Title
Essays on the Distributional Effects of Carbon Taxes and Tipping

Author
Cross, Jeffrey Robert

Publication Date
2021

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Essays on the Distributional Effects of Carbon Taxes and Tipping

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

Jeffrey Robert Cross

Committee in charge:

Professor Olivier Deschenes, Chair
Professor Kelsey Jack
Professor Kyle Meng

June 2021

The Dissertation of Jeffrey Robert Cross is approved.

Professor Kelsey Jack

Professor Kyle Meng

Professor Olivier Deschenes, Committee Chair

June 2021

Essays on the Distributional Effects of Carbon Taxes and Tipping

To my late grandfather Stuart and my parents, Bob and Linda.

Acknowledgements

This has been an incredible journey that would not be possible without the help of many people. First and foremost, I would like to acknowledge my committee, Olivier Deschenes, Kyle Meng, and Kelsey Jack. Their encouragement and mentorship taught me how to ask questions and find answers. This dissertation would not have been possible without lessons that I learned from each of them.

I would also like to thank my colleagues that have had a remarkable influence on my research and experience at UCSB. I appreciate my officemates Hakan, Yongwook, and Guangli, and the rest of my cohort, who made this challenging process significantly easier to overcome. I am also incredibly indebted to the members of the Environmental Reading Group, particularly Danae, Dan, Juliana, Sahaab, Emily, Di, Chris, Jacob, and Vincent. Our weekly meetings provided a healthy dose of friendship and feedback that significantly improved the last few years of this journey.

Furthermore, I am forever grateful to my friends and family. Specifically, I want to thank my parents, Bob and Linda, who nurtured my curiosity from a young age. I would not be the person or researcher that I am today without their constant support, even if that meant frequent stops at random book stores on our road trips. In addition, I would like to thank my lifelong friends Sean Rodich, Kevin Graves, Austin Crochetiere, and Austin Acres. They have taught me that being driven and enjoying life are not mutually exclusive. Lastly, I am incredibly indebted to my partner, Claire Hoch-Frohman. These last few years have been some of the most difficult of my life. It is only with her unwavering support that I was able to reach this point. This dissertation is our success.

Curriculum Vitæ

Jeffrey Robert Cross

Education

2021	Ph.D. in Economics, University of California, Santa Barbara.
2016	M.A. in Economics, University of California, Santa Barbara.
2014	B.A. in Economics and Political Science, Texas Christian University

Research Fields

Environmental Economics, Energy Economics

Working Papers

"Who Bears the Load? Carbon Taxes and Electricity Markets"
"Paying for Integers" (*with Guangli Zhang*)
"Asymmetric Pass-through from Carbon Taxes to Electricity Prices"
"Do Fans Impact Sports Outcomes? A COVID-19 Natural Experiment" (*with Richard Uhrig*)

Works in Progress

"Equity Implications of Local Markets and Carbon Price"
"Cutting Fewer Corners: Taxicabs and Tipping Policies" (*with Guangli Zhang*)
"Linked accounts, more revenue? Collecting municipal debts through prepaid electricity meters" (*with Kelsey Jack and Anders Jensen*)

Research Experience

(RA) Professor Kelsey Jack, UCSB	Spring 2018-2020
(RA) Professor Olivier Deschenes, UC Santa Barbara	2020
(RA) Professor Kiril Tochkov, Texas Christian University	2014-2015
(RA) Professor Vanessa Bouche, Texas Christian University	2014-2015

Teaching Experience

(TA) Intro Econometrics, UC, Santa Barbara.	Fall 2020
	Winter 2021
	Spring 2021
(TA) International Finance, UC, Santa Barbara.	Winter 2020
	Spring 2020

(TA) Intermediate Macro Theory, UC, Santa Barbara.	Fall 2019
	Spring 2019
	Spring 2017
(TA) Intermediate Micro Theory,UC, Santa Barbara.	Summer 2019
	Spring 2018
	Winter 2018
	Summer 2017
(TA) Intro to Game Theory, UC, Santa Barbara.	Fall 2018
(TA) Principles of Microeconomics, UC, Santa Barbara.	Fall 2016
	Winter 2017

Other Employment

Research Assistant Thomas Rooney LLC Fort Worth, Texas	2014-2015
Student Assistant California Air Resources Board (Carl Moyer Program) El Monte, California	2012

Presentations

Federal Reserve Bank of Dallas Undergraduate Research Conference, 2014.
Southern California Graduate Conference in Applied Economics (SOCAE), 2020.

Awards and Fellowships

Outstanding Teaching Assistant Award. UC Santa Barbara	Spring 2019
White Family Graduate Fellowship, UC Santa Barbara	Winter 2019
Andron Fellowship in Economics. UC Santa Barbara	Summer 2018
Academic Senate Outstanding Teaching Assistant Award Nominee	2020
Graduate Student Association Excellence in Teaching Nominee	2017, 2020
SOCAE Best Paper Award	2020

Service

Professional Development Committee	2019-2020
Graduate Student Mentor	2019-2021
Graduate Student Environmental Reading Group (founder)	2018-2021

Abstract

Essays on the Distributional Effects of Carbon Taxes and Tipping

by

Jeffrey Robert Cross

This dissertation consists of three works that estimate the distributional effects of changes to market conditions. Each essay utilizes different techniques to identify behavioral responses from consumers and firms, which then plays a key role in the distribution of consumer and producer surplus.

In the first essay, I study how a carbon tax differentially affects the welfare of electricity producers and consumers, also known as incidence. In so doing, I develop a new framework to estimate the incidence of input taxes that accounts for incomplete pass-through to retail prices, imperfect competition, and heterogeneity in marginal costs of producers. Leveraging exogenous variation in the level of the Australian carbon tax, I then apply this framework to the context of the National Electricity Market in Australia. I find that 95 to 97 percent of the welfare cost from changes in the carbon tax falls on electricity consumers. The large burden borne by consumers is driven by full pass-through of increases in wholesale electricity prices to retail electricity prices in combination with negligible decreases in aggregate profits for electricity producers. Simulations show that incorporating heterogeneous firms is particularly important when market demand is elastic.

The second essay, joint with Guangli Zhang, asks: Do customers respond differently to tip suggestions based on whether or not the suggested tip amount is an integer and, if so, what does this reveal about human behavior? With the advent of cashless payment systems, customers are increasingly being presented convenient, low-stake tip suggestions

following purchases. Despite the rising frequency of these interactions, we still know little about the preferences underlying tipping behavior. Previous research, for example, has documented that customers tend to tip integer amounts, but it is unclear if this is due to smaller cognitive costs associated with tipping an integer amount or direct utility benefits from integer tips. Utilizing plausibly exogenous variation in the occurrence of integer tip suggestion in New York City taxi rides, we find that passengers presented with an integer tip suggestion are 21 percentage points more likely to tip that amount. In our setting where the suggested tip percent is larger than the average passengers give, this leads to a 0.6 percentage point increase in tip rates and an approximately 2.38 million dollars transfer from riders to drivers as a result of a 2012 rate fare change that increased the probability of integer tip suggestions. Using a theoretical model for customer tipping behavior, we find that these results are consistent with direct utility benefits from tipping integer amounts.

In the third essay, I return to the context of the Australian carbon tax to estimate if wholesale electricity prices respond differently to the implementation of the carbon tax compared to when it is repealed. I find that when the carbon tax is implemented prices rise by more than double the amount they decrease when the carbon tax is repealed. I then utilize one of the key advantages of wholesale electricity markets – detailed data on the offer (supply) curve that is used to set prices. I identify that that the pattern evident in prices directly reflects asymmetric changes in how plants respond to the implementation and repeal of the carbon tax. When the carbon tax is implemented, plants that could potentially set prices increase their willingness-to-accept by significantly more than they decrease them when the carbon tax is repealed. These findings provide the first direct evidence that asymmetric price changes are driven by asymmetric supply-side responses.

Contents

Chapter 1

Who Bears the Load? Carbon Taxes and Electricity Markets

1.1 Introduction

Incidence, defined as the relative change in consumer to producer surplus, is a central tool for assessing the distributional consequences of taxes. It has played a crucial role in our understanding of who bears the burden from excise taxes on a wide variety of goods such as cigarettes (Evans, Ringel, and Stech 1999) and gasoline (Doyle and Samphantharak 2008). Many observed policies, however, do not directly change the cost of outputs, but instead change the cost of inputs. For example, carbon prices increase the cost of fossil-fuels and minimum wage policies increase the cost of labor. Political debates about the extent to which these policies hurt consumers or firms can make them difficult to enact, even when there is a strong consensus about the efficiency of the policy, as is the case with a carbon tax.[1] Despite the political focus on the incidence of policies that

[1]There are over 3500 signatories of the Economists' Statement on Carbon Dividends that says "a carbon tax offers the most cost-effective lever to reduce carbon emissions at the scale and speed that is necessary". On the other hand, the Heritage foundation review of climate regulation states that any

change input prices, there is relatively little research on it, particularly in imperfectly competitive markets.

This paper uses variation in the short-lived Australian carbon tax to estimate the burden borne by consumers relative to producers in an imperfectly competitive electricity market. In the process, I offer a novel partial equilibrium approach to analyzing the incidence of input taxes that incorporates heterogeneous producers, various forms of imperfect competition, and heterogeneous price changes in the intermediary (wholesale) and output (retail) markets. Compared to the case where producers are assumed to be homogeneous (Ganapati, Shapiro, and Walker 2020), I find that incidence depends crucially on the interaction between individual firm changes in production in response to the carbon tax and firm-level markups. Leveraging plausibly exogenous changes in the Australian carbon tax over time and detailed firm-level data on production technologies, I estimate the impact of the carbon tax on production and prices. Incorporating these estimates into the incidence framework, I find that electricity consumers bear nearly all of the tax burden due to negligible changes in markups, highly inelastic demand, and increases in retail electricity prices that reflect the change in average marginal costs of electricity generators.

The partial equilibrium approach I develop has several key features. First, in the context of carbon policies it aligns with the theoretical efficiency gains of a carbon price, which primarily comes from heterogeneity in producers. It is only in the case where producers have heterogeneous marginal abatement costs that there are significant efficiency gains relative to emissions standards. Second, it matches many settings where carbon prices have been introduced, which often feature markets with high fixed costs and imperfect competition. The approach that I develop in this paper allows for differences

absorbing of costs would shrink profit margins and consumers will be "hit repeatedly with higher prices" (Loris and Jolevski 2014).

in production technologies between firms in imperfectly competitive markets, while still remaining empirically tractable. Lastly, although I focus on the context of a carbon tax and electricity markets, this approach can also be applied to other policies that change the costs of inputs, such as minimum wage laws, in imperfectly competitive markets with heterogeneous producers.

Crucial to the credible identification of the parameters in my analysis is the exogenous variation in the carbon price over time in the Australian setting. By using changes in a carbon tax over time, I circumvent many of the key concerns associated with estimating the effect of carbon prices on wholesale electricity markets. Electricity generators tend to account for a large fraction of total emissions in carbon markets with tradable permits, which would likely make the permit price endogenous.[2] To overcome this concern, previous research has focused on small electricity markets within a large emissions trading scheme, such as Spain and the EU ETS (Fabra and Reguant 2014), but it is still possible that the emissions price is endogenous due to shared macroeconomics trends.[3] Changes in the Australian carbon tax, however, are a result of a predetermined schedule of carbon prices and thus do not respond to changes in the electricity market or macroeconomic trends. Controlling for other factors that impact the electricity market equilibrium such as seasonality, weather, and global fuel prices, I identify the effect of changes in the carbon tax on electricity market outcomes, even in a context where 12 of the top 15 emitters are electricity generators.

Incorporating estimates of key parameters into the partial equilibrium framework, I have two main empirical findings. First, customers bear 95 to 97 percent of the total burden from changes in the carbon tax, with firms bearing the rest. Separately estimating

[2]For example, from 2015-2017, California's electricity generators accounted for over 60 percent of total permits surrendered under the California cap-and-trade, AB 32.

[3]Fabra and Reguant (2014) alleviate this concern through a rich set of controls including month of sample fixed effects in their estimation strategy.

the impact of the carbon tax on retail electricity prices has a negligible effect on incidence calculations, with customers still bearing over 95 percent of the burden. Compared to United States industries analyzed by Ganapati et al. (2020), the burden borne by customers is 31 to 64 percent larger.

The second empirical analysis in this paper focuses on disentangling potential mechanisms behind the large burden borne by consumers. Intuition would suggest that this is purely a result of highly inelastic demand, but in the context of an input tax and an electricity market with heterogeneous producers this is not the case. Electricity markets are cleared with an auction where the willingness to accept of the last producer determines the market-clearing price. When there is a heterogeneous cost shock, such as a carbon price, the impact on producer surplus depends crucially on how these price-setting (marginal) plants are impacted relative to all other producers. How these plants adjust their markups depends on their own change in marginal costs and any adjustments to markups. To determine whether increases in markups are leading to a large burden borne by consumers, I use an augmented local linear regression discontinuity design (Hausman and Rapson 2018) and a change in the carbon tax to empirically estimate markup adjustments. I find no evidence of markup adjustments as price-setting plants increased their willingness to accept by approximately the change in marginal costs. These results highlight that the large burden borne by consumers relative to producers is not a result of upward adjustments to markups, but rather similarities in the average changes in marginal costs between price-setting and all other plants that led to an increase in prices that covered almost all of the cost increase.

To explore the importance of applying the framework developed in this paper to contexts besides electricity markets, I compare estimates of incidence under the assumption of homogeneous and heterogeneous producers in an imperfectly competitive market. The major challenge with this approach, however, is that I only observe the responses of electricity generators in the observed state of the world, which has highly inelastic demand. To simulate incidence under an alternative demand elasticity, I need to make an assumption about which producers would be most impacted by changes in the demand elasticity. Whether I assume that all producers are impacted equally or only those that decrease production per dollar of the carbon tax when demand is inelastic, I find the same pattern; when demand is more elastic, firm-level estimates of markups become relatively more important. At a demand elasticity of 4, roughly equivalent to the demand elasticity for concrete (Ganapati et al. 2020), the percent difference between estimates of incidence using heterogeneous and homogeneous frameworks ranges from 9 to 40 percent depending on the assumed changes in producer quantities. Given the relatively weak relationship between markups and changes in quantity that I find in my context, this likely understates the importance of applying the partial equilibrium approach developed in this paper for other markets.

A primary contribution of this paper is the development of a framework for estimating incidence in the context of input taxes in imperfectly competitive markets with *heterogeneous* producers. In doing so, this study contributes to an extensive literature that theoretically analyzes incidence. Many of these studies historically (Marshall 1890) and in recent theoretical work (Weyl and Fabinger 2013) focus on incidence of taxes on a firm's outputs. A notable exception is recent work by Ganapati et al. (2020), who develop a framework for estimating incidence in imperfectly competitive markets with input taxes and homogeneous producers. In contrast to their work, however, the framework I develop does not assume producers are homogeneous, but instead allows for differences in marginal costs and changes in marginal costs from the input tax. Harasztosi and Lindner (2019) estimate incidence from a minimum wage change in Hungary without this framework using detailed firm-level data on employment, revenue, costs, and profitability over time. The application of this paper, however, shows that incidence of an input tax can be estimated absent detailed data on input usage or profits.

This paper is also the first to empirically estimate the burden borne by electricity consumers relative to electricity producers for a carbon tax. In doing so, I contribute to the literature that estimates incidence of a carbon tax for particular sectors. Andersson (2019) finds a significant effect of carbon taxes on emissions and full pass-through of the carbon tax to gasoline prices. Ganapati et al. (2020) utilize regional variation in energy prices to estimate the welfare loss for consumers relative to producers in the United States industries of boxes, bread, cement, concrete, gasoline, and plywood. The importance of electricity markets for estimating the distributional effects of a carbon tax, however, has been highlighted by previous work that used a general equilibrium approach (Goulder, Hafstead, Kim, and Long 2019) or input-output matrices (Grainger and Kolstad 2010) to estimate the regressivity or progressivity of a hypothetical carbon tax. My finding that almost all of the increase in costs ends up in retail electricity prices provides some support for the typical assumption that pass-through is complete in input-output approaches.

By estimating the components of incidence, this work also contributes to two different strands of literature that study the impact of carbon prices on electricity markets. The first strand is an extensive literature estimating the impacts of carbon prices on wholesale electricity prices. Much of this literature focuses on quantifying the relationship between emissions costs and prices. Fabra and Reguant (2014), for example, study the Spanish wholesale electricity market and find that changes in marginal emissions costs are passed on fully to wholesale electricity prices. Other papers have studied the relationship between changes in the average marginal emissions cost and wholesale electricity prices, even in the context of Australia (Nazifi, Trück, and Zhu 2017). The approach of this paper, however, highlights that estimating the direct effect of the carbon tax on electricity prices is crucial for estimating incidence.[4] Alternatively, a second strand of the

[4]To discuss competition, one would want to relate changes in marginal costs to changes in prices.

literature examines the impact of carbon prices on electricity generator returns. In contrast to the previous literature that estimates the change in profits through simulations (Sijm, Neuhoff, and Chen 2006) or stock prices (Bushnell, Chong, and Mansur 2013), I empirically estimate the change in profits using estimates of changes in prices, quantities, and markups.

Lastly, this paper contributes to a small literature estimating the impact of cost shocks on markups in wholesale electricity markets. Fabra and Reguant (2014) find that the high level of correlation across firms mitigated incentives to adjust markups in the context of the Spanish wholesale electricity market. Analyzing a shock in the natural gas price on New England's electricity market, Kim (2021) finds a large effect on markups, particularly for largely impacted firms. However, each of these papers analyzes the effect on markups using a simulation where all other factors except for the shock are held constant. These simulations use a first-order approach that implements a small cost perturbation and estimates changes in markups with the first-order condition. This approach estimates how firms should adjust markups, but electricity generators exhibit significant strategic limitations, even in highly sophisticated markets (Hortaçsu, Luco, Puller, and Zhu 2019). With this in mind, I use an augmented local linear regression discontinuity design in combination with an increase in the carbon tax to provide the first causal estimates a heterogeneous cost shock has on markups.

The paper proceeds as follows. Section 1.2 describes the Australian context. In Section 1.3, I derive the formula for incidence. In Section 1.4, I describe the empirical framework and data. Section 1.5 presents the key empirical results of this paper including pass-through, incidence, and markups. Section 1.6 presents the results of simulations that highlight the importance of the theoretical framework and Section 1.7 concludes.

Since I am relating changes in the carbon price to changes in the price, a pass-through rate larger than 1 is not sufficient to reject perfect competition. Intuitively, this is because the change in marginal costs per dollar change in the carbon tax could be larger than 1.

1.2 Background

In this section, I provide background on Australia's carbon price, wholesale electricity market, and retail electricity market.

1.2.1 Carbon Pricing Mechanism

Under the *Clean Energy Act 2011*, starting on July 1, 2012, all entities that emit over 25,000 tonnes of carbon dioxide equivalent (CO_2e) per year, and which were not in the transport or agriculture sectors, were required to obtain emissions permits (carbon units) for each tonne of carbon dioxide equivalent (CO_2e) that they emit.[5] For the entirety of the program's existence, the price of permits was fixed and the quantity unlimited.[6] Importantly, these two features of the policy meant that it behaved identically to a carbon tax and so I will refer to it as such.

Figure 1.1 shows the carbon price history from implementation, July 1, 2012, to repeal, July 1, 2014.[7] When implemented, the price of carbon was set at $23 Australian dollars (AUD) per tonne of CO_2e. The carbon price rose to 24.15 (AUD) on July 1, 2013 where it remained until it was legally revoked on July 1, 2014. All three of the changes in the carbon price were either predetermined or foreshadowed. Both increases in the carbon price in 2012 and 2013 were laid out in the *Clean Energy Act 2011*, which was passed on November 8, 2011. Even the repeal of the carbon price was, to some extent, foreshadowed following the 2013 Federal Election, which took place on September 7, 2013. The centre-right Liberal/National Coalition defeated the incumbent Labor party,

[5] All but the smallest thermal electricity generators meet the minimum eligibility requirements. In fact, the top 8 payers of the carbon tax in the 2012-2013 fiscal year were all electricity generators according to LEPID data.

[6] The ultimate intent was to transition to a flexible cap-and-trade (ETS), but the *Clean Energy Act 2011* was repealed before this could happen.

[7] Following the 2013 Federal Election it was known that the Labor government would likely repeal in the 2014-2015 fiscal year.

Figure 1.1: Carbon Price History

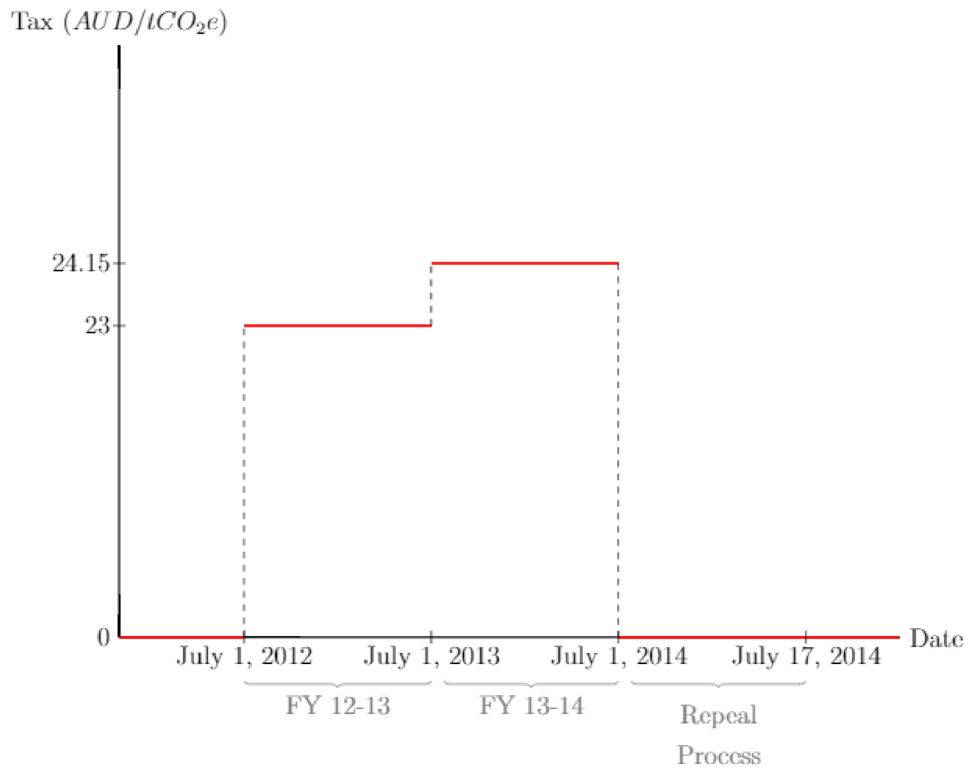

Note: This figure shows the nominal price of carbon emissions over time (carbon tax). FY stands for fiscal year.

which signaled that the *Clean Energy Act 2011* was likely to be repealed after the Senate changed on July 1, 2014. The process to repeal following the Senate change took a little over two weeks with the repeal bill being passed on July 17, 2014, which was backdated to July 1, 2014. When discussing my empirical strategy, I will return to the predetermined nature of these changes in the carbon price and how this will impact my estimates.

1.2.2 National Electricity Market

The National Electricity Market (NEM) began operation as a wholesale spot market in December 1998. It incorporates around 40,000 km of transmission lines and provides electricity to approximately 9 million consumers across five regional market jurisdictions: Queensland, New South Wales, Victoria, South Australia, and Tasmania.[8] Efficient transmission of electricity from generators to consumers is primarily facilitated through a wholesale market operated by the Australian Energy Market Operator (AEMO).

The procedure for determining the market-clearing price in the NEM is similar to many electricity markets throughout the world. Electricity producers and wholesale retailers submit 10 "bands" for each half-hour of the day, each consisting of a price and quantity pair. For electricity producers, the price represents the willingness to accept (WTA) for the paired quantity, while for electricity retailers, the price represents the willingness to pay (WTP) for the paired quantity. AEMO then uses an automated system to rank each electricity producer's band in ascending order, based on WTA, to form a supply curve. A similar process is done with electricity consumers, but their bands are ranked in descending order. The market clears by choosing the lowest-cost combination of producers to satisfy demand, subject to the grid's physical constraints.

The NEM is a fully zonal market, which means that for each period (half-hour), there is a separate market-clearing price for each zone. In the case of the NEM, there are five zones, which are the five regional market jurisdictions. The price for each zone during a given period of the day represents the average change in system cost for one additional unit to be consumed within that zone.[9] In the presence of binding transmission constraints (congestion), zonal prices can, and do, exhibit significant heterogeneity as different electricity generators are required to meet an additional unit of demand for different zones.

[8] Electricity markets in Western Australia and the Northern Territory are not connected to the National Electricity Market and operate separately.

[9] There is a separate financial trading market for electricity that is conducted through over-the-counter trading and through exchange trading via the Sydney Futures Exchange. One example of such a financial contract is a contract for differences where the purchaser agrees to purchase a specified physical quantity of energy from the spot market at a set price.

Summary Statistics

Table 1.1: Summary Statistics for National Electricity Market

	Mean Generation		Capacity		Emissions	
	GWh	%	GW	%	tCO_2e/MWh	Units
Panel A. Energy Source						
Black Coal	11.8	52.3	21.5	41.4	1.0	51
Brown Coal	5.9	26.3	7.9	15.3	1.4	25
Natural Gas	2.4	10.5	9.4	18.0	0.7	64
Hydro	1.7	7.6	9.1	17.6	0	46
Wind	0.4	1.9	2.4	4.5	0	20
Other	0.4	1.9	1.6	3.2	0.9	16
Total	22.5	100	51.9	100		222
Panel B. Firms						
AGL	5.6	24.7	9.9	19	0.7	30
EnergyAus	2.6	11.5	3.7	7.2	1.0	9
Origen	2.1	9.2	5.2	10	0.8	24
Other	12.3	54.5	33	63.7	0.6	159
Total	22.5	100	51.9	100		222

Notes: This table displays generation, capacity, emissions, and units. The top panel displays these variables by energy source and the second panel describes these characteristics for three of the largest electricity generating firms. Over 90 percent of electricity generated in the National Electricity Market comes from fossil fuels. The emissions rates vary widely depending on the fuel and plant, but units that use brown coal tend to have the highest emissions rate and natural gas tends to have the lowest emissions rate. Electricity generation from 3 firms accounts for almost half of total electricity generated and less than 40 percent of total capacity.

The impact of a carbon price on the wholesale electricity market depends on the mix of energy sources used to generate electricity and the extent of market power held by generators. Table 1.1 shows summary statistics for the electricity generated, capacity and emissions rates by type of plant. Approximately 90 percent of the electricity generated in the National Electricity Market uses fossil fuels as the primary energy source with the rest largely being provided by wind and hydropower. Emissions rates vary largely by fossil fuel with average emissions rates of approximately 1.4 tonnes of CO_2e per MWh for brown coal compared to 0.7 for natural gas.

Table 1.1 also highlights the summary statistics by firm for the three largest firms in this market. Combining electricity generation from AGL, EnergyAustrlia, and Origen Energy, these three firms account for more than 40 percent of electricity generated in the market and just less than 40 percent of capacity. Looking at the aggregate market in Table 1.1, however, does not show the market power that these three firms have in the electricity zones of South Australia, New South Wales, and Victoria. In these three electricity zones, the combined generation capacity of the three largest firms account for more than 60 percent of total capacity. This is perhaps not surprising give the high fixed costs of electricity generation.[10]

Retail Electricity Market

Most residential electricity consumers do not interact directly with the wholesale market, but instead purchase from retailers. During the time frame of this study, 2009 to 2014, most residential customers served by the National Electricity Market were in retail markets with full retail contestability, but regulated prices.[11] Retail contestability allows customers to choose between retail providers, but the prices that retailers were allowed to offer were capped by regulations.[12]

The three firms that play a large role in the wholesale electricity market are also

[10] Relative to other contexts, such as Spain (Fabra and Reguant 2014), this is relatively competitive when looking at the entire market. In the presence of transmission congestion, however, market power is exacerbated due to the location of the plants owned by these 3 firms.

[11] The primary exceptions are Tasmania, which did not have a fully contestable market, and Victoria which had full retail contestability and deregulated prices. New South Wales also switched to unregulated prices starting late in 2014.

[12] The maximum cap on retail electricity prices means that there is an upper limit on the rate wholesale electricity prices can be passed on to retail electricity prices. If the cap is binding, then any increases in the wholesale electricity price will have no impact on retail electricity prices.

heavily involved in the retail electricity market. As of 2014, most small consumers in all regions, except Tasmania, had 10 to 16 retail companies to choose from, but much of the electricity demand was met by AGL, EnergyAustralia, and Origen Energy. Market concentration in the retail electricity market has been declining over time, but as of 2016-2017 these 3 firms still accounted for 69 percent of the market share of residential electricity customers (Australian Energy Regulator 2018). If the level of imperfect competition is different across zones then wholesale electricity costs might be passed on at different rates to retail electricity prices across zones. I explore this in my incidence estimates by estimating the relationship between wholesale and retail electricity prices separately for each zone.[13]

1.3 Theoretical Framework

In this section, I derive a formula for incidence in the context of a market with asymmetric (heterogeneous) firms and imperfect competition. I then further extend this framework to allow for differential changes across electricity zones in wholesale and retail electricity prices. Lastly, I discuss some of the key parameters of incidence and how markup adjustments to changes in the carbon price can impact the key parameters.

1.3.1 Incidence

The impact of imperfect competition on incidence is particularly important for a carbon price, which has a large impact on energy-intensive industries that are often characterized by high fixed costs that can lead to market power. In the context of an imperfectly competitive market with symmetric (homogeneous) producers, Ganapati,

[13] This accounts for inter-zonal differences in competition, but it does not account for variation in competition within zones.

Shapiro, and Walker (2020) extend formulas of incidence derived by Weyl and Fabinger (2013).[14] As Table 1.1 highlighted, however, there are large differences in the emissions rates of electricity generators, even among those that burn fossil fuels. To build intuition, however, I will first describe the case of symmetric firms in an imperfectly competitive market before discussing the asymmetric case.

The analysis that follows is partial equilibrium and assumes that all goods outside of the industry are supplied perfectly competitively and the taxed input has perfectly elastic supply.[15] Define the incidence of a change in the marginal input tax rate τ as the marginal effect on consumer surplus (CS) relative to producer surplus (PS):[16]

$$I \equiv \frac{dCS/d\tau}{dPS/d\tau}$$

Unlike taxes on outputs, the impact of a change in the input tax on marginal cost is not one-to-one, but instead depends on the role of that input in production and the availability of substitutes. This is captured by the cost-shift rate $\gamma \equiv dMC/d\tau$. Returning to the case of electricity generation and a carbon price, this term represents the emissions rates of each firm, which can be lowered through input substitution or carbon capture and storage. The pass-through rate of the input tax to output prices, $\rho \equiv dP/d\tau$, captures the impact of changes in the tax on output prices. This is distinct from the marginal cost pass-through rate, $\rho_{mc} \equiv dP/dMC$, which represents the impact of an increase in marginal costs on output prices. Throughout the rest of the paper when mentioning pass-through, I refer to the change in output prices from a change in the input tax.

[14] Applying this to the case of US industries, they find that accounting for the imperfectly competitive nature of the industry leads to 25-75 percent smaller estimates of the total welfare cost that consumers bear.

[15] The latter is consistent with a carbon tax.

[16] When interpreting my results I will focus on the share of the total burden borne by consumers: $\frac{I}{I+1} = \frac{dCS/d\tau}{dCS/d\tau + dPS/d\tau}$

Consider the case of a generalized oligopoly, with N symmetric firms and constant marginal costs such that $AVC = MC$. Ganapati, Shapiro, and Walker (2020) show that, in this case, incidence takes the form:[17]

$$I = \frac{\rho}{\gamma - \rho - \underbrace{\frac{dQ}{d\tau} L \frac{P}{Q}}_{\alpha_{hom}}} = \frac{\rho}{\gamma + \underbrace{\rho L \epsilon_D}_{\alpha_{hom}} - \rho} \tag{1.1}$$

where Q is the total market quantity and $dQ/d\tau$ is the change in the market quantity from a change in the input tax. The market[18] elasticity of demand is denoted by $\epsilon_D \equiv -[dQ/dP][P/Q]$ and the Lerner (1934) Index is represented by $L \equiv (P - MC)/P$. The numerator captures the change in costs for customers through output prices and the denominator shows that the impact on producer surplus depends on the change in marginal costs and prices $(\gamma - \rho)$ as well as the value of lost output (α_{hom}). In the case of perfect competition, price equals marginal cost and $L = 0$ which reduces incidence to:

$$I = \frac{\rho}{\gamma - \rho}$$

where the change in producer surplus depends only on the difference between the cost-shift and pass-through rates.

Asymmetric Imperfect Competition

Homogeneous costs, however, are unrealistic in a wide-variety of contexts and electricity markets, in particular, so I will now extend this framework to allow for heterogeneity in marginal costs and the cost-shift rate. Let $\bar{\gamma} = (\sum_{i=1}^{n} q_i \gamma_i)/Q$ denote the quantity-

[17]Derivation in Ganapati et al. (2020), but also shown in Appendix A.1.1.

[18]Since all firms are homogeneous with a single price and quantity, the firm and market demand elasticities are equivalent (Genesove and Mullin 1998).

weighted average shift-rate. This term captures the average increase in marginal costs per unit of output. Define the firm-specific Lerner Index as $L_i \equiv (P - MC_i)/P$, which captures the difference between the output price and marginal costs for arbitrary firm i.

Proposition 1: *Under generalized oligopoly, with N heterogeneous, constant marginal cost firms, and a single output market, incidence from an input tax is*

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{\rho}{\bar{\gamma} - \rho - \underbrace{\frac{P}{Q}\sum_{i=1}^{n}\frac{dq_i}{d\tau}L_i}_{\alpha_{het}}} \tag{1.2}$$

Proof: Shown in Appendix A.1.2. ■

The result when firms are symmetric is nested within the asymmetric case shown in equation 1.2. Specifically, note that if firms are symmetric then the change in marginal costs (γ_i), the Lerner index (L_i), and the change in quantities ($dq_i/d\tau$) are the same for all firms. In this case the value of lost production, α_{het}, is identical in the heterogeneous and homogeneous cases, α_{hom}. Moreover, in the particular case where $\gamma_i = 1$ for all firms, this reduces to the case of an output tax.[19]

This formulation highlights the two major changes when allowing for firm heterogeneity with input taxes. First, the change in marginal costs for the market is now given by the quantity-weighted average of the cost-shift rate. This captures the average change in marginal costs per unit produced from an increase in the tax. Second, the impact of a change in the input tax on producer surplus depends on the impact on each firm's quantity produced and that firm's markup. Unlike the symmetric case, when firms are asymmetric it now matters *which firms* reduce production, since firms are no longer identical with the same difference between prices and marginal costs.

[19]See Section A.1.3 in the Appendix for the derivation using Weyl and Fabinger (2013) as the starting point.

Extension to Electricity Markets

I now discuss how I will adapt the incidence formula in equation 1.2 to further capture the features of the market that I will analyze, electricity. Applying the standard short-run marginal cost decomposition used in electricity markets (e.g., Fabra and Reguant 2014), define marginal cost for thermal plant i as:

$$c_i = FP_i \cdot hr_i + e_i\tau \tag{1.3}$$

where e_i is the emissions rate, τ is the carbon price, FP_i is the fossil fuel price, and hr_i is the heat rate.[20] Since marginal input and emissions costs are additively separable, the cost shift-rate is equal to the emissions rate and the quantity-weighted average shift-rate is equal to the quantity-weighted average emissions rate:

$$\bar{\gamma} = \bar{e} = \frac{\sum_{i=1}^{n} q_i e_i}{Q}$$

By using this decomposition I am assuming that the change in marginal costs from a carbon price is fully captured in the emissions rate of the electricity generator. This would be violated if, for example, electricity generators reduce emissions by changing fossil-fuel inputs from coal to natural gas. To some extent, the ability to substitute from coal to natural gas is limited by existing plant capabilities, but it is worth noting that coal plants can substitute to coal higher with higher calorific values (more carbon content).[21]

Although the assumption that the change in marginal costs is equal to the emissions rate is more plausible in the short-run, which is the focus of this analysis, I want to

[20] I focus only on defining marginal costs for thermal plants since I am looking at a carbon price. Renewable electricity generators have an emissions rate of 0.

[21] In general electricity plants in Australia are not purchasing electricity on the open market, but instead they use coal mined nearby and transported on conveyors to the plant for use in electricity generation.

highlight that even in the case where it is violated the emissions rate prior to the carbon tax change acts as an upper bound of the change in marginal costs. Intuitively, this is because if a producer makes no changes in inputs and does not invest in emissions reduction technologies then the change in marginal costs is the emissions rate. Any adjustments that producers make to minimize costs in response to the carbon tax thus must lower the marginal costs of production. Returning to the expression for incidence shown in equation 1.2 this means that any violations of this assumption will downward bias my estimates of incidence.

A second aspect of the electricity market not captured in equation 1.2 is that there are multiple market-clearing prices in electricity markets. Wholesale electricity prices in the NEM often differ across each of the 5 electricity zones. To incorporate this added dimension into the model of incidence, I will introduce additional definitions. First, for zone m define the zonal quantity demanded and supplied (generated) as q_{dm} and q_{sm}. Since quantity demanded need not equal quantity supplied in each electricity zone as transmission lines connect each zone,[22] define the quantity-weighted average pass-through rates for producers and consumers as $\bar{\rho_d} \equiv \sum_{m=1}^{5} q_{dm}\rho_m/Q$ and $\bar{\rho_s} \equiv \sum_{m=1}^{5} q_{sm}\rho_m/Q$, respectively. The price each firm experiences within a zone is identical, $P_i = P_m$ for all i in zone m, so I will redefine the firm-specific Lerner Index as $L_i \equiv (P_i - MC_i)/P_i$. Intuitively, this means that two identical firms in the same region will have the same Lerner Index, while two identical firms in different zones can have different Lerner indices.[23] In

[22]Demand and supply need to clear in the aggregate, but need not clear within each zone. Electricity generated in one zone could, hypothetically, be used to meet demand in all zones.

[23]I keep the notation of firm-specific prices primarily for notational simplicity. In reality, underlying every firm's P_i is the zonal price P_m.

