sea? Brexit, shipping, and the UK land-bridge**', I trade the wide-angle viewpoint of macroeconomics for the more telescopic lens of microeconometrics and subject this question to a difference-in-differences analysis.

Several studies have analysed the implications of Brexit for trade values (for example, Kren and Lawless 2022; Freeman et al. 2022; Flynn, Kren, and Lawless 2021a). I aim to furnish this literature with evidence from the perspective of the transport sector by conducting a difference-in-differences analysis of maritime cargo volumes. First, I examine quarterly Eurostat port-level data from 2013 to 2022 using pseudo-Poisson maximum likelihood regressions to assess the effect of Brexit on short sea shipping. Most cargo shipped between the UK and the European Union (EU) is roll-on roll-off (Ro-Ro) or liquid bulk cargo. While I detect no effect on EU-UK liquid bulk volumes, I find a 22 per cent decrease in EU-UK Ro-Ro cargo volumes that can be attributed to Brexit. This effect is more pronounced for Ireland-UK cargo volumes, with a 54 per cent decrease in Ro-Ro cargo between Ireland and the UK due to Brexit. Interestingly, I find a concurrent increase of 147 per cent in this form of cargo between Ireland and France, indicating a diversion from the land-bridge to direct short sea shipping routes.

Of course, this shift has implications for the energy consumption and carbon emissions associated with transporting goods between Ireland and continental Europe. Second, therefore, I present rough calculations, based on some simplifying assumptions, comparing the two routes in terms of the energy consumption and emissions stemming from transporting 1 tonne of cargo. I estimate that energy consumption and emissions would be approximately 60 per cent lower on the direct route than on the land-bridge, highlighting the greater sustainability of

transporting cargo using short sea shipping instead of road transport.

Third, I model a firm exporting goods from Ireland to continental Europe and its decision to allocate goods between the land-bridge and the direct route. I calibrate this model to reflect observed cargo volumes prior to Brexit, and then use the model to derive an ‘implicit tariff’ imposed by Brexit via non-tariff trade barriers such as customs checks. With this approach, I show that the implicit tariff on the land-bridge route is dependent on how easy the firm can switch between the two routes, or its elasticity of substitution, with higher implicit tariffs where substitution is more difficult. For example, the implicit Brexit tariff would be higher for time-sensitive goods that are particularly dependent on the superior transit time of the land-bridge, such as agri-food products.

In Chapter 3, I return to passenger transport. Chapter 1 contributes to the peak car debate from a global perspective. At an individual level, the way we own and use private cars is also evolving. For example, ride hailing platforms such as Uber have been increasing in popularity over the past decade. Whether the emergence of these platforms can be viewed as helpful in transitioning to sustainable transport hinges on their impact on other transport modes. Does ride hailing substitute or complement public transport? Ride hailing could represent an alternative that tempts individuals away from public transport (Webb 2019), or it could help alleviate public transport’s ‘last-mile’ problem by reducing the time it takes to get to the nearest train station (Stiglic et al. 2018). Using microeconomic tools, I turn attention to this question in Chapter 3, **‘All hail? The impact of ride hailing platforms on the use of other transport modes’**.

Drawing on travel diary data from 16,712 individuals in the Scottish Household Survey between 2012 and 2019, I employ multinomial logistic regressions to compare the average change in mode choice in Glasgow and Edinburgh, where

Uber became available in 2015, with the average change in two cities without ride hailing, Dundee and Aberdeen. I find a small complementary effect on the use of public transport relative to driving a car in Glasgow, although this is not reflected in the smaller city of Edinburgh. The complementary effect in Glasgow appears to be more pronounced among individuals who are younger, male, employed and with higher levels of household education and income. This research contributes to a literature that is important in determining whether innovations like ride hailing are useful in enhancing the sustainability of transport (Tirachini and Gomez-Lobo 2019).

Global warming and climate change present us with unprecedented challenges. Given its generous contribution to emissions, the transport sector can play a pivotal role in reducing the sizeable carbon footprint associated with the way we live. Drawing on methods from the arsenals of both macroeconomics and microeconometrics, this thesis aims to broaden our understanding of several topics, including peak car, short sea shipping and ride hailing, that will be important in the journey towards sustainable transport.

DRIVING OVER THE HILL?

Car intensity during structural transformation

1.1. Introduction

Transport accounted for 34 per cent of all carbon dioxide emissions in the United Kingdom (UK) in 2022, with the majority of this stemming from road transport (UK Department for Energy Security and Net Zero [2023a](#)). This is thus a key sector where a reduction in emissions is required to meet the Paris Climate Change Agreement's target of limiting the global temperature rise to 1.5 degrees Celsius (Masson-Delmotte et al. [2018](#)). The United Nations (UN [2021c](#)) has also identified sustainable transport as a key enabler of a range of Sustainable Development Goals. Reductions in motor vehicle use will be required alongside advances in fuel efficiency and in electric vehicle technology to ensure a decrease in carbon emissions.¹ In addition to reduced carbon emissions, a reduction in car use can

¹While the fuel efficiency of motor vehicles has improved substantially over time, it has been pointed out that such improvements can often be negated by increases in vehicle power, vehicle weight or distance travelled in the form of a rebound effect (Stapleton, Sorrell, and Schwanen [2017](#)). The advent of hybrid and electric vehicles can also contribute to lower carbon emissions, although this is dependent on the decarbonisation of electricity generation (Bahamonde-Birke [2020](#)).

1. DRIVING OVER THE HILL?

contribute to congestion relief, improved air quality, increased active travelling, increased wellbeing and improved accessibility (Goodwin 2020).

This study is focused on the level of car use as a country develops, and asks the following questions: How does the dependency of an economy on car use evolve as the economy develops? What role does structural transformation, the redistribution of economic activity among the broad agriculture, manufacturing and services sectors (Herrendorf, Rogerson, and Valentinyi 2014), play in this relationship? In considering these questions, I analyse trends in car use per unit of gross domestic product (GDP), or the level of car use required to produce one unit of GDP, as GDP per capita rises. This can be referred to as the intensity of car use. Uniquely, I will show that among developed economies, the intensity of car use displays an inverted U-shaped or hump-shaped relationship with GDP per capita, increasing initially before peaking and later decreasing. The study will then explore this dynamic by developing a theoretical model of structural transformation capable of reflecting the hump-shape in car use per GDP.

1.1.1. Peak car

While specifically considering car use per unit of GDP, this study is closely related to the debate in transport literature surrounding 'peak car'. The peak car phenomenon has been described as a levelling off in various measures of car use, such as annual per-capita distance travelled in cars and per-capita car ownership, which has been becoming evident in many developed cities and countries (Webb 2019; Bastian, Börjesson, and Eliasson 2016; Goodwin and Van Dender 2013; Millard-Ball and Schipper 2011). Studies including Puentes and Tomer (2008), Lucas and Jones (2009) and Schipper (2009) began to note a stabilisation of growth in car use per capita, and Millard-Ball and Schipper (2011) coined the phrase 'peak car' to describe this emerging trend.

Millard-Ball and Schipper (2011) collated data from 1970 to 2007 from various sources to show that private vehicle use across 8 industrialised countries had plateaued. A descriptive analysis of household travel survey microdata indicated that car use per capita peaked in Paris, Berlin, London and Vienna in the 1990s, and in Copenhagen in the late 2000s (Wittwer, Gerike, and Hubrich 2019). In most developed countries, car use per capita has been viewed as having stagnated around the 2000s (Bussière, Madre, and Tapia-Villarreal 2019), and Goodwin (2020) noted that there was increasing evidence by around 2010 that per-capita car use was reaching a peak level. King (2020), however, regarded the concept of peak car as an open debate rather than an established phenomenon, suggesting that declining car vehicle kilometres (VKM)² in the early 2010s could have been related to economic activity.

Naturally, while peak car is a phenomenon that has become increasingly evident among many developed countries, there has been some interest in whether such patterns could be expected to emerge in countries or cities at an earlier stage of development. Gao and Newman (2018) showed descriptive statistics of car use in the emerging city of Beijing and indicated that car use reduced during the 2010s while GDP continued to increase. Using household survey and population projection data in the developed cities of Lille and Montreal, and in the developing cities of Juarez and Puebla, Bussière, Madre, and Tapia-Villarreal (2019) projected that peak car may be reached in Juarez and Puebla in 2030 based on previous trends in Lille and Montreal.

Numerous factors have been proposed in the literature to explain peak car. One

²A common measure of distance travelled by transport mode is vehicle kilometres (VKM), which represents the movement of one vehicle over one kilometre and can be calculated by multiplying the number of vehicles on a road by the average length of their trips. This may be estimated from traffic volume data, or from survey data related to vehicles or their drivers (Jamroz and Wachnicka 2015). An alternative measure is passenger kilometres (PKM), employed as a measure of passenger distance. This is calculated by multiplying the total distance travelled by the number of passengers transported.

1. DRIVING OVER THE HILL?

popular factor has been that there must exist a saturation level of car use, either through a saturation of travel demand or via technological constraints. While Millard-Ball and Schipper (2011) did not empirically investigate possible factors, they pointed out that there must be some saturation level for car use given time constraints. Historically, new infrastructure or new technologies have led to increased travel speeds, allowing increased car use for a given time allocation, whereas this increase may have since levelled off (Millard-Ball and Schipper 2011).

Metz (2010) had also made this argument using descriptive evidence from the UK National Travel Survey. The existence of a saturation level of car use was also discussed as a possible factor by Newman and Kenworthy (2011), while Metz (2013) argued that a saturation level of travel demand must exist on the basis that the benefits of travelling longer distances within a given time constraint to widen destination choice would be subject to diminishing marginal returns. Again, however, these hypotheses were not empirically tested in these studies. An analysis of car travel in France was conducted by Grimal, Collet, and Madre (2013) in which saturation was assumed to be a function increasing in income and decreasing in fuel price and population density.

Given the clear link between VKM and car ownership identified by many studies (for example Jamroz and Wachnicka 2015; Van Acker and Witlox 2010; De Jong et al. 2004), the proposed existence of a saturation level of car use is closely related to a previously hypothesised saturation point in car ownership. Tanner (1978) discussed a theoretical limit of one car per person among the population old enough to drive a car, and pointed to evidence that growth rates in car ownership declined as car ownership increased. That car ownership continued to predict vehicle kilometres travelled by car was also highlighted by Tanner (1978) as an indication that this saturation point existed, as individuals owning multiple cars would imply that average VKM per car would decrease. Tanner (1978) also noted,

however, that there was evidence of different levels of saturation in different locations.

Another peak car factor discussed in the literature has been public transport. Newman, Kenworthy, and Glazebrook (2013) presented descriptive statistics using global city level data to argue that rail systems in urban areas have been expanding, increasing in both service levels and utilisation. The study also suggested that the speed of rail in urban areas had been increasing relative to car transport, and that this may be a factor in peak car patterns.

Peak car patterns have also been attributed to a decline in car use among young adults. Delbosc and Currie (2013) reviewed several papers that identified a decrease in young adults holding a driving license. Kuhnimhof et al. (2012) presented National Travel Survey data from 6 industrialised countries that showed a fall in driving license holding and car ownership among young adults, with a more pronounced decline among young men. In an empirical analysis using census data for England, Melia, Chatterjee, and Stokes (2018) found that increasing population density was leading to lower car ownership and more public transport utilisation among young adults. Other factors in these trends among young adults that have been suggested include increasing enrolment in tertiary education and a corresponding lower labour force participation (Delbosc and Currie 2013; Kuhnimhof et al. 2012). However, Kuhnimhof, Zumkeller, and Chlond (2013) noted that while younger cohorts appeared to be contributing most to peak car trends, increased car ownership and driving license holding among older adults was having the opposite effect.

Urbanisation has also been identified as a more general factor in peak car by other studies. An econometric analysis of aggregate time series data since 1970 for Great Britain conducted by Stapleton, Sorrell, and Schwanen (2017) found that the

1. DRIVING OVER THE HILL?

levelling off in car VKM could be explained by increasing levels of urbanisation, in addition to increasing fuel prices and decreases in income following the 2008 financial crisis. Headicar (2013) analysed census and household travel survey data in England from 1970 to show a high level of variation in distances travelled by car between areas depending on their level of urbanisation, and noted that peak car patterns appeared to coincide with a reversal of previous migration trends from urban to suburban locations.

The role of more fundamental economic variables such as income or fuel prices in peak car patterns has not been universally settled in this literature. Goodwin and Van Dender (2013) noted a theme that aggregate models of car travel demand that focused solely on GDP per capita and fuel prices were dismissed as too crude and incapable of explaining more recent trends. Several early studies of peak car including Puentes and Tomer (2008), Metz (2010), Newman and Kenworthy (2011) and Millard-Ball and Schipper (2011) pointed out that the plateau in car use preceded the Great Recession and significant increases in oil prices in 2008, and that other forces must therefore be at play. However, Grimal, Collet, and Madre (2013) found that peak car patterns were evident to some extent across all income groups and levels of urbanisation in France, and argued that this suggested that the main factor must have been a significant and widespread change in economic conditions, such as increases in fuel prices. Using data from 6 developed countries for the period between 1980 and 2013, Bastian, Börjesson, and Eliasson (2016) presented regression modelling that indicated peak car patterns could, in fact, be predicted using only income and fuel price as independent variables, while stressing that this did not preclude other factors from also being influential. The study also found evidence of a declining elasticity of car travel with respect to GDP as GDP per capita increased, which was viewed as consistent with the existence of a saturation level of car use or car ownership. Similar results were also shown

specifically for Sweden by Bastian and Börjesson (2015).

Some studies have revealed within-country regional differences in peak car patterns. Metz (2015) showed a stronger decline in car use per capita in London than in the rest of the UK, and suggested that this could have been a reflection of congestion levels or differences in local transport policy. In Sweden, Bastian and Börjesson (2015) highlighted differences in the sensitivity of car use to GDP and fuel prices between the two main cities and the rest of the country, with a higher elasticity evident in cities. The study pointed out that this was consistent with cities offering more public transport and a higher density of possible destinations.

A plethora of other explanations for peak car have been suggested in the literature, including the substitution of car travel by electronic communication (van Wee 2015), lower subsidisation of company cars (Le Vine, Jones, and Polak 2013) and a switch to air travel (Millard-Ball and Schipper 2011). In summary, a consensus in the literature on the factors behind the apparent levelling off in car use per capita among developing countries has remained elusive. As a result of this, the related debates over whether car use has in fact reached a peak level in developed economies and whether it can be expected to decrease into the future have also persisted without agreement.

Considering distance travelled per unit of GDP, as I do in this study, is another lens through which to observe the peak car phenomenon since a persistent idea in the literature has been that car use is strongly related to GDP. The level of car use required in an economy to produce a single unit of GDP could also be thought of as a fundamental factor underlying trends in per-capita car use. Kenworthy (2013) presented a descriptive analysis of city-level data from 1995-1996 and 2005-2006 across the United States (US), Australia, Canada, Europe and Asia that showed reductions in car VKM per unit of GDP, indicating a decoupling of urban car use

from GDP, although the study did not empirically test factors behind these trends. Kenworthy (2013) also found that GDP per capita was actually a poor predictor of car VKM per capita among the cities included in the dataset. A high level of cross-sectional variation in VKM per GDP was also apparent, with cities in the US and Australia showing higher levels than cities in Europe, Canada and Asia, despite the general downward trend across all cities (Kenworthy 2013).

1.1.2. Structural transformation

A separate literature that is relevant to this study has focused on the global phenomenon of structural transformation. This is the reallocation of economic activity among the broad agriculture, manufacturing and services sectors (Herrendorf, Rogerson, and Valentinyi 2014), and has been highlighted as a key feature of modern economic growth (Kuznets 1973). Rogerson (2008) and Herrendorf, Rogerson, and Valentinyi (2014) noted two distinct approaches to modelling structural transformation in multi-sector general equilibrium models that are founded on representative agents. First, in the demand approach, consumer preferences over goods are specified as being non-homothetic (using a form of Stone-Geary preferences), and this gives rise to an ‘income effect’ with consumers devoting a decreasing share of their budget to agriculture goods as they get richer (for example, Comin, Lashkari, and Mestieri 2021; Boppart 2014; Gollin, Parente, and Rogerson 2002; Kongsamut, Rebelo, and Xie 2001; Echevarria 1997). Second, in the productivity approach, technological growth is considered to be uneven across sectors, leading to a ‘substitution effect’ or ‘price effect’ with labour moving between sectors (for example, Acemoglu and Guerrieri 2008; Ngai and Pissarides 2007).

Duarte and Restuccia (2010) proposed a three-sector model including agriculture, manufacturing and services that allowed both of these structural transformation

channels to exist and that was calibrated to US data between 1956 and 2004. Their model was a sequence of static problems, abstracting from household intertemporal choice. They showed that a reallocation of labour could be generated by both channels working in tandem, or, under certain conditions, by either channel in isolation. Studies of structural transformation such as Michaels, Rauch, and Redding (2012), Rogerson (2008) and Gollin, Parente, and Rogerson (2002) have proposed two-sector models by considering the reallocation of labour from agriculture into a combined 'non-agriculture' sector that includes both manufacturing and services. Other studies that developed multi-sector models of structural transformation include Ngai and Pissarides (2007) and Kongsamut, Rebelo, and Xie (2001).

Motor vehicles have not been explicitly considered in these multi-sector general equilibrium models, with labour, land or capital typically included as production inputs and agriculture, manufacturing or services goods available to consumers. To broaden understanding on the relationship between structural transformation and oil prices, Stefanski (2014) proposed a multi-sector, multi-country model that incorporated oil as a production input alongside labour, while Stefanski (2017) included 'modern energy' as an input to production.

It is also important to consider urbanisation in any analysis involving economic development, given a very strong correlation between income per capita and urbanisation among developing countries (Gollin, Jedwab, and Vollrath 2015; Michaels, Rauch, and Redding 2012). Urbanisation has been typically considered as a companion to industrialisation, with historical patterns of cities developing where factories were located. Michaels, Rauch, and Redding (2012) provided empirical evidence in relation to urbanisation in the US between 1880 and 2000 and used a theoretical model to demonstrate the central influence of structural transformation in this process. Their model of structural transformation, including

1. DRIVING OVER THE HILL?

production of an agriculture and a non-agriculture good, was able to predict that only the most densely populated locations produced the non-agriculture good. This was due to a combination of two forces. First, there were weaker agglomeration forces in agriculture than in non-agriculture due to the land-intensive nature of agriculture. Second, a location's population density increased in its productivity, and non-agriculture was characterised by lower mean reversion in productivity (Michaels, Rauch, and Redding 2012).

Furthermore, Brunt and García-Peñalosa (2021) developed a model, calibrated to reflect developments in England between 1550 and 1850, to argue that urbanisation can itself induce structural transformation, with greater innovation and knowledge exchange occurring in denser locations leading to a positive feedback loop between structural change and urbanisation. Using a model calibrated to data on cities in France from 1995 to 2018, Chen et al. (2023) demonstrated that the later structural change from manufacturing to services was more pronounced in cities with higher population density.

Gollin, Jedwab, and Vollrath (2015) presented a three-sector structural transformation model in which urbanisation could also occur in the absence of industrialisation, motivated by a tendency among developing countries to experience high rates of urbanisation in a context of little industrialisation. In addition to industrialisation, their model showed that urbanisation can be driven by increases in income from the exporting of natural resources, known as a 'Dutch disease' or Balassa-Samuelson effect (Gollin, Jedwab, and Vollrath 2015). In their model, rural production consisted only of food, or an agriculture good, while urban production consisted of 'tradeable' and 'non-tradeable' sectors of non-agriculture goods. While the study noted that it was simplistic to simply equate rural production with agriculture and urban production with non-agriculture, it pointed to cross-country empirical evidence from 2000-2010 that non-agriculture accounted for only 30 per

cent of rural employment, but 92 per cent of urban employment (Gollin, Jedwab, and Vollrath 2015).

While much of this structural transformation literature has focused on models of a closed economy, structural change can also stem from factors such as international borrowing or trade. To specifically explore these elements, some studies have modelled structural transformation in an open economy (for example, Chen et al. 2023; Cravino and Sotelo 2019; Kehoe, Ruhl, and Steinberg 2018; Świącki 2017; Uy, Yi, and Zhang 2013).

1.1.3. Environmental Kuznets Curve

This study is also related to literature on the 'Environmental Kuznets Curve' (EKC). This literature stemmed from studies in the early 1970s that began to focus on linking environmental emissions to economic development, such as Forster (1973), Solow (1973) and Stiglitz (1974). Grossman and Krueger (1995) identified a hump shape in emissions per capita over the course of economic development, with emissions increasing during early stages of development but later decreasing. This pattern became known as the EKC, after the Kuznets Curve in inequality due to market forces as an economy develops. Rather than examining emissions per capita as in the standard EKC, several studies have focused on emissions per unit of GDP, or the intensity of emissions, and identified a hump-shaped pattern in intensity during economic development (Stefanski 2013; Bartoletto and Rubio 2008; Tol, Pacala, and Socolow 2006; Kander and Lindmark 2004; Lindmark 2004, 2002).

As part of their study, Copeland and Taylor (2004) conducted a review of EKC literature. They noted four main theoretical explanations for the EKC. First, as the primary source of economic growth changes over time, so too does the impact of

this growth on emissions. This possibility was discussed in an empirical analysis by Grossman and Krueger (1995), for example. Second, the EKC arises due to income effects, with demand for environmental quality changing as incomes rise, for example as in López (1994). Third, the pattern is a result of a threshold effect, with either policy being implemented or coming into effect once a certain threshold of emissions is reached. It has been argued that this effect can emerge from either political processes, as in Jones and Manuelli (1995), or opportunities for emissions abatement, as in Stokey (1998). Fourth, increasing returns to abatement have been proposed as an explanation for the EKC, as in Andreoni and Levinson (2001) for example.

However, the existence of the EKC has been disputed in the literature. For example, Stern and Common (2001) presented empirical evidence using a global sample that did not support a common EKC for sulphur emissions across countries. Other notable criticisms of the EKC literature include Stern (2004) and Carson (2010).

1.1.4. This study

This research contributes to these strands of literature with an analysis of trends in car use per unit of GDP as an economy develops. The lack of research specifically into car use per unit of GDP, potentially an elementary factor in peak car trends, represents a gap in the literature. From a methodological perspective, there also appears to be a dearth of studies that employ general equilibrium macroeconomic models based on representative agents to analyse peak car trends. Meanwhile, motor vehicles have not previously been explicitly included in general equilibrium models of structural transformation. This study offers a novel perspective on peak car by essentially linking literature on car use and car ownership with literature on general equilibrium models of structural transformation.

In Section 2, I demonstrate that car use per unit of GDP evolves in a hump-shaped pattern over the course of economic development, and that this dynamic stems mostly from car intensity. In other words, as GDP per capita increases, car intensity initially increases sharply before reaching a peak and later decreasing. I then argue that structural transformation can have a fundamental role in producing this relationship. Section 3 develops a model of structural transformation that formalises this argument, and the calibration of this model is described in Section 4.

In Section 5, baseline simulations demonstrate that my model can account for just under a quarter of observed variation in car intensity among a cross section of 54 countries. I conduct counterfactual simulations in Section 6 to demonstrate the role of structural transformation in generating a hump shape in intensity as the economy develops. I find that economies that develop later reach a peak level of car intensity at a later stage, but that this peak occurs at a slightly lower level than that experienced by countries that shifted away from agriculture earlier. I follow up this theoretical approach with a semi-parametric econometric analysis in Section 7 to provide further empirical evidence of the hump-shaped relationship between car intensity and structural transformation. Finally, conclusions from the study are summarised in Section 8.

1.2. The facts

1.2.1. Car use per GDP

First, I will demonstrate a hump-shaped relationship between passenger car use per unit of real GDP and real GDP per capita. In other words, as income rises in an economy, car use per GDP increases at first before levelling off and subsequently decreasing. Comparing economies simply using measures such

1. DRIVING OVER THE HILL?

as car use per capita cannot account for the fact that economies produce vastly different levels of GDP. For example, comparing Great Britain and Thailand, two countries with similar population sizes, will reveal substantial differences in car use per capita, but much of this variation would simply reflect differences in GDP. Instead, comparing economies in terms of car use per GDP can offer additional insight on the intrinsic role of car use in economies.

Figure 1.1 summarises data on distances travelled by passenger car, measured as PKM, per purchasing power parity (PPP) US dollar (USD) of real GDP from 1970 to 2019 among 26 Organization for Economic Cooperation and Development (OECD) countries in addition to Argentina, Bulgaria and Russia. Data sources and all included countries are detailed in Appendix A.1.³ I split country-year pairs into deciles according to real PPP GDP per capita and calculated average levels of car PKM per GDP for each decile. Figure 1.1 shows a distinct hump shape in the relationship between car PKM per GDP and GDP per capita, with average car PKM per GDP highest in the fourth decile of GDP per capita.




Figure 1.1: Car passenger kilometres per GDP and real GDP per capita, 29 countries 1970-2019. Axes transformed to \log_e scales. Points show decile averages. Sources: Author's analysis; OECD 2017; Feenstra, Inklaar, and Timmer 2015.

³These 29 countries are a subset of my 'main panel' for which PKM data was available.

One previous study found the level of car use (measured in VKM) per GDP to be lower in 2005 than in 1995 in 39 out of 42 cities studied across the US, Australia, Canada, Europe and Asia (Kenworthy 2013). The hump shape I show in Figure 1.1 using data on 29 countries spanning 50 years is a striking pattern that is worth investigating further. Focusing on a single developed economy, I combined data on distance travelled by car⁴, measured in VKM in this instance, and the number of licensed private cars from 1950 to 2019 in Great Britain (encompassing England, Scotland and Wales) from the UK Department for Transport (2021b, 2021a) with data from the Penn World Tables on GDP and population (Feenstra, Inklaar, and Timmer 2015). VKM can be decomposed into two components, average VKM per car and the number of registered cars, as follows:

$$VKM = \overline{VKM} \times N_{cars} \quad (1.1)$$

In Equation 1.1, \overline{VKM} denotes average VKM per car, while N_{cars} denotes the number of licensed cars. These components could be thought of as intensive (\overline{VKM}) and extensive margins (N_{cars}). Both sides of Equation 1.1 can then be divided by GDP, showing that VKM per unit of GDP can be expressed as the average VKM per car multiplied by the number of licensed cars per unit of GDP as in Equation 1.2. I will call VKM per GDP ‘VKM intensity’, and the number of cars per GDP ‘car intensity’, to highlight them as measures of the level of car use or the number of cars required in the economy to produce one unit of GDP.

$$\frac{VKM}{GDP} = \overline{VKM} \times \frac{N_{cars}}{GDP} \quad (1.2)$$

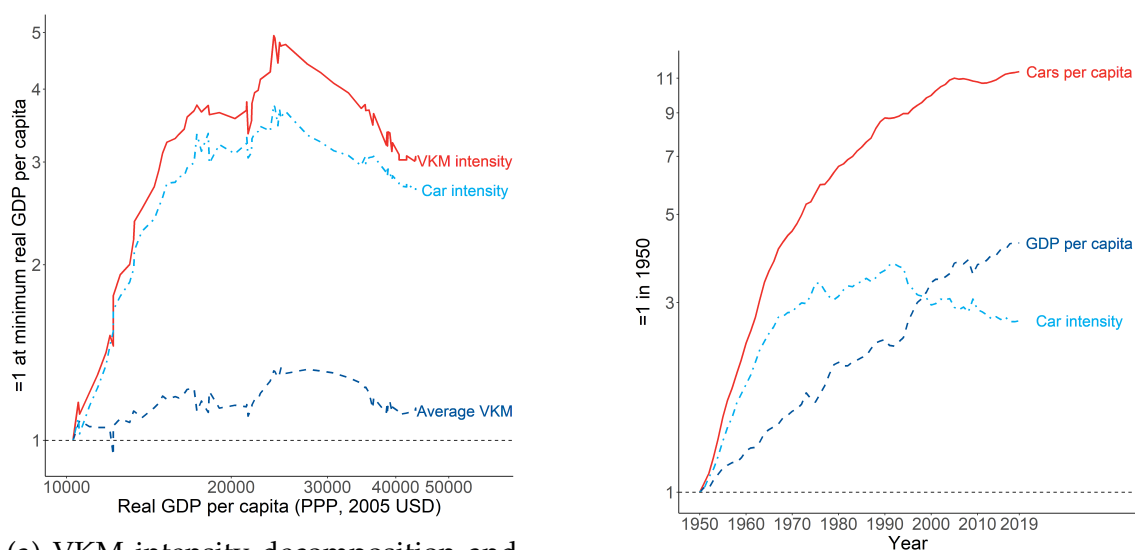
$$intensity_{VKM} = \overline{VKM} \times intensity_{car} \quad (1.3)$$

⁴This data series on distance travelled by car also included taxis.

1. DRIVING OVER THE HILL?

Figure 1.2a plots VKM intensity, average VKM per car and car intensity in Great Britain against real GDP per capita, each normalised to 1 at the minimum (and earliest) level of GDP per capita. First, the plotted VKM intensity line confirms that the hump-shaped pattern identified among a panel of countries in Figure 1.1 is also evident when analysing a single developed economy. Second, Figure 1.2a illustrates that the hump shape in the relationship between VKM intensity and GDP per capita derives mainly from the extensive margin, car intensity, with comparatively little change evident in the intensive margin, average VKM per car, during this period. While the extensive margin increased to a peak level that was 3.8 times higher than in 1950, the highest level of the intensive margin over the period was a mere 1.3 times higher than in 1950. This reveals that the bulk of the hump-shaped dynamic derives from the extensive margin. Rather than the average distance travelled by cars changing substantially, it is car intensity that initially increases but then levels off and decreases as income rises. Using survey data, Metz (2010) found that average distance travelled per person had been changing significantly in the UK between 1972 and 2008, although this variable was not limited to car travel. The finding in Figure 1.2a is consistent with previous research that found VKM to depend heavily on car ownership (Jamroz and Wachnicka 2015; Van Acker and Witlox 2010; De Jong et al. 2004).

Similarly, car intensity itself can be further decomposed into two per capita components: the number of cars per capita and GDP per capita. Figure 1.2b shows the evolution of these variables over time in Great Britain, each normalised to 1 in 1950. The dramatic rise in the number of licensed private cars per capita, increasing more than ten-fold over 50 years, is an illustration of the period of 'mass motorisation' that followed the Second World War. During this period in Great Britain, cars were viewed as a marker of affluence and models such as the Mini and Ford Cortina became accessible to a wider population including manual



(a) VKM intensity decomposition and real GDP per capita. Axes transformed to \log_e scales.

(b) Car intensity decomposition. Y-axis transformed to \log_e scale.

Figure 1.2: Car VKM intensity and private car intensity, Great Britain 1950-2019. Sources: Author's analysis; UK Department for Transport [2021b](#), [2021a](#); Feenstra, Inklaar, and Timmer [2015](#).


workers (Gunn [2018](#)). The hump shape in car intensity, with a peak level reached in the early 1990s, appears to stem from changes in the rates of increase of cars per capita and GDP per capita.

1.2.2. Car intensity

The hump-shaped pattern in car intensity, shown in Figures 1.2a and 1.2b for Great Britain, can also be found among a larger sample of countries. Combining Penn World Table data on GDP and population (Feenstra, Inklaar, and Timmer [2015](#)) with data on the number of private cars in use from Mitchell's *International Historical Statistics* (Palgrave Macmillan Ltd [2013b](#)) for an extended panel of 88 countries over the 1950-2010 period (henceforth referred to as 'extended panel'), I again split country-year pairs into deciles based on GDP per capita and calculated

1. DRIVING OVER THE HILL?

the average car intensity for each decile.⁵ Figure 1.3a plots this calculated data, revealing a similar hump-shaped pattern with car intensity highest in the eighth decile of GDP per capita. Drawing on a cross section of this 88-country sample, Figure 1.3b plots car intensity against real GDP per capita in 2010, confirming that a hump-shaped pattern is also reflected in a cross section. These patterns in car intensity across 88 countries for the years 1950-2010 also echo the pattern in PKM per GDP (or PKM intensity) across 29 countries for 1970-2019 shown in Figure 1.1. This is consistent with the finding in Figure 1.2a that the hump-shaped pattern in VKM intensity flows mostly from its extensive margin, car intensity.



(a) 1950-2010. Points show decile averages.

(b) 2010. Dashed line shows quadratic regression.

Figure 1.3: Private car intensity and real GDP per capita, extended panel. Axes transformed to \log_e scales. Sources: Author's analysis; Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.

To investigate these car intensity patterns further, I ran the following regression specifications using my extended 88-country panel for the years 1950-2010:

$$intensity_{i,t} = \alpha + \beta_1 real_{i,t} + \beta_2 real_{i,t}^2 + \varepsilon_{i,t} \quad (1.4)$$

⁵The additional countries included in this extended panel, and details of data sources, are outlined in Appendix A.1

$$intensity_{i,t} = \alpha + \beta_1 real_{i,t} + \beta_2 real_{i,t}^2 + \sum_{i=1}^{M-1} G_i + \varepsilon_{i,t} \quad (1.5)$$

The outcome variable in both specifications, $intensity_{i,t}$, is private car intensity in country i and year t . Both models include real GDP per capita, $real_{i,t}$, in country i and year t and its square as independent variables, while the specification in Equation 1.5 additionally includes country fixed effects. These fixed effects control for time-constant, country-specific factors that may be influencing car intensity. I also clustered standard errors at the country level in all three specifications, allowing for possible correlations between a country's error terms. If there is a hump-shaped relationship between car intensity and GDP per capita, that is independent of fixed country effects in the case of Equation 1.5, the β_1 coefficient would be positive and the β_2 coefficient negative.

Table 1.1: Private motor vehicles per GDP and real GDP per capita in extended panel 1950-2010

	(1)	(2)	(3)
Real GDP per capita	0.478*** (0.091)	0.398*** (0.081)	0.346*** (0.100)
Squared real GDP per capita	-0.007*** (0.001)	-0.005*** (0.001)	-0.005*** (0.001)
Observations	2998	2998	88
Adjusted R^2	0.151	0.165	0.106
Fixed effects	No	Yes	No

Standard errors in parentheses

Standard errors clustered at country level

Within R-squared reported for fixed effects regression

Real GDP per capita in thousand 2005 PPP USD

Sources: Author's analysis; Feenstra, Inklaar, and Timmer 2015;

Palgrave Macmillan Ltd 2013b

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Results for these regressions are presented in Table 1.1. Columns 1 and 2 show regression coefficients for Equations 1.4 and 1.5 respectively when employing all available country-year pairs across the extended panel of 88 countries for the 1950-2010 period, while column 3 shows coefficients for Equation 1.4 run on a 2010

cross section. In all three models, including the fixed effects model (column 2), car intensity is indeed positively associated with GDP per capita and negatively associated with the square of GDP per capita. This signifies a hump-shaped relationship between car intensity and real GDP per capita that is independent of time-constant, country-specific factors.

These facts indicate that as GDP per capita rises, the level of private car intensity in the economy initially increases before peaking and subsequently decreasing. This implies that as countries reach high levels of economic development, they become less car-dependent, with fewer cars required for each unit of GDP produced. We could also infer from these patterns among developed countries that countries at earlier stages of development may be expected to become more car dependent in the short term, but to experience a decrease in car intensity in the longer run as they continue to develop.

1.2.3. Carbon emissions intensity

I opened this paper by highlighting the key role of the transport sector in carbon dioxide emissions. Therefore, it is worth considering what these patterns in private car intensity imply for the intensity of carbon dioxide emissions, or emissions per unit of GDP. Linking estimates of total carbon emissions deriving from all fuel combustion from the International Energy Agency (IEA [2022c](#)) with data on real GDP per capita (Feenstra, Inklaar, and Timmer [2015](#)) for my extended panel between 1971 and 2019, I again separated country-year pairs into deciles based on GDP per capita and calculated the average emissions intensity for each decile. Figure 1.4 plots this data, essentially revealing the latter stages of a hump shape. This is consistent with the hump shape shown in car intensity (Figure 1.3a), and a peak level of emissions intensity occurring at a lower level of GDP per capita than the peak level of car intensity reflects the separate dynamic of increasing fuel

efficiency in motor vehicles over time.




Figure 1.4: Carbon emissions intensity and real GDP per capita, extended panel 1971-2019. Axes transformed to \log_e scales. Points show decile averages. Sources: Author's analysis; IEA 2022c; Feenstra, Inklaar, and Timmer 2015.

1.2.4. Explaining these facts

Could these patterns in private car intensity be explained by developed economies having reached a saturation level of car ownership (Tanner 1978)? That some level of saturation in car ownership exists seems plausible. However, exactly what this level is remains unclear. Figure 1.5 plots the peak number of private cars per capita among countries that had reached such a peak before 2010 and reveals a wide range in these peak levels, for example 0.25 cars per person in Taiwan compared with 0.79 cars per person in the US. In addition, decreases in car ownership among younger cohorts (Melia, Chatterjee, and Stokes 2018; Delbosc and Currie 2013; Kuhnimhof et al. 2012) occurring alongside increases among older adults (Kuhnimhof, Zumkeller, and Chlond 2013) point to a more complicated picture than there simply being an underlying level of saturation. The wide range between countries depicted in Figure 1.5 suggests that while saturated

1. DRIVING OVER THE HILL?

car ownership may be a factor in hump-shaped car intensity patterns, there were also other elements involved.




Figure 1.5: Peak level of private cars per capita and real GDP per capita, extended panel 1950-2010. X-axis transformed to \log_e scale. Sample limited to countries with GDP per capita $>20,000$ PPP USD and where annual growth levels of vehicles per capita were below 2 per cent in each of the final 3 years of the sample period. Source: Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.

Alternatively, could these car intensity patterns be simply explained by shifts to public transport (Newman, Kenworthy, and Glazebrook 2013)? Descriptive statistics of passenger transport from the UK Department for Transport (2013) suggest that this is not the case. Figure 1.6a plots the level of PKM intensity by passenger transport mode against GDP per capita in Great Britain from 1952 to 2019, with each mode normalised to 1 in 1952. This shows that while there appears to be a minor recovery in rail's PKM intensity in recent years, which would be consistent with findings by Newman, Kenworthy, and Glazebrook (2013), the PKM intensity of other modes have continued to decrease and that the PKM intensity of each mode other than the car remains substantially lower in 2019 than in 1952.

It is also worth pointing out that since the period of mass motorisation, the share

of total PKM accounted for by cars has been overwhelmingly large, peaking at 86.9 per cent in 1994 and only reducing to 84.5 per cent by 2019 (UK Department for Transport 2013). As a result of this, plotting total PKM intensity of all passenger transport modes against GDP per capita also produces a hump-shaped pattern in intensity, as depicted in Figure 1.6b. Based on this evidence, a shift to public transport alone cannot explain these car intensity patterns.

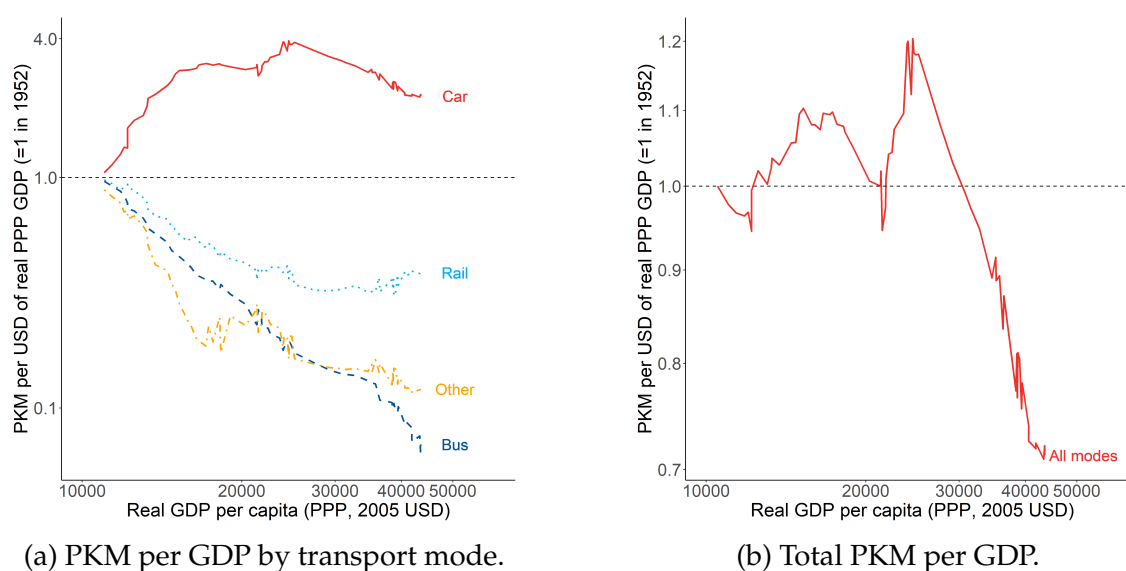


Figure 1.6: PKM per GDP across all transport modes and real GDP per capita, Great Britain 1952-2019. Axes transformed to log_e scales. Source: Author's analysis; Feenstra, Inklaar, and Timmer 2015; UK Department for Transport 2013.

What about road fuel prices (Bastian, Börjesson, and Eliasson 2016; Bastian and Börjesson 2015; Grimal, Collet, and Madre 2013)? Figure 1.7 displays typical retail prices for petrol and diesel along with private car intensity in the UK for the 1955-2010 period. Road fuel prices, which constitute a significant portion of the running costs of motor vehicles, were substantially higher in 2010 than they were at the height of mass motorisation in 1955 in the UK, and this must be regarded as a factor in peak car trends as highlighted by Bastian, Börjesson, and Eliasson (2016), Bastian and Börjesson (2015) and Grimal, Collet, and Madre (2013). However, Figure 1.7 is also consistent with the argument of Puentes and Tomer (2008), Metz

1. DRIVING OVER THE HILL?

(2010), Newman and Kenworthy (2011) and Millard-Ball and Schipper (2011), in that car intensity appeared to be levelling off in the UK prior to the major spikes in fuel prices associated with the 1973 oil crisis. This suggests that while higher fuel prices must be considered to have contributed to the hump shape in car intensity, they are not the whole story either.




Figure 1.7: Private cars per GDP, road fuel prices and real GDP per capita, Great Britain 1955-2010. Y-axis transformed to \log_e scale. Road fuel prices are typical January retail prices in current pence per litre, indexed relative to 1955. Petrol price is for lead replacement petrol 1955-1988, and premium unleaded petrol 1989-2010. Source: Author's analysis; UK Department for Energy Security and Net Zero 2023b; Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.

1.2.5. Structural transformation

How else could we explain this hump-shaped pattern in car intensity? In this paper, I argue that as a country develops, a systematic change in car intensity occurs because of structural transformation, the reallocation of economic activity between broad sectors. Consider a fictional economy and its experience of economic development. When the economy is in its infancy in terms of development, it is dominated by an agriculture sector that is characterised by traditional and labour-intensive processes. At this stage, agriculture is highly localised with households

and communities growing and consuming their own food, and the production of these agriculture goods provides the bulk of employment. These early beginnings are therefore marked by a low level of car intensity.

As the economy grows, technological improvements such as better drainage and systems of crop rotation increase the productivity of labour in the agriculture sector. Moreover, individuals and communities begin to specialise and trade with each other, leading to a relative increase in car intensity in the agriculture sector. As this is happening, the economy is also experiencing a seismic change as it gradually re-orientates itself towards a non-agriculture sector. Initially, this change can be described as industrialisation as the non-agriculture sector may initially be largely comprised of more car-intensive industry and manufacturing sub-sectors. These sectors would require workers to commute to factory locations rather than work locally as was the case in the early, traditional agriculture sector. This industrialisation can also be associated with a process of urbanisation. For example, Gollin, Jedwab, and Vollrath (2015) defined urbanisation in their model economy as an increasing labour share in non-agriculture sectors. The increasing car intensity in agriculture, coupled with the broader shift of economic activity towards the more car-intensive non-agriculture sector, gives rise to a substantial increase in aggregate car intensity in the economy.

Later, as the economy continues to mature and industrialisation is completed, the non-agriculture sector takes the place of agriculture as the dominant sector in the economy. As this occurs, however, non-agriculture itself may shift towards services, and this could facilitate a reduction in its sectoral car intensity if the production of services employs cars less intensively than industry and manufacturing. In addition, industry and manufacturing sectors within non-agriculture may themselves become more efficient in the use of cars as the economy develops, the population becomes more urbanised and technology improves, and

this could further contribute to a decrease in non-agriculture's car intensity. Due to the now-pivotal role of non-agriculture within the economy, this decrease in its sectoral intensity would be expected to propel aggregate car intensity in the economy downwards, following the sharp increase in intensity during the earlier stages of development. A similar theory to this based on dynamic sectoral oil intensities was proposed by Stefanski (2014) in exploring the effect of structural transformation on oil prices. This argument can also be considered distantly related to the growth source explanation for the EKC (Copeland and Taylor 2004; Grossman and Krueger 1995).

To explore this theory in relation to car intensity, I collated a panel dataset of 54 countries over the 1995-2016 period (henceforth referred to as 'main panel') to assess changes in sectoral motor vehicle intensity over the structural transformation process. Based on data availability, this main panel included 34 OECD members, 5 countries designated by the OECD as 'key partners', and 15 non-members. This data was sourced from OECD Input-Output Tables (OECD 2021), the United Nations (UN) national accounts main aggregates database (UN 2021a), International Labour Organization (ILO) modelled employment estimates (ILO 2020), and the Penn World Tables (Feenstra, Inklaar, and Timmer 2015). Further details of these data sources and how I collated them are outlined in Appendix A.1.

I used data from the OECD Input-Output Tables (OECD 2021) to measure the value of sectoral motor vehicle inputs to production. These tables include a category of inputs labelled 'manufacture of motor vehicle, trailers and semi-trailers' (division 29), which I employed when calculating the share of sectoral value added attributable to car transport. More disaggregated data would naturally be more useful in the context of this study, as inputs from division 29 incorporate the manufacture of commercial motor vehicles and trailers in addition to private

cars. Separately, however, I decomposed total motor vehicle intensity in Great Britain over the 1950-2010 period into private car intensity and commercial vehicle intensity, and found that the hump-shaped pattern can also be found when plotting the combined intensity of all motor vehicles against GDP per capita (see Appendix A.1). I also confirmed that the hump-shaped pattern in intensity could be found using aggregate motor vehicle inputs per GDP, calculated using Input-Output data (OECD 2021) data (see Appendix A.1).

To measure structural transformation, I relied on the share of total employment accounted for by the agriculture sector using ILO (2020) data. Figure 1.8a illustrates the well-documented process of structural transformation (for example, see Herrendorf, Rogerson, and Valentinyi 2014; Rogerson 2008; Kuznets 1973), where among the 54 countries included in my main panel, the share of total employment in the agriculture sector was lower in countries with a higher level of GDP per capita. This shows that the agriculture labour share can be considered a measure of a country's structural transformation.

Using my collated main panel, I ran the following linear regression separately for the agriculture sector and for the 'non-agriculture' sector (an aggregation of all other sectors other than the transport sector):

$$VehInputs_{i,t} = \alpha + \beta AgriLabour_{i,t} + \varepsilon_{i,t} \quad (1.6)$$

In Equation 1.6, the dependent variable $VehInputs_{i,t}$ measures the share of sectoral value added (with value added measured in constant 2005 PPP USD) in country i in year t that is accounted for by motor vehicle inputs. The independent variable $AgriLabour_{i,t}$ represents the share of total employment in agriculture in country i and year t . This regression essentially relates sectoral motor vehicle intensity to the

1. DRIVING OVER THE HILL?

extent of structural transformation in each country over time, with the coefficient of interest β capturing the linear relationship between sectoral vehicle intensity and structural transformation.

Table 1.2 presents results of this regression for agriculture (see column 1) and non-agriculture (see column 2). Since the agriculture labour share decreases during the process of structural transformation, Figure 1.8b depicts regression lines for each of these specifications with the x-axis reversed and extended over the full domain of the agriculture labour share to illustrate progress in structural transformation.

Table 1.2: Changes in sectoral motor vehicle inputs over process of structural transformation for main panel 1995-2016

	(1) Agriculture	(2) Non-agriculture
Agriculture labour share	-0.009*** (0.001)	0.004*** (0.001)
Observations	1188	1188
Adjusted R^2	0.031	0.009

Standard errors in parentheses

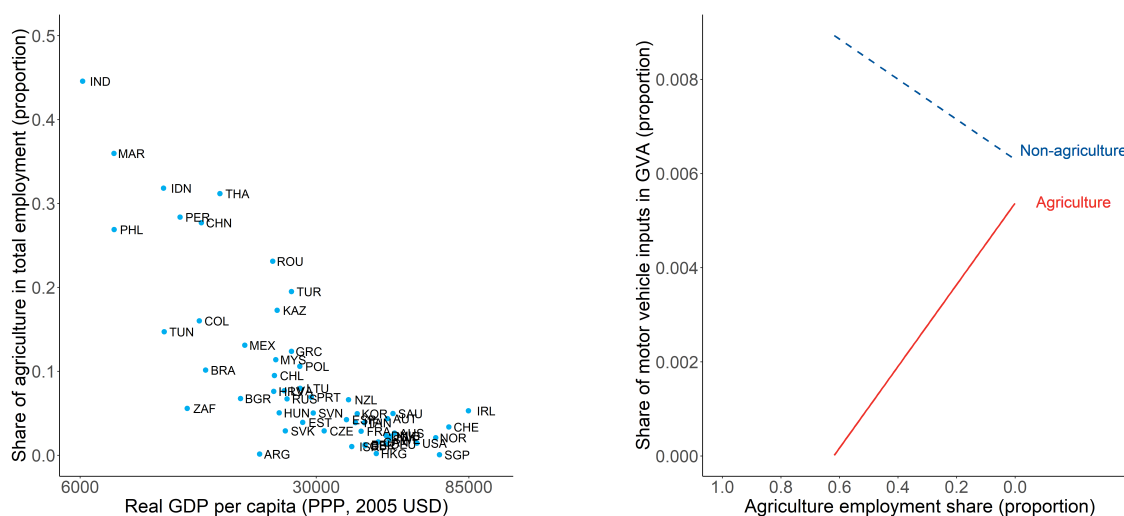
Sources: Author's analysis; OECD [2021](#); ILO [2020](#);

Feenstra, Inklaar, and Timmer [2015](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.2 reveals a negative coefficient on the agriculture labour share variable in the agriculture sector regression, indicating that a higher share of total employment in agriculture was associated with a lower share of transport inputs in the value added of the agriculture sector. In other words, as structural transformation progressed and the agriculture labour share reduced, more motor vehicle transport was required in agriculture to produce one unit of output, as summarised in Figure 1.8b. Intuitively, this can be explained by traditional, localised agriculture processes making way for modern, more car-intensive agriculture processes.

Conversely, a positive coefficient was found on the agriculture labour share variable in the non-agriculture regression. As illustrated in Figure 1.8b, this



(a) Agriculture employment shares and real GDP per capita, 2016. X-axis transformed to \log_e scale.

(b) Regression lines, 1995-2016. X-axis extended over full domain and reversed for illustration.

Figure 1.8: Changes in sectoral motor vehicle intensity over structural transformation, main panel. Sources: Author's analysis; OECD 2021; ILO 2020; Feenstra, Inklaar, and Timmer 2015.

suggests that over the course of structural transformation, less motor vehicle transport was needed to produce one unit of non-agriculture output. This also makes intuitive sense, as this sector may have started life as a transport-intensive sector based on industry and manufacturing, but later morphed into a relatively less intensive sector characterised by an increasingly efficient services sector.⁶

These processes could be fundamental to producing a hump-shaped pattern in aggregate car intensity as a country develops. The aggregate car intensity of the economy is essentially a weighted average of sectoral car intensities. During the early stages of the structural transformation process, the economy is beginning to shift from agriculture to a considerably more car-intensive non-agriculture sector, thus increasing aggregate intensity. However, this non-agriculture sector is

⁶I also ran these regressions separately for the industry and services sectors, and found a positive coefficient on the agriculture labour share in the services regression, while the corresponding coefficient was not significantly different to 0 in the industry regression (see Appendix A.2).

gradually becoming less car-intensive over time, which would lead to a decrease in aggregate car intensity as the sector comes to dominate the maturing economy.

Clearly, that this process of structural transformation could produce a hump-shaped pattern in aggregate car intensity over the course of economic development is dependent on underlying parameters in the economy. How quickly does the structural transformation into non-agriculture occur? What is the rate of decrease in non-agriculture's sectoral car intensity? To help explore this, I developed and calibrated a theoretical model of structural transformation from agriculture to non-agriculture to show that a two-sector model could produce a hump-shaped pattern in car intensity, but that a one-sector model could not do so due to the absence of any structural transformation. I describe this model in the following section.

1.3. The model

This simplified general equilibrium model of structural transformation represents a mathematical formalisation of my theory of why we observe a hump-shaped pattern in car intensity in the real world. I treat this closed economy model as a sequence of static problems that vary over time only via exogenous changes in the productivity of firms. This approach, which abstracts from intertemporal choice, has previously been adopted in studies of structural transformation such as Stefanski (2014) and Duarte and Restuccia (2010). I will describe this model for a single period and thus omit a time subscript t for simplicity. Additional theory in relation to the model is outlined in Appendix A.3.

1.3.1. Consumers

This model economy consists of representative consumers and firms interacting with each other in a perfectly competitive market. The representative household decides on its consumption levels, c_s , of the two final goods: an agriculture good $s = A$ and a composite non-agriculture good $s = N$. As in Stone-Geary preferences, I assume that there exists a subsistence level of the agriculture good, denoted by \bar{c}_A , such that the household will devote all of its income to consuming this good until \bar{c}_A is reached. Intuitively, the household requires \bar{c}_A of the agriculture good, which can be thought of as food, to survive and will thus spend all of its budget on this good until it has acquired \bar{c}_A .

The household takes the wage rate w and prices p_s in the economy as exogenously given and seeks to maximise its utility from consuming these two goods subject to a budget constraint. I assume this consumer optimisation problem takes the form of a log utility function:

$$\begin{aligned} \max_{c_A, c_N} \quad & \phi \log(c_A - \bar{c}_A) + (1 - \phi) \log(c_N) \\ \text{s.t.} \quad & p_A c_A + p_N c_N = w \end{aligned} \tag{1.7}$$

In this constrained maximisation problem, $\phi \in (0, 1)$ represents the utility weight on the agriculture good. The subsistence level of the agriculture good is given by $\bar{c}_A > 0$.

1.3.2. Firms

Meanwhile, production in the model takes place across three perfectly competitive sectors $s = A, N, M$: sector A producing the final agriculture good, sector N producing the final non-agriculture good, and sector M producing motor vehicles as an intermediate good. Motor vehicles thus appear in my model as an intermediate

1. DRIVING OVER THE HILL?

good that is used as an input in the production of the final goods sectors A and N alongside labour. The price of motor vehicles in the model, p_M , is a variable cost that can be thought of as encompassing costs of operating a single vehicle, such as fuel. As a simplification, the model abstracts away from any fixed costs, and does not incorporate possible alternatives to motor vehicles such as public transport.

The firms take the wage rate w and prices p_s in the economy as exogenously given and select levels of motor vehicle and labour inputs, M_s and L_s , that maximise profits Π_s . I assume the sector-specific firm profit maximisation problems for the agriculture and non-agriculture sectors A and N both take the form of a constant elasticity of substitution function:

$$\max_{M_s, L_s} p_s B_s \left(\eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}} + (1 - \eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s-1}} - p_M M_s - w L_s, \quad s = A, N \quad (1.8)$$

In this maximisation problem, total factor productivity B_s differs between the two sectors. Similar to Duarte and Restuccia (2010), Stefanski (2014) and Gollin, Jedwab, and Vollrath (2015), this parameter captures labour productivity in addition to any effects of capital or land in production. Differences in labour productivity across countries and over time can stem from factors such as endowments and capital intensity. However, my model follows the approach of Duarte and Restuccia (2010) in abstracting from these factors as the focus is not explicitly on the productivity of labour.

The parameter η_s is the sector-specific motor vehicle share parameter, while σ_s represents the sector-specific elasticity of substitution between motor vehicle and labour inputs. This elasticity of substitution measures how easy it is for the firm to switch between motor vehicles and labour in production. Formally, it is defined as the percentage change in the firm's marginal rate of technical substitution due to a 1 per cent change in the ratio of motor vehicles to labour inputs, $\frac{M_s}{L_s}$, holding the

level of output constant. A positive elasticity between two factors of production indicates that the factors have some degree of substitutability, and lower values of elasticity are associated with factors of production that are more difficult to swap for each other. In the extreme case, perfect substitutes have an elasticity of infinite.

A constant elasticity of substitution production function is a special case in which the elasticity is simplified to be constant for any combination of production inputs.⁷ The two-factor constant elasticity of substitution production function was first employed by Solow (1956) before becoming widespread in economic literature. As I will discuss later, these constant sectoral elasticities play a significant role in determining the evolution of motor vehicle intensity within sectors over time.

In this model economy, I assume labour can freely move between sectors and thus equalise wage rates to produce a common economy-wide wage rate w . Similarly, I assume motor vehicles produced by firms in sector M can be used in the production of both final goods at a common cost p_M .

An input demand function for M_s in terms of L_s can then be derived from the first-order conditions of this maximisation problem (see Appendix A.3) and solving for M_s :

$$M_s = \left(\frac{w}{p_M} \right)^{\sigma_s} \left(\frac{\eta_s}{1 - \eta_s} \right)^{\sigma_s} L_s, \quad s = A, N \quad (1.9)$$

Meanwhile, I assume the firm profit maximisation problem for the intermediate motor vehicle sector M takes a linear form:

$$\max_{L_M} p_M B_M L_M - w L_M \quad (1.10)$$

⁷A further simplification of this production function that is common in economic literature is the case where this constant elasticity of substitution is unitary, known as the Cobb-Douglas production function. I do not rely on Cobb-Douglas production in my model so that while the elasticity of substitution is constant within sectors, I can allow it to differ between sectors.

1. DRIVING OVER THE HILL?

In this sector, a representative firm operating in perfect competition takes the economy-wide wage rate w and the price of a motor vehicle p_M as exogenously given and chooses a level of labour inputs L_M in order to maximise profit Π_M . Total factor productivity in this sector is given by B_M .

It should be noted that some model simulations involved extending the model backwards to a period before structural transformation began, where the non-agriculture labour share was 0. For these periods, I assumed that the model economy consisted only of a traditional agriculture sector that used labour as the sole production input. Car intensity in my model economy thus remained at 0 during this period until structural transformation began.

1.3.3. Carbon emissions

The use of motor vehicle inputs pollutes the environment in my model economy. Specifically, the deployment of M_s motor vehicles produces P_s units of carbon emissions as follows:

$$P_s = \rho M_s, \quad s = A, N \quad (1.11)$$

In Equation 1.11, the parameter ρ represents the economy-wide emissions factor of motor vehicles. I assume that this parameter is exogenously changing over time to reflect developments in the fuel efficiency of motor vehicles. These emissions are simply a side effect of production, and do not affect the utility of consumers or the productivity of firms in my model economy. This is a fair characterisation of pollution, as it is typically regarded as a negative externality and individual concerns in relation to carbon emissions and global warming deriving from consumption are a relatively recent phenomenon.

1.3.4. Competitive equilibrium

The competitive equilibrium in this model consists of: final goods prices $\{p_A, p_N\}$, taken as exogenous by consumers and firms; a common wage rate w and motor vehicle p_M , also taken as exogenous by agents; a consumption bundle $\{c_A, c_N\}$, chosen by the representative household to maximise utility subject to its budget constraint, given final goods prices and the wage rate; labour $\{L_A, L_N, L_M\}$ and motor vehicle $\{M_A, M_N\}$ inputs, set by the representative firms to maximise profits given final goods prices, the common wage rate and motor vehicle cost; and consequent levels of final output $\{Y_A, Y_N\}$, based on the chosen inputs and sector-specific productivity.

I specify four market-clearing conditions for the economy to satisfy to reach this competitive equilibrium. First, the labour market clears such that all labour provided by the household is split between the three goods sectors: $L = L_A + L_N + L_M$. Second, the market for motor vehicles clears such that all vehicles produced by the firm in sector M are split between the two final goods sectors as production inputs: $B_M L_M = Y_M = M_A + M_N$. Finally, both final goods markets clear such that output in each sector equals the quantity demanded: $Y_A = c_A L$ and $Y_N = c_N L$.

1.3.5. Characterisation

The first step in characterising this model is to normalise the price of one good to 1. This good acts as the ‘numeraire’ in the model, and all other good prices and the wage rate are then relative to the numeraire. I choose motor vehicles for this task, so I set $p_M = 1$. This means that the first order condition of the motor vehicle firm’s maximisation problem (see Appendix A.3) can be rewritten as $w = B_M$, showing that the wage rate in my model economy is equal to the level of total factor productivity in the motor vehicle sector. Armed with this fact, Equation 1.9

can be rewritten as follows:

$$M_s = B_M^{\sigma_s} \left(\frac{\eta_s}{1 - \eta_s} \right)^{\sigma_s} L_s, \quad s = A, N \quad (1.12)$$

This can then be substituted into the second of the firm first-order conditions (see Appendix A.3), giving an equation which in turn can be simplified and solved for p_s to give the price of each the agriculture and non-agriculture goods in terms of model parameters:

$$p_s = \frac{B_M^{\frac{\sigma_s}{\sigma_s-1}}}{B_s} (\eta_s^{\sigma_s} B_M^{\sigma_s} + (1 - \eta_s)^{\sigma_s} B_M)^{\frac{1}{1-\sigma_s}}, \quad s = A, N \quad (1.13)$$

Equation 1.13 for each $s = A$ and $s = N$ can then be substituted into the respective consumer demand functions (see Appendix A.3) to express the functions in terms of model parameters:

$$c_A = \phi B_A B_M^{\frac{1}{1-\sigma_A}} (\eta_A^{\sigma_A} B_M^{\sigma_A} + (1 - \eta_A)^{\sigma_A} B_M)^{\frac{1}{\sigma_A-1}} + (1 - \phi) \bar{c}_A \quad (1.14)$$

$$c_N = B_N B_M^{\frac{-\sigma_N}{\sigma_N-1}} \left(B_M - \frac{B_M^{\frac{\sigma_A}{\sigma_A-1}}}{B_A} \bar{c}_A (\eta_A^{\sigma_A} B_M^{\sigma_A} + (1 - \eta_A)^{\sigma_A} B_M)^{\frac{1}{1-\sigma_A}} \right) \dots \quad (1.15)$$

$$\dots (\eta_N^{\sigma_N} B_M^{\sigma_N} + (1 - \eta_N)^{\sigma_N} B_M)^{\frac{1}{\sigma_N-1}} (1 - \phi)$$

Finally, to identify the sectoral labour inputs L_A , L_N and L_M analytically, equations 1.12, 1.13, 1.14 and 1.15 can be substituted into the market clearing equations and solved for each L_A , L_N and L_M .

1.3.6. Model discussion

1.3.6.1. Structural transformation

In my model, structural transformation in the form of a shift in labour from the agriculture to the non-agriculture sector is generated by an income effect

stemming from non-homothetic preferences, which is the demand approach of the two broad techniques in the literature identified by Herrendorf, Rogerson, and Valentinyi (2014) and Rogerson (2008). Due to the agriculture good in my model having a Stone-Geary subsistence level, \bar{c}_A , the share of expenditure devoted to the agriculture good can be shown to be decreasing in income w :

$$\frac{p_{ACA}}{w} = \phi + (1 - \phi) \frac{p_{A\bar{C}_A}}{w} \quad (1.16)$$

This implies an income elasticity of less than one for the agriculture good: as income rises, consumers will spend a lower proportion of their income on the agriculture good and increasingly favour the non-agriculture good. The result of this dynamic is a shift in labour from agriculture to non-agriculture to meet this increase in demand as income rises.

1.3.6.2. Urbanisation

Michaels, Rauch, and Redding (2012) presented a model demonstrating the key role of structural transformation in generating urbanisation. Based on this, similar to the approach of Gollin, Jedwab, and Vollrath (2015), I assume that the agriculture good is produced only in rural locations, while the non-agriculture goods is produced only in urban locations, and that goods can be traded freely between locations. This is essentially a reduced-form way of modelling the relationship between structural transformation and urbanisation that was more comprehensively modelled by Michaels, Rauch, and Redding (2012). Based on this, urbanisation in my model corresponds with the increase in labour in the non-agriculture sector.

1.3.6.3. Motor vehicle intensity

I am interested in car intensity in this study, or in my model, motor vehicle intensity. Motor vehicle intensity is given by the total value of vehicle inputs, $p_M M_A + p_M M_N$ divided by the total value of final goods, $p_A Y_A + p_N Y_N$, which can be simplified to the following:

$$\frac{p_M M_A + p_M M_N}{p_A Y_A + p_N Y_N} = \left(\frac{\eta_A}{1 - \eta_A} \right)^{\sigma_A} B_M^{\sigma_A - 1} \frac{L_A}{L} + \left(\frac{\eta_N}{1 - \eta_N} \right)^{\sigma_N} B_M^{\sigma_N - 1} \frac{L_N}{L} \quad (1.17)$$

Equation 1.17 shows that aggregate motor vehicle intensity in the economy is essentially a weighted average of sectoral intensities. As labour shifts from agriculture to non-agriculture during structural transformation, with $\frac{L_N}{L}$ increasing at the expense of $\frac{L_A}{L}$, the weight of the non-agriculture sector in aggregate intensity becomes larger. Equation 1.17 also highlights the central role played by the sectoral elasticities of substitution, σ_s for $s = A, N$, in this intensity. During the structural transformation process in my model, total factor productivity in the motor vehicle sector, B_M , is exogenously increasing. As shown in Equation 1.17, the impact of this dynamic on aggregate intensity is dependent on these elasticities. Different sectoral elasticities mean the aggregate elasticity of substitution across the whole economy will be endogenously shifting during structural transformation, and this is a key mechanism in generating aggregate intensity patterns.

I illustrated in Table 1.2 and Figure 1.8b that motor vehicle intensity has been increasing in agriculture but decreasing in non-agriculture among my main panel of countries. In my model, agriculture and non-agriculture sectoral motor vehicle intensities (in current prices for illustration) can also be expressed as follows:

$$\frac{p_M M_s}{w L_s} = \frac{\eta_s}{1 - \eta_s} \left(\frac{M_s}{L_s} \right)^{\frac{\sigma_s - 1}{\sigma_s}}, \quad s = A, N \quad (1.18)$$

Equation 1.18 further emphasises the fact that the sector-specific elasticities of substitution between vehicle and labour inputs, σ_s for $s = A, N$, are key in determining sectoral motor vehicle intensity over time. For example, if the vehicle-labour ratio $\frac{M_s}{L_s}$ was increasing over time, an elasticity of substitution that is greater than 1 in agriculture and less than 1 in non-agriculture, $\sigma_A > 1$ and $\sigma_N < 1$, would produce the increasing motor vehicle intensity in agriculture and decreasing intensity in non-agriculture as observed. Conversely, setting $\sigma_A < 1$ and $\sigma_N > 1$ would produce these patterns if the vehicle-labour ratio was decreasing over time. Considering the number of registered cars per capita in an economy as an estimate of the aggregate motor vehicle-labour ratio, the substantial increase over time in cars per capita in Great Britain shown in Figure 1.2b suggests that $\frac{M_s}{L_s}$ has most likely been increasing. This indicates that when calibrating the model, the elasticities of substitution should be set at $\sigma_A > 1$ and $\sigma_N < 1$, which I confirm in the following section.

1.4. Model calibration

I calibrated my model to data aggregated across my main country panel between 1995 and 2016. This involved setting model parameters such that the model matches key features of structural transformation and sector-specific use of motor vehicle inputs across these countries during this period. I essentially treated the 54 countries in my main panel as a single large economy, and therefore abstracted away from international trade between countries. As my model economy is closed to international trade for simplicity, calibrating the model to the trends observed in this aggregate data had the benefit of reducing any influence of international trade on these trends. In addition, while individual countries in the panel completed the process of structural transformation from agriculture to non-agriculture much earlier than 1995, when considering the aggregated data,

1. DRIVING OVER THE HILL?

the agriculture labour share decreased steadily from 39.0 per cent in 1995 to 23.8 per cent in 2016. The aggregate data, therefore, provided a useful reference for a structural transformation out of agriculture over a period of 22 years. Details on all data sources and the construction of variables are provided in Appendix A.1.

Using ILO (2020) data, I normalized the size of the labour force to 1 in 1995, and calculated the annualised growth rate of the labour force between 1995 and 2016.

For sectoral total factor productivity, I calculated gross value added (GVA) per worker in constant 2005 USD for the agriculture, non-agriculture (excluding transport) and transport sectors using United Nations (UN) data (UN 2021a). I then normalised each sector to 1 in 1995 and calculated sector-specific annualised growth rates. In the absence of more disaggregated data, I employed GVA per worker in the transport sector as a proxy for total factor productivity in my motor vehicle sector.

Similarly, for the economy wide motor vehicle emissions factor, I calculated the average carbon emissions from motor gasoline combustion (excluding biofuels) on roads per private motor vehicle using IEA (2022c) and Palgrave Macmillan Ltd (2013b) data across my main panel for each year. I then normalised this factor to 1 in 1995, and calculated the annualised growth rate to 2016.

Next, I calibrated sector-specific motor vehicle-labour elasticities of substitution σ_A and σ_N , and sector-specific motor vehicle share parameters η_A and η_N . First, Equation 1.9 can be rearranged as a ratio of motor vehicle to labour inputs for the agriculture and non-agriculture sectors:

$$\frac{M_s}{L_s} = \left(\frac{\eta_s}{1 - \eta_s} \right)^{\sigma_s} \left(\frac{w}{p_M} \right)^{\sigma_s}, \quad s = A, N \quad (1.19)$$

The log of Equation 1.19 is:

$$\log\left(\frac{M_s}{L_s}\right) = \sigma_s \log\left(\frac{\eta_s}{1 - \eta_s}\right) + \sigma_s \log\left(\frac{w}{p_M}\right), \quad s = A, N \quad (1.20)$$

Using sectoral input data from the OECD (2021) and ILO (2020) data on employment, I filled in values for M_A , L_A , w and p_M for 2016 and then estimated Equation 1.20 as a regression using ordinary least squares (OLS). As the term $\frac{\eta_s}{1 - \eta_s}$ is a constant, the slope parameter on the log of the wage-to-vehicle price ratio, estimated by the OLS coefficient on $\log\left(\frac{w}{p_M}\right)$, is the sectoral elasticity of substitution between motor vehicles and labour, σ_s . Results of this regression for each agriculture and non-agriculture are shown in Table 1.3. As shown in Table 1.3, as expected, the estimated elasticity of substitution between motor vehicle and labour inputs was greater than 1 in agriculture but less than 1 in non-agriculture. In other words, it was easier to swap between motor vehicles and labour in agriculture than in non-agriculture.

Table 1.3: Estimation of sectoral motor vehicle-labour elasticities of substitution in main panel 2016

	(1) Agriculture	(2) Non-agriculture
Log wage-vehicle price ratio	1.144*** (0.300)	0.476*** (0.135)
Observations	54	54
Adjusted R^2	0.180	0.161

Standard errors in parentheses

Standard errors clustered at country level

Sources: Author's analysis; OECD 2021; ILO 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Second, it can be shown from the market clearing equation $B_M L_M = M_A + M_N$ that the equilibrium share of labour in the motor vehicle sector in terms of the other

sectoral labour shares is as follows:

$$\frac{L_M}{L} = B_M^{\sigma_A-1} \left(\frac{\eta_A}{1-\eta_A} \right)^{\sigma_A} \frac{L_A}{L} + B_M^{\sigma_N-1} \left(\frac{\eta_N}{1-\eta_N} \right)^{\sigma_N} \frac{L_N}{L} \quad (1.21)$$

Using Equation 1.21, I plugged in sectoral labour shares for each 1995 and 2016 using ILO (2020) data along with the values estimated using OLS in Table 1.3 for σ_A and σ_N , and finally solved for η_A and η_N .

This approach to calibrating the production parameters σ_s and η_s allowed me to pin them down using cross-country variation over 54 countries, rather than variation over time between 1995 and 2016. In the short run, when the technology of capital in the economy can be of a fixed nature, substitution between inputs could be restricted by the form of capital in stock. Rather than explicitly including capital in my model, any effect of capital on production is implicitly captured in my total factor productivity parameters. However, Stefanski (2014) argued that the alternative time series approach to calibrating these production parameters using a relatively short time period may pick up shorter-run responses to price changes before capital has had time to upgrade in the long run. Therefore, calibrating production parameters to cross-sectional data was more appropriate for measuring the long-run ease of substitution between labour and motor vehicles.

This left only two further parameters to be calibrated, the utility weight on the agriculture good ϕ and the subsistence level of the agriculture good \bar{c}_A . Starting from market clearing condition $Y_A = Lc_A$, substituting in agriculture production technology for Y_A and Equation 1.14 for c_A and then re-arranging, it can be shown

that the share of total labour employed in the agriculture sector is given by:

$$\frac{L_A}{L} = \frac{B_M^{\frac{\sigma_A}{\sigma_A-1}}}{B_A} (1 - \eta_A)^{\sigma_A} (\eta_A^{\sigma_A} B_M^{\sigma_A} + (1 - \eta_A)^{\sigma_A} B_M)^{\frac{-\sigma_A}{\sigma_A-1}} \dots \quad (1.22)$$

$$\dots (\phi B_A B_M^{\frac{1}{1-\sigma_A}} (\eta_A^{\sigma_A} B_M^{\sigma_A} + (1 - \eta_A)^{\sigma_A} B_M) + (1 - \phi) \bar{c}_A)^{\frac{1}{\sigma_A-1}}$$

I filled in the agriculture labour share using ILO (2020) data and productivity values using UN (2021a) data for each 1995 and 2016, in addition to the σ_A and η_A parameters calibrated previously, and finally solved these two simultaneous equations for \bar{c}_A and ϕ .

Table 1.4: Calibrated parameter values: main panel aggregates 1995-2016

Parameter	Parameter description	Value	Target
$L_0, B_{s,0}$	Labour force and productivity, 1995	1.000	Normalisation
g_L	Labour force growth	1.004	Labour force growth
g_A	Agriculture TFP growth	5.031	Productivity growth in A
g_N	Non-agriculture TFP growth	1.095	Productivity growth in N
g_M	Motor vehicle TFP growth	1.603	Productivity growth in M
ρ_0	Motor vehicle emissions factor, 1995	1.000	Normalisation
g_ρ	Motor vehicle emissions factor growth	-4.941	Emissions factor growth
σ_A	Agriculture elasticity of substitution	1.144	Vehicle inputs in A , 2016
σ_N	Non-agriculture elasticity of substitution	0.476	Vehicle inputs in N , 2016
η_A	Agriculture vehicle share parameter	0.003	Vehicle inputs in A , 2016
η_N	Non-agriculture vehicle share parameter	0.004	Labour share in M , 1995
\bar{c}_A	Subsistence level of agriculture good	0.272	Labour share in A , 1995
ϕ	Agriculture utility weight	0.159	Labour share in A , 2016

TFP denotes total factor productivity. Annualised growth rates reported as percentages.

Table 1.4 displays the parameters selected by this calibration procedure. This shows that total factor productivity was growing in each sector between 1995 and

2016, but that this growth was almost 5 times higher in agriculture than in non-agriculture. The calibrated annualised growth rate in motor vehicle total factor productivity of 1.6 per cent as shown in Table 1.4 is close to an estimate of 1.09 per cent by Bogart (2014) for annual total factor productivity growth in overland passenger transport during industrialisation in Great Britain (between 1700 and 1870).⁸ This suggests that my approach of drawing on total factor productivity growth in the wider transport sector rather than motor vehicles specifically is reasonable in the absence of more disaggregated data for my model of structural transformation. As expected, Table 1.4 confirms that the motor vehicle emissions factor was decreasing between 1995 and 2016 with improved fuel efficiency, with a calibrated annualised growth rate of -4.9 per cent.

Table 1.4 also shows that, as expected, the elasticity of substitution between motor vehicle and labour inputs was greater than 1 in agriculture but less than 1 in non-agriculture, while the motor vehicle share parameter was slightly higher in non-agriculture than in agriculture. It is also evident from Table 1.4 that the majority of the representative consumer's utility weight was placed on the composite non-agriculture good.

1.5. Baseline simulation

This section presents results from the baseline simulation of my calibrated model and compares these simulations to observed data. First, I simulated the model over the 1995-2016 period to which it was calibrated to determine its performance in reflecting real-world trends. Second, I extended the simulation period back to 1970 and compared this with observed trends as an assessment of the external validity of the model. Third, I utilised the model to simulate motor vehicle intensity

⁸Meanwhile, Bogart (2014) estimated annual total factor productivity growth in overland freight transport to be 2.06 per cent during this period in Great Britain.

among a cross section of my main panel. Finally, I used the model to examine three economies at different stages of the structural transformation process, France, South Korea and China, as case studies.

1.5.1. Calibrated period: 1995-2016

My model was calibrated to the 1995-2016 period, and the first model simulation task involved assessing the ability of the model to reflect various features of structural transformation and car intensity during this period. Figure 1.9 compares simulated data from the calibrated baseline model with data observed from 1995 to 2016. Despite the relative simplicity of the model, Figure 1.9 illustrates that the model was constructed to capture the observed process of structural transformation very well in terms of labour shares and the price of the non-agriculture good relative to that of the motor vehicle good (this good took on the role of the numeraire in the model).

While the model reflected the initial decline in the relative price of the agriculture good during the early stages of the calibration period, it was unable to accurately reflect the later increase (see Figure 1.9g). It may be that this observed increase was driven by factors outside the scope of my model, such as the Great Recession. A dip in labour productivity (measured here as GVA per worker) from 2007 into 2008, during the initial stages of the Great Recession, appears to have been particularly pronounced in the motor vehicle sector (measured here as the entire transport sector) in Figure 1.9c. As shown by the simulated data in Figure 1.9c, however, my model is not designed to capture such short-term shocks, but rather longer-term trends.

Figure 1.10 illustrates that the model implied a decrease in the levels of motor vehicle intensity and carbon emissions intensity over the 1995-2016 period. As

1. DRIVING OVER THE HILL?

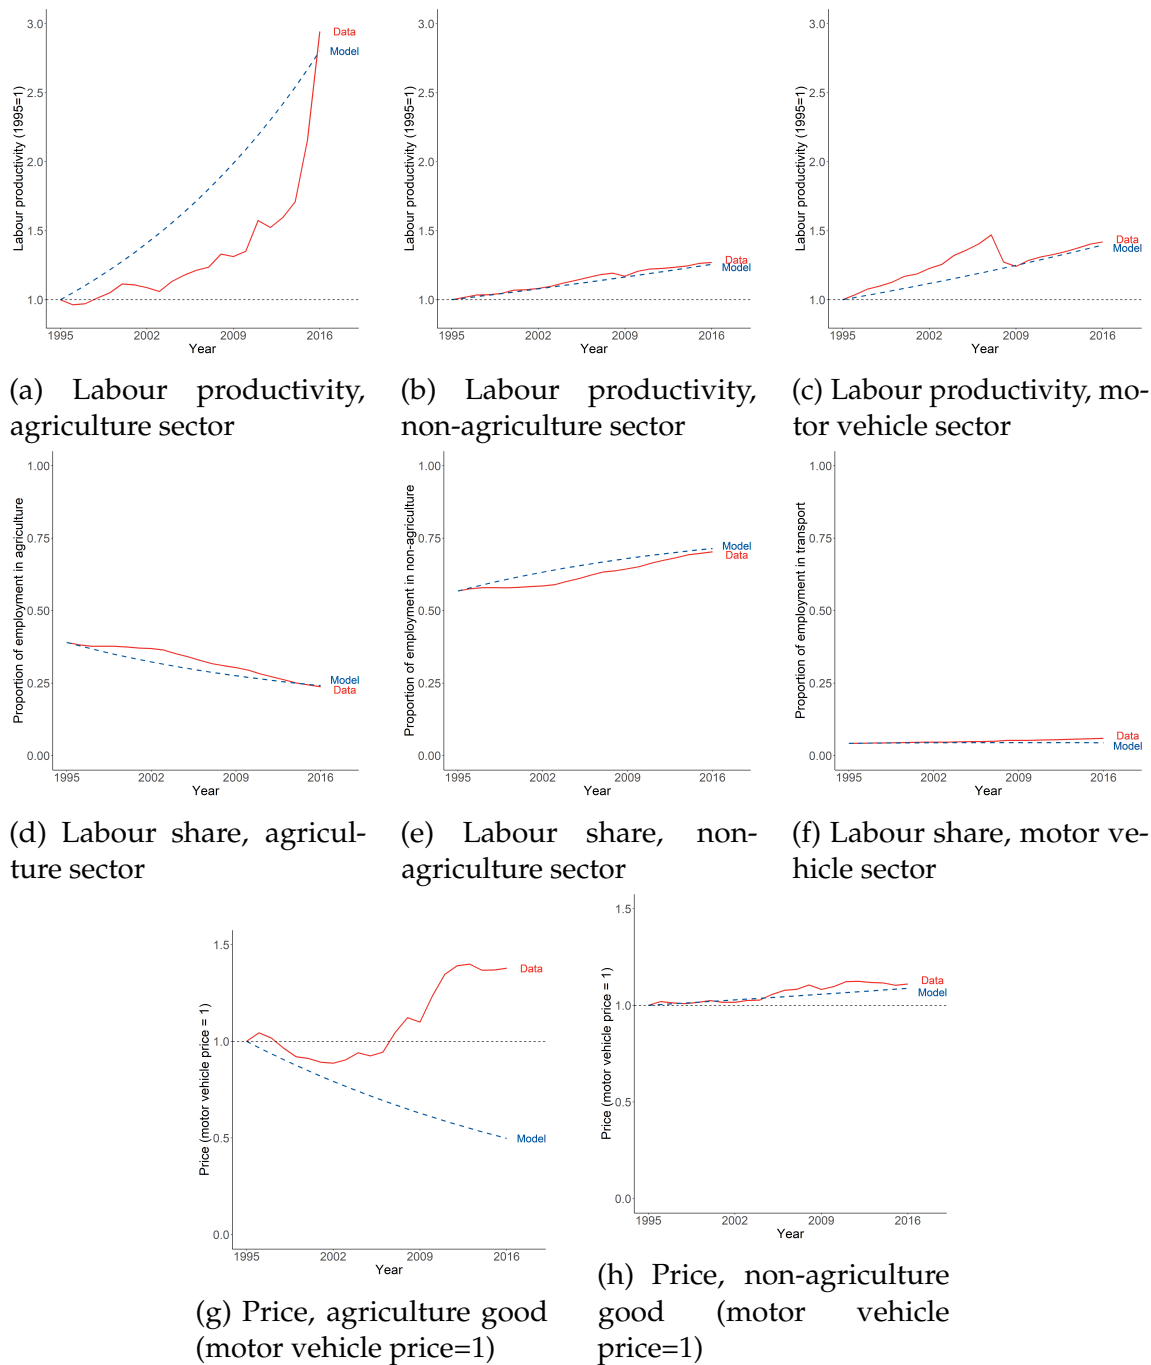



Figure 1.9: Structural transformation, baseline model simulation (blue dashed line) and data (red solid line), 1995-2016 calibration period.

these intensities would begin from 0 if the model was extended backwards to the start of the structural transformation process before increasing, the model is capturing the later decrease in intensity for the 1995-2016 period. The mechanism



(a) Motor vehicle intensity in constant prices

(b) Carbon emissions intensity in constant prices

Figure 1.10: Intensity, baseline model simulation (blue dashed line) and data (red solid line), 1995-2016 calibration period.

for the pattern in motor vehicle intensity is shown in Figure 1.11, which displays the current-price sectoral intensities that correspond to Equation 1.18 implied by the model over the calibration period. Motor vehicle intensity increased slightly in the agriculture sector, but decreased from a much higher base in the non-agriculture sector, producing the decrease in aggregate motor vehicle intensity over the calibration period. Given that the model calibration procedure set the elasticity of substitution parameters as $\sigma_A > 1$ and $\sigma_N < 1$, these sectoral intensity patterns imply that as expected, the vehicle-labour ratio was increasing in both sectors over the calibration period.

1.5.2. External validity: 1970-2016

How well was this model able to reflect structural transformation and intensity trends observed prior to the period to which the model is calibrated? Simulated data from the baseline model is compared with data observed from 1970 to 2016 in Figure 1.12, with the 1995-2016 calibration period shaded in grey. This exercise

1. DRIVING OVER THE HILL?




Figure 1.11: Sectoral motor vehicle intensity in current prices, baseline model simulation for agriculture (red solid line) and non-agriculture (blue dashed line), 1995-2016.

essentially tested the external validity of the calibrated model, and as shown in Figure 1.12, the model was able to capture key features of structural transformation observed prior to 1995. Interestingly, while the model could not capture the increase in the relative price of the agriculture good in the latter stages of the calibration period (see Figure 1.9g), it accurately predicted that this relative price had been mostly decreasing since 1970 (see Figure 1.12b).

In addition, Figure 1.13 indicates that the model correctly implied that motor vehicle intensity had increased sharply prior to the calibration period, during which it peaked and began to decline. This represents clear evidence of the model's ability to produce a hump shape in motor vehicle intensity over the course of economic development. It is also a positive reflection of the external validity of my model that it was successful in capturing the broad trend in car intensity over a 25-year period prior to the period to which the model was calibrated.

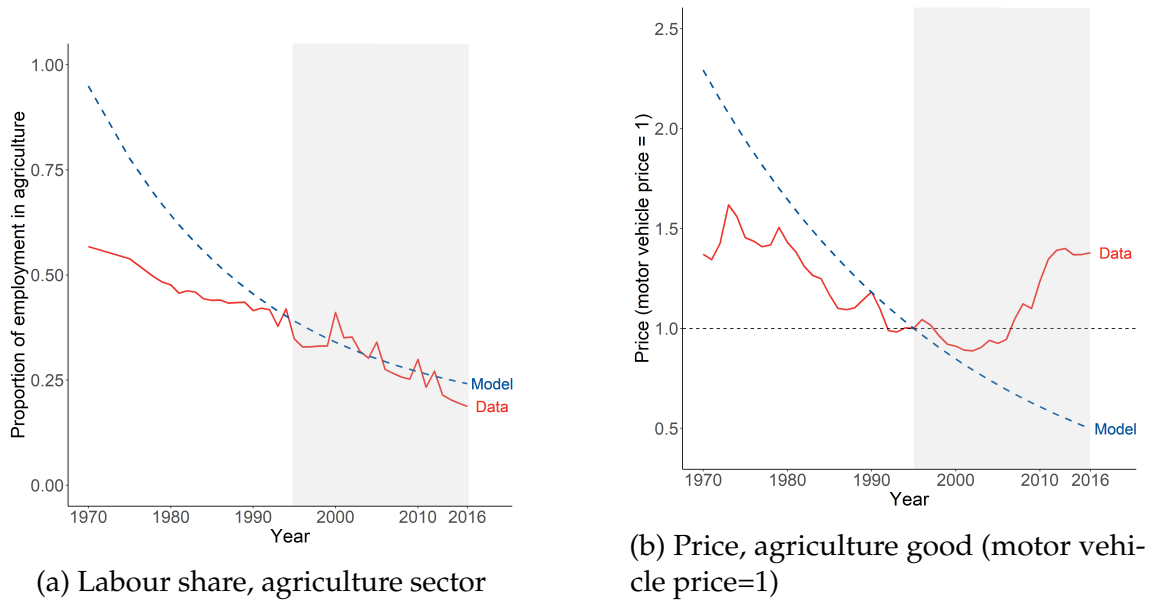


Figure 1.12: Structural transformation, baseline model simulation (blue dashed line) and data (red solid line), 1970-2016. Grey shaded area shows model calibration period.




Figure 1.13: Motor vehicle intensity in constant prices, baseline model simulation (blue dashed line) and data (red solid line), 1970-2016. Grey shaded area shows model calibration period.

1.5.3. Cross-sectional model fit

It is also worth determining whether this model of structural transformation can explain any of the variation in car intensity among a cross section of countries

1. DRIVING OVER THE HILL?

in a single year. I used my model to simulate levels of motor vehicle intensity among my main country panel in the year 2010. I calculated sectoral GVA per worker for each country relative to Great Britain using UN data (UN 2021a), and employed these values as sectoral total factor productivity parameters (B_A , B_N and B_M) for each country, leaving all other model parameters unchanged. Figure 1.14 plots modelled motor vehicle intensity values against observed car intensity values, while Table 1.5 shows results of a linear regression (column 1) and a log-log regression (column 2) of observed values on modelled values.




Figure 1.14: Simulated and observed motor vehicle intensity in constant prices, main panel 2010 (Great Britain = 1). Red dashed line shows 45-degree line.

Table 1.5 shows that the linear model produced a lower value for Akaike's Information Criterion (AIC) than the log-log model, an indication that the linear regression was more appropriate for the data. Table 1.5 reports an adjusted R-squared value of 0.24 for the linear model (and 0.26 for the log-log model), indicating that in this 2010 cross section of 54 countries, the structural transformation model alone was able to account for approximately 24 per cent of observed variation in car intensity.

However, Figure 1.14 indicates that the model generally performs better among higher-intensity countries. A good model fit would be evident through points

Table 1.5: Cross-sectional model fit for main panel 2010

	(1) Observed intensity	(2) Log observed intensity
Modelled intensity	0.651*** (0.154)	
Log modelled intensity		0.324*** (0.072)
Observations	54	54
Adjusted R^2	0.242	0.264
AIC	61	92

Standard errors in parentheses

AIC denotes Akaike's Information Criterion

Sources: Author's analysis; UN [2021a](#); Feenstra, Inklaar, and Timmer [2015](#);

Palgrave Macmillan Ltd [2013b](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

being consistently close to the 45-degree line. In Figure 1.14, while some higher-intensity countries such as Lithuania and Bulgaria deviate from the 45-degree line, more countries appear to deviate from the line at lower levels of intensity. This may partly be due to the aggregated data to which the model was calibrated, as this data is likely to be representative of the economies of some countries more than others. The mixed performance of the model may also stem to some extent from the configuration of the model itself. For example, my model does not include a government sector and thus abstracts away from taxes and subsidies. This would affect its performance in simulating intensity, particularly in economies with higher levels of taxes or subsidies.

1.5.4. Case studies: France, South Korea and China

I also used my baseline model to simulate motor vehicle intensity levels in France, South Korea and China over time. To achieve this, I recalibrated initial levels of productivity in the agriculture sector, $B_{A,0}$, in each of the three countries as a proportion of the productivity aggregated across my main panel in 1995, while leaving all other model parameters unchanged. These observed proportions

1. DRIVING OVER THE HILL?

were 1.4 for France, 0.41 for South Korea and 0.05 for China. For each country, I valued quantities in constant 2005 prices from the France model when calculating simulated intensity levels.

These countries represent three different stories of structural transformation, as shown in Figure 1.15. France and South Korea are both highly developed economies, but South Korea experienced its structural transformation away from agriculture much more recently, with the agriculture labour share still at 49 per cent in 1970 compared with just under 14 per cent in France. China, meanwhile, remains at an earlier stage of the structural transformation process, with just under 38 per cent of total employment still accounted for by the agriculture sector in 2010 (ILO 2020). Based on this, these three economies are interesting case studies to compare against each other in simulating motor vehicle intensity patterns over time.




Figure 1.15: Agriculture employment shares, 1970-2010, France (red solid line), South Korea (blue dashed line) and China (light blue dot-dashed line). Source: Author's analysis; ILO 2020.


Figures 1.16a, 1.16b and 1.16c compare simulated levels of motor vehicle intensity with observed car intensity patterns for France, South Korea and China from 1970 to 2010. These show that the model performed well in capturing the broad

changes in intensity in France and South Korea over the period. In Figure 1.16c, however, the model did not reflect the increase in intensity observed in China, as it implied that structural transformation would not begin until much later than in reality. This further suggests that my model as currently calibrated is more suited to demonstrating the structural transformation process of countries that have already largely completed the move away from agriculture or that are at an advanced stage of this shift, such as France or South Korea, than countries in an earlier phase of development, such as China.


As previously discussed, Figure 1.14 also indicates that the model performed better among countries with intensity levels similar to that of South Korea. A model calibrated to target data from a developing country, rather than data aggregated across a panel of mostly OECD members, may better capture observed patterns for a country such as China. For example, sectoral levels of growth in total factor productivity may differ for developing countries. Alternatively, the model itself could be further developed to include a government sector that incorporates taxes and subsidies. Another possible approach would be to introduce 'wedges' that capture the observed country- and year-specific deviations from the baseline model simulation. These wedges could be interpreted as distortions that arise due to differences between countries in terms of a myriad of factors such as taxation, institutions, climate, natural resources or geography, and would appear in the model in a similar fashion to a tax or productivity shock. Such an approach has been followed by Stefanski (2017), Duarte and Restuccia (2010) and Gollin, Parente, and Rogerson (2002), for example.

Having demonstrated the performance of my model in capturing observed intensity patterns, I then extended the simulation out to 2120 and compared the predicted intensity levels of the three countries. The results of this exercise are displayed in Figure 1.16d.


1. DRIVING OVER THE HILL?




(a) Baseline simulation (blue dashed line) and data (red solid line), France 1970-2010.



(b) Baseline simulation (blue dashed line) and data (red solid line), South Korea 1970-2010.



(c) Baseline simulation (blue dashed line) and data (red solid line), China 1970-2010.



(d) Simulation of motor vehicle intensity, France (red solid line), South Korea (blue dashed line) and China (light blue dot-dashed line), 1970-2120. Horizontal dashed line shows peak intensity for France.

Figure 1.16: Baseline simulations of motor vehicle intensity for France, South Korea and China. Sources: Author's analysis; Palgrave Macmillan Ltd (2013b)

The model correctly implied that China would begin the process of structural transformation at a later stage than in South Korea, which in turn would begin this process later than France, although it clearly overestimated how long it would take

before China began shifting to non-agriculture. As illustrated in Figure 1.16d, in terms of simulated motor vehicle intensity, the model suggested that South Korea would rapidly catch up with France during the early stages of transformation when the economy was shifting from agriculture to the much more intensive non-agriculture sector. Interestingly, however, it also suggested that intensity would peak at a marginally lower level in South Korea than in France. Relative to 2010, the model predicted that car intensity would be 35 per cent lower in France and 15 per cent lower in South Korea by 2040.

Meanwhile, the model implied that China would also eventually experience a high level of car intensity growth when transitioning from agriculture to non-agriculture, but that this too would level off and decrease. It also implied that this levelling off would occur at a lower level of intensity than in South Korea. However, given the failure of the model to reflect earlier observed levels of car intensity in China, the future simulation for China should be treated with caution.

1.6. Role of structural transformation

Having established my model's ability to capture key observed trends in structural transformation and car intensity, I used the model to determine the role played by structural transformation in generating a hump-shaped pattern in intensity. This involved the simulation of two counterfactual models and comparing them to the baseline model. The first counterfactual model imposed that structural transformation did not happen at all, while in the second counterfactual, the process of structural transformation was delayed. For illustrative purposes, I extended these simulations backwards to 1970 and forwards to 2120 under the assumption that model parameters and growth rates in sectoral productivity remain at their calibrated 1995-2016 levels. When calculating simulated intensity

levels in each model, I valued quantities in constant 2005 baseline model prices.

1.6.1. Counterfactual 1: No structural transformation

What is the role of structural transformation in generating this hump-shaped pattern in intensity? To answer this question, a counterfactual must be posed: what would happen in an economy where structural transformation did not occur? In my first counterfactual model, I reduced the multi-sector model to a one-sector model, producing a model economy where structural transformation between sectors was impossible. To achieve this, I essentially merged the agriculture and non-agriculture sectors into a single 'non-agriculture' sector. This involved re-calibrating the remaining model parameters after the removal of a separate agriculture sector, and details of this procedure are provided in Appendix A.3. As a result, in this experiment, no labour was ever allocated to a separate agriculture sector and, therefore, no structural transformation occurred.

The 'One-sector' scenario in Figure 1.17a shows the simulated evolution of motor vehicle intensity in this counterfactual model. Compared with the baseline model, intensity was initially higher in the one-sector model given the non-agriculture nature of the single final sector, and crucially, it decreased at a gradual but relatively uniform rate for the duration of the simulation rather than increasing sharply before levelling off and decreasing. This illustrates the key role of structural transformation in generating a hump-shaped pattern in intensity in my baseline model: a one-sector model was unable to echo the hump-shaped pattern in intensity that a two-sector model could.

The key mechanism for this difference in the two-sector model was the endogenous decline in the aggregate elasticity of substitution between motor vehicles and labour inputs. This aggregate elasticity was essentially an average of sectoral

elasticities in agriculture and non-agriculture, weighted by the respective labour shares. As shown in Figure 1.17b, given a sectoral elasticity of greater than 1 in agriculture and less than 1 in non-agriculture in the two-sector baseline model, the aggregate elasticity was endogenously falling during structural transformation. In the one-sector model, however, without structural transformation this aggregate elasticity was constant (see Figure 1.17b).

This one-sector counterfactual simulation showed that in the absence of a structural transformation from one sector to another, the hump-shaped pattern did not emerge. Therefore, in my model economy, the transition between sectors that were characterised by different production processes was crucial in generating a hump-shaped trend in intensity. This represents further evidence in favour of the theory that structural transformation is a factor in producing a hump shape in car intensity.

It is also worth noting that the different intensity patterns between the baseline model and the 'One-sector' counterfactual had implications for simulated carbon emissions. Total carbon emissions over the 1970-2120 simulation period were 19 per cent higher in the one-sector counterfactual model than in the baseline model.

1.6.2. Counterfactual 2: Later structural transformation

How would aggregate motor vehicle intensity evolve if the structural transformation process occurred at a later stage? This question is particularly relevant for countries that remain dependent on an agriculture sector and that are still increasing in car intensity. In this second counterfactual exercise, I reduced the initial level of total factor productivity in the agriculture sector, $B_{A,0}$, from 1 to 0.6 and left all other model parameters, including the growth rate in agriculture total factor productivity g_A , unchanged. This experiment allowed me to explore

the effect of a delayed structural transformation on motor vehicle intensity in my model economy.


The simulated levels of motor vehicle intensity over time in this second counterfactual model are illustrated by the 'Later' scenario in Figure 1.17a. The model implied that intensity would still evolve in a hump-shaped manner as in the baseline model but with the initial increase starting later, but interestingly, that intensity would peak at a slightly lower level than in the baseline model.

This finding could be explained intuitively by the switch away from agriculture occurring at a later stage in the development of the non-agriculture sector, by which time it has decreased in intensity to a lower level than was the case during the transformation in the baseline simulation. This result suggests that countries that are still experiencing an increase in car intensity will ultimately reach a peak as they structurally transform, but that this plateau may be found at a lower altitude than that experienced by economies that industrialised earlier.

This counterfactual model also had interesting implications for carbon emissions. Relative to the baseline model, over the entire 1970-2120 simulation period, total carbon emissions were 37 per cent lower in the 'Later' counterfactual scenario. This lower level of pollution was due to a combination of intensity peaking at a slightly lower level and the increase in intensity occurring at a later stage when the fuel efficiency of vehicles had improved.

1.7. Empirical analysis

Having proposed a theoretical model and conducted counterfactual exercises to argue that structural transformation has played a role in generating observed hump-shaped patterns in car intensity, I conducted a semi-parametric regression analysis using observed data to empirically test this theory.



(a) Transport intensity in constant prices. Horizontal dashed line shows peak intensity for baseline.

(b) Aggregate elasticity of substitution between motor vehicle and labour inputs.

Figure 1.17: Simulations 1970-2120 using baseline model (red solid line), one-sector model (dark blue dot-dashed line), and later structural transformation model (light blue dashed line).

1.7.1. Method

To empirically test for a relationship between car intensity and structural transformation, I employed locally weighted scatter plot smoothing (LOWESS), which represents a non-parametric approach to examining a relationship between two variables. LOWESS, a modelling method proposed by Cleveland (1979) and further advanced by Cleveland and Devlin (1988), calculates a smoothed value y_i^s for each observed value y_i of a dependent variable by fitting simple regression models to local subsets of the data. Subsets of the data are determined by a nearest neighbour algorithm. Assuming the independent variable x_i is ordered across N total observations, such that $x_i \geq x_{i+1}$ for $i = 1, \dots, N-1$, each smoothed value y_i^s is calculated using observations from $i_- = \max(1, i-k)$ to $i_+ = \min(i+k, N)$, where $k = \left\lfloor \frac{N \times \text{bandwidth} - 0.5}{2} \right\rfloor$. The 'bandwidth' of this algorithm is set by the researcher, with a greater bandwidth imposing a greater degree of smoothing. Within these local subsets, each included observation j is weighted according to

1. DRIVING OVER THE HILL?

a ‘tricube’ weight function with weights calculated as $w_j = \left\{ 1 - \left(\frac{|x_j - x_i|}{\Delta} \right)^3 \right\}^3$, where $\Delta = 1.0001 \times \max(x_{i_+} - x_i, x_i - x_{i_-})$. Each value of y_i^s is thus a weighted least squares regression prediction of y_i at x_i based on a local subset of data.

Crucially, the researcher imposes no global functional form to fit a model to the data. The localised nature of this modelling method, with models being fit only to segments of the data, means that it tends to follow the data rather than assumptions imposed by the researcher. In other words, LOWESS allows the data to do the talking.

Of course, there are other factors that may also affect car intensity. To account for these possible factors, I first ran a fixed effects regression of car intensity:

$$intensity_{i,t} = \alpha + \sum_{t=1971}^{2010} D_t + \sum_{i=1}^{M-1} G_i + X_{i,t} + \varepsilon_{i,t} \quad (1.23)$$

In Equation 1.23, the dependent variable $intensity_{i,t}$ was car intensity in country i and year t . Country and year fixed effects were both included as D_t and G_i respectively, with the number of countries in the panel given by M , to capture any factors that were specific to a particular country or year.

I also included a vector of four control variables, $X_{i,t}$, to capture factors suggested in the literature that may have been both country-specific and time-varying, as these would not have been accounted for by the fixed effects. First, I controlled for enrolment in tertiary education as higher enrolment may lead more younger people to delay car ownership (Delbosc and Currie 2013; Kuhnimhof et al. 2012). Second, I included the old-age dependency ratio to account for higher car ownership among older adults (Kuhnimhof, Zumkeller, and Chlond 2013). Third, I adjusted for road gasoline (petrol) prices to account for their influence on car intensity (Bastian, Börjesson, and Eliasson 2016; Bastian and Börjesson 2015;

Grimal, Collet, and Madre 2013). I followed Bastian, Börjesson, and Eliasson (2016) in including the log of road gasoline prices in the regression model. Fourth, I included railway density as a proxy variable for the extent of public transport infrastructure on the basis that higher public transport availability may reduce car use (Newman, Kenworthy, and Glazebrook 2013).

In addition, I clustered standard errors at the country level. This allowed for potential correlations between the standard errors of observations from the same country, acknowledging possible factors such as country-specific measurement error when quantifying uncertainty.

Having run the fixed effects regression in Equation 1.23 and estimated residuals, I employed LOWESS to test for a relationship between these residuals and a measure of structural transformation that I calculated as 1 minus the agriculture labour share. This is based on the fact that countries with higher levels of GDP per capita have smaller shares of total employment devoted to agriculture, as previously shown in Figure 1.8a. This essentially tested for a relationship between structural transformation and the remaining variation in car intensity that could not be explained by the fixed effects and control variables included in Equation 1.23. While Equation 1.23 imposed parametric assumptions on the relationships between car intensity and the fixed effects and other independent variables, deploying LOWESS laid down no such assumptions on the relationship with structural transformation. In implementing LOWESS, I specifically used running-line least-squares smoothing with the bandwidth set to 0.8, 0.6 and 0.4 to assess different levels of smoothing.

Finally, I ran four further regressions of car intensity based on my LOWESS results to empirically test the hypothesis that structural transformation was associated with car intensity. First, using the functional form suggested by the LOWESS

results, I regressed car intensity on structural transformation without including fixed effects or any other independent variables. Second, I again regressed car intensity on structural transformation, this time including country fixed effects but not including year fixed effects. Third, I added year fixed effects to the car intensity regression. Finally, I then added structural transformation to the full fixed effects regression with controls shown in Equation 1.23, again using the functional form suggested by the LOWESS results. I conducted this empirical analysis using Stata/MP 16.1.

1.7.2. Data

I employed my main panel dataset for this empirical analysis, which included data on car intensity (Palgrave Macmillan Ltd 2013b) and the agriculture labour share (ILO 2020). I added data for my four control variables to this collated panel. The old-age dependency ratio, available from The World Bank (2019), measures the number of persons in the population older than 64 for every 100 persons aged between 15 and 64. Tertiary education enrolment, sourced from The World Bank (2022), is the ratio of total enrolment in tertiary education (regardless of age) to the population in the age group typically corresponding with tertiary education. For road gasoline prices, I used the country-level end-use real price index (2015=100) for unleaded motor gasoline in the IEA Energy Prices and Taxes database (IEA 2022a). I calculated railway density, in kilometres of open railway per square kilometre of land area, by combining data on railway length from Palgrave Macmillan Ltd (2013a) with land area data from The World Bank (2023). Further details on these variables, including definitions, data sources and descriptive statistics, are available in Appendix A.1.

Of course, there was a possibility that some of these independent variables may have been correlated with each other, raising the spectre of multicollinearity. A

strong correlation, indicated by a correlation coefficient of approximately 0.7 or higher, between any of these variables would have been of particular concern. In Appendix A.1, however, I include a table of correlation coefficients showing that the highest correlation was only moderate at 0.56 (between the old-age dependency ratio and tertiary education enrolment).

1.7.3. Results

Column 1 of Table 1.6 displays results for the fixed effects regression described in Equation 1.23, while Figure 1.18 illustrates the country fixed effects estimated for this regression. The directions of the coefficients on the control variables in column 1 were generally as expected based on the literature. A higher old-age dependency ratio was associated with higher car intensity, consistent with the hypothesis that older people are more likely to own a car. While only significant at the 10 per cent level, log fuel prices were negatively associated with car intensity. The coefficients on the tertiary enrolment and railway density variables were not statistically significant. In terms of the country fixed effects, Figure 1.18 indicates that, for example, car intensity tended to be lower in countries such as South Korea and Norway, but higher in countries such as Poland and Bulgaria. This country-specific variation, along with year-specific factors that did not vary across countries and the control variables, was accounted for in estimating Equation 1.23, with the remaining variation in car intensity that could not be explained by these factors comprising the residual.

Using LOWESS, Figure 1.19 indicates a hump-shaped non-linear relationship between structural transformation and this residual estimated from the fixed effects regression in Equation 1.23. By applying LOWESS to this residual and structural transformation, I assessed the relationship between car intensity and structural transformation that was independent of country and year fixed effects,

1. DRIVING OVER THE HILL?




Figure 1.18: Estimated country fixed effects and 95 per cent confidence intervals from regression model of car intensity, main panel 1970-2010. Standard errors clustered at country level. Sources: Author's analysis; Palgrave Macmillan Ltd 2013b.

the dependency ratio, tertiary enrolment, fuel prices and railway density. Figure 1.19 is thus illustrating a hump-shaped relationship between car intensity and structural transformation, accounting for time-invariant country-specific factors, year-specific factors across countries, and the four additional control variables. Reassuringly, this relationship was consistent across varying levels of smoothing, with similar results produced whether the nearest neighbour algorithm's bandwidth was set to 0.8, 0.6 or 0.4. Given the semi-parametric nature of this approach, combining a parametric fixed effects regression with a non-parametric LOWESS method, Figure 1.19 represents clear empirical evidence of a hump-shaped relationship between car intensity and structural transformation.

Given the hump-shaped relationship revealed in Figure 1.19, I then ran quadratic regressions of car intensity on structural transformation, first without including any fixed effects or other independent variables, to test the hypothesis that there was a non-linear relationship between these two variables. Column 2 in Table 1.6 confirms that car intensity was positively associated with structural transfor-




Figure 1.19: LOWESS, car intensity and structural transformation, main panel 1970-2010. LOWESS carried out using running-line least-squares smoothing. Points show country-year observations. Red solid line shows smoother with bandwidth = 0.8. Blue dashed line shows bandwidth = 0.6. Orange dot-dashed line shows bandwidth = 0.4. Sources: Author's analysis; ILO 2020; Palgrave Macmillan Ltd 2013b.

mation, but negatively associated with the square of structural transformation, indicating a hump-shaped relationship. This specification registered an adjusted R-squared of 0.24, indicating that 24 per cent of the variation in car intensity could be explained by structural transformation and its square alone. Interestingly, this echoed the 24 per cent of variation in observed car intensity that my structural transformation model was able to account for in a 2010 cross section (see Table 1.5). Columns 3 and 4 in Table 1.6 indicate that the quadratic association between car intensity and structural transformation held when accounting for country fixed effects and additionally for year fixed effects respectively.

Finally, column 5 in Table 1.6 presents results of the fixed effects regression in Equation 1.23 with structural transformation and its square added as independent variables based on the LOWESS results. The hump-shaped relationship between car intensity and structural transformation also persisted in this specification. Reassuringly, the associations with the dependency ratio and fuel prices remained

1. DRIVING OVER THE HILL?

Table 1.6: Regression model of car intensity across main panel 1970-2010

	(1)	(2)	(3)	(4)	(5)
Structural transformation		92.911*** (31.629)	25.678*** (8.509)	26.406*** (9.287)	32.752*** (9.189)
Squared structural transformation		-53.247** (21.854)	-14.253** (7.014)	-13.800* (7.497)	-18.295** (6.963)
Old-age dependency ratio	37.112** (18.208)				39.762** (17.695)
Tertiary enrolment	-0.496 (1.682)				-0.778 (1.388)
Log gasoline price	-2.199* (1.131)				-2.060* (1.063)
Railway density	10.474 (14.701)				8.639 (13.688)
Observations	1406	1406	1406	1406	1406
Adjusted R^2	0.095	0.239	0.042	0.075	0.150
Country fixed effects	Yes	No	Yes	Yes	Yes
Year fixed effects	Yes	No	No	Yes	Yes

Standard errors in parentheses

Standard errors clustered at country level

Within R-squared reported for fixed effects regressions (columns 1, 3, 4, 5)

Sources: Author's analysis; ILO 2020; The World Bank 2023, 2022, 2019; IEA 2022a;

Palgrave Macmillan Ltd 2013b, 2013a

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

consistent with the results shown in column 1.

As I measured structural transformation in this empirical analysis as 1 minus the proportion of total labour in the agriculture sector, the coefficient on the structural transformation variable of 32.8 in column 5 of Table 1.6 indicates that a 10 per cent increase in structural transformation (measured as a 10 per cent decrease in the agriculture labour share) was associated with an increase in car intensity of 3.28 cars per million USD of GDP. A 10 per cent increase in squared structural transformation was associated with 1.83 fewer cars per million GDP. Similarly, in column 5 the old-age dependency ratio was included in the regression model as a proportion, so a 10 per cent increase in this ratio was associated with 3.98 more cars per million GDP. Meanwhile, an increase in the log fuel price of 1 USD was associated with 2.06 fewer cars per million GDP in column 5, albeit only

statistically significant at the 10 per cent level.

It should be noted that due to possible endogeneity, the results of this empirical analysis cannot be interpreted as casual effects. However, the results do provide further evidence of a hump-shaped relationship between car intensity and structural transformation that was independent of country and year fixed effects and the control variables.