Appendix A.1.4, I derive the following formula that I will empirically estimate:[24]

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{\bar{\rho}_d}{\bar{e} - \bar{\rho}_s - \underbrace{\sum_{i=1}^{n} \frac{dq_i}{d\tau} L_i \frac{P_i}{Q}}_{\alpha_{het}}} \tag{1.4}$$

Comparing equations 1.2 and 1.4, there are three key differences. First, for a marginal change in the carbon price, the change in marginal costs is equal to the marginal emissions rate for each firm. Second, prices can differ across each zone of the electricity market which, in combination with the fact that supply does not necessarily equal demand in each zone, means that the average pass-through rate experienced by consumers need not be equal to that experienced by electricity generators (i.e., $\bar{\rho}_d \neq \bar{\rho}_s$). Lastly, the value of decreases in output depends on the difference between marginal cost and price in the zone where the electricity is generated. Using the average price in the market would overweight electricity generators in zones that, in the absence of transmission constraints, would export electricity (lower prices) compared to those that import electricity (higher prices).

Incorporating Retail Electricity Markets

Taking equation 1.4 to the data, it is worth noting that both sides of the market are firms. Electricity in the wholesale market is purchased, for the most part, by electricity retailers. These retailers then sell electricity to residential consumers, where the prices are

[24]The homogeneous case is given by

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{\bar{\rho}_d}{\bar{e} - \bar{\rho}_s - \underbrace{\sum_{m=1}^{5} \frac{dq_m}{d\tau} L_m \frac{P_m}{Q}}_{\alpha_{hom}}}$$

where q_m is the total quantity generated in zone m.

often fixed over large periods of time. Depending on the pass-through rate from wholesale to retail, the burden by actual consumers could be more or less than what is estimated in equation 1.4. In order to capture this, define the welfare for actual electricity consumers as CS^{ret} and retail prices in zone m as P_m^{ret}. Assuming that all electricity demanded in each zone is purchased at retail price P_m^{ret}, the incidence formula for consumers compared to electricity generators is a minor adaptation of equation 1.4:

$$I = \frac{dCS^{ret}/d\tau}{dPS/d\tau} = \frac{\bar{\rho}^{ret}}{\bar{e} - \bar{\rho}_s - \sum\limits_{i=1}^{n} \frac{dq_i}{d\tau} L_i \frac{P_i}{Q}} \tag{1.5}$$

The denominator is identical, but now the loss in consumer surplus reflects the average pass-through to retail prices per unit consumed. This is defined as $\bar{\rho}^{ret} \equiv \sum_{m=1}^{5} q_{dm} \rho_m^{ret}/Q$ where $\rho_m^{ret} \equiv dP_m^{ret}/d\tau$ is the pass-through rate to retail prices. The impact of the carbon tax on consumers now depends on the relationship between the input tax and the price of electricity purchased by customers. Intuitively, the pass-through rate to retail $(dP^{ret}/d\tau)$ is equal to the pass-through rate to wholesale $(dP/d\tau)$ multiplied by the relationship between wholesale and retail electricity prices (dP^{ret}/dP).

As equation 1.5 highlights, the change in producer surplus depends, in large part, on the wedge between the change in marginal costs $(\bar{e})$ relative to the change in wholesale electricity prices $(\bar{p}_s)$. Since electricity markets function as an auction, there are two mechanisms which could drive this wedge, both of which are shown in Figure 1.2. First, if marginal (price-setting) plants have emissions rates that are different than infra-marginal plants emissions rates then, even if marginal plants increase their bid price by the change in marginal costs, the change in price will be different than the change in average marginal costs. Second, as panel (b) shows, marginal plants can increase their bid by more or less than their change in marginal costs. In other words, marginal plants can change their

Figure 1.2: Stylized Example - Change in Producer Surplus

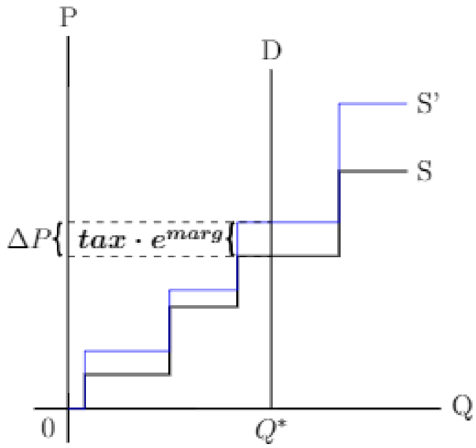

(a) Bidding Marginal Cost (b) Markup Adjustment for Marginal Plant

Note: This figure shows the change in producer surplus from a change in the carbon tax, assuming all firms set their offer price equal to marginal costs. The first panel shows that if the marginal plant does not adjust markups then the change in producer surplus depends on the emissions of the marginal plant, e^{marg}, compared to the emissions rates of infra-marginal plants. In this example, the infra-marginal plants have emissions rates smaller than the marginal plant, which is evident in small adjustments to the bid price from a change in the carbon tax. The second panel shows the importance of markups for the marginal plant. Adjustments to markups by marginal plants directly impacts the change in producer surplus through changes in the price. In this example, marginal plant markups increase which increases "windfall profits" for all plants.

markups in response to a change in the carbon price.

As I show in Appendix A.2, the effect of a change in the carbon price is theoretically ambiguous both in magnitude and sign depending on the level of heterogeneity in emissions rates. Determining whether estimates of incidence are a result of behavioral changes by electricity generators or purely the location of electricity generators along the supply curve is crucial to understanding the mechanism behind the effect of the carbon tax on producer surplus. In a later section, I will empirically estimate each of these components to determine what is driving any observed differences between average emissions rates and the change in prices.

1.4 Empirical Framework and Data

In this section, I describe the empirical strategy and data I will use to estimate the necessary parameters for incidence.

1.4.1 Wholesale Empirical Framework

There are two key parameters that I need in order to calculate equation 1.4 - pass-through for each region ($dp_m/d\tau$) and the change in quantity produced for each electricity generator ($dq_i/d\tau$). I estimate the impact of the carbon tax on each of these wholesale market equilibrium outcomes separately with:

$$y_{jt} = \beta\tau_t + \sum_{j=1}^{J} \beta_j(\tau_t \cdot R_j) + \mathbf{S}_{jt} + \gamma\mathbf{D}_{jt} + \psi\mathbf{X}_t + \delta I_{jt} + \epsilon_{jt} \tag{1.6}$$

where τ_t is the price on carbon and R_j is indicator for each location. Depending on the regression, the dependent variable y_{jt} is either the price in zone j on date t or the electricity generated by plant j on date t. In all specifications, the βs are the parameters of interest with β estimating the relationship between τ and y for the dropped zone or plant and β_j estimating the differential impact of τ on y for zone or plant j. The error term ϵ_{jt} is clustered at the month of sample in all regressions.[25]

Although the level and timing of changes in the carbon tax are exogenous to fluctuations in the wholesale electricity market, there are two key concerns when identifying the

[25]Standard errors are clustered at the month level to match results from applying the automatic bandwidth selection procedure of Newey and West (1994), where the automatic lag length was approximately 30.

effect of the carbon tax on wholesale electricity market outcomes. First, if changes in the carbon price are correlated with unobserved shocks in the wholesale electricity market then this would bias estimates of β. I address the potential omitted variable bias concern by controlling for exogenous shifts in demand ($\boldsymbol{D_t}$), supply ($\boldsymbol{S_t}$), general shocks ($\boldsymbol{X_t}$), and fluctuations across regions over time ($\boldsymbol{I_{jt}}$). In my preferred specification I include location (zone or plant) by month, location by day of the week, and year fixed effects to control for unobserved fluctuations across regions and over time. As common controls ($\boldsymbol{X_t}$), I include global fossil-fuel prices (coal, gas, and oil), which are major determinants of demand and supply. To control for changes in electricity demand due to weather fluctuations ($\boldsymbol{D_t}$), I include average temperature (minimum and maximum) and rainfall as well as each of these variables squared. Lastly, I include average lagged water levels for electricity providing dams in each region ($\boldsymbol{S_t}$) to account for differences in available water for electricity generation over time. Crucially, when estimating equation 1.6 with electricity generator production as the dependent variable, I allow fossil fuel prices to impact each electricity generator differently, capturing differences in heat rates as well as energy sources for all plants.

The second key concern when estimating equation 1.6 is that, although the carbon tax is exogenous to the wholesale electricity market, the timing of changes in the carbon tax were known beforehand. The predetermined or foreshadowed, in the case of the repeal, nature of the carbon tax changes could lead to responses by electricity generators in advance of the date the tax changes, which could bias my estimates of the impact that the carbon tax has on the wholesale electricity market.

Focusing on estimating the relationship between the carbon tax and wholesale electricity prices, knowledge of future carbon tax changes will bias my estimates downwards. For most plants, short-run production decisions are not influenced by distant changes in marginal costs or increases in the wholesale electricity price, except through adoption of emissions reductions technologies. The primary exception is hydropower, which operates with a binding intertemporal budget constraint on production. If a hydropower plant knows that the price is going to increase in the near future from an increase in the carbon tax, then they will decrease current production and increase production in the future. Intuitively, this effect drives up wholesale electricity prices in advance of the carbon tax as decreases in hydropower electricity are compensated by potentially more expensive peaking plants. Alternatively, when the carbon tax is implemented there is now more water available for these plants, which can drive down wholesale electricity prices. I attempt to account for this channel directly by controlling for lagged dam levels of hydropower plants in each region, but it is worth noting that this channel downward biases my estimates of the relationship between the carbon tax and wholesale electricity prices.

When estimating the relationship between the carbon tax and individual electricity generator production, knowledge of future carbon tax changes could bias my estimates through two key channels. First, for thermal electricity generators, forewarning of an impending carbon tax change could lead to technology adoption in advance of the carbon tax implementation. This would lead to a smaller change in marginal costs for these plants, thereby leading to, potentially, smaller decreases in production. If this was a large factor, then one would expect a consistent pattern of downward revisions in emissions rates from 2011 to 2012, but this is not evident in updated emissions rates in a report prepared for the Australian Energy Market Operator (A. Allen 2014). Given that there is little evidence of decreases in emissions rates, this suggests that the effect of this channel is likely small.

Second, for hydropower plants, intertemporal reallocation of production would directly lead to changes in production away from when there is no carbon tax towards when there is. Insofar as this is not captured by the lagged dam level controls, this would bias my estimates of the relationship between the carbon tax and hydro production upwards. Moreover, if production is reallocated by infra-marginal plants, such as hydro, then this will lead to an identical magnitude, but opposite sign reallocation for marginal plants. As I discuss in depth in Appendix A.3, this would bias my estimates of incidence upward since markups are higher for hydropower plants compared to those operating at the margin. To ensure that my results are not driven by this, I will compare my incidence estimates to those where I set the estimated change in production per dollar of the carbon tax equal to 0, which is the response one would expect from hydropower if the carbon tax was implemented without forewarning.

1.4.2 Retail Empirical Framework

Moving from estimates of incidence that focus only on the wholesale electricity market to ones that incorporate the retail electricity component, which is shown in equation 1.5, I need to estimate the relationship between retail electricity prices and the carbon tax. I do this with the following specification:

$$r_{jt} = \alpha + \beta\tau_t + \sum_{j=1}^{5} \beta_j(\tau_t \cdot R_j) + \gamma I_{jt} + \epsilon_{jt} \tag{1.7}$$

where r_{jt} is the average retail electricity price in zone j and quarter of the year t and the carbon tax is τ_t. The electricity zone is represented by R_j, which is an indicator for each of the 5 electricity zones. The parameters of interest are again β and β_j with the former identifying the effect for the omitted zone and the latter identifying the differential effect for zone j compared to the omitted zone. If the carbon tax has the same impact on all electricity zones then β_j will be equal to 0.

The retail electricity prices used in this analysis reflect annual consumption applied to each quarter's consumption. As Figure A.2 shows, this leads to minimal seasonality throughout the year as variation in average retail prices is primarily annual. These annual changes in prices reflect changes to four components of the retail electricity prices: wholesale electricity prices, operating costs, retailer markups, and distribution and network costs. To control for variation in the latter three components, my preferred specification includes zone fixed effects and a linear year time trend, I_{it}. Zone fixed effects control for time invariant differences across zones while time trends control for changes in retail prices over time due to factors such as increasing network costs. In my primary specifications, I use linear time trends, but results are nearly identical using year fixed effects. I cluster the standard errors at the quarter of the year and estimate p-values using the wild bootstrap (Roodman, Nielsen, MacKinnon, and Webb 2019).

Crucial to the interpretation of β is whether or not one would expect the carbon tax to impact the retail electricity prices, except through wholesale electricity prices. For example, if the carbon price has a significant impact on the prices of other goods that are complements or substitutes for electricity, then β would not be representative of just the effect that the tax has on the electricity market. This concern is mitigated, to some extent, by the low levels of price and income elasticities that have been found in electricity markets across a wide-variety of contexts (e.g.,Ito 2014, Deryugina, MacKay, and Reif 2020, Vesterberg 2016). More importantly, however, the carbon tax was not broadly applied to goods that one would expect could impact electricity significantly, such as gasoline. These two factors, highly inelastic demand and limited coverage of the carbon tax, provide suggestive evidence that the interpretation of β is likely to reflect the impact of the carbon tax on the electricity market and not the effect of the carbon tax on other markets.

1.4.3 Markups Empirical Framework

Much of the literature that has analyzed the impact of heterogeneous cost shocks on markups in electricity markets conduct a simulation where all other factors except for the shock can be controlled (Fabra and Reguant 2014, Kim 2021). The simulation is based on a first-order approach where a small cost perturbation is implemented to the electricity market, and then the changes in the markups of each firm are estimated based on the first-order condition. Importantly, this approach provides an estimate for how firms should, if they are behaving optimally, adjust markups in response to the cost shock. Even in highly sophisticated markets, however, electricity generators have been shown to exhibit significant strategic limitations (Hortaçsu et al. 2019). Instead of estimating changes in how markups *should* change, I estimate observed changes in markups in response to a small change in the carbon price from 23 to 24.15 on July 1, 2013.

The optimal change in markups as a result of a change in the carbon price is only valid for units which, ex-ante, believe they set the price with some positive probability and are not capacity constrained. To estimate the effect of the carbon price on markups I thus must identify the effect that the carbon price has on bands (steps) of electricity generator supply curves which there is reason to believe are (i) set with the belief that it might be marginal and (ii) not responding to factors like start-up costs or capacity constraints. To do this, I leverage one of the unique features of the Australian electricity market. Since electricity generators only have 10 price bands per day, they often use each band for different periods throughout the day and for different purposes, such as price-setting. I leverage this fact to identify not only the electricity generator, but also the price band, for which there is reason to believe ex-ante might be marginal.[26] Specifically, I focus

[26] I focus on ex-ante potentially marginal primarily because the ex-post marginal plant will likely depend on how they respond to the change in costs.

on electricity generator bands that are marginal more than 95 percent of other bands throughout 2012,[27] that are also not the first or last step of the supply function. These 37 offer price bands from 27 electricity generators account for more than 20 percent of *all* price-setting bands from 2012 to 2014 and the month before and after the carbon tax change that I analyze in 2013. The high frequency at which these price-setting bands identified are marginal over long periods of time highlights the consistent strategic behavior of these electricity generators.

The impact of the carbon price on frequently marginal offer prices for two different electricity generators is shown in Figure 1.3. When the carbon tax is implemented there is a large increase in the price that they are willing to accept. From 2012 to 2014 this fluctuates over time before decreasing sharply in response to the repeal of the carbon tax. At each of the tax changes in July of 2012, 2013, and 2014 there is a large, fixed change in the carbon tax which is constant until the next year. As long as the change in marginal costs at that point in time is not shared with other changes in the electricity market, then the impact of the carbon price on markups can be identified through a regression discontinuity design (RD).

Formally, define the difference between the band offer price and marginal emissions costs as $V_{it} = b_{it} - mec_{it}$ where mec_{it} and b_{it} represent the marginal emissions cost and bid price for plant i on date t. One can estimate the impact of the carbon tax on markups for marginal electricity generators as:

$$V_{it} = \alpha + \gamma \mathbb{1}(t \geq c) + \beta_1 f(t) + \beta_2 f(t) \mathbb{1}(t \geq c) + \epsilon_{it} \tag{1.8}$$

The parameter of interest is γ, which identifies if the tax changes the level of markups as a result of a change in the carbon tax. If the bid price increases by approximately the

[27] I concentrate on this period since it is the year before the cost shock that I focus on here, which is in the middle of 2013.

Figure 1.3: Frequently Marginal Offer Prices over Time

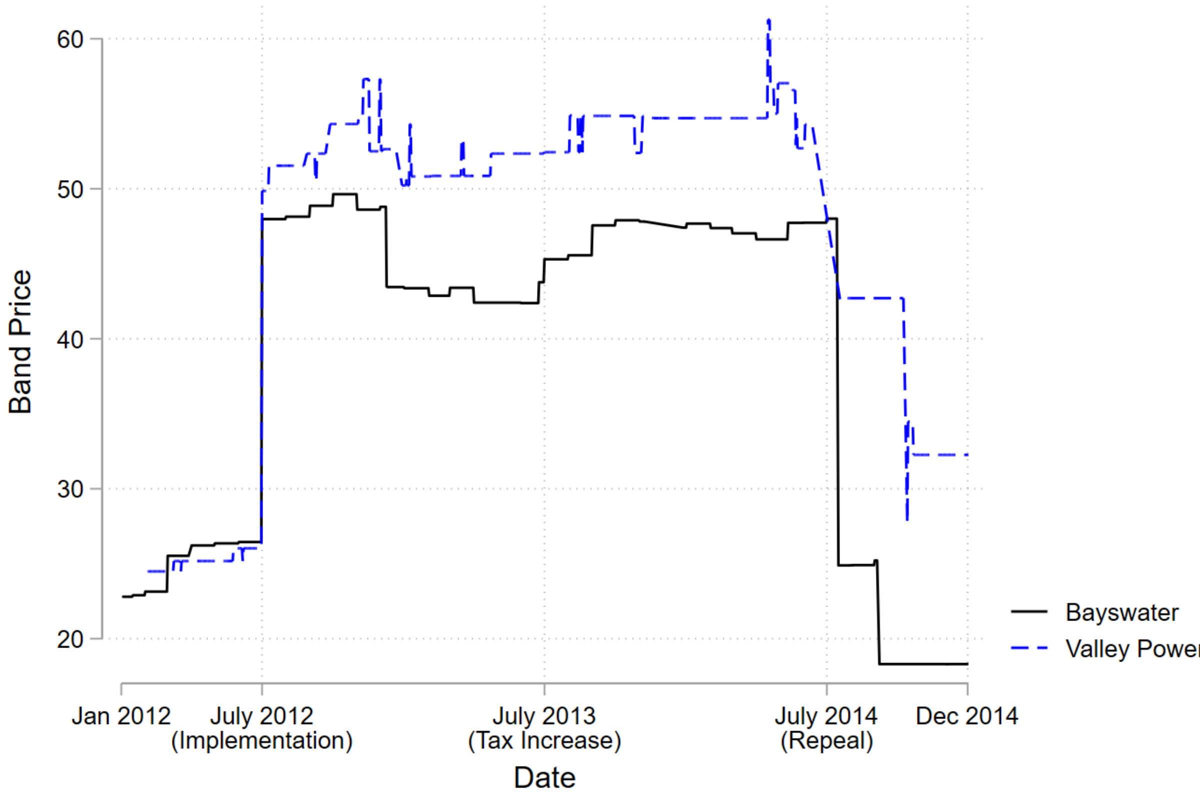

Note: This figure plots a single band's offer price over time for Bayswater and Vales Point "B" power stations. Both of the bands shown are in the top 5 percent in terms of how frequent that it is marginal. The line in blue plots the offer price for Valley Power over time, while the line in black represents Bayswater. Both plants use black coal as the primary energy source.

change in marginal emissions costs, then γ will be insignificantly different from 0, while if markups are adjusted upwards (downwards) then γ will be more (less) than 0.

There are two key concerns when estimating a regression discontinuity in time. First, in order to interpret γ as a change in markups, marginal costs and electricity demand must be continuous at the cutoff date. If there are unobserved plant-specific shocks to input prices when the carbon tax changes, then this would bias estimates of the markup

(γ). Alternatively, shocks in demand would change the interpretation of γ as shifts in demand impact the incentive for markup, independent of the cost shock. To test both of these factors, I estimate equation 1.8 using demand and short-term natural gas prices as dependent variables.[28] Plotting the natural gas prices and quantity demand in Figures A.4 and A.3, there appears to be no evidence of a discontinuity at the cutoff. This is supported by the regression discontinuity estimates in Table A.1, which show no significant effect of the tax change on natural gas prices or demand. These results provide evidence in favor of interpreting the RD estimate in equation 1.8 as the change in markups.

Second, although there are no significant difference in demand at the cutoff, there is large amounts of variation over time which can lead to noise in the offer price data. One challenge with the local linear specification is that estimating equation 1.8 with controls can be difficult without increasing the bandwidth. To overcome this, I adopt an alternative estimation strategy put forth in Hausman and Rapson (2018), the augmented local linear approach. This is a two-step process where I first estimate the impact of factors such as coal prices on the outcome of interest, which in my case is the difference between the band price and marginal emissions costs. Next, I use the residuals from the first regression as the dependent variable in a regression discontinuity design with a local linear specifications and a narrow bandwidth, such as 30 days. Standard errors are recovered with a bootstrap that allows variance from the first-stage to be reflected in the second stage. Following Lee and Card (2008), the standard errors are clustered according to the running variable, which in this context is the date.

In my preferred augmented local linear RD results, I first regress the difference between the band price and marginal emissions costs, V_{it}, on fixed effects and fuel price

[28]Natural gas prices are the only fuel for which the temporal variation is sufficient to estimate a regression discontinuity design. Data on natural gas prices is at the day level, while the best available coal data is at the month level.

controls. Specifically, I control for differential band prices by year and day of the week through band by year and band by day of the week fixed effects. To control for differential impacts of fuel prices on electricity generators and specific bands for each electricity generator, I allow for a heterogeneous relationship between coal and natural gas prices for each price band. Lastly, the Australian Energy Market Operator (AEMO) gives projected demand to market participants before they must submit their supply curve. I allow for a quadratic relationship between electricity demand projections on each potentially marginal price band that I analyze.[29] Using the residuals from this regression, I then estimate a regression discontinuity on the date the carbon tax changes, July 1, 2013.

1.4.4 Data

Wholesale Electricity Prices

The primary dependent variable that I will utilize in my analysis is the quantity-weighted average wholesale electricity price for each day, which I construct using wholesale electricity price and quantity data for every half-hour from July 1, 2009 to December 31, 2014.[30] That is, for region j and date t, I calculate the daily wholesale electricity price as the quantity-weighted average:

$$p_{jt} = \frac{\sum_{h=1}^{48} p_{htj} \cdot q_{htj}}{\sum_{h=1}^{48} q_{htj}} \tag{1.9}$$

[29] I use electricity demand projections since that is what electricity generators likely respond to. Weather conditions used as controls in equation 1.6 are ex-post, but electricity generators set their bid prices prior to the weather realization.

[30] Electricity generators are required to submit all of the bands for a given day by the afternoon of the previous day. Since all of the bands are submitted daily, most of the results presented throughout this paper will use day of the year as the unit of time. They are are robust, however, to using a half-hour (trading period) as the unit of time and including period fixed effects.

Figure 1.4: Wholesale Electricity Price Distribution

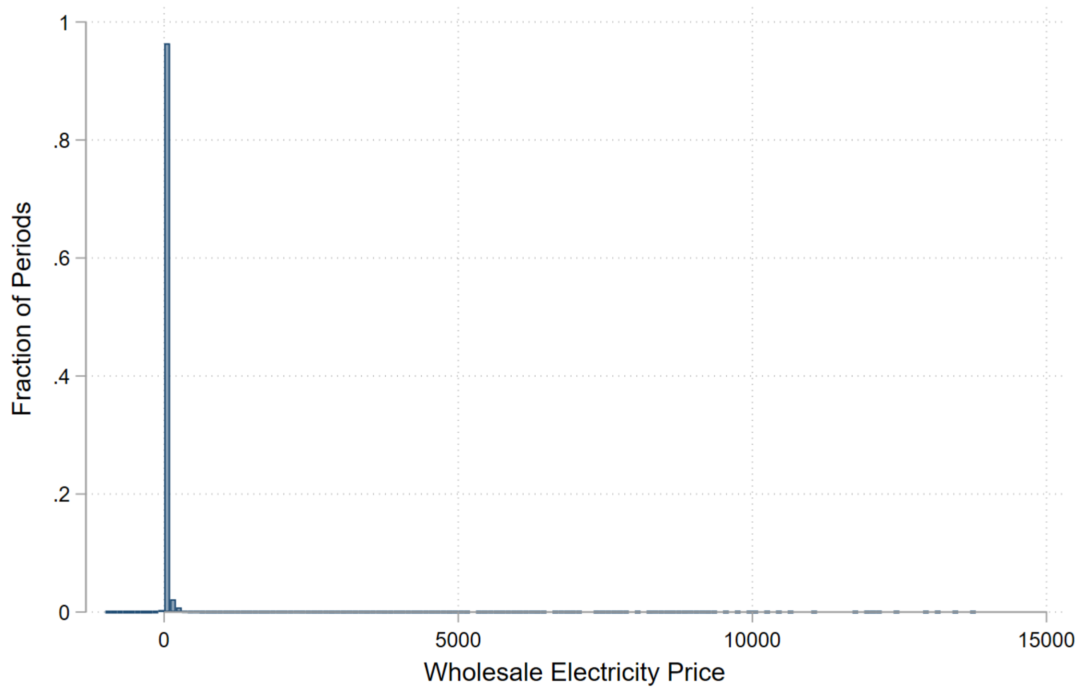

Note: This figure plots the distribution of electricity prices, which highlights the large outlier wholesale electricity prices. Bins are width 100 with well over 90 percent of observations between 0 and 100 AUD/MWh.

where h represent every half-hour (trading period) of the day. The daily wholesale electricity price is thus a weighted average of each half-hour price, p_{htj}, where the weights depend on the fraction of the total quantity consumed during that half-hour, $q_{htj}/\sum_{h=1}^{48} q_{htj}$. By putting larger weights on prices during times when more electricity is consumed, the quantity-weighted average price reflects the average price for a single unit of electricity in a given day.

One of the major problems with wholesale electricity price data is the presence of event-driven outliers. These days are evident in Figure 1.4, which plots the distribution of wholesale electricity prices for all half-hour periods from July 1, 2009 to December 31,

2014. Large spikes in electricity prices are often driven by factors, such as line repair or extreme weather events, which are difficult to capture with controls. To ensure that my results are not driven by these events, I will adopt two different approaches to dealing with extreme values. First, I show the results where I treat extremely large values as outliers. Specifically, I drop prices larger than 1000 AUD/MWh, which is less the 0.01% of observations, and show that the primary conclusions do not change. Second, I will also show results where I estimate a median regression, which minimizes least absolute distance and is thus robust to the presence of extreme values.

Retail Electricity Prices

The primary source of retail data used in this paper is quarterly data from the Australian Consumer Price Index (CPI). The Australian CPI separately tracks residential electricity prices in the capital city of each region (zone) over time. These prices are obtained quarterly from energy providers and local councils and they include both concessional and standard rates.[31] Current charges are applied to estimates of annual consumption of electricity, gas and water to derive the annual payment in the current quarter's prices. Connection fees, delivery and similar charges are included as part of the price of the utility service.[32]

Retail electricity price data in the CPI is not presented in levels, but instead as an index. I convert the retail index into real prices by using annual Australian Energy Market Operator retail price data for each region. Since some years have constant retail electricity prices across all quarters in the CPI, I use these years to recover the rest of

[31] The retail prices include connection fees, delivery and similar charges that are normally included as part of the price for the utility service.

[32] Governments and councils occasionally impose levies on customers of these services as a means of raising money for some possibly unrelated services such as ambulance services. As these levies are considered an inescapable cost of obtaining the original service they are counted as a part of the cost of the original service.

the retail electricity prices for that region from the third quarter of 2009 to the fourth quarter of 2014.

Weather Data

Extreme weather events are a major factor in some of the price extremes in the National Electricity Market, but they are also a major determinant for some of the more subtle variation in day-to-day wholesale prices. To control for these fluctuations, I use weather station data from the Bureau of Meteorology to control for weather across each electricity zone. Since I am primarily concerned with weather conditions for electricity consumers, I combine the location of all weather stations with a detailed map of all high voltage lines in the National Electricity Market. I then put a 1 degree buffer, approximately 100 kilometres, around all of the high voltage lines and exclude weather stations that are not near electricity customers. Figures A.5 and A.6 show this process for parts of the states of Victoria, New South Wales, and South Australia. All of the gray dots that disappear between the two figures are weather stations farther than 1 degree from any high voltage transmission line. Figure A.5 illustrates that the choice of the buffer size has some impact on the inclusion of a few weather stations towards the interior of the country, but it has little impact for the vast majority of weather stations that are along the coast. To get a measure of the average weather for each region, I take a simple average of the daily weather measures for all stations that were not dropped. As Figure A.5 shows, weather stations tend to be focused around populated areas, which are on the coast, so the average across all stations naturally gives a higher weight towards more populated areas.

Although renewable energy accounts for a small portion of electricity generation (Table 1.1), over half of this comes from hydro generation. I control for changes in the availability of water directly by using Bureau of Meteorology daily data on all dam levels to calculate the daily average for each region over time.

Input Prices

To control for changes in the marginal input costs, I include Federal Reserve Economic Data (FRED) data on global prices for coal, natural gas, and crude oil in my primary specifications.[33] Since all of the fuel prices are in US dollars, I convert all of the prices to Australian dollars using daily exchange rate data from the Reserve Bank of Australia. Figure A.7 shows the international gas price compared to local natural gas prices. Prior to the end of 2014, the prices of international and local gas prices are different in levels, but they follow on similar trends. Starting in 2015, however, the international and local natural gas prices appear to disentangle. This was driven in large part by exporters of natural gas who agreed to contracts with punitive damages. As exporters struggled to meet the agreed on quota, domestic supply for natural gas was heavily constrained and prices began to rise even while global prices fell. With this in mind, I focus on the period before 2015 primarily since international gas prices are a bad control for the changes in gas prices experienced by electricity generators from 2015 onward.

1.5 Incidence Estimation

This section proceeds in two parts. During the first part, I estimate the key parameters for incidence using a variety of specifications. In the second part, I detail how I measure the remaining parameters for incidence and calculate incidence using estimates of the key parameters.

[33]Data retrieved from the following FRED commodity series: PCOALAUUSDM, PNGASJPUSDM, and POILWTIUSDM.

1.5.1 Wholesale Pass-through Estimates ($\frac{dP_m}{d\tau}$)

I estimate the impact of the carbon price on wholesale electricity prices using equation 1.6 where the dependent variable is daily quantity-weighted wholesale electricity prices. The total effect for each electricity zone, $\beta + \beta_j$, from estimating equation 1.6 using my preferred specification and various methods of addressing outliers are shown as red dots in Figure 1.5. All other dots in the figure represent alternative specifications, which are shown in Appendix Tables A.2, A.3, and A.4. In general, they show a clear pattern where, for a dollar increase in the carbon price, wholesale electricity prices in Queensland and South Australia increase by a dollar, Victoria and New South Wales by slightly less than a dollar, and Tasmania by approximately 50 cents.

The results of my preferred specification are also shown in Table 1.2. Each column represents a separate regression and each row shows the total effect ($\beta+\beta_j$) for electricity zone j. All estimates are statistically different from 0, highlighting the fact that the carbon tax has a significant impact on all electricity zone prices. Moreover, across all three regressions there is a clear pattern where the highest pass-through rates are around 1 in the Queensland and South Australia zones. Estimates of the pass-through rate, however, fall in Tasmania and Victoria when outliers are better accounted for by dropping prices larger than 1000 or estimating a median regression. Conversely, New South Wales pass-through rate increases significantly when these observations are dropped as a result of some large events prior to the carbon tax implementation. Overall, the results do change when the right tail of the wholesale electricity price distribution is dropped, but it is clear that the large estimated pass-through rates in South Australia and Queensland are not driven by outlier events. Since exact pass-through rates vary by specification, however, I will estimate incidence with pass-through estimates from alternative estimation strategies and specifications.

Figure 1.5: Pass-through Rate by Zone

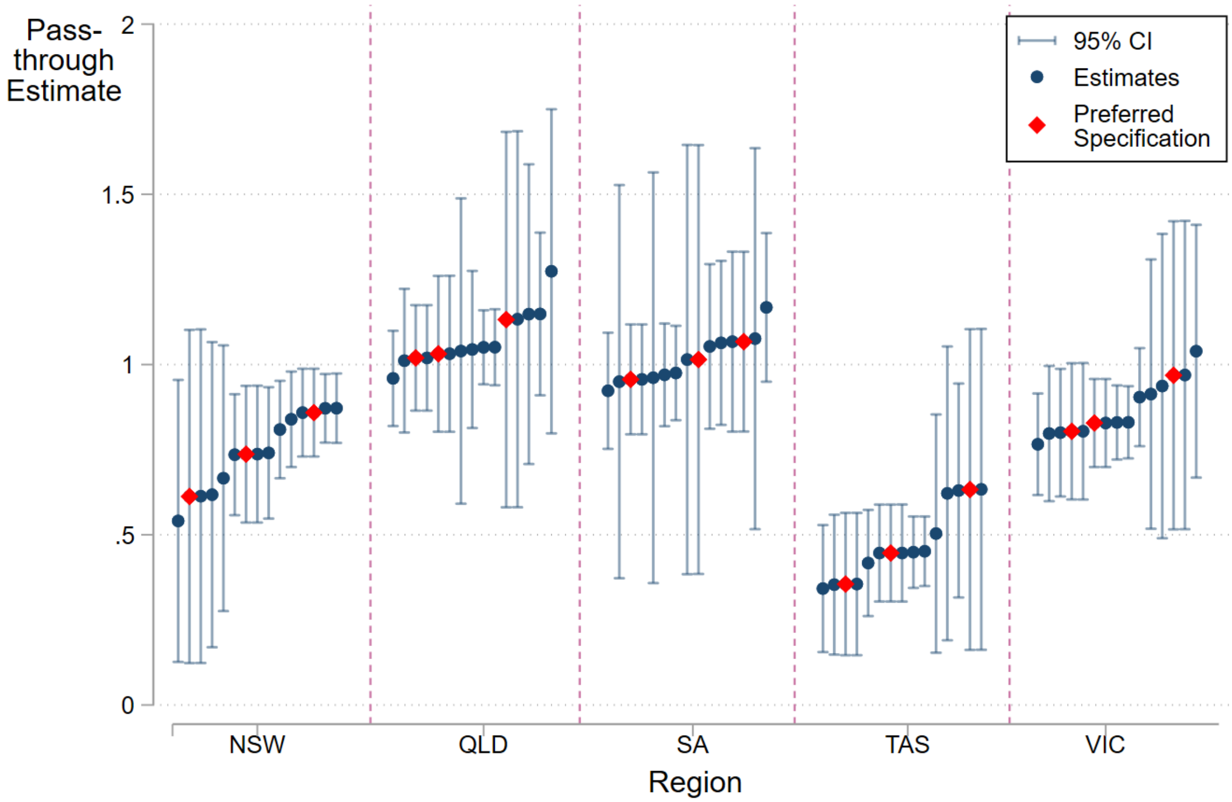

Note: This figure shows the estimated pass-through rate for each zone for all results shown in Tables A.2- A.4. All 95 percent confidence intervals have standard errors clustered at the month of sample level. Estimates labeled as being from the "preferred specification" represent the results from column 5. These includes year fixed effects, zone by month, zone by day of the week, and electricity zone fixed effects. It controls for fossil-fuel prices, previous day dam levels, and weather and allows a quadratic relationship between weather and dam levels. There are three preferred specifications signifying the preferred specification estimated with the raw data, dropped outliers ($1000 AUD/MWh$), and a median regression. Although not labeled, the regression results using the raw data and ordinary least squares are the results with the widest confidence intervals.

Table 1.2: Wholesale Pass-through by Region

	Wholesale Electricity Price		
	(1)	(2)	(3)
NSW	0.618***	0.809***	0.741***
	[0.229]	[0.073]	[0.098]
Queensland	1.040***	0.960***	1.011***
	[0.279]	[0.071]	[0.108]
South Aus	0.961***	0.922***	1.053***
	[0.273]	[0.087]	[0.123]
Tasmania	0.622***	0.417***	0.354***
	[0.199]	[0.080]	[0.105]
Victoria	0.937***	0.766***	0.797***
	[0.190]	[0.076]	[0.101]
Observations	10,049	10,049	10,049
Clusters	66	66	66
Preferred Specification	Yes	Yes	Yes
Median Reg	No	Yes	No
Outliers Dropped	No	No	Yes

Notes: Each column represents a separate regression where the dependent variable is wholesale electricity prices with the second column showing results using a median regression and the last column drops periods with prices larger than 1000 AUD/Mwh. The key independent variable is the price on carbon emissions. All columns include the most restrictive specification (5) estimated in Tables A.2- A.4. This includes zone by day of the week and zone by month fixed effects along with controlling for fossil fuels and allowing for a quadratic relationship between weather and dam levels on prices. Every row represents the total pass-through rate for that region. Statistical significance is based on testing if the total effect with a null hypothesis of 0. Standard errors are in brackets clustered at the month of sample level with the median regression clustering following Parente and Silva (2016). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

1.5.2 Firm-Level Changes in Production Due to the Tax ($\frac{dq_i}{d\tau}$)

One of the primary short-run outcomes from a carbon price in wholesale electricity markets is that it leads to reordering along the supply curve as firms with different emissions rates internalize their marginal emissions costs. Similar to the equilibrium outcomes estimated at the market level, to causally identify the impact of the carbon price on electricity generation, one must control for all other factors that could lead to changes in electricity generator production.

I estimate equation 1.6 simultaneously for all firms where the outcome variable is average electricity generation per half-hour period of the day. The base specification includes electricity generator, month, day of the week, and year fixed effects. The most restrictive specification allows for differences in monthly and day of week production for each electricity generator through generator by month fixed effects and generator by day of the week fixed effects. Moreover, since electricity generators use different fossil fuels, I allow fuel prices to impact each electricity generator differently.

The results for all specifications are shown graphically in Figure 1.6. Each dot represents one of the 220 electricity generator's estimated quantity response to a change in price of carbon emissions with those that are statistically significant at the 5% level in red. Across all five specifications, the pattern of results are remarkable similar, with relatively small changes for most firms, but large changes for a few. This is not surprising since much of the change in production is likely to occur for firms located around the equilibrium quantities as they are likely to be the ones most impacted by the reordering along the supply curve. All else equal, however, we would expect electricity generators with higher emissions rates, CO_2e/MWh, to be decreasing electricity generation more than those with lower emissions rates. Figure 1.7 plots the coefficient estimates by emissions rates using the base and most restrictive specifications. For both specifications

Figure 1.6: Quantity Changes by Specification

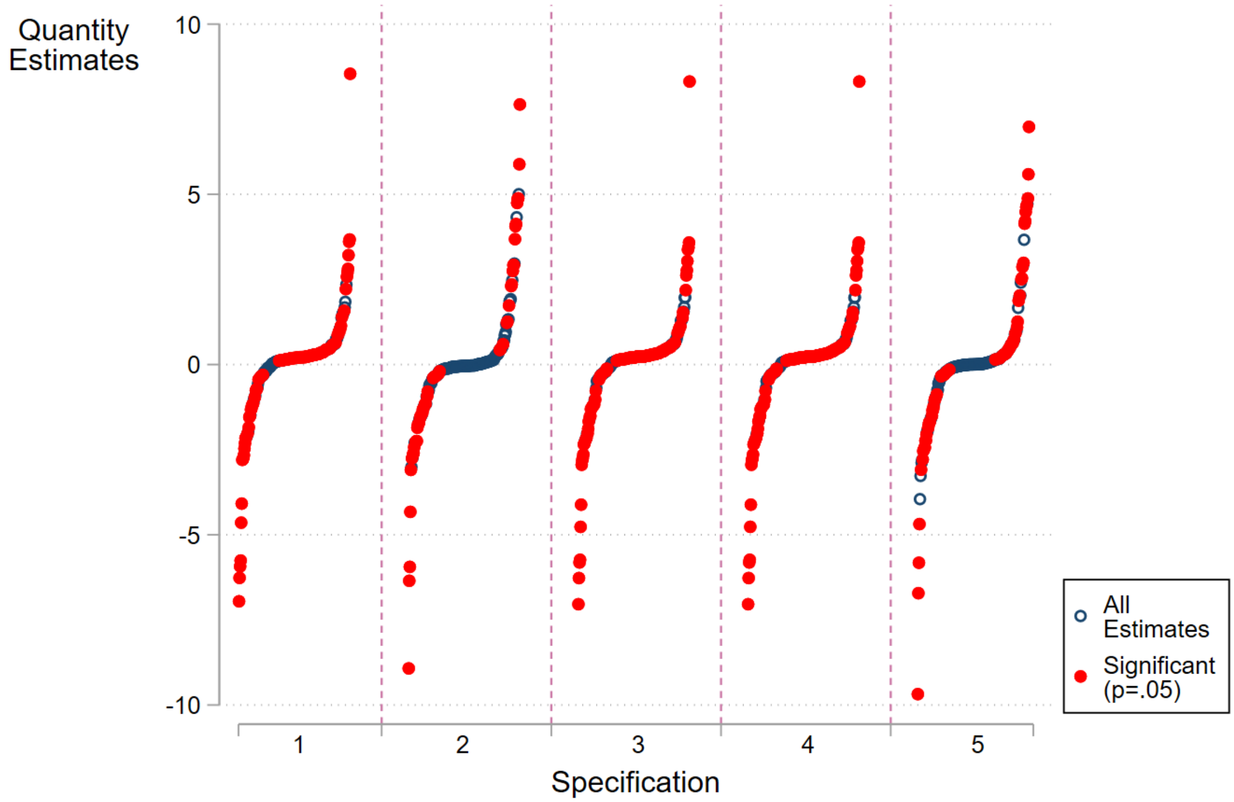

Note: This figure plots estimates of the changes in plant generatio per dollar of the carbon price. Standard errors clustered at the month of sample level. All dots in red represent estimates that reject 0 at a 5 percent level. Each specification represents a separate regression. Specification 1 includes generator, month, year, and day of the week fixed effects and also controls for weather and fuel prices. Specification 2 includes the same fixed effects, but allows each control variable to have a separate effect on each plant's production. Specification 3 is the same as specification 1, but changes month fixed effects to month by plant fixed effects. Specification 4 is the same as specification 3, but changes day of the week fixed effects to month by day of the week. Specification 4 is the same as specification 3, but allows each control variable to have a separate effect on each plant's production.

I find that plants with emissions rates larger than 1 decreased electricity generation and

Figure 1.7: Quantity Changes by Emissions Rates

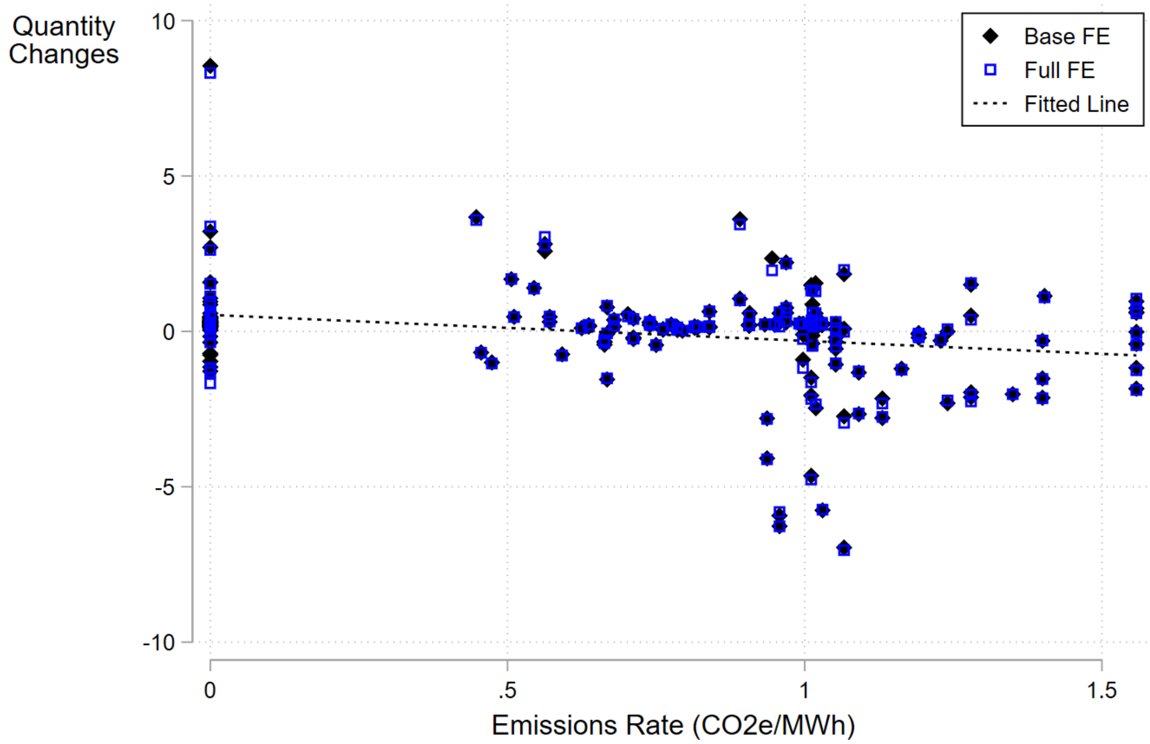

Note: This figure plots estimates the changes in plant generation per dollar of the carbon price depending on the plant's emissions rate. Estimates that are black diamonds show estimates with the "Base FE", which is specification 1 of Figure 1.6. This includes generator, month, year, and day of the week fixed effects and also controls for weather and fuel prices. Estimates shown with blue boxes represent estimates with the "Full FE", which is specification 5 of Figure 1.6. This builds on the basic model by also including generator by month and generator by day of the week fixed effects. It also allows each control variable to have a separate effect on each plant's production. The fitted line shown is for the estimates with the Full FE, but the line is nearly identical using the Base FE. It shows that increases in the price on carbon emissions tends to decrease generation in high polluting plants and increase generation for lower emitting plants.

"cleaner" plants with emissions rates less than 1 increased electricity generation.

The results shown in Figures 1.6 and 1.7 highlight that estimates of the change in quantities are relatively consistent across specifications. When incorporating these estimates into the calculation of incidence, I will focus on the results using the most restrictive specification, but the results are nearly identical when using estimates from any of the other specifications.

1.5.3 Impact of Carbon Tax on Retail Electricity Prices ($\frac{dP_m^{ret}}{d\tau}$)

Figure A.8 plots wholesale and retail electricity prices per kWh for each quarter from the third quarter of 2009 to the end of 2014. The carbon tax is implemented at the very start of the third quarter in 2012, which leads to an immediate increase in both wholesale and retail electricity prices. As Figure A.8 shows, however, retail electricity prices are not stationary and appear to be increasing linearly over time. To control for this trend in the data, my preferred specifications use either year fixed effects or a linear time trend.

The results from estimating equation 1.7 where the carbon tax interacted with a region indicator variable are the coefficients of interest is shown in Table 1.3. The coefficient in each row represents the total effect for each electricity zone, $\beta + \beta_j$, and p-values for all estimates are shown in brackets. As one would expect, retail electricity prices rise by more in most regions that have higher pass-through rates from the carbon tax to wholesale electricity prices. The sole exception is Queensland, which experienced smaller increases in retail electricity prices, despite a large increase in wholesale electricity prices. This is due, in large part, to the fact that there was a freeze on the retail electricity price during a part of the time period when the carbon tax was implemented. This limited the ability for retailers to pass on changes in the wholesale electricity prices to retail prices.

Importantly, the results in Table 1.3 highlight that the choice of linear annual time trends or year fixed effects does not change my results significantly. When calculating incidence I will use the estimates that use a linear time trend, but using either specifica-

Table 1.3: Impact of the Carbon Tax on Retail Prices

	Retail Electricity Price		
	(1)	(2)	(3)
New South Wales	2.812***	1.262***	1.162***
	[0.000]	[0.000]	[0.000]
Queensland	2.257***	0.707*	0.607
	[0.001]	[0.062]	[0.141]
South Aus	3.074***	1.512***	1.412***
	[0.000]	[0.000]	[0.000]
Tasmania	2.277***	0.715*	0.615*
	[0.000]	[0.063]	[0.084]
Victoria	2.853***	1.291***	1.191***
	[0.000]	[0.000]	[0.000]
Observations	110	110	110
Clusters	22	22	22
Zone FE	Yes	Yes	Yes
Linear Time Trend	No	Yes	No
Year FE	No	No	Yes

Notes: Each column represents a separate regression where the dependent variable is the retail electricity price and the key independent variable is the wholesale electricity price, all in cents per kWh. The wholesale electricity price is instrumented by the carbon tax interacted with a region indicator variable. All estimates represent the impact of a change in the wholesale electricity price on the retail electricity price. Standard errors are clustered at the month of sample level. Wild bootstrap with quarter of the year clusters (Roodman et al. 2019). P-values are reported inside square brackets.

tion leads to nearly identical incidence estimates.

1.5.4 Incidence

The empirical estimates of pass-through ($dP/d\tau$), electricity generator production ($dq_i/d\tau$), and the impact of the carbon tax on retail electricity prices ($dP^{ret}/d\tau$) are all key parameters when estimating incidence using equations 1.4 and 1.5, but they are not sufficient. In order to estimate incidence, one must also determine: average demand and supply of electricity in each electricity zone, emissions rates for each firm, Lerner Index for each electricity generator, and the zonal price over total quantity. For the quantity supplied and demanded as well as the zonal price over total quantity, I use the average from 2012 to 2014. For each electricity generator, I use detailed data on the emissions rates of each plant, e_i, from the National Transmission Network Development Plan. Lastly, I calculate marginal input costs for each firm by using heat rate and fuel price data from ACIL Tasman's report for AEMO and combine that with emissions rates and the carbon price to construct an estimate of each plant's marginal costs from 2012 to 2014.[34] For wind electricity generators, I assume a marginal cost of 0 and for hydroelectric electricity generators I use marginal cost estimates from the same ACIL Tasman report. Using data on the location of each plant and cleared generation for each period of the day, I calculate the average Lerner Index, $L_{im} = (P_i - MC_{im})/P_i$, for each plant when they are generating electricity.[35]

The results for incidence are shown in Table 1.4. Each column uses parameters estimated from a different specification outlined in the previous sections as inputs in the calculation of incidence. The parameter estimates used are presented in Panel A, while the specifications used to calculate them are shown in Panels C and D. Panel B shows

[34] I focus on this window due to data limitations prior to 2012.
[35] The price experienced by each plant i is equal to the price in their zone P_m.

the estimates of the consumer share of the burden defined as:

$$\frac{I}{I+1} = \frac{dCS/d\tau}{dCS/d\tau + dPS/d\tau}$$

If this is equal to one then customers bear all of the burden and if it is more (less) than one then customers bear more (less) than 100 percent of the burden.

Table 1.4 shows that all of the burden from the carbon price is borne by customers in the electricity market. Estimates of the consumer share of the burden range from 0.949 to 0.971, with nearly identical estimates when the impact of the carbon tax on retail electricity prices is included. Importantly, these results are consistent across various specifications and robust to when extreme wholesale electricity price events are dropped when calculating the daily average price. If anything, Figure 1.8 shows that these estimates are conservative relative to less restrictive specifications. Overall, these results shows a consistent pattern - the loss in producer surplus from increases in the carbon price is negligible compared to the loss in consumer surplus.

1.5.5 Why do customers bear all of the burden?

The results on incidence show that customers bear nearly all of the burden from the carbon price. This is the case even when incorporating pass-through from wholesale to retail electricity prices. Intuitively, this captures the fact that electricity prices rose by more for customers compared to the loss in average per-unit profit for producers. Table 1.4 highlights that a key reasons for the small loss in profits for producers is because wholesale electricity prices increased by almost the same amount as the increase in marginal costs ($\bar{e} - \bar{\rho}_s$). As previously discussed, this difference depends on two key factors - markups of marginal plants and differences in emissions between marginal and infra-marginal plants.

Table 1.4: Incidence Estimates using Tax to Retail Estimates

	(1)	(2)	(3)	(4)	(5)
Panel A. Incidence components					
Demand-weighted pass-through ($\bar{\rho}_{dw}$)	0.833	0.833	0.827	0.825	0.831
Demand-weighted retail pass-through ($\bar{\rho}_{ret}$)	1.114	1.114	1.114	1.114	1.114
Generation-weighted pass-through ($\bar{\rho}_{gw}$)	0.834	0.834	0.827	0.823	0.842
Emissions Intensity ($\bar{e}$)	0.885	0.885	0.885	0.885	0.885
Symm. Imperf. Competition Parameter (α_{hom})	-0.019	-0.013	-0.013	-0.013	-0.013
Asymm. Imperf. Competition Parameter (α_{het})	0.006	0.018	0.018	0.018	0.018
Panel B. Consumer share of burden					
Symmetric Oligopoly (Wholesale)	0.923	0.928	0.920	0.917	0.937
Asymmetric Oligopoly (Wholesale)	0.949	0.962	0.954	0.950	0.971
Asymmetric Oligopoly (Retail)	0.961	0.972	0.966	0.962	0.978
Panel C. Wholesale Pass-through Specification					
Base FE	Yes	Yes	No	No	No
Full FE	No	No	Yes	Yes	Yes
Outliers Dropped	Yes	Yes	Yes	No	No
Median Regression	No	No	No	Yes	No
Raw Prices	No	No	No	No	Yes
Panel D. Quantity regression specification					
Base FE	Yes	No	No	No	No
Full FE and Gen*Fuel	No	Yes	Yes	Yes	Yes

Notes: This table presents results for welfare incidence for electricity by component specification. Incidence is defined as the change in the consumer surplus as a share of the change in consumer and producer surplus. The imperfect competition parameter, α, represents the last term in the denominator of equation 1.4. Instead of estimating the pass-through from wholesale to retail in order to calculate the demand-weighted retail pass-through rate, this table calculates incidence using the estimates from directly estimating the impact of the tax on retail prices. Specifically, the demand-weighted retail pass-through rate is calculated using column (2) of Table 1.3. Panel A shows the estimates for each incidence component depending on the combination of regression specifications shown in Panels C and D. Panel B shows the estimate of incidence based on the components in Panel A. The first row of panel B shows the incidence estimate if pass-through from wholesale to retail is assumed to be 1, while the second row uses wholesale to retail pass-through estimates. Each column represents a separate specification used for an estimated parameter in the incidence formula. Base fixed effects are month, year, and day of week, while full fixed effects allows month and day of the week fixed effects to vary by location. For pass-through estimates this is a region, while for the quantity regressions this is the electricity generator. For the specification where outliers are dropped the cutoff is 1000 AUD/MWh, which represents less than 0.1% of prices.

Figure 1.8: All Incidence Estimates

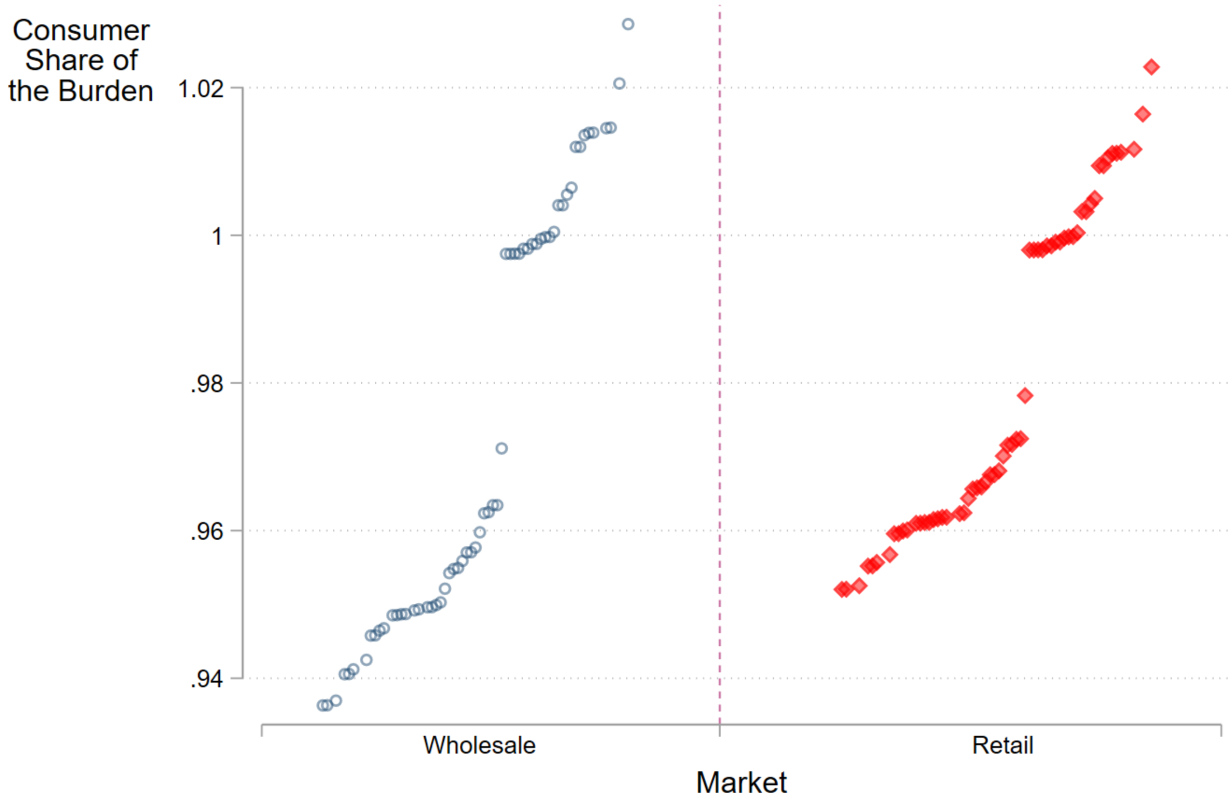

Note: This figure shows the share of the welfare loss borne by customers if the market that they purchase from is the wholesale market compared to the retail market. Incidence estimates for every combination of specifications used to estimate the impact of the carbon tax on wholesale electricity markets is shown here. This includes all specifications using a median regression compared to when the raw price data or trimmed price data are used in an OLS regression. The estimated relationship between the carbon tax and retail prices shown here uses linear time trends and region fixed effects.

To determine the role of markup adjustment, I estimate the augmented local linear regression design similar to equation 1.8, but where the dependent variable is the residual from a regression that regresses V_{it} on controls. Figure 1.9 plots the results from my preferred specification, which controls for electricity generator changes over time through generator offer price by year and day of the week fixed effects and also accounts for fluctuations in demand and supply costs through generator by projected demand and generator by coal and natural gas price controls. There is little evidence of a significant change in markups in response to the carbon tax change on July 1. This is corroborated by the full set of results shown in Table A.5, where I vary bandwidths, polynomials, and estimating equations in the first stage. I find consistent evidence of a small effect that is not statistically different from 0. In my preferred specification with a bandwidth of 30 days on each side of the cutoff and a local linear regression, the estimated effect on markups is -0.186, but far from statistically significant. These results highlight that the effect of the small increase in the carbon price on markups appears to be negligible and, if anything, slightly negative.

The fact that electricity generators did not change markups appears to be inconsistent with the level of heterogeneity among marginal electricity generators. As Figure A.9 illustrates, a wide variety of plants are marginal in the National Electricity Market with emissions ranging from 0 to well over 1. In Figure A.10, however, I split the emissions rates of marginal plants by electricity zone. Although emissions rates of marginal plants are different between electricity zones, within each zone marginal plants are frequently quite similar.[36] This local competition between marginal plants is more important when there is transmission congestion as the residual demand faced by the electricity generators does not depend on plants outside of the local, unconstrained market. The fact that markups did not change significantly in response to a change in the carbon price is thus consistent with the importance of local heterogeneity in cost shocks for markups rather than in the aggregate market.[37] The insignificant effect of the carbon tax on

[36]In New South Wales and Queensland, the marginal electricity generators are primarily black coal with emissions rates of 1, while in Tasmania the marginal plants are primarily hydro and in Victoria they are brown coal.

[37]Given the similarities between marginal electricity generators it is likely that changes in marginal plants induced by the tax are minimal. One would thus expect that the marginal emissions costs pass-through rate would yield a similar finding. Employing nearly an identical regression to Fabra and Reguant (2014), I do find this to be the case in Table A.6.

Figure 1.9: Markup RD Estimates

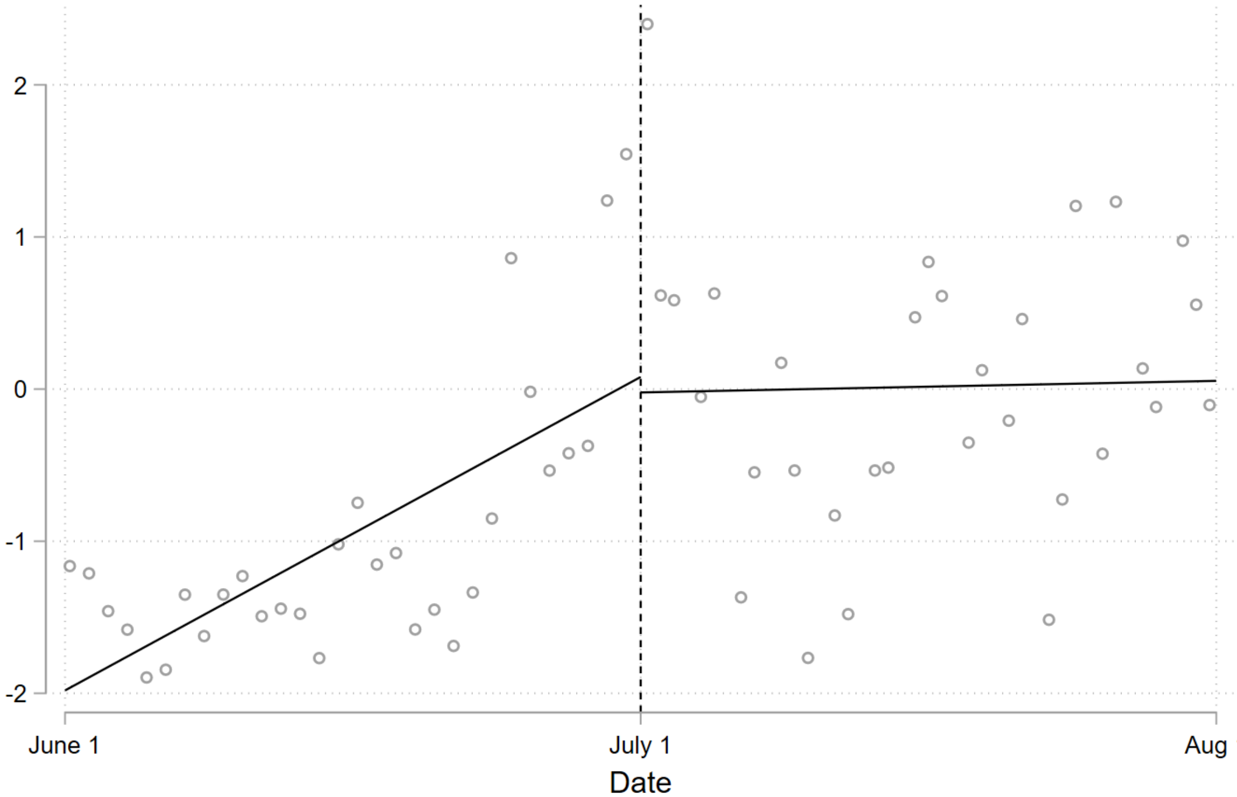

Note: This figure plots the results from an augmented local linear regression discontinuity design. In the first stage, I regress the difference between frequently marginal offer prices and their marginal emissions costs on a rich set of controls. Specifically, I control for: coal and natural gas prices by generator and offer price, quadratic relationship between projected (by the Australian Energy Market Operator) demand and each generator's offer price, and lastly generator by day of the week and generator by year fixed effects. In the second stage, I estimate a regression discontinuity with a bandwidth of 30 days where the residuals from the first stage are the dependent variable, which is what is shown here.

markups in this context thus highlights that when carbon taxes are implemented in other markets, a key factor to consider is the local heterogeneity in marginal electricity generator emissions.

Evidence of no effect of the carbon tax on markups thus suggests that the average emissions rates of marginal plants are likely similar to the emissions rates of infra-marginal plants. This is confirmed by Figure A.10, which shows an average emissions rate almost identical to that shown for the whole market in Table 1.4. The increase in prices thus matched the increase in average marginal costs because markups did not change and average emissions rates of marginal and infra-marginal plants were roughly similar.

1.5.6 Robustness Checks

Estimates of incidence shown in Table 1.4 seem to indicate that increases in wholesale electricity prices are passed on fully, on average, to retail electricity prices. To investigate this, I estimate the relationship between wholesale and retail electricity prices as:

$$r_{jt} = \alpha + \beta p_{jt} + \gamma I_{jt} + \epsilon_{jt} \tag{1.10}$$

One potential confound with estimating this equation, however, is that retail prices can have a direct impact on wholesale prices. This classic simultaneity problem means that estimates of β would be biased. To causally identify the impact of wholesale prices on retail prices, I use the carbon tax as an instrument for the wholesale electricity price. If the carbon tax has a strong effect on wholesale electricity prices and does not impact retail electricity prices through any channel then it is a valid instrument for the carbon tax.

As I discussed previously, the carbon tax did not cover markets that one would expect to be strong complements or substitutes for electricity consumption and so the exclusion restriction is likely to hold. Moreover, as Table A.7 shows, the carbon tax is a strong predictor for increases in the wholesale electricity prices and thus satisfies the relevance restriction. Although this is an instrumental variables regression with a very small sample size, the results in Table A.7 highlight that increases in the wholesale electricity price do seem to be passed on to retail electricity prices fully. This is important since it matches Table 1.4 and thus shows that the results appear to be internally consistent.

A potential concern is the fact that electricity generators knew when changes in the carbon tax were going to take place. This forewarning allows hydropower to reduce production prior to the tax change in order to increase production when wholesale electricity prices are higher after the tax change. This would bias downwards estimates of the impact that the carbon tax has on wholesale electricity prices (pass-through) and marginal electricity generator production. Alternatively, this would bias upwards estimates of the impact that the carbon tax has on hydropower electricity generation.

To ensure that the preemptive behavior by hydropower is not driving my results, I estimate a lower bound of incidence. To do this, I note that if the intertemporal constraint on production is binding for hydropower then a surprisingly implemented carbon tax would have no change on production. In other words, for all hydropower electricity generators I set $dq_i/d\tau$ equal to 0. No change in production is the lower bound for how one would expect a clean energy source to respond to a change in the carbon tax. Combined with any potential downward bias in pass-through and marginal electricity generator production, this will provide a lower bound for estimates of the consumer share of the burden from the carbon tax. As Table A.8 shows, even in this case, my estimates show consumers bearing approximately 90 percent of the burden. Importantly, this highlights that even in this case where I estimate a lower bound, the results are qualitatively similar with customers bearing nearly all of the burden from the carbon tax.

1.6 Difference between Asymmetric and Symmetric Frameworks

In the preferred estimates of incidence, Table 1.4, the estimated share of the burden borne by consumers is nearly identical when applying the symmetric framework and asymmetric frameworks. In this section, I explore whether this is a result of the particular setting of this paper or if this is generally the case.

1.6.1 Simulation Approach

The only difference between the symmetric and asymmetric frameworks is through the value of lost production, α. Abstracting from multiple market prices, in the symmetric case, α, is calculated as:

$$\alpha_{hom} = \frac{dQ}{d\tau} L \frac{P}{Q} = -\frac{dP}{d\tau} L \epsilon_D$$

while in the asymmetric case it depends on firm-specific changes in production:

$$\alpha_{het} = \frac{P}{Q} \sum_{i=1}^{n} \frac{dq_i}{d\tau} L_i$$

Comparing these two expressions, it is only in the case where there is a wedge between $\sum\limits_{i=1}^{n} \frac{dq_i}{d\tau} L_i$ and $\frac{dQ}{d\tau} L$ that estimates of incidence will differ.

Two unique features of wholesale electricity markets are that it is centrally cleared in an auction and demand is almost perfectly inelastic. Both of these features of the market could be driving the small difference between estimates of incidence when applying the symmetric and asymmetric frameworks. Returning to the expressions for α, this is for two distinct reasons. First, highly inelastic demand implies that $\sum_{i=1}^{n} dq_i/d\tau = dQ/d\tau$, is nearly 0 and, in most applications, one would expect individual changes in production to be similarly small.[38] The value of lost production, which is represented by α, is thus likely to be small in both the symmetric and asymmetric cases as there is no drop in production. Second, since the market is cleared through a central auction this implies that many of the increases and decreases in production are likely for similar, often marginal, generators. Importantly, this means that firms with increases and decreases in production likely have roughly similar ex-ante marginal costs.[39] This is evident in Appendix Figure A.11, which plots the relationship between the Lerner index and estimates of individual firm's changes in quantity in response to changes in the carbon tax. The correlation between the Lerner index and changes in quantity is so low (0.0975) that a regression between the two variables, without controls, does not find a statistically significant effect of a firm's Lerner index on the change in quantity. The weak relationship between individual firm changes in production ($dq_i/d\tau$) and the Lerner index (L_i) means that, combined with highly inelastic demand, the symmetric and asymmetric estimates of α are likely going to be similar.

The previous discussion explains why the difference between applying the two frameworks is negligible in this application. It cannot address, however, whether one would expect there to be a difference in other applications. To explore this, I will keep all other parameters in the setting constant except the change in overall and individual quantities in response to a change in the tax. Importantly, this implicitly corresponds with changes in the demand elasticity which can be written as:

$$\epsilon_D = \frac{dQ}{d\tau} \cdot \frac{1}{\frac{dP}{d\tau}} \cdot \frac{P}{Q}$$

[38]Only in particular cases will this not be true. You could get large differences between α_{hom} and α_{het} if, for example, there are off-setting large changes in production and large differences in L_i. One example of this would be all firms with large decreases in production had $L_i = 1$ while those that had large increases in production had $L_i = 0$. This is unlikely to be the case in most applications.

[39]Even plants with the same marginal costs and marginal emissions costs could respond differently to changes in the tax. The response by the firm depends also on the impact that the tax has on all of the firm's other plants.

By increasing the total change in quantity in response to a change in the tax, $\frac{dQ}{d\tau}$, while keeping all else constant, I am able to simulate the difference between the estimates under alternative distributions of changes in quantities.

The primary challenge with varying the total change in quantity in response to a change in the tax is allocating the simulated reduction in production between firms. In my primary results, I assume that for all electricity generators that experience decreases in production per dollar of the tax this would be larger if demand was more elastic. Specifically, I assume that for all electricity generators with $dq_i/d\tau < 0$ the new change in production is:

$$\frac{dq_i^{new}}{d\tau} = b \cdot \frac{dq_i}{d\tau} \tag{1.11}$$

where $dq_i/d\tau$ are my empirical estimates from specification 5 in Figure 1.6 and b is a scalar larger than 1. When b is equal to 1, the demand elasticity is that implied by original estimates, which is approximately -0.03, and as I increase b above 1 the demand becomes more elastic. One of the key advantages of this approach is that it rests on the intuitive assumption that as demand becomes more elastic, firms that had large decreases when it was inelastic will have an even larger decrease.

As an alternative, I also show results under three assumptions. First, I assume that all producers decrease their production by the same amount such that:

$$\frac{dq_i^{new}}{d\tau} = \frac{dq_i}{d\tau} - a \tag{1.12}$$

where a is a constant greater than 0. Intuitively, this amounts to a shift in the entire distribution of $\frac{dq_i}{d\tau}$ regardless of the initial estimates. As a minor change to this approach, I also estimate the difference in incidence estimates when I instead assume that all firms decrease production, but firms that have smaller markups decrease by more. In other words, I assume:

$$\frac{dq_i^{new}}{d\tau} = \frac{dq_i}{d\tau} - a - (1 - L_i) \tag{1.13}$$

so that all firms decrease by the same quantity a, but marginal plants decrease by more. Lastly, I adopt a similar approach, but only for those that decrease production. Specifically, for all electricity generators that decrease production in response to the tax, I let the decrease be larger if they are more likely to be marginal:

$$\frac{dq_i^{new}}{d\tau} = (c - L_i) \cdot \frac{dq_i}{d\tau} \tag{1.14}$$

where c is a constant greater than 2. To simulate the impact of higher demand elasticity, I will increases the constants a, b, and c, which will lead to larger changes in aggregate quantity ($dQ/d\tau$) and implied demand elasticity (ϵ_D).

1.6.2 Simulation Results

There are two key outcomes of interest from the simulations. First, taking as given heterogeneity in the Lerner Index for each firm (L_i), are the symmetric and asymmetric estimates of α still similar if demand was more elastic? Second, do the differences, if any, in α have an economically significant effect on estimates of the consumer share of the burden?

Figure 1.10 shows the effect of increases in demand elasticity on the value of lost production, α, depending on whether firms are assumed to be symmetric (α_{hom}) or asymmetric (α_{het}) under the assumption stated in equation 1.11. When demand is inelastic, the difference between the two estimates are nearly identical, but as demand becomes more elastic the wedge between the estimates increases. At an assumed demand elasticity of -3, the difference in estimates of the value of lost production, α, is approximately 0.2

Figure 1.10: Simulated Values of Lost Production

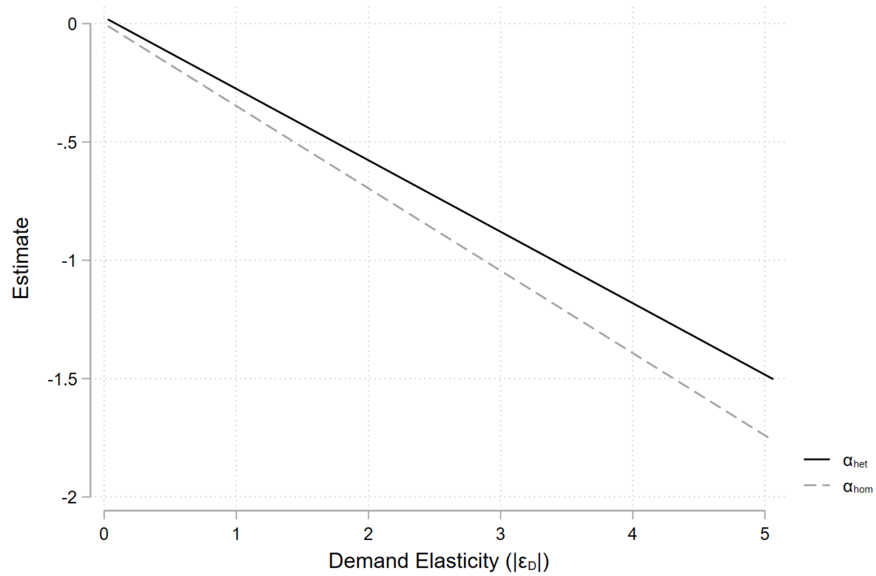

Note: This figure shows estimates of the value of lost production (α) when the implied demand elasticity is increased. This is done by simulating a larger decrease in production in response to the carbon tax ($dQ/d\tau$) and then estimating the values of lost production when firms are assumed to be symmetric, gray line, compared to asymmetric, black line. All estimates of the changes in quantities for each firm use as a baseline the preferred specification, Figure 1.6 specification 5, and the preferred pass-through estimates with the raw price data, Table 1.2 column 1.

with the symmetric case approximately 20% larger in magnitude compared to the asymmetric case.