1.8. Discussion

In a global context of global warming and climate change, establishing a clear understanding of factors behind peak car is important. In this study, I focused on the number of registered cars in the economy per unit of GDP, or car intensity, as an intrinsic component of car use that may be underlying the peak car phenomenon.

In a previous study, Kenworthy (2013) found car use per unit of GDP to have reduced among a sample of cities in developed countries. Building on this, I showed that car use per GDP evolves in a hump-shaped pattern over the full course of economic development, and that this dynamic stems mostly from car intensity. In other words, as GDP per capita increases, car intensity initially increases before reaching a peak and later decreasing. Car intensity reached a peak in Great Britain in 1992, for example, and has been gradually decreasing ever since. This phenomenon is relevant to the peak car debate regarding the apparent levelling off in annual per-capita distance travelled in cars and per capita car ownership among developed countries, as it offers insight into intrinsic factors that may be underlying peak car trends.

The central argument in this paper is that structural transformation, the reallocation of economic activity among the broad agriculture, manufacturing and services sectors, has a fundamental role to play in producing this relationship

between car intensity and GDP per capita. A budding economy that is based on a labour-intensive agriculture sector may industrialise by veering towards a relatively more car-intensive non-agriculture sector as it grows, giving rise to an increase in car intensity. As the economy blossoms, however, the now dominant non-agriculture sector itself may become more capable of reducing its car intensity, leading the mature economy to experience a decrease in aggregate car intensity. This is a similar argument to that made by Stefanski (2014) in a study of structural transformation and oil prices.

I proposed a theoretical economic model characterised by a transition from a labour-intensive agriculture sector to a relatively more vehicle-intensive non-agriculture sector that was able to reflect the hump-shaped pattern in intensity. Structural transformation was generated by an income effect in this model, with non-homothetic preferences specified for the agriculture good. The calibrated model, with an endogenously decreasing aggregate elasticity of substitution between motor vehicle and labour inputs in production, was able to account for just under a quarter of observed variation in car intensity among a cross section of countries.

In addition, I conducted a counterfactual exercise using this model and an empirical semi-parametric regression analysis to demonstrate how structural transformation can generate a hump shape in intensity as the economy develops. While my baseline two-sector model characterised by structural transformation could reflect this pattern in intensity, a one-sector model could not account for it.

In a further counterfactual exercise, I found that economies that develop later reach a peak level of car intensity at a later stage, but that this peak occurs at a lower level than that experienced by countries that shifted away from agriculture earlier. This had important implications for carbon emissions, with total emissions over

the entire simulation period 37 per cent lower in my counterfactual simulation than in my baseline model.

1.8.1. Policy implications

The findings of this paper have several implications for public policy. First, car intensity should continue to fall among developed countries. For example, my model predicted that car intensity will be 35 per cent and 15 per cent lower in France and South Korea respectively in 2040 than in 2010. A falling level of car intensity, with fewer cars required to produce economic output, suggests that developed countries have indeed reached peak car.

Second, as developing economies structurally transform, it can be expected that their levels of car intensity will increase. This increase is unlikely to be permanent, however, as intensity in these countries should eventually reach a peak before later decreasing. Due to the delayed increase in intensity, a lower peak level of intensity, and improvements in fuel efficiency of motor vehicles, it can also be expected that cumulative carbon emissions from cars during the process of structural transformation should be lower in these countries than in countries that industrialised earlier.

Third, how soon this peak is reached, and the subsequent rate of decrease, could potentially be influenced by targeting the productivity with which the non-agriculture sector deploys motor vehicle inputs, as a declining intensity in the non-agriculture sector is the channel through which aggregate intensity eventually falls in my model. Precisely how best to achieve this with levers available to public policy, such as taxation, is a potential avenue for future research.

1.8.2. Limitations and strengths

Certain limitations must be acknowledged in relation to this study. First, while the paper is focused on car intensity, the closest data available on sectoral inputs were categorised as ‘manufacture of motor vehicle, trailers and semi-trailers’ in the OECD Input-Output Tables (OECD 2021). Meanwhile, in the absence of more disaggregated data, I used ‘transport’ categories as proxies for the motor vehicle sector in GVA (UN 2021a) and employment (ILO 2020) data when calibrating my model. These categories are clearly broader and cannot be considered direct equivalents to private motor vehicles, and this is a limitation of the study due to data availability. Second, the calibrated model performed poorly in reproducing intensity patterns for developing economies such as China. In particular, the model overestimated the delay before structural transformation occurred. This could indicate that the data the model was calibrated to was more suited to developed economies. Third, while my regression analysis provided empirical evidence of a hump-shaped association between car intensity and structural transformation that was independent of country and year fixed effects and several control variables, these results could not be interpreted as causal due to possible endogeneity.

On the other hand, this paper can boast several key strengths. First, to my knowledge, this is the first study to empirically demonstrate a hump shape in car intensity. Second, despite the relative parsimony of my economic model, it was able to reflect this hump-shaped intensity pattern and performed well in reproducing intensity patterns among developed economies, indicating that structural transformation can explain some of the hump-shaped pattern in intensity. Third, the results of hypothesis tests using semi-parametric methods with observed data added further weight to the argument that structural transformation can play a role in generating this pattern. The paper, therefore, presents both theoretical and empirical arguments in favour of structural transformation influencing car

intensity, and by extension, being an elementary factor that is underlying observed peak car trends. Based on these strengths, this paper represents a valuable contribution to peak car and structural transformation literature.



CHAPTER TWO

ALL AT SEA?

Brexit, shipping, and the UK land-bridge

2.1. Introduction

In June 2016, the United Kingdom of Great Britain and Northern Ireland (UK) voted to leave the European Union (EU). The UK formally withdrew from the EU on 31 January 2020, but high-profile negotiations continued between the UK and the EU until the UK's exit from the EU Single Market and Customs Union on 31 December 2020 (Walker [2021](#)). As the new year of 2021 was rung in around the world, a new trade relationship commenced between the EU and its first former member.

Maritime shipping transports over 80 per cent of the world's trade by volume, and in doing so employs approximately 1.9 million seafarers globally (UN [2021b](#)). This study assesses the impact of Brexit on the shipping of goods between EU ports and the UK. It also specifically examines the effect on the UK 'land-bridge' route for goods being transported between Ireland and the rest of the EU, and considers the implications of this for energy consumption and carbon emissions. Transporting cargo between Ireland and the European continent via Great Britain has traditionally been a popular route for trade, known as the UK land-bridge.

2.1.1. How has Brexit affected trade?

While the UK's newly-appointed Secretary of State for Exiting the EU proclaimed in 2016 that Brexit would allow the UK to become 'a beacon for free trade across the world' (House of Commons 2016), the withdrawal from the Single Market fundamentally represented a departure from a large area of free trade and thus the construction of additional barriers to trade. There is an extensive literature on the effects on international trade of joining a free trade agreement or reducing trade barriers, but Brexit steered the UK and EU into relatively uncharted territory by casting aside such an agreement. Empirical studies of the impact of Brexit on international trade are thus emerging in a developing area of research. The subsequent EU-UK Trade and Cooperation Agreement, provisionally applied from 1 January 2021 before formally coming into force in May 2021 (Walker 2021), avoided the imposition of tariffs or quotas on EU-UK trade. However, this trade did not escape the establishment of non-tariff barriers to trade such as customs inspections or changes in documentation requirements (Flynn, Kren, and Lawless 2021a). This makes Brexit a fascinating case study in leaving a large-scale free trade agreement and consequently increasing non-tariff trade barriers.

Baier and Bergstrand (2007) highlighted an endogeneity problem when studying the impact of free trade agreements using econometrics. Realistically, countries select into such agreements for reasons that may not be observed by the econometrician but that are correlated with trade flows. Therefore, a strategy to isolate and identify effects of free trade agreements using applied microeconomic methods is required. Baier and Bergstrand (2007) recommended the use of panel data approaches rather than cross-sectional instrumental variable or control function methods when tackling this endogeneity concern.

To analyse the effect of Brexit on international trade, several studies have employed

a difference-in-differences methodology using a Poisson pseudo-maximum likelihood (PPML) estimator, which was recommended by Santos Silva and Tenreyro (2006) for data on trade flows. Using this methodology, Kren and Lawless (2022) estimated the effect on goods trade flows between the EU and the UK at an individual product level. To isolate the impact of Brexit from other factors such as the COVID-19 pandemic, Kren and Lawless (2022) included both product-time and product-partner fixed effects to control for other changes in trade patterns. The study found a decline in EU-UK trade as a result of Brexit that was close to 20 per cent in both directions. The PPML difference-in-differences approach was also utilised for analysing the initial impact of Brexit on EU-UK trade by Flynn, Kren, and Lawless (2021a), who found an asymmetric effect between imports and exports. The authors suggested that this may be due to the more gradual introduction of customs checks by the UK compared with the EU (Flynn, Kren, and Lawless 2021a).

When utilising this difference-in-differences methodology, adequately capturing the counterfactual trend is crucial: what would have happened to EU-UK trade had Brexit not occurred? Kren and Lawless (2022) highlighted the importance of the choice of control group for estimating the effect of Brexit on trade by showing different results depending on the counterfactual. Freeman et al. (2022) compared UK-EU trade with the UK's trade with the rest of the world when studying the impact of Brexit on UK trade, and found a sudden and persistent 25 per cent decrease in UK imports from the EU from the start of 2021, but a smaller and only temporary decline in UK exports to the EU. However, Kren and Lawless (2022) argued that global EU trade was more appropriate as a counterfactual than global UK trade, particularly due to possible spill-over effects into global UK trade since almost half of UK trade was with the EU.

When studying the effect of Brexit on trade, the Republic of Ireland represents an

2. ALL AT SEA?

interesting case study as it is a clear outlier among EU member states in terms of the share of its trade accounted for by the UK (Kren and Lawless 2022). Flynn, Kren, and Lawless (2021b) focused on trade between Ireland and the UK and analysed the initial impact of Brexit, also using a difference-in-differences approach with a PPML estimator. Their central research question was how the effect on Ireland-UK trade patterns differed between Great Britain (encompassing England, Scotland and Wales) and Northern Ireland. Interestingly, they found a decrease in trade between Ireland and Great Britain but an increase in trade between Ireland and Northern Ireland. As with Flynn, Kren, and Lawless (2021a), Flynn, Kren, and Lawless (2021b) found an asymmetric impact on trade flows, with a larger effect on imports from Great Britain than on Irish exports to Great Britain (Flynn, Kren, and Lawless 2021b).

Given the high-profile and protracted nature of Brexit, it is possible that trade flows were affected before the UK formally left the EU Single Market and Customs Union at the end of 2020. Betting markets have been highlighted as evidence that the result of the Brexit referendum in June 2016 came as a genuine surprise (Douch and Edwards 2021; Graziano, Handley, and Limão 2020; Davies and Studnicka 2018), but from that point onwards trade patterns were subject to possible uncertainty and anticipatory effects. Using stock return data, Davies and Studnicka (2018) showed that most firms experienced negative returns immediately after the referendum. Nearly USD3 trillion was wiped off global markets, and firms with global value chains based in the UK or EU were most affected (Davies and Studnicka 2018). Despite the claims of the UK Secretary of State for Exiting the EU in 2016 that Brexit represented ‘a positive outcome for those who wish to trade in goods and services’ (House of Commons 2016), this finding clearly indicated that financial markets anticipated that the UK withdrawing from the Single Market would negatively affect international trade.

While Freeman et al. (2022) found no evidence of uncertainty or anticipatory effects during the interim period between the June 2016 vote and the 31 December 2020 exit from the Single Market, some studies did find such effects. Kren and Lawless (2022) found separate negative effects on trade during the interim period and after the Single Market exit. Douch and Edwards (2021) employed a synthetic control methodology with data on trade between the UK and 14 EU member states and 14 other non-EU trade partners up until March 2018, almost three years before the UK's exit from the Single Market. Despite there being no new trade barriers in place during this period, the study found a 20 per cent decrease in UK exports to the EU and a 15 per cent decrease in exports to the non-EU partners, and argued that this represented an anticipatory effect. The study found that this effect began to emerge before the June 2016 referendum, and suggested that this reflected the May 2015 election of David Cameron as Prime Minister having promised to hold a referendum on EU membership. Graziano, Handley, and Limão (2020) also tested for uncertainty effects associated with Brexit before the referendum, and found that increases in the probability of Brexit occurring, measured using prediction markets, reduced UK exports to the EU, with a larger effect on products where a reversion to World Trade Organization (WTO) tariffs would have been most severe.

A plethora of earlier ex-ante studies estimated the effect various Brexit scenarios may have on future trade flows. For example, Lawless and Morgenroth (2019) estimated the product-level impact of a reversion to WTO tariffs on UK-EU trade and showed that products such as foods and textiles would be most affected. Vandenbussche, Connell, and Simons (2022) utilised a network trade model to analyse the impact of various possible Brexit trade shock scenarios, and argued that there would be substantial direct and indirect negative effects on both value added and employment in the UK and EU. The study also identified that these

effects would be heterogeneous across industries.

2.1.2. How has Brexit affected maritime transport?

Given its integral role in international trade, it is also worth determining how maritime transport has been affected by Brexit. Ke et al. (2022) constructed a simulation model to forecast the potential impact of Brexit on port congestion and supply chain costs in the largest port in Europe for roll-on roll-off (Ro-Ro) cargo, the port of Dover. The study showed that based on the contemporary port infrastructure, customs checkpoints would lead to congestion, and that supply chain costs of a case study firm would increase by between 7 and 20 per cent depending on the scenario. Ke et al. (2022) further suggested that this congestion may increase the relative appeal of the nearby lift-on lift-off (Lo-Lo) port in Southampton for lower-value cargo.

Given its geographical location as an island to the west of Great Britain, Ireland represents a particularly interesting case study when studying Brexit's effect on maritime shipping. The UK accounts for a major share of Irish port traffic, with over a third of goods received in Irish ports arriving from, and over half of goods forwarded bound for, UK ports in 2016 (Lawless and Morgenroth 2017). In addition to direct trade between Ireland and the UK, moving goods between Ireland and the European continent via Great Britain, known as the UK 'land-bridge', has been a very popular route for trade. While the UK land-bridge has typically costed more than direct routes to the European continent, it has typically offered Irish exporters to Europe better transit times, greater reliability and a higher frequency of shipping connections (Vega, Feo-Valero, and Espino-Espino 2018; Breen et al. 2018). Transit times between Ireland and mainland Europe were previously around 20 hours on the land-bridge, but could be up to 40 hours on direct routes (Breen et al. 2018). The land-bridge route has predominantly been a

route for Ro-Ro cargo, and time-sensitive goods such as agri-food products have been particularly reliant on it (Vega, Feo-Valero, and Espino-Espino 2018; Breen et al. 2018). Lawless and Morgenroth (2017) estimated that 11 per cent of Irish imports from the rest of the world excluding the UK, and 53 per cent of Irish exports to the rest of the world, was transported through Great Britain in 2016.

The effect of Brexit on this trade route is worth exploring. Vega, Feo-Valero, and Espino-Espino (2018) analysed descriptive statistics and conducted a stated choice experiment, through interviews with representatives from 49 firms exporting from Ireland, to predict the potential impact of Brexit on this trade route. This ex-ante study suggested that the probability of using direct routes from Ireland to the European continent that bypass Great Britain would increase by up to 14 per cent due to Brexit, and that this may also involve a shift from Ro-Ro to Lo-Lo cargo.

2.1.3. This study

In this study, I contribute to these areas of literature with an ex-post empirical analysis of the impact of Brexit on goods shipping between EU ports and partner countries. First, I explored data from all EU main ports and asked the following research question: Was the reduction in trade flows due to Brexit reflected in maritime cargo volumes? Second, I focused on data from Irish main ports and asked: Has Brexit caused a diversion of maritime cargo from the UK land-bridge trade route to direct routes?

In addition to the importance of assessing the impact of Brexit on trade patterns, these research questions are relevant to European transport policy. The UN (2021c) highlighted sustainable transport as a key enabler of a range of its Sustainable Development Goals. The European Commission (2016) outlined targets for low-emission mobility in Europe. One element of reducing transport emissions is a

2. ALL AT SEA?

European Commission policy aim to transition from road freight transport to more efficient modes including ‘short sea shipping’ (European Commission 2011). Short sea shipping encompasses inter- and intra-regional transport of both bulk and container cargo by coastal waters or inland waterways, and in some instances can be an alternative to land-based transport modes such as road or rail (Frost and Brooks 2018). While short sea shipping tends to be slower than road freight transport, it typically costs less, is less energy intensive, can exploit economies of scale by carrying the equivalent cargo of hundreds of trucks on a single vessel, and avoids road congestion. Vega, Feo-Valero, and Espino-Espino (2018) pointed out that moving UK land-bridge goods traffic from UK roads to direct short sea shipping routes offers a clear avenue for moving towards achieving these policy goals.

In Section 2, I formulate the two main hypotheses to be empirically tested, outline the difference-in-differences methodology, and describe the quarterly Eurostat (2022a) data on the volume of goods transported via European ports between 2013 and 2022 employed in this study. Section 3 presents results of the empirical analysis, revealing a 22 per cent decrease in EU-UK Ro-Ro cargo volumes, and a 54 per cent decrease in Ireland-UK Ro-Ro cargo, that could be causally attributed to Brexit. However, I determined that this occurred alongside a 147 per cent increase in Ro-Ro cargo volumes between Ireland and France that was caused by Brexit. Results of various robustness checks are described in Section 4.

These results represented evidence of a Brexit-induced shift from the UK land-bridge route to the direct Ireland-France short sea shipping route. According to rough calculations in Section 5, this could have led to a reduction of approximately 60 per cent in the energy consumption and carbon emissions associated with this freight transport. In Section 6, results from the empirical analysis are combined with a two-factor production function to demonstrate that the magnitude of

the ‘implicit tariff’ imposed by Brexit depended on exporting firms’ elasticity of substitution between the land-bridge and the direct route, ranging from approximately 3 per cent for goods characterised by a high elasticity to around 29 per cent for goods with unitary elasticity. Finally, a discussion of results and conclusions from the study are summarised in Section 7.

2.2. Methods

2.2.1. Theoretical framework

The first research question of this study stemmed from literature on the negative impact of Brexit on trade values and needs little further introduction. Specifically, the first hypothesis tested in this study was to determine whether this effect extended to maritime cargo volumes:

Hypothesis 1: The observed reduction in EU-UK trade flows due to Brexit was reflected in maritime cargo volumes.

The second research question, focusing on the effect on the UK land-bridge, is worth briefly considering in a theoretical context. Consider the problem facing a representative firm exporting a good from Ireland to France with a choice of two alternative transport routes, the road-based land-bridge route, R , and the direct short sea shipping route, S . In a static model, the firm takes the costs of each transport route, v_R and v_S , as exogenously given and allocates goods between the two routes to maximise profit, Π . If we assume that the elasticity of substitution between these two routes is constant, the firm’s constrained maximisation problem can be expressed as a constant elasticity of substitution production function:

$$\max_{R,S} pA \left(\eta R^{\frac{\sigma-1}{\sigma}} + (1-\eta) S^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - v_R R - v_S S \quad (2.1)$$

2. ALL AT SEA?

In Equation 2.1, the price of the firm's good is denoted by p , A measures the firm's total factor productivity, and the constant $\eta \in (0, 1)$ parameter is the share parameter on the land-bridge route. The constant elasticity of substitution, denoted by σ , measures how easy it is for the firm to switch between the two transport routes. Formally, it is defined as the percentage change in the firm's marginal rate of technical substitution due to a 1 per cent change in the ratio of land-bridge to direct goods, $\frac{R}{S}$, holding the volume of output constant. A positive elasticity between the two routes indicates that they have some degree of substitutability, and lower values of elasticity indicate that they are more difficult to swap for each other. In the extreme case, an infinite elasticity would imply that the routes are perfect substitutes.¹

From Equation 2.1, and introducing a time subscript of $t = 0$ to specify the period before Brexit, it is possible to derive the ratio of the volume of goods transported on the road-based route to the volume of goods transported on the short sea shipping route (see Appendix B.1 for further details):

$$\frac{R_0}{S_0} = \left(\frac{v_S}{v_R} \right)^\sigma \left(\frac{\eta}{1 - \eta} \right)^\sigma \quad (2.2)$$

Equation 2.2 shows that the allocation of goods between the two routes is determined by the relative cost of the routes, $\frac{v_S}{v_R}$, the elasticity of the substitution between the two routes, σ , and the land-bridge share parameter, η .

The introduction of non-tariff trade barriers due to Brexit (such as customs checks) can be modelled as an 'implicit tariff', denoted τ , on the land-bridge route in this

¹A constant elasticity of substitution production function is a special case in which the elasticity is simplified to be constant for any combination of production inputs. The two-factor constant elasticity of substitution production function became widespread in economic literature following its use by Solow (1956).

ratio when $t = 1$:

$$\frac{R_1}{S_1} = \left(\frac{v_S}{v_R(1 + \tau)} \right)^\sigma \left(\frac{\eta}{1 - \eta} \right)^\sigma \quad (2.3)$$

The implicit tariff, $\tau > 0$, in Equation 2.3 increases the relative cost of the land-bridge route, and holding all other model parameters constant, would result in the firm adjusting its allocation of goods between the two routes in favour of the direct route. This was the basis of the second hypothesis tested in this study:

Hypothesis 2: Brexit caused a diversion of maritime cargo from the UK land-bridge trade route to direct short sea shipping routes.

2.2.2. Study design

In this study, I sought to identify the effect on maritime cargo volumes that can be causally attributed to Brexit. When the UK voted to leave the EU in June 2016, it essentially set in motion an unprecedented natural experiment in reversing a free trade arrangement and constructing non-tariff barriers to trade. This facilitated an identification strategy using applied microeconomic methods to overcome the endogeneity concerns surrounding the effects of free trade agreements highlighted by Baier and Bergstrand (2007). Similar to several studies in this emerging literature (Kren and Lawless 2022; Freeman et al. 2022; Flynn, Kren, and Lawless 2021b, 2021a), to identify and isolate this Brexit effect from other factors such as the COVID-19 pandemic or exchange rate fluctuations, I employed a difference-in-differences methodology where I compared the average change over time between 2013 and 2022 in EU-UK cargo volumes with the average change over time in EU global cargo volumes. Specifically, I estimated the following regression:

$$weight_{i,j,t} = \exp [\lambda_t + \gamma_j + \delta Brexit_{j,t}] + \varepsilon_{i,j,t} \quad (2.4)$$

In Equation 2.4, the outcome variable $weight_{i,j,t}$ was a continuous variable recording the weight of goods transported to or from port i , to or from partner country j , in period t , and a period was a quarter-year. Equation 2.4 included a full set of quarter-year fixed effects λ_t and partner country fixed effects γ_j . The coefficient of interest was δ , the constant difference-in-differences parameter on the $Brexit_{j,t}$ dummy variable that was equal to 1 if the partner country was the UK and the period was the first quarter of 2021 or later.

Following Kren and Lawless (2022), Flynn, Kren, and Lawless (2021b) and Flynn, Kren, and Lawless (2021a), I estimated the regression model in Equation 2.4 using a pseudo-Poisson maximum likelihood (PPML) estimator. The PPML estimator, which is also known as the Poisson quasi-maximum likelihood estimator (Wooldridge 2008), is equivalent to a Poisson regression that calculates standard errors that are robust to heteroskedasticity and was initially proposed by Davies and Guy (1987) for modelling spatial flow data. Santos Silva and Tenreyro (2006) further recommended PPML for estimating trade gravity equations by demonstrating its more consistent performance in the presence of heteroskedasticity compared with more conventional log-linearised OLS regressions. Moreover, log transformation is problematic in a trade setting where there are likely to be many zero-flows, and this was another reason for the popularisation of PPML in international trade literature (Correia, Guimarães, and Zylkin 2020). In the case of the PPML estimator, Santos Silva and Tenreyro (2006) pointed out that the data does not necessarily need to be Poisson for the Poisson likelihood function to be consistent, an attribute originally noted by Gourieroux, Monfort, and Trognon (1984).

When estimating Equation 2.4, I used data from all main ports among EU-27 member states. This specification essentially compared the average change over time between 2013 and 2022 in cargo volumes between these EU ports and the

UK with the average change in cargo volumes between these ports and all other countries. This followed the recommendation of Kren and Lawless (2022) in relying on EU global trade to establish a counterfactual trend for EU-UK trade, rather than relying on UK global trade. A decrease in EU-UK cargo volumes due to Brexit would be confirmed by a negative δ coefficient.

To determine the effect of Brexit on the UK land-bridge between Ireland and the European continent, I instead estimated the following regression:

$$weight_{i,j,t} = \exp [\lambda_t + \gamma_j + \delta_{UK} Brexit_{j,t} + \delta_{FR} France_{j,t}] + \varepsilon_{i,j,t} \quad (2.5)$$

As with Equation 2.4, the outcome variable $weight_{i,j,t}$ in Equation 2.5 was a continuous variable recording the weight of goods transported to or from port i , to or from partner country j , in period t . Again, a period was a quarter-year, and a full set of quarter-year fixed effects λ_t and partner country fixed effects γ_j were included as independent variables. Equation 2.5 contained two coefficients of interest. First, as with Equation 2.4, δ_{UK} represented the constant difference-in-differences parameter on the $Brexit_{j,t}$ dummy variable. Second, δ_{FR} was the constant difference-in-differences parameter on the $France_{j,t}$ dummy variable that was equal to 1 if the partner country was France and the period was the first quarter of 2021 or later. I also estimated this regression using PPML.

To estimate the regression model in Equation 2.5, I included only data on Irish main ports. Therefore, this specification compared the average change over time between 2013 and 2022 in cargo volumes between Irish ports and the UK, the average change in cargo volumes between Irish ports and France, and the average change in cargo volumes with all other countries excluding the UK and France. This allowed the estimation of separate effects for Ireland-UK volumes and Ireland-France volumes, with Irish global cargo volumes establishing a

2. ALL AT SEA?

counterfactual trend. The shortest direct maritime transport route between Ireland and continental Europe is to France, so I considered this as the main alternative route to the UK land-bridge route. A shift in the shipping of cargo from the UK land-bridge route to the direct route as a result of Brexit would be shown by a negative δ_{UK} and a positive δ_{FR} coefficient, although it should be noted that a negative δ_{UK} coefficient would also reflect other effects of Brexit such as changes in demand for imports in the UK.

In difference-in-differences design of Equations 2.4 and 2.5, the crucial identifying assumption is that the outcome in the absence of treatment can be captured by the additive structure including a partner country component that does not change over time, γ_j , and a time component that does not change across partner countries, λ_t . This is known as the parallel trends assumption. In other words, identification relies on trends in the outcome variable being the same in the treatment and control groups if the treatment did not occur, once the fixed effects are accounted for. Provided this assumption holds, the trend in the control group can be utilised to impose a counterfactual trend on the treatment groups, with a treatment effect identified as a deviation from this counterfactual trend.

A significant factor in cargo volumes during my study period was the COVID-19 pandemic. For my difference-in-differences study design to separate the effect of COVID-19 on cargo volumes from the effect of Brexit, it was necessary to make the assumption that the pandemic affected all cargo volumes during the same quarters, for example the second and third quarters of 2020. If this was the case, the COVID-19 effect would have been accounted for by the quarter fixed effect, λ_t .

Conducting a difference-in-differences analysis using regressions facilitates the calculation of standard errors as a measure of the uncertainty attached to estimates. Bertrand, Duflo, and Mullainathan (2004) outlined various issues with standard

error calculation that can emanate from difference-in-differences settings. In my study design, for example, one potential problem was that there may have been common unobserved factors that affected all cargo volumes associated with a particular partner country j . Standard errors, even if they are robust to heteroskedasticity as facilitated by the PPML estimator, may be biased if this was the case. A common solution recommended by Bertrand, Duflo, and Mullainathan (2004) is to cluster standard errors. In all my regression analysis, therefore, I clustered standard errors at the partner country level, which allowed error terms to be correlated between all cargo volumes with the same partner country. I compared results for Equation 2.5 clustering on various different groups to assess the validity of this approach (see Appendix B.2).

PPML regression coefficients estimating δ in Equations 2.4 and 2.5 would represent the change in the logarithm of the cargo weight due to Brexit, controlling for all included fixed effects. These coefficients are easier to interpret when exponentiated, whereby they can be interpreted as percentage changes in cargo weight. For example, a PPML regression coefficient $\hat{\delta}$ can be converted into a percentage change using the transformation $e^{\hat{\delta}} - 1$. Therefore, I report all results as exponentiated coefficients (which can also be termed incident-rate ratios). An exponentiated coefficient greater than 1 indicates an increase in cargo volumes, while an exponentiated coefficient of less than 1 indicates reduced cargo volumes, as a result of Brexit. I also report the transformed standard errors associated with these exponentiated coefficients.² In addition, McFadden's pseudo-R squared is reported for each PPML regression.³ For all results, untransformed PPML regression coefficients and associated standard errors are also available in

²These were transformed using the delta rule. Specifically, the transformed standard error associated with an exponentiated PPML regression coefficient $e^{\hat{\delta}}$ was estimated as $e^{\hat{\delta}} \times SE(\hat{\delta})$.

³This is calculated as $1 - \frac{ll(model)}{ll(null)}$, where ll denotes log likelihood.

Appendix B.2.

I conducted all regression analysis using Stata/MP 16.1.

2.2.3. Alternative specifications

Correia, Guimarães, and Zylkin (2020) discussed the estimation of PPML models involving high-dimensional covariates, for example high-dimensional fixed effects. Santos Silva and Tenreyro (2010) and Correia, Guimarães, and Zylkin (2019) outlined how maximum likelihood estimates may not always exist in non-linear models due to a problem known as statistical separation, where a predictor is associated with a single outcome value when the predictor is split at a certain threshold value. The maximisation of the likelihood function using separated data can lead to a lack of convergence, ‘infinite’ estimates and problems with calculating standard errors, and Correia, Guimarães, and Zylkin (2019) indicated that this issue can be amplified in settings with high-dimensional covariates.

As a solution to this, to ensure the existence of consistent estimates, Correia, Guimarães, and Zylkin (2019) proposed an algorithm for detecting such separated observations to be dropped from the sample for the purposes of estimation. As a test of the robustness of my main results that were estimated using the ‘poisson’ command in Stata, I instead employed the ‘ppmlhdfe’ command, developed by Correia, Guimarães, and Zylkin (2020) for Stata, which implements a PPML regression estimator with high-dimensional fixed effects that is robust to statistical separation. Kren and Lawless (2022) employed this estimator when studying the effect of Brexit on trade values.

When making causal inference, it is common for studies to assess the assumption of parallel trends by conducting a ‘placebo’ falsification test. This tests for an effect in a setting where it should not occur, similar to the ‘control experiment’ conducted

by Duflo (2001) in a study of economic returns to education. For example, there would be clear reason to doubt the central parallel trends assumption if, using the same control group, I detected a Brexit ‘effect’ on cargo volumes between non-UK partners that should not have been affected by Brexit. To gain confidence in my assumption of parallel trends, I therefore conducted a placebo test by running my main difference-in-differences regression (Equation 2.4) again, removing any EU-UK cargo volumes and instead designating cargo volumes between the EU and EU-27 member Sweden as being ‘treated’ from the first quarter of 2021 onwards. This instead compared the average change over time between 2013 and 2022 in EU-Sweden cargo volumes with the average change in global EU cargo volumes. Sweden accounted for 1.3 per cent of total cargo volumes transported via EU-27 main ports in 2018, while the UK accounted for 6.4 per cent. The trade relationship between Sweden and the rest of the EU-27 did not significantly change during this period, and Sweden is sufficiently distant from the UK in geographical terms that potential spillover effects into EU-Sweden shipping due to Brexit were not of concern. Therefore, no Brexit ‘effect’ should be detected in this test.

In difference-in-differences settings, an important issue that demands attention is the possibility of pre-existing trends. This was discussed in detail by Wolfers (2006). If present, partner-specific pre-existing trends in the outcome variable would undermine the parallel trends assumption. For example, if the EU-UK cargo volumes were generally decreasing before Brexit, adequately separating further decreases from a Brexit effect would be problematic. Angrist and Pischke (2008) suggested additionally controlling for panel-specific trends as a test for this, pointing to a study of labour regulation in India by Besley and Burgess (2004) as an example of this approach in econometric literature. Therefore, I adapted Equation 2.4 as follows to assess the robustness of my main results to the inclusion

of linear trends:

$$weight_{i,j,t} = \exp [\lambda_t + \gamma_{0,j} + \gamma_{1,j}t + \delta Brexit_{j,t}] + \varepsilon_{i,j,t} \quad (2.6)$$

In Equation 2.6, the outcome variable $weight_{i,j,t}$ again measured the weight of goods transported to or from port i , to or from partner country j , in period t , and a full set of quarter-year fixed effects λ_t and partner country fixed effects $\gamma_{0,j}$ were included. In this specification, $\gamma_{1,j}$ captured partner-specific linear trends in the outcome variable. The main coefficient of interest was δ , the constant difference-in-differences parameter on the $Brexit_{j,t}$ dummy variable. In this case, identification of an effect derived from whether Brexit led to deviations from any existing partner-specific trends. Due to the inclusion of linear trends, this represented a more restrictive version of Equation 2.4 and was therefore expected to increase standard errors. I also adapted Equation 2.5 in a similar manner to control for partner-specific trends in cargo volumes handled by Irish ports.

To test whether there were anticipatory effects on cargo volumes between the June 2016 referendum result and the 31 December 2020 exit from the EU Single Market, similar to studies of international trade values (Kren and Lawless 2022; Freeman et al. 2022), I added a dummy variable to Equation 2.4 as follows:

$$weight_{i,j,t} = \exp [\lambda_t + \gamma_j + \rho Postref_{j,t} + \delta Brexit_{j,t}] + \varepsilon_{i,j,t} \quad (2.7)$$

In Equation 2.7, the outcome variable $weight_{i,j,t}$ still measured the weight of goods transported to or from port i , to or from partner country j , in period t , and a full set of quarter-year fixed effects λ_t and partner country fixed effects γ_j were still included. The main coefficient of interest was δ , the constant difference-in-differences parameter on the $Brexit_{j,t}$ dummy variable. This time, a second coefficient of interest was ρ on the added $Postref_{j,t}$ dummy variable, a post-June

2016 referendum indicator that was equal to 1 if the partner country was the UK and the period was after the second quarter of 2016 and before 2021. This was also a more restrictive specification, allowing for separate post-treatment and anticipatory effects, and this could be expected to reduce the precision of results. The ρ coefficient could also be labelled as an estimate of a ‘post-referendum’ effect, but Douch and Edwards (2021) argued that any effects before 2021 should be viewed specifically as anticipatory effects as no trade barriers were actually imposed until after the exit from the Single Market on 31 December 2020. I also adjusted Equation 2.5 in a similar manner to test for anticipatory effects in Ireland-UK and Ireland-France cargo volumes.

When testing the parallel trends assumption in difference-in-differences settings, Angrist and Pischke (2008) also recommended the inclusion of treatment lags and leads in a more generalised model. This approach was followed by Duflo (2001) and Autor (2003) in econometric literature, for example. Instead of estimating a single constant treatment effect δ , this method involves estimating separate δ coefficients for different quarters. These coefficients can then be plotted as a test of causality in the spirit of Granger (1969). Specifically, to further assess the robustness of my main results, I ran the following regression specification:

$$weight_{i,j,t} = \exp \left[\lambda_t + \gamma_j + \sum_{\tau=0}^m \delta_{-\tau} Brexit_{j,t-\tau} + \sum_{\tau=1}^q \delta_{+\tau} Brexit_{j,t+\tau} \right] + \varepsilon_{i,j,t} \quad (2.8)$$

Equation 2.8 included the same outcome variable $weight_{i,j,t}$ and fixed effects λ_t and γ_j . What set this specification apart from Equation 2.4 was that rather than obtaining a single estimate for a constant δ , separate estimates for $\delta_{-\tau}$ were obtained for each period after the UK’s exit from the Single Market on 31 December 2020, and estimates for $\delta_{+\tau}$ were obtained for periods prior to this exit. The set of $\delta_{-\tau}$ parameters are estimates of post-treatment effects, while the set of $\delta_{+\tau}$

2. ALL AT SEA?

parameters represent estimates of anticipatory effects. As with Equations 2.6 and 2.7, this was expected to produce higher standard errors as it was a more data-intensive version of Equation 2.4.

Essentially, this generalised model could offer evidence on whether causes happened before consequences, rather than the other way around. While this form of Granger causality test falls short of definitively proving causality, it can provide additional confidence in the parallel trends assumption (Angrist and Pischke 2008). Given the high profile and protracted nature of Brexit negotiations, there was a distinct possibility of anticipatory effects between the June 2016 referendum and the 31 December 2020 exit from the Single Market. However, if the parallel trends assumption held and a significant post-treatment effect was evident in cargo volumes, there should not be significant anticipatory effects evident for periods prior to Brexit becoming a serious possibility. The cut-off period before which anticipatory effects should not be expected in this setting could be defined as the third quarter of 2016, following the June 2016 referendum, although Douch and Edwards (2021) suggested that anticipatory effects were possible from the middle of 2015 onwards following the May 2015 general election in the UK.

2.2.4. Data

For this difference-in-differences analysis, I drew on quarterly Eurostat (2022a) data on the gross weight of goods transported through European ports by direction, cargo type and partner country. I analysed data on total cargo and on the four main cargo types: liquid bulk, dry bulk, large containers and Ro-Ro containers. In this Eurostat (2022a) data, container cargo was classified by the method used to move it from quay to ship. First, a large container category incorporated containers lifted on and off the ship, known as 'Lo-Lo' cargo. Second, a 'Ro-Ro' category encompassed containers rolled on and off the ship. Ro-Ro is considered to be

a faster, more flexible method of transferring from quay to ship than Lo-Lo (Ke et al. 2022), and can be either accompanied by a road goods vehicle (in other words, a lorry) or unaccompanied. While accompanied Ro-Ro cargo is more autonomous, the handling of unaccompanied Ro-Ro cargo requires ports to be equipped with towing machinery and storage space (Eurostat 2022b). For the purposes of my empirical analysis, however, I combined these categories into a single Ro-Ro cargo category.

The Eurostat (2022a) data covered all main ports from the 22 EU-27 member states that have a coastline.⁴ Kren and Lawless (2022) noted that one of the many effects of Brexit was a break in the measurement of trade flows in EU data. The Eurostat (2022a) shipping data was not immune to this, with data available on all main UK ports only up until the end of 2020. Based on this issue, I did not include UK origin ports in my sample. For all main ports from the other 22 countries, data on quarterly trade flows with all partner countries, including the UK, was available up until the second quarter of 2022, which I set as the end of my sample period to maximise the use of post-Brexit data. I set the third quarter of 2013, the time of the most recent EU border change before Brexit due to the accession of Croatia on 1 July 2013, as the start of my sample period to provide sufficient data to establish trade patterns prior to the Brexit referendum in 2016.

Partner entities were available at the country level in this Eurostat (2022a) data. This meant that trade flows between EU-27 main ports and the UK included trade flows with Northern Ireland. Flynn, Kren, and Lawless (2021b) discussed how the Brexit withdrawal agreement between the EU and UK provided Northern Ireland with differential treatment to Great Britain in terms of customs and immigration arrangements on the border between Northern Ireland and the

⁴These countries are: Belgium, Croatia, Cyprus, Czechia, Denmark, Estonia, Finland, France, Germany, Greece, Ireland, Italy, Latvia, Lithuania, Malta, Netherlands, Poland, Portugal, Romania, Slovenia, Spain and Sweden.

2. ALL AT SEA?

Republic of Ireland, the EU's only land border with the UK. This special status also involved specific arrangements for trade between Northern Ireland and Great Britain, which includes a well-established route for Ro-Ro traffic. Unfortunately, Northern Ireland could not be separated from Great Britain as a partner entity in the Eurostat (2022a) data. However, using pre-2020 Eurostat (2022a) data on UK main ports, it was possible to determine that the 4 main ports in Northern Ireland accounted for only 6 per cent of cargo handled by all UK main ports as of 2018, and that all Ro-Ro cargo handled by Northern Irish main ports in 2018 came from, or was bound for, Great Britain. Therefore, this issue is unlikely to have affected results of my difference-in-differences analysis.

Eurostat defines a main port as a one that handles over one million tonnes of goods or records over 200,000 passengers annually (Eurostat 2022b). Data on some smaller ports, known as minor ports, was also included in the quarterly Eurostat (2022a) data. However, as these ports were not required to report trade flows on a quarterly basis, I excluded these from my sample and focused on 265 main ports.⁵

In the Eurostat (2022a) data, it was possible that not all zero-cargo flows were directly reported by origin ports, resulting in excess missing observations. For example, if no cargo was shipped between Dublin Port and Cyprus in a given quarter, this may appear in the dataset as a missing observation rather than a genuine zero-flow. To deal with this, similar to the approach of Kren and Lawless (2022) with data on international trade values, I replaced missing observations with zeroes for any port-partner-direction-cargo type series that included at least one non-missing observation during my study period.

⁵I identified main ports using the Eurostat definitions based on 2018 cargo and passenger volumes. For example, the 6 main ports included from Ireland were: Drogheda, Dublin, Limerick, Cork, Rosslare Harbour and Waterford (see Appendix B.2 for a map of these ports).

2.2.5. Outcome variable

Studies of international trade typically analyse the value of trade. However, Lawless and Morgenroth (2017) pointed out that when considering the transport implications of international trade, the volume of cargo measured as weight is more relevant than the value of this cargo. The number of vehicles or vessels required to move goods will mostly be determined by the volume, not the value, of cargo. Other studies examining trade from a transport perspective have also focused on cargo volume rather than value, such as Vega, Feo-Valero, and Espino-Espino (2018). Therefore, for my analysis of the effect of Brexit on maritime transport, my outcome variable was the gross weight of cargo handled by a main port, measured in tonnes in the Eurostat (2022a) data.

Table 2.1: Descriptive statistics for total trade cargo weights of EU-27 main ports 2013-2022

	N	Mean	S.D.	Min.	Max.
Total weight (thousand tonnes)	598392	48.77	271.91	0.00	17933.00
Liquid bulk weight (thousand tonnes)	251604	42.46	233.49	0.00	14382.00
Dry bulk weight (thousand tonnes)	324180	20.09	98.45	0.00	7322.00
Large containers weight (thousand tonnes)	303696	23.06	133.53	0.00	6165.00
Ro-Ro weight (thousand tonnes)	134568	24.42	171.30	0.00	7793.00

N denotes observations. S.D. denotes standard deviation

Sources: Author's analysis; Eurostat 2022a

Table 2.2: Descriptive statistics for total trade cargo weights of Irish main ports 2013-2022

	N	Mean	S.D.	Min.	Max.
Total weight (thousand tonnes)	10,116	43.27	244.00	0.00	4350.00
Liquid bulk weight (thousand tonnes)	3,708	25.31	97.20	0.00	997.00
Dry bulk weight (thousand tonnes)	7,020	19.08	61.68	0.00	879.00
Large containers weight (thousand tonnes)	2,592	26.15	102.15	0.00	1083.00
Ro-Ro weight (thousand tonnes)	2,772	48.31	312.54	0.00	3345.00

N denotes observations. S.D. denotes standard deviation

Sources: Author's analysis; Eurostat 2022a

2. ALL AT SEA?

Table 2.1 displays descriptive statistics for the gross weight of total trade (aggregating imports and exports) transported via EU-27 main ports by cargo type during my study period, while Table 2.2 shows the corresponding statistics for the specific subset of Irish main ports. Figure 2.1 illustrates the mix of cargo types in cargo volumes through EU-27 ports and specifically Irish ports over the study period, both with all partner countries and with the UK only. First, this indicates that Ro-Ro cargo accounts for a much higher proportion of total cargo handled by Irish ports than by ports in the EU-27 as a whole, where liquid bulk makes up a substantial proportion of total cargo. Second, when focusing on the cargo mix for maritime trade with the UK, it is clear that the proportion of cargo accounted for by Ro-Ro is much higher for EU-27 ports and Irish ports, with Ireland-UK cargo volumes in particular dominated by Ro-Ro cargo.

Over the course of my study period, EU-27 ports traded with 199 different partner countries, while Irish ports traded with 108 countries. Figure 2.2 depicts the different partner countries connected via maritime transport to main ports in Ireland during 2018, and also lays out the volume of all cargo shipped with these partner countries (see Appendix B.2 for a full map extent of Figure 2.2). This reveals that Ireland is highly connected with the rest of the world through maritime transport, but also highlights the importance of the UK to Ireland as a trade partner, with 41.69 per cent of total cargo volumes handled by Irish ports in 2018 transported to or from the UK. The corresponding figure for Ro-Ro cargo is particularly high, at 86.31 per cent. In EU-27 ports, meanwhile, the UK accounted for 6.38 per cent of total cargo volumes and 19.53 per cent of Ro-Ro cargo volumes. These shares succinctly underline the greater exposure of Ireland to Brexit relative to the rest of the EU-27 member states. Figure 2.2 is also an illustration of the UK land-bridge trade route, as much of the Ro-Ro cargo shipped between the UK and Ireland will have originated in, or have been ultimately bound for, mainland

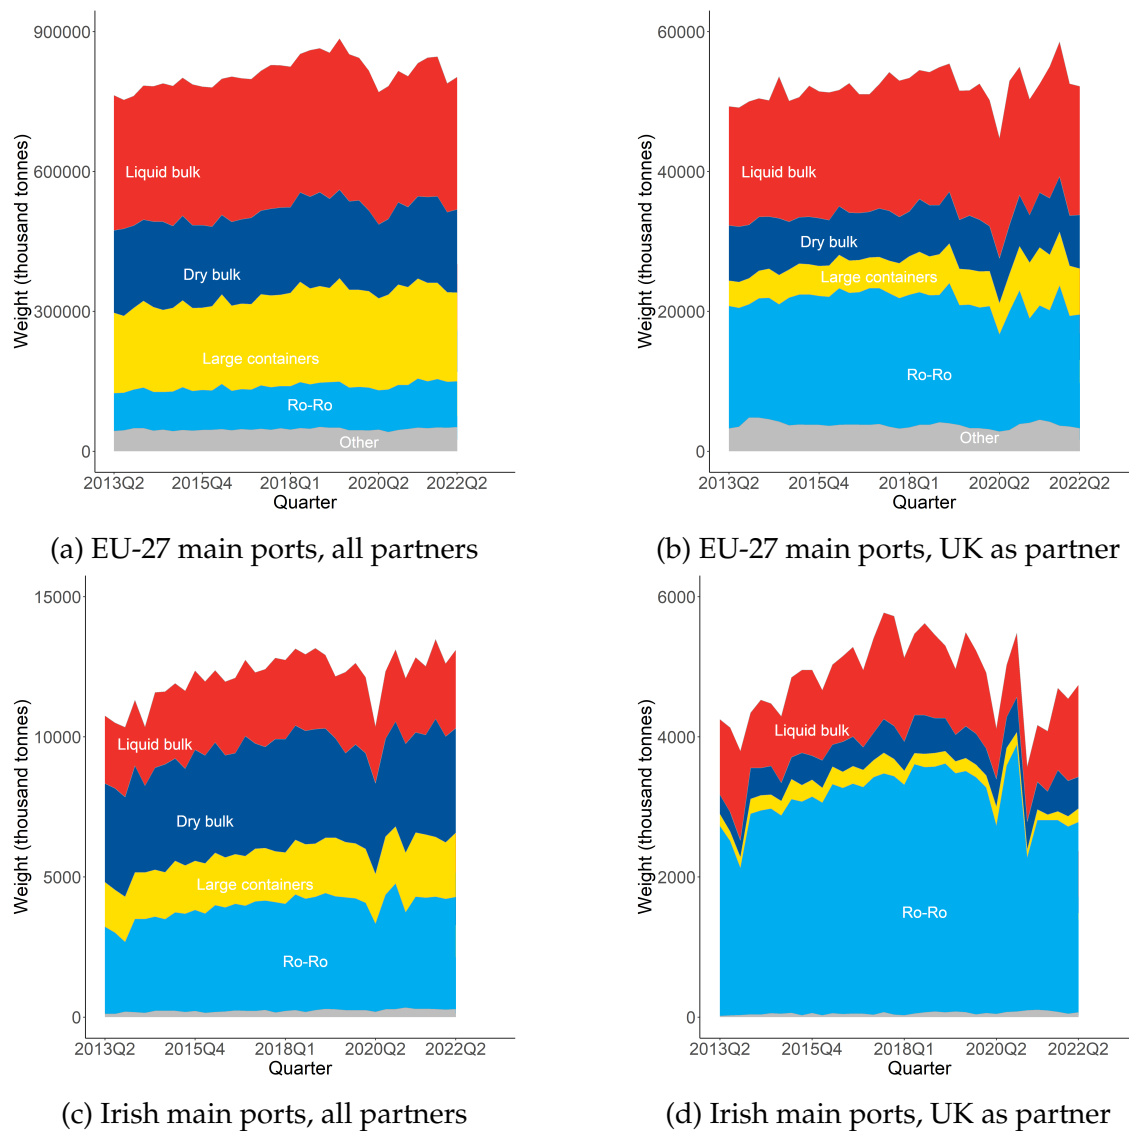


Figure 2.1: Cargo volumes by cargo type, 2013-2022. Sources: Author's analysis; Eurostat [2022a](#).

Europe.

Figure 2.3 compares the UK and the rest of the world in terms of total cargo volumes handled by EU-27 ports over my study period. This reflects the fact that the UK only accounts for a small portion of goods in EU-27 ports. However, Figure 2.3 shows again that the UK accounts for a much larger proportion of total cargo transported via Irish ports, and compares these volumes with the rest

2. ALL AT SEA?

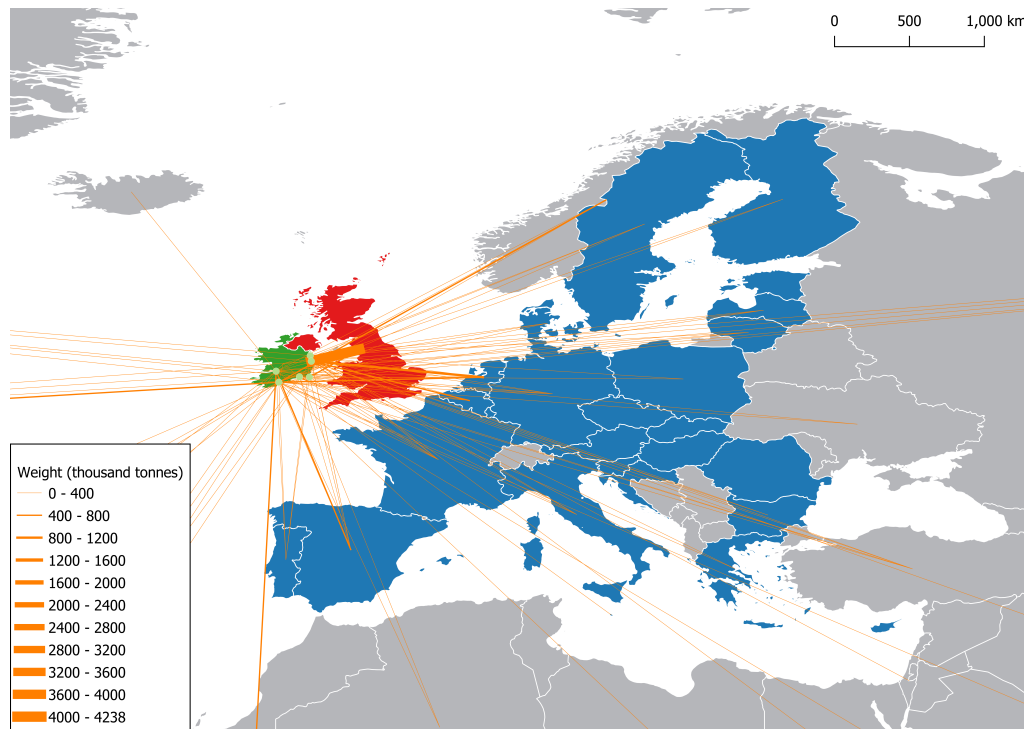


Figure 2.2: Total cargo volumes by partner country, Irish main ports 2018. Projection: WGS84. Sources: Author's analysis; Eurostat [2022a](#); Esri [2022](#).