Although the value of lost production is a key parameters in the calculation of incidence, the differences may not be large enough to have an impact on estimates of the consumer burden. Figure 1.11 shows that this wedge in estimated value of lost production has significant impacts on the difference in estimates of the consumer burden. Figure 1.11 plots the percent difference between estimates of the consumer share of the burden when applying the symmetric framework, equation 1.1, compared to the asymmetric framework, equation 1.2. The gray lines provide points of references for the numerical values

Figure 1.11: Simulated Percent Difference in Estimates of Consumer Burden

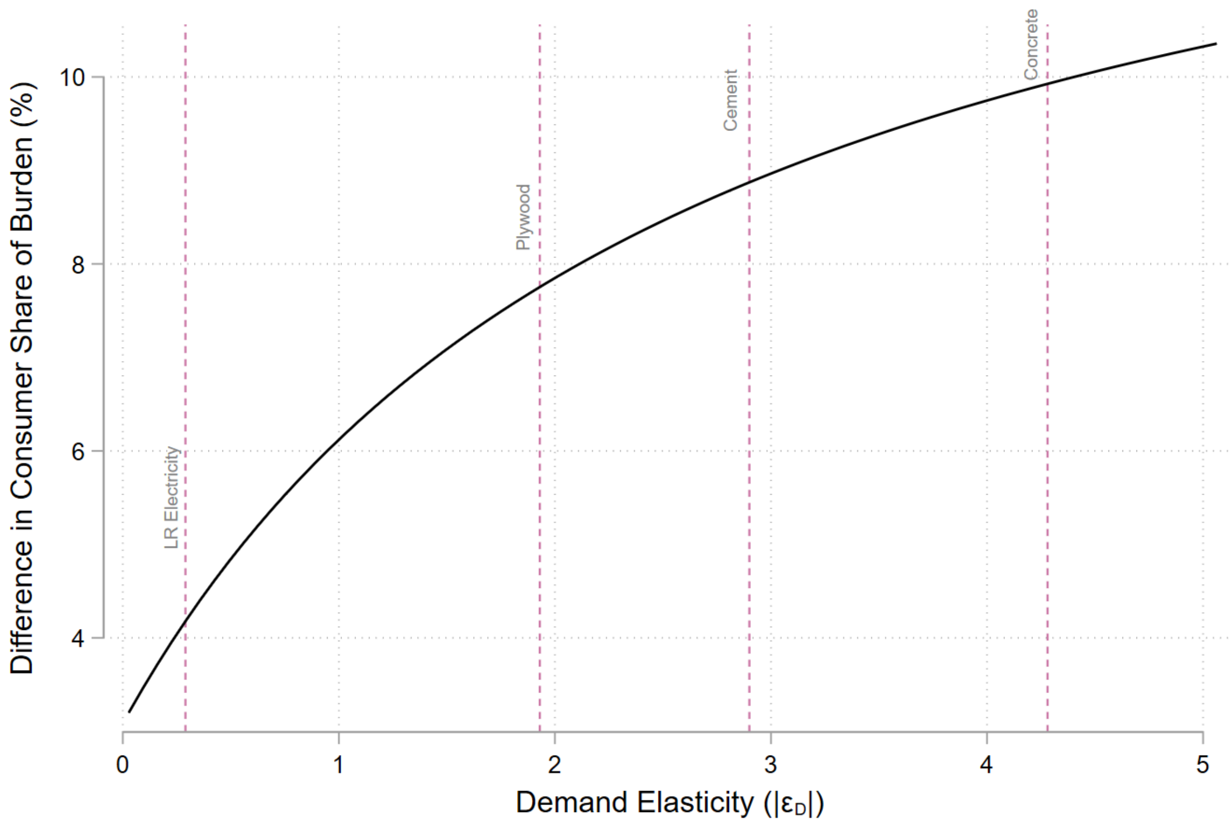

Note: This figure shows the percent difference between estimates of the consumer share of the burden when producers are assumed to be symmetric compared to when producers are allowed to be asymmetric. All estimates where the elasticity is larger than 0.03 are the result of simulations where I increase the demand elasticity from the baseline by simulating larger decrease in production in response to the carbon tax ($dQ/d\tau$). This is done by assuming that all electricity generators that decrease production in response to the carbon tax, $\frac{dq_i}{d\tau} < 0$, now decrease by an even larger amount such that $\frac{dq_i^{new}}{d\tau} = b\frac{dq_i}{d\tau}$ where b is a constant greater then 1. As a increases, the implied demand elasticity shown in the x-axis increases. I then estimate the values of lost production when firms are assumed to be symmetric (α_{hom}) compared to asymmetric (α_{het}) and incidence in both cases. All estimates of the changes in quantities for each firm use as a baseline the preferred specification, Figure 1.6 specification 5, and the preferred pass-through estimates with the raw price data, Table 1.2 column 1. All estimates in the figure assume the same market conditions estimated in the context of the Australian electricity market, but with larger quantity changes in response to changes in the carbon tax. Gray lines act only as reference points for estimates of demand elasticity in various markets. The LR electricity demand elasticity is -0.36, which lies in the middle of estimates from Deryugina et al. (2020). Plywood, cement, and concrete demand elasticity estimates are from Ganapati et al. (2020).

of the demand elasticity shown along the x-axis, but I want to emphasize that the estimates shown at those points are assuming the parameter estimates of the electricity market. Importantly, I want to highlight that as demand becomes more elastic, the difference between estimates using the two frameworks increases. At an implied demand elasticity of -4, slightly less than the demand elasticity for concrete (Ganapati et al. 2020), estimates of the consumer share of the burden in the asymmetric framework are approximately 10 percent different from the symmetric framework. The pattern of results shown in Figure 1.11 are further emphasized by similar, or larger, results under the assumptions in equations 1.12, 1.13, and 1.14, which are shown in Appendix Figures A.12, A.13, and A.14. If one assumes that all firms decrease production, with larger decreases for those that are likely to be marginal, the difference between incidence estimates can be as large as 40 percent even at an implied demand elasticity of -3. Intuitively, this reflects a higher correlation between L_i and $dq_i/d\tau$ of approximately 0.4, which leads to a larger difference between the two frameworks.

The results from these simulations highlight a key fact. In the context of this paper, electricity, estimating equation 1.1 or 1.2 leads to nearly identical results, but this is not necessarily the case in other contexts. The result of Figures 1.10 and 1.11 show that, even with the low levels of correlation between the Lerner index (L_i) and changes in quantity ($dq_i/d\tau$), the framework developed in this paper is more important in markets where demand is elastic.

1.7 Conclusion

In this paper, I empirically estimate the welfare change of electricity consumers relative to electricity generators, also known as incidence. I first derive a general framework for incidence in the context of imperfect competition with a carbon tax and heterogeneous marginal costs. I then expand on this to incorporate pass-through in zonal market and from the intermediary market, wholesale in this case, to the final goods market of retail electricity. Leveraging the exogeneity of the Australian carbon tax in combination with detailed micro data on each electricity generator, I then estimate the key parameters of the incidence formula. Using these estimate, I find that electricity customers bear 95 percent of the welfare loss from the carbon tax.

There are important caveats to this analysis. The analysis in this paper does not focus on the overall distribution of the tax burden from the carbon tax, but instead focuses only on the electricity sector. To fully estimate the distribution of the tax burden between producers and consumers, one must shift from the partial equilibrium approach adopted here to an economy wide general equilibrium approach. Moreover, this paper also does not account for the effect of the carbon tax on consumer welfare through emissions reductions or revenue recycling. Some electricity customers in Australia received direct transfers as well as reductions in the income tax, both of which mitigate the welfare loss from the carbon tax. On the other hand, large brown coal electricity generators also received a large number of free permits, which directly subsidized their operation throughout the carbon tax. If these free permits prevented the plants from shutting down then the impact that they have on incidence is ambiguous. Incorporating the impact of revenue recycling and free permits, in particular, deserves the attention of future work.

Through the framework and estimates of incidence in this paper, a few key elements are highlighted when considering other contexts. First, increases in wholesale electricity prices are likely to be passed on fully to retail electricity prices in the absence of regulatory constraints. Increases in wholesale electricity prices are reflected directly in cost increases for electricity generators. Since residential electricity demand is highly inelastic, full pass-through of the changes in marginal costs will likely occur in most settings. Regulators can institute strict controls over retail electricity price changes, but this limits the behavioral changes for electricity customers as the experienced consumer price does not fully internalize the externality cost of generation. Second, the distribution of electricity generator emissions rates plays a key role in the relative welfare loss of consumers compared to producers. Oftentimes the differences between marginal and infra-marginal profits is highlighted since this can lead to windfall profits for renewable energy generators. A key contribution of this paper, however, is that it highlights that differences between emissions rates of competing potentially marginal plants also plays a key role in incidence through the incentive to adjust markups.

Chapter 2

Paying for Integers

with Guangli Zhang

2.1 Introduction

In many markets, consumers are faced with situations where they are expected to voluntarily pay extra in the form of a tip for no additional good or service. Historically, tipping is prevalent in particular markets within the United States, such as the restaurant industry, where tips have accounted for more than $40 billion of revenue (Azar, 2008). In the early 2010s, however, cloud-based point-of-sale systems like Square, Inc. were introduced. These systems allow firms to present and customize suggested tip functions in their payment interface. As a result, consumers are increasingly encountering formal prompts and suggestions for tips in settings like coffee houses, where previously there were none. Despite the increasing prevalence in new markets and large revenue in traditional 'tipping markets', economists still understand little about the determinants of consumer tipping behavior.

In this paper, we exploit the unique setting of New York City taxi rides, where we observe high frequency, trip-level responses to preset tip suggestions. Similar to previous work on tips, we document that customers do respond to default tip suggestions. Despite the fact that default tip suggestions do not cluster at integer tip amounts, however, we find that customers have a tendency to tip integer amounts. Furthermore, customers exhibit this behavior despite the fact that tips in this setting are automatically incorporated into final prices by the credit card machine. We thus ask: do customers respond differently to tip suggestions based on whether or not the suggested tip amount is an integer and, if so, what does this reveal about human behavior?

To theoretically explain the tendency for passengers to give integer tips, we use the model of Donkor (2020) as a starting point. In this model, a passenger's preferred tip rate absent a menu (i.e., custom tip) is where the marginal costs associated with tipping more is equal to the marginal gains from smaller norm-deviation costs. When presented with a menu, she then decides if it is worth paying the cognitive costs to tip her preferred tip rate or if she would instead like to select an option from the menu, which has no cognitive costs associated with it. In this model there is no reason for clustering at integer tip amounts, so we extend the decision that passengers make by incorporating lower cognitive costs when giving custom tips that are integers and lump-sum utility gains when giving an integer tip. Both mechanisms, differential cognitive costs and lump-sum utility gains for integer tips, lead to increases in the frequency of integer custom tips. Only in the presence of lump-sum utility gains when giving integer tips, however, are customers more likely to give the suggested tip amount if it is an integer. We leverage this implication of our model, in combination with plausibly random variation in whether a customer is presented an integer tip suggestion, to provide evidence on whether the pattern of clustering at integers is driven, in part, by customers experiencing lump-sum utility gains from tipping an integer amount.

Endogeneity in prices, tipping schemes, and consumer purchasing decisions can make studying consumers' tipping behavior challenging. It could be the case, for example, that integer tip suggestions only occur when customers purchase a certain combination of goods. Alternatively, it is possible that customers are more likely to select integer tip suggestions because it is easier to calculate the total amount that they must pay. If these customers differ from those that purchase other combinations of goods, then this would bias estimates of the relationship between integer tip suggestions and tipping behavior.

In the context of New York City taxi trips, however, we are able to overcome many of the endogeneity concerns due to the fact that 1) tips are automatically added to the fare amount when selected off the menu and 2) plausibly random variation in tip suggestions depending on the credit card payment machine and surcharges throughout the day. Every passenger that pays with a credit card during the time period that we study is faced with a menu of three tip suggestions: 20, 25, and 30 percent. The total that is used to calculate these suggestions, however, varies between the two credit card payment machines as one does not include the Metropolitan Transportation Authority (MTA) tax of $0.50 while the other does. Importantly, this means that two customers with identical trips, i.e., same date and distance, will receive slightly different tip suggestions depending on the credit card payment machine the taxi is using, which is not evident from the exterior. One of these customers could thus be "treated" with an integer tip suggestion, while the other is presented a nearly identical non-integer tip suggestion. Since surcharges change by day of the week and time, the credit card payment machine that presents integer tips to customers changes thereby allowing us to isolate the effect of an integer tip suggestion from potential confounds like differences in tipping behavior throughout the day or by credit card machine.

We leverage this variation in the occurrence of integer tip suggestions to examine whether customers' behaviors are consistent with a model where they experience direct utility benefits from giving integer tips. Across a variety of specifications and estimation strategies we find consistent support for this hypothesis in the form of increases in take-up of default tip suggestions and tip rates. In our preferred specification where we control for average differences in tipping behavior by driver, hour of that date, and pickup and drop-off census blocks, we find that the probability a passenger selects the default option increases by more than 21 percentage points and tip rates increase by more than 0.6 percentage points. Intuitively, the increase in tip rates is due to the fact that, in our context, passengers tend to give custom tips smaller than all menu options. As passengers switch from custom tips to selecting an option from the menu, this leads to an increase in average tip rates.

The likelihood that a customer faces an integer tip suggestion is jointly determined by the interaction between the fare rate and the tip rate used for the tip suggestions. Given that customers' tipping behaviors respond to integer tip suggestions, a change in prices can indirectly impact the likelihood of integer tip suggestions and with this, revenue. We explore the magnitude of this effect using an increase in the fare rate from 40 to 50 cents in September 2012 that increased the probability of integer tip suggestions from 3% to 21%. When we decompose the effect of the fare rate change on revenue, our estimates suggest that the increase in integer tip suggestions after the fare change led to an increase in revenue of approximately 1.4 cent per trip. With over 170 million taxi trips and 41,000 unique drivers this leads to a transfer of 2.38 million dollars from riders to drivers in the year following the policy change.

Our paper is closely related to the literature that documents clustering around integers or round numbers in other domains of individual decision making (E. J. Allen, Dechow, Pope, & Wu, 2017; Lynn, Flynn, & Helion, 2013).[1] People's tendency to use integer or round numbers is commonly associated with lower cognitive cost (Isaac, Wang, & Schindler, 2020; Schindler & Wiman, 1989) or lower trading negotiation cost (Harris,

[1] Round numbers refer to integers that end with 5 or 0.

1991). Although several studies have suggested that this clustering pattern can be rationalized with people's direct preference towards integers or round numbers, there is limited causal evidence. Our paper contributes to this literature in two ways: first, we document a similar pattern of clustering at integer values in the context of taxi tipping; second, we provide theoretical underpinning and causal evidence for this behavior. Specifically, our finding is consistent with a previously under-explored mechanism that suggests people derive direct utility gains from giving integer tips.

Our paper also relates to the strand of literature that examines the potential drivers of tipping behaviors. The literature has offered causal evidence for a number of mechanisms: for example, customers' tipping decisions can be affected by the default suggestions (Alexander, Boone, & Lynn, 2021; Haggag & Paci, 2014; Hoover, 2019), their compliance to social norms (Donkor, 2020; Thakral & Tô, 2019) and their degrees of social preferences (Azar, 2007; Chandar, Gneezy, List, & Muir, 2019). Similar to this literature, we offer causal evidence for an underlying determinant of consumer tip behavior, utility gains from tipping integers. The modeling approach of this paper, however, is mostly related to Donkor (2020) who focus on estimating parameters for the optimal tipping menu in the presence of customers who conforms to social norms. We extend his model by introducing utility costs/ benefits that are associated with integer tips. Our model enables us to further decompose the mechanism that affects tipping behavior, particularly as it relates to the tendency to tip integers.[2] Our paper thus contributes

[2]In a related note, our paper offers a new insight that could potentially reconcile the conflicting findings on the magnitude of the default effect in the context of taxi tipping. Using quasi-experimental variation in tip suggestion, Haggag and Paci (2014) documents increase in default raises average tip rate significantly. On the other hand, Chandar et al. (2019) ran a field experiment that manipulates default tip suggestions via Uber and only finds a moderate effect from increasing default options. Chandar et al. (2019) attributes this difference to the reduction in norm compliance under no monitoring (in the case of Uber tipping). The integer effect uncovered in our paper offers an additional explanation to this seemingly conflicting evidence. Specifically, the default change in Haggag and Paci (2014) is accompanied by the increase in the occurrences of integer tip suggestions, whereas all default options offered in Chandar et al. (2019)'s study are in integer terms already. Therefore, the estimated default effect from Haggag and Paci (2014) is a combination of integer and default effect which is indeed greater

to our understanding of the determinants of tipping behavior by providing additional causal evidence that passengers respond to integer tip amount suggestions in a manner consistent with direct utility gains from integer tips.

The rest of our paper is structured as follows. Section 2.2 describes the institutional setting, our dataset and the sampling restrictions. Section 2.3 presents descriptive evidence and two models of tipping behavior. Section 2.4 describes the variations tip suggestions and our main econometric specification. Section 2.5 presents customers' estimated responses to integer tip suggestions. Section 2.6 discusses the implications of varying fare rate and tip suggestions on revenue. Section 2.7 concludes.

2.2 Context and Data

We use data provided by the Taxi and Limousine Commission (TLC) of New York City to estimate the effect of integer tip suggestions on tipping behavior and driver revenue. As of 2008, the entire taxi fleet was outfitted with new equipment that allowed customers to pay using credit cards and also the electronic collection of trip data. Nearly the entire fleet used equipment provided by either Creative Mobile Technologies (CMT) or VeriFone Incorporation (VTS).[3] Taxi cabs equipped by both of these vendors had a Passenger Information Monitor (PIM) which, at the end of a trip, displayed a payment screen. At this point, the devices show a tip menu to passengers who pay with credit cards. Passengers can then choose to give a tip based off the menu options, manually enter in an amount, or provide a separate cash tip.

than Chandar et al. (2019)'s estimated 'net' default effect.

[3]We do not use data from a third vendor, Digital Dispatch Systems, which accounted for less than 5% of electronic transmission devices in use in 2010.

2.2.1 Context

For standard rate fares, passengers are charged \$2.50 and a \$0.50 Metropolitan Transportation Authority (MTA) tax upon entering the cab. The fare increases by an additional \$0.40, or \$0.50 after September 4, 2012, for every fifth of a mile or for every minute where the vehicle travels less than 12 miles per hour. Throughout the period of our analysis, there is a night surcharge of \$0.50 for trips between 8 PM and 6 AM and a \$1.00 surcharge for trips between 4 and 8 PM on weekdays.

At the end of each trip, passengers are shown trip expenses through the touch-screen payment device. Passengers that pay with a credit card are then presented with a tip menu that varied by vendor over time. An example of this screen for a CMT outfitted vehicle in 2012 is shown in Appendix B1. Based on the selection of the passenger for the tip, the total is calculated and the passenger proceeds with payment. If the taxi uses a CMT device, the tip menu calculates tips on the total fare, which includes the base fare, MTA tax, tolls, and any surcharges. Alternatively, for a VTS device, tips are calculated using the base fare and the surcharge, but does not include tolls or MTA tax. In Figure 2.1, we show the menu of tip suggestions for CMT and VTS devices over time. Prior to February 9, 2011, customers in taxicabs with CMT devices were presented with tip suggestions that were 15, 20, and 25 percent. From February 9, 2011, onward all options on the CMT menu went up to higher tip percentages of 20, 25, and 30. For VTS devices, tip suggestions changed in January of 2012. Prior to that month, they offered a tip menu of dollar amounts (\$2, \$3, and \$4) if the base fare and surcharge was under \$15, and suggestions of 20, 25, and 30 percent for larger fares. After that month, VTS offered only the percentage choices (20, 25, and 30), regardless of the trip fare.

Figure 2.1: Timeline of Fare and Tip Suggestion Changes

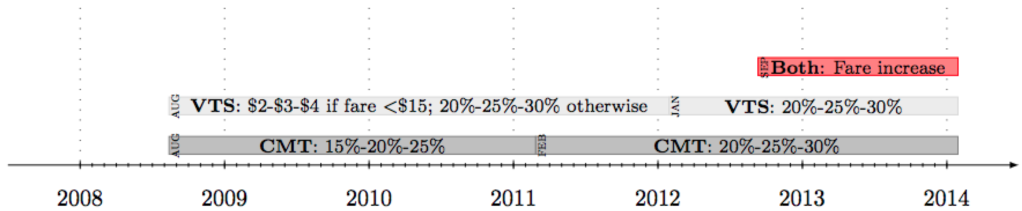

Notes: This figure shows the timing of tip and fare rate changes. NYC taxi cabs were equipped with electronic payment systems around August, 2008. At the beginning, VTS implemented a \$-% hybrid tip suggestion menu: the tip prompt is programmed to display 2, 3, and 4 dollars of suggestions if the rate fare (surcharge + fare) is less than \$15, and 20, 25, 30 percent if otherwise. On the other hand, the default menu for CMT was 15, 20, and 25 percent. On February 9, 2011, CMT increased their default suggestion to 20, 25, and 30 percent. On the week of January 22, 2012, VTS removed the \$ tip suggestions for rate fare below \$15 and set their tip suggestion to 20, 25, and 30 percent. On September, 2012, fare rate increased from 40 cents to 50 cents per one fifth of a mile.

2.2.2 Data

Our data consists of trip (ride) level data on all tax rides in New York City and surrounding counties from 2010 to 2013. For each trip, our data records the date, time, and geographic location of the pickup and drop-off. Each observation is recorded with a unique medallion number and a taxi driver license number. These numbers identify a unique cab and driver for any given year, but cannot be used to identify drivers or cabs across years. In addition, the equipment records information on trip time, trip distance, fare amount, tolls, tax, surcharge, rate code, and payment method. For all customers that pay digitally when using a credit card, we observe the tip entered into the credit card machine. Importantly, however, we do not observe tips for trips paid with cash, and we cannot interpret manually entered tips of 0 when paying with a credit card as a tip of 0.

To account for potential differences in customer characteristics, we use data from the American Community Survey's 5 year estimates (2006-2010), which consists of census tract level summary statistics. We leverage the GPS coordinates for each pickup and drop-off location to assign each trip pickup and drop-off census tracts. We then merge this with the ACS census tract variables so that we can characterize the median income of where a customer is picked up and dropped off.

We take many of the same steps to cleaning the data that have been used in the previous literature, see Haggag and Paci (2014). Since we do not observe tip information for trips or tips paid by cash, we drop these and focus on trips paid with credit cards that have positive tips in our analysis.[4] In addition, our primary analysis focuses on all trips that use standard rate fares. We do this in large part, since our primary results leverage plausibly exogenous variation in tip suggestions present in the standard rate fare, which is not present with all other rate fares.[5] The conclusions from our analysis, however, do not change when including trips with all rate fares. To ensure that our results are not influenced by drivers changing between vendors, we drop all drivers that change vendors within the same year. Similar to Farber (2015), for simplicity we focus on a random sample of drivers in all of the analysis that follows. Specifically, since we cannot track drivers across years, we use a sample of 2,000 random taxi driver and car pairs for each year.[6]

In our primary analysis, we utilize variation in the decimal places of a constant menu of tip suggestions, 20, 25, and 30 percent. Our preferred subsample focuses on the time

[4]In our sample, about 55% of the payments were made by cash. The differences between trips with cash and credit payments are shown in Table B1. The table does not highlight the fraction of zero tip trips that are paid with credit cards. To highlight that this is a small fraction, we include these trips in our primary descriptive figures but we will exclude them in the regression analysis.

[5]The rate for trips between JFK to Manhattan, for example, is fixed and would introduce non-random variation in tip suggestions.

[6]We include the detailed data refinement procedure in Appendix B.1. To ensure that our results are robust to the larger dataset, we drew a second random sample. The conclusions are identical regardless of the sample we use.

window from February to August of 2012, where all standard fare rides were subject to the same rate fare and menu of tip suggestions, regardless of vendor. This offers the key advantage of a single distribution relating rate fare to tip suggestions that all customers are subject to for a significant length of time.

2.3 Tipping Behavior

Previous research (e.g., Haggag and Paci 2014) has documented that customers respond to tipping menus by changing tipping behavior. Little attention, however, has been paid towards patterns of customer tipping behavior and the implications for tipping rates. In this section, we will first document descriptive evidence on patterns of customer tipping behavior. We will then introduce a theoretical model to guide our empirical analyses.

2.3.1 Descriptive Evidence of Tipping Behavior

A typical taxi ride experience ends with tip payments. On the payment screen, taxi passengers are often prompted with three tip suggestions and a number pad that allows them to enter any non-negative custom tip amounts. It is well documented that passengers' tip decisions are influenced by defaults and menus. Before analyzing the distribution of these choices, we first want to analyze whether there are any differences between VTS and CMT trips. Table 2.1 presents the summary statistics at the trip level for our preferred subsample, split by CMT and VTS. Although there do not appear to be differences in the details of the trips (e.g., distance or time), there are differences in average tipping behavior. Passengers that ride in vehicles with VTS equipment tend to give a higher tip rate in part due to selecting options from the menu at a higher rate. This is likely due to differences in equipment or presentation of the tip suggestions, which

Table 2.1: Summary Statistics by Trip (Ride): Feb–Aug 2012

	(1) VTS	(2) CMT	(3) Difference
Fare Amount	9.48	9.48	-0.00
	(5.00)	(4.87)	(0.01)
Tip Amount	1.86	1.93	0.07***
	(1.35)	(1.27)	(0.00)
Tip Rate	0.20	0.19	-0.00***
	(0.12)	(0.09)	(0.00)
Trip Length (in minutes)	12.08	12.13	0.05**
	(7.27)	(7.69)	(0.02)
Trip Distance (in miles)	2.54	2.54	-0.01
	(2.10)	(2.12)	(0.01)
Zero Tip	0.03	0.02	-0.02***
	(0.18)	(0.13)	(0.00)
Pr(Select 'low' default)	0.42	0.37	-0.05***
	(0.49)	(0.48)	(0.00)
Pr(Select 'middle' default)	0.13	0.11	-0.02***
	(0.33)	(0.31)	(0.00)
Pr(Select 'high' default)	0.05	0.04	-0.02***
	(0.23)	(0.19)	(0.00)
Observations	358,416	351,643	710,059

Notes: This table presents the summary statistics for the random sample of 2,000 taxi drivers during the time period of our main study: February to August 2012. During this period of time, the % tip suggestions are identical to CMT and VTS in all trips: 20, 25 and 30 percent. Tip rate is defined as the tip amount divided by the total fare excluding the tipped amount. *rate fare* (the value F_i used for tip computation) is defined differently for CMT and VTS. For CMT: Rate Fare = fare + surcharge + mta tax + tolls; For VTS: Rate Fare = fare + surcharge. Standard deviations are in parenthesis.

has been highlighted by previous research Hoover (2019) and will not be the focus of our analysis.

As illustrated in Figure 2.2, over 40% of the tips were made at the default options under a menu of tip suggestions at 20, 25, and 30 percent. However, the distribution of tip rate does not provide a holistic perspective of customers' tipping behaviors as it

Figure 2.2: Distribution of Tip Rate and Tip Amounts: Feb – Aug 2012

(a) Distribution of Tip Rate

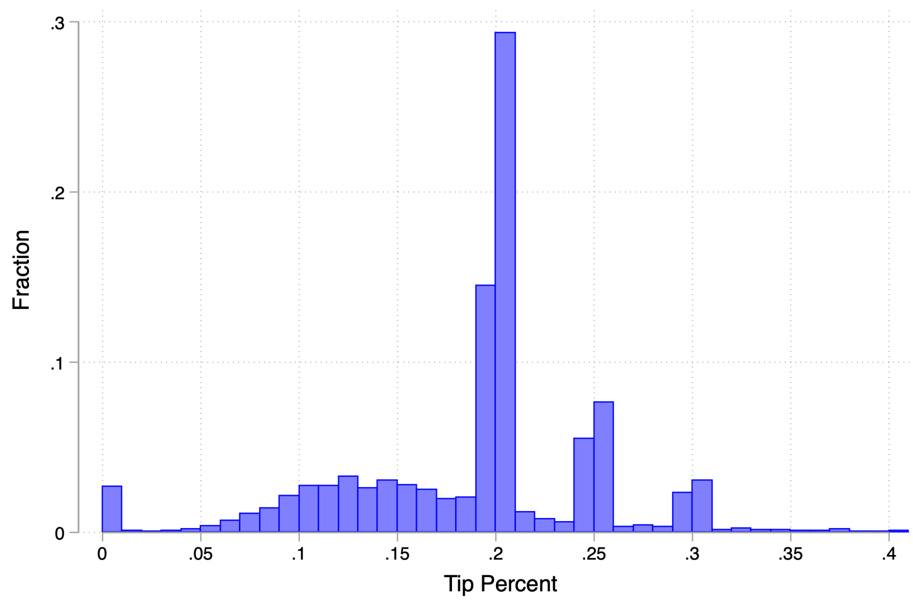

(b) Distribution of Tip Amount

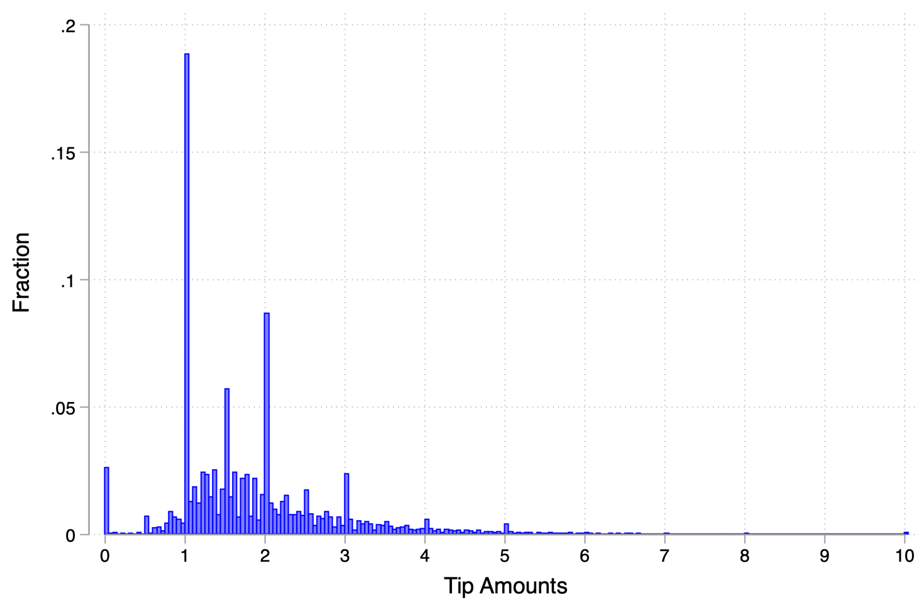

Notes: Panel (a) shows the distribution of tip % for all non-airport trips that were paid by credit card. Panel (b) figure shows the distribution of tip amounts for all non-airport trips that were paid by credit card. Extreme tip rate ($> 99^{th}$ percentile) are excluded from the figure. Tip rate is defined as the tip amount divided by the total rate fare.

ignores potential patterns that might exist in the nominal tip values. Indeed, when we plot the raw distribution of tip amounts in panel (b) of Figure 2.2, we observe clustering of tips at integer values. This is unlikely since the second decimal place of tip suggestions are equally likely to be an even number as we show in panel (a) Figure 2.3.[7] Although the second decimal places are equally likely, it is evident in panel (b) of Figure 2.3 that passengers are more likely to tip the suggested amount when the low suggestion is an integer. The clustering of tip amounts at integers appears to be driven, in part, by this increased tendency for passengers to select default integers when they are integers in combination with custom tips at integer values, as is evident in Figure B3.

Overall, a visual inspection of aggregate tipping behavior suggests that (1) customers respond to default suggestions, (2) customers tend to tip at integer values, and (3) the tendency to give integer tips is evident in custom and, to a lesser extent, default tips.

2.3.2 Models of Tipping Behavior

Given the large number of taxi drivers, we will model tipping behavior as being primarily influenced by the pressure of social norms Azar (2007) instead of strategic incentives Azar (2008).[8] Following Donkor (2020), consider a passenger i that gives a tip of $t_i\%$ at the end of her taxi ride that costs F_i. She believes that the socially accepted tipping rate to give based on the ride is $T_i\%$, which can vary by passenger. If her chosen tip rate is different than what she believes is the socially accepted tipping rate, then she incurs a norm-deviation cost of $v(T_i, t_i)$. Assume that for any fixed T_i, the norm-deviation cost of $v(T_i, t_i)$ is convex with respect to t_i with a minimum at $t_i = T_i$. When making her tipping decision she is presented a menu of tipping options D, which consists of a variety of suggested tipping percentages. Without loss of generality, denote the preferred

[7] The distribution of tip suggestions for each of the options is shown in Figure B2

[8] There are over 10,000 Yellow taxis in New York City, which minimizes the potential for repeated passenger and driver interactions.

Figure 2.3: Distribution of Second Decimal Places

(a) Low Tip Suggestion

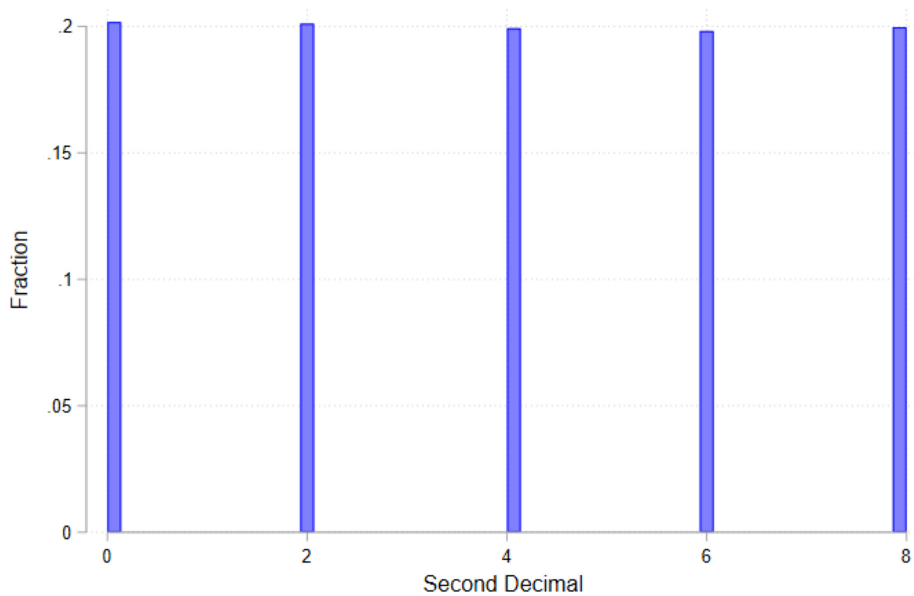

(b) Fraction Tipping the Suggested Amount

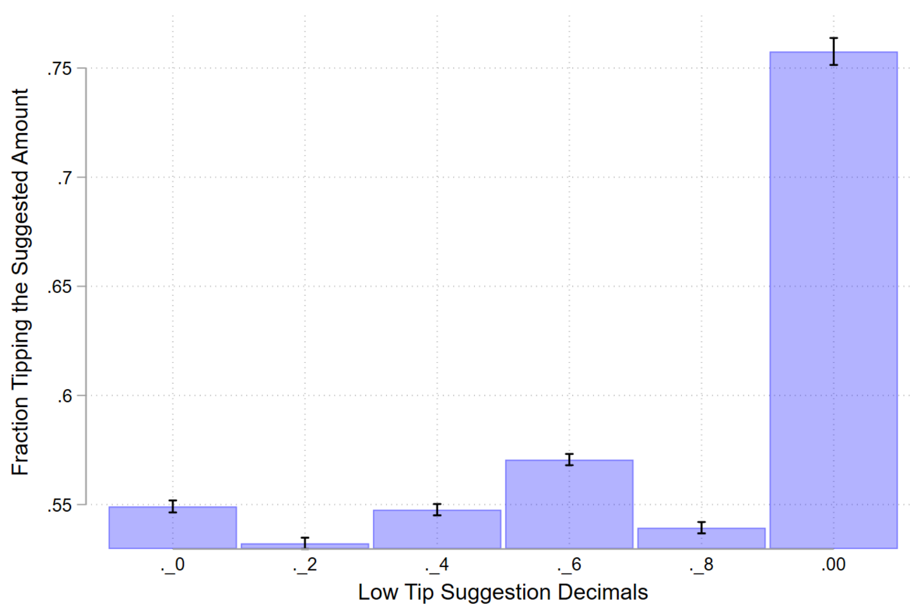

Notes: Panel (a) presents the distribution of second decimal places for tip amounts provided by the default suggestions. The pattern indicates that 0, 2, 4, 6, and 8 are approximately equally likely to appear in the tip suggestions. Panel (b) shows the fraction of customers that tip the suggested amount for each second decimal place of the low tip suggestion. The pattern indicates that the tip rate is significantly higher when the low tip suggestion ends with ".00", i.e., is an integer, compared to other suggestions.

option out of the menu for customer i as t_i^D. In order to choose an option that is not on the menu, she incurs a cost c_i that reflects the cognitive costs associated with finding her ideal tip percentage. The utility maximization problem for passenger i can be written as:

$$Max_{t_i} U = -t_i F_i - v(T_i, t_i) - c_i \cdot \mathbb{1}\{t_i \neq t_i^D\} \tag{2.1}$$

The first term represents the passenger's expenditure. The second term represents the cost of deviating from her perceived socially accepted tipping rate, T_i, and the last term captures the cost of computing a tip not presented in the tip menu.

Since the cognitive cost for all custom tip options are the same, the utility-maximizing custom tip rate is the tip rate that maximizes the first two terms. Given the assumption on the functional form of $v(T_i, t_i)$ the utility-maximizing custom tip satisfies:

$$\frac{\partial v}{\partial t_i} = -F_i \tag{2.2}$$

Intuitively, this shows that she will increase her tip rate until the marginal return of reducing the norm deviation cost, $\frac{\partial v}{\partial t_i}$, is equal to the marginal cost of increasing the tipping rate, $-F_i$. For now, denote the custom tip that solves equation (2.2) as t_i^C. As the left panel from Figure 2.4 shows, this means that, even absent cognitive costs, she will not give the socially accepted tipping rate, but will instead "shade" downwards and give a custom tip rate less than T_i.

Working backwards, the passenger then decides if she will give the custom tip or instead choose a default option from the menu. It is only worth the cognitive cost of

Figure 2.4: Individual's Utility Maximization under Baseline and Extended Models

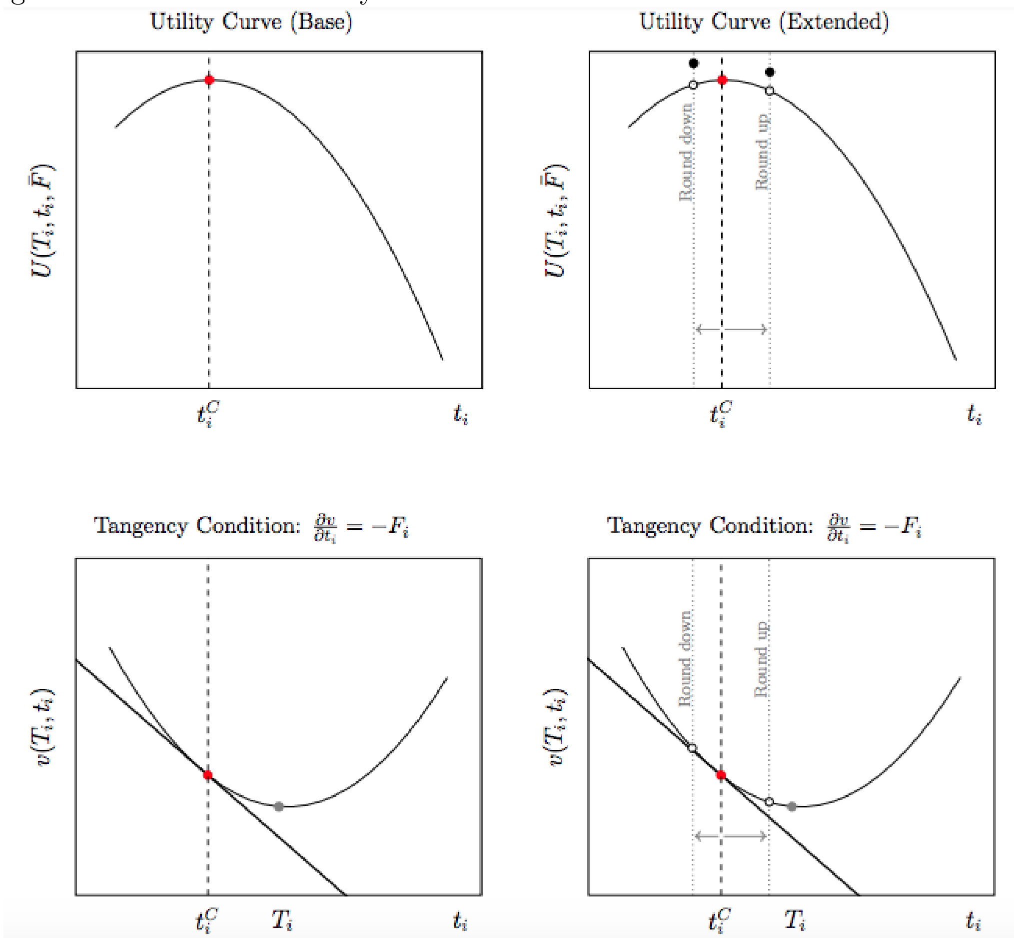

Notes: The left and right panels compare and contrast a typical passenger's utility maximization decision under the baseline model (Donkor, 2020) and the extended model. The top panel presents utility curves. The bottom panel presents the corresponding tangency condition. In particular, the convex function represents norm deviation cost ($v(T_i, t_i)$) and the downward sloping line represents the cost of increasing tipping rate. Under the baseline model, the passenger solves the utility maximization by choosing $t_i = t_i^C$ at the tangent point. Under the extended model, given the utility function is not continuous at integer tip amounts, the tangency condition might not lead to a global maximum. In this specific example, it is optimal for the passenger to *Round* t_i^C *down* to the nearest integer.

calculating and manually entering the custom tip if:

$$\underbrace{[-t_i^C F_i - v(T_i, t_i^C)]}_{U(t_i^C) \text{ if } c_i = 0} - \underbrace{[-t_i^D F_i - v(T_i, t_i^D)]}_{U(t_i^D)} > c_i \tag{2.3}$$

The left side of the equation captures the utility gains from manually entering a custom tip relative to selecting a default option if cognitive costs were 0. A passenger compares this to the cognitive costs on the right side, and then decides if the custom tip is worth calculating. All else equal, passengers are more likely to select custom tips if their cognitive costs are low or, alternatively, if they strongly prefer the custom tip to the default options.