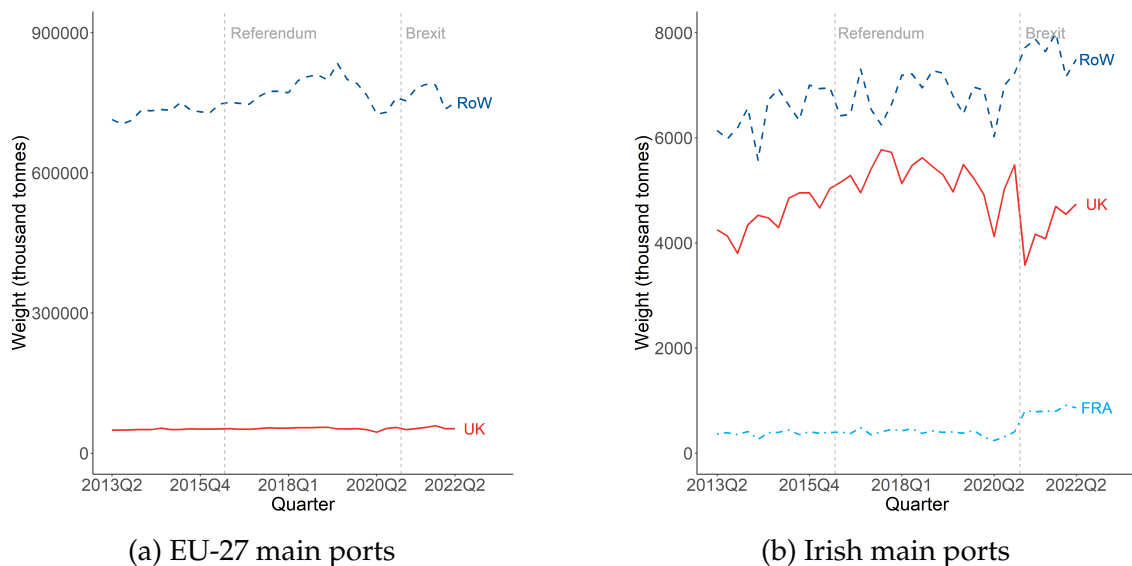
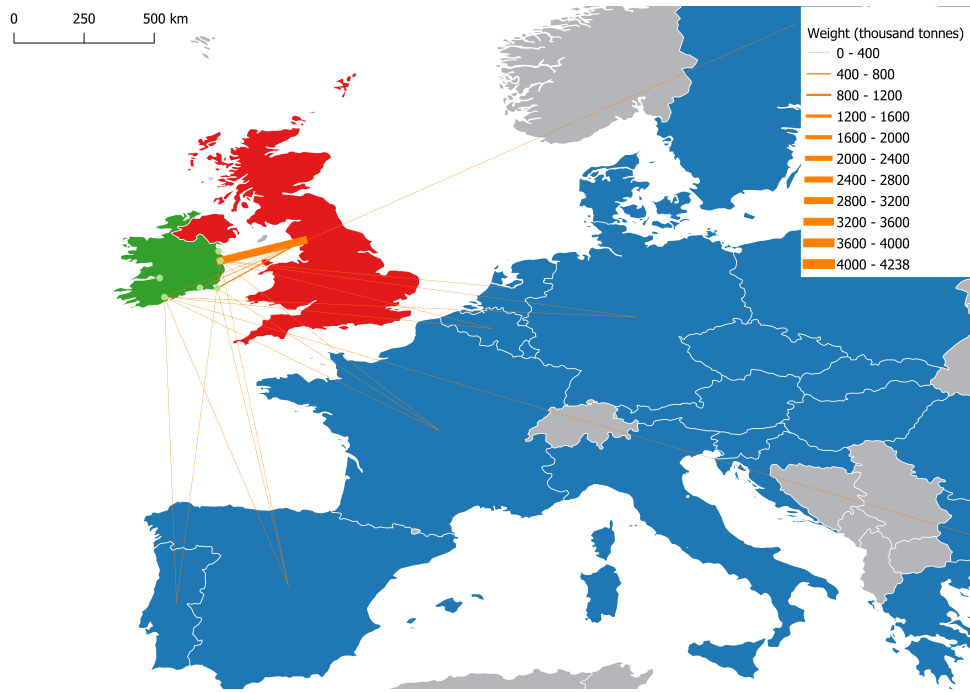
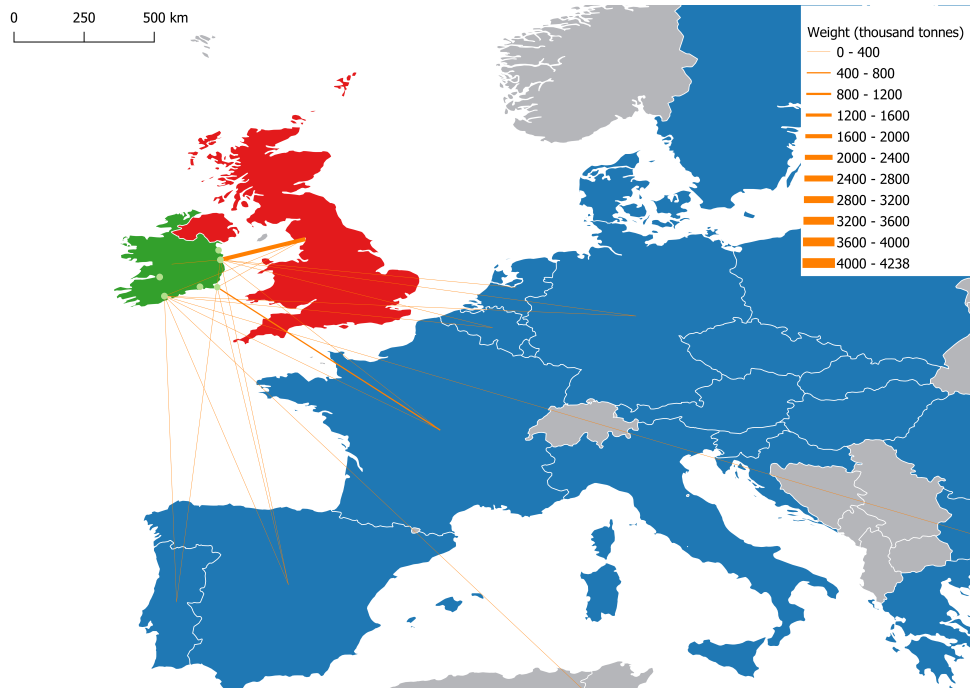


Figure 2.3: Total cargo volumes by partner entity, 2013-2022. Sources: Author's analysis; Eurostat [2022a](#).

of the world and with France, the shortest direct route to the continent. These plots suggest that Brexit may have had a much more pronounced effect on cargo



(a) 2020



(b) 2021

Figure 2.4: Ro-Ro cargo volumes by partner country, Irish main ports. Projection: WGS84. Sources: Author's analysis; Eurostat [2022a](#); Esri [2022](#).

volumes in Irish ports than across the wider EU-27, with a substantial decrease in cargo volumes between Ireland and the UK in the first quarter of 2021 compared with little discernible change in cargo volumes between the EU-27 and the UK. A difference-in-differences framework can help to assess whether Brexit can be isolated as the cause of these changes.

A shift in cargo from the UK land-bridge trade route to the direct Ireland-France route is also hinted at in Figure 2.3, with the first quarter of 2021 also experiencing a considerable increase in cargo volumes between Irish ports and France that persisted in subsequent quarters. As the UK land-bridge is mainly a route for Ro-Ro cargo, Figure 2.4 maps the volume of Ro-Ro cargo transported between Irish main ports and partner countries in 2020 and again in 2021. This reveals a decrease from 2020 to 2021 in Ireland-UK volumes, particularly via Dublin Port on the east coast of Ireland, and a concurrent increase in Ireland-France volumes, particularly via Rosslare Harbour on the south-east corner of Ireland, a Ro-Ro port that is geographically closer to France than Dublin Port. Figures 2.3 and 2.4 represent descriptive evidence of a shift from the land-bridge route to the direct route. These patterns are also worthy of further analysis in a difference-in-differences framework to determine if this is an effect that can be causally linked to Brexit.

2.3. Results

2.3.1. Cargo volumes

Table 2.3 displays exponentiated regression coefficients as my difference-in-differences estimates $e^{\hat{\delta}}$ for the effect of Brexit on all cargo (aggregating imports and exports) handled in EU-27 main ports (see Equation 2.4). Percentage changes in cargo volumes due to Brexit can be derived by subtracting 1 from these

exponentiated coefficients. Column 1 shows the effect on total cargo volumes, where volumes were aggregated across all cargo types, while columns 2-5 show results for regressions that used data disaggregated by cargo type. Across EU-27 ports, these results indicated no significant effect on total cargo volumes with the UK. However, this masked considerable heterogeneity across cargo types. Liquid bulk and Ro-Ro cargo account for the majority of EU-UK cargo volumes (see Figure 2.1). I found no effect of Brexit on liquid bulk, but significant positive effects on dry bulk and large containers. Column 3 of Table 2.3 reports an exponentiated regression coefficient for my difference-in-differences parameter of 1.094, suggesting that EU-UK dry bulk volumes increased by 9.4 per cent due to Brexit. Column 4 shows that large container cargo volumes increased by 53.3 per cent, while in Column 5 Ro-Ro volumes decreased by 21.7 per cent. Therefore, there was evidence in support of Hypothesis 1 in EU-UK Ro-Ro cargo volumes.

Table 2.3: Regression difference-in-differences estimates of Brexit effect on cargo in EU-27 main ports 2013-2022

	(1) Total	(2) Liquid bulk	(3) Dry bulk	(4) Large containers	(5) Ro-Ro
UK post-Brexit	1.020 (0.023)	1.024 (0.050)	1.094** (0.045)	1.533*** (0.030)	0.783*** (0.020)
Observations	598392	251604	324180	303696	134568
Pseudo R^2	0.258	0.223	0.202	0.275	0.377
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at partner country level
and transformed using delta method

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.4 reports results for my difference-in-differences analysis of Ireland-UK cargo volumes using data from Irish ports (see Equation 2.5). Unsurprisingly, given the UK's larger share in Irish cargo volumes than in wider EU-27 volumes, Table 2.4 indicates that Brexit dealt a heavier blow to cargo volumes in Irish ports.

Total cargo volumes between Irish ports and the UK (column 1) decreased by 24.1 per cent due to Brexit. However, this also caused an 87.7 per cent increase in total cargo volumes between Irish ports and France, offering clear evidence in favour of Hypothesis 2. Ireland-UK maritime cargo is predominantly in the form of Ro-Ro cargo, with liquid bulk also representing a significant share of cargo volumes (see Figure 2.1). When exploring heterogeneity in results between cargo types in columns 2-5 in Table 2.4, it was apparent that the result for total cargo was driven primarily by Ro-Ro cargo, with a 53.8 per cent decrease in Ireland-UK Ro-Ro cargo volumes and a concurrent 147 per cent increase in Ireland-France Ro-Ro cargo volumes due to Brexit, albeit from a much lower base. I found no significant effect on liquid bulk or dry bulk cargo volumes, though there was a 50 per cent decrease in large container volumes between Irish ports and the UK. While this reduction mirrored the result for Ireland-UK Ro-Ro cargo, no increase in Ireland-France large container cargo was detected.

Table 2.5 shows that the significant impact on Irish Ro-Ro cargo volumes was apparent in both imports and exports, further supporting Hypothesis 2. Column 1 of Table 2.5 repeats my main Ireland-UK Ro-Ro result, while columns 2 and 3 indicate a greater Ireland-UK impact on exports to the UK but a larger Ireland-France impact on imports from France, although a substantial effect was evident across imports and exports in both cases.

2.3.2. Robustness

I confirmed that these results were not sensitive to the inclusion of high-dimensional fixed effects, as the results were also evident using the PPML estimator developed by Correia, Guimarães, and Zylkin (2020) that identified and dropped separated observations (see Appendix B.2).

Table 2.4: Regression difference-in-differences estimates of Brexit effect on cargo in Irish main ports 2013-2022

	(1) Total	(2) Liquid bulk	(3) Dry bulk	(4) Large containers	(5) Ro-Ro
UK post-Brexit	0.759*** (0.060)	0.838 (0.185)	1.083 (0.115)	0.500*** (0.038)	0.462*** (0.026)
FRA post-Brexit	1.877*** (0.148)	0.793 (0.175)	0.893 (0.094)	1.060 (0.080)	2.470*** (0.138)
Observations	10116	3708	7020	2592	2772
Pseudo R^2	0.665	0.538	0.615	0.829	0.716
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at partner country level

and transformed using delta method

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.5: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022

	(1) Ro-Ro total	(2) Ro-Ro imports	(3) Ro-Ro exports
UK post-Brexit	0.462*** (0.026)	0.490*** (0.018)	0.423*** (0.044)
FRA post-Brexit	2.470*** (0.138)	2.656*** (0.097)	2.279*** (0.239)
Observations	2772	1512	2268
Pseudo R^2	0.716	0.725	0.717
Quarter fixed effects	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at partner country level

and transformed using delta method

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

To gain confidence in my assumption of parallel trends, I determined that there was no Brexit 'effect' on Ro-Ro cargo volumes between Sweden and the rest of the EU-27 in a placebo test (see Appendix B.2). Reassuringly, I also determined that my EU-UK and Ireland-UK Ro-Ro cargo results were robust to the inclusion of partner-specific linear trends (see Equation 2.6) in the outcome variable. As

2. ALL AT SEA?

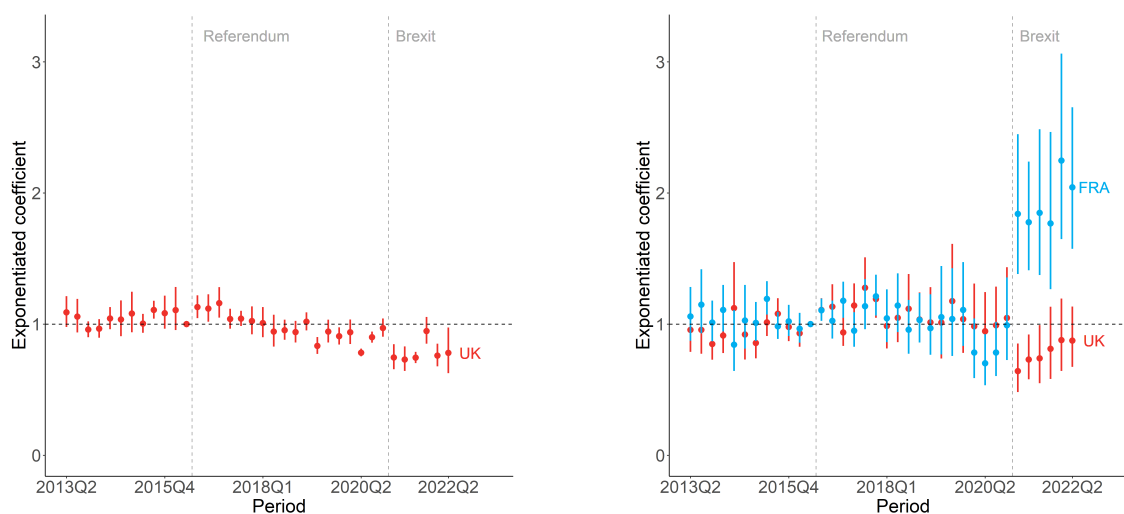
expected given the more restrictive model specification, precision was slightly lower in my EU-UK linear trend specification but treatment effects nonetheless remained statistically significant (see Appendix B.2).

Results for my post-referendum specification (see Equation 2.7) are also included in Appendix B.2. While standard errors were again marginally higher due to the more restrictive specification, my EU-UK and Ireland-UK results for Ro-Ro cargo persisted, with an additional negative effect on both EU-UK cargo volumes and Ireland-UK cargo volumes evident in the post-referendum period. The magnitude of the post-Brexit effects were larger than the magnitude of post-referendum effects in both EU-UK and Ireland-UK cargo volumes.

Finally, Figure 2.5 illustrates exponentiated coefficients for my more generalised model that included treatment leads and lags (see Equation 2.8) for EU-UK Ro-Ro cargo volumes and Ireland-UK total cargo volumes.⁶ In both settings, anticipatory effects prior to the June 2016 referendum were mostly insignificant, a finding that offered further credibility to the parallel trends assumption. There was some evidence of negative anticipatory effects in EU-UK cargo volumes during the latter stages of the period between the referendum and the exit of the Single Market, echoing the results of my EU-UK post-referendum regression specification (see Appendix B.2). In the case of cargo volumes in Irish ports, there was little evidence in Figure 2.5 of anticipatory effects between the referendum and the Single Market exit.

Clear post-treatment effects were evident in both settings, which was reflective of my main results (see Tables 2.3 and 2.4). In EU-UK cargo volumes, negative post-treatment effects were significant in all but one quarter from January 2021 onwards. In the first quarter of 2021, the post-treatment effect was estimated as a

⁶There were not enough observations to obtain consistent estimates for this generalised model using Ro-Ro cargo volumes in Irish ports.



(a) Ro-Ro cargo, EU-27 main ports

(b) Total cargo, Irish main ports

Figure 2.5: Regression difference-in-differences estimates and 95 per cent confidence intervals for Brexit effect on quarterly cargo volumes, generalised models 2013-2022. Exponentiated coefficients. Robust standard errors clustered at partner country level and transformed using delta method. Sources: Author's analysis; Eurostat [2022a](#).

25.5 per cent decrease in cargo volumes.

There were also strong post-treatment effects on both Ireland-UK and Ireland-France cargo volumes from the first quarter of 2021. For Ireland-UK cargo, the estimated treatment effect in the first quarter of 2021 was a 35.9 per cent decrease in cargo volumes. Figure 2.5 suggests that this shock may not have been persistent, with the effect levelling off and becoming insignificant by the fourth quarter of 2021. In the case of Ireland-France cargo, however, the Brexit effect appeared to be more persistent. The first post-treatment effect was an 84 per cent increase in cargo volumes, and this did not dissipate in later quarters, but rather increased in magnitude in the first half of 2022.

2.4. Energy consumption and emissions

What were the implications of this shift to the direct short sea shipping route between Ireland and France for energy consumption and emissions in the transport sector? A comprehensive inventory of the energy and emissions associated with transporting freight on each route would be a study in itself, but a rough approximation was possible with some simplifying assumptions.

One of the advantages of short sea shipping over road freight is that it is less energy intensive. Energy intensity for freight transport is typically measured in energy consumption per unit of weight and distance transported, for example in megajoules per tonne-kilometre (MJ/tkm). According to CE Delft (2021), the energy intensity of a truck with a gross vehicle weight (a measure of how much weight a truck can carry) of 10 to 20 tonnes towing a medium load on a motorway was 1.2 MJ/tkm in 2018.⁷ Meanwhile, combining figures from CE Delft (2021) and the International Maritime Organization (2021), I estimated that the energy intensity of a Ro-Ro vessel with a deadweight tonnage (a measure of how much weight a ship can carry) of between 5,000 and 9,999 and a medium to heavy load was 0.7 MJ/tkm in 2018.⁸ These energy intensities are summarised in Figure 2.7.⁹

For illustrative purposes, consider typical routes between Ireland and France using the land-bridge and using the direct route (see Figure 2.6). For goods being

⁷Deutsche Bahn (2019) also calculated an energy intensity of 1.2 MJ/tkm for road freight transport across their logistics network in 2018. For context, Deutsche Bahn (2019) calculated energy intensities of 0.3 MJ/tkm for rail freight and 9.8 MJ/tkm for air freight.

⁸While intensity figures specifically for Ro-Ro vessels were not available in CE Delft (2021), the study calculated an energy intensity of 0.3 MJ/tkm, and an emissions intensity of 19.1 g/tkm, for a general cargo ship with a deadweight tonnage of 5,000-9,999 and a medium to heavy load. This implied a coefficient of proportionality between energy consumption and carbon emissions of 63.7. Meanwhile, the International Maritime Organization (2021) calculated that the emissions intensity of a Ro-Ro vessel with a deadweight tonnage of 5,000-9,999 was 36.3 g/tkm in the same year, considerably higher than a general cargo vessel of similar size. Applying the same coefficient of 63.7, I estimated an energy intensity of 0.7 MJ/tkm for Ro-Ro vessels of this size.

⁹Given the scale economies of shipping, a vessel with a greater deadweight tonnage would be even less energy intensive. Both Deutsche Bahn (2019) and the International Energy Agency (2020) indicated that the energy intensity of international or ocean shipping was only 0.1 MJ/tkm in 2018.

exported from Ireland to France, the UK land-bridge route would involve three main stages. First, a Ro-Ro vessel from Dublin Port to Holyhead Port in the UK, a distance of 113km consuming 0.7 MJ/tkm of energy. Second, a truck from Holyhead to the Port of Dover which is 600km by road, consuming 1.2 MJ/tkm. Third, a second Ro-Ro vessel from Dover to the Port of Calais in France, consuming 0.7 MJ/tkm for 40km. Transporting a single tonne of goods along this route would therefore consume over 807 MJ of energy. Meanwhile, the direct route would involve a Ro-Ro vessel from Rosslare Harbour in Ireland to Cherbourg in France, a distance of around 600km by sea, consuming 0.7 MJ/tkm. The movement of 1 tonne of goods on this route would therefore involve the consumption of approximately 342 MJ of energy. This suggests that shifting cargo from the UK land-bridge to the direct route led to a decrease of 57.6 per cent in the energy consumption associated with transporting that cargo between Ireland and France.

A similar exercise can be conducted for carbon emissions. In freight transport, emissions intensity can be measured in grams per tonne-kilometre (g/tkm). According to the CE Delft (2021), a truck with a gross vehicle weight of 10 to 20 tonnes towing a medium load on a motorway emitted a mean of 83.1 g/tkm of carbon dioxide in 2018. Meanwhile, the International Maritime Organization (2021) estimated that a Ro-Ro vessel with a deadweight tonnage of between 5,000 and 9,999 was characterised by an average emissions intensity of 36.3 g/tkm as of 2018. These emissions intensities are summarised in Figure 2.7. Applying these figures over the two routes as with energy consumption, transporting 1 tonne of cargo via the land-bridge route would incur 55.4kg of carbon emissions, while moving that tonne of cargo using only short sea shipping directly would emit 21.8kg, or 60.7 per cent less.

Only a portion of the 53.8 per cent decrease in Ireland-UK Ro-Ro cargo volumes

2. ALL AT SEA?

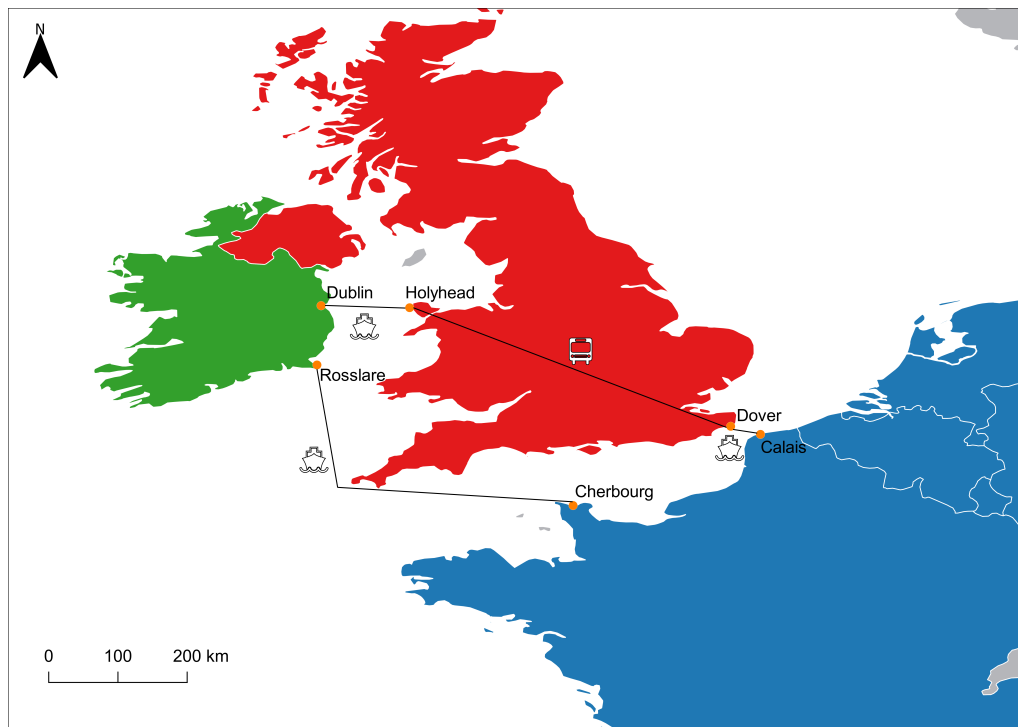


Figure 2.6: Typical routes for Ireland-France Ro-Ro cargo. Land-bridge route involves Ro-Ro vessel from Dublin to Holyhead, road truck from Holyhead to Dover, and Ro-Ro vessel from Dover to Calais. Direct route involves Ro-Ro vessel from Rosslare to Cherbourg. Projection: WGS84. Sources: Author's analysis; Eurostat [2022a](#).

due to Brexit will have shifted to the direct short sea shipping route, as some of this decrease will have been in direct Ireland-UK trade that never involved the rest of the EU-27. A previous study estimated that 53 per cent of Irish exports to the rest of the world was transported through Great Britain in 2016 (Lawless and Morgenroth [2017](#)). However, given the trade relationship between Ireland and France was not directly changed by Brexit, it seems likely that most of the Brexit-induced 147 per cent increase in Ireland-France Ro-Ro cargo volumes stemmed from freight transport diverting to avoid the UK land-bridge route. The mean quarterly volume of Ro-Ro cargo between Irish ports and France over the 2013 to 2020 period was 70,200 tonnes, indicating that this 147 per cent increase equated to roughly 33,000 tonnes of additional quarterly Ro-Ro cargo on average in 2021-2022.

2.4. Energy consumption and emissions

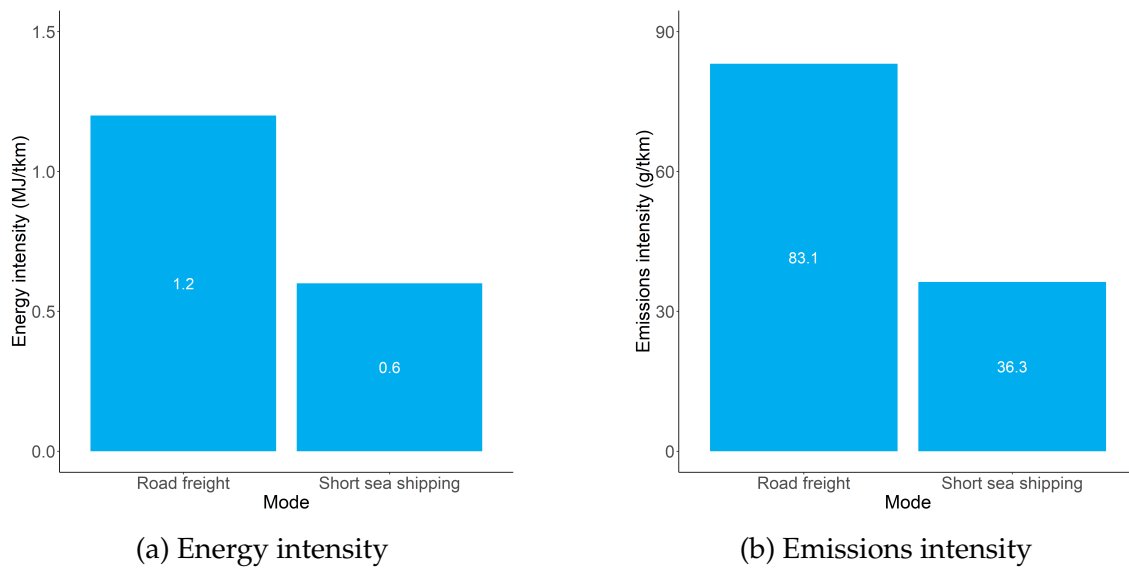


Figure 2.7: Intensity factors of road freight (truck with gross vehicle weight of 10-20 tonnes towing a medium load on a motorway) and short sea shipping (Ro-Ro vessel with deadweight tonnage of 5,000-9,999 and medium to heavy load). Emissions figures are for emissions arising from fuel combustion during vehicle or vessel use. Sources: Author's analysis; CE Delft 2021; International Maritime Organization 2021.

The above energy consumption and emissions calculations are based on comparing mean intensity levels by transport mode and on making several assumptions. For example, an individual ship's intensity would vary over time due to navigational or meteorological conditions beyond its control. The intensity of a truck travelling by road may be less sensitive to meteorological conditions but would be at the mercy of variable levels of congestion. Intensity levels while loading or waiting in port would be different. While I compared the shortest land-bridge route with the shortest direct route, there were also other land-bridge and direct routes in operation. Therefore, these calculations should only be interpreted as very rough approximations. In addition, these are approximations of energy consumption and carbon emissions reductions arising specifically from shifting goods from the land-bridge to the direct route; they do not consider the overall, general equilibrium effect of Brexit on energy consumption or carbon emissions, as this was beyond

the scope of this study.

However, these rough calculations serve to illustrate the underlying point: the transportation of freight using the direct short sea shipping route between Ireland and France required a considerably lower consumption of energy and level of emissions than the UK land-bridge route. If anywhere close to 33,000 tonnes of quarterly Ro-Ro cargo shifted to the short sea shipping route as a result of Brexit, the energy and emissions savings were substantial.

2.5. Implicit tariff on land-bridge

Non-tariff barriers to trade, such as customs checks, are more difficult to quantify than tariffs or quotas. Armed with results from my difference-in-differences analysis, it is worth revisiting the model of a representative exporter's transport route decision outlined in Equations 2.1, 2.2 and 2.3 to derive an approximation of the 'implicit tariff' Brexit imposed on the UK land-bridge route. Essentially, this exercise asked how large a tariff would have been required to produce the shift from the land-bridge to the direct short sea shipping route caused by the non-tariff barriers introduced by Brexit.

As demonstrated by Equation 2.3, the elasticity of substitution (the σ parameter) between the two routes plays a crucial role in determining the size of the implicit tariff (the τ parameter). First, for various possible values of σ , I used Equation 2.2 to back out the implied values of the land-bridge share parameter, η , for the 2013-2020 period prior to Brexit. This involved plugging observed values for all other variables into Equation 2.2. Of course, only a proportion of cargo volumes between Irish ports and the UK was accounted for by goods being shipped between Ireland and mainland Europe. Lawless and Morgenroth (2017) previously estimated that in 2016, 53 per cent of Irish exports to the UK was bound for the European

continent. Drawing on this estimate, for the pre-Brexit volume of cargo on the land-bridge route (the variable R_0 in Equation 2.2), I took the mean volume of Irish Ro-Ro exports to the UK over the 2013-2020 period, multiplied by 0.53. The volume of cargo on the direct route (the variable S_0) was more straightforward, as I simply used the mean volume of Irish Ro-Ro exports to France over the 2013-2020 period.

For the relative route cost, I relied on my energy consumption approximations from the previous section. These calculations suggested an energy consumption of around 807 MJ per tonne of cargo on the land-bridge route, and roughly 342 MJ per tonne on the direct route. Assuming a litre of diesel produces an average of 38 MJ of energy, I calculated a cost per tonne of cargo by multiplying diesel consumption levels by the mean of annual typical retail prices in pence per litre of ultra-low sulphur diesel in the UK over the 2013-2020 period using data from the UK Department for Energy Security and Net Zero (2023a). This suggested that the cost of the land-bridge route was higher than that of the direct route roughly by a factor of 2.36. For various realistic values of σ , it was then possible to back out values for the η parameter from Equation 2.2.

Second, holding these values for the σ and η parameters constant, I used Equation 2.3 with updated values for R_1 and S_1 to calculate values of the implicit tariff, τ , for each value of σ . Assuming both routes were faced with the same changes in the price of diesel, the relative route cost, $\frac{v_S}{v_R}$, was unchanged apart from the implicit tariff. For post-Brexit cargo volumes on each route, I adjusted the pre-Brexit volumes using the France post-Brexit effect on Irish Ro-Ro exports from Table 2.5. The decrease in Irish exports directly to the UK due to Brexit would have incorporated factors such as reduced UK demand in addition to cargo shifting away from the land-bridge route. Given the trade relationship between Ireland and France was unchanged by Brexit, however, it is more reasonable to assume

that most of the increase in Irish Ro-Ro exports to France was accounted for by displaced land-bridge cargo volumes. Therefore, I increased the cargo volume on the direct route by 147 per cent, such that $S_1 = S_0(2.47)$, and subtracted the level of this difference from the cargo volume on the land-bridge route, such that $R_1 = R_0 - (S_1 - S_0)$.

It was then possible to derive values of the implicit tariff, τ , for various realistic values of σ . Figure 2.8 plots these implicit tariff values against a range of substitution elasticities. If there was a high degree of substitutability between the land-bridge and direct routes, with an elasticity of substitution of between 5 and 10, the implicit tariff was calculated to be approximately 3 to 7 per cent of the cost of the land-bridge route. However, some Ro-Ro products that were particularly time-sensitive, such as agri-food products, may have been characterised by a lower level of substitutability between the routes given the superior transit time, flexibility and reliability of the land-bridge route. Figure 2.8 suggests that the implicit tariff would have been substantially higher for such products, for example with an implicit tariff of approximately 41 per cent of the land-bridge route cost calculated for a unitary elasticity of substitution.

2.6. Discussion

The withdrawal of the UK from the EU and Single Market is an intriguing case study in leaving a free trade agreement and increasing trade barriers. As the spectre of trade tariffs or quotas was averted by the EU-UK Trade and Cooperation Agreement, additional barriers to trade imposed by Brexit were specifically of the non-tariff variety. This study employed Eurostat ([2022a](#)) data on cargo volumes in all EU-27 main ports between 2013 and 2022 in a difference-in-differences investigation of the effect of Brexit on the shipping of goods between EU ports

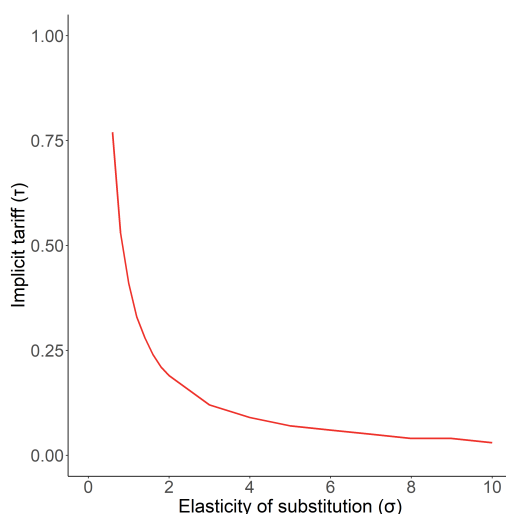


Figure 2.8: Implicit tariff on UK land-bridge route due to Brexit depending on elasticity of substitution between land-bridge and direct route. Source: Author's analysis.

and partner countries.

First, I posed the following research question: Was the reduction in trade flows due to Brexit reflected in maritime cargo volumes? My regression analysis indicated that a reduction in EU-UK cargo volumes was evident in roll on roll off (Ro-Ro) cargo, with a 21.7 per cent decrease in EU-UK Ro-Ro cargo volumes that could be causally attributed to Brexit based on my identification strategy. This result was robust to the inclusion of partner-specific linear trends in the regression model, and also persisted in more generalised models that included treatment leads and lags. While I focused on maritime cargo volumes, this result was roughly in line with a decrease of around 20 per cent in the value of all EU-UK trade as a result of Brexit found by Kren and Lawless (2022).

Ro-Ro shipping, where wheeled containers are rolled on and off ships, offers greater flexibility and speed in moving cargo from the quay to the ship and is predominantly used for shorter shipping routes. I showed that EU-UK maritime cargo consists mainly of Ro-Ro and liquid bulk, but found no effect on EU-

2. ALL AT SEA?

UK liquid bulk cargo volumes. That Ro-Ro was most negatively affected by Brexit suggests that congestion in Ro-Ro ports due to additional customs or documentation checks has represented the most significant non-tariff barrier to trade imposed by Brexit, as such congestion would detract from the speed advantages offered by Ro-Ro over other forms of shipping. Ke et al. (2022) had previously shown in simulation models that supply chain costs would increase due to congestion in the UK Ro-Ro port of Dover if customs checkpoints were established.

Conversely, I found a 9.4 per cent increase in dry bulk, and a 53.3 per cent increase in large container, EU-UK cargo volumes because of Brexit. These cargo types account for a much smaller proportion of EU-UK cargo volumes than Ro-Ro, and these effects could be interpreted as some substitution from Ro-Ro to large container shipping, and to a lesser extent, to dry bulk shipping.¹⁰ Ke et al. (2022) had speculated that increased congestion in the port of Dover could lead to a shift in lower-value cargo to the nearby large container port in Southampton. My finding of an increase in EU-UK large container cargo volumes due to Brexit could be interpreted as evidence in favour of this hypothesis.

Second, I explored the following research question: Has Brexit caused a diversion of maritime cargo from the UK land-bridge trade route between Ireland and mainland Europe to direct routes? For this question, I focused on cargo volumes handled by main ports in Ireland. I showed that the vast majority of Ireland-UK maritime cargo is accounted for by Ro-Ro cargo, which is unsurprising given the geographical proximity of the two countries. As with wide EU-UK patterns, my regression results revealed that Ireland-UK Ro-Ro cargo was most affected by Brexit, with a 53.8 per cent decrease in Ro-Ro cargo volumes. This was similar in

¹⁰In Appendix B.2, I provide additional results showing an increase in the share of large containers in total cargo due to Brexit.

magnitude to the 49 per cent decrease found by Flynn, Kren, and Lawless (2021b) in the value of all imports to Ireland from Great Britain in the early months of 2021. However, while Flynn, Kren, and Lawless (2021b) did not find an effect on the value of total Irish exports to Great Britain in early 2021, I found that the negative effect on Ro-Ro cargo volumes was evident in both imports and exports. Of course, it should be noted that while some of this decrease in Ireland-UK Ro-Ro cargo volumes can be accounted for by cargo shifting away from the UK land-bridge, much of the decrease will have stemmed from other factors such as reduced demand in the UK for imports from Ireland due to Brexit. Lawless and Morgenroth (2017) estimated that in 2016, 53 per cent of Irish exports to the rest of the world was transported via Great Britain, and Vega, Feo-Valero, and Espino-Espino (2018) pointed out that this UK land-bridge was mainly a route for Ro-Ro traffic.

The main alternative to this UK land-bridge route is the shortest direct shipping route between Ireland and mainland Europe, which is between Ireland and France. Prior to Brexit, the land-bridge route was more expensive in monetary terms for exporters than this direct route but boasted superior transit times and more frequent shipping connections (Vega, Feo-Valero, and Espino-Espino 2018). Increased congestion due to additional documentation checks at ports would have reduced the transit time advantage offered by the land-bridge route, thus boosting the relative appeal of the direct route. I determined that the substantial negative effect on Ireland-UK Ro-Ro cargo occurred alongside a 147 per cent increase in Ro-Ro cargo volumes between Ireland and France as a result of Brexit.

Given the trade relationship between Ireland and France was not changed by Brexit, this result represented clear evidence of a shift from the UK-land-bridge route to the direct Ireland-France route that occurred due to Brexit. These Ireland-UK and Ireland-France Ro-Ro results also held in more generalised regression

2. ALL AT SEA?

models that included partner-specific linear trends or treatment leads and lags. The differences in percentage changes between Ireland-UK and Ireland-France cargo volumes can largely be explained by the fact that the UK accounts for a much higher proportion of cargo in Irish ports than France, and the increase in Ireland-France volumes was therefore from a significantly lower base.

My generalised model including treatment leads and lags suggested that the negative effect on Ireland-UK cargo volumes dissipated over the course of 2021, with no significant post-treatment effect found after the third quarter of 2021. However, the results from this model indicated that the positive shock to Ireland-France cargo volumes was persistent, with the positive effect only increasing in magnitude in 2022. This pointed to some recovery in Ireland-UK cargo volumes over time, but also to the more long-term establishment of the direct short sea shipping route for Ro-Ro cargo between Ireland and France.

Unlike my results for wider EU-UK cargo volumes, I found no evidence of substitution between Ro-Ro and large container cargo due to Brexit, with Ireland-UK large container cargo volumes also negatively affected. I found no significant effect on liquid bulk or dry bulk volumes between Ireland and the UK. While there may have been some substitution towards large container shipping for EU-UK cargo, the main shift in Irish ports appears to have been towards direct Ro-Ro routes to mainland Europe.

Given the lower energy and emissions intensity of short sea shipping relative to road freight transport, the diversion of some cargo from the UK land-bridge route to the direct route would have led to reductions in the energy consumption and carbon emissions associated with this transport. Based on several assumptions, I calculated that energy consumption and emissions would have been approximately 60 per cent lower on the direct short sea shipping route. This rough figure

applies specifically to goods that shifted from the land-bridge to the direct route, and is not an estimate of a general equilibrium effect of Brexit on total energy consumption and carbon emissions.

Finally, drawing on the results of my difference-in-differences analysis, I used a constant elasticity of substitution production function to derive the ‘implicit tariff’ that was imposed on the land-bridge route by non-tariff trade barriers such as customs checks post-Brexit. I showed that the magnitude of this implicit tariff depended on the elasticity of substitution between the land-bridge and the direct route, ranging from approximately 3 per cent for goods characterised by a high elasticity to around 41 per cent for goods with unitary elasticity.

2.6.1. Policy implications

The findings of this study concur with previous studies that Brexit had a negative impact on trade, reflected here in cargo volumes. This offers additional empirical evidence, if any more was needed, to show that withdrawing from a free trade agreement is detrimental from the perspective of international trade. Importantly, it also highlights the fact that this negative effect is felt in the transport sector, and specifically in maritime shipping.

Findings in relation to the shift away from the UK land-bridge route for cargo between Ireland and the European continent are highly context-specific, in that they are based on Ireland’s geographic position as an EU-27 island member state to the west of Great Britain. However, there is a more general lesson for policy that can be taken from this finding: improving the transit time associated with short sea shipping relative to land-based alternatives can increase its attractiveness. In the context of this study, while the monetary cost of the direct route was always less than that of the land-bridge route, congestion due to additional customs checks at

ports due to Brexit increased the transit time of the land-bridge route, thus making the direct route relatively more appealing. As policy makers such as the European Commission seek to promote short sea shipping over land-based alternatives as part of making transport more sustainable, this unintended consequence of Brexit highlights transit times as a key factor to focus on. Transit times on short sea shipping routes could be targeted through investment in port infrastructure, for example.

2.6.2. Limitations and strengths

Some limitations of this study should be noted when interpreting its results and findings. First, cargo volume data was not available at the product level, and the econometric analysis abstracted away from possible heterogeneity in effects between product types. Second, while the difference-in-differences identification strategy followed the approach of similar studies in the literature, the identification of effects was based solely on temporal variation. Therefore, the possibility of some remaining omitted variable bias cannot be ruled out entirely, although the inclusion of a full set of period and partner country fixed effects should have minimised this risk. Third, while I focused on the UK land-bridge in this study, the Brexit effect on Ireland-UK Ro-Ro cargo volumes will have reflected factors such as reduced UK demand in addition to cargo shifting away from the land-bridge. Fourth, the energy consumption, carbon emissions and implicit tariff calculations relied on making several simplifying assumptions, and should be treated as rough approximations used to make a broad point rather than definitive figures. Fifth, due to a lack of quarterly port-level data from UK ports after 2020, it was not possible to examine the specific experience of Northern Ireland. As Northern Ireland was treated differently to the rest of the UK in terms of customs requirements following Brexit, subject to data requirements, assessing whether

the effects on cargo volumes in Northern Irish ports were different would be a fascinating avenue for further research.

However, this study can boast several key strengths. First, while other studies have focused on the effects on international trade more broadly, this study specifically considered Brexit from the perspective of transport and maritime shipping, including implications for transport sector energy consumption and carbon emissions. Second, it drew on Eurostat ([2022a](#)) quarterly port-level cargo volume data by cargo type and partner country to provide a highly detailed picture of the impact of Brexit on maritime shipping. Third, this rich data facilitated the use of a difference-in-differences identification strategy using a PPML estimator to furnish the literature with robust causal effects.

ALL HAIL?

The impact of ride hailing platforms on the use of other transport modes

3.1. Introduction

The modern advent of ride hailing offers an alternative form of transportation and a novel approach to car ownership, and thus represents an interesting development in the economics of transport. In this study, I examine the impact of ride hailing on other modes of transport in the cities of Glasgow and Edinburgh in Scotland.

In 2015, the Paris Climate Change Agreement established the target of limiting the global temperature rise to 1.5 degrees Celsius (Masson-Delmotte et al. [2018](#)). Having been responsible for 34 per cent of all carbon dioxide emissions in the United Kingdom (UK) in 2022 (UK Department for Energy Security and Net Zero [2023a](#)), transport is an important sector where a reduction in emissions is required to meet this target. The United Nations (UN [2021c](#)) has also identified sustainable transport as a key enabler of a range of Sustainable Development Goals. The key to whether ride hailing can be viewed as a positive development in the transition to sustainable transport lies in its impact on other transport modes.

3.1.1. What is ride hailing?

The term 'ride hailing' is used to describe an arrangement where an individual requests a specific trip and is matched via a mobile application with a driver willing to meet that demand with a private car (Tirachini 2019). It is among a range of mobility services linked to an emerging sharing economy that is based on internet platforms and smartphone apps (Miramontes et al. 2017). Ride hailing is also linked to an emerging concept of 'transport as a service', characterised by an increasingly blurred line between private and public transport with emphasis on the use rather than ownership of vehicles (Crozet 2020; Webb 2019; Miramontes et al. 2017). In the case of ride hailing, companies sign up car owners as drivers such that neither the company nor the consumer owns the vehicle in use (Crozet 2020).

Where available, ride hailing involves an easy-to-use app that empowers the consumer to request a trip whenever and wherever they want, with no need to find a local taxi phone number or hail on the street. The consumer can check the price of the trip in advance, avoiding the need to haggle with the driver or worry about being scammed. The payment itself is also handled through the app, removing the need to pay in cash. The app also facilitates checking the journey duration and route in advance, reducing the need to provide directions or worry about expensive detours. The consumer can follow journey progress on the app in real time, offering the security of knowing the journey is always proceeding towards the requested destination. In addition to these benefits, it has been argued that ride hailing boasts the potential to match passengers and drivers more efficiently than the street hailing of traditional taxis and can leverage internet-based technology to adjust prices in real time to consolidate supply and demand (Tirachini 2019). Dynamic pricing means that ride hailing can sometimes be cheaper than hailing a taxi, although this will not always be the case, particularly during periods of high

demand.

In a review of literature, Tirachini (2019) noted that other terms for ride hailing have included ‘ride sourcing’, ‘app-based ride services’, ‘ride booking’ and ‘on-demand ride services’, and that companies facilitating this with mobile application technology have been referred to as ‘transportation network companies (TNCs)’. Examples of such companies include Uber, Lyft, Cabify, Ola, DiDi Chuxing and RideAustin. Uber, having launched in 2009, had established a presence in approximately 800 cities within less than a decade (Tirachini 2019). Figure 3.1 presents Google search frequency data (Google Trends 2022) to illustrate the rapid rise of Uber between 2012 and 2019, both at a global level and in Scotland.




Figure 3.1: Frequency of Google searches for ‘uber’ search term, 2012-2019. Source: Author’s analysis; Google Trends 2022.

Although the terms have sometimes been used interchangeably in the literature, the concept of ride hailing is distinct from ‘ride sharing’, which is a platform used for carpooling rather than requesting a driver and private car, for example BlaBlaCar (Mitropoulos, Kortsari, and Ayfantopoulou 2021). TNCs often provide ride sharing in addition to ride hailing, for example Uber’s UberPool, Lyft’s LyftSharing or DiDi’s Hitch. In a survey of ride hailing journeys in Denver,

3. ALL HAIL?

Colorado, Henao and Marshall (2018) found that the average vehicle occupancy was 1.4, indicating that ride hailing tended to be utilised as a driver and private car service rather than a car pooling arrangement. The present study will focus on ride hailing.

3.1.2. Who uses ride hailing?

Ride hailing has been found to attract individuals seeking affordable, fast, point-to-point travel for journeys of between 10 and 30 minutes in duration (Tang et al. 2019). Many studies have examined the characteristics of ride hailing users, and factors in their decision to avail of ride hailing (for example Loa and Habib 2021; Wang and Noland 2021; Acheampong et al. 2020; Dias et al. 2019; Gehrke, Felix, and Reardon 2019; Mitra, Bae, and Ritchie 2019; Sikder 2019; Young and Farber 2019; Henao and Marshall 2018; Miramontes et al. 2017). The literature review conducted by Tirachini (2019) found that ride hailing users tend to be younger, more highly educated and wealthier than the general population, although Gehrke, Felix, and Reardon (2019) found a more balanced income distribution in a survey of ride hailing passengers in Boston, Massachusetts. In addition, in the US, Sikder (2019) presented evidence that African-Americans were less likely to use ride hailing, and that individuals working full-time with flexible hours were more likely than other workers or non-workers to utilise it.

Using US National Household Travel Survey 2017 data, Mitra, Bae, and Ritchie (2019) assessed the use of ride hailing among older adults. They found that those availing of ride hailing were more likely to be at the younger end of the age distribution, living alone, living in an urban area, more highly educated, wealthier and in possession of a smartphone (Mitra, Bae, and Ritchie 2019). Meanwhile, Wang and Noland (2021) employed data on ride hailing trips over the period of one month in Chengdu, China to show a positive association between ride hailing

usage and several spatial factors, including population density, housing prices and subway proximity.

The Tirachini (2019) review found that reasons for choosing ride hailing as a transport mode include cost, travel time and comfort. Another factor Tirachini (2019) noted in the literature was avoiding the need for a designated driver, which was consistent with their finding that ride hailing appears to be used predominantly for occasional leisure trips rather than regular commuting, with the majority of users only turning to ride hailing a few times a month (Acheampong et al. 2020; Tirachini 2019; Tirachini and Ríó 2019). The Denver survey by Henao and Marshall (2018) found that 94.5 per cent of passengers stated they were using ride hailing for the entire trip, rather than combining it with another transport mode. Peak demand times for ride hailing have been found to be Friday and Saturday nights in addition to typical 'rush hour' periods (Wang and Noland 2021; Henao and Marshall 2018).

3.1.3. How does ride-hailing affect other forms of transport?

In relation to public transport, does ride hailing represent a complementary or a substitute good? Tirachini (2019) highlighted this question as being crucial to whether ride hailing can be considered a sustainable alternative mode of transport. Stiglic et al. (2018) pointed to literature discussing the first- and last-mile problem in connecting a public transport network with trip origin and destination points, and contended that promoting the integration of ride hailing and public transport could potentially help to alleviate this. Equally, however, in a literature review on the effects of evolving urban transport systems, Webb (2019) noted that ride hailing could alternatively tempt individuals away from public transport in favour of a new, on-demand form of personalised transport.

3. ALL HAIL?

The existence of substitution or complementary effects of ride hailing on public transport is an active debate in academic literature. The Tirachini (2019) literature review found mixed results on such effects. The review indicated that across multiple cities, studies have tended to find the substitution effect of ride hailing on public transport to be greater than the complementary effect, and that it has thus added to congestion (Tirachini 2019). In addition, Tirachini (2019) found the relationship between car ownership and ride hailing to be disputed.

In terms of methodology, many studies of ride hailing's effect on other transport modes have relied on cross-sectional data and, therefore, could not furnish the literature with any causal relationships. These studies have also been typically based on stated preference information. For example, Acheampong et al. (2020) employed structural equation modelling using a cross-sectional survey of commuters in Ghana to show that ride hailing users stated they would have used traditional taxis, public transport or private cars to complete the reference trip in the absence of ride hailing. The study pointed to this as evidence of substitution, although with many surveyed ride hailing users stating they also used other transport modes on the same day, the overall impact of ride hailing was unclear. The Gehrke, Felix, and Reardon (2019) Boston survey found cross-sectional evidence that respondents with good access to public transport were more likely to replace public transport and thus increase the number of car trips, although there was some heterogeneity in this finding by household income. Henao and Marshall (2019) and Henao and Marshall (2018) both drew on the same cross-sectional survey data for Denver, with Henao and Marshall (2019) finding evidence that ride hailing was replacing driving trips to reduce demand for parking, and Henao and Marshall (2018) arguing that ride hailing led to an 83.5 per cent increase in vehicle kilometres. These two findings can be considered consistent with each other when accounting for factors such as 'deadheading', whereby a ride hailing

driver is driving without a passenger in between journeys (Henao and Marshall 2018). Shi et al. (2021) showed that 16.8 per cent of respondents in a cross-sectional survey of ride hailing users in Chengdu, China increased travel frequency due to ride hailing, suggesting additional induced demand, but that around half of the respondents stated they had substituted ride hailing for public transport, cycling or walking.

Evidence of substitution, particularly with ride hailing replacing public transport and taxis, was also found by Tirachini and R  o (2019) in a cross-sectional survey of residents of Santiago, Chile. Tang et al. (2019) found that many respondents in a cross-sectional survey of ride hailing users across Chinese cities stated they would have taken a taxi, public transport or private car in the absence of ride hailing, suggesting a substitution effect. In a large 2017 cross section, Mitra, Bae, and Ritchie (2019) found that among older adults in the US, ride hailing users make more public transport trips than non-users, suggesting a possible complementarity between the two modes for this age group. There has also been some cross-sectional evidence that an individual's level of ride hailing use matters to the impact on other modes, with Wang, Shi, and Chen (2021) finding that relative to occasional users, regular ride hailing users tended to own fewer vehicles among a cross-sectional sample of households in the US.

Contreras and Paz (2018) analysed a time series dataset of monthly travel by transport mode in Las Vegas from 2010 to 2016 using linear regressions, and found a negative association between ride hailing and the use of traditional taxis. However, as acknowledged by the authors, this methodology could be improved by employing a model more appropriate to count data as the dependent variable was a count of taxi pick-ups and drop-offs. Other studies have employed longitudinal or panel data to move further in the direction of causality in tackling these questions. For example, using longitudinal data on passenger numbers from

3. ALL HAIL?

2002 to 2015 for Canada and the US, Boisjoly et al. (2018) found no significant effect for the presence of ride hailing on the number of passenger trips on public transport.

A difference-in-differences methodology has become popular in this literature, with several studies adopting this approach in an effort to unearth causal effects of ride hailing on other transport modes (for example, Zhong et al. 2022; Shi, Li, and Xia 2021; Paundra et al. 2020; Zhong, Lin, and Yang 2020; Ward et al. 2019; Guo et al. 2018). These studies have typically deployed a dummy variable for the presence of a ride hailing platform as a treatment variable. Across Chinese cities, Guo et al. (2018) found a short-term increase in car sales due to the entry of DiDi, while Zhong, Lin, and Yang (2020) detected a negative impact on private car ownership. Zhong et al. (2022) showed a negative impact of ride hailing on the use of traditional taxis across cities in China, and Shi, Li, and Xia (2021) found that the presence of ride hailing reduced the number of bus passengers but increased the number of rail passengers. Using US data, Ward et al. (2019) found that the entry of ride hailing platforms into states reduced vehicle registrations. In Indonesia, Paundra et al. (2020) distinguished between a short-term negative impact, but a longer-term positive impact on new vehicle registrations.

The Tirachini (2019) review also highlighted literature on traffic externalities associated with ride hailing, including impacts on congestion and road safety. The number of road traffic collisions typically increases with traffic volumes, and a ride hailing effect on traffic could therefore be felt in road safety outcomes. The review noted an emerging literature using difference-in-differences approaches to establish a causal relationship between ride hailing and road traffic collisions, particularly alcohol-related collisions, in which results remain very mixed (Tirachini 2019).

Some studies have considered whether ride hailing affects congestion levels (Fageda 2021; Agarwal, Mani, and Telang 2019), which can be considered a more indirect way of assessing the impact on other transport modes. Agarwal, Mani, and Telang (2019) exploited ride hailing driver strikes in New Delhi, Bangalore and Mumbai in India in a difference-in-differences analysis of the effect of ride hailing on congestion, and found that the absence of ride hailing on strike days reduced travel times in all three cities. This finding that ride hailing increases congestion suggests that ride hailing is acting as a substitute for more sustainable travel modes. Conversely, however, using congestion data across 130 cities in Europe, Fageda (2021) found that the presence of Uber reduced average congestion, suggesting that ride hailing may be substituting private cars rather than public transport.

A related literature to these studies has considered cross-elasticities between transport modes, or the demand effect on one mode when an attribute of another mode is changed marginally. This cross elasticity can depend on a mode's own-elasticity of demand, each mode's relative market shares, and a diversion factor (Dodgson 1986). Fearnley et al. (2018) collated cross elasticity estimates for bus and rail from over 20 different sources, and while the review found low levels of substitution between modes, passengers were found to be more sensitive to time (such as in-vehicle time, waiting time) than fare variations when selecting a mode. Rose and Hensher (2013) analysed factors in the demand for taxi services in Melbourne and developed choice models to derive cross elasticities for taxis.

3.1.4. How has regulation responded to ride hailing?

It should be noted that the rise of ride hailing has not been without controversy. TNCs providing ride hailing emerged in a grey area of the regulatory environment, with regulation forced to catch up over time. This is in contrast with traditional

3. ALL HAIL?

taxi with whom ride hailing competes, which operate in a much more regulated environment. DiDi has been operating in China since 2012, but was only formally regulated and licensed from 2016 (Shi, Li, and Xia 2021). In Austin, Texas, RideAustin began providing a ride hailing platform in May 2016 after Uber and Lyft had withdrawn from the city due to disputes over local regulations, representing an example of tensions arising from regulation attempting to catch up (Dias et al. 2019). When Cabify and Uber entered the market in Chile in 2012 and 2014 respectively, despite being illegal, they thrived as a result of weak law enforcement coupled with high demand for the platforms. The central government began a process of legalisation and regulation in 2016 following violent clashes between Uber drivers and taxi drivers, another indication of the tensions created by the rise of ride hailing (Tirachini and Río 2019). In addition, there have also been serious safety concerns in relation to ride hailing and ride sharing, for example with DiDi suspending its Hitch platform in the wake of two separate female passenger deaths (Shah et al. 2021). However, while these issues are important in providing context to ride hailing, this study focuses on the impact of ride hailing on the use of other transport modes rather than on regulation, accountability, and safety.

3.1.5. This study

There remains clear disagreement in the literature as to the impact of ride hailing on the use of other transport modes such as public transport, and it is this question that I focus on in this paper. I contribute to this debate with an empirical analysis of the effect of the entry of Uber into the cities of Glasgow and Edinburgh in Scotland during 2015. In Section 2, I describe the difference-in-differences methodology I adopt using travel diary data from repeated cross sections of the Scottish Household Survey from 2012 to 2019. Results in Section 3 show that the

availability of ride hailing increased the use of public transport by just under 75 per cent relative to driving a car in Glasgow, although this effect was not reflected in Edinburgh. The increase in public transport use in Glasgow was more pronounced among respondents who were younger, male, and employed. Conclusions and policy implications are discussed in Section 4.

3.2. Methods

3.2.1. Theoretical framework

First, it is worth conceptualising the entry of a ride hailing platform into a city. Consider a general model of the market for taxis, based on the model proposed by Douglas (1972). To the discerning passenger, both time and money are valuable. In this model, therefore, for a given level of demand for journeys, demand for taxis is assumed to be a function of time and monetary cost:

$$Q = f_Q(w_m + w_p, P) \quad (3.1)$$

Specifically, demand for taxis Q is a function decreasing in the taxi fare P and decreasing in the sum of the time spent matching a driver to a passenger, w_m , and the pick-up time, w_p . These waiting times, in turn, are each functions of the number of vacant taxi vehicles, V , and the number of waiting passengers, $W = Qw_m$:

$$w_m = f_m(V, W), \quad w_p = f_p(V, W) \quad (3.2)$$

In the street hailing model of Douglas (1972), where matching occurs when a vacant taxi passes a passenger, the bulk of the waiting time is spent matching a passenger with a driver. Formally, the matching time is inversely proportional to

3. ALL HAIL?

the number of vacant vehicles while the pick-up time is close to 0 (Douglas 1972).

$$w_m \propto \frac{1}{V}, \quad w_p \rightarrow 0 \quad (3.3)$$

Arnott (1996) instead specified a radio dispatch model, where matching occurs when a passenger sends a request to a central dispatch centre. In this alternative model, it is the pick-up time that contributes most to the waiting time. Formally, the matching time is close to 0 while the pick-up time is inversely proportional to the square root of the number of vacant vehicles (Arnott 1996):

$$w_m \rightarrow 0, \quad w_p \propto \frac{1}{\sqrt{V}} \quad (3.4)$$

A ride hailing market could be thought of as being similar to the radio dispatch model of Arnott (1996), with an online platform taking the place of a central dispatcher. Compared with a human dispatcher armed with a telephone and radio system, the enhanced efficiency of the online matching platform could further reduce the matching time w_m towards 0. Moreover, the introduction of ride hailing in a city could also be viewed as an increase in the combined supply of vacant taxi and ride hailing vehicles V , particularly in settings where taxis are subject to regulation and licensing requirements. This would decrease both the matching time w_m and pick-up time w_p , and thus increase the combined demand for taxis and ride hailing. This represents the first hypothesis that was tested as part of this study:

Hypothesis 1: The introduction of ride hailing increased the proportion of journeys where the main transport mode was car as passenger or taxi.

To ensure that any such effect would not simply pick up a concurrent increase in

the supply of traditional licensed taxis, I confirmed that the number of licensed taxi vehicles in each Glasgow and Edinburgh was largely static between 2012 and 2019 (Transport Scotland 2020, see Appendix C.3).

Whether ride hailing helps or hinders the transition to sustainable transport hinges on another question: What impact did ride hailing have on the use of public transport (Tirachini 2019)? Bates (2018) described the distribution and modal split (DMS) model that forms an integral part of the widely used four-step model of transport demand. Bates (2018) outlined how the DMS model typically took the form of a conditional indirect utility function:

$$U_{j,m|i} = \log(A_j) + \alpha_m + \beta_c C_{i,j,m} + \beta^m t \left(t_{i,j,m}^1 + \sum_2^k w_k t_{i,j,m}^k \right) + \varepsilon_{i,j,m} \quad (3.5)$$

Equation 3.5 describes a function of the indirect utility $U_{j,m|i}$ derived from travelling to destination j using transport mode m , conditional on the journey starting in origin i . The attraction of destination j is measured by A_j , while α_m represents a constant for mode m . The monetary cost of the journey from i to j using mode m is given by $C_{i,j,m}$, while time is denoted by t . Specifically, the time spent travelling on the journey's main mode m is given by $t_{i,j,m}^1$, while all other travel time components numbered 2 to k are denoted by t^k and weighted by w_k . These other time components may include time spent walking to a bus stop, getting another secondary mode to the bus stop, or waiting at the bus stop, for example. Based on the law of demand, the β coefficients on cost and time would all be expected to be negative (Bates 2018).

The introduction of ride hailing in a city could be viewed as having the potential to reduce the overall time cost associated with a journey where the main mode is public transport. By acting as a secondary transport mode in a public transport journey, ride hailing could reduce the t^k that represents the time spent travelling

3. ALL HAIL?

to or from the nearest bus stop or train station. This is the essence of the last-mile argument made by Stiglic et al. (2018). If this were the case, ride hailing could be expected to have a complementary effect on public transport, and this was the second hypothesis tested in this study:

Hypothesis 2: The introduction of ride hailing increased the proportion of journeys where the main transport mode was public transport.

Alternatively, however, the introduction of ride hailing in a city could be regarded as the development of a mode of transport that competes with public transport to be chosen as the main mode for journeys. If ride hailing reduces the time or monetary cost associated with getting a taxi, this would increase the relative time t or monetary cost C associated with public transport. If this were true, ride hailing could be expected to act as a substitute for public transport.

Hypothesis 3: The introduction of ride hailing decreased the proportion of journeys where the main transport mode was public transport.

Increased road congestion is a possible externality that has been associated with ride hailing (Tirachini 2019). Drivers sign up to ride hailing platforms with their private cars to supply journeys, and while an increase in vacant vehicles V in Equation 3.4 may reduce the pick-up time w_p , it may also represent an increase in the total number of motor vehicles using a city's road network. Agarwal, Mani, and Telang (2019) analysed travel times to assess the effect of ride hailing on congestion. A final hypothesis that could be tested in the present study stemmed from this:

Hypothesis 4: The introduction of ride hailing reduced the

average journey speed of road-based journeys.

All four hypotheses were tested against the null hypothesis of ride hailing having no impact on either the choice of main mode for journeys or average journey speed respectively. Of course, to uncover any causal effects, an identification strategy was required to overcome issues related to unobserved confounding variables.

3.2.2. Study design

The causal relationship I explored in this study was the effect of ride hailing availability from 2015 on the use of other transport modes in Glasgow and Edinburgh. An ideal experiment to reveal this effect might have been to randomly allocate the population of each city into two groups, one that could avail of ride hailing and one that could not, and to compare changes in mode choice between the groups. However, as such an experiment was clearly not feasible, an alternative strategy to identify this causal relationship using applied microeconomic methods was required.

To test Hypotheses 1 to 3, following several studies in this literature (Zhong et al. 2022; Shi, Li, and Xia 2021; Paundra et al. 2020; Zhong, Lin, and Yang 2020; Ward et al. 2019; Guo et al. 2018), I employed a difference-in-differences methodology to identify any causal relationships between the presence of ride hailing platforms and the use of other transport modes. This method essentially involved comparing the average change over time between 2012 and 2019 in Glasgow and Edinburgh, where ride hailing became available, with the average change over time in two Scottish cities without ride hailing platforms, Dundee and Aberdeen. Specifically, I used the following regression design:

$$mode_{i,j,c,t} = \lambda_t + \gamma_c + \delta_c treated_{c,t} + X_{j,t} + \varepsilon_{i,j,c,t} \quad (3.6)$$

3. ALL HAIL?

In Equation 3.6, the outcome variable $mode_{i,j,c,t}$ was a categorical variable recording the main mode of transport used for journey i , undertaken by individual j , in city c and year t . I specified driving a car or van as the reference category of this outcome variable for the purposes of the regression. Year fixed effects were accounted for by λ_t , while city fixed effects relative to the control group were captured by γ_c . I also included a vector of individual control variables $X_{j,t}$, namely gender, age group, education and household income, in this specification. The parameter of interest was δ_c , as this was the difference-in-differences parameter on a $treated_{c,t}$ categorical variable, which specified whether the journey occurred in 2016 or later and either started or ended in Glasgow or Edinburgh. This specification allowed the estimation of two separate treatment effects, one for Glasgow and one for Edinburgh, but assumed a constant treatment effect within each city.

Given the nominal, categorical nature of the outcome variable $mode_{i,j,c,t}$, I estimated Equation 3.6 as a multinomial logistic regression using maximum likelihood. A multinomial logistic regression is a generalisation of a logistic regression that can be used to predict probabilities of different outcomes of a categorical variable with more than 2 possible discrete outcomes, given independent variables. Generally, given an outcome variable with 3 categories and vectors of control variables X and corresponding coefficients β , a multinomial logistic regression involves estimating a set of coefficients $\beta^{(1)}$, $\beta^{(2)}$ and $\beta^{(3)}$ for each possible outcome. To identify this model, it is necessary to arbitrarily set one outcome as the reference outcome. If $y = 1$ is set as the reference outcome such that $\beta^{(1)} = 0$, the multinomial logistic

regression model includes the following equations:

$$\begin{aligned}\Pr(y = 1) &= \frac{1}{1 + e^{X\beta^{(2)}} + e^{X\beta^{(3)}}} \\ \Pr(y = 2) &= \frac{e^{X\beta^{(2)}}}{1 + e^{X\beta^{(2)}} + e^{X\beta^{(3)}}} \\ \Pr(y = 3) &= \frac{e^{X\beta^{(3)}}}{1 + e^{X\beta^{(2)}} + e^{X\beta^{(3)}}}\end{aligned}\tag{3.7}$$

The relative probability of $y = 2$ can then be derived from this model:

$$\frac{\Pr(y = 2)}{\Pr(y = 1)} = e^{X\beta^{(2)}}\tag{3.8}$$

This relative probability is known as the ‘relative risk ratio’. Returning to my study design summarised in Equation 3.6, a relative risk ratio for a given transport mode indicates how the ‘risk’ of that mode being chosen as the journey’s main mode, relative to the risk of it being the reference mode, changes with a one-unit change in the respective independent variable, controlling for all other included independent variables.

Therefore, for a given transport mode, a multinomial logistic regression estimate for δ_c in Equation 3.6, denoted $\hat{\delta}_c$, would show the change in the multinomial log-odds of that mode being chosen over the reference transport mode (which I set as driving a car, the most common mode) due to ride hailing becoming available in city c . It is easier to interpret this coefficient when exponentiated (in other words, as the relative risk ratio), as the transformation $e^{\hat{\delta}_c} - 1$ is a percentage change in the relative probability of the respective mode being chosen over the reference mode. On this basis, I report all multinomial logistic regression results as exponentiated coefficients. For example, when examining results for the public transport category, a relative risk ratio that is greater than 1 tells us that an increase in the respective independent variable is associated with the chosen transport

3. ALL HAIL?

mode becoming more likely to be public transport. The standard errors reported for these exponentiated coefficients were transformed using the delta rule.¹ For each multinomial logistic regression, I also report McFadden's pseudo-R squared.² Untransformed regression coefficients and associated standard errors are also provided in Appendix C.2 for all multinomial regressions.

As individuals will typically select a transport mode from a range of options, the multinomial logistic regression is a more appropriate specification for mode choice than a binary logistic regression. Of course, as a decision, mode choice is also intertwined with destination choice. The DMS model specified in Equation 3.5 is a multinomial logistic regression as long as no partition between mode and destination choice is assumed and a Gumbel distribution is assumed for the random term $\varepsilon_{i,j,m}$ (Bates 2018). However, Bates (2018) highlighted the nested logit model as a more appropriate way of combining these decisions, with either mode or destination choice assumed to be conditional on the other. A nested logit would have required individual- and year-specific characteristics about each possible mode that were not available to this study, and this approach was therefore not feasible. Instead, the multinomial logistic regression required an assumption known as 'independence of irrelevant alternatives', which imposed that the odds of selecting one mode over another did not depend on the presence or absence of other 'irrelevant' alternative modes not explicitly included in the $mode_{i,j,c,t}$ outcome variable.

As established in setting Hypothesis 4, another potential aspect to the impact of ride hailing on other transport modes is the effect on congestion. Similar to Agarwal, Mani, and Telang (2019), I utilised a difference-in-differences approach

¹Specifically, the transformed standard error associated with an exponentiated multinomial logistic regression coefficient $e^{\hat{\delta}_c}$ was estimated as $e^{\hat{\delta}_c} \times SE(\hat{\delta}_c)$.

²This is calculated as $1 - \frac{ll(model)}{ll(null)}$, where ll denotes log likelihood, and is not a direct equivalent of the R squared statistic calculated for OLS regressions.

to detect any causal impact on average journey speed as a measure of congestion, in this case by fitting a linear regression specification using ordinary least squares (OLS):

$$speed_{i,j,c,t} = \lambda_t + \gamma_c + \delta_c treated_{c,t} + X_{j,t} + \varepsilon_{i,j,c,t} \quad (3.9)$$

The outcome variable in Equation 3.9, $speed_{i,j,c,t}$, was a continuous variable measuring the average speed in kilometres per hour of journey i , undertaken by individual j , in city c and year t . I calculated this variable as the journey distance divided by the journey duration. The right-hand-side of Equation 3.9 was the same as in Equation 3.6. I also ran two further linear regressions with journey distance and journey duration as the outcome variables instead of journey speed.

I report linear regression results for treatment effect estimates as regression coefficients. These coefficients show the change in the outcome variable, for example speed in kilometres per hour, as a result of ride hailing becoming available, again controlling for all other independent variables.

The crucial identifying assumption in the difference-in-differences design in Equations 3.6 and 3.9 is that the outcome in the absence of treatment can be captured with an additive structure that includes a city component that does not change over time, and a time component that does not change across cities, conditional on control variables. In Equations 3.6 and 3.9, these components were represented by γ_c and λ_t respectively. This is known as the parallel trends assumption. In other words, identification relies on trends in the outcome variable being the same in the treatment and control groups if the treatment did not occur, conditional on control variables and once the fixed effects are accounted for. This essentially allows the trend in the control group to be employed to impose a counterfactual trend on the treatment groups, with a treatment effect identified as a deviation from this counterfactual trend.

3. ALL HAIL?

This assumption would fall down if ride hailing became available in Glasgow and Edinburgh, but not Dundee or Aberdeen, due to some time-varying factor specific to Glasgow and Edinburgh. It is most likely, however, that Glasgow and Edinburgh were chosen by ride hailing companies simply due to them being larger population centres. As shown in Figure 3.2, while all four cities experienced a small increase in population between 2012 and 2019, the cities of Glasgow and Edinburgh boast much larger populations. The general population increase over time would have been accounted for by λ_t , and time invariant population differences across cities by γ_c , in Equations 3.6 and 3.9. It should also be noted that in this difference-in-differences design, identification was based solely on temporal variation (ride hailing only being available after 2015), in which case the possibility of an omitted variable bias could not be entirely ruled out. In Appendix C.3, however, I also discuss transport infrastructure in each of the 4 cities in my study setting, in addition to changes in fuel prices over my study period.

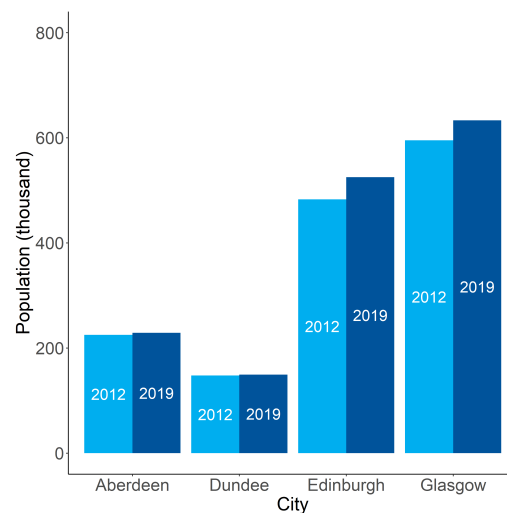


Figure 3.2: Population by city, 2012 and 2019. Sources: Author’s analysis; Office for National Statistics [2020](#).

A regression represents a convenient way of conducting a difference-in-differences analysis, as it allows for the calculation of standard errors and can also facilitate

the inclusion of control variables if required. Issues can arise with the calculation of standard errors in a difference-in-differences setting, however, as discussed in detail by Bertrand, Duflo, and Mullainathan (2004). One potential problem that may be relevant to my study design is that there may have been common unobserved factors that affected all journeys i made by the same individual j . A common solution suggested by Bertrand, Duflo, and Mullainathan (2004) is to cluster standard errors. When running the regressions described in Equations 3.6 and 3.9, I therefore clustered standard errors at the individual level j . This allowed for error terms to be correlated between journeys recorded by the same individual, albeit at the possible expense of precision. I assessed the merits of this approach by comparing standard errors calculated using various variance estimators (see Appendix C.2). Another possible problem is that my treatment variable, $treated_{c,t}$, changed very little over time within a city, and this could have produced serial correlation. However, in this case, clustering at the city level c with only 4 cities would likely lead to bias in calculating standard errors (Bertrand, Duflo, and Mullainathan 2004).

I conducted this regression analysis using Stata/MP 16.1.

3.2.3. Alternative specifications

To assess whether results were sensitive to the choice of combining Dundee and Aberdeen to construct a control group, I ran my main difference-in-differences regression (Equation 3.6) again using only Dundee-based journeys as the control group, and then using only journeys based in Aberdeen as the control group. Each of these alternative regressions involved dropping journeys based in the other city from the sample, thus reducing sample size in addition to altering the control group.

3. ALL HAIL?

A common diagnostic in studies seeking to make causal inference is a ‘placebo’ test, or falsification test, which involves testing for the presence of the effect being studied in a setting where it should not occur. If I detected a ride hailing ‘effect’ among journeys in settings where ride hailing did not become available, this would raise serious concerns about the integrity of the key parallel trend assumption and thus about inferring causality. I conducted a placebo test by running my main difference-in-differences regression (Equation 3.6) again, maintaining Dundee and Aberdeen as the control group, but with journeys based in Glasgow or Edinburgh replaced with all non-city journeys that had been dropped from my main sample. Therefore, this test compared the average change over time between 2012 and 2019 in all mainland non-city local authorities with the average change over time in the cities of Dundee and Aberdeen. As ride hailing never became available in these areas, no ‘effect’ should be detected in this test. This placebo test is a similar procedure to the ‘control experiment’ conducted by Duflo (2001) in a study of economic returns to education.

The parallel trends assumption underpinning the difference-in-differences methodology may also come unstuck if a city-specific pre-existing trend in the outcome variable is present. For example, if the proportion of journeys using public transport in Glasgow or Edinburgh was generally increasing prior to the entry of ride hailing, it would be difficult to properly disentangle any further increase due to ride hailing from this general trend. This is an issue with difference-in-differences that was discussed in detail by Wolfers (2006). One useful test for the presence of pre-existing trends is to additionally control for panel-specific trends (Angrist and Pischke 2008). Angrist and Pischke (2008) highlighted a study of labour regulation in India by Besley and Burgess (2004) as an example of this difference-in-differences approach in econometric literature. Specifically, I applied the following design to a multinomial logistic regression model to assess the

robustness of my main results to the inclusion of linear trends:

$$mode_{i,j,c,t} = \lambda_t + \gamma_{0,c} + \gamma_{1,c}t + \delta_c treated_{c,t} + X_{j,t} + \varepsilon_{i,j,c,t} \quad (3.10)$$

In Equation 3.10, the λ_t and $\gamma_{0,c}$ parameters still accounted for year and city fixed effects respectively and $X_{j,t}$ still represented a vector of individual-level controls, while $\gamma_{1,c}$ captured city-specific linear trends in the outcome variable. The difference-in-differences parameter was again δ_c on the $treated_{c,t}$ variable. In this specification, the identification of an effect derived from whether ride hailing led to deviations from existing city-specific trends. This represented a more restrictive version of Equation 3.6 due to the inclusion of linear trends, and was therefore expected to increase standard errors.

However, as Wolfers (2006) pointed out, controlling for linear trends in this manner is problematic in any context where the treatment effect might be dynamic. For example, if the effect of ride hailing on mode choice increased over time as more drivers registered with the ride hailing platform, controlling for a linear trend while not modelling this impact dynamically could result in the linear trend picking up the post-treatment pattern. While it is not clear that ride hailing has a dynamic rather than constant impact on mode choice, this cannot be ruled out and regression results based on Equation 3.10 should therefore be interpreted with caution.

Another method for testing the robustness of the parallel trends assumption in a difference-in-differences framework suggested by Angrist and Pischke (2008) is to include lags and leads of the treatment in a more generalised model, as was done by Duflo (2001) and Autor (2003) in econometric literature, for example. This involves estimating δ_c coefficients for different years, and these can then be plotted as an additional test of causality in the spirit of Granger (1969). Essentially, this

3. ALL HAIL?

method could furnish the study with some evidence on whether causes happened before consequences, rather than the other way around. Specifically, to further assess the robustness of my main results, I applied the following design to a multinomial logistic regression model:

$$mode_{i,j,c,t} = \lambda_t + \gamma_c + \sum_{\tau=0}^3 \delta_{c,-\tau} treated_{c,t-\tau} + \sum_{\tau=1}^3 \delta_{c,+\tau} treated_{c,t+\tau} + X_{j,t} + \varepsilon_{i,j,c,t} \quad (3.11)$$

In Equation 3.11, λ_t and γ_c again represented year and city fixed effects respectively and $X_{j,t}$ denoted a vector of individual-level controls. This time, rather than obtaining a single estimate for each constant δ_c , separate estimates for $\delta_{c,-\tau}$ were obtained for 4 years after ride hailing becoming available, and estimates for $\delta_{c,+\tau}$ were obtained for 3 years prior to the introduction of ride hailing. The set of $\delta_{c,-\tau}$ parameters are known as post-treatment effects, while the set of $\delta_{c,+\tau}$ parameters are known as anticipatory effects. As with Equation 3.10, this was a more data-intensive version of Equation 3.6, and was thus expected to produce higher standard errors.

If there was an effect of ride hailing on mode choice and the parallel trends assumption held, post-treatment effect coefficients should be statistically significant while anticipatory effect coefficients should not be significant. In other words, there should only be evidence of an effect once ride hailing was available, assuming individuals were not altering their mode choices in anticipation of ride hailing becoming available. While this form of Granger causality test is still not conclusive in definitively proving causality, it can provide confidence in the parallel trends assumption (Angrist and Pischke 2008).

3.2.4. Data

In this study, I employed data from the Scottish Household Survey (SHS) for each year from 2012 to 2019. This granular data offered a comprehensive picture of journeys undertaken by a representative sample of individuals over an 8-year period in cities with and without access to ride hailing, in addition to a detailed account of the socio-demographic characteristics of these individuals.

The SHS is an annual, cross-sectional survey of the characteristics, attitudes and behaviours of households and individuals across Scotland. The primary objective of the survey is to make representative estimates for the country. Each year, the survey targets a large sample size of 10,450 households, with a minimum of 250 households in each local authority to facilitate an analysis of all local authority areas. The Royal Mail's (UK postal service) Postcode Address File is used as the sample frame for address selection, and addresses selected for the survey are then removed from the sample frame for a period of at least 4 years (SHS 2020).

The survey is conducted primarily via computer assisted personal interviewing (CAPI), which involves face-to-face interviews in respondents' homes that are supported by a computer. The first component of the interview, which captures data on the composition and characteristics of the household, is completed by the highest income householder or their partner. The second component is completed by a random adult in the household (aged 16 or over), and this gathers information on the attitudes and experiences of the random adult. For each annual survey during the 2012-2019 period, fieldwork for the survey was completed between January and either February or March (SHS 2020).

I sourced journey data from the travel diary component of the SHS, which is completed by the random adult. During the CAPI interview, the random adult is requested to complete a diary of their travel behaviour during the day prior to the

3. ALL HAIL?

interview. Any journeys undertaken by the random adult over the course of the day are recorded, including details of the start and end local authority area, the main mode of transport used, any other modes used during different stages of the journey, the purpose of the journey, and the journey distance and duration (SHS 2020).

In total, 152,219 journeys were recorded in this manner by the SHS during the 2012-2019 period. To focus the analysis on the four main cities in Scotland, I reduced this to a sub-sample of 43,169 journeys that either started or ended in the local authority areas of Glasgow City, City of Edinburgh, Dundee City or Aberdeen City. Further details on how I reduced the sample size are provided in Appendix C.1. An illustration of these journeys by local authority area is presented in Figure 3.3, and descriptive statistics for journey transport modes and journey purposes are provided in Table 3.1.

Table 3.1: Descriptive statistics of included journeys, 2012-2019

Variable	Category	Count	Percentage
Main mode	Car as driver	18786	43.52
	Public transport	7096	16.44
	Car as passenger	5029	11.65
	Walk	10431	24.16
	Other	1827	4.23
Journey purpose	Work	11224	26
	Education	2438	5.65
	Shopping	9900	22.93
	Health	1116	2.59
	Leisure	8526	19.75
	Returning home	5258	12.18
	Going for walk	1510	3.5
	Other	3197	7.41
Total		43169	100.00

Car as passenger includes taxi.

Sources: Author's analysis; SHS 2020

These journeys were spread across 16,712 respondents, implying that on average, each random adult recorded 2.6 journeys in the travel diary. The socio-

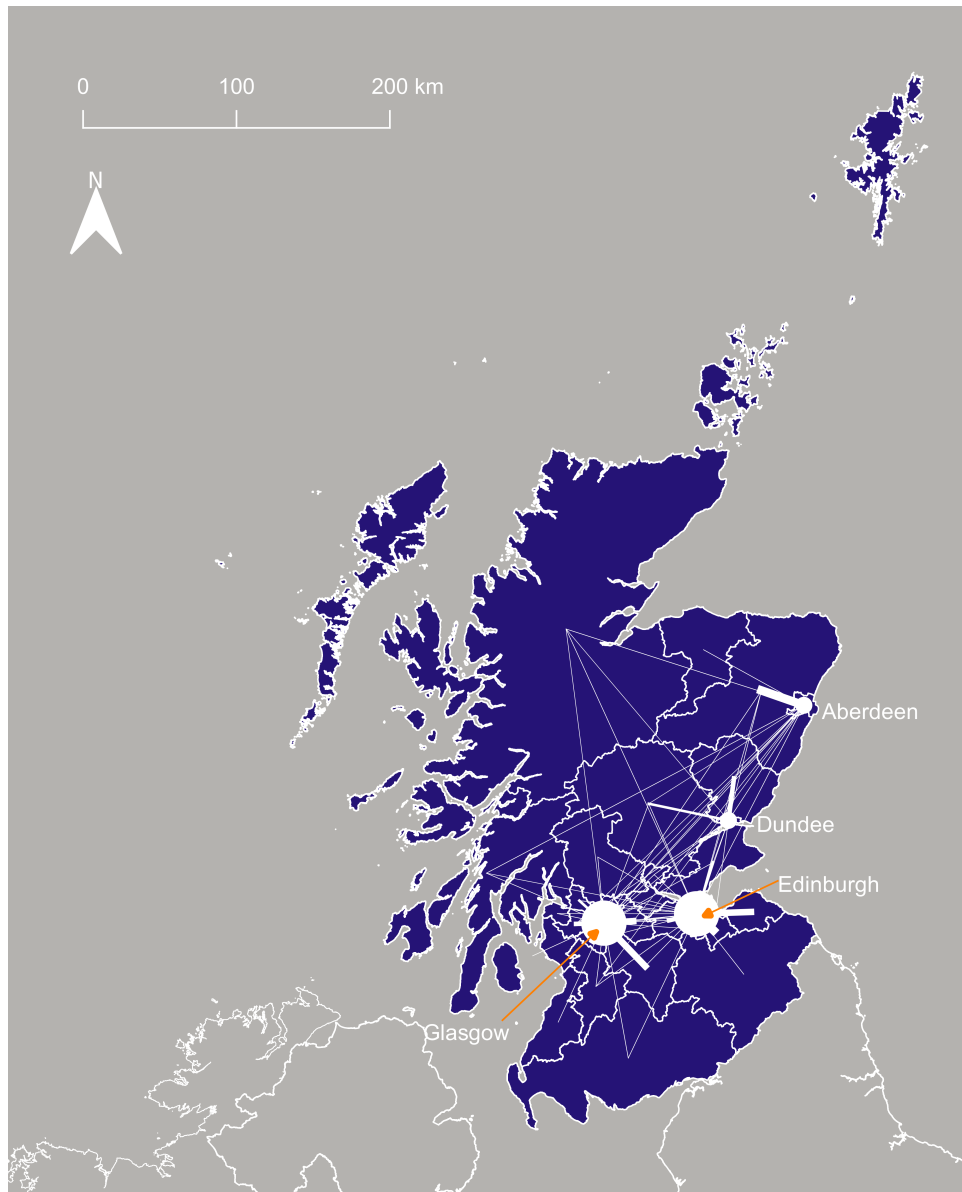


Figure 3.3: Journeys starting or ending in Edinburgh, Glasgow, Dundee or Aberdeen, 2012-2019. Line weight is defined by the number of journeys between two local authorities. Point size is defined by the number of journeys within a city. Source: Author's analysis; SHS 2020.

demographic characteristics, including age, gender and economic status, of these random adults are summarised in Table 3.2.

The socio-demographic characteristics of the random adults were broadly similar between the control group and each Glasgow and Edinburgh, with a higher

3. ALL HAIL?

Table 3.2: Descriptive statistics of included random adults, 2012-2019

Variable	Category	Count	Percentage
Age group	16-24 years	1679	10.05
	25-44 years	6097	36.48
	45-64 years	5602	33.52
	65+ years	3291	19.69
	Missing	45	0.27
Gender	Male	7847	46.95
	Female	8822	52.78
	Missing	45	0.27
Economic status	Self-employed	979	5.86
	Employed	9063	54.22
	Looking after home/family	649	3.88
	Retired	3610	21.6
	Unemployed	563	3.37
	In education	1156	6.92
	Not working due to illness/injury	583	3.49
	Other	66	0.39
	Missing	45	0.27
	Marital status	Never married	6546
Married		6645	39.76
Separated		544	3.25
Divorced		1640	9.81
Widowed		1293	7.74
Missing		46	0.28
Highest education	Secondary, National	2601	15.56
	Secondary, Higher	2609	15.61
	Further education	1836	10.98
	Degree or higher	6651	39.79
	Other	823	4.92
	No qualification	2149	12.86
Household income	Missing	45	0.27
	£0-10,000	1841	11.01
	£10,000-20,000	4618	27.63
	£20,000-30,000	3468	20.75
	£30,000-40,000	2298	13.75
	£40,000-50,000	1689	10.11
	>£50,000	2197	13.14
Missing	603	3.61	
Total		16712	100.00

Highest education and total household annual net income measured at household level.

Sources: Author's analysis; SHS [2020](#)

level of education evident in Edinburgh (see Appendix C.2). I also compared distributions of age and household income between the control groups and each treatment group, further indicating largely similar demographic and socio-

economic characteristics (see Appendix C.2). Any time invariant, city-specific socio-demographic factors that affected my outcome variables were picked up by γ_c in Equations 3.6 and 3.9. I ran t-tests of differences between 2012 and 2019 in socio-demographic characteristics of the random adults in each group to determine whether any of these factors were time-varying (see Appendix C.2). Household income increased in each group over the period, and there was also evidence of increasing education levels in Glasgow and the control group. General time-varying socio-demographic factors that were not city-specific and that affected my outcomes were accounted for by λ_t in Equations 3.6 and 3.9. The t-tests revealed a 6 per cent increase in the average age among random adults making journeys to or from Edinburgh that was not reflected in Glasgow or the control group. The inclusion of individual-level control variables, including age group, helped account for this Edinburgh-specific increase.

Surveys that aim to glean information about a target population will typically apply survey weights to collected data to account for response rate differences between groups and for unequal selection probabilities. When making inferences about the population of Scotland, the SHS (2020) applies survey weights at the household, random adult and travel diary level. Throughout my difference-in-differences analysis, I applied travel diary weights to journeys. During fieldwork for all cross sections of the SHS (2020), disproportionately fewer interviews were conducted on Friday, Saturday and Sunday, and disproportionately more adults in full-time employment were interviewed over the weekend. Based on this, to calculate travel diary weights, the SHS (2020) rescaled the random adult weights to ensure travel diaries were representative of travel patterns over the entire week and of working status across each day of the week (see SHS 2020 for further details). As a robustness check, I also assessed whether my results were sensitive to the inclusion of these travel diary weights.

It is worth noting that as with any self-reported data, my outcome variables were likely subject to a small degree of measurement error. For example, respondents may have inaccurately recalled some details of their travel patterns, although this possibility was minimised by travel diaries being recorded for the day immediately prior to the interview. Alternatively, it is possible that respondents could have deliberately omitted or misrepresented some information on their travel patterns. The possibility of small errors in capturing travel diary data cannot be ruled out either, although the use of CAPI would have minimised this risk. Such measurement error would have increased noise in my regression specifications, making the detection of any true effect more difficult.

3.2.5. Outcome variables

My main outcome variable of interest was the main mode of transport used for the recorded journey (see Equation 3.6). I collapsed this into a five-category variable (see Appendix C.1 for details) for the purposes of this empirical analysis, for example by combining bus and train categories into a public transport category. Category frequencies are included in Table 3.1, and Figure 3.4 illustrates this modal split by destination city. As shown in Figure 3.4, driving a car boasted the largest share of journeys across all four cities, with walking the second largest across all cities.

I also employed average journey speed, as well as journey duration and journey distance, as outcome variables (see Equation 3.9). For each recorded journey, the respondent specified the distance of the journey in kilometres and the duration in minutes. These variables were available as continuous variables in the SHS (2020) data. I calculated average journey speed as the distance divided by the duration and converted this speed to kilometres per hour. Descriptive statistics for these three variables are provided in Table 3.3.

Table 3.3: Descriptive statistics for continuous variables of included journeys 2012-2019

	N	Mean	S.D.	Min.	Max.
Journey distance (km)	43,169	7.07	9.49	0.00	57.33
Journey duration (mins)	43,169	21.96	15.19	1.00	70.00
Journey average speed (km/h)	43,169	17.14	15.71	0.00	207.98

N denotes observations. S.D. denotes standard deviation

Sources: Author's analysis; SHS 2020




Figure 3.4: Journey main mode choice by city, 2012-2019. Car as passenger includes taxi. Other includes bicycle, motorcycle/moped, ferry, air, horse-riding and tram. Sources: Author's analysis; SHS 2020.

My difference-in-differences methodology compared the average change over time between 2012 and 2019 in Glasgow and Edinburgh with the average change over time in the control group comprised of Dundee and Aberdeen. Figure 3.5 illustrates the proportion of journeys where the main transport mode was car as driver, public transport and car as passenger³, as well as the mean of my calculated average journey speeds, by year for each Glasgow and Edinburgh and the combined control group of Dundee and Aberdeen. One of the most striking

³As shown in Appendix C.1, I collapsed the 'Taxi/minicab' category into the 'Car/van as passenger' category as the taxi category was simply too small to be reliable in my empirical analysis. In addition, it is possible that a ride hailing journey could be recorded in either of these two categories by respondents.

3. ALL HAIL?

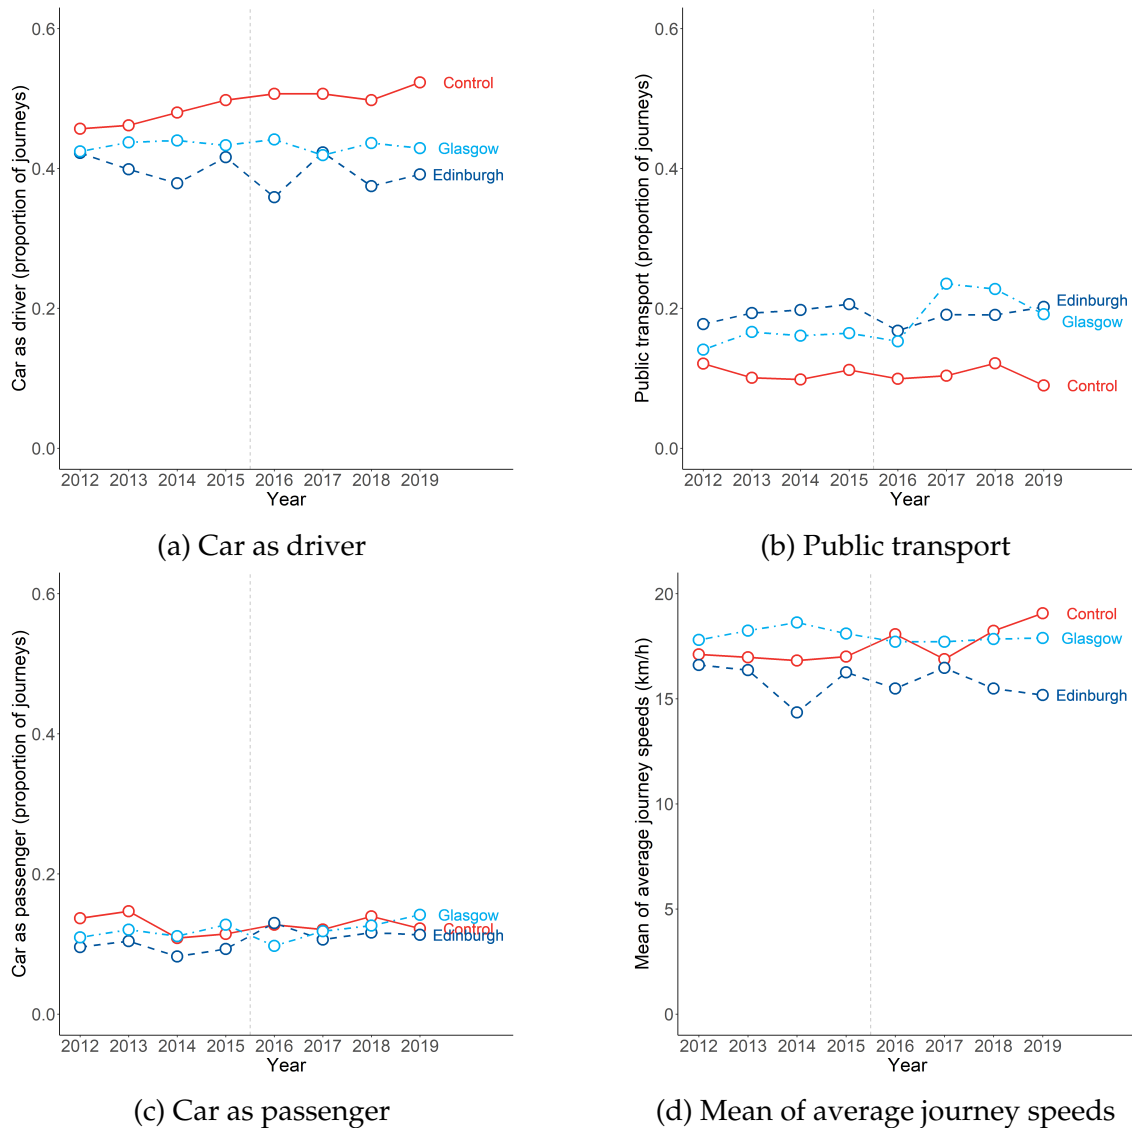


Figure 3.5: Journey main mode choice and mean of calculated average journey speeds, 2012-2019. Control group comprised of Dundee and Aberdeen. Ride hailing became available in Glasgow and Edinburgh between 2015 and 2016 surveys (grey dashed line). Sources: Author’s analysis; SHS 2020.

dynamics across these outcomes was a dramatic increase in the proportion of journeys where public transport was the main mode in Glasgow between 2016 and 2017, relative to a much smaller increase in the control group. This may indicate a treatment effect from the introduction of ride hailing that was delayed by around a year, as Uber launched in Glasgow in late 2015. Figure 3.5 also suggests a decrease in average journey speeds (and, by extension, an increase in

congestion) in Glasgow and Edinburgh in more recent years that was not reflected in the control group.

These raw differences in mean outcomes were worthy of further examination in a difference-in-differences framework.

3.2.6. Difference-in-differences variables

I assigned each of the 43,169 journeys i that started or ended in Glasgow, Edinburgh, Dundee or Aberdeen to a single city c . First, if the journey started in one of the cities, I assigned the journey to that city. For example, journeys from Glasgow to Edinburgh or from Glasgow to the local authority area of East Dunbartonshire were assigned to Glasgow. Second, for the remaining journeys that started from a non-city local authority area but ended in one of the cities, I assigned the journey to that city. For example, a journey from the local authority area of Fife to Edinburgh was assigned to Edinburgh. Figure 3.6 illustrates the percentage of journeys assigned to each city following this approach.




Figure 3.6: Percentage of journeys by assigned city, before and after availability of ride hailing. Sources: Author's analysis; SHS 2020.

Once journeys were assigned to cities, the next task involved the generation of a

treatment variable. Uber entered Glasgow and Edinburgh to provide a ride hailing platform in during October and November 2015, after the fieldwork for the 2015 cross section of the SHS was completed.⁴ As of the end of 2019, a ride hailing platform had not entered either Dundee or Aberdeen. Therefore, I considered journeys assigned to Edinburgh or Glasgow to represent my treatment groups, and my control group to consist of journeys assigned to Dundee or Aberdeen.

To allow for the estimation of treatment effects for each Glasgow and Edinburgh, I generated a categorical $treated_{c,t}$ variable. This variable was equal to 1 if the journey year t was 2016 or later and the assigned city c was Glasgow, equal to 2 if the year t was 2016 or later and the city c was Edinburgh, and equal to 0 otherwise.

3.3. Results

3.3.1. Main transport mode

The first column of Table 3.4 presents summary results for my difference-in-differences multinomial logistic regression of the journey transport mode (Equation 3.6), with results reported as exponentiated regression coefficients, or relative risk ratios. Relative to driving a car (the reference category for my outcome variable), ride hailing increased the probability of public transport being used as the main transport mode in Glasgow by almost 75 per cent. This result was not reflected in Edinburgh, where the effect on public transport was positive but not statistically significant.

Taking a car as a passenger (including taxi journeys) became more likely than driving a car in Edinburgh as a result to ride hailing, although this was not reflected in Glasgow, and there was some evidence that walking became more

⁴Another ride hailing platform, MyTaxi (later Free Now), launched in Edinburgh in May 2018, by which time ride hailing was already available through Uber.

likely relative to driving a car in both cities. In Glasgow, the ‘other’ category was much less likely to be chosen relative to driving a car as a result of ride hailing. This result is difficult to interpret given the miscellaneous nature of the category, which in any case only accounts for a small proportion of total journeys.

Table 3.4: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Under 45	(3) Male	(4) Employed	(5) High income	(6) Degree
Public						
Treated Glasgow	1.749*** (0.242)	1.898*** (0.398)	2.315*** (0.486)	2.074*** (0.392)	2.987*** (0.960)	1.807** (0.465)
Treated Edinburgh	1.138 (0.166)	1.014 (0.221)	1.286 (0.277)	1.339 (0.258)	1.391 (0.447)	1.131 (0.287)
Passenger						
Treated Glasgow	1.211 (0.166)	1.032 (0.209)	0.942 (0.214)	1.208 (0.226)	1.382 (0.381)	0.980 (0.221)
Treated Edinburgh	1.450** (0.220)	1.162 (0.264)	1.515 (0.389)	1.280 (0.266)	1.589 (0.480)	1.271 (0.316)
Walk						
Treated Glasgow	1.381*** (0.167)	1.312 (0.223)	1.237 (0.214)	1.335* (0.214)	1.713** (0.453)	1.074 (0.206)
Treated Edinburgh	1.385** (0.177)	1.194 (0.214)	1.330 (0.243)	1.276 (0.203)	1.379 (0.361)	1.288 (0.237)
Other						
Treated Glasgow	0.424*** (0.134)	0.610 (0.233)	0.751 (0.282)	0.415** (0.160)	0.677 (0.505)	0.297*** (0.135)
Treated Edinburgh	0.889 (0.287)	1.049 (0.414)	1.259 (0.486)	0.778 (0.305)	1.341 (0.999)	0.857 (0.393)
Observations	43169	20290	19978	25805	10333	17758
Pseudo R^2	0.091	0.089	0.100	0.067	0.071	0.071
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author’s analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

I ran this same regression (Equation 3.6) on various sub-samples of the data

3. ALL HAIL?

to assess whether there was any heterogeneity in effects by socio-demographic characteristics, and columns 2 to 6 of Table 3.4 provides results of this sub-sample analysis. This shows that the effect on public transport was more pronounced among younger respondents under the age of 45 (column 2) and male respondents (column 3). The effect was also stronger among respondents who were employed (column 4) and whose household held at least a degree (column 6). The result appeared to be strongest among respondents with a total annual net household income of greater than £30,000, with public transport more likely by almost 200 per cent because of ride hailing (column 5). The consistency of this result across all sub-samples represented evidence in favour of Hypothesis 2 (and against Hypothesis 3).

The result in relation to taking a car as a passenger for Edinburgh in my full-sample regression (column 1) was less clear as it was not reflected in most sub-samples. Therefore, there was insufficient evidence to reject the null in the case of Hypothesis 1. The results suggesting that walking became more likely, and 'other' becoming less likely, relative to driving a car in both cities were generally not reflected in my sub-sample analysis either.

I also conducted a sub-sample analysis to explore heterogeneity in results by journey purpose. Column 1 of Table 3.5 repeats the results of my full-sample regression, and columns 2 and 3 show results for sub-samples of work-related and leisure journeys respectively. The effect on public transport relative to driving a car in Glasgow was considerably more pronounced among work-related journeys, with the relative probability increased by 177 per cent, but statistically insignificant among leisure journeys. In addition, among work-related journeys, I detected a positive effect on the probability of using public transport in Edinburgh, with an increased relative probability of just under 91 per cent.

Table 3.5: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Work	(3) Leisure
Public			
Treated Glasgow	1.749*** (0.242)	2.770*** (0.656)	1.387 (0.395)
Treated Edinburgh	1.138 (0.166)	1.907*** (0.462)	0.775 (0.224)
Passenger			
Treated Glasgow	1.211 (0.166)	0.969 (0.323)	1.568* (0.366)
Treated Edinburgh	1.450** (0.220)	0.991 (0.384)	1.792** (0.457)
Walk			
Treated Glasgow	1.381*** (0.167)	1.124 (0.277)	1.650** (0.386)
Treated Edinburgh	1.385** (0.177)	1.285 (0.304)	1.123 (0.270)
Other			
Treated Glasgow	0.424*** (0.134)	0.456* (0.211)	0.265** (0.143)
Treated Edinburgh	0.889 (0.287)	1.007 (0.474)	0.397* (0.222)
Observations	43169	11224	8526
Pseudo R^2	0.091	0.094	0.098
Year fixed effects	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

3.3.2. Average journey speed

Column 3 of Table 3.6 shows summary results for my difference-in-differences linear regression of average journey speed (Equation 3.9) using a sub-sample of road-based journeys, in addition to results for linear regressions of journey duration (column 1) and journey distance (column 2). I found some evidence

3. ALL HAIL?

of a negative effect on average journey speed, and by extension an increase in congestion, in Glasgow. Over the 2012-2015 period, before ride hailing became available, the mean of average journeys speeds was 18.19 and 15.87 kilometres per hour in Glasgow and Edinburgh respectively. Table 3.6 suggests that the introduction of ride hailing led to a reduction in average journey speed in Glasgow of 1.3 kilometres per hour, or 7.15 per cent, and represents some evidence in favour of Hypothesis 4. Again, however, this effect was not reflected in Edinburgh. I found no evidence of an effect on journey distance, but some evidence journey duration was longer by 1.23 minutes in Edinburgh due the introduction of ride hailing.

Table 3.6: Regression difference-in-differences estimates of ride hailing effect on journey duration, distance and speed

	(1) Duration	(2) Distance	(3) Speed
Treated Glasgow	0.142 (0.701)	-0.629 (0.505)	-1.296* (0.768)
Treated Edinburgh	1.231* (0.744)	0.227 (0.520)	-0.096 (0.792)
Observations	29732	29732	29732
Adjusted R^2	0.023	0.031	0.042
Year fixed effects	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at individual level

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

3.3.3. Robustness

As shown in Appendix C.2, I determined that the ride hailing effect on public transport in Glasgow persisted when limiting the control group solely to journeys from Dundee, and solely to Aberdeen journeys. I also confirmed using a placebo test that no 'effect' could be detected among local authorities where ride hailing did

not become available (see Appendix C.2). I also provide results of my alternative regression specification that included city-specific linear trends (Equation 3.10) in Appendix C.2. This regression tested whether my results in relation to public transport becoming more likely relative to driving a car were robust to the inclusion of city-specific linear trends. As expected, standard errors were increased due to the more restrictive specification. The public transport result did not hold in this specification, and was thus not robust to the inclusion of linear trends.

However, Figure 3.7 illustrates the exponentiated regression coefficients for the effect of ride hailing on public transport, relative to driving a car, in each Glasgow and Edinburgh from my more generalised model that included lags and leads of the treatment (Equation 3.11). This shows that in Glasgow, no treatment was detected in the years before ride hailing became available, while an effect was evident in each 2017, 2018 and 2019. In addition to suggesting that the effect on public transport in Glasgow was in fact robust and that the null hypothesis could be rejected in the case of Hypothesis 2, this provides reassurance that cause occurred before consequence rather than the other way around. It also suggests that it may have taken some time from the launch of Uber in late 2015 to produce an effect on other transport modes from early 2017 onwards, rather than there being an immediate effect.

In the case of Edinburgh, Figure 3.7 shows that there was no evidence of an effect on public transport, which is consistent with the weaker public transport results for Edinburgh in other regression specifications.

My results for average journey speed were not reflected in the generalised model of treatment lags and leads, as shown in Figure 3.8. While there was no evidence of anticipatory effects, I also found little evidence of post-treatment effects, indicating that the results in Table 3.6 cannot be considered robust. Based on this, there was

3. ALL HAIL?




Figure 3.7: Regression difference-in-differences estimates coefficients and 95 percent confidence intervals for effect of ride hailing effect on public transport, generalised model 2012-2019. Exponentiated coefficients. Robust standard errors clustered at individual level and transformed using delta method. Uber operating in Glasgow and Edinburgh from 2016 onwards (red dashed line). Sources: Author's analysis; SHS 2020.

insufficient evidence to reject the null for Hypothesis 4.

3.3.4. Mechanism

To dig deeper into this finding of a positive effect on the use of public transport, I ran my main regression specification (Equation 3.6) again with the mode choice outcome variable amended to split the public transport category into separate bus and rail categories. This tested whether the effect on public transport mainly affected bus journeys or rail journeys. Results for this regression, provided in Appendix C.2, indicated that the public transport effect stemmed from rail journeys, with significant positive effects found on the use of rail in both Glasgow and Edinburgh. Meanwhile, no effect was found among bus journeys in either city.

How could ride hailing complement public transport? I formed Hypothesis 2




Figure 3.8: Regression difference-in-differences estimates coefficients and 95 per cent confidence intervals for effect of ride hailing effect on average journey speed of road-based journeys, generalised model 2012-2019. Robust standard errors clustered at individual level. Uber operating in Glasgow and Edinburgh from 2016 onwards (red dashed line). Sources: Author's analysis; SHS 2020.

on the basis that ride hailing could, by acting as a secondary mode of transport for a journey where the main mode was public transport, reduce the time spent travelling to or from the nearest train station or bus stop (Stiglic et al. 2018). Figure 3.9 thus compares the percentage of all public transport journeys that also involved the use of a car as a passenger, during another stage of the same journey, before and after the launch of ride hailing in Scotland in 2015. While this percentage was very small in both periods, it was perceptibly higher after 2015 at 1.4 per cent compared with 0.1 per cent before.⁵ To test this difference, I ran a logistic regression of this percentage on a dummy variable that was equal to 1 if the journey occurred in 2016 or later, and 0 otherwise. I found a significant increase in the probability that a public transport journey also involved the use of a car as a passenger in the period after the launch of ride hailing (see Appendix C.2). This result cannot be interpreted as a causal effect as the number of journeys

⁵Specifically among rail journeys, this percentage increased from 0 to 3.9 per cent.

3. ALL HAIL?

involving both public transport and a car as a passenger was too small to permit a difference-in-differences analysis, but it nonetheless provides some descriptive evidence on how ride hailing may have affected public transport.




Figure 3.9: Use of car as passenger as journey stage among public transport journeys, before and after availability of ride hailing. Car as passenger includes taxi. Sources: Author’s analysis; SHS [2020](#).

3.4. Discussion

Ride hailing has been increasing in popularity over the past decade (see Figure 3.1), and yet a consensus in the literature as to its impact on the use of other transport modes remains elusive. Does ride hailing complement or substitute other forms of transport, such as public transport? This question is key to whether ride hailing should be viewed as helpful in transitioning to sustainable transport. This paper contributes to this debate with an empirical analysis of the effect of the entry of Uber into the cities of Glasgow and Edinburgh in Scotland during 2015, using a difference-in-differences methodology and travel diary data from repeated cross sections of the SHS ([2020](#)) from 2012 to 2019.

The results of my difference-in-differences analysis revealed that the availability of ride hailing increased the use of public transport relative to driving a car in Glasgow. The magnitude of this effect was considerable, increasing the probability of using public transport by approximately 75 per cent relative to driving a car. There was insufficient evidence of this effect being reflected in Edinburgh, however.

I determined that the public transport result for Glasgow was not robust to the inclusion of city-specific linear trends in the regression model. This may have been a reflection of the issue discussed by Wolfers (2006), that controlling for linear trends may actually contaminate results in a setting where the treatment effect might be dynamic rather than immediate and constant. It is possible that as more drivers registered with the ride hailing platform and more passengers became aware of its availability, its effect on other transport modes increased. In a more generalised model that included treatment leads and lags to test for anticipatory and post-treatment effects, I found that a positive post-treatment effect on public transport was evident in Glasgow for each year other than the first year in which ride hailing was available, 2016. This result suggested that there was indeed a dynamic effect that took time to emerge rather than being immediate. This may explain why the result did not appear in my linear trends specification, as the linear trend could have soaked up some of this dynamic effect. I also found no evidence of anticipatory effects in Glasgow to provide reassurance that consequence occurred after cause rather than vice versa. In the case of Edinburgh, I found no evidence of anticipatory or post-treatment effects.

The public transport result for Glasgow represents evidence in favour of the argument that ride hailing has a complementary effect on public transport, in contrast with the finding of the Tirachini (2019) review that found a substitution effect to be more common among studies. Many of these studies on the effects of

3. ALL HAIL?

ride hailing on other transport modes have relied on cross-sectional data (Tirachini 2019), however, and thus fell short of establishing causal relationships.

I found that the effect on public transport was more pronounced among respondents who were younger, male, employed and were members of a household that held at least a degree. The effect was particularly pronounced among respondents with higher levels of household income, with the probability of using public transport increased almost by almost 200 per cent among this sub-sample. These results were in line with the finding of the Tirachini (2019) review that ride hailing users tended to be younger, more highly educated and wealthier. They were also consistent with evidence provided by Sikder (2019) that individuals working full-time with flexible hours were more likely than non-workers to utilise ride hailing.

I then ascertained that the positive effect on public transport stemmed from an effect on rail transport, with no effect on bus transport found. This effect on rail was evident in both Glasgow and Edinburgh. This was consistent with a difference-in-differences study in China that found that ride hailing increased the number of rail passengers, but did not concur with its finding that ride hailing reduced the number of bus passengers (Shi, Li, and Xia 2021). The mixed results in the literature on the effect of ride hailing on public transport, as highlighted by Tirachini (2019), likely reflect different study settings, contexts and empirical methods. However, my finding that there was an effect on rail but not on bus transport, in addition to the findings of Shi, Li, and Xia (2021), indicate that public transport should not be treated solely as a homogeneous category in this literature.

There are several possible explanations for why ride hailing may have a complementary effect on public transport. While Tirachini (2019) found that ride hailing was predominantly used for occasional leisure trips, my results showed

that the effect on public transport was stronger among work-related journeys, including either commutes or journeys undertaken in the course of work. Among these journeys, a positive effect of ride hailing on the use of public transport was found in both cities, rather than in Glasgow alone. Meanwhile, the result was not reflected in either city among journeys taken for leisure purposes in either city. These findings suggest that any mechanism for a positive effect on public transport resides in work-related rather than leisure journeys. This represents evidence in favour of the hypothesis discussed by Stiglic et al. (2018) that ride hailing has a complementary effect on public transport by helping to overcome the 'last mile' problem of connecting the home or workplace of individuals to the public transport network. Individuals may be utilising ride hailing to transport themselves between their home or workplace and the nearest train station, for example.