Given the utility problem presented in equation (2.1), it is difficult to explain the pattern of tips at integer values. Tipping integer values represent different tipping rates across trips so t_i^D does not naturally cluster at integer values. In addition, it is unlikely that utility-maximizing custom tips that satisfy equation (2.2) would lead to disproportionately more integer custom tips relative to non-integer custom tips. To better explain the concentration of tips at integer values, we will now propose two extensions to the model.

First, it is possible that passengers give integer tips because it decreases the cognitive costs associated with computing the ideal tip. In other words, if a passenger believes that the suggested options are too high, she might choose a lower tip that is close to the ideal tip percentage, but is an integer and is thus less cognitively costly. To incorporate this into the passenger's problem, let there be a lower cognitive cost $c_i^{int} < c_i^{non}$ when the selected tipping choice is an integer. Define the difference in cognitive costs as $\alpha_i = c_i^{non} - c_i^{int} > 0$, which is passenger specific. The second potential mechanism behind customers tipping integer amounts is that passengers, in general, feel more comfortable giving integer tips. We model this as a lump-sum utility gain, b_i, whenever a passenger tips an integer, regardless of whether it is on the menu or not. We can then write the utility maximization problem for passenger i as:

$$Max_{t_i}U = -t_iF_i - v(T_i,t_i) - \mathbb{1}\{t_i \neq t_i^D\}[c_i^{non} - \underbrace{\alpha_i \cdot \mathbb{1}\{t_iF_i \in \mathbb{Z}\}]}_{\text{Reduced Cognitive Costs}} + \underbrace{b_i \cdot \mathbb{1}\{t_iF_i \in \mathbb{Z}\}}_{\text{Integer Utility Gain}} \quad (2.4)$$

which nests the previous model of passenger utility, shown in equation (2.1), but we now allow for integer tips to directly impact utility as a lump-sum utility gain and lower cognitive costs when giving a custom tip.[9]

The inclusion of differential cognitive costs and lump-sum utility gains when giving an integer tips impact the utility-maximization problem in two key ways. First, in the model presented in equation (2.1), the choice of custom tip rate is where the marginal return of reducing the norm deviation cost, $\frac{\partial v}{\partial t_i}$, is equal to the marginal cost of increasing the tipping rate, $-F_i$. As the right panel from Figure 2.4 shows, however, this need not be the case in the extended model. The tip rate that satisfies equation (2.2) might not be utility-maximizing if the tip amount is not an integer. Intuitively, this is because the benefits from an integer tip suggestion, $b_i + \alpha_i$, can outweigh the lower utility from not equating the marginal return of reducing the norm deviation cost to the marginal cost of increasing the tipping rate.

To show this more formally, define the custom tip rate that satisfies equation (2.2) as t_i^{non} and the preferred custom integer tip rate of t_i^{int}, which has a lower cognitive cost.[10] For arbitrary benefits and additional cognitive costs b_i and α_i, she will choose to give a

[9]Figure B4 illustrates an passenger's decision process under the extended model.

[10]Intuitively, given the functional form assumptions on $v(T_i,t_i)$ the customer will have a preferred integer custom tip rate that rounds up or down from $F_it_i^{non}$.

non-integer custom tip that satisfies equation (2.2) if:

$$\underbrace{[-t_i^{non}F_i - v(T_i, t_i^{non})]}_{U(t_i^{non}) \text{ if } c_i^{non}=0} - \underbrace{[-t_i^{int}F_i - v(T_i, t_i^{int})]}_{U(t_i^{int}) \text{ if } c_i^{int}=0} > \underbrace{c_i^{non} - c_i^{int}}_{\alpha_i} + b_i \tag{2.5}$$

The left-hand side represents utility gains from giving the preferred non-integer tip, which satisfies equation (2.2), relative to the integer tip. If this outweighs the benefit of giving an integer tip, shown on the right-hand side, then she will give the non-integer tip. As the benefits of the integer tip increase, she is increasingly likely to prefer an integer custom tip. Alternatively, as the benefits from an integer tip approach 0, she is more likely to give the custom tip rate that satisfies equation (2.2), t_i^{non}.

The second way that α_i and b_i impact how a passenger tips is through the decision between custom and default tips. Denote the preferred custom tip as t_i^C, which need not satisfy equation (2.2), and define I^C and I^D as indicator variables equal to one if the custom and default tip rates lead to integer tip amounts. When choosing between the custom and default tip rate, she will give the custom tip rate if:

$$\underbrace{[-t_i^C F_i - v(T_i, t_i^C) + I^C \cdot b_i]}_{U(t_i^C) - c_i^{non}} - \underbrace{[-t_i^D F_i - v(T_i, t_i^D) + I^D \cdot b_i]}_{U(t_i^D)} > \underbrace{c_i^{non} - \alpha_i \cdot I^C}_{\text{Cognitive Costs}} \tag{2.6}$$

where the left-hand side represents gains from giving custom tips without considering the cognitive costs. If this is larger than the cognitive costs on the right-hand side, then she will choose to "pay" the cognitive cost for the custom tip rate, t_i^C.

The gains and cognitive costs associated with the custom tip rate now depend on whether the preferred custom tip rate leads to an integer tip, but also, importantly, on whether the default tip suggestion is an integer. For a given default tip rate, a small change in the fare that leads to an integer default tip suggestion sharply increases the utility of the default tip option based on the magnitude of b_i. If b_i is small, then whether or not the default tip is an integer will likely have no effect on the choice between a custom or default tip rate. If it is large enough, however, then this could switch customers away from a custom tip rate towards the default tip option.[11] Importantly, this highlights that the impact of integer default tip suggestions depend on the magnitude of b_i instead of α_i.

In summary, assuming b_i is larger than 0 on average, we have the following hypothesis:

Hypothesis 1: *When percent-based tip suggestions lead to integer tip amounts, passengers are more likely to choose the default option.*

The intuition for this hypothesis is evident in equation (2.6). If $b_i > 0$, then when $I^D = 1$ (integer tip amount suggestion) the left-side decreases thereby lowering the likelihood the passenger gives a custom tip and, in turn, increasing the likelihood a default option is chosen. A natural implication of our first hypothesis in this setting where custom tip rates, t_i^C, tend to be lower than the default tip rates, t_i^D, i:

Hypothesis 2: *Average tip rates of passengers are higher when percent-based tip suggestions lead to integer tip amounts.*

This hypothesis falls naturally from Hypothesis 1 combined with the fact that custom tip rates tend to be lower than the default tip suggestions. Passengers that normally would give a lower custom tip rate switch to the default tip suggestion, which leads to a higher average tip rate. In the following section, we detail how we leverage our empirical setting to test these hypotheses.

[11]There is no obvious reason that utility from custom tips would have a similarly sharp change to integer default tip suggestions. We explore this, however, in Appendix B.2 where we parameterize the utility function. We find no response in custom tip utility, including cognitive costs, when there are integer default tip suggestions.

2.4 Identification Strategy

To identify what drives the clustering of integer tips, we need to combine the predictions obtained from the theoretical models with random variation in the frequency of integer tip suggestions that we observe in the raw data. In this section, we first describe the quasi-random variation in the frequency of integer tip suggestions across the two vendors. We then present the econometric specifications for our main analysis.

2.4.1 Variation in Tip Suggestions

From February to August of 2012 all standard fare rides where customers paid with a credit card, regardless of the vendor or fare, were presented with tip suggestions of 20, 25, and 30 percent. This means that for every fifth of a mile or for every minute where the vehicle travels less than 12 miles per hour, the three options on the menu of suggested tips increase by \$0.08, \$0.10, and \$0.12, respectively. With a base fare of \$2.50, this means that the lowest tip suggestion on the VTS menu, $\gamma_{i,j}^{VTS} = 0.2 \cdot F_i$, between pickup location i and drop-off location j can be defined as:[12]

$$\gamma_{i,j}^{VTS} = 0.50 + 0.08x(d, mph) + 0.2s \tag{2.7}$$

where x is a function of distance (d) and speed (mph), which implicitly depend on the locations i and j. In addition, the tip suggestion depends on s, which is a categorical variable equal to \$0.00 if there is no surcharge, \$0.50 if it is a night surcharge, and \$1.00 if there is a peak weekday surcharge. Alternatively, since CMT trips include the MTA tax (\$0.50) and tolls when calculating the tip percentage, the lowest tip suggestion, $\gamma_{i,j}^{CMT}$,

[12] We focus on the lowest option on the menus here since it is the option that passengers select most frequently.

can be defined as:

$$\gamma_{i,j}^{CMT} = 0.60 + 0.08x(d, mph) + 0.2(s + \tau) \tag{2.8}$$

where the base tip suggestion increases by \$0.10, reflecting the MTA tax, and tip suggestions now depend on the cost of tolls, τ.

A key implication from these two formulas is that, at any given point in time the probability of an integer is different depending on the vendor and whether or not there is a surcharge. This is shown in Figure 2.5, which plots the low tip suggestions by $x(d, mph)$, surcharges, and vendor. Points shown in black represent integer tip suggestions. For taxicabs using VTS, regardless of the pickup and drop-off locations, the probability that γ_{ij}^{VTS} is an integer is 0, except when there is a night surcharge, $s = \$0.50$. For VTS trips that travel at least a fifth of a mile, it is only in the case where $s = \$0.50$ and $x = 25y + 5$ where y is an integer greater than or equal to 0. On the other hand, trips in CMT taxicabs have an integer γ_{ij}^{CMT} when $x = 25y + 5$, and there is no surcharge, $s = \$0.00$. If there is a night surcharge, then the probability of a low integer tip suggestion is 0. Alternatively, if it there is a peak weekday surcharge, $s = \$1.00$, then a CMT trip will give an integer tip suggestion for the lowest option if $x = 25y + 15$. In summary, for the majority of the day CMT trips have a positive probability of an integer low tip suggestion, but this changes at night when only VTS trips have a positive probability of an integer low tip suggestion.

Assume that for any given passenger, $x(d, mph)$, is fixed and known. This raises two concerns when estimating the impact of integer tips on tipping behavior based on the tipping suggestion formulas shown in equations (2.7) and (2.8). First, if customers can sort on vendors then this could lead to non-random variation in probability of an integer tip. This concern of selection by the rider is mitigated by the fact that, prior to entry,

Figure 2.5: Low Tip Suggestion by $x(d, mph)$, Surcharges and Vendor

(a) VTS

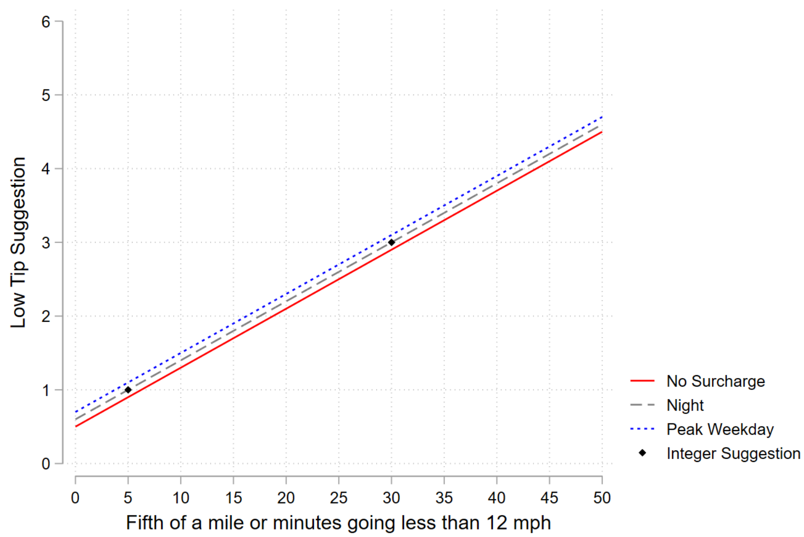

(b) CMT

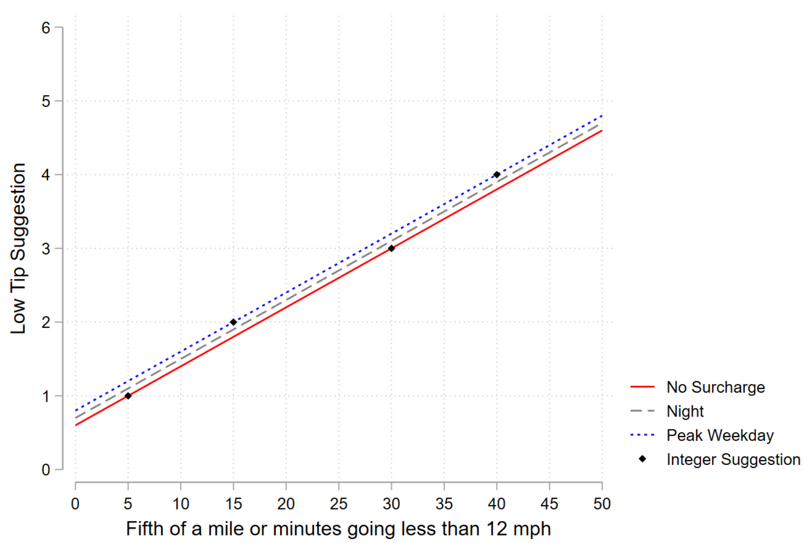

Notes: Figures show the mapping from $x(d, mph)$ and surcharges to Low tip suggestions for each vendors from Feb – Aug 2012. For VTS, integer suggestions only appears in weekdays during peak hours. For CMT, integer suggestions could appear when there's no surcharge, or during the peak weekdays (surcharge = $1.00 from 4pm-8pm on weekdays).

taxicabs with VTS or CMT credit card machines appear essentially identical. Second, it is possible that, absent an integer tip suggestion, customers (trips) that are likely to have an integer tip suggestion are different in tipping behavior compared to those that are unlikely to have an integer tip suggestion. For example, in this case where $x(d, mph)$ is deterministic, even though there is not sorting by vendor, all customers except those with $x = 25y+5$ or $x = 25y+15$ will never be presented with a low integer tip suggestion. Although there is no reason to believe ex-ante that these customers should differ from those that have a low, or zero, probability of receiving an integer tip suggestion, we mitigate this concern by controlling for average tipping behavior in pickup and drop-off locations.

In reality, even when taking a trip between the same locations, $x(d, mph)$ is likely to vary due to exogenous factors like rainfall. There is also, however, the concern that drivers could manipulate $x(d, mph)$ in a way that leads to an unobserved correlation between tipping behavior and the probability of an integer tip suggestion. If there was manipulation by drivers, to induce more frequent integer occurrences then this should be apparent in more frequent tip suggestions ending with a 0. One would expect that this would show up as a higher frequency of 0 in the second decimal place, however, which is not what we see in panel (a) of Figure 2.3.

2.4.2 Econometric Specifications

Each passenger that pays with a credit card faces a set of tip suggestions at the end of the trip. None of these options have an integer tip suggestion in the vast majority of the time, but, as we detailed in the previous section, the treatment and control groups change throughout the day depending on the vendor and distance of the trip. To isolate the effect of the integer tip suggestion, we attempt to control for each of these factors, such as time and distance, which could influence tipping behavior.

Let D_{ijcdhm} denote whether a trip from location i to location j in taxi c on date d, pickup hour h, and pickup minutes m has a nominal tip suggestion in the menu that is an integer. We estimate the effect of D on the probability a customer selects an option from the tipping menu and the tipping rate, defined as the fraction of the fare rate tipped, as:

$$y_{ijcdhm} = \alpha + \beta D_{ijcdhm} + \delta x_{ijcdhm} + \gamma I_{ijcdh} + \epsilon_{ijcdhm} \tag{2.9}$$

where y is either an indicator for whether a passenger gives a suggested tip or the tipping rate, which we define as the tip amount divided by the rate fare used when determining the tip suggestion. Our coefficient of interest is β, which estimates the effect of a an integer tip suggestion on selecting a default tip option or the tipping rate. According to our hypotheses based on the behavioral model, if customers experience increased utility from tipping integers ($b_i > 0$) then we will find that integer tip suggestions increase the probability of selecting the default tip option and the average tipping rate ($\beta > 0$). Our model also shows that tipping behavior varies by the fare, so we linearly control for $x(d, mph)$. In addition, we control for average differences in tipping by driver, location, and over time with driver, date by hour, pickup census block, and drop-off census block fixed effects, I_{ijcdh}, in our preferred specification. Although this is our preferred specification, we will vary the controls to ensure the robustness of our results to alternative specifications. In all specifications standard errors are two-way clustered at the driver and date level.

To test for heterogeneous effects of integer tip suggestion, we define D^1_{ijcdhm} if the lowest of the three tip suggestions is an integer nominal tip suggestion, and D^2_{ijcdhm} if either of the other two options are an integer nominal tip suggestion. We then estimate:

$$y_{ijcdhm} = \alpha + \beta_1 D^1_{ijcdhm} + \beta_2 D^2_{ijcdhm} + \delta x_{ijcdhm} + \gamma I_{ijcdh} + \epsilon_{ijcdhm} \tag{2.10}$$

where β_1 estimates the impact of an integer tip suggestion if it is the lowest option on the menu, and β_2 shows the effect if it is either of the other options. If the effect of an integer tip suggestion is identical regardless of its place on the menu, then our estimate for β_1 and β_2 should be the same. We include the same controls in our preferred specification and cluster the standard errors at the driver level.

2.5 Results

We present our empirical findings in this section. In Section 2.5.1, we show baseline results from regression analyses that utilize quasi-random variation in the occurrences of integer tip suggestions. In Section 2.5.2, we complement the baseline findings with a series of robustness checks. In Section 2.5.3, we provide additional results on the tip-rounding behaviors for individuals who opted-in for custom tips.

2.5.1 Regression Analysis

Table 2.2 presents our primary results for the effect of integer tip suggestions on tipping an option from the menu. We find that, when presented with an integer tip suggestion, passengers are approximately 21 percentage points more likely to give a tip equal to one of the suggested options. The last three columns highlight that this response can be largely attributed to when the lowest suggestion is an integer. Although this provides some evidence that clustering around integer values, we cannot rule out that a part of the effect we are estimating is due to some customers always tipping an integer value, regardless of the suggestion or percentage. If this behavior is at lower integer values, then this could also potentially explain that the effect is largely driven by the low option.

If the mechanism behind the result shown in Table 2.2 is customers always tipping

Table 2.2: Impact of Integer Tip Suggestions on Selecting a Default Suggestions

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.19506***	0.21361***	0.21462***			
	[0.00590]	[0.00739]	[0.00746]			
Low Option Integer				0.23487***	0.23239***	0.23388***
				[0.00569]	[0.00642]	[0.00642]
Mid or High Option Integer				-0.00530	0.05978***	0.05806***
				[0.00670]	[0.01357]	[0.01357]
Outcome Mean	.554	.554	.554	.554	.554	.554
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,658	1,656	1,655	1,658	1,656	1,655
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the probability that a custom tips a suggested amount. The results shown here are for all standard rate fare trips from February to August of 2012, which did not involve a pickup or drop-off at an airport. All estimates are from a linear probability model with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

integer values, regardless of tip suggestions, there should be no difference in the average tipping rate when the suggestions are an integer. This strategy would be, by definition, independent of the tipping suggestion, and thus should not lead to a larger (or smaller) tipping rate based on integer tip suggestions on the menu. Table 2.3 presents our main results for the effect of integer tip suggestions on passenger tipping rates. In our preferred

Table 2.3: Impact of Integer Tip Suggestions on Tip Rates

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00796***	0.00716***	0.00631***			
	[0.00080]	[0.00067]	[0.00064]			
Low Option Integer				0.01334***	0.01013***	0.00826***
				[0.00074]	[0.00073]	[0.00073]
Mid or High Option Integer				-0.01215***	-0.00568***	-0.00300**
				[0.00121]	[0.00124]	[0.00118]
Outcome Mean	.193	.193	.193	.193	.193	.193
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,658	1,656	1,655	1,658	1,656	1,655
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the rate fare used when determining the tip suggestion. The results shown here are for all standard rate fare trips from February to August of 2012, which did not involve a pickup or drop-off at an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

specification, we find that passengers increase their tipping rates by 0.006 (0.6 percentage points) when they are presented with an integer tip suggestion. Column 6 shows that this is due to integer tip suggestions for the lowest option, and not the other two options.

The results from Table 2.2 and 2.3 show that customers are, in fact, responding differently to the presentation of integer tip suggestions, particularly if it is the lowest option. From the visual evidence shown in Figure 2.2, it was evident that customers tend to give integer tips. These results match our hypotheses, which suggests that based on our behavioral model passengers do experience additional utility from giving an integer tip. Specifically, only when $b_i > 0$ would our model suggest this pattern of behavior. Even when there are cognitive costs associated with tipping, there is no reason that cognitive costs associated with giving custom tips change discontinuously when the tip suggestion is an integer.

2.5.2 Robustness Checks

Leveraging plausibly exogenous variation in the decimals of tip suggestions in 2012, we found that customers increase the their tipping rate by 0.6 percentage points when presented a tip suggestion that is an integer. This is driven almost entirely by integer tip suggestions in the lowest menu option. To ensure that our findings are not a result of how we defined the tipping rate, the sample restrictions, or the identification strategy used, we will take alternative approaches to each of these aspects of our analysis.

First, in our primary results we defined the tipping rate as the tip amount divided by the rate fare used when determining the tip suggestion. This varies by vendor since CMT includes tolls and the MTA tax, while VTS does not. To make the denominator comparable, we calculate the total cost of the trip, except for the tip, and use this as the denominator when defining the tip rate. The results from estimating equations (2.9) and (2.10) with this alternative method of calculating the tip rate are shown in Table 2.4. Our results with this method are similar to our primary results, albeit slightly smaller with our preferred specification showing a change in the tip rate of approximately 0.56 percentage points.

Second, to check that our results are not driven by our sample restrictions, we conduct

Table 2.4: Impact of Integer Tip Suggestions on Tip Rates (Alternative Definition)

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00935***	0.00624***	0.00556***			
	[0.00075]	[0.00066]	[0.00063]			
Low Option Integer				0.01258***	0.00901***	0.00747***
				[0.00077]	[0.00072]	[0.00073]
Mid or High Option Integer				-0.00559***	-0.00556***	-0.00330***
				[0.00122]	[0.00124]	[0.00118]
Outcome Mean	.187	.187	.187	.187	.187	.187
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,658	1,656	1,655	1,658	1,656	1,655
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the total fare excluding the tipped amount. The results shown here are for all standard rate fare trips from February to August of 2012, which did not involve a pickup or drop-off at an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

two additional robustness checks. We first expand the sample to include all trips, not just standard rate trips during the time period of February to August 2012.[13] As we show in Table 2.5, including these trips increases our estimates of the impact on customer tip

Table 2.5: Impact of Integer Tip Suggestions on Tip Rates: Including non-Standard Rate Trips

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00506***	0.01106***	0.00900***			
	[0.00062]	[0.00065]	[0.00066]			
Low Option Integer				0.00845***	0.01420***	0.01146***
				[0.00064]	[0.00061]	[0.00073]
Mid or High Option Integer				-0.01080***	-0.00640***	-0.00445***
				[0.00114]	[0.00114]	[0.00111]
Outcome Mean	.193	.193	.193	.193	.193	.193
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	1,663	1,662	1,661	1,663	1,662	1,661
Clusters (Date)	213	213	213	213	213	213

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the rate fare used when determining the tip suggestion. The results shown here are for all trips from February to August of 2012, even if the drop-off or pickup location is an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

rates. Instead of increasing the trips during this period into our analysis, we could also

[13]The restriction to standard rate trips is because other trips, such as from JFK Airport to Manhattan, are charged at a fixed rate leading to non-random spikes in our data at certain tip suggestions. JFK to Manhattan, for example, is charged at a fixed rate of $45, which leads to, absent any surcharges or tolls, a 9 tip suggestion.

expand the time period that we analyze. Specifically, we can use all CMT trips starting in February 2011 until August 2012, which faced the same tip suggestions of 20, 25, and 30%. If customers do not select into a specific vendors, then we should find similar estimates focusing on CMT over this period compared to our primary results. Table 2.6

Table 2.6: Impact of Integer Tip Suggestions on Tip Rates: CMT (2011 Feb – 2012 Aug)

	(1)	(2)	(3)	(4)	(5)	(6)
Any Default Integer	0.00804***	0.00642***	0.00591***			
	[0.00068]	[0.00056]	[0.00053]			
Low Option Integer				0.01595***	0.01198***	0.00988***
				[0.00064]	[0.00058]	[0.00059]
Mid or High Option Integer				-0.01100***	-0.00792***	-0.00522***
				[0.00077]	[0.00079]	[0.00079]
Outcome Mean	.189	.189	.189	.189	.189	.189
x(d,mph) Control	No	Yes	Yes	No	Yes	Yes
Date FE	No	Yes	No	No	Yes	No
Pick-up Hour FE	No	Yes	No	No	Yes	No
Driver FE	No	Yes	Yes	No	Yes	Yes
Drop-off Block FE	No	No	Yes	No	No	Yes
Pickup Block FE	No	No	Yes	No	No	Yes
Date by Hour FE	No	No	Yes	No	No	Yes
Clusters (Driver)	2,317	2,311	2,310	2,317	2,311	2,310
Clusters (Date)	573	573	573	573	573	573

Notes: This table shows the estimated impact of having an integer tip suggestion option on the tipping rate, defined as the tip amount divided by the rate fare used when determining the tip suggestion. The results shown here are for all standard rate fare trips from February 2011 to August of 2012 that use CMT equipment and did not involve a pickup or drop-off at an airport. All estimates are from an ordinary least squares regression with specifications varying by column. The specifications of the first three columns are repeated in the next 3 columns. The first three columns show the effect if any of the options on the menu are an integer, while the last three columns show the effect if the lowest option is an integer or the other two options are integers. The first (and fourth) column has no controls or fixed effects. The second (and fifth) column includes date, driver, and hour fixed effects. The third (and sixth) column includes date by hour, driver, pickup census block, drop-off census block, and nearest integer to the low suggestion fixed effects. Apart from the first and the fourth column, we control for the linear relationship between $x(d, mph)$ and the outcome of interest. Standard errors in brackets are two-way clustered at the driver level and the pick-up date level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

shows that this is the case. We find a very similar pattern of results with an even higher estimated effect of integer tip suggestions on the tip rate.

Lastly, we verify that our results are not driven by our identification strategy in a couple ways. First, we estimate our preferred specifications with placebo treatment effects. To do this, we compare our preferred estimates with a series of placebo treatment effects in which we randomly assign treatment status to each trip for the main sampling period.[14] The placebo estimation follows the baseline specification (2.9). We repeat the process for 1,000 times to obtain a distribution of estimated placebo effects. We define the p-value in this context as the probability that the baseline estimate is obtained purely by chance. The results from placebo tests are shown in Figure B5. We find our baseline estimations are significantly different from their respective random benchmarks (p-value< 0.001).

In addition, we utilize variation over time in VTS tip suggestions. From January 22 to January 26 2012, VTS updated the tip suggestions for fares under \$15 from a fixed menu of \$2, \$3, and \$4 to 20%, 25%, and 30%. The difference between the tip suggestions following the policy change varied based on how far the fare was from \$10. For a fare of exactly \$10, two of the suggestions were identical (\$2 and \$3) with the only change being replacing the rarely used \$4 option with a \$2.50 option.[15] Outside of this option, this means that if a fare is slightly above or below \$10, the two primary tip suggestions are thus nearly identical before and after the menu change, except they are no longer integers. In addition, customers are no longer presented an option at the right-tail of the distribution shown in Figure 2.2, but instead have a 25% tip rate option.

We use the change in the VTS menu and two alternative empirical strategies to analyze the impact of this menu change on selecting an option from the menu and customer tip rates. Limiting our sample to trips with total fares in the range of \$9 to \$11 to

[14] We maintain the proportion of trips with integer suggestions when assign our placebo treatments.

[15] Prior to the change, customers used the option to give a \$4 tip for a fare ranging from 9 to 11 dollars less than 4 percent of the time.

ensure that the percent and dollar values are similar for two of the options, we estimate a regression discontinuity in time and an event study design. Our regression discontinuity estimates are local linear regressions using a triangular kernel and a bandwidth of 30 days, while our event study estimates control for changes in rate fares in September 2012, tipping percentages for CMT in February 2011, vendor and date fixed effects, and controlling linearly for $x(d, mph)$. When we estimate a regression discontinuity in time, we find a statistically significant decrease in the tip rate of 0.7 percentage points, which is shown in Figure B6.[16] This is very similar to the results from our event study in Figure B7. We find that, consistent with our primary results, switching away from integer tip suggestions decreased the probability that a passenger chose to tip an option from the menu and it decreased the average tip rate. In terms of magnitude, the estimated impact on tip rate is similar, but the estimated effect on selecting default tip options is much smaller than our primary results, likely due to the replacement of the $4 option with a $2.50 option.

2.5.3 Who are the Switchers?

Our results highlight that tip options that are integers lead to an increased probability of selecting an option from the tip menu and an overall increase in the tip rate. One question that we have not addressed, however, is what passengers that switch to a menu option do if there are no integer options on the menu. From the theory, we would expect these to be customers with high b_i and/or $T_i \approx 20\%$. This matches what we find in the data, which shows that over half of customers that choose to give a tip not on the menu give either a tip of one dollar, or they round (up or down) from the lowest tip suggestion

[16] Although we do estimate a decrease in the selection of options from the tipping menu using a regression discontinuity design, we do not present them here since we are not able to confidently classify the tipping menu around the cutoff. It took more than half a week for changes to the tip menu to occur for all vehicles, which is a process we do not observe. For this reason, we do not present these results here.

to the nearest dollar or fifty cents.

To better understand how customers are giving custom tips, we create bins based on the decimals of the lowest tip suggestion. We then calculate what fraction of custom tips within each bin exhibit rounding up or down to the nearest integer or rounding up or

Figure 2.6: Fraction of Custom Tips that Round to Nearby Values

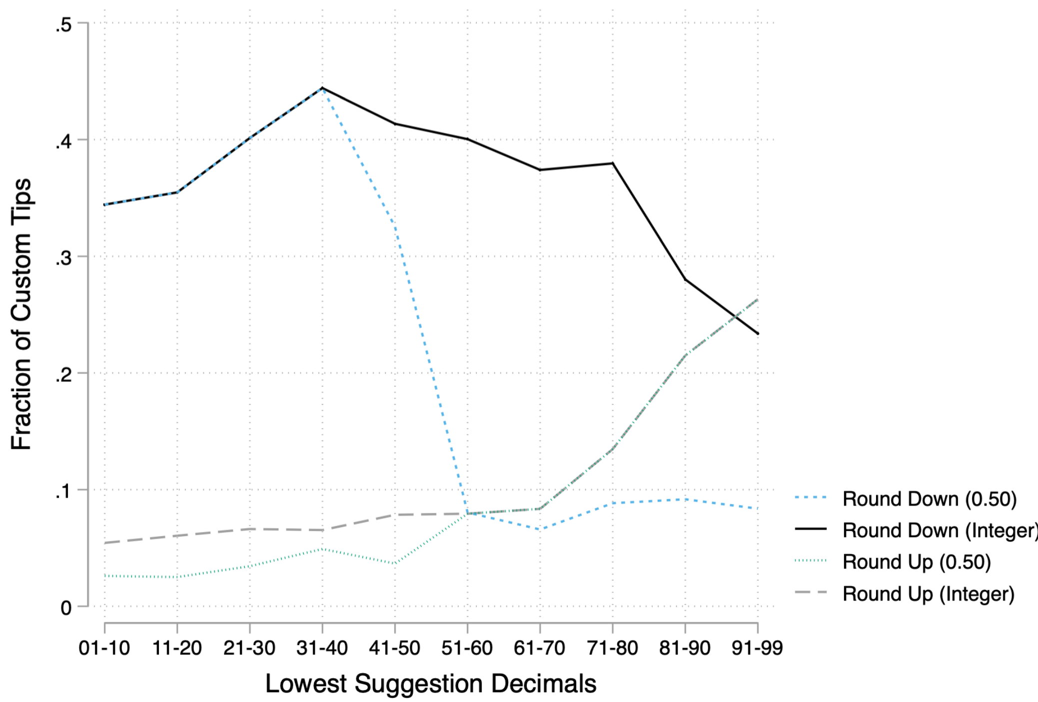

Notes: This figure shows rounding heuristics used by customers when they opt-in for custom tips. The horizontal axis presents decimal places for the lowest tip suggestion amount. The decimal places are divided into 10 equally spaced bins. The vertical axis represents the fraction of custom tips that either rounds the lowest default suggestion up or down to the nearest dollar or 50 cents. Overall, we observe that customers are more likely to round down regardless of the decimal places of the suggestions.

down to the nearest 50 cents. We plot the pattern of passenger behavior by decimal bin in Figure 2.6. Regardless of the decimal place of the suggestions, customers tend to round down to the nearest dollar or 50 cents far more than they round up. In fact, only at the point where the decimal places of the tip suggestion are in the range of 91-99 do more customers round up to the nearest integer more than they round down. Empirically, this pattern highlights the mechanism behind our results showing an increase in the tip percentage at the integer values. Customers, when choosing to give a custom tip, tend to choose a lower integer rather than one larger than the lowest option on the tip menu. When presented with an integer tip, however, customers that are rounding for an integer value are less likely to do so since they can select the tip suggestion and avoid the cognitive cost. Since most of these customers tend to round down, this leads to an increase in the tip rate.

2.6 Implications

In the previous sections, we first documented that customers tend to tip integer values. We then presented a model of tipping behavior that can explain this observed pattern in the data. On the one hand, when customers choose to give a tip that is not on the menu the cognitive cost could be lower when choosing an integer. On the other hand, it could be the case that customers actually experience a lump-sum increase in utility when tipping an integer, regardless of if the tip is from the menu or a custom tip. Importantly, both of these mechanisms can explain clustering of tip amounts at integer values, but they differ in one key dimension – the response when presented an integer tip on the menu.

To understand the mechanism underlying the tendency for customers to tip integers, we leverage plausibly exogenous variation in the occurrence of integer tip suggestions. We find consistent evidence that customers do experience a lump-sum increase in utility, which is shown through increases in the tip rate and the probability that a customer chooses a tip from the menu when presented with an integer tip suggestion. In the data, this is driven by customers that, instead of tipping an integer below or above the tip suggestion, are now choosing the option from the menu.

In this section, we consider the implications of customer's preference for integers on the impact that price changes (rate fares) and tip suggestions have on revenue.

2.6.1 Implications for the Impact of Price Changes on Revenue

Let a taxi-drivers revenue for a trip be given by $R = F + t \cdot F$, where F is the total fare and t is the optimal tip rate from the passenger that is maximizing the utility function in equation (2.4). Denote the average tip rate if customers do not face an integer suggestion as $\bar{t}^{non}$. Similarly, define the average tip rate of those that do face an integer suggestion as $\bar{t}^{int}$, where $\bar{t}^{int} - \bar{t}^{non} = \eta$. Assume that the fraction of trips that face an integer suggestion is z and the fraction that do not is $1 - z$. The average taxi-drivers revenue can then be rewritten:

$$\bar{R} = \bar{F} + \bar{F}[z\bar{t}^{int} + (1-z)\bar{t}^{non}] = \bar{F} + \bar{F}[\bar{t}^{non} + z\eta] \tag{2.11}$$

the revenue from taxi-drivers now depends on how often the tip suggestion presented is an integer and the impact this has on average tipping rates.

To examine the importance of this channel, we will now consider the impact of increasing the rate fare from 0.40 to 0.50, which occurred on September 4, 2012. Assuming that η is unchanged, the change in revenue is:

$$\Delta\bar{R} = \bar{R}_1 - \bar{R}_0 = \Delta F + \underbrace{\bar{F}_1\bar{t}_1^{non} - \bar{F}_0\bar{t}_0^{non}}_{\text{Non-integer Tip Change}} + \underbrace{\eta[\bar{F}_1 z_1 - \bar{F}_0 z_0]}_{\text{Integer Tip Change}} \tag{2.12}$$

where the last term, representing the integer tip change, can simplify further if the likelihood of an integer suggestion after the fare change, z_1, is the same as before the fare change, z_0. In the case of this fare change, however, z_1 increased significantly as we show in Figure 2.7. Average total fares increase from approximately 10.36 ($\bar{F}_0$) to 12 ($\bar{F}_1$) along with a large increase in the probability of an integer tip suggestion from 3.11% (z_0) to 21.27% (z_1). By combining these parameters with our preferred estimate for η from column (3) of Table 2.3, we are able to calculate the last component of the

Figure 2.7: The Probability of Integer Tip Suggestions by Month-Year

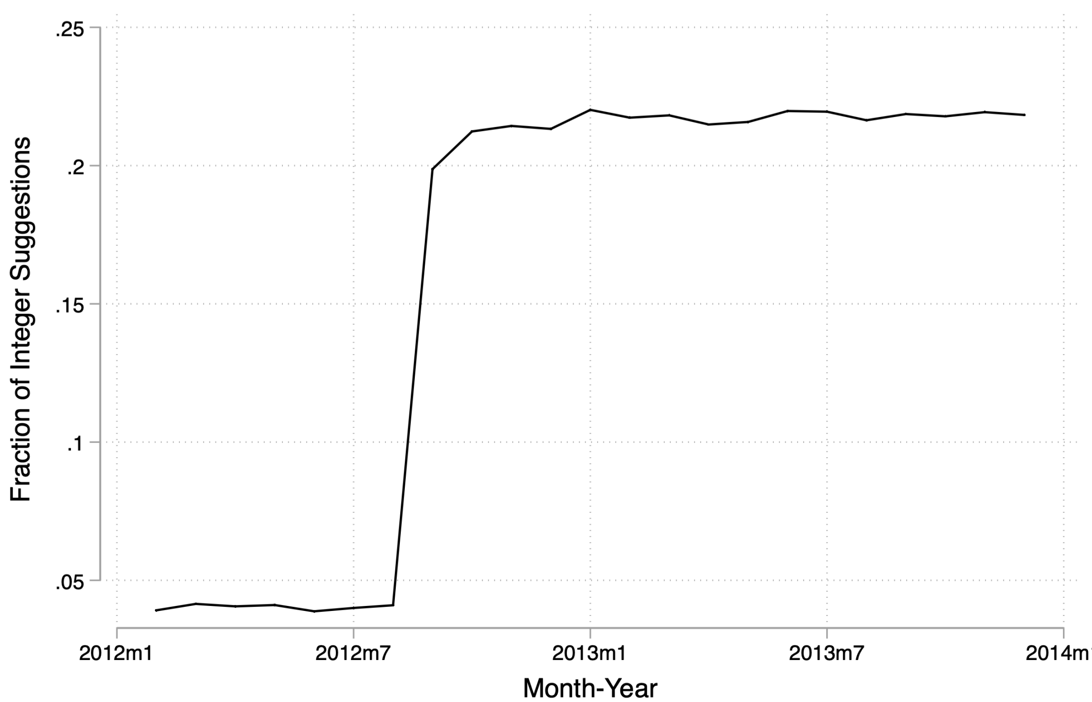

Notes: This figure shows the average probability of having an integer tip suggestion overtime from 2012.Feb onward. Around 2012.Sept, per-unit fare rate increased from \$0.40 to \$0.50. Given a tip suggestion menu: 20, 25, and 30 percent, this fare increase has significantly raised the probability of having an integer tip suggestion.

change in revenue. In other words, we can calculate how much the average revenue of a trip increased as a result of the tendency for customers to tip a higher percentage when presented with integer tip suggestions. Plugging in all of the aforementioned parameters into the last component of equation (2.12) we find that this led to an approximately 1.4 cent increase in revenue per trip. With over 170 million taxi trips and 41,000 unique drivers this leads to a transfer of 2.38 million dollars from riders to drivers.

The previous example highlights that when the tip menu is based on percentages, changes in prices (rate fares) can significantly change the nominal values of the options presented to customers. Switching from a rate fare of 40 cents to 50 cents increased the likelihood of an integer tip suggestion by approximately 18 percentage points. Given the differential response of customers to integer tip suggestions, this led to an estimated transfer of 2.38 million dollars from riders to drivers in the year following the policy change. This result emphasizes the key role that the interaction between prices and tip suggestions can have on revenue.

2.6.2 Implications for Default Tip Suggestions

The differential response of passengers based on the tip suggestion decimal places is evident in Figure 2.6. Even when the decimals are in the .91 to .99 range, passengers round down to the 50 cent or dollar at a higher rate than they round up. This leads to the natural question: should drivers round up tip suggestions to the nearest dollar? Intuitively, this seems like it would increase revenue as customers that would have rounded down will now tip the suggested option.

In order to analyze whether "rounding upwards" of tip suggestions would lead to higher tips, on average, we return to our theoretical framework shown in equation (2.4). An intuitive way to think about the impact that this would have on tips is to consider the removal of the lowest option on the tip menu, which is replaced by a new option, denoted as $t_i^{D'}$. Assume that this option is a larger tip rate, $t_i^{D'} \geq t_i^D$ that amounts to rounding up to the nearest integer, $t_i^{D'} F_i \in \mathbb{Z}$.[17]

Unfortunately, our estimates cannot definitively answer how such a menu change would impact average customer tips. We are limited by our identification strategy, which

[17] It is worth noting that $t_i^{D'}$ is not a single value, but changes based on the total fare.

leverages plausibly random tip suggestions that are integer tips at the current menu of suggested tip rates: 20, 25, and 30%. This means that although we find a higher tip rate when 20% tip suggestions are integers, we cannot conclusively say that this will be the

Figure 2.8: Fraction of Customers Choose a $2 Default, by Fare+Surcharge: VTS (pre-2012)

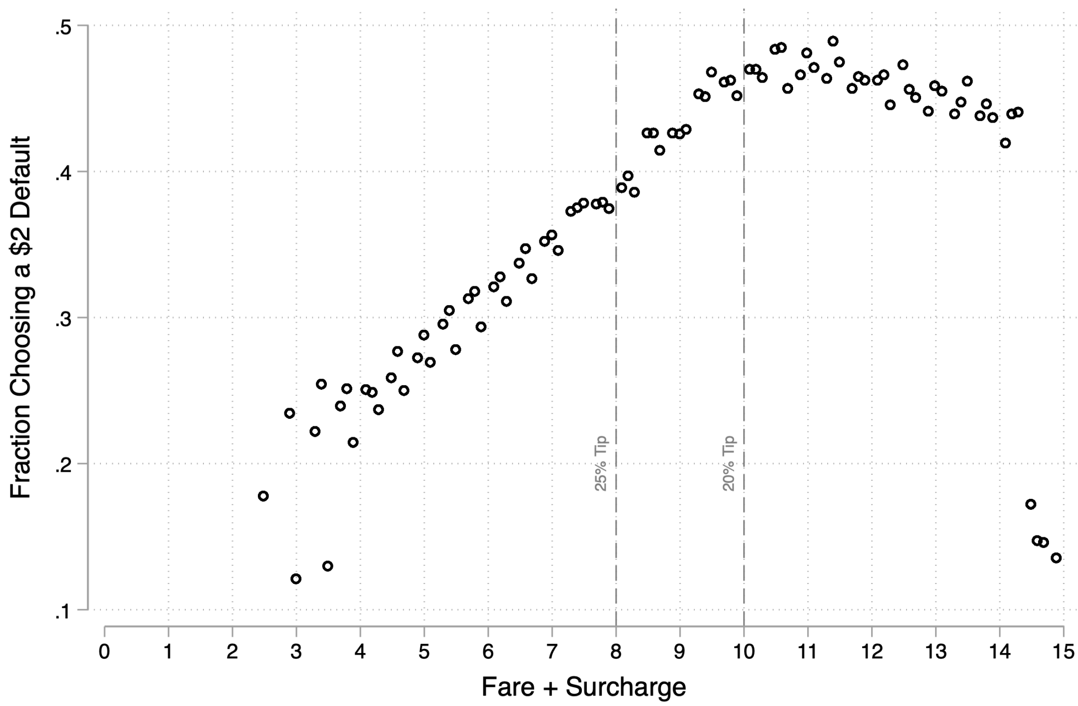

Notes: This figure shows the fraction of VTS customers choose a $2 default before 2012.Jan. we divide the *rate fare* (Fare + Surcharge) into 150 equally sized bins, then we compute the average $2 tip take-up rate for each bin. The 20% and 25% marks represents *rate fare* values that are equivalent to a $2 tip.

case if they are presented with, for example, a 21% suggested tip rate that is an integer. It is possible that increasing the tip rate leads to a large increase in the fraction of customers that choose a custom tip instead of the option from the menu. We can provide some suggestive evidence that customers are unlikely to switch away from the menu option to a custom tip by plotting the selection of the $2 dollar tip suggestion for VTS trips under $15 dollars prior to the menu change in January of 2012. As Figure 2.8 illustrates, customers are unlikely to take-up this suggested tip option when it is a higher percentage of the total fare. However, in the region of the 20% tip it is relatively flat, which suggests that a small change in the tip rate in order to round up the tip suggestion is unlikely to decrease selection of that option significantly.

In summary, our results point towards an alternative tip menu where tip suggestions are rounded up to the nearest dollar. Although we cannot definitively answer whether this would increase revenue, we use the case of 2, 3, and 4 dollar VTS tip suggestions for all trips prior to the January 2012 to show that the down-side from such a policy depends on how large of a shift in the suggested tip rate is. Figure 2.8 shows a small decrease in selecting an integer tip suggestion when it is approximately 20%, but as the tip rate approaches 25 to 30% customers give custom tips at a much higher rate. Importantly, this non-linear relationship between selecting the menu option and tip rates highlight that it may be the case that rounding up tip suggestions can increase revenue, but only when it does not increase the tip rate significantly as this can push passengers away from selecting the tip suggestion.

2.7 Conclusions

Previous research has highlighted that the menu of tip suggestions presented to customers impact the amount that they tip. Using detailed data on millions of trips in New York taxicabs, however, we document that customers tend to tip at integer values even when the menu of tip suggestions are rarely integer values. By extending the model of Donkor (2020), we show that this behavior can be explained by decreasing cognitive costs associated with computing a custom tip that is an integer or from additional utility gains when giving an integer tip. Despite the fact that both model extensions can lead to clustering at integer tips, only in the presence of utility gains at integer values do customers tip differently when presented with an integer suggestion on the tipping menu.

To estimate if customers respond differentially to integer tip suggestions, we leverage plausibly exogenous variation in when integer tip suggestions are given to customers. Across a variety of sample restrictions, specifications, and estimation strategies, we find that when customers are presented with an integer they are more likely to give a tip equal to the suggested value and they tip at a higher rate. That customers respond differentially to the presentation of an integer tip suggestion provides evidence that the tendency to tip integer values is driven, at least in part, by a preference to give integer tips.

Customers' differential responses to integer tip suggestions has natural implications for how prices and tip menus impact revenue. Specifically, our estimates of how customers respond to integer tip suggestions imply that the rate fare change in September 2012 increased average annual revenue for the NYC taxi industry by approximately 2.38 million dollars per year due to the fact that it increased the probability of integer tip suggestions by 18 percentage points. In addition, our finding that customers experience utility gains from integer suggestions could have implications for revenue-maximizing tip menus. Incorporating a model of consumer utility that accounts for the differential response of consumers to integer prices and tip suggestions is an area that deserves the attention of future work.

Chapter 3

Asymmetric Pass-through from Carbon Taxes to Electricity Prices

3.1 Introduction

Asymmetric price transmission is the phenomenon that the change in prices depends on the direction of the change in costs (Peltzman 2000; Borenstein, Cameron, and Gilbert 1997). This is colloquially referred to as "Rockets and Feathers" since the general finding is that prices rise like rockets when input prices (costs) go up, but fall like feathers when input prices go down. Determining the source of this phenomenon, however, is often challenging. Many theoretical explanations have been suggested based on factors such as consumer search (Benabou and Gertner 1993; Tappata 2009; Cabral and Fishman 2012), customer fairness beliefs (Eyster, Madarász, & Michaillat, 2020), menu costs (Ball and Mankiw 1994), and even simply the shape of demand and supply curves (Ritz 2015). Disentangling the mechanisms behind evidence of asymmetric price transmission is often challenging since we do not observe many of these factors (e.g., fairness) or even supply and demand curves.

In this paper, I utilize the unique setting of wholesale electricity markets to directly analyze whether plants do, in fact, respond differently to increases in costs compared to decreases in costs. Leveraging the implementation and repeal of the short-lived Australian carbon tax, I first analyze the impact that this has on equilibrium wholesale electricity prices. Similar to the previous literature, I find prices increase by more when costs rise (implementation) than they decrease when costs fall (repeal). By focusing on wholesale electricity markets, I am able to move beyond the equilibrium outcomes to directly analyze how electricity generators responded to the implementation and repeal of the carbon tax. Wholesale electricity markets clear using plant-level supply curves during each half-hour of the day, which are aggregated into the market supply curve. Analyzing the change in the supply curves submitted by each plant, I find that the price plants are willing to accept to generate electricity increases by more when the carbon tax is implemented compared to when it is repealed. In so doing, I provide the first evidence directly linked asymmetric responses of suppliers to evidence of asymmetric price transmission.

There are several unique advantages of this setting for analyzing asymmetric pass-through of cost shocks. First, the National Electricity Market in Australia is a multi-unit uniform price auction with a homogeneous good (electricity) where electricity generators and retailers submit step-wise supply and demand curves, respectively. These are aggregated by the market operator and the market equilibrium is determined in order to maximize economic efficiency subject to real world constraints. Importantly, this means that even though there is market power in wholesale electricity markets, other common explanations like consumer search, customer fairness beliefs, and menu costs do not apply to this context.[1] Equilibrium prices, in this context, are set electronically with produc-

[1] The role of market power in asymmetric pass-through of costs is still unresolved. Loy, Weiss, and Glauben (2016), for example, find a negative relationship between asymmetric pass-through of costs and market power.

ers costlessly setting their supply curves and retailers submitting, essentially, perfectly inelastic demand curves throughout the day. Second, I am able to directly observe the individual supply curves (and in turn the aggregate supply curve) that is used to clear the market. This means that, unlike other settings, I am able to directly analyze supply-side responses to changes in costs. Lastly, the Australian carbon tax had minimal impact on investment due to policy uncertainty (Teeter and Sandberg 2017), was implemented and repealed at the same time of year only 2 years apart, and can be linked directly to marginal costs using emissions rates. These features mitigate identification concerns and also allow me to link the shock in input (carbon) prices to changes in marginal costs and, in turn, change in willingness-to-accept.

The analysis of this paper proceeds in three parts. First, I use a difference-in-regression discontinuity design to estimate the impact of the implementation of the carbon tax compared to its repeal only two years later. Even when controlling for fuel prices, weather, and fluctuations in demand by day of the week, I find that wholesale electricity prices increase by more when the carbon tax is implemented compared to the decrease when it is repealed. Specifically, I find that the 23 dollar increase in the carbon tax (implementation) immediately increases wholesale electricity prices by approximately 30 dollars, while a 24.15 decrease in the carbon tax (repeal) immediately decreases wholesale electricity prices by approximately 15 dollars.

Although the change in prices is twice as large when the carbon tax is implemented compared to when it is repealed, it is possible that this is due to changing economic conditions (e.g., unobserved demand shocks) or different changes in marginal costs associated with the carbon tax changes. The last two parts of the analysis utilize information on firm emissions rates and the supply curve to analyze these two mechanisms.

Prices in wholesale electricity markets reflect the willingness-to-accept of the last plant called on to meet demand. Although I find a very similar pattern of emissions intensities for these plants around the implementation and repeal of the carbon tax, it is possible that the response in prices is reflecting the behavior of only a single plant. In the second part of the analysis, I explore if the results for prices reflect average changes in the price plants are willing to accept to produce electricity in an area around the equilibrium. Specifically, the Australian Energy Market Operator (AEMO) releases a range of demand projections the day before that includes the 10th and 90th percentile of their demand estimates. Using this as the range that plants believe has a positive percent of being the true demand, I estimate the average willingness-to-accept in this region. Using an identical identification strategy to the analysis focused on prices, I find that there is a similar asymmetry evident in average changes to willingness-to-accept as with prices. Importantly, evidence of a similar asymmetry in changes to willingness-to-accept highlights that the asymmetry in price changes is not purely due to unobserved demand shocks, but are instead a result of asymmetric supply-side responses.

Lastly, I directly analyze the behavior of plants by focusing on the supply curves that each plant submits to AEMO. To simplify the analysis, I focus on the part of the supply curve for each plant that, once aggregated, is most likely to impact prices. I define this as the step of each plant's step-wide supply function that is closest to the projected demand for that part of the day.[2] I then estimate how, on average, plant's adjust the willingness-to-accept in response to the implementation and repeal of the carbon tax, controlling for each plant's average prior to the tax change in 2012 and 2014. I find that plants do respond asymmetrically with large increases in the price they are willing to accept when the carbon tax is implemented, but smaller decreases when the carbon tax is repealed. This provides the first evidence that directly shows the asymmetric changes in equilibrium prices is a result of asymmetric firm-level responses.

[2]Since some plants have no bands near the projected equilibrium, I further restrict the analysis to steps of supply functions that are reasonably close to the equilibrium. I define this as being within 1000 MW, but the results are robust to alternative bandwidths.

A primary contribution of this paper is directly identifying asymmetric responses of producers to cost shocks. Previous work has documented asymmetric responses to changes in input costs (Peltzman 2000), value added taxes (Benzarti, Carloni, Harju, and Kosonen 2020), and even wholesale electricity markets and carbon prices (Zachmann and Von Hirschhausen 2008) but cannot directly identify shifts in the supply curve in response to cost shocks. The results of this paper are most closely related to Benzarti et al. (2020), who show through a variety of tests that standard theory cannot explain why prices rise more with VAT increases than they fall with VAT decreases. Although I cannot speak to dynamics past the short-run like they can, this paper uses the unique advantages of wholesale electricity markets to provide the first evidence of asymmetric supply-side responses to a cost-shock.

By estimating the short-run changes in wholesale electricity prices in response to a carbon tax, this work contributes to the literature that study the impact of carbon prices on electricity markets. Much of this literature focuses on quantifying the relationship between emissions costs and prices. Fabra and Reguant (2014), for example, study the Spanish wholesale electricity market and find that changes in marginal emissions costs are passed on fully to wholesale electricity prices. Other papers have studied the relationship between changes in the average marginal emissions cost and wholesale electricity prices, even in the context of Australia (Nazifi et al. 2017). Similar to Zachmann and Von Hirschhausen (2008), this paper highlights that pass-through estimates can vary based on the direction of the change in carbon prices. In contrast to previous work, however, this paper provides direct evidence of asymmetric supply-side responses depending on the direction of the change in carbon prices.

By estimating asymmetric responses of firms to changes in the carbon price, this paper also contributes to the literature focused on estimating the equity implications of carbon pricing. Previous work has shown the importance of revenue recycling (Goulder et al. 2019) and market structure (Ganapati et al. 2020) on tax incidence. The results of this paper highlight a previously unexplored factor that could impact who bears the burden of a carbon price – the direction of the change in the carbon price.

The rest of the paper is organized as follows. Section 2 presents the background and institutional details. Section 3 analyzes how prices responded to the implementation and repeal of the carbon tax. Section 4 analyzes the supply-side response to the carbon tax and section 5 concludes.

3.2 Background

In this section, I provide background on the National Electricity Market and the *Clean Energy Act 2011*.

3.2.1 National Electricity Market

The National Electricity Market incorporates around 40,000 km of transmission lines that connects electricity generators and consumers across five regional market jurisdictions: Queensland, New South Wales, Victoria, South Australia, and Tasmania. Efficient transmission of electricity from generators to consumers is primarily facilitated through a wholesale market operated by the Australian Energy Market Operator (AEMO).[3] The procedure for determining the market-clearing price in the NEM is similar to many electricity markets. Electricity producers and wholesale retailers submit 10 "bands" for each half-hour of the day, each consisting of a price and quantity pair. For electricity producers, the price represents the willingness to accept (WTA) for the paired quantity, while

[3] There is a separate financial trading market for electricity that is conducted through over-the-counter trading and through exchange trading via the Sydney Futures Exchange. One example of such a financial contract is a contract for differences where the purchaser agrees to purchase a specified physical quantity of energy from the spot market at a set price.

for electricity retailers, the price represents the willingness to pay (WTP) for the paired quantity. AEMO then uses an automated system to rank each electricity producer's band in ascending order, based on WTA, to form a supply curve. A similar process is done with electricity consumers, but their bands are ranked in descending order. The market clears by choosing the lowest-cost combination of producers to satisfy demand, subject to the grid's physical constraints.

Prices in the National Electricity market are different for each of the five regional market jurisdictions. For each half-hour of the day, the price in a zone represents the average change in system cost for one additional unit to be consumed with that zone. In the presence of binding transmission constraints, prices can exhibit significant heterogeneity as different electricity generators are required to meet an additional unit of demand for different regional market jurisdictions.

Electricity Generator Summary

Approximately 90 percent of the electricity generated in the National Electricity Market uses fossil fuels as the primary energy source with the rest largely being provided by wind and hydropower. Common to many electricity markets, there is a high concentration of output across a few firms in the National Electricity Market. Combining electricity generation from AGL, EnergyAustrlia, and Origen Energy, these three firms account for more than 40 percent of electricity generated in the market and just less than 40 percent of capacity. This is perhaps not surprising give the high fixed costs of electricity generation.[4]

[4] Relative to other contexts, such as Spain (Fabra and Reguant 2014), this is relatively competitive when looking at the entire market. In the presence of transmission congestion, however, market power is exacerbated due to the location of the plants owned by these 3 firms.

3.2.2 Clean Energy Act 2011

Under the *Clean Energy Act 2011*, starting on July 1, 2012, all entities that emit over 25,000 tonnes of carbon dioxide equivalent (CO_2e) per year, and which were not in the transport or agriculture sectors, were required to obtain emissions permits (carbon units) for each tonne of carbon dioxide equivalent (CO_2e) that they emit.[5] Throughout the policy's existence, these permits were purchased from the government at a fixed price, or issued for free as a part of industry assistance measures, and the quantity available to each producer was unlimited. Although firms purchased permits, since the price was fixed and the quantity unlimited, the policy behaved identically to a carbon tax, so I will refer to it as such.

When implemented, the price of carbon was set at $23 Australian dollars (AUD) per tonne of CO_2e. The carbon price rose to 24.15 (AUD) on July 1, 2013 where it remained until July 1, 2014 when the carbon price returned to 0 when it was repealed.

3.2.3 Clean Energy Act Repeal

The price on carbon was repealed on July 1, 2014, which is the start of the 2014-2015 fiscal year. It is worth noting, however, that this is not the date that the *Clean Energy Legislation (Carbon Tax Repeal) Act 2014* was physically passed. The repeal bill was passed two and a half weeks later on July 17, but it was backdated to July 1. Backdating, in this context, means that the repeal bill legally comes into effect on July 1, even though it was passed on July 17. The inclusion of this feature in the repeal bill meant that from July 1 to July 17 of 2014 firms were unsure about the current cost of carbon emissions. If the repeal bill passed with backdating, then the cost of carbon

[5] All but the smallest thermal electricity generators meet the minimum eligibility requirements. In fact, the top 8 payers of the carbon tax in the 2012-2013 fiscal year were all electricity generators (LEPID 2012-2013).

during this time period was 0. If the repeal bill did not pass or backdating was removed, then the cost of carbon during this period was the cost of a pollution permit.

In summary, before July 1, 2014, electricity generators were responsible for pollution permits for their emissions. After July 17, 2014, electricity generators knew they did not have to pay for their emissions. During the two and a half-weeks between these two dates, electricity generators were unsure if they were going to have to surrender pollution permits or not.

3.3 Asymmetric Price Changes

In this section, I estimate the impact that the carbon tax implementation and its subsequent repeal had on wholesale electricity prices.

3.3.1 Graphical Evidence

Although it is not causal evidence of a differential response to the implementation and repeal, wholesale electricity prices do appear to respond differently to the implementation and repeal of the carbon tax. In Figure 3.1, I overlay wholesale electricity prices around both events. Blue dots represent the wholesale electricity price in 2012 and black dots represent the wholesale electricity price in 2014. Both gray dashed lines correspond to important dates for the carbon tax. The gray dashed line on July 1 represents the implementation and legal repeal date and the gray dashed line on July 17 shows the date that the repeal bill was physically passed. The period between these two lines thus represents the period of uncertainty in 2014.

To compare the impact of the carbon tax changes, I am going to focus on how prices change from June 30 to July 1 (or July 18) in 2012 compared to June 30 to July 18 in 2014. Focusing on these two time periods in Figure 3.1, there appears to be a difference

Figure 3.1: Price Changes over Time

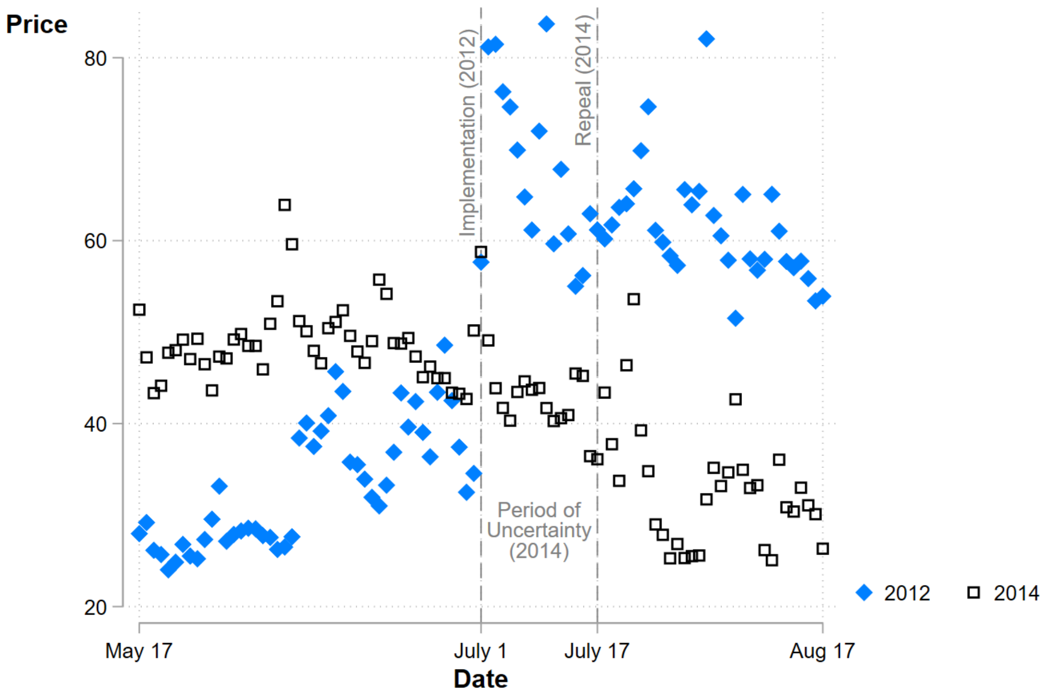

Note: This figure plots estimated quantity-weighted average daily prices from May 17 to August 17 for the years 2012 and 2014. Blue dots represent 2012 prices and black dots represent 2014 prices. Gray lines show the implementation date (July 1), and the date the bill was actually repealed (July 17). The area between these two gray lines is the period of uncertainty in 2014 when the repeal bill was in the process of being passed.

in short run price responses when the carbon tax is implemented compared to when it is repealed. Immediately after the carbon tax is implemented wholesale electricity prices increase sharply by well over 20 dollars. When the carbon tax is repealed, however, wholesale electricity prices appear to decrease slowly over time. Although this is not controlling for factors that impact prices such as weather and fuel costs, the price changes shown in Figure 3.1 suggest that there may be differential responses to the carbon tax implementation compared to its eventual repeal.

Figure 3.2: Emissions over Time

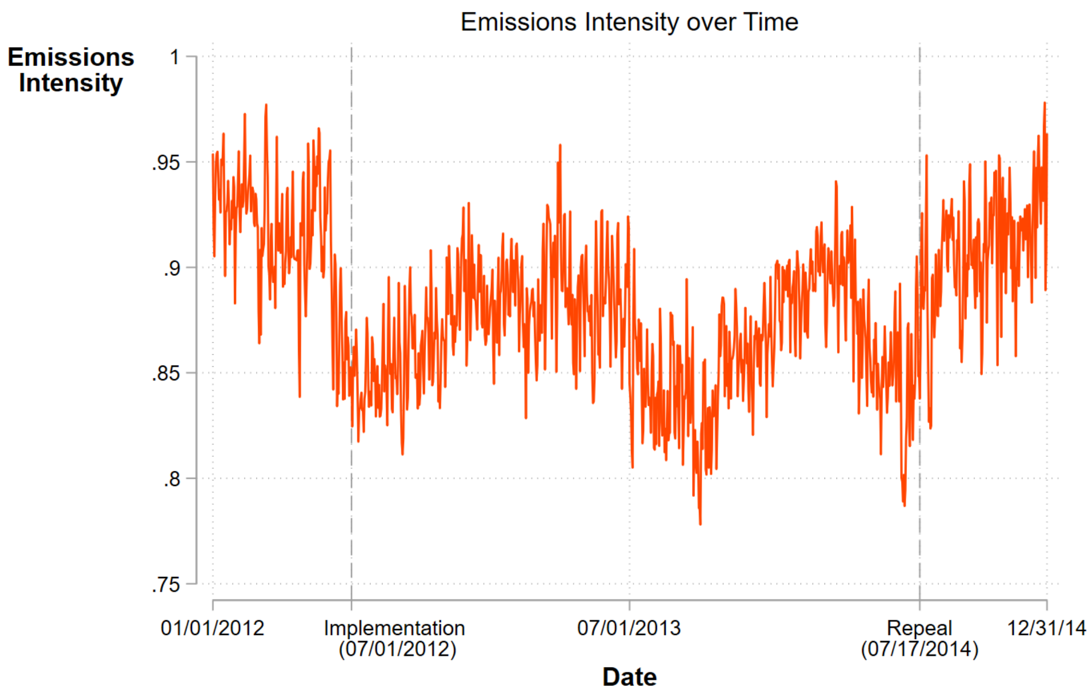

Note: This figure plots estimated quantity-weighted average daily emissions intensities over time. Dashed lines show the date the carbon price was implemented and repealed. The date of July 17, 2014 is shown as the repeal date since this is when the repeal bill was passed, although it was backdated to July 1, 2014.

A natural explanation for this would be wide-scale changes in emissions in the electricity market. It is possible, for example, that many of the high emitting coal plants shutdown during these two years which, along with investment in renewable electricity generators lowered emissions in the electricity markets. AEMO estimates of market emissions intensity over time that are shown in Figure 3.2, however, seem to indicate that this was not the case. Emissions appear to return to previous levels following the repeal of the carbon price, which suggests that the fuel mix did not change significantly during this time period. This is not surprising since evidence from surveys suggest that policy uncertainty undercut investment (Teeter and Sandberg 2017).

3.3.2 Empirical Framework

On July 1, 2012 the carbon tax changed from 0 to 23, and on July 1, 2014 the carbon tax was legally repealed implying a decrease from 24.15 to 0. I utilize this to first estimate separate regression discontinuities in time (RDiT) for the implementation and repeal dates where I vary the size of the bandwidth, controls, and the smoothing parameter. I then estimate a difference-in-regression discontinuities (DiRD) (Murnane and Willett 2010) with a local linear regression and the optimal data-driven procedure for bandwidth selection suggested by Calonico, Cattaneo, and Titiunik (2014). With the DiRD approach, I simultaneously estimate regression discontinuities for both cutoffs, which allows me to empirically test if the price changes are different for the implementation and repeal. It comes at the cost of using the same bandwidths for both the implementation and repeal, but the results from the RDiT will highlight whether this restriction is driving my DiRD estimates.

I separately identify the impact of the implementation and repeal of the carbon tax on the quantity-weighted average daily wholesale electricity price, p_t, by estimating the following RDiT:[6]

$$P_t = \alpha + \rho 1(t \geq c) + \beta f(t) + \gamma f(t) 1(t \geq c) + \theta X_t + \epsilon_t \tag{3.2}$$

[6] Despite the fact that the wholesale electricity market clears every half hour, in this paper I am primarily going to focus on changes in the wholesale electricity market across days. The primary reason for this is that the prices for electricity generator supply functions are limited within a given day. Instead, I calculate the quantity-weighted average price for date t as:

$$p_t = \frac{\sum\limits_{j=1}^{5}\sum\limits_{h=1}^{48} p_{htj} \cdot q_{htj}}{\sum\limits_{j=1}^{5}\sum\limits_{h=1}^{48} q_{htj}} \tag{3.1}$$

where j is a region and h is a half-hour period.

The primary coefficient of interest is ρ, which represents the change in the wholesale electricity price at the cutoff. To test the robustness of my results, I vary the kernel, bandwidth, controls (X_t), and polynomial ($f(t)$) across specifications. My primary results where I include controls, X_t, account for common demand and supply shifters like weather, fuel prices, and changes in electricity demand by day of the week. Specifically, for weather controls I use Bureau of Meteorology data on average water levels in hydro relevant dams, rainfall, and maximum and minimum temperature.[7] I also account for fluctuations in demand with day of the week fixed effects, and changes in fuel costs with FRED data on global prices for coal, natural gas, and crude oil. Although my preferred specification is a local linear regression with a triangular kernel and data-driven bandwidths (Calonico et al. 2014), I also use uniform (rectangular) and epanechnikov kernels, alternative polynomials of order 0 and 2, and I manually vary the size of the bandwidths. Lastly, due to concerns over the role of event-driven outlier days (Figures C1 and C2), I will show estimates when periods with extreme prices larger than the 99 percentile or smaller than the 1st percentile are dropped when calculating the quantity-weighed average daily price.

Estimating separate RDiT means that I can separately identify optimal bandwidths for each cutoff, but I cannot directly test if prices responded similarly to the implementation and repeal of the carbon tax. To test for this differential price response, I estimate

[7] I am only interested in the weather in areas that consume electricity, so I use QGIS to determine which weather stations are within 1 degree, approximately 100 kilometres, of high voltage lines. From here, I take the average of the daily weather measures for all of these stations. Since the location of weather stations is endogenous and tend to be focused around populated areas, the average across all stations will give a higher weight towards more populated areas.

the following DiRD:

$$P_{dy} = \alpha_0 + \alpha_1 R_y + \underbrace{\beta_0 d + \beta_1 d R_y}_{\text{Slopes Below Cutoff}} + \underbrace{\gamma_0 d T_d + \gamma_1 d Post_y T_d}_{\text{Slopes Above Cutoff}} + \underbrace{\rho_0 I_y T_d + \rho_1 R_y T_d}_{\text{Cutoff Estimates}} + \theta X_t + \epsilon_t \tag{3.3}$$

where d is day of the year and y is year. Treatment is represented by T_d, which is an indicator for if the day of the year is after the cutoff,[8] while R_y and I_y are indicator variables equal to 1 if the year is 2014 or 2012, respectively. Bandwidths are determined following Calonico, Cattaneo, and Titiunik (2015) for the uniform and triangular kernels without controls (X_t), which I also, for simplicity, use when including controls. If prices respond symmetrically to the implementation and repeal, then, under certain identification assumptions, estimates of the change in prices at the cutoff should be the same for the implementation, ρ_0, and repeal, ρ_1.

3.3.3 Identification Assumptions

There are two sets of identification assumptions underlying the results from the DiRD. First, there are the standard identification assumptions underlying estimating an RDiT as in equation (3.1). Although the assumptions on expected outcomes are fundamentally untestable, suggestive evidence on their validity can still be gathered using tests (Hausman and Rapson 2018) like placebo regression discontinuities, which are shown in Figure C3. Second, the policy change for the implementation and repeal dates need to be identical. If this is not the case, then differences in policy could drive differences between ρ_0 and ρ_1 even if prices would have responded identically to the implementation and repeal of the carbon tax.

[8] All optimal bandwidths are less than half a year.

There are two sources of concern about policy differences between the implementation and repeal of the carbon tax. First, on the date that the carbon tax was implemented there was a cut to federal income tax rates and a mineral resources tax was enacted. No similar changes were made on the repeal date which could, insofar as RD estimates on the implementation date are impacted by these policies, lead to different price changes between the two dates. If the cuts to federal income tax rates increased demand and the mineral resources tax increased marginal costs for electricity producers, then the change in prices on the implementation will be larger than if the carbon tax was enacted in isolation. I test the effect of each of these policies by leveraging the perfectly inelastic demand of electricity in the short-run and the repeal of the mineral resources tax late in 2014. I find no evidence of an increase in demand on the implementation date (Figure C4) or a decrease in wholesale electricity prices when the mineral resources tax is repealed (Figure C5), which suggests that the effect that they had on short-run electricity prices was likely minimal.[9]

The second source of concern about policy differences is that the repeal and implementation of the carbon tax were fundamentally different in nature, despite being similar in magnitude. While the implementation of the carbon tax represents a sharp regression discontinuity, the legal repeal date of July 1, 2014 is, in practice, a fuzzy regression discontinuity. During the time period from July 1 to July 17 of 2014 the carbon tax was actually in the process of being repealed and backdated, but was not yet actually repealed. The perceived marginal costs of producers during this period were thus based on their beliefs that the carbon price would be successfully repealed and backdated. To make the policies comparable, I thus implement a "half-dont RD" where I drop the dates from July 1 to July 17 of 2014 when estimating the RDiT and DiRD approaches.[10]

[9] In Appendix C.1, I discuss these policies in more depth and some reasons why I find they have minimal effects on wholesale electricity prices in the short-run.

[10] Unlike other applications of the "donut RD" (Barreca, Guldi, Lindo, and Waddell 2011) where

Figure 3.3: First Principal Component of Weather (Implementation and Repeal)

(a) Implementation Weather RD

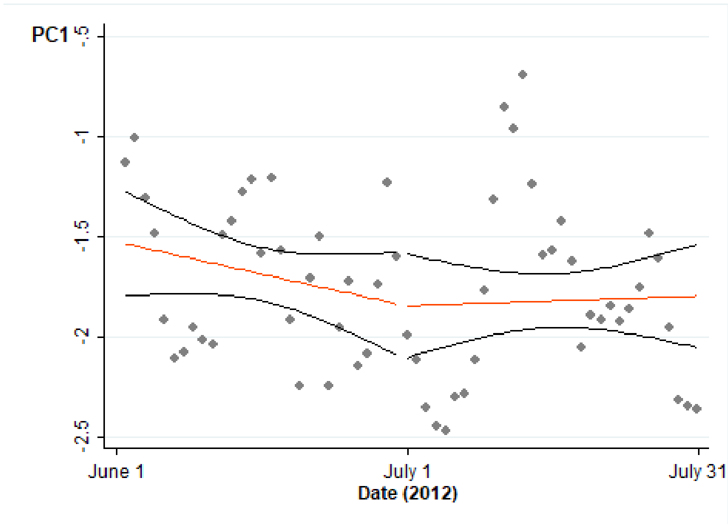

(b) Repeal Weather RD

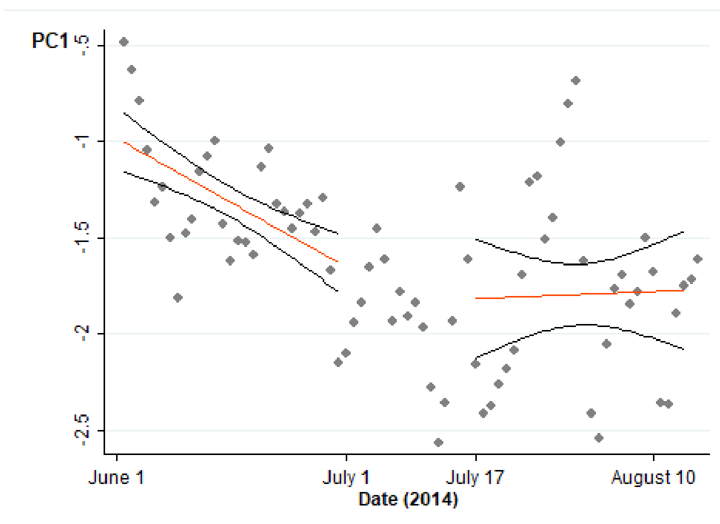

Note: This figure shows the first principal component of weather (i.e., rainfall, minimum temperature, and maximum temperature) around the implementation and repeal dates.

concerns are based on sorting around the cutoff, I drop observations only on one side of the cutoff to account for carbon price uncertainty.