Based on this, I also showed evidence that the proportion of public transport journeys that also involved the use of a car as a passenger in a separate stage of the journey was higher after the introduction of ride hailing in Glasgow and Edinburgh. While there were not enough such journeys in my control group of Dundee and Aberdeen to establish this as a causal relationship, this increase at least offers some additional descriptive evidence to support the last mile theory. The increase in journeys combining public transport and car as passenger was significant in a logistic regression, but it should be noted that the increase was from a very low base, with 1.4 per cent of public transport journeys after the introduction of ride hailing, compared with 0.1 per cent before, also involving car as passenger.

What could explain the apparent differences in treatment effects between Glasgow and Edinburgh? In most specifications, while the estimated treatment effect for Edinburgh was positive, it could not be deemed statistically significant. While

3. ALL HAIL?

both cities are larger population centres than Dundee or Aberdeen (see Figure 3.2), population density and the density of transport infrastructure are higher in Glasgow, with 3,374 and 1,768 persons per square kilometre living in Glasgow and Edinburgh respectively in 2012 (Office for National Statistics 2020, see Appendix C.3 for further details). In addition, prior to ride hailing becoming available, the mean distance of Glasgow journeys was 7.41 kilometres compared with 6.77 kilometres for Edinburgh journeys among my SHS (2020) journeys. I showed in Figure 3.4 that while around 20 per cent of journeys in both Glasgow and Edinburgh were undertaken using public transport, driving a car was more popular in Glasgow than in Edinburgh, while walking was relatively more popular in Edinburgh. These figures suggest that Edinburgh is a smaller city that is more conducive to walking than Glasgow. The difference in treatment effects may indicate that there is more scope for ride hailing to complement public transport in a larger city setting where many journeys are undertaken using a car and where there is a relatively high density of transport infrastructure.

I also found some evidence of an effect on average journey speed among road-based journeys, which can be used as a measure of traffic congestion, but this effect did not prove to be sufficiently robust to be considered causal. Previously, Fageda (2021) found that the presence of ride hailing reduced average congestion in European cities, while Agarwal, Mani, and Telang (2019) found ride hailing increased congestion in Indian cities.

3.4.1. Policy implications

Overall, the results of my empirical analysis showed that ride hailing had a complementary effect on public transport in Glasgow, and that this specifically affected rail transport. However, while significant, the proportion of total journeys affected by ride hailing appeared to be very small and, therefore, the effect of ride

hailing on the overall transport system should not be overstated.

These results indicate that ride hailing can contribute to the move towards sustainable transport. On the basis that ride hailing can help overcome the last-mile problem, the complementarity with public transport could be strengthened by facilitating ride hailing at public transport stations and hubs, for example by ensuring the reliable provision of internet access or facilitating the use of dedicated collection or drop-off spaces alongside registered taxis.

3.4.2. Limitations and strengths

The findings of this study should be interpreted in the context of certain limitations. First, I had no data on the actual use of ride hailing in Scotland over the study period. My empirical analysis could be improved by a measure of the use of ride hailing in each city, as this would allow an examination of effects by the intensity of ride hailing use. This could also further allay concerns over a possible omitted variable bias that cannot be ruled out as my identification strategy was based purely on temporal variation. Second, although the multinomial logistic regression was a more appropriate model choice than a binary logistic regression, it relied on the assumption of independence of irrelevant alternatives. Richer, more granular data could facilitate the use of a nested logit model which would improve on this approach. Third, it should be highlighted that the use of data from the Scottish Household Survey focused the study on residents of Scotland and their travel behaviour. It is possible that the impact of ride hailing on the travel behaviour of non-residents, such as tourists or people travelling to Scotland in the course of work, is different and this presents an interesting avenue for future research. Fourth, this study centred on Glasgow and Edinburgh, two cities in a developed country. Further research would be required to determine if ride hailing can also have a complementary effect on public transport in a developing

3. ALL HAIL?

city.

This study can boast several key strengths. First, I employed a difference-in-differences methodology, which allowed me to identify a causal relationship rather than an association between ride hailing and public transport. This approach improved upon many of the studies in this literature that have used cross-sectional data. Second, I drew on travel diary data from a large, representative survey of households in Scotland, the SHS (2020). Several studies in this literature have relied on cross-sectional stated preference surveys of ride hailing users, whereas this travel diary data provided a detailed picture of the journeys made by a representative sample of individuals for each year between 2012 and 2019. This granular data facilitated the difference-in-differences identification strategy using multinomial logistic regressions with a large sample size before and after the introduction of ride hailing, covering cities with and without access to ride hailing. The travel diary data, while still self-reported, provided revealed preference rather than stated preference information, with respondents recording their actual mode choice rather than stating what they would have chosen in a hypothetical scenario. Third, in addition to providing information on mode choice, this travel diary data recorded the purpose of each journey, allowing me to assess heterogeneity in results between work-related and leisure journeys. Furthermore, the travel diary data was also linked to information on the socio-demographic characteristics of survey respondents and their households. This meant that I could control for these characteristics and also analyse heterogeneity in results by demographic and socio-economic status, and thus provide a clearer and more detailed account of the impact of ride hailing on other transport modes.

CONCLUSIONS

In this thesis, I set out to enhance our understanding of several questions that are relevant in plotting a course towards sustainable transport. Have we reached 'peak car'? Did Brexit inadvertently cause a shift from road freight transport to the more energy efficient mode of short sea shipping? Can ride hailing platforms advance the sustainability of transport?

In Chapter 1, I demonstrated a hump-shaped relationship between the number of licensed cars required in the economy to produce one unit of GDP, or car intensity, and GDP per capita among a panel of 88 countries from 1950 to 2010. In Great Britain, for example, car intensity increased sharply before reaching a peak in 1992 and has been gradually decreasing ever since. This phenomenon is relevant to the peak car debate regarding the apparent levelling off in car use per-capita among developed countries.

My core argument in this chapter was that structural transformation, the reallocation of economic activity between the broad agriculture, manufacturing and services sectors, can have an influence in producing this relationship between car intensity and GDP per capita. A budding economy that is based on a labour-intensive agriculture sector may industrialise and realign towards a relatively more car-intensive non-agriculture sector as it grows, inducing a rise in car intensity. As the economy blossoms and becomes centred on non-agriculture, however, the non-agriculture sector itself may become more capable of reducing its car intensity,

allowing the mature economy to experience a decrease in aggregate car intensity.

My calibrated model that formalised this theory was successful in reflecting the hump shape in car intensity and was able to account for almost a quarter of observed variation in car intensity among a cross section of 54 countries in 2010. In counterfactual exercises using this model, I found economies that develop later reach a peak level of car intensity at a later stage, but that this peak occurs at a lower level than that experienced by countries that shifted away from agriculture earlier. With intensity peaking at a lower level and at a stage when the fuel efficiency of vehicles had improved, cumulative carbon emissions over a 1970-2120 simulation period were 37 per cent lower in my later-developer counterfactual than in my baseline model.

The findings in Chapter 1 have implications for public policy. First, we can expect car intensity to continue its descent among developed countries. For example, my model predicted that car intensity will be around 35 per cent and 15 per cent lower in France and South Korea respectively in 2040 than in 2010. A falling level of car intensity, with fewer cars required to produce economic output, represents evidence that developed countries have indeed entered peak car. Second, as developing economies structurally change, we can expect their levels of car intensity to increase. However, this increase is unlikely to be permanent as intensity in these countries should eventually reach a peak before later decreasing.

I switched from a macroeconomic to a microeconometric angle and turned away from passenger transport to focus on the transport of freight in Chapter 2. Using a difference-in-differences analysis of EU-27 port-level cargo volumes between 2013 and 2022, I identified a 22 per cent decline in EU-UK roll-on roll-off (Ro-Ro) cargo volumes that could be causally attributed to Brexit. Homing in on cargo volumes in Irish ports, my results revealed that Ireland-UK Ro-Ro cargo was most affected

by Brexit, with a 54 per cent decrease in Ro-Ro cargo volumes. Interestingly, I showed that this occurred alongside a 147 per cent increase in Ro-Ro cargo volumes between Ireland and France, albeit from a much lower base, because of Brexit.

These results pointed to a diversion of cargo being transported between Ireland and mainland Europe from the UK land-bridge route to direct short sea shipping routes. This shift would have had substantial implications for the energy consumption and emissions associated with transporting these goods. Based on several simplifying assumptions, I calculated that the energy consumption and emissions of transporting cargo on the direct short sea shipping route would have been approximately 60 per cent lower than on the land-bridge route.

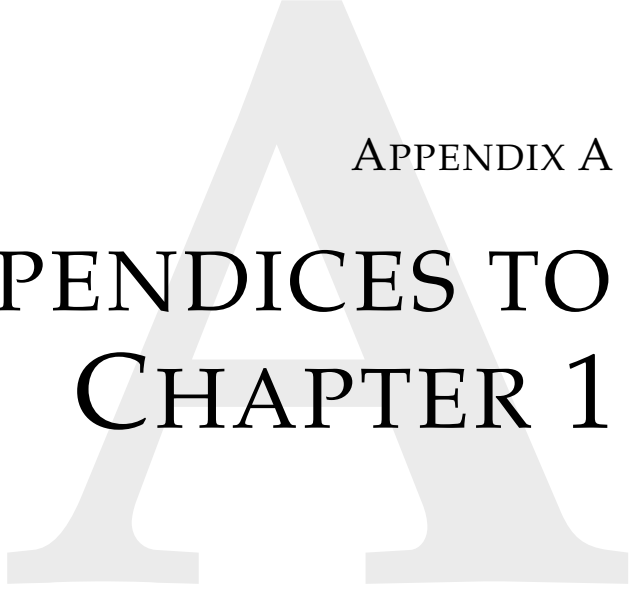
Of course, these results are highly specific to the geographic context of Ireland. However, they contain a more general lesson for policy: the appeal of short sea shipping can be advanced by reducing its transit time relative to land-based alternatives. For moving cargo between Ireland and continental Europe, the monetary cost of the direct short sea shipping route was always lower than the land-bridge route. However, additional customs checks at ports due to Brexit increased the transit time of the land-bridge route, thus making the direct route relatively more attractive. As policy makers such as the European Commission endeavour to encourage the use of short sea shipping over land-based alternatives as part of making transport more sustainable, this unintended consequence of Brexit highlights transit times as a key indicator to focus on. For example, investment in port infrastructure could be a channel through which short sea shipping transit times can be targeted.

While I continued to inspect through the lens of microeconometrics, I returned to passenger transport in Chapter 3 to focus on the question of whether ride hailing

platforms substitute or complement public transport. My empirical analysis, which scrutinised travel diary data from Scotland between 2012 and 2019 in a difference-in-differences framework, revealed that the availability of ride hailing increased the use of public transport relative to driving a car by around 75 per cent in Glasgow, although this effect was not reflected in Edinburgh.

I showed descriptive evidence suggesting that this complementary effect may have stemmed from ride hailing helping to mitigate public transport's 'last mile' problem. For example, individuals may utilise ride hailing for getting to the nearest train station. However, the proportion of total journeys affected by ride hailing appeared to be small and, therefore, the effect of ride hailing on the overall transport system should not be overstated. I also found insufficient evidence of any effect on congestion levels in Glasgow or Edinburgh. This analysis suggested that while it has not been a game-changer in promoting sustainable transport in Scotland, the advent of ride hailing need not be viewed as a regressive step on the journey to sustainability.

If the aim of limiting the global temperature rise to 1.5 degrees Celsius is to be met, transport must reduce its emissions and become more sustainable. This thesis has shed light on a range of topics in this regard, including peak car, short sea shipping and ride hailing. In doing so, it has demonstrated how methods and perspectives from the fields of macroeconomics and microeconometrics can be deployed in building a foundation of evidence that can support the passage to a more sustainable transport sector.



APPENDIX A

APPENDICES TO CHAPTER 1

A.1. Data

A.1.1. Main country panel: 54 countries

Using a range of different sources, I constructed a panel dataset for 54 countries covering 22 years from 1995 to 2016 that could be employed to calibrate the structural transformation model. I also used this data to illustrate patterns in PKM intensity, and in conducting an empirical analysis of car intensity using regression models.

Based on data availability, the following 54 countries were included in my main country panel: Argentina, Australia, Austria, Belgium, Brazil, Bulgaria, Canada, Chile, China, China (Hong Kong SAR), Colombia, Croatia, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, India, Indonesia, Ireland, Israel, Italy, Kazakhstan, Latvia, Lithuania, Malaysia, Mexico, Morocco, Netherlands, New Zealand, Norway, Peru, Philippines, Poland, Portugal, Republic of Korea, Romania, Russian Federation, Saudi Arabia, Singapore, Slovakia, Slovenia, South Africa, Spain, Sweden, Switzerland, Thailand, Tunisia, Turkey,

United Kingdom, United States. This includes 34 OECD members, 5 countries designated by the OECD as ‘key partners’, and 15 non-members. The agriculture labour shares and values of real GDP per capita of the 54 included countries in 2016 are illustrated in Figure 1.8a.

This section provides details on the various data sources collated in the panel. As shown in Table A.1, I considered ‘agriculture’ to correspond to section A in the International Standard Industrial Classification of All Economic Activities (ISIC) Revision 4 (UN 2008) and ‘non-agriculture’ to correspond to the aggregation of sections B-G and I-U, with section H being considered ‘transport’ as a proxy for private motor vehicles where more disaggregated data was not available.

Employment

I relied on data from the International Labour Organization (ILO) for levels of employment by economic activity (ILO 2020). I aggregated the sectoral data up to three broad sectors: agriculture, non-agriculture and transport. I then calculated sectoral labour shares as proportions of total employment. Without more disaggregated data, I used the transport labour share to correspond with the motor vehicle sector in my model. When aggregated across the 54 countries in my collated main panel, the transport labour share was just over 4 per cent in 1995 and had increased to just under 6 per cent by 2016. As the primary labour dynamic in my model is a large-scale shift from agriculture to non-agriculture, I contend that using this transport labour share as a proxy for employment in my motor vehicle sector is a reasonable approach given data limitations.

Value added and prices

For sectoral gross value added (GVA) from 1995 to 2016, both in constant and current prices, I sourced data from the United Nations (UN) National Accounts

Table A.1: Mapping International Standard Industrial Classification of All Economic Activities (Revision 4) sections to broad sectors

Agriculture	Non-agriculture	Transport
A: Agriculture, forestry and fishing	B: Mining and quarrying C: Manufacturing D: Electricity, gas, steam and air conditioning supply E: Water supply; sewerage, waste management and remediation activities F: Construction G: Wholesale and retail trade; repair of motor vehicles and motorcycles I: Accommodation and food service activities J: Information and communication K: Financial and insurance activities L: Real estate activities M: Professional, scientific and technical activities N: Administrative and support service activities O: Public administration and defence; compulsory social security P: Education Q: Human health and social work activities R: Arts, entertainment and recreation S: Other service activities T: Activities of households as employers; undifferentiated goods- and services-producing activities of households for own use U: Activities of extraterritorial organisations and bodies	H: Transportation and storage

Sources: Author's analysis; UN [2008](#)

Main Aggregates Database (UN 2021a). This database provided sectoral GVA in constant 2015 US dollars (USD), and I re-based this to 2005 prices. I aggregated the GVA data up to the broad agriculture, non-agriculture and transport sectors, and again used the transport sector data for my model's motor vehicle sector.

Price indices for the agriculture, non-agriculture and transport sectors could then be calculated as the ratio of sectoral GVA in current prices to sectoral GVA in constant prices:

$$\frac{p_{s,t}}{p_{s,2005}} = \frac{GVA_{s,t}}{GVA_{s,2005}}, \quad s = A, N, T \quad (\text{A.1})$$

The resulting indices were relative to the base year of the constant price GVA data, 2005 in the case of my collated dataset.

When comparing sectoral productivity across countries, however, it is important to account for any sectoral price differences between countries. For example, the relative price of the agriculture good may be different in Germany compared to India. Simply converting GVA data from local currencies to USD using market exchange rates ignores these differences in relative prices between economies.

To account for this, I calculated sectoral price levels using data from the International Comparison Programme (ICP), a database that provides values of final expenditure by category in 2005 (The World Bank 2005). This data is available as 'nominal' expenditure in current USD, converted from local currencies using market exchange rates, as well as 'real' expenditure in purchasing power parity (PPP) USD (The World Bank 2008), which allowed me to extract the relative price levels of sectors relative to the corresponding sector in the US in 2005. Specifically, I calculated the relative price level of sector s in country i as follows:

$$\frac{p_s^i}{p_s^{PPP}} = \frac{E_s^i}{E_s^{PPP}} \quad (\text{A.2})$$

In Equation A.2, E_s^i and E_s^{PPP} denote expenditure in country i and sector s in current USD and in PPP USD respectively.

The publicly available ICP data (The World Bank 2005) is disaggregated across several expenditure categories that can be mapped to the broad agriculture, non-agriculture and transport sectors in my model. Table A.2 displays how I mapped ICP categories to these sectors. I aggregated expenditure across all corresponding categories for each broad sector before extracting relative price levels.

Table A.2: Mapping International Comparison Programme 2005 categories to broad sectors

Agriculture	Non-agriculture	Transport	Not mapped
110100: Food and non-alcoholic beverages	110300: Clothing and footwear	110700: Transport	111300: Balance of expenditures of residents abroad and expenditures of non-residents in the economic territory
110200: Alcoholic beverages, tobacco and narcotics	110400: Housing, water, electricity, gas, and other fuels*		130000: Individual consumption expenditure by government
	110500: Furnishing, household equipment and routine maintenance of the house		140000: Collective consumption expenditure by government
	110600: Health - HHC*		170000: Balance of exports and imports
	110800: Communication		160100: Changes in inventories
	111100: Restaurants and hotels		160200: Acquisitions less disposals of valuables
	150100: Machinery and equipment		
	150200: Construction		
	150300: Other products		

* Expenditure data in USD not available in 2005.

Sources: Author's analysis; The World Bank 2005

Having extracted the 2005 relative price levels of sectors across countries as in Equation A.2, I converted the sectoral GVA data in constant 2005 USD (having re-based this data from 2015 prices) to 2005 PPP USD by dividing it by the corresponding relative price level.

Finally, I also used these 2005 relative price levels to adjust the 1995-2016 sectoral price indices that I calculated using UN (2021a) data. Specifically, I multiplied the price indices from Equation A.1 by the relative price levels in 2005 from Equation A.2 to get 1995-2016 indices of prices levels by sector across countries.


Sectoral motor vehicle inputs

I sourced data on sector-specific motor vehicle inputs from the OECD Input-Output Tables 2021 (OECD 2021), which provided annual inputs in current USD from 1995 disaggregated across 2-digit ISIC (Rev. 4) divisions. Using these categorisations, I assumed motor vehicle inputs to include inputs to each the broad agriculture and non-agriculture categories specifically from division 29 ('manufacture of motor vehicle, trailers and semi-trailers'). I converted these inputs into shares of total value added (also from the Input-Output Tables in current USD) for each agriculture and non-agriculture, and multiplied these shares by sectoral GVA in 2005 PPP USD using the value added data to convert the input values into constant price values.


Of course, more disaggregated data would be more useful in the context of this study, as inputs from division 29 incorporate the manufacture of commercial motor vehicles and trailers in addition to private cars. However, using data from Mitchell's *International Historical Statistics* (Palgrave Macmillan Ltd 2013b) and from the Penn World Tables (Feenstra, Inklaar, and Timmer 2015), in Figure A.1a I decomposed total motor vehicle intensity in Great Britain over the 1950-2010 period into private car intensity and commercial vehicle intensity. First, this

illustrates that private car intensity in Great Britain since 1950 has been consistently and substantially higher than commercial vehicle intensity. Second, it indicates that there has been relatively little change in the level of commercial vehicle intensity since 1950 in Great Britain compared with the change evident in private car intensity. Third, it demonstrates that, given these two facts, a hump-shaped pattern can also be found when plotting the combined intensity of all motor vehicles against GDP per capita.

I also confirmed in Figure A.1b that the hump-shaped pattern in intensity can be found using aggregate motor vehicle inputs per GDP, calculated using Input-Output data (OECD 2021), even though this data was only available from 1995 rather than from 1950. Based on this, while it is a data limitation of this study, I contend that it was reasonable to use this Input-Output data to calculate a proxy measure of sectoral car intensity in the absence of more disaggregated data.



(a) Motor vehicles per GDP and real GDP per capita, Great Britain 1950-2010. Axes transformed to \log_e scales. Sources: Author's analysis; Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.



(b) Aggregate motor vehicle inputs per GDP and real GDP per capita, main panel 1995-2016. X-axis transformed to \log_e scale. Points show decile averages. Sources: Author's analysis; OECD 2021; Feenstra, Inklaar, and Timmer 2015.

Figure A.1: Using OECD Input-Output Table data on motor vehicle inputs to calculate a proxy measure of sectoral car intensity.

Wage

In calibrating my model of structural transformation, I considered the economy wage (common across sectors) to be total GVA per worker in 2005 PPP USD. This was calculated using the collated employment and value added data.

Passenger kilometres by transport mode

For 29 of the 54 countries in my main panel, annual data on passenger kilometres (PKM) for passenger vehicles, rail and bus from 1970 to 2019 were also available from the International Transport Forum's Transport Statistics database (OECD 2017). The 29 countries were: Argentina, Australia, Austria, Belgium, Bulgaria, Canada, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Republic of Korea, Lithuania, Netherlands, Norway, New Zealand, Poland, Portugal, Russian Federation, Slovakia, Slovenia, Spain, Sweden, Switzerland, United Kingdom, United States. Given the data source, this 29-country subset mainly consisted of OECD members, with Argentina, Bulgaria and Russia also included.

While I used my main panel covering 54 countries over the 1995-2016 period for calibrating my model, I used this 29-country subset from 1970 to 2019 to illustrate the observed hump-shaped pattern in PKM intensity. To calculate PKM intensity over time, I divided PKM by real GDP in constant 2005 PPP USD.

Old-age dependency ratio

For the empirical analysis involving regression models, I also sourced data from 1970 to 2010 on the old-age dependency ratio from The World Bank (2019). The old-age dependency ratio is defined as the number of persons in the population older than 64 for every 100 persons aged between 15 and 64. This ratio is available

as a percentage, and I converted this to a proportion for the purposes on my empirical analysis.

Tertiary education enrolment

As a measure of enrolment in higher and further education, I employed the gross enrolment ratio in tertiary education, available from The World Bank (2022). This variable is defined as the ratio of total enrolment in tertiary education, regardless of age, to the population in the age group that would typically correspond with tertiary education. It is designed as a measure of the capacity of a country's tertiary education system (The World Bank 2022). The ratio is available as a percentage, and I converted this to a proportion for the empirical analysis.

Road gasoline price

I sourced data on road fuel prices from the International Energy Agency's (IEA) Energy Prices and Taxes database (IEA 2022a). Specifically, I employed the country-level household sector end-use real price index (2015 = 100) for unleaded motor gasoline (petrol). This index was calculated from a representative price of a combination of the most-consumed unleaded motor gasoline products (or leaded motor gasoline for earlier periods), with the nominal price index deflated using national Consumer Price Indices (IEA 2022b).

Railway density

Mitchell's *International Historical Statistics* included historical data on the length of railway open by country in kilometres. Data was available from 1970 to 2010 for my main panel, although this time range varied by country. I combined this data with data on land area from The World Bank (2023), which measured the total area of a country, excluding inland water bodies and national claims to coastal waters,

in square kilometres. I divided railway length by land area to derive the density of railways in kilometres of railway per square kilometre of land.

A.1.2. Extended country panel: 88 countries

My 54-country panel also included annual data on GDP, the number of registered motor vehicles, and carbon emissions, which were available for a wider and more diverse panel of 88 countries. I used this extended 88-country panel for identifying patterns in car intensity. This panel included all countries from my main panel, in addition to the following: Albania, Algeria, Azerbaijan, Bahrain, Belarus, Bolivia (Plurinational State of), Bosnia and Herzegovina, Botswana, Congo, Costa Rica, Dominican Republic, Ecuador, Egypt, El Salvador, Eswatini, Georgia, Iran (Islamic Republic of), Iraq, Jamaica, Jordan, Kuwait, Lebanon, Mauritius, North Macedonia, Panama, Paraguay, Republic of Moldova, Serbia, Syrian Arab Republic, Taiwan, Trinidad and Tobago, Ukraine, Uruguay, Venezuela (Bolivarian Republic of). Car intensity levels and values of real GDP per capita for 2010 of the 88 countries included in my extended panel are illustrated in Figure 1.3b.

GDP

For a time series on GDP, I used real output-side GDP at chained PPPs in 2017 USD, sourced from the Penn World Tables (Feenstra, Inklaar, and Timmer 2015). For consistency with the rest of my collated data, I re-based this real GDP data to 2005 PPP USD.

Registered motor vehicles

I drew on Mitchell's *International Historical Statistics* for a time series on the number of private cars (Palgrave Macmillan Ltd 2013b). Data was available from 1950 to 2010, although this time range varied by country. This data included time series

on 'private cars' and 'commercial vehicles', allowing the total number of motor vehicles to be calculated. However, the data was more widely available for private cars, so I used this in calculating car intensity. To calculate car intensity over time in the economy, I simply divided the number of private cars by real GDP in constant 2005 million PPP USD.

Carbon emissions

I obtained detailed estimates of greenhouse gas emissions among my extended panel for the 1971-2019 period from the IEA (2022c). Using this data, I added country-specific time series on total carbon dioxide emissions from fuel combustion, and also more specifically on emissions from the combustion of motor gasoline excluding biofuels for road transport, to my extended panel. The IEA estimated these emissions in kilo-tonnes of carbon dioxide using IEA energy databases and emissions factors. To calculate carbon emissions intensity for total fuel combustion and for road motor gasoline combustion, I divided these estimates by real GDP in constant 2005 million PPP USD.

A.1.3. Great Britain

I also collated data for a single developed economy, namely Great Britain (incorporating England, Scotland and Wales). I used this data to illustrate patterns in VKM and car intensity.

Vehicle kilometres by car

I sourced data on vehicle kilometres (VKM) by vehicle type, including a category for 'cars and taxis', from 1950 to 2019 from the UK Department for Transport (2021a). The UK Department for Transport (2021a) calculate annual traffic estimates for the UK using a combination of manual traffic counts, automatic

traffic counters and data on road lengths. I combined this data on VKM by cars and taxis with Penn World Table data on real GDP (Feenstra, Inklaar, and Timmer 2015) to calculate VKM intensity by car.

Number of licensed cars

Data on the number of licensed vehicles by tax class, including a category for 'private cars', from 1950 to 2019 was also available from the UK Department for Transport (2021b). To calculate car intensity, I combined this data with Penn World Table data on real GDP (Feenstra, Inklaar, and Timmer 2015).

A.1.4. Empirical analysis

Table A.3 presents descriptive statistics for the variables included in my empirical hypothesis tests. Country-year pairs missing observed data for car intensity (the outcome variable in this empirical analysis) were dropped, leaving 1,406 observations across the 54 countries of my main panel. For the remaining observations missing data for tertiary education enrolment, gasoline price or railway density, I reset values to 0 to preserve the observation. I generated a dummy variable corresponding to each of these three variables that was equal to 1 if its value had been missing and reset to 0, and included these dummy variables in any regression specifications that involved these variables. As shown in Table A.3, missing data was a particular issue for my gasoline price variable, with data available for only 790 of 1,406 observations.

Correlation coefficients between structural transformation (calculated as 1 minus the agriculture labour share), the old-age dependency ratio, tertiary education enrolment, the road gasoline price index and railway density are displayed in Table A.4. This shows the highest correlation between variables to be moderate at 0.56, specifically between the dependency ratio and tertiary education enrolment.

Table A.3: Descriptive statistics of selected variables for main panel 1970-2010

	N	Mean	S.D.	Min.	Max.
Car intensity (cars/million USD)	1,406	11.06	5.77	0.16	29.82
Structural transformation (proportion)	1,406	0.86	0.16	0.23	1.00
Old-age dependency ratio (proportion)	1,406	0.17	0.07	0.05	0.31
Tertiary enrolment (proportion)	1,172	0.38	0.22	0.01	1.04
Regular unleaded gasoline price (2015=100)	790	93.64	20.18	53.15	174.06
Railway density (km of rail per sq. km)	1,223	0.04	0.03	0.00	0.15

N denotes observations. S.D. denotes standard deviation

Sources: Author's analysis; ILO [2020](#); The World Bank [2022](#), [2019](#); IEA [2022a](#); The World Bank [2023](#); Palgrave Macmillan Ltd [2013b](#), [2013a](#)

Table A.4: Correlation coefficients of selected variables for main panel 1970-2010

	Transformation	Dependency ratio	Tertiary	Gasoline price
Transformation	1			
Dependency ratio	0.539	1		
Tertiary	0.506	0.562	1	
Gasoline price	-0.143	-0.0556	-0.0840	1

Variable names abbreviated for illustrative purposes

Sources: Author's analysis; ILO [2020](#); The World Bank [2023](#), [2022](#), [2019](#); IEA [2022a](#); Palgrave Macmillan Ltd [2013a](#)

A.2. Disaggregating non-agriculture

The 'non-agriculture' sector in my model consists of two principal components, industry and services. I considered 'industry' to correspond to the aggregation of sections B-E, and 'services' to correspond to the aggregation of section G and sections I-U, in the ISIC Revision 4 (UN [2008](#)), with section H still kept aside as my motor vehicles sector. Figure A.2 illustrates how observed car intensity evolves as the economies transition from agriculture to industry, and later to services. This provides some evidence that the initial increase in car intensity stems from a shift away from agriculture, but that the later decrease may only set in once the services sector assumes a greater role in the economy.

In Section 2 of the paper, I ran the following linear regression separately for the




Figure A.2: Ternary plot of sectoral labour shares and aggregate car intensity, main panel 1970-2019. Triangle sides show percentages of total employment in agriculture (agri.), industry (ind.) and services (ser.) sectors. Points, representing country-year pairs, are coloured by aggregate car intensity in private motor vehicles per million USD of real PPP GDP. Sources: Author’s analysis; ILO 2020; Feenstra, Inklaar, and Timmer 2015; Palgrave Macmillan Ltd 2013b.

agriculture sector and for the ‘non-agriculture’ sector:

$$VehInputs_{i,t} = \alpha + \beta AgriLabour_{i,t} + \varepsilon_{i,t} \quad (A.3)$$

The dependent variable $VehInputs_{i,t}$ measures the share of sectoral value added (with value added measured in constant 2005 PPP USD) in country i in year t that is accounted for by motor vehicle inputs. The independent variable $AgriLabour_{i,t}$ represents the share of total employment in agriculture in country i and year t .

A.2. Disaggregating non-agriculture

This regression essentially relates sectoral motor vehicle intensity to the extent of structural transformation in each country over time, with the coefficient of interest β capturing the linear relationship between sectoral vehicle intensity and structural transformation.

Table A.5 re-produces the results of this regression for agriculture (column 1) and non-agriculture (column 2) shown in the paper. But what if I disaggregated non-agriculture into its two components, industry and services? Columns 3 and 4 in Table A.5 provide regression results for industry and services, showing that the non-agriculture sector appears to more closely resemble the services sector.

Table A.5: Changes in sectoral motor vehicle inputs over process of structural transformation for main panel 1995-2016

	(1) Agriculture	(2) Non-agriculture	(3) Industry	(4) Services
Agriculture labour share	-0.009*** (0.001)	0.004*** (0.001)	-0.000 (0.002)	0.006*** (0.001)
Observations	1188	1188	1188	1188
Adjusted R^2	0.031	0.009	-0.001	0.018


Standard errors in parentheses

Sources: Author's analysis; OECD [2021](#); ILO [2020](#); Feenstra, Inklaar, and Timmer [2015](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Since the agriculture labour share decreases during the process of structural transformation, Figures A.3a and A.3b depict regression lines for each of these specifications with the x-axis reversed and extended over the full domain of the agriculture labour share to illustrate progress in structural transformation. Figure A.3a simply re-produces the 2-sector regression lines shown in the paper, while Figure A.3b presents regression lines for the 3-sector version.

Table A.5 reveals a negative but statistically insignificant coefficient on the agriculture labour share variable in the industry sector's regression, indicating that a higher share of total employment in agriculture was associated with a lower share of motor vehicle transport inputs in the value added of the agriculture



(a) Agriculture and non-agriculture.

(b) Agriculture, industry and services.

Figure A.3: Changes in sectoral motor vehicle intensity over structural transformation as regression lines, main panel 1995-2016. X-axis extended over full domain and reversed for illustration. Sources: Author's analysis; OECD 2021; ILO 2020; Feenstra, Inklaar, and Timmer 2015.

sector. Meanwhile, a positive coefficient was found on the agriculture labour share variable in the services regression. As illustrated in Figure A.3b, this suggests that over the course of structural transformation, less motor vehicle transport was needed to produce one unit of output in services.

These additional results may shed some light on why motor vehicle intensity appears to decrease in non-agriculture during structural transformation. Rather than the industry sector becoming more efficient in its use of motor vehicles, these results suggest that the decrease actually stems from non-agriculture shifting from industry to services, which itself is initially more intensive but is becoming more efficient in its use of motor vehicles. This is in contrast a similar analysis of oil intensity by Stefanski (2014), who found oil intensity to be decreasing in both industry and services while increasing in agriculture among a sample of OECD countries between 1970 and 2000.

A.3. Model theory

I have assumed that my model economy operates in perfect competition. Figure A.4 provides a stylised illustration of how agents interact in the market of this economy.

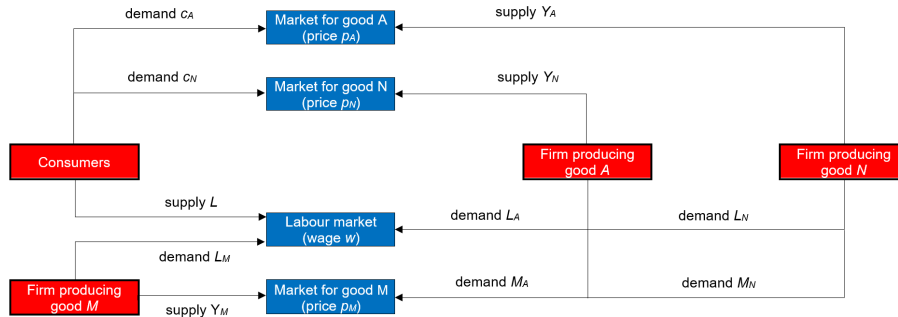


Figure A.4: Market system in my economic model.

A.3.1. Consumers

I assume the consumer optimisation problem takes the form of a log utility function:

$$\begin{aligned} \max_{c_A, c_N} \quad & \phi \log(c_A - \bar{c}_A) + (1 - \phi) \log(c_N) \\ \text{s.t.} \quad & p_A c_A + p_N c_N = w \end{aligned} \tag{A.4}$$

In this constrained maximisation problem, $\phi \in (0, 1)$ represents the utility weight on the agriculture good. A subsistence level of the agriculture good is given by $\bar{c}_A > 0$. Figure A.5 provides a graphical illustration of this problem.

The Lagrangian method can be used to solve the household's problem. The Lagrangian can be specified as follows:

$$\mathcal{L} = \phi \log(c_A - \bar{c}_A) + (1 - \phi) \log(c_N) - \lambda (p_A c_A + p_N c_N - w) \tag{A.5}$$

This yields the following three first-order conditions:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial c_A} &= \frac{\phi}{c_A - \bar{c}_A} - \lambda p_A = 0 \\ \frac{\partial \mathcal{L}}{\partial c_N} &= \frac{1 - \phi}{c_N} - \lambda p_N = 0 \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= p_A c_A + p_N c_N - w = 0\end{aligned}\tag{A.6}$$

Together, these first-order conditions can be solved for c_A and c_N to give demand functions for each consumption good:

$$c_A = \phi \frac{w - p_A \bar{c}_A}{p_A} + \bar{c}_A\tag{A.7}$$

$$c_N = (1 - \phi) \frac{w - p_A \bar{c}_A}{p_N}\tag{A.8}$$

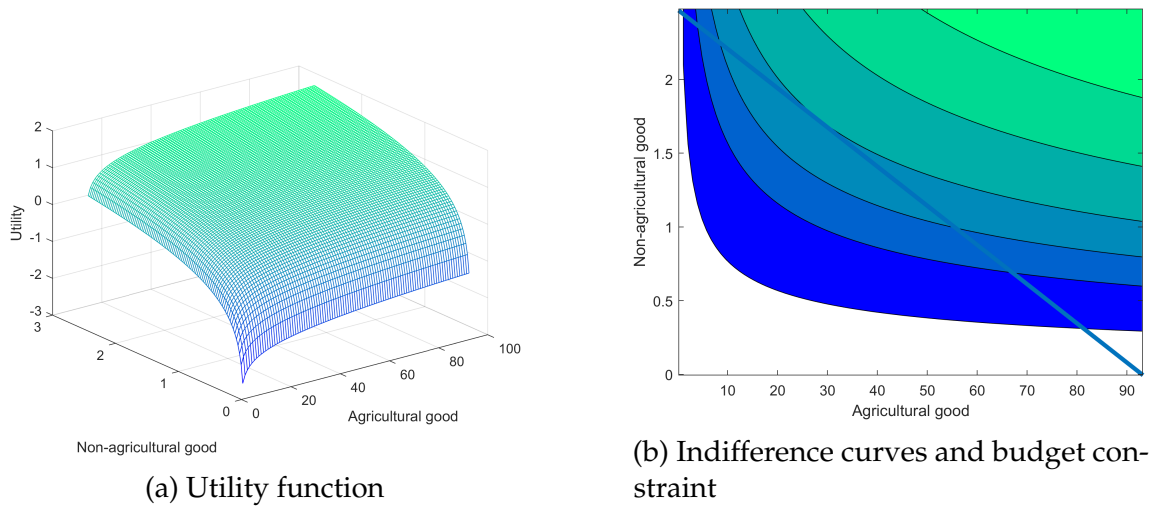


Figure A.5: Consumer's utility maximisation problem.

A.3.2. Firms

I assume the sector-specific firm profit maximisation problems for the agriculture and non-agriculture sectors A and N both take the form of a constant elasticity of

substitution function:

$$\max_{M_s, L_s} p_s B_s \left(\eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}} + (1 - \eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s-1}} - p_M M_s - w L_s, \quad s = A, N \quad (\text{A.9})$$

In this maximisation problem, η_s is the sector-specific motor vehicle share parameter, while σ_s represents the sector-specific elasticity of substitution between motor vehicle and labour inputs. Total factor productivity B_s also differs between the two sectors. This form of production function allows me to set constant sector-specific elasticities, σ_s , as part of the model calibration process. Figure A.6 illustrates this production function.




Figure A.6: Firm's production function.

The firm maximisation problem yields two first-order conditions:

$$\begin{aligned} \frac{\partial \Pi_s}{\partial M_s} &= p_s B_s \eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}-1} \left(\eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}} + (1 - \eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s-1}} - p_M = 0, \quad s = A, N \\ \frac{\partial \Pi_s}{\partial L_s} &= p_s B_s (1 - \eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}-1} \left(\eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}} + (1 - \eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s-1}} - w = 0, \quad s = A, N \end{aligned} \quad (\text{A.10})$$

These first-order conditions can be divided by each other as follows:

$$\frac{p_M}{w} = \frac{p_s B_s \eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}-1} \left(\eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}} + (1-\eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s-1}}}{p_s B_s (1-\eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}-1} \left(\eta_s M_s^{\frac{\sigma_s-1}{\sigma_s}} + (1-\eta_s) L_s^{\frac{\sigma_s-1}{\sigma_s}} \right)^{\frac{\sigma_s}{\sigma_s-1}}}, \quad s = A, N \quad (\text{A.11})$$

An input demand function for M_s in terms of L_s can then be derived by simplifying this equation and solving for M_s :

$$M_s = \left(\frac{w}{p_M} \right)^{\sigma_s} \left(\frac{\eta_s}{1-\eta_s} \right)^{\sigma_s} L_s, \quad s = A, N \quad (\text{A.12})$$

Meanwhile, I assume the firm profit maximisation problem for the intermediate motor vehicle sector M takes a linear form:

$$\max_{L_M} p_M B_M L_M - w L_M \quad (\text{A.13})$$

In this sector, a representative firm operating in perfect competition takes the economy-wide wage rate w and the price of motor vehicles p_M as exogenously given and chooses a level of labour inputs L_M in order to maximise profit Π_M . Total factor productivity in this sector is given by B_M . This maximisation problem yields the following first-order condition:

$$\frac{\partial \Pi_M}{\partial L_M} = p_M B_M - w = 0 \quad (\text{A.14})$$

This shows that $p_M B_M = w$ in my model.

A.3.3. Counterfactual 1: No structural transformation

In my first counterfactual model, I imposed that no structural transformation occurred in the economy. To achieve this, I merged the agriculture and non-

agriculture sectors into a single final good sector as follows.

First, I again assume the consumer optimisation problem takes the form of a log utility function but with consumers simply choosing a level of consumption of the single non-agriculture good that maximises utility subject to their budget constraint:

$$\begin{aligned} \max_{c_N} \quad & \log(c_N) \\ \text{s.t.} \quad & p_N c_N = w \end{aligned} \tag{A.15}$$

This optimisation problem yields a more straightforward demand function than my main model:

$$c_N = \frac{w}{p_N} \tag{A.16}$$

Second, I assume the non-agriculture firm profit maximisation problem again takes the form of a constant elasticity of substitution function, while the firm profit maximisation problem for the intermediate motor vehicle sector M continues to take a linear form. Given the absence of a separate agriculture sector, the agriculture-related parameters B_A , ϕ , \bar{c}_A , σ_A and η_A all disappear from the model, while all other parameters remained.

Therefore, it was necessary to re-calibrate the remaining parameters in this one-sector model using the same data and methodology as with my main model, but with agriculture and non-agriculture sectors combined into a single ‘non-agriculture’ sector in terms of inputs and value added. Table A.6 reports these re-calibrated values.

Table A.6: Re-calibrated one-sector parameter values: main panel aggregates 1995-2016

Parameter	Parameter description	Value	Target
$L_0, B_{s,0}$	Labour force and productivity, 1995	1.000	Normalisation
g_L	Labour force growth rate	1.004	Labour force growth
g_N	Non-agriculture TFP growth	2.777	Productivity growth in N
g_M	Motor vehicle TFP growth	1.603	Productivity growth in M
σ_N	Non-agriculture elasticity of substitution	0.604	Vehicle inputs in N , 2016
η_N	Non-agriculture vehicle share parameter	0.006	Labour share in M , 1995

TFP denotes total factor productivity. Annualised growth rates reported as percentages.

APPENDICES TO CHAPTER 2

B.1. Model theory

To establish this paper's theoretical framework, I outlined the problem facing a representative firm exporting a good from Ireland to France with a choice of two alternative transport routes, the road-based land-bridge route, R , and the direct short sea shipping route, S . In a static model, the firm takes the costs of each transport route, v_R and v_S , as exogenously given and allocates goods between the two routes to maximise profit Π . Assuming that the elasticity of substitution between these two routes is constant, the firm's constrained maximisation problem can be expressed as a constant elasticity of substitution production function:

$$\max_{R,S} pA \left(\eta R^{\frac{\sigma-1}{\sigma}} + (1-\eta) S^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - v_R R - v_S S \quad (\text{B.1})$$

In this problem, the price of the firm's good is denoted by p , A measures the firm's total factor productivity, and the constant $0 < \eta < 1$ parameter is the share parameter on the land-bridge route. The constant elasticity of substitution, denoted

by σ , measures how easy it is for the firm to switch between the two transport routes.

The firm's maximisation problem yields two first-order conditions:

$$\begin{aligned}\frac{\partial \Pi}{\partial R} &= pA\eta R^{\frac{\sigma-1}{\sigma}-1} \left(\eta R^{\frac{\sigma-1}{\sigma}} + (1-\eta)S^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - v_R = 0 \\ \frac{\partial \Pi}{\partial S} &= pA(1-\eta)S^{\frac{\sigma-1}{\sigma}-1} \left(\eta R^{\frac{\sigma-1}{\sigma}} + (1-\eta)S^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - v_S = 0\end{aligned}\tag{B.2}$$

These first-order conditions can be divided by each other as follows:

$$\frac{v_R}{v_S} = \frac{pA\eta R^{\frac{\sigma-1}{\sigma}-1} \left(\eta R^{\frac{\sigma-1}{\sigma}} + (1-\eta)S^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}}{pA(1-\eta)S^{\frac{\sigma-1}{\sigma}-1} \left(\eta R^{\frac{\sigma-1}{\sigma}} + (1-\eta)S^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}}\tag{B.3}$$

This can then be simplified and re-arranged to derive the ratio of the volume of goods transported on the road-based route to the volume of goods transported on the short sea shipping route in terms of relative route costs and constant model parameters. Introducing a time subscript $t = 0$ to denote the period prior to Brexit, this ratio is:

$$\frac{R_0}{S_0} = \left(\frac{v_S}{v_R} \right)^\sigma \left(\frac{\eta}{1-\eta} \right)^\sigma\tag{B.4}$$

This shows that the allocation of goods between the two routes is determined by the relative cost of the routes, $\frac{v_S}{v_R}$, the elasticity of the substitution between the two routes, σ , and the land-bridge share parameter, η . Post-Brexit, with $t = 1$, the introduction of non-tariff trade barriers due to Brexit can then be incorporated into the problem as an 'implicit tariff', denoted $\tau > 0$, on the land-bridge route, thus increasing its relative cost:

$$\frac{R_1}{S_1} = \left(\frac{v_S}{v_R(1+\tau)} \right)^\sigma \left(\frac{\eta}{1-\eta} \right)^\sigma\tag{B.5}$$

B.2. Additional tables and figures

B.2.1. Descriptive statistics

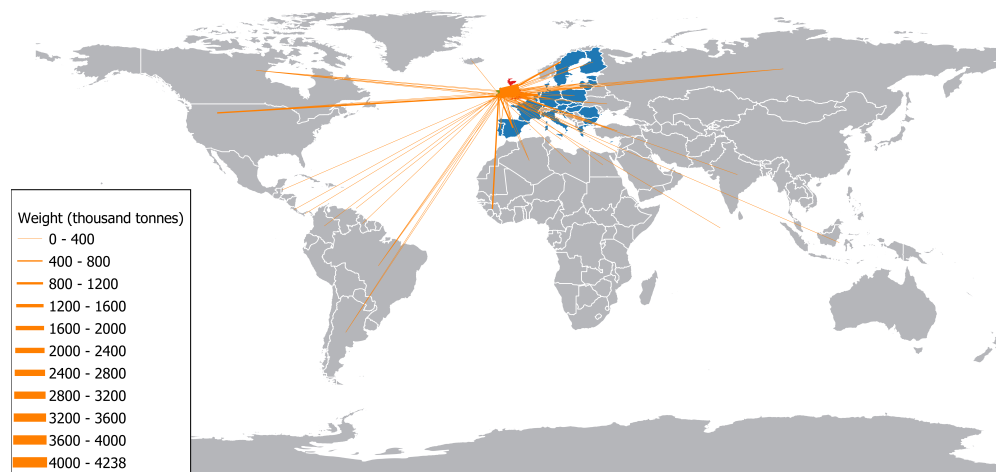


Figure B.1: Total cargo volumes by partner country, Irish main ports 2018. Line width determined by cargo volume. Projection: WGS84. Sources: Author's analysis; Eurostat [2022a](#); Esri [2022](#).

Figure B.1 illustrates the total cargo volumes (imports and exports combined, all cargo types aggregated) handled by Irish main ports by partner country in 2018. Line widths are determined by the cargo volume between the respective main port and partner country.

Figure B.2 displays total cargo volumes (imports and exports aggregated) shipped between EU-27 main ports and both the UK and the rest of the world (including other EU-27 countries) over the study period by cargo type. These were essentially the cargo volumes compared in my difference-in-differences analysis of EU-UK shipping (Equation 2.4). Any effects on EU-UK cargo volumes due to Brexit are

difficult to detect from observing Figure B.2.

Similarly, Figure B.3 shows the total cargo volumes shipped between Irish main ports and each the UK, France, and the rest of the world (excluding France, but including all other EU-27 countries) during the study period by cargo type. These cargo volumes were compared in my difference-in-differences analysis of Ireland-UK shipping and the UK land-bridge (Equation 2.5). As confirmed by my difference-in-differences results, a large impact due to Brexit was evident in Ro-Ro cargo.

The locations of the 6 main ports in Ireland are mapped in Figure B.4. Dublin, Rosslare Harbour and Cork are all equipped with berths to accommodate Ro-Ro cargo.

B.2.2. Additional results

I compared various methods for calculating standard errors for my difference-in-differences PPML regression models of cargo volumes in Table B.1. Note that while I report results as exponentiated coefficients in the paper, the coefficient in Table B.1 is the untransformed PPML coefficient rather than the exponentiated coefficient, and standard errors are similarly untransformed for greater ease in comparing them.

First, I obtained a variance-covariance matrix using a sandwich estimator (using Stata's `vce(robust)` option, also known as a Huber/White estimator), which allowed for heteroskedasticity and was implied in the use of PPML (see 'Robust' in Table B.1). This approach would have been valid if errors were identically distributed. However, correlation of errors between certain groups, such as cargo volumes associated with the same partner country, was a distinct possibility. Therefore, I then used various robust clustered variance estimators (using Stata's

B.2. Additional tables and figures

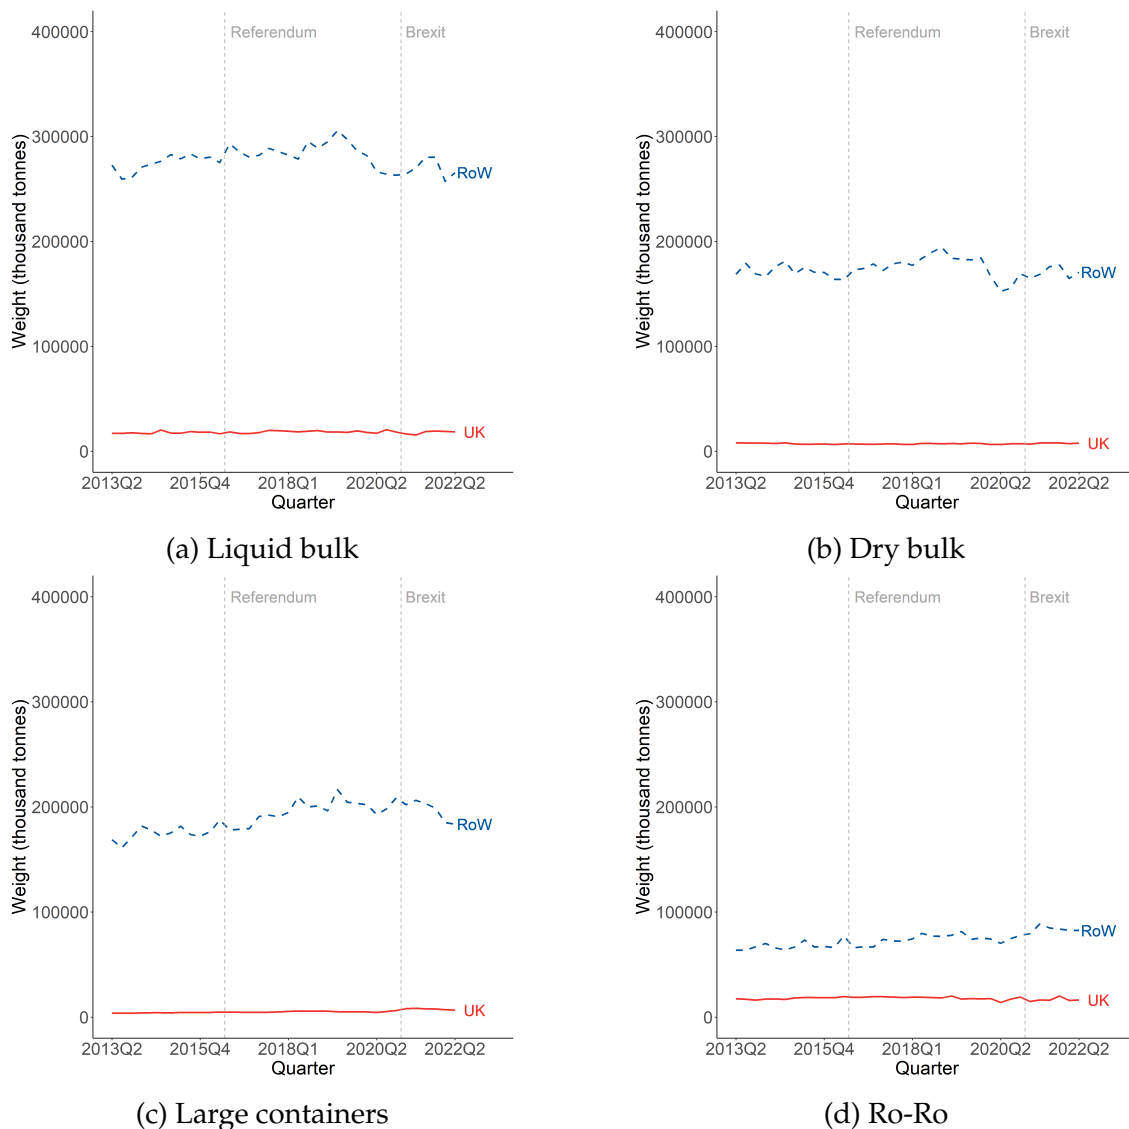


Figure B.2: Cargo volumes by partner entity, EU main ports 2013-2022. Sources: Author's analysis; Eurostat [2022a](#).

vce(cluster *cluster variable*) option), which relaxed this independence of errors assumption by additionally allowing for errors to be correlated within specified groups. First, I clustered at the origin port-partner country pair level (there were 16,622 different pairs, see 'Port-country'), allowing errors to be correlated between cargo volumes between a particular port and partner country. Second, I clustered at the country-country pair level (2,997 pairs, see 'Country-country'), which instead accommodated error correlation in cargo volumes between a particular

B. APPENDICES TO CHAPTER 2



Figure B.3: Cargo volumes by partner entity, Irish main ports 2013-2022. Sources: Author's analysis; Eurostat [2022a](#).

pair of countries. Third, I clustered at the partner country level (199 partner countries, see 'Partner country'), allowing correlation in errors between cargo volumes linked to the same partner country. Reassuringly, the results in Table B.1 held across all variance estimators that allowed for error correlation between groups. In the rest of my econometric analysis, I proceeded using robust standard errors clustered at the partner country level.