Although implementing a half-donut RD allows me to overcome challenges with comparing a fuzzy and sharp regression discontinuity with the DiRD approach, one of the biggest potential issues is that other factors impacting electricity markets can change during this period. I show this in Figure 3.3 where I estimate equation (3.1) with the first principal component of weather as the dependent variable.[11] Despite the fact that weather is not being impacted by the policy, there is a "discontinuity" in weather conditions when the half-donut is included for the repeal date. This concern is mitigated, to some extent, by the fact that I am able to directly control for these factors, but I also test how big a concern this is more generally by estimating the sharp and half-donut regression discontinuities for the implementation in 2012. Specifically, I will drop a similar number of days around the cutoff for the implementation and compare results from the sharp regression discontinuity with the half-donut RD. If they are similar then this suggests that any differences I find between the repeal and implementation is likely not driven solely by the implementation of the half-donut RD.

3.3.4 Results

I estimate equation (3.1) separately for the carbon tax implementation and repeal dates. The results from a single specification where I do include controls, use a triangular kernel, polynomial order 1, and the data-driven optimal bandwidth is shown in Figure 3.4. The top panel shows that wholesale electricity prices increase sharply when the carbon price is implemented, while the bottom panel shows a smaller, sharp decrease in response to the repeal of the carbon tax.

The results in Figure 3.4 show a single specification, but as I show in Tables 3.1 and 3.2 the conclusions are the same when using a variety of kernels, bandwidths, and

[11] Principal component analysis uses an orthogonal transformation to convert a set of observations into a set of linearly uncorrelated variables, the first of which accounts for the most variation in the data.

Figure 3.4: Impact of Implementation and Repeal on Wholesale Electricity Prices

(a) Implementation RD

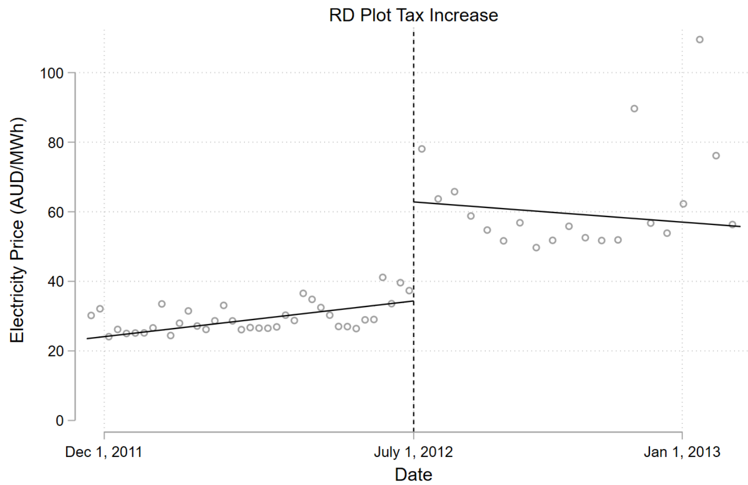

(b) Repeal RD

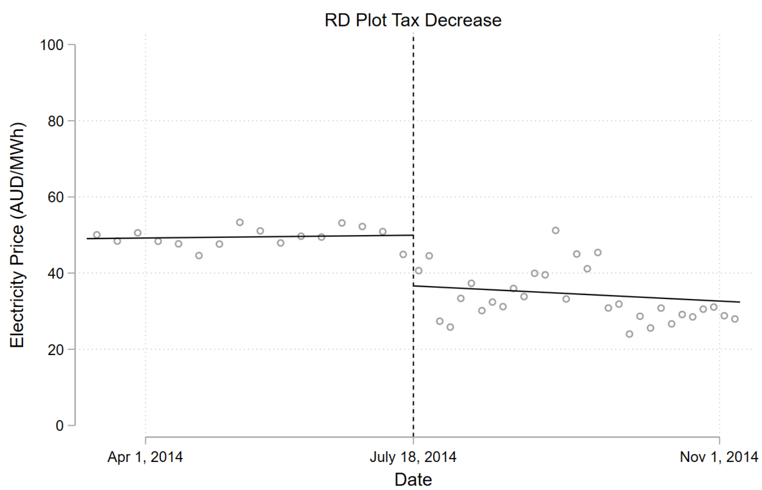

Note: This figure shows the change in wholesale electricity prices in response to the implementation and repeal of the carbon tax. Each panel is estimated separately using a triangular kernel with the data-driven optimal bandwidth and polynomial order 1.

controls. Estimates from the tax implementation and tax repeal are represented by columns labeled "Up" and "Down", respectively. Each cell represents a separate spec-

Table 3.1: RD Estimates with Varying Kernels

		Triangular		Rectangular		Epanechnikov	
Model	Polynomial order	Up	Down	Up	Down	Up	Down
A. No Controls							
	0	27.89	-14.55	28.88	-15.13	24.41	-14.66
		(2.628)	(1.844)	(3.885)	(1.692)	(2.577)	(1.764)
	1	25.11	-10.53	25.89	-9.289	24.20	-8.601
		(3.732)	(4.38)	(3.607)	(3.345)	(3.536)	(3.959)
	2	26.39	-7.469	24.22	-7.899	25.32	-7.448
		(5.52)	(5.035)	(4.982)	(5.966)	(5.621)	(5.363)
B. Controls							
	1	35.17	-13.65	31.19	-10.19	35.97	-13.19
		(2.850)	(2.957)	(3.592)	(3.718)	(3.233)	(2.846)
	2	37.96	-23.80	28.68	-29.00	37.80	-24.37
		(3.654)	(4.936)	(3.447)	(6.06)	(3.75)	(5.225)
C. Trimmed with Controls							
	1	30.90	-15.70	12.19	-13.94	29.92	-16.19
		(1.641)	(2.194)	(2.140)	(2.307)	(1.812)	(2.280)
	2	32.01	-20.90	31.60	-20.28	31.95	-19.94
		(1.718)	(4.225)	(1.891)	(3.937)	(1.873)	(4.539)

Notes: Each coefficient shows a separate regression discontinuity estimate. Specifications vary by panel. Polynomial order for the running variable (date) varies by rows within each panel. The kernel used varies across column pairs and all allow for separate bandwidths on each side of the cutoff. All estimates have an optimal bandwidth determined as put forth by Calonico et al (2015). Columns within column pairs correspond to estimating the change in wholesale electricity price depending on the direction of the carbon tax price change. Estimates of the decrease exclude the dates from July 1, 2014 to July 17, 2014. Standard errors in brackets are clustered at the week of the year level.

ification, which is defined by the event (implementation or repeal), model, polynomial order, bandwidth, and kernel that corresponds with that cell. The difference between Table 3.1 and Table 3.2 is that the former varies the kernel, while the latter varies bandwidths using the triangular kernel. Panel A for both tables show estimates of equation (3.1) without controls, X_t, while panels B and C include controls for weather variables, fossil fuel prices, and the day of the week. Panel C is identical to Panel B, but with outliers dropped when calculating the daily electricity price.

Focusing on Table 3.1, there are a few key takeaways. First, specifications that include controls tend to provide larger estimates of the price changes across both carbon

Table 3.2: RD Estimates with Varying Bandwidths

		Optimal BW		50% Optimal		150% Optimal	
Model	Polynomial order	Up	Down	Up	Down	Up	Down
A. No Controls							
	0	27.89	-14.55	27.95	-14.45	29.56	-14.02
		(2.628)	(1.844)	(3.453)	(2.795)	(2.516)	(1.782)
	1	25.11	-10.53	29.30	-11.88	27.07	-12.46
		(3.732)	(4.38)	(4.52)	(5.96)	(3.2)	(3.164)
	2	26.39	-7.469	36.98	-13.01	24.28	-11.00
		(5.52)	(5.035)	(5.096)	(6.706)	(4.397)	(4.146)
B. Controls							
	1	35.17	-13.65	37.23	-4.236	26.67	-20.99
		(2.850)	(2.957)	(4.943)	(3.451)	(2.778)	(3.906)
	2	37.96	-23.80	42.26	-1.69	26.91	-21.93
		(3.654)	(4.936)	(6.00)	(4.535)	(3.74)	(4.09)
C. Trimmed with Controls							
	1	30.90	-15.70	14.79	-9.343	30.21	-16.90
		(1.641)	(2.194)	(1.802)	(2.485)	(1.573)	(1.865)
	2	32.01	-20.90	18.04	1.021	29.26	-17.07
		(1.718)	(4.225)	(1.622)	(1.078)	(1.682)	(2.877)

Notes: Each coefficient shows a separate regression discontinuity estimate. Specifications vary by panel. Polynomial order for the running variable (date) varies by rows within each panel. Bandwidths vary across column pairs, but all optimal bandwidths are estimated with a triangular kernel, separate bandwidths on either side of the cutoff and optimal bandwidth determined by the procedure suggested by Calonico et al. (2015). Columns within column pairs correspond to estimating the change in wholesale electricity price depending on the direction of the carbon tax price change. Estimates of the decrease exclude the dates from July 1, 2014 to July 17, 2014. Standard errors in brackets are clustered at the week of the year level.

tax changes.[12] This tends to hold even when outliers are dropped, which shows that the pattern evident in the raw wholesale electricity price data is not a result of outlier periods or other factors like fuel prices. Second, estimates with the triangular and epanechnikov kernels are remarkably consistent across nearly all specifications and polynomials. They show a pattern of prices rising by significantly more when the carbon tax is implemented compared to when it is repealed. This result is generally similar with the rectangular kernel, but results appear to be much more volatile. Since the rectangular kernel puts

[12]One difficulty with a local linear specification is that controls can be difficult to include or influenced by some outlier observations given the limited observations. One proposed solution by Hausman and Rapson (2018) is the augmented local linear approach, which first estimates the impact of the control using the entire window. Then, using these residuals a local linear specification is estimated with a narrow bandwidth (e.g., 30 days).

Table 3.3: Difference-in-Regression Discontinuity Estimates (Prices)

	Price			
	(1)	(2)	(3)	(4)
Increase (2012)	38.611***	38.879***	43.815***	38.879***
	[4.952]	[5.164]	[3.759]	[4.130]
Decrease (2014)	-16.090***	-14.338***	-9.390	-11.970*
	[5.001]	[4.525]	[5.743]	[6.591]
Observations	238	342	238	342
Clusters	36	52	36	52
F-test (Increase = -Decrease)	10.24	12.78	23.61	10.48
Kernel	Rectangular	Triangular	Rectangular	Triangular
Controls	No	No	Yes	Yes

Notes: Each column represents results from estimating equation (3.3) with alternative specifications. For all specifications I estimate a local linear regression with optimal bandwidth determined by the procedure suggested in Calonico et al. (2015). Columns alternate between the rectangular (uniform) and triangular kernels, with the first two showing the results without controls and the last two the results with controls. All F-tests indicate a statistically significant difference between the magnitude of the price change when the carbon tax is implemented compared to when it is repealed. Estimates of the decrease exclude the dates from July 1, 2014 to July 17, 2014. Standard errors in brackets are clustered at the week of the year level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

equal weights on all observations in the bandwidth, this may reflect high leverage observations far from the cutoff.

Table 3.2 shows estimates using the triangular kernel and alternative bandwidths. The first column pair of Table 2 displays estimates identical to the first column pair of Table 3.1 so that these can be compared to the results obtained under alternative bandwidths. The second and third column pairs are estimates where the optimal bandwidth is decreased and increased by 50%, respectively. The results in Table 3.1 are broadly corroborated by Table 3.2. Despite large bandwidth changes, across all specifications the magnitude of the price change is larger when the carbon tax is implemented compared to when it is repealed. In summary, the results from the regression discontinuities in time approach seem to indicate that prices rose by more when the carbon tax is implemented compared to when it is repealed.

I estimate if these differences are statistically significant with equation (3.3), where I test if the price change from implementation (ρ_0) is equal in magnitude to the price change from the repeal (ρ_1). The results from this approach are shown in Table 3.3, which show a pattern similar to those from the RDiT approach. Across all specifications, I find prices increase by more at the implementation date compared to the decrease in prices when the carbon tax is repealed. In other words, prices seemed to rise faster on implementation of the carbon tax than they fell when the carbon tax was repealed.

3.3.5 Robustness Checks

There are two identification concerns that I have not addressed thus far. First, I am implementing a half-donut RD only for the repeal date, which could be driving these results. To explore whether this is the case, I estimate a half-donut RD for the implementation date, the results of which are shown in Figure C6. Regardless of the kernel or inclusion of controls, I find a similar pattern of estimates with all half-donut RD estimates lying within the confidence intervals for the sharp RD estimates. Although this does not fully address the concern, it does indicate that implementing a half-donut RD for the implementation date would not change the conclusions from Tables 3.1 to 3.3.

Second, even dropping the period from July 1 to July 17 in 2014, the repeal date could still present a fuzzy discontinuity if plants had a positive belief that they would not have to pay the carbon price for emissions before June 30. The explanatory memorandum for the *Clean Energy Legislation (Carbon Tax Repeal) Bill 2013* was clear that any bills passed after July 1, 2014 would be backdated to July 1, but it is still possible that plants believed a repeal may pass in advance of this date. To investigate this, I estimate a regression discontinuity in time with the date the *Clean Energy Legislation (Carbon Tax*

Repeal) Bill 2013 was blocked in the Senate as the cutoff. If plants maintained a positive belief on the probability of repeal before July 1, 2014, then this event should present a downward shock to beliefs that then appears as an increase in prices. The results from estimating a regression discontinuity in time using this date as a cutoff are shown in Figure C7. I find no evidence of any change in prices on this date, which provides some evidence that plants knew that only when the new Senate came to power on July 1, 2014 would there be any chance of a repeal bill being passed.

3.4 Did producers respond asymmetrically?

In the previous section I found a consistent pattern of immediate price increases from the implementation being larger in magnitude than the immediate decrease from the carbon tax repeal. This result holds across a wide variety of specifications and it does not change when I estimate a half-donut RD on the implementation date. These results, however, do not provide conclusive evidence that firms are responding differently to the carbon tax implementation compared to the repeal for two reasons. First, I have not ruled out shifts in demand that increase prices when the carbon tax is repealed. In fact, the differences in weather evident in Figure 3.3 seem to indicate that demand likely shifted during this time period. Second, prices could be responding differently to the implementation compared to the repeal due to differences in emissions intensities of price-setting plants. This could be a result of investment in emissions reductions technologies or, alternatively, differences in which plants are setting prices.

In this section, I leverage a key advantage of wholesale electricity markets – detailed information on the supply curve used to set prices. Specifically, I use AEMO data on daily price offers and period quantity offers to form plant level bid (supply) curves. I then sort each band from all bids in order of price to form an aggregate offer (supply)

Figure 3.5: Example Offer Curves around the Cutoff Dates

(a) Offer Curves 2012

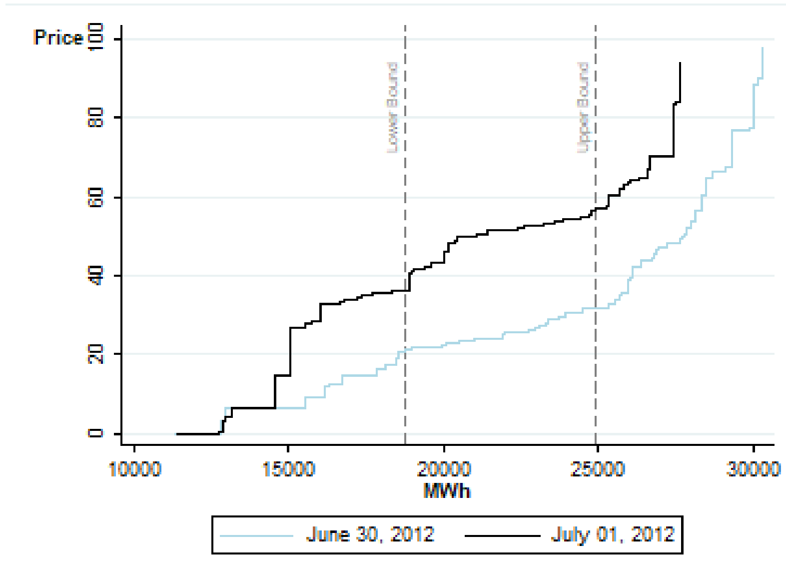

(b) Offer Curves 2014

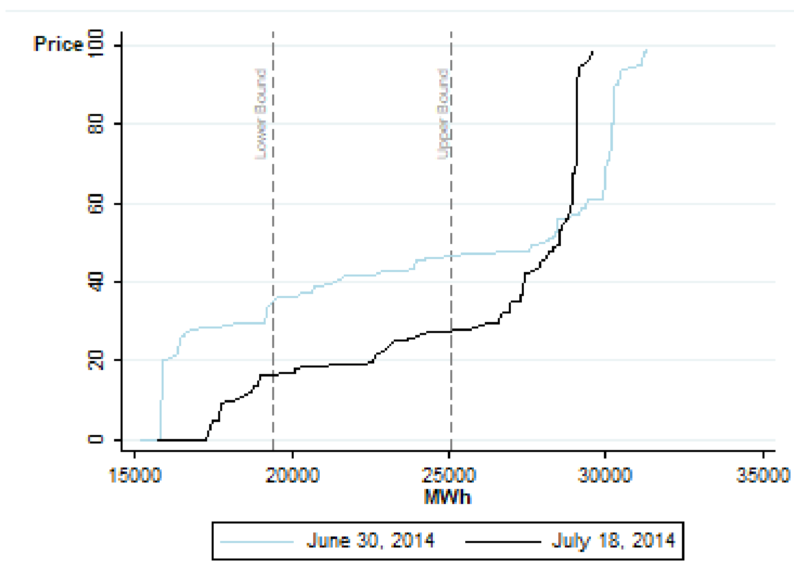

Note: This figure plot the offer curves for the 7-7:30 a.m. period for June 30 and July 1, 2012 in panel (a) and June 30 and July 18, 2014 in panel (b). The dashed vertical lines provide reference points for the minimum and maximum equilibrium quantities in that period (7-7:30 a.m.) within two weeks of the tax change.

curve. Examples of offer curves are shown in Figure 3.5, which plot the offer curves for the 7-7:30 a.m. period for dates around the tax changes. The dashed vertical lines provide reference points for the minimum and maximum equilibrium quantities in that period within two weeks of the tax change.

With this information, I ask the following questions. First, is the asymmetry in price changes a result of lower emissions rates for price-setting plants when the carbon price is repealed? In other words, is the change in prices reflecting differential impacts of the same carbon price on marginal costs? Second, how did average willingness to accept in the area where plants are potentially setting prices respond to the implementation compared to the repeal of the carbon tax? Lastly, did individual plants, on average, change their willingness to accept differently with the two carbon tax changes?

3.4.1 Marginal Plant Emissions

Combining AEMO data on which plants are setting prices with their emissions rates, I am able to calculate the daily average marginal plant emissions rates. If plants setting prices when the carbon price is implemented have higher emissions rates, then this could drive differences in price changes. The asymmetry in prices would thus not reflect differential responses by producers, but differences in emissions.

In Figure 3.6, I plot the average and 95 percent confidence interval for marginal plant emissions rates before and after the implementation and repeal of the carbon tax. Specifically, pre and post periods are defined as June 1 to June 30 and July 1 to August 17 (1 month after repeal) for 2012 (implementation) and 2014 (repeal). The key takeaway from this figure is that the marginal plants during this period were approximately similar with, if anything, higher average emissions rates of price-setters when the carbon tax is repealed compared to implemented. Importantly, this means that the asymmetric

Figure 3.6: Daily Average Marginal CO_2e Emissions Rates

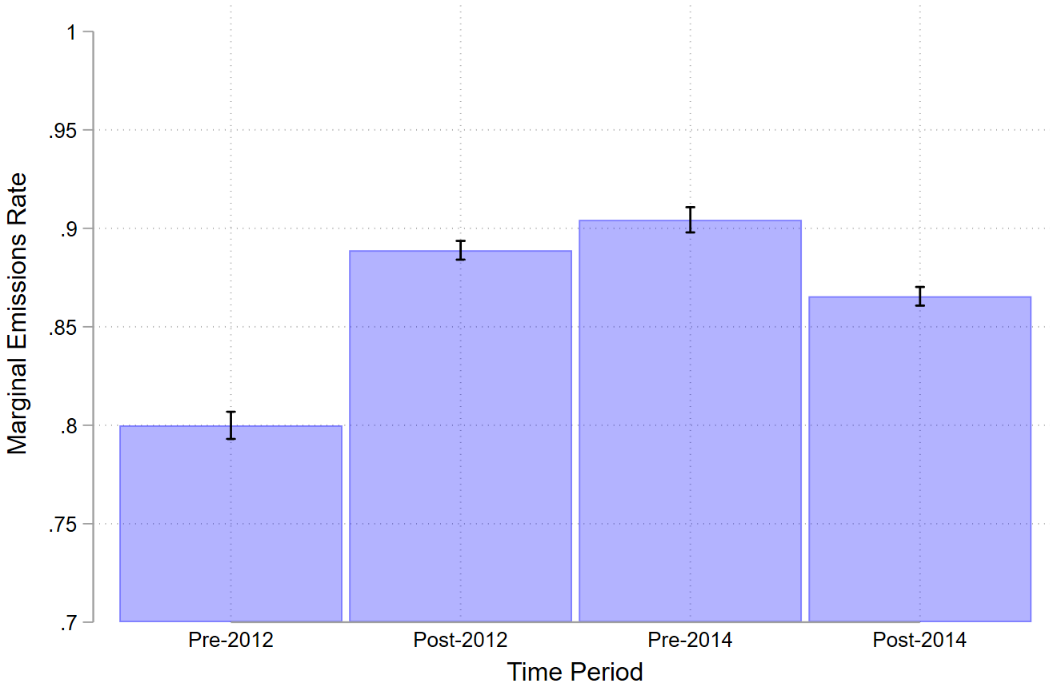

Note: This figure plots the average emissions rates of marginal (price-setting) plants, which are shown with the blue bars, for four different periods. Each of the pre-periods are June 1 to June 30, while the post-period is from July 1 to August 17. For 2014 I drop dates from July 1 to July 17 to match the regression analysis. The range shown with the black bars represent 95 percent confidence intervals.

responses in prices does not reflect significantly lower average emissions rates of marginal plants when the carbon tax is repealed.

3.4.2 Average Changes in Willingness to Accept

Although the asymmetry in price changes is not driven by differences in emissions rates, this does not necessarily mean that producers are responding differently to the implementation and repeal of the carbon price. Unobserved shocks to demand, for example, could be driving up prices during the time period of the carbon tax repeal. To test if this is driving the difference in price responses, I use demand projections that AEMO provides to electricity generators. These include 10th and 90th percentile demand projections, which I use as lower and upper bounds for the range at which firms can be potentially marginal on the supply curve. I then construct the average willingness-to-accept in this region with the contemporaneous demand projections, using the 2012 demand projections around both cutoffs, and lastly using the 2014 demand projection around both cutoffs. Intuitively, what this allows me to do is analyze the average changes in the supply curve based on contemporaneous demand, 2012 demand, and 2014 demand.

Regardless of the demand projections that I use (i.e., region of the supply curve I analyze), I find asymmetries in changes to willingness-to-accept in response to the implementation compared to the repeal. This is evident graphically in Figures C8, C9, and C10, which plots the daily averages from June 1 to August 1 for 2012 and 2014. Across all three figures, there is a clear pattern of sharp increases in 2012 that are not mirrored by the repeal in 2014. To test this empirically, I estimate equation (3.3) using the daily averages shown in the Figures. Table 3.4 shows that these results mirror the findings from analyzing prices in Table 3.3. Analyzing behavior along the supply curve, not just marginal producers, holding demand constant for the implementation and repeal dates thus yields similar results to the analysis using equilibrium prices. Unobserved demand shocks thus do not appear to be driving the results in equilibrium prices, since the asymmetry remains even when holding demand constant for the implementation and repeal dates.

Table 3.4: Difference-in-Regression Discontinuity Estimates (WTA)

	Average Willingness-to-Accept					
	(1)	(2)	(3)	(4)	(5)	(6)
Increase (2012)	45.215***	44.979***	45.728***	44.473***	38.452***	37.379***
	[2.935]	[2.486]	[2.963]	[2.075]	[3.126]	[2.258]
Decrease (2014)	-9.744***	-8.184***	-10.966	-1.301	-8.211***	-7.430***
	[2.472]	[2.017]	[6.922]	[6.428]	[2.085]	[1.801]
Observations	206	234	214	166	166	214
Clusters	32	36	32	24	24	32
F-test (Increase = -Decrease)	85.45	132.1	21.32	40.85	64.79	107.54
Demand Proj.	Cont.	Cont.	2012	2012	2014	2014
Kernel	Rectangular	Triangular	Rectangular	Triangular	Rectangular	Triangular

Notes: Each column represents results from estimating equation (3.3) with daily average willingness-to-accept in the region of projected demand using alternative specifications. For all specifications I estimate a local linear regression with optimal bandwidth determined by the procedure suggested in Calonico et al. (2015). Columns alternate between the rectangular (uniform) and triangular kernels. The first two columns use the region of projected demand during the realized dates and periods. The second two columns use projected demand for the day of the year and period from 2012 for both events in 2012 and 2014. The last two columns use projected demand for the day of the year and period from 2014 for both events in 2012 and 2014. All F-tests indicate a statistically significant difference between the magnitude of the price change when the carbon tax is implemented compared to when it is repealed. Estimates of the decrease exclude the dates from July 1, 2014 to July 17, 2014. Standard errors in brackets are clustered at the week of the year level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.4.3 Individual Firm Responses

Instead of estimating the response of market average willingness-to-accept in the region of the equilibrium quantity, one can examine the response of individual electricity generators. Each plant submits a supply function for each period of the day, which I simplify by focusing on the band (step) in the supply function that is closest to the AEMO projected demand on the aggregate supply curve. I then calculate the daily average willingness-to-accept for this, most likely to be marginal, band for each plant.[13]

[13]Some plants are unlikely to be marginal during certain periods of the day, so I do not incorporate bands that are more than 1000 MW away from the projected demand.

Table 3.5: Difference-in-Regression Discontinuity Estimates (Plant WTA)

	Plant Average Willingness-to-Accept					
	(1)	(2)	(3)	(4)	(5)	(6)
Increase (2012)	40.237***	40.658***	34.642***	36.880***	33.971***	35.878***
	[2.406]	[2.146]	[1.377]	[1.183]	[1.461]	[2.123]
Decrease (2014)	-7.812***	-10.022***	1.029	1.223	-3.169**	-3.828***
	[2.052]	[2.908]	[4.228]	[4.587]	[1.164]	[1.085]
Observations	12,322	17,448	5,405	8,539	5,782	8,971
Clusters	32	44	12	20	16	24
F-test (Increase = -Decrease)	105.17	71.84	64.35	64.7	271.93	180.71
Plant by Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Demand Proj.	Cont.	Cont.	2012	2012	2014	2014
Kernel	Rectangular	Triangular	Rectangular	Triangular	Rectangular	Triangular

Notes: Each column represents results from estimating equation (3.3) where the dependent variable is the daily average of a plant's willingness-to-accept for the band (step) that is closest to projected demand. If a plant's closest band is more than 1000 MW away from the projected demand, then that period is not included in the daily average. For all specifications I estimate a local linear regression where I include plant by year fixed effects with optimal bandwidth determined by the procedure suggested in Calonico et al. (2015). Columns alternate between the rectangular (uniform) and triangular kernels. The first two columns use the region of projected demand during the realized dates and periods. The second two columns use projected demand for the day of the year and period from 2012 for both events in 2012 and 2014. The last two columns use projected demand for the day of the year and period from 2014 for both events in 2012 and 2014. All F-tests indicate a statistically significant difference between the magnitude of the price change when the carbon tax is implemented compared to when it is repealed. Estimates of the decrease exclude the dates from July 1, 2014 to July 17, 2014. Standard errors in brackets are clustered at the week of the year level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

With this measure, I am then able to repeat the analysis shown in Table 3.4, but with plant by year fixed effects. Importantly, this means that my estimates of how the implementation and repeal of the carbon tax impacted willingness to accepts is now driven by average changes in individual plant behavior.

The results from including individual plant by year fixed effects are shown in Table 3.5. Including plant by year fixed effects yields nearly identical results with large increases in willingness-to-accept when the carbon tax is implemented, but small changes following its repeal. Importantly, this indicates that the results focused on the aggregate supply curve willingness-to-accept are not driven by changes in which plants are in that part of the supply curve.

3.5 Conclusion

In this paper, I estimate the impact that the implementation and repeal of the Australian carbon tax has on wholesale electricity prices and electricity generator behavior. Leveraging regression discontinuities in time I find that prices increase more when the carbon tax is implemented than they decrease when the carbon tax is repealed. I then utilize one of the unique features of wholesale electricity markets – information on aggregate and individual supply curves. With this information, I am able to identify if the asymmetric responses is a result of asymmetric shifts in the supply curve, and not unobserved shocks to demand. I find that the willingness-to-accept for steps of the supply curve that are most likely to be marginal (i.e., set prices) increase more when the tax is implemented compared to the decrease when it is repealed.

There are important caveats to this analysis. The analysis in this paper leverages regression discontinuities for identification so the asymmetry that is identified is inherently short-run. Prices appear to decrease over time following the repeal, but I cannot identify the length of time for which the asymmetry remains. Moreover, focusing on wholesale electricity markets allows me to examine the behavior of individual plants, but it does not allow me to identify the impact on the retail electricity prices that consumers pay. If electricity retailers also respond asymmetrically to changes in costs, asymmetric responses in the wholesale electricity market may have significant impacts on retail electricity prices. Alternatively, since retail electricity contracts often offer a fixed price for a length of time, such as a year, short-run asymmetric pass-through in the wholesale electricity market may have minimal impacts on retail electricity prices. Estimating asymmetric pass-through in wholesale and retail electricity markets, in particular, deserves the attention of future work.

Appendix A

Appendix: Who Bears the Load? Carbon Taxes and Electricity Markets

A.1 Theory and Derivations

A.1.1 Symmetric Oligopoly Incidence

All firms are symmetric so I drop the subscript on the marginal cost in the profit function:

$$\pi_i = (P - MC)q_i$$

Differentiating with respect to the tax on inputs and substituting in definitions of ρ and γ:

$$\frac{d\pi_i}{d\tau} = \frac{dq_i}{d\tau}(P - MC) + q_i[\rho - \gamma]$$

The change in producer surplus is the sum of the change in all firms profits:

$$\frac{dPS}{d\tau} = \sum_{i=1}^{n} \frac{d\pi_i}{d\tau} = \frac{dQ}{d\tau}(P - MC) + Q[\rho - \gamma]$$

which comes from defining $Q = \sum_{i=1}^{n} q_i$ and noting that $\frac{dQ}{d\tau} = \sum_{i=1}^{n} \frac{dq_i}{d\tau}$. Lastly, multiply (P-MC) by P/P, apply the definition of the Lerner index, and note that the change in consumer surplus is $dCS/d\tau = -Q\rho$:

$$I = \frac{\rho}{\gamma - \rho - L\frac{dQ}{d\tau}\frac{P}{Q}}$$

A.1.2 Asymmetric Oligopoly Incidence

Proof: Each of the N asymmetric producers maximize profits by selling q_i units of output for price P:

$$\pi_i = (P - MC_i)q_i$$

Differentiating profits with respect to τ and substituting the definitions for γ_i and ρ:

$$\frac{d\pi_i}{d\tau} = \frac{dq_i}{d\tau}[P - MC_i] + q_i[\rho - \gamma_i]$$

The change in producer surplus due to the tax is given by the sum of the change in profits for all firms:

$$\frac{dPS}{d\tau} = \sum_{i=1}^{n} \frac{d\pi_i}{d\tau} = \sum_{i=1}^{n} \frac{dq_i}{d\tau}[P - MC_i] + Q\rho - \sum_{i=1}^{n} q_i\gamma_i$$

Multiplying the producer suplus by Q/Q and noting that the change in consumer surplus is given by $dCS/d\tau = -Q\rho$:

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{\rho}{\frac{\sum_{i=1}^{n} q_i \gamma_i}{Q} - \rho - \frac{\sum_{i=1}^{n} \frac{dq_i}{d\tau}[P - MC_i]}{Q}}$$

Multiplying the last term on the denominator by P/P, and plugging in the definitions for L_i and $\bar{\gamma}$ gives:

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{\rho}{\bar{\gamma} - \rho - \frac{P}{Q} \sum_{i=1}^{n} \frac{dq_i}{d\tau} L_i}$$

which is equation 1.2. ■

A.1.3 Asymmetric Output Tax

In this section, I first lay out the most general model in Weyl and Fabinger (2013) that allows for asymmetric firms, imperfect competition, differentiated products, and firm-specific output taxes. I will then simplify assume a single output tax and market and derive incidence in this setting.

Following Weyl and Fabinger (2013), allow for firm-specific taxes $\boldsymbol{\tau}$, assuming by way of normalization: $(\boldsymbol{\tau} \cdot \boldsymbol{q})/Q = 1$. The size of the tax imposed is denoted by a scalar t_τ. To form a basis for all τ, assume that either $d\boldsymbol{p}/d\sigma_i$ or $d\boldsymbol{q}/d\sigma_i$ are of full rank with σ_i defined as a single-dimensional strategic variable that determines the firms actions. By linear independence, any τ is a linear combination of τ_i (ith component of vector $\boldsymbol{\tau}$) and the collection of the coefficients of the linear combination are λ^τ. In this general case, as Weyl and Fabinger (2013) show:

$$\frac{dPS}{d\tau} = \sum_i \lambda_i^\tau \left[\frac{d\boldsymbol{p}}{dt_{\tau_i}} \cdot \boldsymbol{q} + \frac{d\boldsymbol{q}}{dt_{\tau_i}} (\boldsymbol{m} - \boldsymbol{t}) \right] - Q$$

where $\boldsymbol{m} \equiv \boldsymbol{p} - \boldsymbol{mc}$ and $\boldsymbol{t}$ are the output taxes.

To simplify, note that in the particular case where there is a single output tax, $\boldsymbol{\tau}$ is a vector of ones, by the assumed normalization. The coefficients of the linear combination, λ^{τ}, of components τ_i are now all also equal to 1. The size of the tax, t_{τ}, now reflects the size of the tax t. Assuming all firms operate only in the market with the output tax, the formula reduces to:

$$\frac{dPS}{d\tau} = \sum_i \left[\frac{dP}{dt} \cdot q_i + \frac{dq_i}{dt}(m - t) \right] - Q$$

Plugging in for m and simplifying:

$$\frac{dPS}{d\tau} = \frac{dP}{dt} \sum_i q_i + \sum_i \frac{dq_i}{dt}(P - MC_i - t) - Q$$

Following similar steps to those shown in A.1.2 including multiplying by P/P and plugging in the formula for $L_i = (P - MC_i)/P$ where t is incorporated in MC_i:

$$\frac{dPS}{d\tau} = \frac{dP}{dt} Q + \sum_i \frac{dq_i}{dt} \cdot P \cdot L_i - Q$$

$$\frac{dPS}{d\tau} = Q[\frac{dP}{dt} + \sum_i \frac{dq_i}{dt} \cdot \frac{P}{Q} L_i - 1]$$

Lastly, noting that $dCS/d\tau = -dP/dtQ$ and the formula for incidence:

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{\rho}{1 - \rho - \frac{P}{Q} \sum_{i=1}^{n} \frac{dq_i}{d\tau} L_i}$$

This highlights that the key difference in terms of the formula comes through the term $\bar{\gamma}$, which measures the change in marginal costs per dollar change in the tax.

A.1.4 Incidence and Electricity

Starting with the first order condition where $\gamma_i = e_i$:

$$\frac{d\pi_{im}}{d\tau} = \frac{dq_i}{d\tau}(P_i - MC_i) + q_i(\rho_m - e_i)$$

Since the price change for individual plant i is the same as the zonal (regional) price change in the zone that they are in (m), I have rewritten ρ_i as ρ_m. Summing over all firm profits to get the change in producer surplus:

$$\frac{dPS}{d\tau} = \sum_{i=1}^{n} \frac{d\pi_i}{d\tau} = \sum_{i=1}^{n} \frac{dq_i}{d\tau}(P_i - MC_i) + \sum_{i=1}^{n} q_i\rho_m - \sum_{i=1}^{n} q_i e_i$$

Defining q_{sm} as the total electricity generated in zone m and simplifying:

$$\frac{dPS}{d\tau} = \sum_{i=1}^{n} \frac{dq_i}{d\tau}(P_i - MC_i) + \sum_{m=1}^{5} q_{sm}\rho_m - \sum_{i=1}^{n} q_i e_i$$

Multiplying the first term by PQ/PQ and all other terms by Q/Q and applying the following definitions: $\bar{\rho_s} = \sum_{m=1}^{5} q_{sm}\rho_m/Q$, $\bar{e} = \sum_{i=1}^{n} q_i e_i/Q$, and $L_i = (P_i - MC_i)/P_i$ gives:

$$\frac{dPS}{d\tau} = Q[\sum_{i=1}^{n} \frac{dq_i}{d\tau} L_i \frac{P_i}{Q} + \bar{\rho_s} - \bar{e}]$$

Lastly, the change in consumer surplus is no longer $dCS/d\tau = -\rho Q$ since the change in prices for customers is now different across zones. Instead, it is given by $dCS/d\tau = -\sum\limits_{m=1}^{5} \rho_m q_{dm}$ where q_{dm} equals the quantity consumed (demanded) in zone m. Multiplying consumer surplus by Q/Q, defining $\bar{\rho_d} = \sum\limits_{m=1}^{5} q_{dm}\rho_m/Q$, and simplifying gives equation (4):

$$I = \frac{dCS/d\tau}{dPS/d\tau} = \frac{\bar{\rho_d}}{\bar{e} - \bar{\rho_s} - \sum\limits_{i=1}^{n} \frac{dq_i}{d\tau} L_i \frac{P_i}{Q}}$$

A.2 Markups Theoretical Framework

I now characterize a firm's optimal bidding strategy in the National Electricity Market's real time wholesale electricity market. First, I follow Woerman (2019) and derive the condition for optimal markup in an electricity market with zonal prices. I then use the optimal markup rule to derive an expression for how firms change offers and markups given changes in the price of carbon emissions.