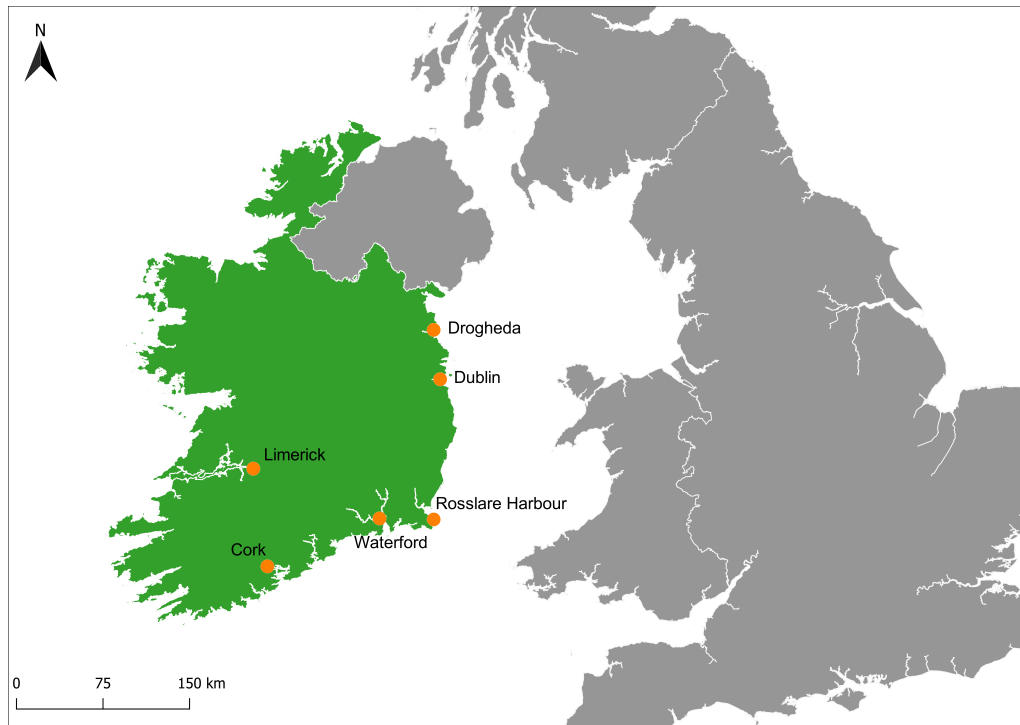


Figure B.4: Main ports in Ireland. Projection: TM65/Irish Grid. Sources: Author's analysis; Marine Institute [2023](#); Eurostat [2022a](#).

Table B.1: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022

	Coefficient	Standard error	Number of clusters
	-0.241		
Robust		0.177	-
Port-country		0.059	16,622
Country-country		0.050	2,997
Partner country		0.026	199
Observations	134568		
Pseudo R^2	0.377		
Quarter fixed effects	Yes		
Partner fixed effects	Yes		

Sources: Author's analysis; Eurostat [2022a](#)

Standard errors calculated using robust unclustered variance estimator in 'Robust'

Robust standard errors clustered at port-country pair level in 'Port-country'

Robust standard errors clustered at country-country pair level in 'Country-country'

Robust standard errors clustered at partner country level in 'Partner country'

In Table B.2, I confirmed that the negative impact on EU-UK Ro-Ro cargo volumes was evident in both imports from the UK and exports to the UK. In column 1, my main result of a 21.7 per cent decrease in total EU-UK Ro-Ro cargo volumes is repeated. EU imports from the UK were more heavily impacted than exports, with a 34.7 per cent decrease in imports (see column 2) compared with a 14.8 per cent decrease in exports (see column 3).

Table B.2: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022

	(1) Ro-Ro total	(2) Ro-Ro imports	(3) Ro-Ro exports
UK post-Brexit	0.783*** (0.020)	0.653*** (0.019)	0.852*** (0.022)
Observations	134568	84060	113868
Pseudo R^2	0.377	0.320	0.388
Quarter fixed effects	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at partner country level
and transformed using delta method

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.3 displays results of PPML regressions using the cargo type share of total cargo by partner country and period as the outcome variable instead of cargo volume to assess whether there was substitution between cargo types due to Brexit. These results provide some additional evidence of a small increase in the share of large containers in total cargo (see Column 3). The Ro-Ro result (see Column 4) suggests a decrease in the Ro-Ro share, although this result was not statistically significant.

Table B.4 presents results for various alternative regression specifications testing the robustness of my main results for Ro-Ro cargo in EU-27 main ports. Column 1 again repeats the main result of a 21.7 per cent decline in EU-UK Ro-Ro cargo

Table B.3: Regression difference-in-differences estimates of Brexit effect on cargo type shares in EU-27 main ports 2013-2022

	(1) Liquid bulk	(2) Dry bulk	(3) Large containers	(4) Ro-Ro
UK post-Brexit	1.012 (0.022)	0.992 (0.015)	1.029** (0.014)	0.992 (0.029)
Observations	251604	324179	303694	134502
Pseudo R^2	0.040	0.044	0.048	0.106
Quarter fixed effects	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at partner country level

and transformed using delta method

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

volumes due to Brexit (Equation 2.4). I confirmed that this result was not sensitive to the inclusion of high-dimensional fixed effects, as it held when instead using the PPML estimator developed by Correia, Guimarães, and Zylkin (2020) that identified and dropped separated observations (see column 2). To gain confidence in my key parallel trends assumption, I determined in a placebo test that there was no significant Brexit 'effect' on cargo volumes between the EU-27 and Sweden (see column 3). Column 4 shows that the result of a decrease in Ro-Ro cargo volumes held in my linear trend specification (Equation 2.6). Finally, column 5 confirms that the result also held in my post-referendum specification (Equation 2.7). I found a 6.3 per cent decrease in EU-UK Ro-Ro cargo volumes during the period between the June 2016 referendum and the UK's exit from the EU at the start of 2021, indicating some anticipatory or uncertainty effect. However, the magnitude of the post-treatment effect was larger than this anticipatory effect, with a 24.8 per cent decrease in cargo volumes after 2020.

Results for various robustness checks of my Ireland-UK findings are displayed in Table B.5. In column 1, the main results of a 53.8 per cent decrease in Ireland-UK Ro-Ro cargo volumes and a concurrent 147 increase in Ro-Ro cargo volumes

Table B.4: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022

	(1) Main	(2) HDFE	(3) Placebo	(4) Linear trend	(5) Post-ref.
UK post-referendum					0.937*** (0.017)
UK post-Brexit	0.783*** (0.020)	0.783*** (0.020)		0.881*** (0.028)	0.752*** (0.021)
SWE post-Brexit			0.974 (0.030)		
Observations	134568	130968	131292	134568	134568
Pseudo R^2	0.377	0.372	0.379	0.378	0.377
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes	Yes
Linear trend	No	No	No	Yes	No

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at partner country level

and transformed using delta method

HDFE denotes high-dimension fixed effects

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

between Irish ports and France are repeated. I employed the PPML estimator developed by Correia, Guimarães, and Zylkin (2020) that identified and dropped 'separated' observations in column 2 to determine that results were not sensitive to the inclusion of high-dimensional fixed effects. The results were robust to the inclusion of partner-specific linear trends (column 3) and also held in my post-referendum specification (column 4). Similar to my EU-UK results (Table B.4), I found a significant anticipatory effect on Ireland-UK Ro-Ro cargo volumes during the post-referendum period in column 4 of Table B.5. Interestingly, while the positive post-Brexit effect on Ireland-France cargo volumes persisted in this specification, column 4 of Table B.5 indicates that there was an initial negative impact on these cargo volumes in the post-referendum period. Given the trade relationship between Ireland and France did not change due to Brexit, this could be regarded as an uncertainty effect more than an anticipatory effect, particularly given the larger post-treatment effect was positive.

Table B.5: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022

	(1) Main	(2) HDFE	(3) Linear trend	(4) Post-ref.
UK post-referendum				0.906** (0.035)
FRA post-referendum				0.782*** (0.030)
UK post-Brexit	0.462*** (0.026)	0.462*** (0.026)	0.547*** (0.012)	0.433*** (0.035)
FRA post-Brexit	2.470*** (0.138)	2.470*** (0.140)	3.447*** (0.072)	2.113*** (0.169)
Observations	2772	1440	1405	2772
Pseudo R^2	0.716	0.647	0.644	0.716
Quarter fixed effects	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes
Linear trend	No	No	Yes	No

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at partner country level
and transformed using delta method

HDFE denotes high-dimension fixed effects (use of ppmlhdfe estimator)
ppmlhdfe estimator also used in linear trend specification

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.2.3. Untransformed PPML coefficients

Throughout the paper, I report results as exponentiated coefficients as these are easier to interpret than raw coefficients in the case of PPML estimators. For the benefit of the reader who prefers to interpret results from raw coefficients, Tables B.6, B.7, B.8, B.9, B.10, B.11 and B.12 and Figure B.5 present these same results as raw PPML coefficients with corresponding untransformed standard errors.

Table B.6: Regression difference-in-differences estimates of Brexit effect on cargo in EU-27 main ports 2013-2022

	(1) Total	(2) Liquid bulk	(3) Dry bulk	(4) Large containers	(5) Ro-Ro
UK post-Brexit	0.020 (0.022)	0.024 (0.049)	0.090** (0.041)	0.427*** (0.020)	-0.245*** (0.026)
Observations	598392	251604	324180	303696	134568
Pseudo R^2	0.258	0.223	0.202	0.275	0.377
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at partner country level

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.7: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022

	(1) Ro-Ro total	(2) Ro-Ro imports	(3) Ro-Ro exports
UK post-Brexit	-0.245*** (0.026)	-0.426*** (0.029)	-0.160*** (0.025)
Observations	134568	84060	113868
Pseudo R^2	0.377	0.320	0.388
Quarter fixed effects	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at partner country level

Sources: Author's analysis; Eurostat [2022a](#)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.8: Regression difference-in-differences estimates of Brexit effect on cargo type shares in EU-27 main ports 2013-2022

	(1) Liquid bulk	(2) Dry bulk	(3) Large containers	(4) Ro-Ro
UK post-Brexit	0.012 (0.022)	-0.008 (0.015)	0.029** (0.013)	-0.008 (0.029)
Observations	251604	324179	303694	134502
Pseudo R^2	0.040	0.044	0.048	0.106
Quarter fixed effects	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at partner country level

Sources: Author's analysis; Eurostat [2022a](#)* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.9: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in EU-27 main ports 2013-2022

	(1) Main	(2) HDFE	(3) Placebo	(4) Linear trend	(5) Post-ref.
UK post-referendum					-0.065*** (0.018)
UK post-Brexit	-0.245*** (0.026)	-0.245*** (0.026)		-0.127*** (0.031)	-0.285*** (0.029)
SWE post-Brexit			-0.026 (0.031)		
Observations	134568	130968	131292	134568	134568
Pseudo R^2	0.377	0.372	0.379	0.378	0.377
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes	Yes
Linear trend	No	No	No	Yes	No

Standard errors in parentheses

Robust standard errors clustered at partner country level

HDFE denotes high-dimension fixed effects

Sources: Author's analysis; Eurostat [2022a](#)* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.10: Regression difference-in-differences estimates of Brexit effect on cargo in Irish main ports 2013-2022

	(1) Total	(2) Liquid bulk	(3) Dry bulk	(4) Large containers	(5) Ro-Ro
UK post-Brexit	-0.276*** (0.079)	-0.177 (0.220)	0.080 (0.106)	-0.693*** (0.076)	-0.773*** (0.056)
FRA post-Brexit	0.630*** (0.079)	-0.232 (0.220)	-0.113 (0.106)	0.058 (0.076)	0.904*** (0.056)
Observations	10116	3708	7020	2592	2772
Pseudo R^2	0.665	0.538	0.615	0.829	0.716
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at partner country level

Sources: Author's analysis; Eurostat [2022a](#)* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.11: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022

	(1) Ro-Ro total	(2) Ro-Ro imports	(3) Ro-Ro exports
UK post-Brexit	-0.773*** (0.056)	-0.714*** (0.036)	-0.860*** (0.105)
FRA post-Brexit	0.904*** (0.056)	0.977*** (0.036)	0.824*** (0.105)
Observations	2772	1512	2268
Pseudo R^2	0.716	0.725	0.717
Quarter fixed effects	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at partner country level

Sources: Author's analysis; Eurostat [2022a](#)* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B.12: Regression difference-in-differences estimates of Brexit effect on Ro-Ro cargo in Irish main ports 2013-2022

	(1) Main	(2) HDFE	(3) Linear trend	(4) Post-ref.
UK post-referendum				-0.099** (0.039)
FRA post-referendum				-0.246*** (0.039)
UK post-Brexit	-0.773*** (0.056)	-0.773*** (0.057)	-0.604*** (0.023)	-0.837*** (0.080)
FRA post-Brexit	0.904*** (0.056)	0.904*** (0.057)	1.237*** (0.021)	0.748*** (0.080)
Observations	2772	1440	1405	2772
Pseudo R^2	0.716	0.647	0.644	0.716
Quarter fixed effects	Yes	Yes	Yes	Yes
Partner fixed effects	Yes	Yes	Yes	Yes
Linear trend	No	No	Yes	No


Standard errors in parentheses

Robust standard errors clustered at partner country level

HDFE denotes high-dimension fixed effects

Sources: Author's analysis; Eurostat 2022a

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$



(a) Ro-Ro cargo, EU-27 main ports

(b) Total cargo, Irish main ports

Figure B.5: Regression difference-in-differences estimates and 95 per cent confidence intervals for Brexit effect on quarterly cargo volumes, generalised models 2013-2022. Robust standard errors clustered at partner country level. Sources: Author's analysis; Eurostat 2022a.

APPENDICES TO CHAPTER 3

C.1. Data

Figure C.1 shows how I reduced the sample size of journeys from the SHS (2020), with all journeys recorded in the annual cross sections from 2012 to 2019 combined, to the sample size of 43,169 journeys used for my empirical analysis. First, I removed journeys that were not attached to a respondent ID, and thus could not be matched back to data on respondent characteristics, and journeys for which the travel diary weight or main transport mode (my primary outcome variable) was missing. Second, I removed any that started or ended in the island-based local authority areas of Shetland Islands, Orkney Islands or Na h-Eileanan Siar (the Western Isles) as these would have required air or water transport, which were not the focus of this study. Third, I removed journeys that were outliers in terms of distance, duration or imputed average speed, with outliers defined as being in excess of 3 standard deviations from the mean. Finally, I reduced the remaining Scotland-wide sample to a city sub-sample of journeys that either started or ended in any of the Glasgow City, City of Edinburgh, Dundee City or Aberdeen City local authority areas. The remaining sample of 43,169 journeys, employed throughout

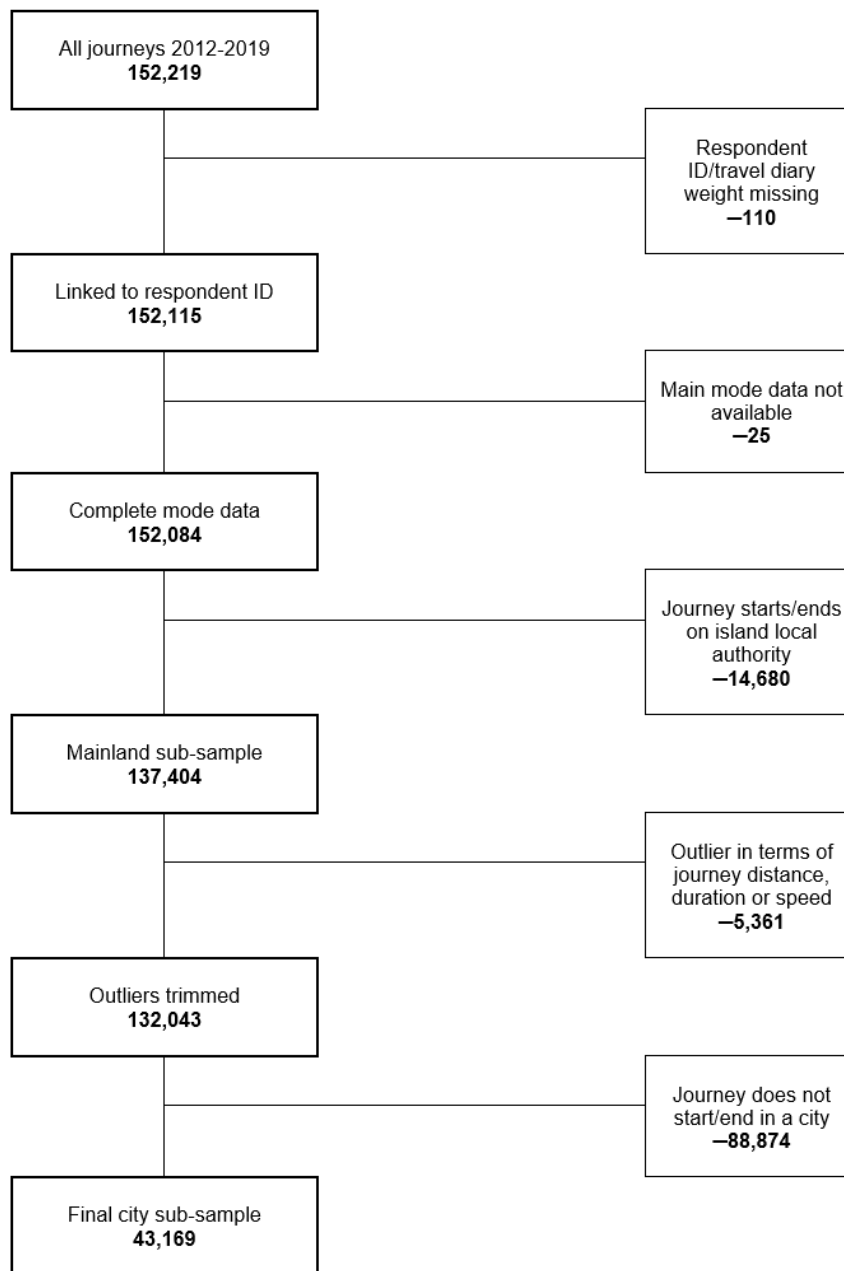


Figure C.1: Flow chart of study sample size from Scottish Household Survey 2012-2019. Sources: Author’s analysis; SHS 2020.

my empirical analysis, thus included only mainland journeys that started or ended in the cities of Glasgow, Edinburgh, Dundee or Aberdeen for which data on transport modes and respondent characteristics were available.

Table C.1 details how I collapsed the categorical variable of main transport mode

choice available from the SHS (2020) to a five-category variable for my primary outcome variable of interest. Bus and train categories were combined into a public transport category, car/van as passenger and taxi categories were combined into a single category due to the small size of the taxi category, and bicycle was added to the 'other' category due to its small size.

Table C.1: Collapsing categorical variable of mode choice

	Drive	Public	Passenger	Walk	Other	Total
Walking	0	0	0	10431	0	10431
Car/Van as driver	18786	0	0	0	0	18786
Car/Van as passenger	0	0	4694	0	0	4694
Bicycle	0	0	0	0	688	688
Bus	0	5917	0	0	0	5917
Train/Underground	0	1179	0	0	0	1179
Other	0	0	0	0	1139	1139
Taxi/minicab	0	0	335	0	0	335
Total	18786	7096	5029	10431	1827	43169

Drive denotes car as driver

Public denotes public transport

Passenger denotes car as passenger

Bus includes school/work and ordinary service

Other includes motorcycle/moped, ferry, air, horse-riding, tram

Sources: Author's analysis; SHS 2020

C.2. Additional figures and tables

C.2.1. Respondent characteristics

I compared the socio-economic and demographic characteristics of the SHS random adults of the control group and each Glasgow (see Table C.2) and Edinburgh (see Table C.3) for 2012-2015, before ride hailing became available. Figures C.2 and C.3 provide graphical illustrations of these characteristics over the entire 2012-2019 study period. Overall, the socio-demographic characteristics of the SHS random adults were broadly similar between the control group and each Glasgow and Edinburgh. Higher levels of education and a higher mean level of total household income were apparent in Edinburgh relative to the control group, while mean household income was lower in Glasgow than in the control group. In addition, a lower proportion of random adults undertaking journeys in Glasgow were married. Time invariant, city-specific socio-demographic factors that may have affected my outcome variables were picked up by city fixed effects γ_c in my regression specifications.

I also employed t-tests to detect differences between 2012 and 2019 in the socio-demographic characteristics of the surveyed random adults for each the Glasgow treatment group (see Table C.4), the Edinburgh treatment group (see Table C.5) and the control group of Dundee and Aberdeen (see Table C.6). These t-tests revealed that the mean of total household income increased in each group over the period. There was also evidence of an increase in education levels in Glasgow and the control group that was not reflected in Edinburgh. The t-tests revealed a 6 per cent increase in the average age among random adults making journeys to or from Edinburgh that was not reflected in Glasgow or the control group. Crucially for the parallel trends assumption of my identification strategy, the linking of these individual-level socio-demographic attributes with the travel diary data allowed

me to control for variation in these characteristics in my regression specifications.

Based on these differences, I included 4 individual-level socio-demographic control variables in all regression specifications. These were categorical variables for gender, age group, household education and household income. These were included as categorical variables to account for possible non-linear relationships between these factors and my outcome variables. Figure C.4 displays each of these variables by city.

Table C.2: Socio-demographic characteristics of Glasgow treatment group and control group 2012-2015

	Glasgow	Control	Difference	p-value
Age	47.11	47.53	-0.43	0.36
Female	0.54	0.54	0.01	0.63
Finished school	0.60	0.62	-0.02	0.17
Degree	0.33	0.33	0.00	0.74
Employed	0.59	0.58	0.01	0.56
Retired	0.20	0.22	-0.02	0.11
Married	0.35	0.42	-0.06	0.00
Household income	25495.52	27556.89	-2061.38	0.00
Observations	6069			

Sources: Author's analysis; SHS [2020](#)

C.2.2. Additional results

To assess the impact of including each of my 4 individual-level variables controlling for socio-demographic characteristics, I first ran my difference-in-differences multinomial logistic regression of mode choice without any control variables, and then proceeded to iteratively add controls to the specification. Column 1 of Table C.7 displays results for the specification devoid of control variables. Columns 2 to 5 then iteratively add controls for gender, age, household education and household income. Table C.7 shows that my main result of a positive effect on the use of public

Table C.3: Socio-demographic characteristics of Edinburgh treatment group and control group 2012-2015

	Edinburgh	Control	Difference	p-value
Age	46.81	47.53	-0.72	0.16
Female	0.51	0.54	-0.02	0.11
Finished school	0.73	0.62	0.11	0.00
Degree	0.48	0.33	0.16	0.00
Employed	0.63	0.58	0.05	0.00
Retired	0.21	0.22	-0.01	0.54
Married	0.43	0.42	0.01	0.50
Household income	28743.19	27556.89	1186.29	0.01
Observations	4949			

Sources: Author's analysis; SHS 2020

Table C.4: Socio-demographic characteristics of Glasgow treatment group in 2012 and 2019

	2019	2012	Difference	p-value
Age	47.62	46.30	1.31	0.10
Female	0.52	0.55	-0.02	0.34
Finished school	0.71	0.59	0.11	0.00
Degree	0.43	0.33	0.10	0.00
Employed	0.62	0.58	0.04	0.11
Retired	0.20	0.21	-0.01	0.63
Married	0.37	0.34	0.03	0.21
Household income	30171.46	25238.15	4933.31	0.00
Observations	1888			

Sources: Author's analysis; SHS 2020

transport in Glasgow held across all of these specifications, and that the pseudo-R squared statistic, measuring each regression model's goodness-of-fit (although it should be noted that the pseudo-R squared statistic is not a direct equivalent of the R squared statistic from OLS), was improved by the addition of each control variable. Based on this, I proceeded in the rest of my difference-in-differences analysis of mode choice and journey speed including all 4 socio-demographic

Table C.5: Socio-demographic characteristics of Edinburgh treatment group in 2012 and 2019

	2019	2012	Difference	p-value
Age	48.76	46.01	2.75	0.00
Female	0.56	0.51	0.04	0.11
Finished school	0.70	0.72	-0.02	0.36
Degree	0.43	0.48	-0.05	0.04
Employed	0.63	0.62	0.01	0.61
Retired	0.25	0.21	0.04	0.09
Married	0.40	0.43	-0.02	0.40
Household income	32027.06	27629.91	4397.15	0.00
Observations	1452			

Sources: Author's analysis; SHS [2020](#)

Table C.6: Socio-demographic characteristics of control group in 2012 and 2019

	2019	2012	Difference	p-value
Age	48.10	47.49	0.61	0.60
Female	0.54	0.56	-0.03	0.43
Finished school	0.69	0.60	0.09	0.00
Degree	0.40	0.31	0.09	0.00
Employed	0.58	0.57	0.01	0.70
Retired	0.23	0.24	-0.01	0.81
Married	0.42	0.42	-0.01	0.85
Household income	31518.12	25892.78	5625.34	0.00
Observations	984			

Control group comprised of Dundee and Aberdeen.
Sources: Author's analysis; SHS [2020](#)

control variables.

In addition, as shown in Table C.8, I compared various methods for calculating standard errors for my difference-in-differences multinomial logistic regression of mode choice. For assessing the calculation of standard errors, I did not apply travel diary weights to the regression summarised in Table C.8. For ease of comparing standard errors, results are reported as untransformed coefficients and standard



Figure C.2: Random adult characteristics, control and treatment groups 2012-2019. Control group comprised of Dundee and Aberdeen. Axis shows percentage of random adults. Sources: Author’s analysis; SHS 2020.

errors in Table C.8. Therefore, the regression summarised in Table C.8 corresponds with the regression results shown in column 4 (‘unweighted’) of Table C.17.

For ‘Default’ in Table C.8, I calculated standard errors using the default observed information matrix variance estimator, which assumed errors were independent and identically distributed normal. I tested this assumption in ‘Robust’ by instead using a robust unclustered Huber/White/sandwich variance estimator (using Stata’s `vce(robust)` option), which allowed for heteroskedasticity, and found little change in standard errors. For ‘Clustered’, as I discussed in the main paper, I then clustered standard errors at the individual level (using Stata’s `vce(cluster cluster variable)` option), which additionally allowed for errors to be correlated between journeys recorded by the same individual. This increased standard errors compared with ‘Default’ and ‘Robust’, although my main results held. I proceeded in the rest of my econometric analysis of mode choice and journey speed using robust standard errors clustered at the individual level (as in ‘Clustered’).




Figure C.3: Age and total household net annual income distributions of random adults, control and treatment groups 2012-2019. Control group comprised of Dundee and Aberdeen. Sources: Author's analysis; SHS 2020.

I assessed the sensitivity of my results to the choice of combining Dundee and Aberdeen as the control group by running my main regression of mode choice again using only Dundee journeys, and then using only Aberdeen journeys, as the control group. Results for these alternative regressions are summarised (columns 2 and 3 respectively) alongside my main results (column 1) in Table C.9, with the ride hailing effect on the use of public transport in Glasgow holding in both

C. APPENDICES TO CHAPTER 3

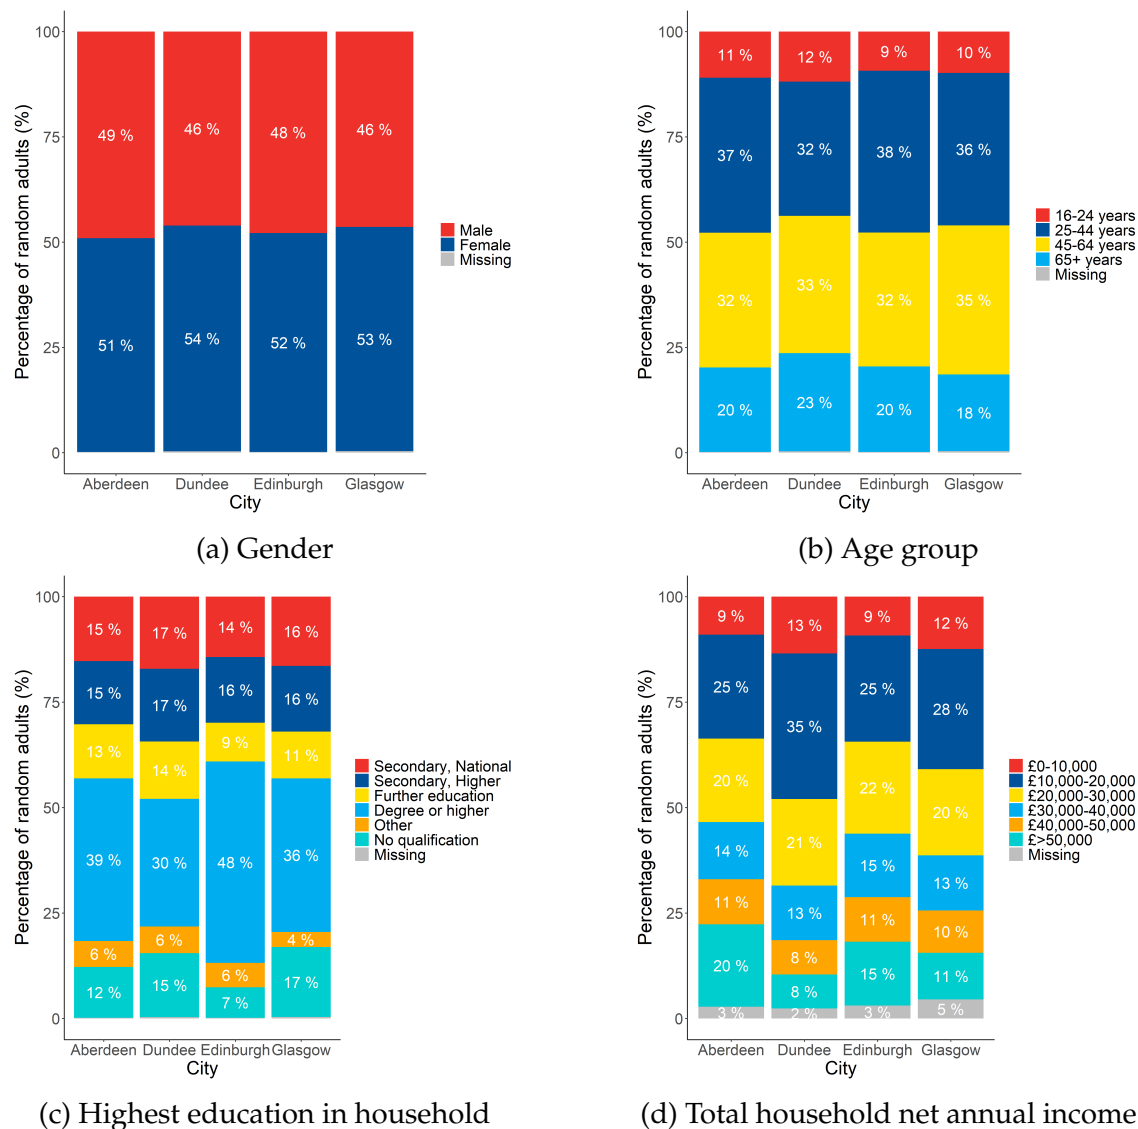


Figure C.4: Socio-demographic characteristics of random adults, 2012-2019. Sources: Author's analysis; SHS 2020.

specifications. In addition, to gain confidence in the validity of my key parallel trends assumption, I conducted a placebo test comparing the average change over time in all mainland non-city local authorities with the average change in Dundee and Aberdeen. Reassuringly, as shown in column 4 of Table C.9, no ride hailing effect on mode choice was found in the placebo test.

Table C.10 summarises results for various alternative regression specifications I

C.2. Additional figures and tables

Table C.7: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) No controls	(2) Controls	(3) Controls	(4) Controls	(5) Controls
Public					
Treated Glasgow	1.463*** (0.199)	1.461*** (0.199)	1.521*** (0.208)	1.609*** (0.220)	1.749*** (0.242)
Treated Edinburgh	1.065 (0.151)	1.056 (0.150)	1.089 (0.156)	1.096 (0.157)	1.138 (0.166)
Passenger					
Treated Glasgow	1.098 (0.147)	1.091 (0.148)	1.143 (0.156)	1.183 (0.162)	1.211 (0.166)
Treated Edinburgh	1.412** (0.209)	1.385** (0.207)	1.428** (0.215)	1.435** (0.218)	1.450** (0.220)
Walk					
Treated Glasgow	1.204 (0.140)	1.204 (0.140)	1.236* (0.146)	1.264** (0.150)	1.381*** (0.167)
Treated Edinburgh	1.298** (0.158)	1.295** (0.157)	1.329** (0.162)	1.337** (0.165)	1.385** (0.177)
Other					
Treated Glasgow	0.407*** (0.129)	0.409*** (0.129)	0.412*** (0.130)	0.402*** (0.127)	0.424*** (0.134)
Treated Edinburgh	0.847 (0.275)	0.853 (0.275)	0.856 (0.277)	0.863 (0.278)	0.889 (0.287)
Observations	43169	43169	43169	43169	43169
Pseudo R^2	0.015	0.023	0.051	0.067	0.091
Year fixed effects	Yes	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes	Yes
Control: female	No	Yes	Yes	Yes	Yes
Control: age group	No	No	Yes	Yes	Yes
Control: education	No	No	No	Yes	Yes
Control: income group	No	No	No	No	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

ran to assess the robustness of my main results. Column 1 repeats my main results for the choice of transport mode. Column 2 shows that these results were largely unchanged by the inclusion of an additional dummy variable in the model that controlled for the opening of Edinburgh Trams. This is a tramway connecting

Table C.8: Regression difference-in-differences estimates of ride hailing effect on choice of public transport as main mode in Glasgow

VCE estimator	Coefficient	Standard error	Number of clusters
	0.523		
Default		0.080	-
Robust		0.079	-
Clustered		0.117	16,712
Observations	43169		
Pseudo R^2	0.084		
Year fixed effects	Yes		
City fixed effects	Yes		
Individual controls	Yes		

Sources: Author's analysis; SHS 2020

VCE denotes variance-covariance matrix of estimator. Standard errors calculated using observed information matrix variance estimator in 'Default'. Robust standard errors calculated using Huber/White/sandwich estimator in 'Robust'. Robust standard errors clustered at individual level in 'Clustered'. Travel diary weights not applied when assessing calculation of standard errors

the city centre to Edinburgh airport, which opened to passengers in May 2014. The additional dummy variable was thus set equal to 1 if the journey started or ended in Edinburgh in 2015 or later. The coefficient on the dummy variable was not statistically significant for any comparison transport mode.

I also ran my main regression specification, with mode choice as the outcome variable, again on a sub-sample of journeys that started and ended in the same city, on the basis that ride hailing may be used mainly for short journeys within an urban area. This sub-sample omitted inter-city journeys, and journeys connecting a city with a peripheral local authority. Column 3 of Table C.10 presents summary results for this sub-sample regression, showing that although standard errors were higher due to the smaller sample size, my main results largely continued to hold. Column 4 of Table C.10 shows summary results for my main regression specification using data that had not been adjusted using any survey weights. These results confirm that my main results were not sensitive to the inclusion of travel

diary weights.

Finally, column 5 of Table C.10 summarises results for my alternative regression specification that additionally controlled for city-specific linear trends. I discuss this specification, including its advantages and disadvantages, in the main paper. As expected given the more restrictive regression specification, standard errors were higher (column 5) than in my main results (column 1). As discussed in the paper, these results indicate that the positive effect on public transport in Glasgow did not persist when linear trends were controlled for.

To test whether the effect on public transport affected bus journeys or rail journeys, I ran my main regression specification again with the mode choice outcome variable amended to split the public transport category into separate bus and rail categories. Results for this regression are displayed alongside my repeated main results in Table C.11, with bus and rail categories separated in column 2. These results clearly indicate that the public transport effect stemmed from rail journeys, with significant positive effects found on the use of rail in both Glasgow and Edinburgh. Meanwhile, no effect was found among bus journeys in either city.

To test the difference in the proportion of public transport journeys that also involved taking a car as a passenger before and after the introduction of ride hailing, I ran a logistic regression of this proportion on a $POST_t$ dummy that was equal to 1 if the journey occurred in 2016 or later, and 0 otherwise. As reported in Table C.12, I found a significant increase in the probability that a public transport journey also involved the use of a car as a passenger in the period after the launch of ride hailing.

C.2.3. Untransformed multinomial logistic regression coefficients

In the paper, for greater ease of interpretation, I report multinomial logistic regression results as exponentiated coefficients, or relative risk ratios. For those more partial to interpreting results from untransformed coefficients, Tables C.13, C.14, C.15, C.16, C.17 and C.18, and Figure C.5 present these same multinomial logistic regression results as raw coefficients with corresponding untransformed standard errors. These coefficients show the change in the multinomial log-odds of the respective mode being chosen over the reference transport mode (driving a car) due to ride hailing becoming available in that city.




Figure C.5: Regression difference-in-differences estimates coefficients and 95 per cent confidence intervals for effect of ride hailing effect on public transport, generalised model 2012-2019. Robust standard errors clustered at individual level. Uber operating in Glasgow and Edinburgh from 2016 onwards (red dashed line). Sources: Author's analysis; SHS 2020.

C.3. Transport in Scotland

Across the UK, the provision of public transport has been characterised by privatisation for some time. Bus services were privatised from 1986, and this was followed by the privatisation of railways in 1993.¹ In Scotland, all commuter rail services are operated by ScotRail, a brand name that has been owned by various private companies but was re-nationalised by the Scottish Government in 2022. Local bus services in Scotland are provided by private operators.

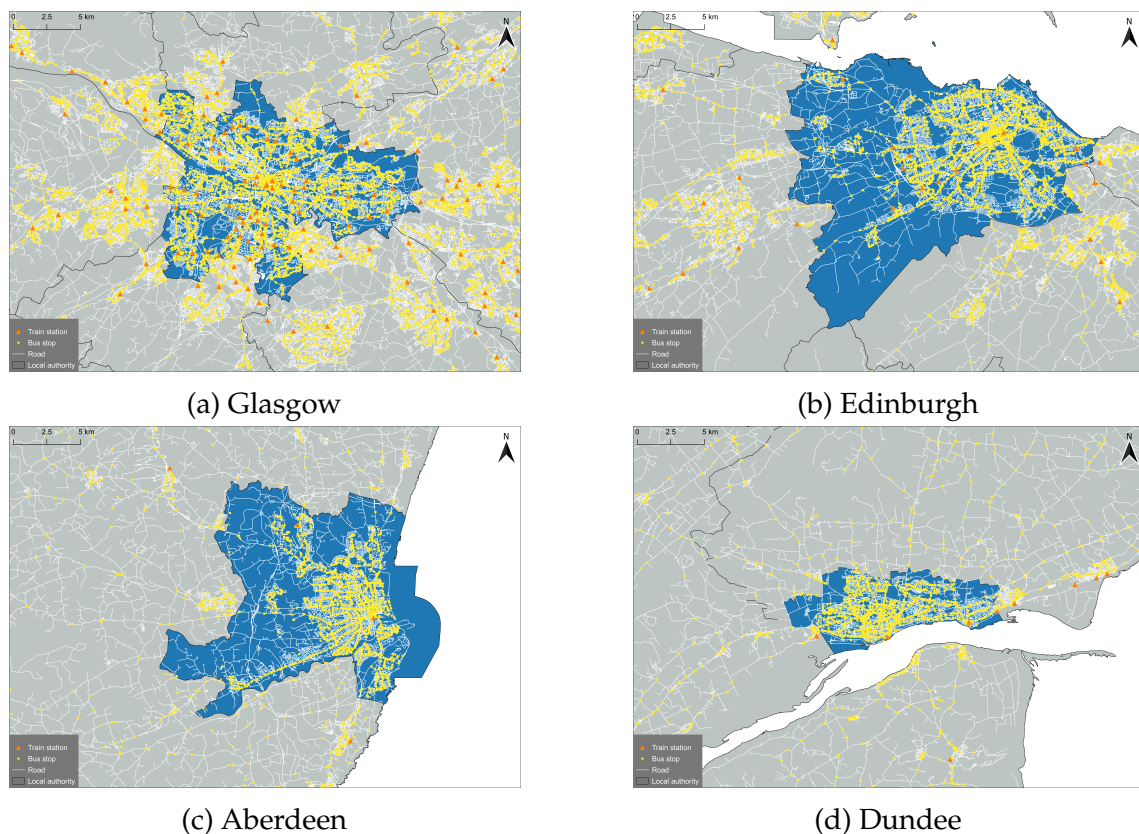


Figure C.6: Roads and public transport stops, 2020. Source: Author's analysis; UK Department for Transport [2020a](#); Ordnance Survey [2020](#).

¹See Gunn (2018) for a history of transport systems in the UK.

C.3.1. City transport profiles

Figure 3.3 in the main paper maps the geographical locations of Glasgow, Edinburgh, Aberdeen and Dundee within Scotland. Figure C.6 illustrates the road networks, bus stops and railway stations in each of the 4 cities as of 2020.

Glasgow has an international airport, and the city is served by an extensive road network, suburban railway lines and a light metro line. Several motorways run through the city, including the M8 connecting with Edinburgh to the east, the M73, the M74 connecting with England to the south, the M77 and the M80.² A low emission zone (LEZ), where all vehicles that do not meet a low-emissions standard are charged a penalty fee, has applied to bus vehicles in central Glasgow since 2018, and to other vehicles since June 2023 (although residents of the LEZ are exempt until 2024). Glasgow's suburban railway is operated by ScotRail and based around the Glasgow Central and Glasgow Queen Street terminus stations. Several private operators provide an extensive bus network in the city. A single circular light metro route, Glasgow Subway, has been in operation since 1896 and was modernised in 1977. There has been no tram network in Glasgow since 1962.

Edinburgh is also served by an international airport, in addition to road and rail networks. The Edinburgh City Bypass (A720) links with major roads including the M8 connecting with Glasgow to the west, the A1 connecting with England to the south, and the M9. There are currently no congestion charges in the city, although a LEZ will be enforced from 2024. The suburban railway is operated by ScotRail and based in the terminus station of Edinburgh Waverley. The main public transport option in Edinburgh is the bus network, which is operated by private companies. Edinburgh's original tram system closed in 1956, but Edinburgh Trams, a single route connecting the airport on the western outskirts of the city to Prince's Street

²In the UK, motorway numbers are prefixed with 'M'. There is also a network of major non-motorway roads that are numbered with the prefix 'A', known as A-roads.

and St Andrews Square in the city centre, was opened in May 2014.

Aberdeen has a small international airport and also boasts a major helicopter terminal that serves offshore oil installations in the North Sea. A network of 6 A-roads connects the city with the rest of Scotland. There are currently no congestion charges in the city, although a LEZ will be enforced from 2024. ScotRail operates the suburban railway, with 2 railway stations in the city. There is also a bus network in the city, offered by private operators. There has been no tram network in Aberdeen since 1958.

There is a small domestic airport in **Dundee**, and the city is served by the A90 road that connects with Aberdeen to the north and to the M90 motorway. There are currently no congestion charges in the city, although a LEZ will be enforced from 2024. As with Aberdeen, there are 2 main railway stations in the city and the suburban railway is operated by ScotRail. Similar to the other three cities, Dundee's tram system was closed in 1956 with routes replaced by diesel buses, and the city now has an extensive bus network operated by private companies.

While statistics on population density or the density of transport infrastructure are partly determined by the boundaries of the city administrative areas, it can be observed from Figure C.6 that infrastructure density is higher in Glasgow than in other cities. To emphasise this, Figure C.7 compares population and infrastructure density across the 4 cities, with Glasgow boasting the highest density in terms of population, roads, bus stops and particularly rail stations. While both Glasgow and Edinburgh are larger population centres than Dundee and Aberdeen, as shown in Figure 3.2 in the main paper, Glasgow is a denser city than Edinburgh, and in particular is served by a more extensive rail network. This may offer some explanation for a difference in treatment effects between these cities.

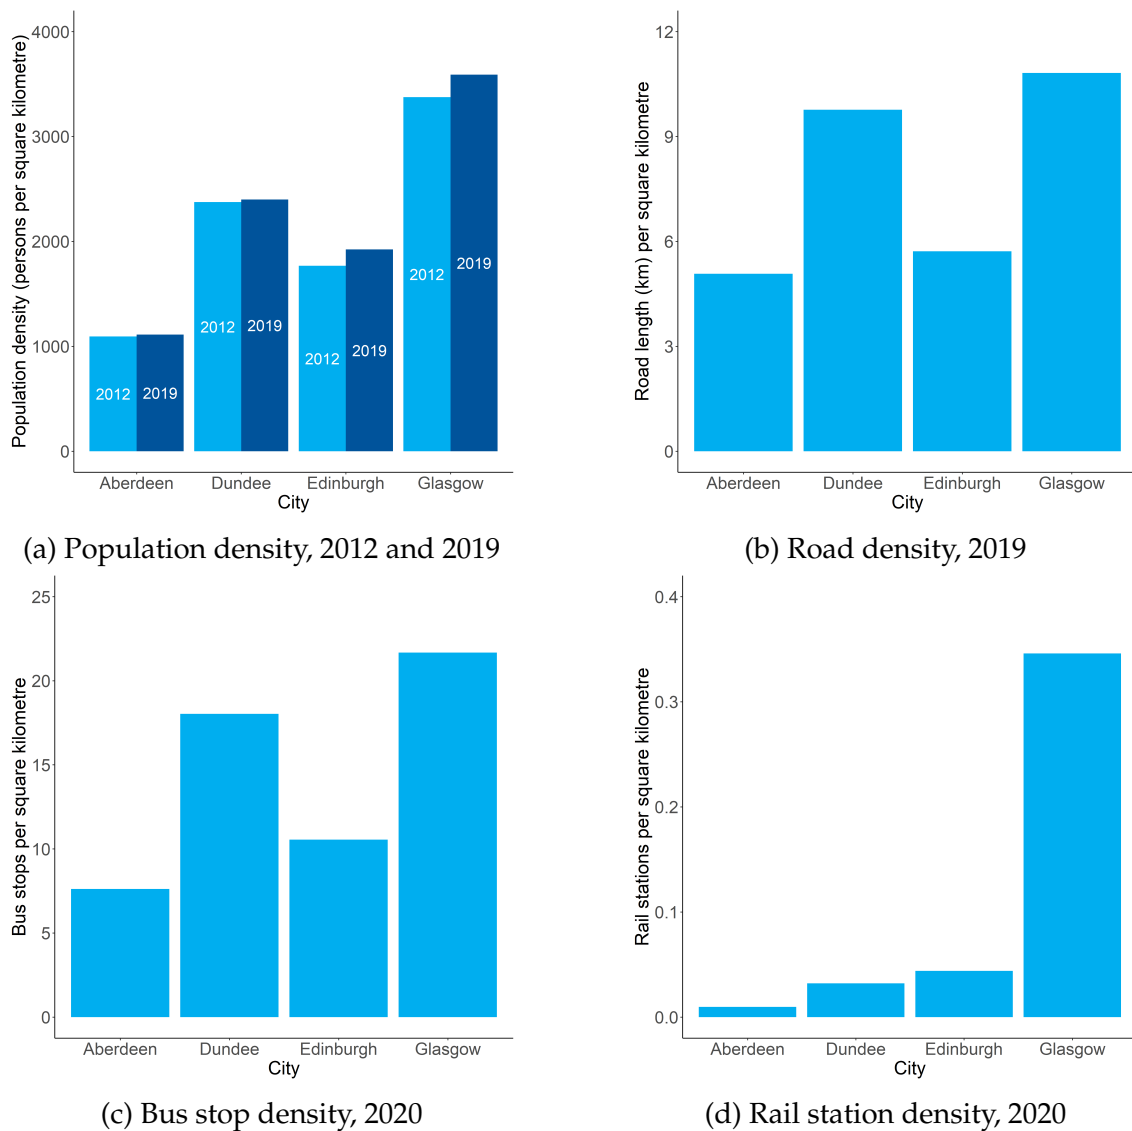


Figure C.7: Population density and density of transport infrastructure by city. Source: Author’s analysis; Office for National Statistics 2020; Ordnance Survey 2020; UK Department for Transport 2020a, 2020b.

C.3.2. Changes in transport infrastructure

Figure C.8 displays some statistics on changes in transport infrastructure in Scotland during my study period between 2012 and 2019. First, there was very little change in the total length of the road network in any of the 4 cities, ranging from a 5.82 per cent increase in Aberdeen to a 2.68 per cent increase in Dundee between 2012 and 2019. Any effect on my outcome variables from this small

increase in road length across all 4 cities would be captured by the year fixed effects, while the fact that total road length is different between cities would be captured by the city fixed effects.

Second, there was virtually no change in the number of passenger rail stations across the 4 cities between 2012 and 2019. One new station, Edinburgh Gateway, opened in Edinburgh in 2016 to provide a connection with Edinburgh Trams near the airport.

Third, there was more change evident in the number of licensed bus vehicles. While there was little change in Aberdeen or Dundee, Glasgow experienced a 15 per cent decline in licensed bus vehicles in 2017 before recovering slightly in 2019, while there was a steady increase of 35.37 per cent between 2012 and 2019 in Edinburgh.

Fourth, the number of licensed taxi vehicles was almost static between 2012 and 2019 in both Glasgow and Edinburgh. However, the number of taxis was 17.35 per cent lower in Aberdeen, and 11.53 per cent lower in Dundee, in 2019 than in 2012.

C.3.3. Changes in transport prices

While the extent of transport infrastructure may not have changed substantially between 2012 and 2019, Figure C.9 paints a different picture for transport prices, namely road fuel prices and public transport fares. However, while there was no city-level data available for these variables, I argue here that there are unlikely to have been significant city-specific changes during this time.

First, local bus fares in Scotland increased by 26.14 per cent over the course of my study period, compared with 21.05 per cent more broadly in Great Britain. Figure C.9 shows that the increase in local bus fares was closely related to the increase in the retail price index for Great Britain. There was a similar increase

C. APPENDICES TO CHAPTER 3




Figure C.8: Transport infrastructure in Scotland, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author’s analysis; Transport Scotland 2020; UK Department for Transport 2020b.

of 17.92 per cent in rail fares across Great Britain. Any effect on my outcome variables of this general increase in public transport fares, which appears to have been largely in line with inflation in retail prices, would have been captured by the year fixed effects, while the city fixed effects would have accounted for any city-specific differences that did not change over time. The possibility of city-specific


changes over time in public transport fares between 2012 and 2019 seems remote, particularly in the case of rail transport since ScotRail is the single operator for the suburban rail network across all 4 cities.

Second, there was considerable fluctuation in fuel prices between 2012 and 2019 in the UK. Figure C.9 indicates very similar patterns in typical retail prices for premium unleaded petrol and diesel, with a significant decrease during 2014 and 2015 before a subsequent steady recovery for the prices of both fuels. Figure C.9 also plots the price index for crude oil acquired by refineries over this period, as I contend that these changes in road fuel prices were mainly influenced by the price of crude oil, a factor exogenous to transport systems in Scotland. In addition, taxes on road fuel are applied at a UK level and thus would not vary between cities. Any effect on my outcome variables from city-specific deviations in fuel prices from these UK-wide typical retail prices, that could arise from other factors such as fuel transport or storage costs, would have been captured by the city fixed effects unless these city-specific differences changed significantly over time between 2012 and 2019, which seems unlikely.


C.3.4. Changes in aggregate transport use

While this study focused on individual journeys using SHS (2020) travel diary data, it is also worth looking at aggregate measures of transport use. In my study setting of Glasgow, Edinburgh, Dundee and Aberdeen between 2012 and 2019, such aggregate data did not afford large enough sample sizes to apply the difference-in-differences methodology employed with travel diary data. However, Figures C.10 and C.11 offer some descriptive statistics of these aggregate measures. Figure C.10 displays annual aggregate statistics on private car ownership, car traffic and road traffic collisions by city. The number of licensed private cars


C. APPENDICES TO CHAPTER 3



(a) Index of local bus fares, Scotland and Great Britain 2012-2019 (2012 = 100). 'RPI' shows retail price index for Great Britain.



(b) Index of rail fares (all ticket types), Great Britain 2012-2019 (2012 = 100). 'RPI' shows retail price index for Great Britain.




(c) Indices of typical retail prices for petrol (premium unleaded) and diesel, UK 2012-2019 (January 2012 = 100). 'Crude oil' shows price index for crude oil acquired by refineries.

Figure C.9: Transport prices in the UK, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author's analysis; UK Department for Energy Security and Net Zero [2023c](#); UK Department for Transport [2023](#); Office of Rail and Road [2022](#).

increased steadily by 10.16 per cent in Edinburgh, 3.73 per cent in Aberdeen and 12.74 per cent in Dundee between 2012 and 2019. A sharp decrease in licensed


cars of 14.07 per cent was evident in Glasgow in 2017, the same year in which I first found a significant positive effect of ride hailing on public transport. Despite this, however, there was relatively little change during the study period in total car traffic (measured in vehicle kilometres) across the 4 cities, apart from a notable increase in Aberdeen in 2019. The number of road traffic collisions declined over the study period, with particularly pronounced declines of 33.78 per cent and 36.31 per cent in Glasgow and Edinburgh respectively between 2016 and 2019. Consistent with my difference-in-differences results using SHS (2020) travel diary data, there is no evidence in Figure C.10 of a substantial increase in car use in Glasgow or Edinburgh following the introduction of ride hailing in late 2015.

Figure C.11, meanwhile, depicts annual aggregate statistics on public transport use during the study period. Between 2012 and 2019, the number of within-Scotland rail passenger journeys (measured by origin or destination city) increased by 13.42 per cent, 21.51 per cent and 23.8 per cent in Glasgow, Edinburgh and Dundee respectively, while journeys actually decreased by 29.8 per cent in Aberdeen. There was a 4.64 per cent increase in rail passenger journeys to or from Glasgow in 2017 alone, reflective of my difference-in-differences results for Glasgow. On the other hand, the number of local bus journeys was decreasing across Scotland over this period, with a decrease of 18.25 per cent in the South-West and Strathclyde region (including Glasgow), 4.73 per cent in the South-East region (including Edinburgh), and 22.8 per cent in the North-East, Tayside and Central region (including Dundee and Aberdeen) between 2012 and 2019. The substantial decrease in the South-West region is reflective of the decline in the number of licensed bus vehicles in Glasgow, while the much smaller decrease in journeys in the South-East region may be linked to the increase in bus vehicles in Edinburgh (see Figure C.8).



(a) Licensed private cars by city

(b) Car traffic by city



(c) Number of road traffic collisions by city

Figure C.10: Aggregate measures of car use by city or region, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author’s analysis; UK Department for Transport 2022; Transport Scotland 2020.

Table C.9: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Dundee	(3) Aberdeen	(4) Placebo
Public				
Treated Glasgow	1.749*** (0.242)	1.810*** (0.325)	1.685*** (0.300)	
Treated Edinburgh	1.138 (0.166)	1.189 (0.220)	1.097 (0.202)	
Placebo treatment				0.960 (0.128)
Passenger				
Treated Glasgow	1.211 (0.166)	1.024 (0.182)	1.397** (0.238)	
Treated Edinburgh	1.450** (0.220)	1.234 (0.234)	1.669*** (0.304)	
Placebo treatment				1.038 (0.120)
Walk				
Treated Glasgow	1.381*** (0.167)	1.381** (0.215)	1.366** (0.205)	
Treated Edinburgh	1.385** (0.177)	1.396** (0.225)	1.372** (0.213)	
Placebo treatment				1.070 (0.102)
Other				
Treated Glasgow	0.424*** (0.134)	0.741 (0.346)	0.257*** (0.099)	
Treated Edinburgh	0.889 (0.287)	1.562 (0.735)	0.534 (0.209)	
Placebo treatment				0.762 (0.236)
Observations	43169	37350	37930	99932
Pseudo R^2	0.091	0.086	0.088	0.076
Year fixed effects	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Control group consist of only Dundee in Column 2

Control group consists of only Aberdeen in Column 3

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.10: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Tram	(3) Within-city	(4) Unweighted	(5) Linear trend
Public					
Treated Glasgow	1.749*** (0.242)	1.754*** (0.243)	1.729*** (0.295)	1.687*** (0.198)	1.077 (0.308)
Treated Edinburgh	1.138 (0.166)	1.054 (0.173)	1.042 (0.176)	1.145 (0.141)	0.799 (0.243)
Passenger					
Treated Glasgow	1.211 (0.166)	1.214 (0.167)	1.515** (0.279)	1.224* (0.144)	0.562** (0.154)
Treated Edinburgh	1.450** (0.220)	1.358* (0.249)	1.355 (0.256)	1.331** (0.170)	1.225 (0.390)
Walk					
Treated Glasgow	1.381*** (0.167)	1.384*** (0.168)	1.480*** (0.211)	1.393*** (0.145)	1.181 (0.292)
Treated Edinburgh	1.385** (0.177)	1.274* (0.187)	1.360** (0.193)	1.359*** (0.148)	1.196 (0.310)
Other					
Treated Glasgow	0.424*** (0.134)	0.424*** (0.134)	0.381** (0.154)	0.351*** (0.095)	1.406 (0.731)
Treated Edinburgh	0.889 (0.287)	0.853 (0.296)	0.804 (0.326)	0.831 (0.234)	1.131 (0.581)
Observations	43169	43169	28328	43169	43169
Pseudo R^2	0.091	0.091	0.084	0.084	0.092
Year fixed effects	Yes	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes
Edinburgh tram control	No	Yes	No	No	No
Travel diary weights	Yes	Yes	Yes	No	Yes
Linear trend	No	No	No	No	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.11: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Rail
Public		
Treated Glasgow	1.749*** (0.242)	
Treated Edinburgh	1.138 (0.166)	
Passenger		
Treated Glasgow	1.211 (0.166)	1.190 (0.163)
Treated Edinburgh	1.450** (0.220)	1.444** (0.220)
Walk		
Treated Glasgow	1.381*** (0.167)	1.341** (0.162)
Treated Edinburgh	1.385** (0.177)	1.378** (0.176)
Other		
Treated Glasgow	0.424*** (0.134)	0.455** (0.142)
Treated Edinburgh	0.889 (0.287)	0.968 (0.311)
Bus		
Treated Glasgow		1.142 (0.168)
Treated Edinburgh		1.048 (0.159)
Rail		
Treated Glasgow		6.205*** (2.464)
Treated Edinburgh		2.372** (1.019)
Observations	43169	43057
Pseudo R^2	0.091	0.098
Year fixed effects	Yes	Yes
City fixed effects	Yes	Yes
Individual controls	Yes	Yes

Exponentiated coefficients; Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$


Table C.12: Post-treatment coefficient for use of car as passenger as part of journey

	(1)
Car as passenger used in stage POST	13.368*** (8.664)
Observations	6017
Pseudo R^2	0.210
Individual controls	Yes

Exponentiated coefficients; Standard errors in parentheses
Standard errors clustered at individual level

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$



(a) Number of within-Scotland passenger rail journeys by origin or destination city

(b) Number of bus passenger journeys by region. Glasgow is in South-West, Strathclyde; Edinburgh is in South-East; Dundee and Aberdeen are both in North-East, Tayside, Central.

Figure C.11: Aggregate measures of public transport use by city or region, 2012-2019. Ride hailing became available in Glasgow and Edinburgh in late 2015 (grey dashed line). Source: Author's analysis; Transport Scotland 2020.

Table C.13: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Under 45	(3) Male	(4) Employed	(5) High income	(6) Degree
Public						
Treated Glasgow	0.559*** (0.139)	0.641*** (0.209)	0.839*** (0.210)	0.729*** (0.189)	1.094*** (0.322)	0.592** (0.257)
Treated Edinburgh	0.129 (0.146)	0.014 (0.218)	0.252 (0.215)	0.292 (0.193)	0.330 (0.321)	0.123 (0.254)
Passenger						
Treated Glasgow	0.191 (0.137)	0.032 (0.202)	-0.060 (0.227)	0.189 (0.187)	0.324 (0.276)	-0.020 (0.226)
Treated Edinburgh	0.372** (0.152)	0.150 (0.227)	0.415 (0.257)	0.247 (0.208)	0.463 (0.302)	0.240 (0.248)
Walk						
Treated Glasgow	0.323*** (0.121)	0.272 (0.170)	0.213 (0.173)	0.289* (0.160)	0.538** (0.264)	0.071 (0.192)
Treated Edinburgh	0.325** (0.128)	0.177 (0.180)	0.285 (0.183)	0.244 (0.159)	0.321 (0.262)	0.253 (0.184)
Other						
Treated Glasgow	-0.859*** (0.316)	-0.494 (0.382)	-0.287 (0.375)	-0.880** (0.385)	-0.390 (0.745)	-1.215*** (0.454)
Treated Edinburgh	-0.118 (0.323)	0.048 (0.394)	0.231 (0.386)	-0.251 (0.392)	0.293 (0.745)	-0.154 (0.458)
Observations	43169	20290	19978	25805	10333	17758
Pseudo R^2	0.091	0.089	0.100	0.067	0.071	0.071
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.14: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Work	(3) Leisure
Public			
Treated Glasgow	0.559*** (0.139)	1.019*** (0.237)	0.327 (0.285)
Treated Edinburgh	0.129 (0.146)	0.646*** (0.242)	-0.255 (0.289)
Passenger			
Treated Glasgow	0.191 (0.137)	-0.032 (0.333)	0.450* (0.234)
Treated Edinburgh	0.372** (0.152)	-0.009 (0.387)	0.583** (0.255)
Walk			
Treated Glasgow	0.323*** (0.121)	0.117 (0.246)	0.501** (0.234)
Treated Edinburgh	0.325** (0.128)	0.251 (0.237)	0.116 (0.240)
Other			
Treated Glasgow	-0.859*** (0.316)	-0.786* (0.463)	-1.328** (0.538)
Treated Edinburgh	-0.118 (0.323)	0.007 (0.470)	-0.925* (0.560)
Observations	43169	11224	8526
Pseudo R^2	0.091	0.094	0.098
Year fixed effects	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.15: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) No controls	(2) Controls	(3) Controls	(4) Controls	(5) Controls
Public					
Treated Glasgow	0.381*** (0.136)	0.379*** (0.136)	0.419*** (0.137)	0.476*** (0.137)	0.559*** (0.139)
Treated Edinburgh	0.063 (0.142)	0.054 (0.142)	0.085 (0.143)	0.092 (0.143)	0.129 (0.146)
Passenger					
Treated Glasgow	0.093 (0.134)	0.087 (0.136)	0.134 (0.136)	0.168 (0.137)	0.191 (0.137)
Treated Edinburgh	0.345** (0.148)	0.325** (0.150)	0.356** (0.151)	0.361** (0.152)	0.372** (0.152)
Walk					
Treated Glasgow	0.185 (0.117)	0.185 (0.116)	0.212* (0.118)	0.234** (0.119)	0.323*** (0.121)
Treated Edinburgh	0.261** (0.121)	0.259** (0.121)	0.284** (0.122)	0.290** (0.123)	0.325** (0.128)
Other					
Treated Glasgow	-0.900*** (0.317)	-0.895*** (0.316)	-0.888*** (0.316)	-0.912*** (0.316)	-0.859*** (0.316)
Treated Edinburgh	-0.166 (0.324)	-0.159 (0.322)	-0.155 (0.323)	-0.147 (0.322)	-0.118 (0.323)
Observations	43169	43169	43169	43169	43169
Pseudo R^2	0.015	0.023	0.051	0.067	0.091
Year fixed effects	Yes	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes	Yes
Control: female	No	Yes	Yes	Yes	Yes
Control: age group	No	No	Yes	Yes	Yes
Control: education	No	No	No	Yes	Yes
Control: income group	No	No	No	No	Yes

Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.16: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Dundee	(3) Aberdeen	(4) Placebo
Public				
Treated Glasgow	0.559*** (0.139)	0.594*** (0.180)	0.522*** (0.178)	
Treated Edinburgh	0.129 (0.146)	0.173 (0.185)	0.093 (0.184)	
Placebo treatment				-0.041 (0.133)
Passenger				
Treated Glasgow	0.191 (0.137)	0.024 (0.178)	0.334** (0.170)	
Treated Edinburgh	0.372** (0.152)	0.210 (0.189)	0.512*** (0.182)	
Placebo treatment				0.037 (0.116)
Walk				
Treated Glasgow	0.323*** (0.121)	0.323** (0.155)	0.312** (0.150)	
Treated Edinburgh	0.325** (0.128)	0.334** (0.161)	0.316** (0.156)	
Placebo treatment				0.068 (0.095)
Other				
Treated Glasgow	-0.859*** (0.316)	-0.300 (0.467)	-1.360*** (0.386)	
Treated Edinburgh	-0.118 (0.323)	0.446 (0.470)	-0.628 (0.392)	
Placebo treatment				-0.272 (0.309)
Observations	43169	37350	37930	99932
Pseudo R^2	0.091	0.086	0.088	0.076
Year fixed effects	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Control group consist of only Dundee in Column 2

Control group consists of only Aberdeen in Column 3

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.17: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Tram	(3) Within-city	(4) Unweighted	(5) Linear trend
Public					
Treated Glasgow	0.559*** (0.139)	0.562*** (0.139)	0.548*** (0.171)	0.523*** (0.117)	0.074 (0.286)
Treated Edinburgh	0.129 (0.146)	0.053 (0.164)	0.041 (0.169)	0.136 (0.123)	-0.224 (0.304)
Passenger					
Treated Glasgow	0.191 (0.137)	0.194 (0.137)	0.416** (0.184)	0.202* (0.117)	-0.576** (0.274)
Treated Edinburgh	0.372** (0.152)	0.306* (0.183)	0.303 (0.189)	0.286** (0.128)	0.203 (0.319)
Walk					
Treated Glasgow	0.323*** (0.121)	0.325*** (0.121)	0.392*** (0.143)	0.332*** (0.104)	0.166 (0.247)
Treated Edinburgh	0.325** (0.128)	0.242* (0.147)	0.307** (0.142)	0.307*** (0.109)	0.179 (0.259)
Other					
Treated Glasgow	-0.859*** (0.316)	-0.857*** (0.316)	-0.966** (0.403)	-1.048*** (0.271)	0.341 (0.520)
Treated Edinburgh	-0.118 (0.323)	-0.159 (0.347)	-0.218 (0.405)	-0.185 (0.281)	0.123 (0.514)
Observations	43169	43169	28328	43169	43169
Pseudo R^2	0.091	0.091	0.084	0.084	0.092
Year fixed effects	Yes	Yes	Yes	Yes	Yes
City fixed effects	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes
Edinburgh tram control	No	Yes	No	No	No
Travel diary weights	Yes	Yes	Yes	No	Yes
Linear trend	No	No	No	No	Yes

Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.18: Regression difference-in-differences estimates of ride hailing effect on choice of main mode

	(1) Main	(2) Rail
<hr/>		
Public		
Treated Glasgow	0.559*** (0.139)	
Treated Edinburgh	0.129 (0.146)	
<hr/>		
Passenger		
Treated Glasgow	0.191 (0.137)	0.174 (0.137)
Treated Edinburgh	0.372** (0.152)	0.368** (0.152)
<hr/>		
Walk		
Treated Glasgow	0.323*** (0.121)	0.293** (0.121)
Treated Edinburgh	0.325** (0.128)	0.321** (0.128)
<hr/>		
Other		
Treated Glasgow	-0.859*** (0.316)	-0.788** (0.313)
Treated Edinburgh	-0.118 (0.323)	-0.033 (0.322)
<hr/>		
Bus		
Treated Glasgow		0.132 (0.147)
Treated Edinburgh		0.047 (0.151)
<hr/>		
Rail		
Treated Glasgow		1.825*** (0.397)
Treated Edinburgh		0.864** (0.430)
<hr/>		
Observations	43169	43057
Pseudo R^2	0.091	0.098
Year fixed effects	Yes	Yes
City fixed effects	Yes	Yes
Individual controls	Yes	Yes

Standard errors in parentheses

Robust standard errors clustered at individual level

Outcome reference category: Car as driver

Public denotes public transport

Passenger denotes car as passenger

Sources: Author's analysis; SHS 2020

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

REFERENCES

- Acemoglu, Daron, and Veronica Guerrieri. 2008. "Capital Deepening and Non-balanced Economic Growth." *Journal of Political Economy* 116 (3): 467–498. <https://doi.org/10.1086/589523>.
- Acheampong, Ransford A., Alhassan Siiba, Dennis K. Okyere, and Justice P. Tuffour. 2020. "Mobility-on-demand: An empirical study of internet-based ride-hailing adoption factors, travel characteristics and mode substitution effects." *Transportation Research Part C: Emerging Technologies* 115:102638. <https://doi.org/10.1016/j.trc.2020.102638>.
- Agarwal, Saharsh, Deepa Mani, and Rahul Telang. 2019. "The Impact of Ridesharing Services on Congestion: Evidence from Indian Cities." *SSRN Electronic Journal*, <https://doi.org/10.2139/ssrn.3410623>.
- Andreoni, James, and Arik Levinson. 2001. "The simple analytics of the environmental Kuznets curve." *Journal of Public Economics* 80 (2): 269–286. [https://doi.org/10.1016/s0047-2727\(00\)00110-9](https://doi.org/10.1016/s0047-2727(00)00110-9).
- Angrist, Joshua D, and Jörn-Steffen Pischke. 2008. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Arnott, Richard. 1996. "Taxi Travel Should Be Subsidized." *Journal of Urban Economics* 40 (3): 316–333. <https://doi.org/10.1006/juec.1996.0035>.

- Autor, David H. 2003. "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing." *Journal of Labor Economics* 21 (1): 1–42. <https://doi.org/10.1086/344122>.
- Bahamonde-Birke, Francisco J. 2020. "Who will bell the cat? On the environmental and sustainability risks of electric vehicles." *Transportation Research Part A: Theory and Practice* 133:79–81. <https://doi.org/10.1016/j.tra.2019.12.001>.
- Baier, Scott L., and Jeffrey H. Bergstrand. 2007. "Do free trade agreements actually increase members' international trade?" *Journal of International Economics* 71 (1): 72–95. <https://doi.org/10.1016/j.jinteco.2006.02.005>.
- Bartoletto, Silvana, and M.d.Mar Rubio. 2008. "Energy Transition and CO2 Emissions in Southern Europe: Italy and Spain (1861-2000)." *Global Environment* 1 (2): 46–81. <https://doi.org/10.3197/ge.2008.010203>.
- Bastian, Anne, and Maria Börjesson. 2015. "Peak car? Drivers of the recent decline in Swedish car use." *Transport Policy* 42:94–102. <https://doi.org/10.1016/j.tranpol.2015.05.005>.
- Bastian, Anne, Maria Börjesson, and Jonas Eliasson. 2016. "Explaining "peak car" with economic variables." *Transportation Research Part A: Policy and Practice* 88:236–250. <https://doi.org/10.1016/j.tra.2016.04.005>.
- Bates, John. 2018. "Forecasting the demand for transport." Chap. 11, First, edited by Jonathan Cowie and Stephen Ison, 157–175. *The Routledge Handbook of Transport Economics*. Oxon: Routledge.
- Bertrand, M., E. Duflo, and S. Mullainathan. 2004. "How Much Should We Trust Differences-In-Differences Estimates?" *The Quarterly Journal of Economics* 119 (1): 249–275. <https://doi.org/10.1162/003355304772839588>.

- Besley, Timothy, and Robin Burgess. 2004. "Can Labor Regulation Hinder Economic Performance? Evidence from India." *The Quarterly Journal of Economics* 119 (1): 91–134. <https://doi.org/10.1162/003355304772839533>.
- Bogart, Dan. 2014. "The Transportation Revolution in Industrialising Britain." Chap. The Cambridge Economic History of Modern Britain, edited by Roderick Floud, Jane Humphries, and Paul Johnson, 1:368–391. Cambridge, United Kingdom: Cambridge University Press.
- Boisjoly, Geneviève, Emily Grisé, Meadhbh Maguire, Marie-Pier Veillette, Robbin Deboosere, Emma Berrebi, and Ahmed El-Geneidy. 2018. "Invest in the ride: A 14-year longitudinal analysis of the determinants of public transport ridership in 25 North American cities." *Transportation Research Part A: Policy and Practice* 116:434–445. <https://doi.org/10.1016/j.tra.2018.07.005>.
- Boppart, Timo. 2014. "Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences." *Econometrica* 82 (6): 2167–2196. <https://doi.org/10.3982/ecta11354>.
- Breen, Ben, Paul Brewster, Andreas Tsakaridis, and Ciaran O'Driscoll. 2018. *The implications of Brexit on the use of the landbridge*. Technical report. Dublin: Irish Maritime Development Office. <https://www.imdo.ie/Home/sites/default/files/IMDOFiles/972918IMDOTheImplicationsofBrexitontheUseoftheLandbridgeReport-DigitalFinal.pdf>.
- Brunt, Liam, and Cecilia García-Peñalosa. 2021. "Urbanisation and the Onset of Modern Economic Growth." *The Economic Journal* 132 (642): 512–545. <https://doi.org/10.1093/ej/ueab050>.
- Bussière, Yves D., Jean-Loup Madre, and Irving Tapia-Villarreal. 2019. "Will peak car observed in the North occur in the South? A demographic approach with

- case studies of Montreal, Lille, Juarez and Puebla." *Economic Analysis and Policy* 61:39–54. <https://doi.org/10.1016/j.eap.2018.06.002>.
- Carson, Richard T. 2010. "The Environmental Kuznets Curve: Seeking Empirical Regularity and Theoretical Structure." *Review of Environmental Economics and Policy* 4 (1): 3–23. <https://doi.org/10.1093/reep/rep021>.
- CE Delft. 2021. *STREAM Freight Transport 2020: Emissions of freight transport modes*. Technical report. Delft: CE Delft, February. https://cedelft.eu/wp-content/uploads/sites/2/2021/03/CE_Delft_190325_STREAM_Freight_Transport_2020_FINAL.pdf.
- Chen, Natalie, Dennis Novy, Carlo Perroni, and Horng Chern Wong. 2023. "Urban-biased structural change." *CEPR Discussion Paper*, no. DP18522 (October). <https://cepr.org/publications/dp18522>.
- Cleveland, William S. 1979. "Robust Locally Weighted Regression and Smoothing Scatterplots." *Journal of the American Statistical Association* 74 (368): 829–836. <https://doi.org/10.1080/01621459.1979.10481038>.
- Cleveland, William S., and Susan J. Devlin. 1988. "Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting." *Journal of the American Statistical Association* 83 (403): 596–610. <https://doi.org/10.1080/01621459.1988.10478639>.
- Comin, Diego, Danial Lashkari, and Martí Mestieri. 2021. "Structural Change With Long-Run Income and Price Effects." *Econometrica* 89 (1): 311–374. <https://doi.org/10.3982/ecta16317>.
- Contreras, Seth D., and Alexander Paz. 2018. "The effects of ride-hailing companies on the taxicab industry in Las Vegas, Nevada." *Transportation Research Part A: Policy and Practice* 115:63–70. <https://doi.org/10.1016/j.tra.2017.11.008>.

- Copeland, Brian R, and M. Scott Taylor. 2004. "Trade, Growth, and the Environment." *Journal of Economic Literature* 42 (1): 7–71. <https://doi.org/10.1257/42.1.7>.
- Correia, Sergio, Paulo Guimarães, and Tom Zylkin. 2020. "Fast Poisson estimation with high-dimensional fixed effects." *The Stata Journal: Promoting communications on statistics and Stata* 20 (1): 95–115. <https://doi.org/10.1177/1536867x20909691>.
- Correia, Sergio, Paulo Guimarães, and Thomas Zylkin. 2019. "Verifying the existence of maximum likelihood estimates for generalized linear models." *arXiv*, <https://doi.org/10.48550/ARXIV.1903.01633>.
- Cravino, Javier, and Sebastian Sotelo. 2019. "Trade-Induced Structural Change and the Skill Premium." *American Economic Journal: Macroeconomics* 11 (3): 289–326. <https://doi.org/10.1257/mac.20170434>.
- Crozet, Yves. 2020. "Cars and Space Consumption: Rethinking the Regulation of Urban Mobility." *International Transport Forum Discussion Papers*, nos. 2020/13, <https://www.itf-oecd.org/sites/default/files/docs/cars-space-consumption-regulation-urban-mobility.pdf>.
- Davies, Richard B., and Clifford M. Guy. 1987. "The Statistical Modeling of Flow Data When the Poisson Assumption Is Violated." *Geographical Analysis* 19 (4): 300–314. <https://doi.org/10.1111/j.1538-4632.1987.tb00132.x>.
- Davies, Ronald B., and Zuzanna Studnicka. 2018. "The heterogeneous impact of Brexit: Early indications from the FTSE." *European Economic Review* 110:1–17. <https://doi.org/10.1016/j.euroecorev.2018.08.003>.

REFERENCES

- De Jong, Gerard, James Fox, Andrew Daly, Marits Pieters, and Remko Smit. 2004. "Comparison of car ownership models." *Transport Reviews* 24 (4): 379–408. <https://doi.org/10.1080/0144164032000138733>.
- Delbosc, Alexa, and Graham Currie. 2013. "Causes of Youth Licensing Decline: A Synthesis of Evidence." *Transport Reviews* 33 (3): 271–290. <https://doi.org/10.1080/01441647.2013.801929>.
- Deutsche Bahn. 2019. *2018 Integrated Report: On track towards a better railway*. Technical report. Berlin: Deutsche Bahn, March. <https://ir.deutschebahn.com/en/reports/archive/details/integrierter-bericht-2018-deutsche-bahn0>.
- Dias, Felipe F., Patrícia S. Lavieri, Taehooie Kim, Chandra R. Bhat, and Ram M. Pendyala. 2019. "Fusing Multiple Sources of Data to Understand Ride-Hailing Use." *Transportation Research Record: Journal of the Transportation Research Board* 2673 (6): 214–224. <https://doi.org/10.1177/0361198119841031>.
- Dodgson, J. S. 1986. "Benefits of Changes in Urban Public Transport Subsidies in the Major Australian Cities." *Economic Record* 62 (2): 224–235. <https://doi.org/10.1111/j.1475-4932.1986.tb00898.x>.
- Douch, Mustapha, and Terence Huw Edwards. 2021. "The bilateral trade effects of announcement shocks: Brexit as a natural field experiment." *Journal of Applied Econometrics* 37 (2): 305–329. <https://doi.org/10.1002/jae.2878>.
- Douglas, George W. 1972. "Price Regulation and Optimal Service Standards: The Taxicab Industry." *Journal of Transport Economics and Policy* 6 (2): 116–127. ISSN: 00225258, accessed August 25, 2023. <http://www.jstor.org/stable/20052271>.
- Duarte, Margarida, and Diego Restuccia. 2010. "The Role of the Structural Transformation in Aggregate Productivity." *Quarterly Journal of Economics* 125 (1): 129–173. <https://doi.org/10.1162/qjec.2010.125.1.129>.

- Duflo, Esther. 2001. "Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment." *American Economic Review* 91 (4): 795–813. <https://doi.org/10.1257/aer.91.4.795>.
- Echevarria, Cristina. 1997. "Changes in Sectoral Composition Associated with Economic Growth." *International Economic Review* 38 (2): 431. <https://doi.org/10.2307/2527382>.
- Esri. 2022. *World countries - generalized*, May.
- European Commission. 2016. *Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: European strategy for low-emission mobility*. Brussels: Publications Office of the European Union, July.
- . 2011. *White Paper on Transport: Roadmap to a Single European Transport Area: Towards a Competitive and Resource-Efficient Transport System*. Brussels: Publications Office of the European Union, March.
- Eurostat. 2022a. *Gross weight of goods transported to/from main ports - quarterly data*, December. <https://ec.europa.eu/eurostat/web/main/data/database>.
- . 2022b. *Reference manual on maritime transport statistics*. Technical report. Luxembourg: Eurostat, January. https://ec.europa.eu/eurostat/documents/29567/3217334/Reference+Manual+Maritime_v4_2017.pdf/341b5fb2-1fff-4d3e-917f-b329525a79e8.
- Fageda, Xavier. 2021. "Measuring the impact of ride-hailing firms on urban congestion: The case of Uber in Europe." *Papers in Regional Science* 100 (5): 1230–1253. <https://doi.org/10.1111/pirs.12607>.

- Fearnley, Nils, Graham Currie, Stefan Flügel, Fredrik A. Gregersen, Marit Killi, Jeremy Toner, and Mark Wardman. 2018. "Competition and substitution between public transport modes." *Research in Transportation Economics* 69:51–58. <https://doi.org/10.1016/j.retrec.2018.05.005>.
- Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2015. "The Next Generation of the Penn World Table." *American Economic Review* 105 (10): 3150–3182. <https://doi.org/10.1257/aer.20130954>.
- Flynn, Eimear, Janez Kren, and Martina Lawless. 2021a. "Early reactions of EU-UK trade flows to Brexit." *ESRI Working Paper 713*, <https://www.esri.ie/publications/early-reactions-of-eu-uk-trade-flows-to-brexit>.
- . 2021b. "Initial impact of Brexit on Ireland-UK trade flows." *ESRI Working Paper 714*, <https://www.esri.ie/publications/initial-impact-of-brexit-on-ireland-uk-trade-flows>.
- Forster, Bruce A. 1973. "Optimal Capital Accumulation in a Polluted Environment." *Southern Economic Journal* 39 (4): 544. <https://doi.org/10.2307/1056705>.
- Freeman, Rebecca, Kalina Manova, Thomas Prayer, Thomas Sampson, et al. 2022. *Unravelling deep integration: UK trade in the wake of Brexit*. Centre for Economic Performance, London School of Economics / Political Science. <https://cep.lse.ac.uk/pubs/download/dp1847.pdf>.
- Frost, James D., and Mary R. Brooks. 2018. "Short sea shipping and ferries." Chap. 20, First, edited by Jonathan Cowie and Stephen Ison, 325–347. *The Routledge Handbook of Transport Economics*. Oxon: Routledge.

- Gao, Yuan, and Peter Newman. 2018. "Beijing's Peak Car Transition: Hope for Emerging Cities in the 1.5 °C Agenda." *Urban Planning* 3 (2): 82–93. <https://doi.org/10.17645/up.v3i2.1246>.
- Gehrke, Steven R., Alison Felix, and Timothy G. Reardon. 2019. "Substitution of Ride-Hailing Services for More Sustainable Travel Options in the Greater Boston Region." *Transportation Research Record: Journal of the Transportation Research Board* 2673 (1): 438–446. <https://doi.org/10.1177/0361198118821903>.
- Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath. 2015. "Urbanization with and without industrialization." *Journal of Economic Growth* 21 (1): 35–70. <https://doi.org/10.1007/s10887-015-9121-4>.
- Gollin, Douglas, Stephen Parente, and Richard Rogerson. 2002. "The Role of Agriculture in Development." *American Economic Review* 92 (2): 160–164. <https://doi.org/10.1257/000282802320189177>.
- Goodwin, Phil. 2020. "Trends in Car Use, Travel Demand and Policy Thinking." *International Transport Forum Discussion Papers*, nos. 2020/27, <https://www.itf-oecd.org/sites/default/files/docs/trends-car-use-travel-demand-policy.pdf>.
- Goodwin, Phil, and Kurt Van Dender. 2013. "'Peak Car' — Themes and Issues." *Transport Reviews* 33 (3): 243–254. <https://doi.org/10.1080/01441647.2013.804133>.
- Google Trends. 2022. <https://www.google.com/trends>.
- Gourieroux, C., A. Monfort, and A. Trognon. 1984. "Pseudo Maximum Likelihood Methods: Applications to Poisson Models." *Econometrica* 52 (3): 701. <https://doi.org/10.2307/1913472>.

- Granger, C. W. J. 1969. "Investigating Causal Relations by Econometric Models and Cross-spectral Methods." *Econometrica* 37 (3): 424. <https://doi.org/10.2307/1912791>.
- Graziano, Alejandro G, Kyle Handley, and Nuno Limão. 2020. "Brexit Uncertainty and Trade Disintegration." *The Economic Journal* 131 (635): 1150–1185. <https://doi.org/10.1093/ej/ueaa113>.
- Grimal, Richard, Roger Collet, and Jean-Loup Madre. 2013. "Is the Stagnation of Individual Car Travel a General Phenomenon in France? A Time-Series Analysis by Zone of Residence and Standard of Living." *Transport Reviews* 33 (3): 291–309. <https://doi.org/10.1080/01441647.2013.801930>.
- Grossman, G. M., and A. B. Krueger. 1995. "Economic Growth and the Environment." *The Quarterly Journal of Economics* 110 (2): 353–377. <https://doi.org/10.2307/2118443>.
- Gunn, Simon. 2018. *The history of transport systems in the UK*. Technical report. Centre for Urban History, University of Leicester. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/761929/Historyoftransport.pdf.
- Guo, Yue, Fu Xin, Stuart J. Barnes, and Xiaotong Li. 2018. "Opportunities or threats: The rise of Online Collaborative Consumption (OCC) and its impact on new car sales." *Electronic Commerce Research and Applications* 29:133–141. <https://doi.org/10.1016/j.elerap.2018.04.005>.
- Headicar, Peter. 2013. "The Changing Spatial Distribution of the Population in England: Its Nature and Significance for 'Peak Car'." *Transport Reviews* 33 (3): 310–324. <https://doi.org/10.1080/01441647.2013.802751>.

- Henao, Alejandro, and Wesley E. Marshall. 2019. "The impact of ride hailing on parking (and vice versa)." *Journal of Transport and Land Use* 12 (1): 127–147. ISSN: 19387849, accessed October 5, 2022. <https://www.jstor.org/stable/26911261>.
- . 2018. "The impact of ride-hailing on vehicle miles traveled." *Transportation* 46 (6): 2173–2194. <https://doi.org/10.1007/s11116-018-9923-2>.
- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi. 2014. "Growth and Structural Transformation." In *Handbook of Economic Growth*, 855–941. Elsevier. <https://doi.org/10.1016/b978-0-444-53540-5.00006-9>.
- House of Commons. 2016. *HC Deb 6 September 2016, vol 614, col 38*, September. <https://hansard.parliament.uk/commons/2016-09-05/debates/580694B7-CD84-4929-8DCF-BAF46D47AF64/ExitingTheEuropeanUnion>.
- IEA. 2022a. *Energy Prices and Taxes (2022Q3 Edition)* [in en]. <https://doi.org/10.5257/IEA/EPT/2022Q3>.
- . 2022b. *Energy prices June 2022 edition: Database documentation*. Technical report. Paris: International Energy Agency (IEA), January.
- . 2022c. *Greenhouse Gas Emissions from Energy (2022 Edition)* [in en]. <https://doi.org/10.5257/IEA/CO2/2022>.
- ILO. 2020. *Employment by sex and economic activity - ILO modelled estimates*. <https://ilostat.ilo.org/data/>.
- International Energy Agency. 2020. *Energy intensity of international shipping under EEDI and Sustainable Development Scenario, 2005-30*, May. <https://www.iea.org/data-and-statistics/charts/energy-intensity-of-international-shipping-under-eedi-and-sustainable-development-scenario-2005-30>.

REFERENCES

- International Maritime Organization. 2021. *Fourth IMO Greenhouse Gas Study*. Technical report. London: International Maritime Organization. <https://wwwcdn.imo.org/localresources/en/OurWork/Environment/Documents/FourthIMOGHGStudy2020-Fullreportandannexes.pdf>.
- Jamroz, Kazimierz, and Joanna Wachnicka. 2015. "Macro Models of Vehicle Kilometres Travelled." *Logistics and Transport* 27:17–24.
- Jones, Larry, and Rodolfo Manuelli. 1995. *A Positive Model of Growth and Pollution Controls*. Technical report. <https://doi.org/10.3386/w5205>.
- Kander, A., and M. Lindmark. 2004. "Energy consumption, pollutant emissions and growth in the long run: Sweden through 200 years." *European Review of Economic History* 8 (3): 297–335. <https://doi.org/10.1017/s1361491604001224>.
- Ke, Luqi, Qing Liu, Ke Han, and Weibing Zhang. 2022. "The impact of Brexit on supply chain cost and Ro-Ro traffic at Dover." *Maritime Policy & Management*, 1–17. <https://doi.org/10.1080/03088839.2022.2158382>.
- Kehoe, Timothy J., Kim J. Ruhl, and Joseph B. Steinberg. 2018. "Global Imbalances and Structural Change in the United States." *Journal of Political Economy* 126 (2): 761–796. <https://doi.org/10.1086/696279>.
- Kenworthy, Jeff. 2013. "Decoupling urban car use and metropolitan GDP growth." *World Transport Policy and Practice* 19:8–21.
- King, David. 2020. "Zero Car Growth: A Challenge for Transport Justice." *International Transport Forum Discussion Papers*, nos. 2020/12, <https://www.itf-oecd.org/sites/default/files/docs/zero-car-growth-transport-justice.pdf>.
- Kongsamut, P., S. Rebelo, and D. Xie. 2001. "Beyond Balanced Growth." *The Review of Economic Studies* 68 (4): 869–882. <https://doi.org/10.1111/1467-937x.00193>.

- Kren, Janez, and Martina Lawless. 2022. "How has Brexit changed EU-UK trade flows?" *ESRI Working Paper 735*, <https://www.esri.ie/publications/how-has-brex-it-changed-eu-uk-trade-flows>.
- Kuhnimhof, Tobias, Jimmy Armoogum, Ralph Buehler, Joyce Dargay, Jon Martin Denstadli, and Toshiyuki Yamamoto. 2012. "Men Shape a Downward Trend in Car Use among Young Adults—Evidence from Six Industrialized Countries." *Transport Reviews* 32 (6): 761–779. <https://doi.org/10.1080/01441647.2012.736426>.
- Kuhnimhof, Tobias, Dirk Zumkeller, and Bastian Chlond. 2013. "Who Made Peak Car, and How? A Breakdown of Trends over Four Decades in Four Countries." *Transport Reviews* 33 (3): 325–342. <https://doi.org/10.1080/01441647.2013.801928>.
- Kuznets, Simon. 1973. "Modern economic growth: findings and reflections." *The American Economic Review* 63 (3): 247–258.
- Lawless, Martina, and Edgar Morgenroth. 2017. "Ireland's international trade and transport connections." *ESRI Working Paper WP573*, <https://www.esri.ie/publications/irelands-international-trade-and-transport-connections>.
- Lawless, Martina, and Edgar L. W. Morgenroth. 2019. "The product and sector level impact of a hard Brexit across the EU." *Contemporary Social Science* 14 (2): 189–207. <https://doi.org/10.1080/21582041.2018.1558276>.
- Le Vine, Scott, Peter Jones, and John Polak. 2013. "The Contribution of Benefit-in-Kind Taxation Policy in Britain to the 'Peak Car' Phenomenon." *Transport Reviews* 33 (5): 526–547. <https://doi.org/10.1080/01441647.2013.827267>.
- Lindmark, Magnus. 2002. "An EKC-pattern in historical perspective: carbon dioxide emissions, technology, fuel prices and growth in Sweden 1870–1997."

REFERENCES

- Ecological Economics* 42 (1-2): 333–347. [https://doi.org/10.1016/S0921-8009\(02\)00108-8](https://doi.org/10.1016/S0921-8009(02)00108-8).
- Lindmark, Magnus. 2004. "Patterns of historical CO2 intensity transitions among high and low-income countries." *Explorations in Economic History* 41 (4): 426–447. <https://doi.org/10.1016/j.eeh.2004.01.003>.
- Loa, Patrick, and Khandker Nurul Habib. 2021. "Examining the influence of attitudinal factors on the use of ride-hailing services in Toronto." *Transportation Research Part A: Policy and Practice* 146:13–28. <https://doi.org/10.1016/j.tra.2021.02.002>.
- López, Ramón. 1994. "The Environment as a Factor of Production: The Effects of Economic Growth and Trade Liberalization." *Journal of Environmental Economics and Management* 27 (2): 163–184. <https://doi.org/10.1006/jeem.1994.1032>.
- Lucas, Karen, and Peter Jones. 2009. *The Car in British Society*. Technical report. London: The Royal Automobile Club Foundation, April. https://www.racfoundation.org/wp-content/uploads/car_in_british_society-lucas_et_al-170409.pdf.
- Marine Institute. 2023. *Ireland's Marine Atlas*.
- Masson-Delmotte, V, P Zhai, HO Pörtner, D Roberts, J Skea, PR Shukla, A Pirani, W Moufouma-Okia, C Péan, R Pidcock, et al. 2018. *Summary for Policymakers. Global Warming of 1.5 C. An IPCC Special Report on the Impacts of Global Warming of 1.5 C Above pre-Industrial Levels., Global Warming of 1.5 C. An IPCC Special Report on the Impacts of Global Warming of 1.5 C Above Pre-Industrial Levels and Related Global Greenhouse gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change.*

- Melia, Steve, Kiron Chatterjee, and Gordon Stokes. 2018. "Is the urbanisation of young adults reducing their driving?" *Transportation Research Part A: Policy and Practice* 118:444–456. <https://doi.org/10.1016/j.tra.2018.09.021>.
- Metz, David. 2013. "Peak Car and Beyond: The Fourth Era of Travel." *Transport Reviews* 33 (3): 255–270. <https://doi.org/10.1080/01441647.2013.800615>.
- . 2015. "Peak Car in the Big City: Reducing London's transport greenhouse gas emissions." *Case Studies on Transport Policy* 3 (4): 367–371. <https://doi.org/10.1016/j.cstp.2015.05.001>.
- . 2010. "Saturation of Demand for Daily Travel." *Transport Reviews* 30 (5): 659–674. <https://doi.org/10.1080/01441640903556361>.
- Michaels, G., F. Rauch, and S. J. Redding. 2012. "Urbanization and Structural Transformation." *The Quarterly Journal of Economics* 127 (2): 535–586. <https://doi.org/10.1093/qje/qjs003>.
- Millard-Ball, Adam, and Lee Schipper. 2011. "Are We Reaching Peak Travel? Trends in Passenger Transport in Eight Industrialized Countries." *Transport Reviews* 31 (3): 357–378. <https://doi.org/10.1080/01441647.2010.518291>.
- Miramontes, Montserrat, Maximilian Pfertner, Hema Sharanya Rayaprolu, Martin Schreiner, and Gebhard Wulfhorst. 2017. "Impacts of a multimodal mobility service on travel behavior and preferences: user insights from Munich's first Mobility Station." *Transportation* 44 (6): 1325–1342. <https://doi.org/10.1007/s11116-017-9806-y>.
- Mitra, Suman Kumar, Youngeun Bae, and Stephen G. Ritchie. 2019. "Use of Ride-Hailing Services among Older Adults in the United States." *Transportation Research Record: Journal of the Transportation Research Board* 2673 (3): 700–710. <https://doi.org/10.1177/0361198119835511>.

- Mitropoulos, Lambros, Annie Kortsari, and Georgia Ayfantopoulou. 2021. "A systematic literature review of ride-sharing platforms, user factors and barriers." *European Transport Research Review* 13 (1). <https://doi.org/10.1186/s12544-021-00522-1>.
- Newman, Peter, and Jeff Kenworthy. 2011. "'Peak car use': understanding the demise of automobile dependence." *World Transport Policy and Practice* 17 (2): 31–42.
- Newman, Peter, Jeffrey Kenworthy, and Garry Glazebrook. 2013. "Peak Car Use and the Rise of Global Rail: Why This Is Happening and What It Means for Large and Small Cities." *Journal of Transportation Technologies* 03 (04): 272–287. <https://doi.org/10.4236/jtts.2013.34029>.
- Ngai, L. Rachel, and Christopher A Pissarides. 2007. "Structural Change in a Multisector Model of Growth." *American Economic Review* 97 (1): 429–443. <https://doi.org/10.1257/aer.97.1.429>.
- OECD. 2021. *Input-Output Tables 2021*. https://stats.oecd.org/Index.aspx?DataSetCode=IOTS_2021.
- . 2017. "Passenger transport." *ITF Transport Statistics*, no. database, <https://doi.org/https://doi.org/https://doi.org/10.1787/g2g5557f-en>. <https://www.oecd-ilibrary.org/content/data/g2g5557f-en>.
- Office for National Statistics. 2020. *Population estimates and components of population change. Detailed time series 2001 to 2019: United Kingdom, local authorities, sex and age*, June.
- Office of Rail and Road. 2022. *Table 7182: Average change in fares by ticket type, Great Britain, 2004 to 2022*, June. <https://dataportal.orr.gov.uk/statistics/finance/rail-fares/>.

- Ordnance Survey. 2020. *OS Open Roads*. <https://osdatahub.os.uk/downloads/open/OpenRoads>.
- Palgrave Macmillan Ltd. 2013a. "F1 Length of Railway Line Open (in kilometres)." In *International Historical Statistics*, 4423–4434. Palgrave Macmillan UK. https://doi.org/10.1057/978-1-137-30568-8_442.
- . 2013b. "F6 Motor Vehicles in Use (thousands)." In *International Historical Statistics*, edited by Palgrave Macmillan Ltd, 4522–4538. Palgrave Macmillan UK. https://doi.org/10.1057/978-1-137-30568-8_447.
- Paundra, Joshua, Jan van Dalen, Laurens Rook, and Wolfgang Ketter. 2020. "Ridesharing platform entry effects on ownership-based consumption in Indonesia." *Journal of Cleaner Production* 265:121535. <https://doi.org/10.1016/j.jclepro.2020.121535>.
- Puentes, Robert, and Adie Tomer. 2008. *The Road Less Travelled: An Analysis of Vehicle Miles Traveled Trends in the U.S.* Technical report. Washington D.C.: Brookings Institution, December.
- Rogerson, Richard. 2008. "Structural Transformation and the Deterioration of European Labor Market Outcomes." *Journal of Political Economy* 116 (2): 235–259. <https://doi.org/10.1086/588029>.
- Rose, John M., and David A. Hensher. 2013. "Demand for taxi services: new elasticity evidence." *Transportation* 41 (4): 717–743. <https://doi.org/10.1007/s11116-013-9482-5>.
- Santos Silva, J.M.C., and Silvana Tenreyro. 2010. "On the existence of the maximum likelihood estimates in Poisson regression." *Economics Letters* 107, no. 2 (May): 310–312. ISSN: 0165-1765. <https://doi.org/10.1016/j.econlet.2010.02.020>.

- Santos Silva, J.M.C., and Silvana Tenreyro. 2006. "The Log of Gravity." *The Review of Economics and Statistics* 88 (4): 641–658. ISSN: 00346535, 15309142, accessed February 28, 2023. <http://www.jstor.org/stable/40043025>.
- Schipper, Lee. 2009. "Fuel economy, vehicle use and other factors affecting CO2 emissions from transport." *Energy Policy* 37 (10): 3711–3713. <https://doi.org/10.1016/j.enpol.2009.07.028>.
- Shah, Syed Hamad Hassan, Saleha Noor, Shen Lei, Atif Saleem Butt, and Muhammad Ali. 2021. "Role of privacy/safety risk and trust on the development of prosumption and value co-creation under the sharing economy: a moderated mediation model." *Information Technology for Development* 27 (4): 718–735. <https://doi.org/10.1080/02681102.2021.1877604>.
- Shi, Kunbo, Rui Shao, Jonas De Vos, Long Cheng, and Frank Witlox. 2021. "The influence of ride-hailing on travel frequency and mode choice." *Transportation Research Part D: Transport and Environment* 101:103125. <https://doi.org/10.1016/j.trd.2021.103125>.
- Shi, Xiaoyang, Zhengquan Li, and Enjun Xia. 2021. "The impact of ride-hailing and shared bikes on public transit: Moderating effect of the legitimacy." *Research in Transportation Economics* 85:100870. <https://doi.org/10.1016/j.retrec.2020.100870>.
- SHS. 2020. *Scottish Household Survey: Methodology and Fieldwork Outcomes 2019*. Technical report. Edinburgh: The Scottish Government, December. <https://www.gov.scot/binaries/content/documents/govscot/publications/statistics/2020/12/scottish-household-survey-2019-methodology-fieldwork-outcomes/documents/methodology-fieldwork-outcomes-2019/methodology-fieldwork-outcomes-2019/govscot:document/methodology-fieldwork-outcomes-2019.pdf>.

- Sikder, Sujan. 2019. "Who Uses Ride-Hailing Services in the United States?" *Transportation Research Record: Journal of the Transportation Research Board* 2673 (12): 40–54. <https://doi.org/10.1177/0361198119859302>.
- Solow, Robert M. 1956. "A Contribution to the Theory of Economic Growth." *The Quarterly Journal of Economics* 70 (1): 65. <https://doi.org/10.2307/1884513>.
- . 1973. "Is the End of the World at Hand?" *Challenge* 16 (1): 39–50. <https://doi.org/10.1080/05775132.1973.11469961>.
- Stapleton, Lee, Steve Sorrell, and Tim Schwanen. 2017. "Peak car and increasing rebound: A closer look at car travel trends in Great Britain." *Transportation Research Part D: Transport and Environment* 53:217–233. <https://doi.org/10.1016/j.trd.2017.03.025>.
- Stefanski, Radoslaw. 2017. "Dirty little secrets: Inferring fossil-fuel subsidies from patterns in emission intensities." *University of St Andrews School of Economics and Finance Discussion Paper*, no. 1705 (March). https://research-repository.st-andrews.ac.uk/bitstream/handle/10023/10412/Stefanski_EconDiscPap_1705.pdf?sequence=1.
- . 2013. "On the mechanics of the Green Solow Model." *OxCarre Research Paper* 47.
- . 2014. "Structural transformation and the oil price." *Review of Economic Dynamics* 17 (3): 484–504. <https://doi.org/10.1016/j.red.2013.09.006>.
- Stern, David I. 2004. "The Rise and Fall of the Environmental Kuznets Curve." *World Development* 32 (8): 1419–1439. <https://doi.org/10.1016/j.worlddev.2004.03.004>.

- Stern, David I., and Michael S. Common. 2001. "Is There an Environmental Kuznets Curve for Sulfur?" *Journal of Environmental Economics and Management* 41 (2): 162–178. <https://doi.org/10.1006/jeem.2000.1132>.
- Stiglic, Mitja, Niels Agatz, Martin Savelsbergh, and Mirko Gradisar. 2018. "Enhancing urban mobility: Integrating ride-sharing and public transit." *Computers & Operations Research* 90:12–21. <https://doi.org/10.1016/j.cor.2017.08.016>.
- Stiglitz, Joseph. 1974. "Growth with Exhaustible Natural Resources: Efficient and Optimal Growth Paths." *The Review of Economic Studies* 41:123. <https://doi.org/10.2307/2296377>.
- Stokey, Nancy L. 1998. "Are There Limits to Growth?" *International Economic Review* 39 (1): 1. <https://doi.org/10.2307/2527228>.
- Świącki, Tomasz. 2017. "Determinants of structural change." *Review of Economic Dynamics* 24:95–131. <https://doi.org/10.1016/j.red.2017.01.007>.
- Tang, Bao-Jun, Xiao-Yi Li, Biying Yu, and Yi-Ming Wei. 2019. "How app-based ride-hailing services influence travel behavior: An empirical study from China." *International Journal of Sustainable Transportation* 14 (7): 554–568. <https://doi.org/10.1080/15568318.2019.1584932>.
- Tanner, J. C. 1978. "Long-Term Forecasting of Vehicle Ownership and Road Traffic." *Journal of the Royal Statistical Society. Series A (General)* 141 (1): 14. <https://doi.org/10.2307/2344775>.
- The World Bank. 2019. *Age dependency ratio, old (percentage of working-age population)*. https://data.worldbank.org/indicator/SP.POP.DPND.OL?name_desc=true.

-
- . 2008. *Global Purchasing Power Parities and Real Expenditures: 2005 International Comparison Programme*. Technical report. Washington D.C.: The World Bank. <https://thedocs.worldbank.org/en/doc/982121487105148964-0050022017/original/2005ICPReportFinalwithNewAppG.pdf>.
- . 2005. *International Comparison Programme (ICP) 2005*. [https://databank.worldbank.org/source/international-comparison-program-\(icp\)-2005](https://databank.worldbank.org/source/international-comparison-program-(icp)-2005).
- . 2023. *Land area (sq. km)*, November. <https://data.worldbank.org/indicator/AG.LND.TOTL.K2>.
- . 2022. *School enrollment, tertiary (percentage of gross)*, June. https://data.worldbank.org/indicator/SE.TER.ENRR?name_desc=true.
- Tirachini, Alejandro. 2019. "Ride-hailing, travel behaviour and sustainable mobility: an international review." *Transportation* 47 (4): 2011–2047. <https://doi.org/10.1007/s11116-019-10070-2>.
- Tirachini, Alejandro, and Andres Gomez-Lobo. 2019. "Does ride-hailing increase or decrease vehicle kilometers traveled (VKT)? A simulation approach for Santiago de Chile." *International Journal of Sustainable Transportation* 14 (3): 187–204. <https://doi.org/10.1080/15568318.2018.1539146>.
- Tirachini, Alejandro, and Mariana del Río. 2019. "Ride-hailing in Santiago de Chile: Users' characterisation and effects on travel behaviour." *Transport Policy* 82:46–57. <https://doi.org/10.1016/j.tranpol.2019.07.008>.
- Tol, Richard S. J., Stephen W. Pacala, and Robert Socolow. 2006. "Understanding Long-Term Energy Use and Carbon Dioxide Emissions in the USA." *SSRN Electronic Journal*, <https://doi.org/10.2139/ssrn.927741>.

- Transport Scotland. 2020. *Scottish Transport Statistics 2019 Edition: A National Statistics Publication for Scotland*. Technical report 38. Edinburgh: National Statistics, March. <https://www.transport.gov.scot/publication/scottish-transport-statistics-no-38-2019-edition/>.
- UK Department for Business, Energy and Industrial Strategy. 2022. *Final UK greenhouse gas emissions national statistics: 1990 to 2020*, February. <https://www.gov.uk/government/statistics/final-uk-greenhouse-gas-emissions-national-statistics-1990-to-2020>.
- UK Department for Energy Security and Net Zero. 2023a. *2022 UK greenhouse gas emissions, provisional figures*. Technical report. London: UK Department for Energy Security and Net Zero, March. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1147372/2022_Provisional_emissions_statistics_report.pdf.
- . 2023b. *Energy prices: road fuels and other petroleum products*, February. <https://www.gov.uk/government/statistical-data-sets/oil-and-petroleum-products-monthly-statistics>.
- . 2023c. *Typical retail prices of petroleum products and crude oil price index*, February. <https://www.gov.uk/government/statistical-data-sets/oil-and-petroleum-products-monthly-statistics>.
- UK Department for Transport. 2021a. *All vehicles VEH0103: Licensed vehicles by tax class: Great Britain and United Kingdom*. <https://www.gov.uk/government/statistical-data-sets/all-vehicles-veh01>.
- . 2023. *Costs, fares and revenue for local bus services: at current prices (not adjusted for inflation)*, January. <https://www.gov.uk/government/statistical-data-sets/bus-statistics-data-tables#costs-fares-and-revenue-bus04>.

-
- . 2013. *Modal comparisons*, December. <https://www.gov.uk/government/statistical-data-sets/tsgb01-modal-comparisons>.
- . 2020a. *National Public Transport Access Nodes (NaPTAN)*. <https://www.data.gov.uk/dataset/ff93ffc1-6656-47d8-9155-85ea0b8f2251/national-public-transport-access-nodes-naptan>.
- . 2022. *Road Safety Data*, November. <https://www.data.gov.uk/dataset/cb7ae6f0-4be6-4935-9277-47e5ce24a11f/road-safety-data>.
- . 2021b. *Road traffic statistics TRA0201: Road traffic (vehicle kilometres) by vehicle type in Great Britain*. <https://www.gov.uk/government/statistical-data-sets/road-traffic-statistics-tra#traffic-volume-in-kilometres-tra02>.
- . 2020b. *Total road length (kilometres) by road type and local authority in Great Britain*. <https://www.gov.uk/government/statistical-data-sets/road-length-statistics-rdl>.
- UN. 2008. *International Standard Industrial Classification of All Economic Activities (ISIC), Rev. 4*. Statistical Papers Series M 4. New York: United Nations. https://unstats.un.org/unsd/demographic-social/census/documents/isic_rev4.pdf.
- . 2021a. *National accounts main aggregates database*. <https://unstats.un.org/unsd/snaama/>.
- . 2021b. *Review of Maritime Transport*. Technical report. Geneva: United Nations, November. https://unctad.org/system/files/official-document/rmt2021_en_0.pdf.
- . 2021c. *Sustainable transport, sustainable development: Interagency report for second Global Sustainable Transport Conference*. Technical report. New York:

- United Nations. https://sdgs.un.org/sites/default/files/2021-10/TransportationReport2021_FullReport_Digital.pdf.
- Uy, Timothy, Kei-Mu Yi, and Jing Zhang. 2013. "Structural change in an open economy." *Journal of Monetary Economics* 60 (6): 667–682. <https://doi.org/10.1016/j.jmoneco.2013.06.002>.
- Van Acker, Veronique, and Frank Witlox. 2010. "Car ownership as a mediating variable in car travel behaviour research using a structural equation modelling approach to identify its dual relationship." *Journal of Transport Geography* 18 (1): 65–74. <https://doi.org/10.1016/j.jtrangeo.2009.05.006>.
- van Wee, Bert. 2015. "Peak car: The first signs of a shift towards ICT-based activities replacing travel? A discussion paper." *Transport Policy* 42:1–3. <https://doi.org/10.1016/j.tranpol.2015.04.002>.
- Vandenbussche, Hylke, William Connell, and Wouter Simons. 2022. "Global value chains, trade shocks and jobs: An application to Brexit." *The World Economy* 45 (8): 2338–2369. <https://doi.org/10.1111/twec.13259>.
- Vega, Amaya, Maria Feo-Valero, and Raquel Espino-Espino. 2018. "The potential impact of Brexit on Ireland's demand for shipping services to continental Europe." *Transport Policy* 71:1–13. <https://doi.org/10.1016/j.tranpol.2018.07.005>.
- Walker, Nigel. 2021. "Brexit timeline: events leading to the UK's exit from the European Union." *House of Commons Library Briefing Paper* (London), no. 7960 (January). <https://researchbriefings.files.parliament.uk/documents/CBP-7960/CBP-7960.pdf>.

- Wang, Sicheng, and Robert B. Noland. 2021. "Variation in ride-hailing trips in Chengdu, China." *Transportation Research Part D: Transport and Environment* 90:102596. <https://doi.org/10.1016/j.trd.2020.102596>.
- Wang, Yanghao, Wei Shi, and Zhenhua Chen. 2021. "Impact of ride-hailing usage on vehicle ownership in the United States." *Transportation Research Part D: Transport and Environment* 101:103085. <https://doi.org/10.1016/j.trd.2021.103085>.
- Ward, Jacob W., Jeremy J. Michalek, Inês L. Azevedo, Constantine Samaras, and Pedro Ferreira. 2019. "Effects of on-demand ridesourcing on vehicle ownership, fuel consumption, vehicle miles traveled, and emissions per capita in U.S. States." *Transportation Research Part C: Emerging Technologies* 108:289–301. <https://doi.org/10.1016/j.trc.2019.07.026>.
- Webb, Jeremy. 2019. "The future of transport: Literature review and overview." *Economic Analysis and Policy* 61:1–6. <https://doi.org/10.1016/j.eap.2019.01.002>.
- Wittwer, Rico, Regine Gerike, and Stefan Hubrich. 2019. "Peak-Car Phenomenon Revisited for Urban Areas: Microdata Analysis of Household Travel Surveys from Five European Capital Cities." *Transportation Research Record* 2673 (3): 686–699. <https://doi.org/10.1177/0361198119835509>.
- Wolfers, Justin. 2006. "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results." *American Economic Review* 96 (5): 1802–1820. <https://doi.org/10.1257/aer.96.5.1802>.
- Wooldridge, Jeffrey M. 2008. *Econometric Analysis of Cross Section and Panel Data*. The MIT Press. ISBN: 9780262232586.

REFERENCES

- Young, Mischa, and Steven Farber. 2019. "The who, why, and when of Uber and other ride-hailing trips: An examination of a large sample household travel survey." *Transportation Research Part A: Policy and Practice* 119:383–392. <https://doi.org/10.1016/j.tra.2018.11.018>.
- Zhong, Jun, Yan Lin, and Siqi Yang. 2020. "The Impact of Ride-Hailing Services on Private Car Use in Urban Areas: An Examination in Chinese Cities." Edited by Michela Le Pira. *Journal of Advanced Transportation* 2020:1–15. <https://doi.org/10.1155/2020/8831674>.
- Zhong, Jun, Huan Zhou, Yan Lin, and Fangxiao Ren. 2022. "The impact of ride-hailing services on the use of traditional taxis: Evidence from Chinese urban panel data." *IET Intelligent Transport Systems*, <https://doi.org/10.1049/itr2.12237>.