Optimal Markup Rule

Let the offer curve for every plant, S_i, be composed of offer prices and quantities given by $\boldsymbol{b}_i$ and $\boldsymbol{q}_i$, respectively. Each plant is located in one of the five electricity zones, which is constant over time. The market that each electricity generator is in, however, is not constant. When no transmission constraints are binding then the market is fully integrated and the market clearing price is the same for all electricity generators. As transmission constraints bind, however, the NEM segments and prices differ across zones. I denote each market as m and the set of feasible markets as $\mathcal{M}$.[1]

Regardless of the levels of transmission congestion, the market clears by setting the price in each market such that production equals demand net of flows between markets:

$$\sum_{i=\mathcal{I}_m} S_i(p_m) = D_m + T_m$$

where $\mathcal{I}_m$ is the set of plants in market m, D_m is the demand in market m, and T_m is the transmission flows out of market m. Both demand and transmission flows are assumed to be perfectly inelastic. This can be rewritten from the perspective of arbitrary plant i

[1] The ways in which the NEM can segment are limited by the physical connections between zones. For example, Tasmania is connected to the rest of the NEM through Victoria so Tasmania is either completely segmented or in a market that also includes Victoria.

as:

$$S_i(p_m) = R_i(p_m) = D_m + T_m - \sum_{\substack{j \in \mathcal{I}_m \\ j \neq i}} S_j(p_m) \tag{A.1}$$

where $R_i(p_m)$ is the residual demand of plant i. Define the cost to produce at plant i in market m as $C_i(q)$ then the profit earned is:

$$\pi_i = S_i(p_i)p_i - C_i(S_i(p_i)) \tag{A.2}$$

where the price for plant i is equal to the price of the market that it is in.

When firms submit their offers they face uncertainty in the offer curves of other firms, the level of demand, and transmission congestion. Importantly, this means that even though they know their forward positions, uncertainty in transmission congestion leads to uncertainty in the distribution of these positions over markets. Each firm maximizes its total profits over all profits and forward positions:

$$\max_{\{S_i | i \in \mathcal{I}_{\{}\}} \mathbb{E}[\sum_{i \in \mathcal{I}_{\{}} (S_i(p_i)p_i - C_i(S_i(p_i))) + \sum_{m \in \mathcal{M}} (P^F_{fm} - p_m)Q^F_{fm}] \tag{A.3}$$

The set of plants owned by firm f is $\mathcal{I}_f$ and the price and quantity of forward contracts are defined as Q^F_{fm} and P^F_{fm}.

Woerman (2019) derives the firm's first-order condition with two assumptions. First, on the margin the firms' offer price does not impact whether a transmission constraint is binding. Second, every plant has constant marginal costs up to the plant's capacity constraint. Defining the plant's marginal cost as c_i and the firm's production net of the forward position as $Q_{fm}(p_m)$ yields the first order condition for plant i's kth offer price, b_{ik}:

$$b_{ik} - c_i = \frac{[Q_{fm}(b_i k)]}{-[\partial R_i(b_{ik})/\partial p_m]} \tag{A.4}$$

The difference between the bid price and the marginal cost depends on expected net production and the expected slope of the residual demand curve when this offer clears. When a firm has more production in a market it has a greater incentive to exercise market power and when residual demand is more inelastic then a firm has a greater ability to exercise market power.

Carbon Price and Markups

In the context of a price on carbon, define marginal costs as in equation 1.3. Since marginal costs for each plant is additively separable in marginal input costs and marginal emissions costs, taking the derivative of the optimal markup rule with respect to a marginal change in the price of carbon emissions:

$$\frac{db_{ik}}{d\tau} = e_i + \frac{\frac{\partial[\partial R_i(b_{ik})/\partial p_m]}{\partial \tau}\cdot[Q_{fm}(b_{ik})] - \frac{\partial[Q_{fm}(b_{ik})]}{\partial\tau}\cdot[\partial R_i(b_{ik})/\partial p_m]}{[\partial R_i(b_{ik})/\partial p_m]^2}$$

Splitting the second-term into two terms and simplifying:

$$\frac{db_{ik}}{d\tau} = e_i + \frac{\frac{\partial[\partial R_i(b_{ik})/\partial p_m]}{\partial \tau}\cdot[Q_{fm}(b_{ik})]}{[\partial R_i(b_{ik})/\partial p_m]^2} - \frac{\frac{\partial[Q_{fm}(b_{ik})]}{\partial\tau}}{[\partial R_i(b_{ik})/\partial p_m]}$$

Multiplying the top and bottom of the second term by $[Q_{fm}(b_{ik})]$, noting the common term, and rearranging:

$$\frac{db_{ik}}{d\tau} = e_i + \left[\frac{[Q_{fm}(b_ik)]}{-[\partial R_i(b_{ik})/\partial p_m]}\right]\cdot\left[\frac{\frac{\partial[Q_{fm}(b_{ik})]}{\partial\tau}}{[Q_{fm}(b_ik)]} - \frac{\frac{\partial[\partial R_i(b_{ik})/\partial p_m]}{\partial \tau}}{[\partial R_i(b_{ik})/\partial p_m]}\right]$$

Lastly, from the optimal bid rule we know:

$$b_{ik} - mc_i = \frac{[Q_{fm}(b_ik)]}{-[\partial R_i(b_{ik})/\partial p_m]}$$

which we can substitute in to get the final expression:

$$\frac{db_{ik}}{d\tau} = e_i + [b_{ik} - mc_i] \cdot \left[\frac{\frac{\partial [Q_{fm}(b_{ik})]}{\partial \tau}}{[Q_{fm}(b_ik)]} - \frac{\frac{\partial [\partial R_i(b_{ik})/\partial p_m]}{\partial \tau}}{[\partial R_i(b_{ik})/\partial p_m]} \right] \tag{A.5}$$

The first term on the right-hand side shows that, all else equal, plants with high emissions rates will increase their offer prices by more than those with low emissions rates. The remainder of the right-hand side represents the impact of the carbon price on markups. The term outside the brackets is the markup before the change in the carbon price, while the terms inside the bracket represent the relative changes in expected net production and the expected slope of the residual demand curve. Intuitively, the difference between these two factors captures the relative impact on the *incentive* to mark up (net quantity) compared to the relative impact on the *ability* to mark up (slope of residual demand).

The intuition of equation A.5 is evident in two simple examples. First, consider the case where all firms have the same emissions rates. An increase in the carbon price in this case has a symmetric impact on costs for all firms. Figure A.1 highlights that the carbon price has no impact on markups. Panel (a) of Figure A.1 shows a hypothetical supply curve before and after the change in the carbon price. Since all firms experience the same increase in costs, the amount of electricity generated by each firm remains constant. Moreover, taking a symmetric increase in offer prices by Firms A and C as given, Panel (b) plots the residual demand curve for the marginal electricity generator, Firm B, before and after the carbon price change. The "slope" of the residual demand curve is identical compared to before the increase in the tax. Returning to equation A.5, this highlights the fact that when emissions rates are homogeneous across all electricity generators the second term on the right-hand side is 0 and offer prices increase by the change in marginal costs, which leaves markups constant.

In contrast, assume that Firms A and B have the same emissions rates, but Firm C has a higher emissions rate. Figure A.1 panel (c) shows the hypothetical supply curve before and after the change in the carbon price, assuming that Firms A and C increase their offer prices by the increase in marginal costs. The increase in marginal costs does not lead to changes in who produces between the three firms, but panel (d) shows that the slope of the residual demand curve becomes more inelastic due to the large increase in marginal costs for Firm C. The marginal firm now has an increased ability to exert market power, which would lead to an increase in markups which is shown in panel (c) and a larger price increase than if Firm C has a lower emissions rate. Although these are both specific, stylized examples, they highlight the important role of heterogeneity in emissions rates on offer prices, markups, and equilibrium prices. This is a factor that I will return to when interpreting the empirical estimates on how markup changes with the price on carbon emissions.

Figure A.1: Carbon Price History

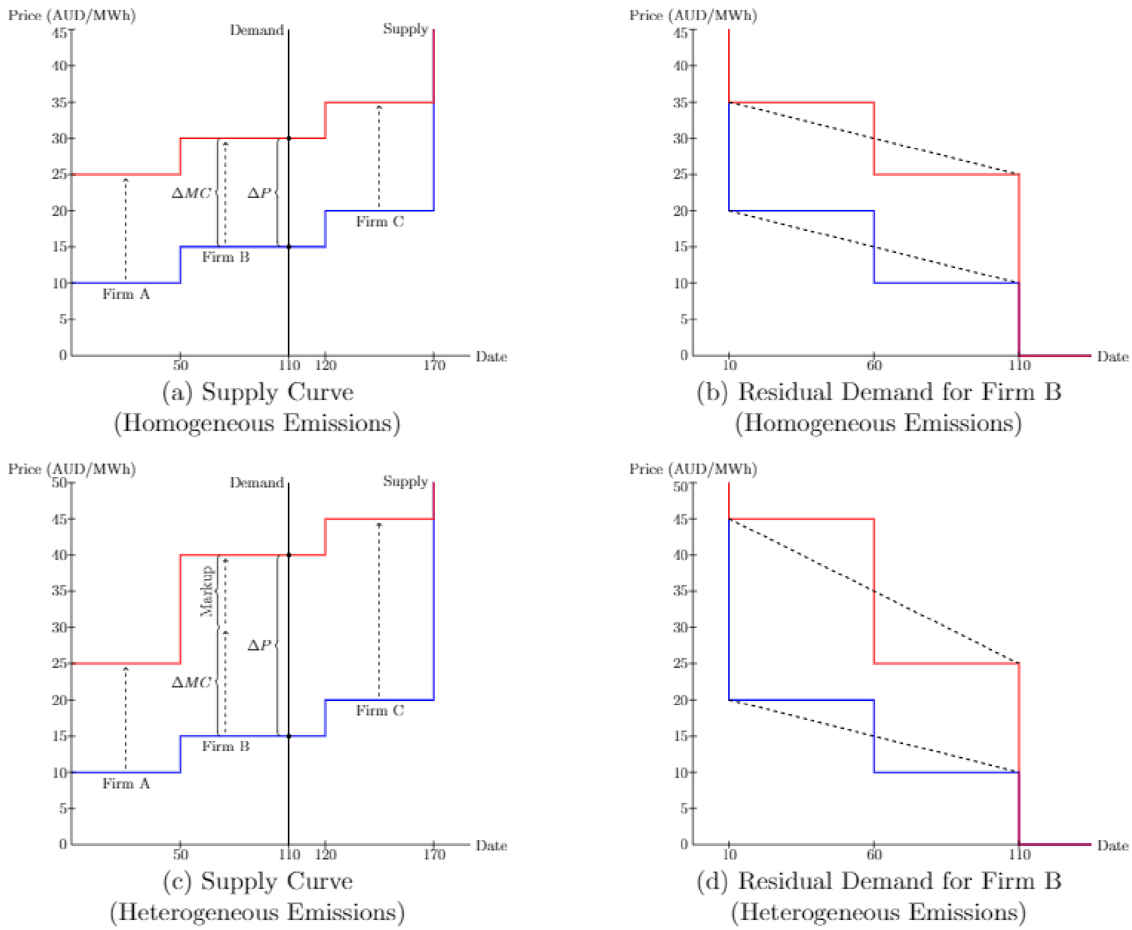

(a) Supply Curve (Homogeneous Emissions)

(b) Residual Demand for Firm B (Homogeneous Emissions)

(c) Supply Curve (Heterogeneous Emissions)

(d) Residual Demand for Firm B (Heterogeneous Emissions)

Note: This figure shows the change in the supply curve, equilibrium, and residual demand from a change in the carbon tax when all firms have the same emissions rates or different emissions rates. Shown in blue is the initial supply curves and residual demand curves for Firm B, which are the same in both the homogeneous and heterogeneous emissions cases. The primary difference between the homogeneous and heterogeneous cases shown is evident in the differential changes in the supply curve for Firms A and C shown in panels (a) and (c). I assume that Firms A and C increase their offer prices by the change in marginal costs. The impact that this has on the residual demand for Firm B is shown in panels (b) and (d). In the homogeneous case, the average slope when the other electricity generators are producing is shown as the dashed line. The slope does not change in panel (b) when firms have homogeneous emissions rates, while the residual demand curve becomes more inelastic in panel (d), highlighting the change in Firm B's ability to exert market power.

A.3 Preemptive Changes in Hydropower Electricity Generation

Despite the fact that any upward bias in estimates of changes in hydropower production due to the carbon tax is reflected in a similar bias for marginal plants, the bias of individual estimates can still have an effect on estimates of incidence. Equation 1.4 shows that the value of lost production, $\alpha_{het} = \sum_{i=1}^{n} \frac{dq_i}{d\tau} L_i \frac{P_m}{Q}$, depends on each firm's Lerner index. As one moves up the supply curve, the difference between the price and marginal costs decreases. Since hydropower is infra-marginal with lower marginal costs compared to marginal plants, the Lerner index will be higher for these firms. Intuitively, this means the upward bias will receive more weight in α_{het} compared to the downward bias - even if the total of the upward and downward biases is 0.

Returning to estimates of incidence, any upward bias in α_{het} will result in a smaller denominator, which in turn implies a larger estimate of incidence. In my robustness checks, I set $dq_i/d\tau = 0$ for all hydropower plants. No change in production for clean energy is the lower bound of what one would expect. Moreover, if hydropower electricity generators are operating witha binding intertemporal budget constraint on production, then this is what one would theoretically expect the impact of a non-anticipated carbon tax to have on hydropower production. Returning to the equation for α_{het}, this acts as a lower bound on estimates since now only the downward bias remains.

A.4 Additional Tables

Table A.1: Impact of the Carbon Tax on Daily Electricity Demand and Gas Prices

	Total Daily Demand		Gas Prices	
RD Estimate	768.15264	-4598.66050	-0.04486	-0.41833
	[6606.22302]	[6087.82407]	[0.29446]	[0.26272]
Polynomial	1	1	1	1
Bandwidth (days)	20	30	20	30

Notes: This Table shows the results from a regression discontinuity design with the cutoff date of July 1, 2013 where the carbon tax increased from 23 to 24.15. The coefficient and standard error in each column represents a separate regression with the polynomial and bandwidths specified at the bottom of the table. Coefficients represent the change in either total daily electricity demand (MWh) or gas prices. Gas prices here are ex-post short-term trading market prices in Sydney, Adelaide, and Brisbane.

Table A.2: Pass-through by Region with Raw Price Data

	Wholesale Electricity Price				
	(1)	(2)	(3)	(4)	(5)
NSW	0.666***	0.541**	0.613**	0.613**	0.618***
	[0.199]	[0.211]	[0.250]	[0.250]	[0.229]
Queensland	1.273***	1.148***	1.132***	1.133***	1.040***
	[0.243]	[0.225]	[0.281]	[0.281]	[0.229]
South Aus	1.076***	0.950***	1.015***	1.015***	0.961***
	[0.285]	[0.295]	[0.321]	[0.322]	[0.310]
Tasmania	0.630***	0.504***	0.633***	0.633***	0.622***
	[0.160]	[0.179]	[0.240]	[0.240]	[0.220]
Victoria	1.039***	0.913***	0.969***	0.969***	0.937***
	[0.189]	[0.202]	[0.231]	[0.231]	[0.228]
Observations	10,049	10,049	10,049	10,049	10,049
Clusters	66	66	66	66	66
Zone FE	Yes	Yes	Yes	No	No
Month FE	Yes	Yes	Yes	Yes	No
DOW FE	Yes	Yes	Yes	No	No
Year FE	No	Yes	Yes	Yes	Yes
Flexible Controls	No	No	Yes	Yes	Yes
Zone*Dow	No	No	No	Yes	Yes
Zone*Month	No	No	No	No	Yes

Notes: Each column represents a separate regression where the dependent variable is wholesale electricity prices and the key independent variable is the price on carbon emissions. All columns include the full set of controls: minimum and maximum temperature, rainfall, hydro water levels, and fossil-fuel prices. Every row represents the total pass-through rate for that region. Statistical significance is based on testing if the total effect with a null hypothesis of 0. Standard errors are in brackets and are clustered at the month of sample level.

Table A.3: Pass-through by Region (Median Regression)

	Wholesale Electricity price				
	(1)	(2)	(3)	(4)	(5)
NSW	0.872***	0.872***	0.859***	0.859***	0.809***
	[0.051]	[0.052]	[0.066]	[0.066]	[0.073]
Queensland	1.050***	1.050***	1.020***	1.020***	0.960***
	[0.055]	[0.057]	[0.079]	[0.079]	[0.071]
South Aus	0.975***	0.970***	0.957***	0.957***	0.922***
	[0.070]	[0.077]	[0.082]	[0.082]	[0.087]
Tasmania	0.452***	0.449***	0.446***	0.446***	0.417***
	[0.052]	[0.054]	[0.073]	[0.073]	[0.080]
Victoria	0.830***	0.830***	0.828***	0.828***	0.766***
	[0.054]	[0.056]	[0.066]	[0.066]	[0.076]
Observations	10,049	10,049	10,049	10,049	10,049
Clusters	66	66	66	66	66
Zone FE	Yes	Yes	Yes	No	No
Month FE	Yes	Yes	Yes	Yes	No
DOW FE	Yes	Yes	Yes	No	No
Year FE	No	Yes	Yes	Yes	Yes
Flexible Controls	No	No	Yes	Yes	Yes
Zone*Dow	No	No	No	Yes	Yes
Zone*Month	No	No	No	No	Yes

Notes: Each column represents a separate regression where the dependent variable is wholesale electricity prices and the key independent variable is the price on carbon emissions. All columns include the full set of controls: minimum and maximum temperature, rainfall, hydro water levels, and fossil-fuel prices. Every row represents the total pass-through rate for that region. Statistical significance is based on testing if the total effect with a null hypothesis of 0. Standard errors are in brackets and clustered at the month of sample level following Parente and Silva (2016). Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.4: Pass-through by Region (Outliers Dropped)

	Wholesale Electricity Price				
	(1)	(2)	(3)	(4)	(5)
NSW	0.839***	0.735***	0.737***	0.737***	0.741***
	[0.072]	[0.091]	[0.102]	[0.103]	[0.098]
Queensland	1.149**	1.044**	1.032**	1.032**	1.011***
	[0.122]	[0.118]	[0.117]	[0.117]	[0.108]
South Aus	1.168***	1.064***	1.067***	1.067***	1.053***
	[0.111]	[0.123]	[0.135]	[0.135]	[0.123]
Tasmania	0.446***	0.342***	0.355***	0.355***	0.354***
	[0.073]	[0.095]	[0.107]	[0.107]	[0.105]
Victoria	0.904***	0.800***	0.804***	0.804***	0.797***
	[0.074]	[0.096]	[0.102]	[0.102]	[0.101]
Observations	10,049	10,049	10,049	10,049	10,049
Clusters	66	66	66	66	66
Region FE	Yes	Yes	Yes	No	No
Month FE	Yes	Yes	Yes	Yes	No
DOW FE	Yes	Yes	Yes	No	No
Year FE	No	Yes	Yes	Yes	Yes
Flexible Controls	No	No	Yes	Yes	Yes
Zone*Dow	No	No	No	Yes	Yes
Zone*Month	No	No	No	No	Yes

Notes: Each column represents a separate regression where the dependent variable is wholesale electricity prices, dropping periods with prices larger than 1000 AUD/MWh, and the key independent variable is the price on carbon emissions. All columns include the full set of controls: minimum and maximum temperature, rainfall, hydro water levels, and fossil-fuel prices. Every row represents the total pass-through rate for that region. Statistical significance is based on testing if the total effect with a null hypothesis of 0. Standard errors are in brackets and are clustered at the month of sample level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.5: Markup Estimates: Augmented Local Linear RD

	Markup Estimates					
A. Preferred Specification						
RD Estimate	-0.21363	-0.18609	0.04996	0.08140	-0.17206	-0.43453
	[0.63636]	[0.58520]	[0.55132]	[0.54749]	[0.61961]	[0.60517]
B. Gen. Band FE						
RD Estimate	-0.50533	-0.67924	-0.37719	-0.00920	-0.33539	-0.88022
	[0.73174]	[0.64039]	[0.55299]	[0.88413]	[0.78622]	[0.76300]
C. Gen. Band by Time FE						
RD Estimate	-0.30235	-0.22068	-0.05469	-0.21925	-0.32796	-0.42220
	[0.66041]	[0.56178]	[0.50063]	[0.76258]	[0.69200]	[0.63382]
Polynomial	1	1	1	2	2	2
Bandwidth (days)	20	30	40	20	30	40

Notes: This Table shows the results from an augmented local linear regression discontinuity with the cutoff date of July 1, 2013 where the carbon tax increased from 23 to 24.15 Australian dollars. The coefficient and standard error in each cell of the table represents a separate regression, with the polynomial and bandwidths specified at the bottom of the table. My preferred specifications are shown in the second column of results which is a local linear regression which includes 30 days on each side of the cutoff. The augmented local linear regression estimates come from two stages. In the first stage, I regress the difference between frequently marginal offer prices and their marginal emissions costs on a rich set of controls. In my preferred specification, shown in panel *A*, I control for: coal and natural gas prices by generator and offer price, quadratic relationship between projected (by the Australian Energy Market Operator) demand and each generator's offer price, and lastly generator by day of the week and generator by year fixed effects. In the second stage, I estimate a regression discontinuity where the residuals from the first stage are the dependent variable. The results shown in panel B and C vary the regression estimated in the first stage. Panel B only includes generator band fixed effects, while panel C also includes generator band by year and generator by day of the week fixed effects.

Table A.6: Marginal Emissions Costs Pass-through Rate

	Price				
	(1)	(2)	(3)	(4)	(5)
MEC	1.091***	0.875***	0.801***	0.834***	0.708***
	[0.183]	[0.190]	[0.201]	[0.195]	[0.232]
Observations	10,005	10,005	10,005	10,005	10,005
Clusters	67	67	67	67	67
F-test	3958.2	2994.3	2730.2	2814.9	2904.3
Region FE	Yes	Yes	Yes	No	No
Month FE	Yes	Yes	Yes	Yes	Yes
DOW FE	Yes	Yes	Yes	No	No
Year FE	No	Yes	Yes	Yes	Yes
Flexible Controls	No	No	Yes	Yes	Yes
Region*Dow	No	No	No	Yes	Yes
Region*Month	No	No	No	No	Yes

Notes: Each column represents a separate regression where the dependent variable is wholesale electricity prices and the key independent variable is the marginal emissions costs of the marginal plant, which is instrumented by the carbon tax. When there are multiple marginal plants, I take the average. All columns include the full set of controls: maximum daily temperature, minimum daily temperature, rainfall, hydro water levels (previous day), and global fossil fuel prices. The F-tests indicate that the instrument has a strong effect on the endogenous variable, the marginal emissions cost. Standard errors are in brackets and are clustered at the month of sample level. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.7: Wholesale to Retail Pass-through: Average Across All Zones

	Retail Electricity Price		
	(1)	(2)	(3)
Wholesale Price	2.962***	1.141**	1.173***
	[0.001]	[0.015]	[0.004]
Observations	110	110	110
Clusters	22	22	22
F-Stat	31.91	49.74	37.339
Zone FE	Yes	Yes	Yes
Linear Time Trend	No	Yes	Yes
Year FE	No	No	Yes

Notes: Each column represents a separate regression where the dependent variable is the retail electricity price and the key independent variable is the wholesale electricity price, all in cents per kWh. Both electricity prices are raw averages for a given quarter of the year. The wholesale electricity price is instrumented by the carbon tax. This table shows the results of the average relationship between the wholesale and retail electricity price across all zones. All estimates represent the impact of a change in the wholesale electricity price on the retail electricity price. Standard errors are clustered at the month of sample level. Wild bootstrap with quarter of the year clusters (Roodman et al. 2019). P-values are reported inside square brackets.

Table A.8: Incidence Estimates Assuming Hydropower does not Change with the Tax

	(1)	(2)	(3)	(4)	(5)
Panel A. Incidence components					
Demand-weighted pass-through ($\bar{\rho}_{dw}$)	0.833	0.833	0.827	0.825	0.831
Demand-weighted retail pass-through ($\bar{\rho}_{ret}$)	1.114	1.114	1.114	1.114	1.114
Generation-weighted pass-through ($\bar{\rho}_{gw}$)	0.834	0.834	0.827	0.823	0.842
Emissions Intensity ($\bar{e}$)	0.885	0.885	0.885	0.885	0.885
Imp. Competition Parameter (α_{het})	0.006	0.018	0.018	0.018	0.018
Assuming $dq_i/d\tau = 0$ for Hydro (α_{het})	-0.052	-0.035	-0.035	-0.035	-0.035
Panel B. Consumer share of burden					
Oligopoly (Baseline)	0.949	0.962	0.954	0.950	0.971
Oligopoly (Assuming $dq_i/d\tau = 0$ for Hydro)	0.890	0.907	0.900	0.895	0.914
Panel C. Wholesale Pass-through Specification					
Base FE	Yes	Yes	No	No	No
Full FE	No	No	Yes	Yes	Yes
Outliers Dropped	Yes	Yes	Yes	No	No
Median Regression	No	No	No	Yes	No
Raw Prices	No	No	No	No	Yes
Panel D. Quantity regression specification					
Base FE	Yes	No	No	No	No
Full FE and Gen*Fuel	No	Yes	Yes	Yes	Yes

Notes: This table presents results for welfare incidence for electricity by component specification. Incidence is defined as the change in the consumer surplus as a share of the change in consumer and producer surplus. Panel A shows the estimates for each incidence component depending on the combination of regression specifications shown in Panels C and D. Panel B shows the estimate of incidence based on the components in Panel A. The imperfect competition parameter, α_{het}, represents the last term in the denominator of equation 1.4. In the baseline the imperfect competition parameter is the second to last row of Panel A, while in the counterfactual that assumes the change in production for hydropower in response to changes in the carbon tax is 0 is shown in the last row of Panel A. The first row of panel B shows the incidence estimate using empirically estimated changes in quantities. The second row uses the same estimates, but replaces all estimates of the change in quantities for hydropower with 0. Each column represents a separate specification used for an estimated parameter in the incidence formula. Base fixed effects are month, year, and day of week, while full fixed effects allows month and day of the week fixed effects to vary by location. For pass-through estimates this is a region, while for the quantity regressions this is the electricity generator. For the specification where outliers are dropped the cutoff is 1000 AUD/MWh, which represents less than 0.1% of prices.

A.5 Additional Figures

Figure A.2: Seasonality in Retail Prices

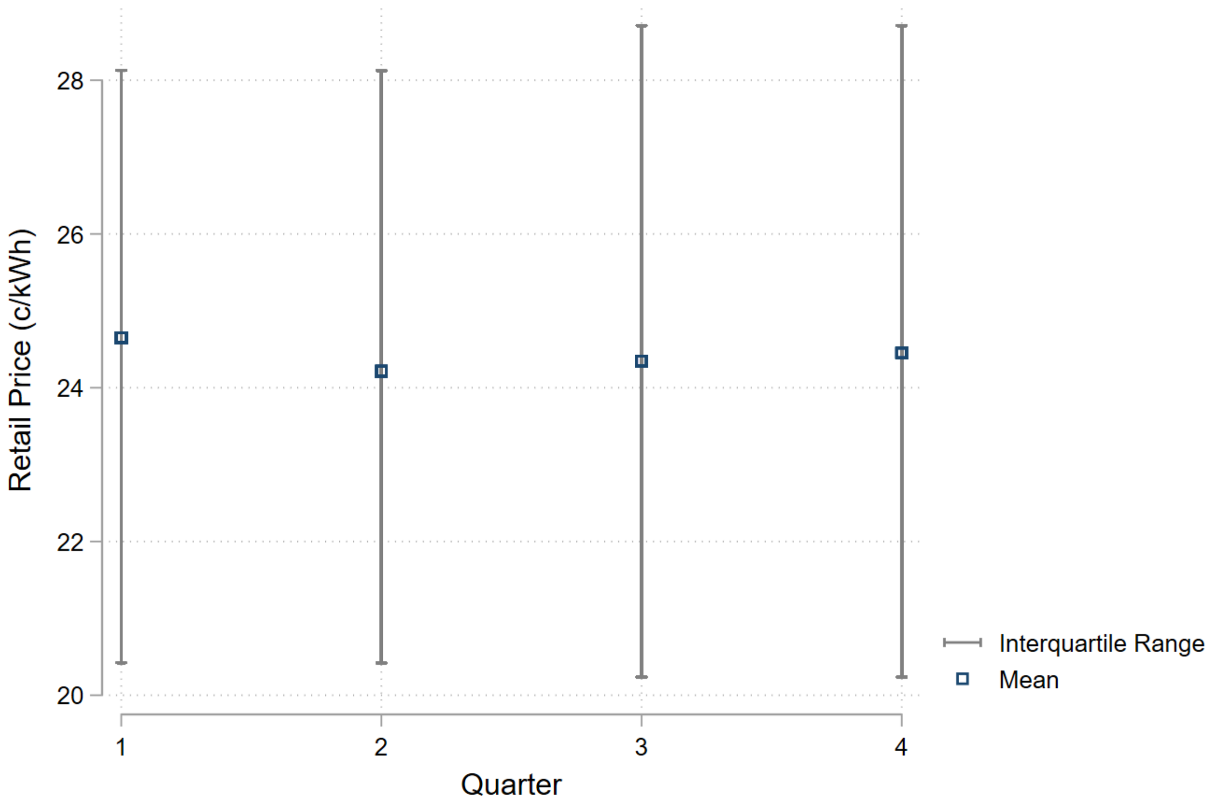

Note: This figure plots retail electricity prices by quarter of the year. The mean is shown as a point and the interquartile range is the bars. Despite seasonality in wholesale electricity prices this figure shows minimal seasonality in retail electricity prices, reflecting the fact that prices are often fixed for a period of time (e.g., annual).

Figure A.3: Daily Demand in June and July of 2013

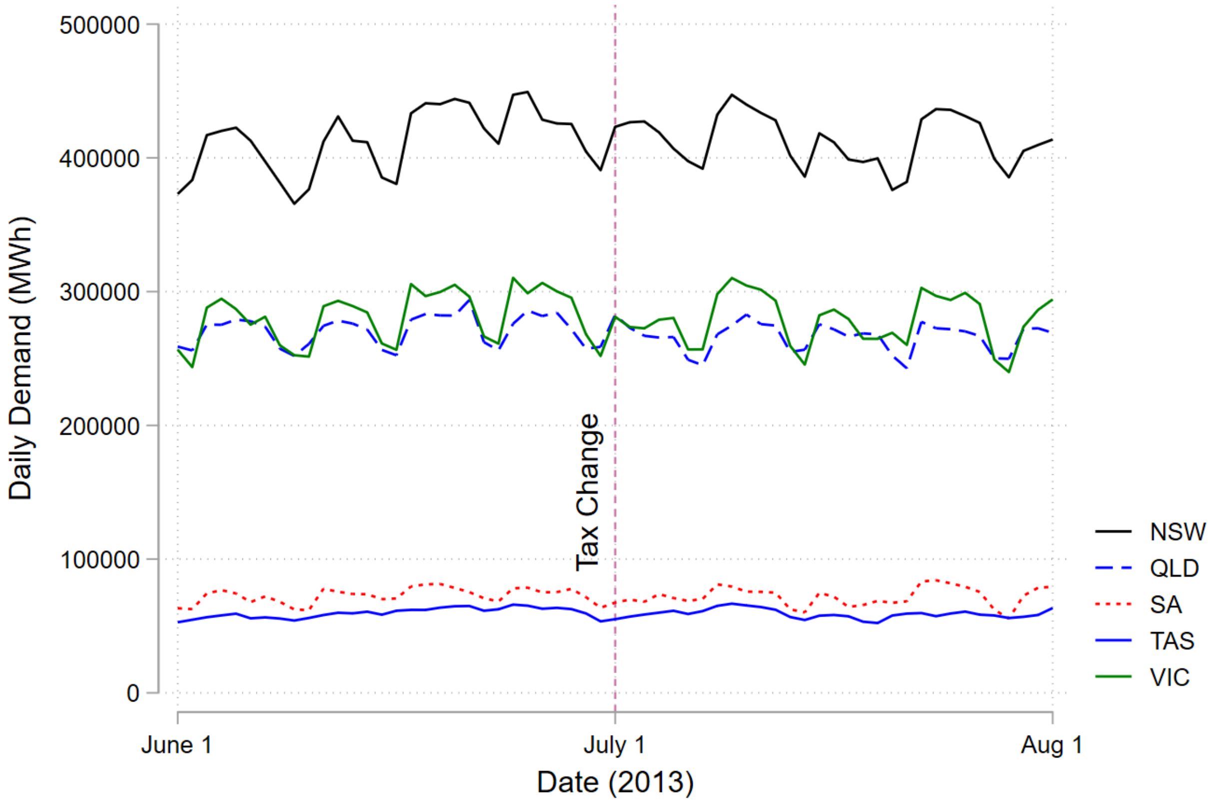

Note: This figure plots the total demand by day for each of the 5 electricity zones. Although there is no clear discontinuity on July 1, this figure does highlight the cyclical nature of electricity demand.

Figure A.4: STTM Natural Gas Prices in June and July of 2013

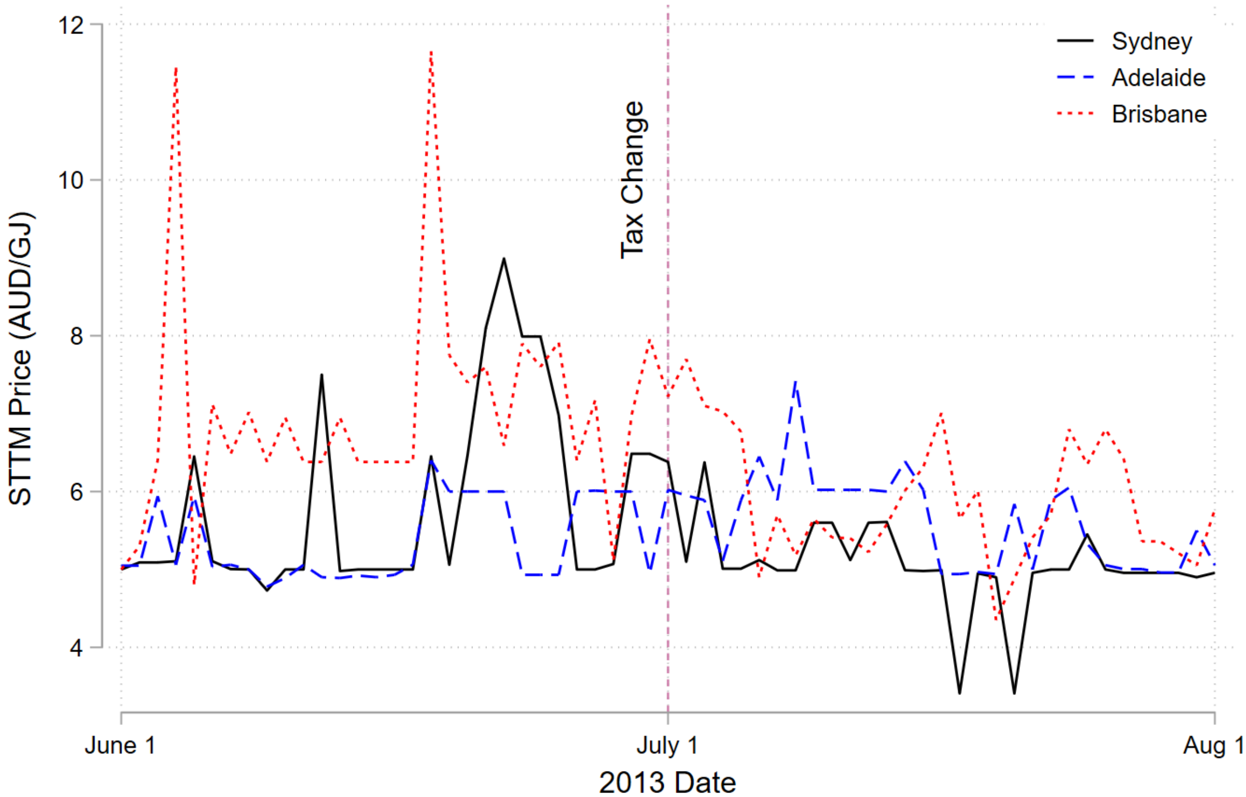

Note: This figure plots the ex-post daily short-term trading market natural gas prices for Sydney, Adelaid, and Brisbane using data from AEMO. All of the natural gas prices appear to be relatively stationary with little change in the natural gas prices from the change in the carbon tax.

Figure A.5: All Weather Stations

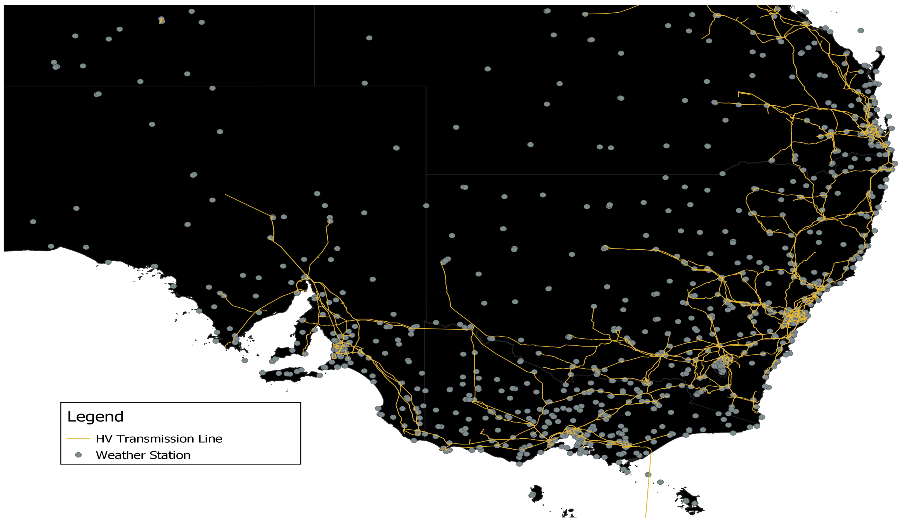

Figure A.6: Relevant Stations

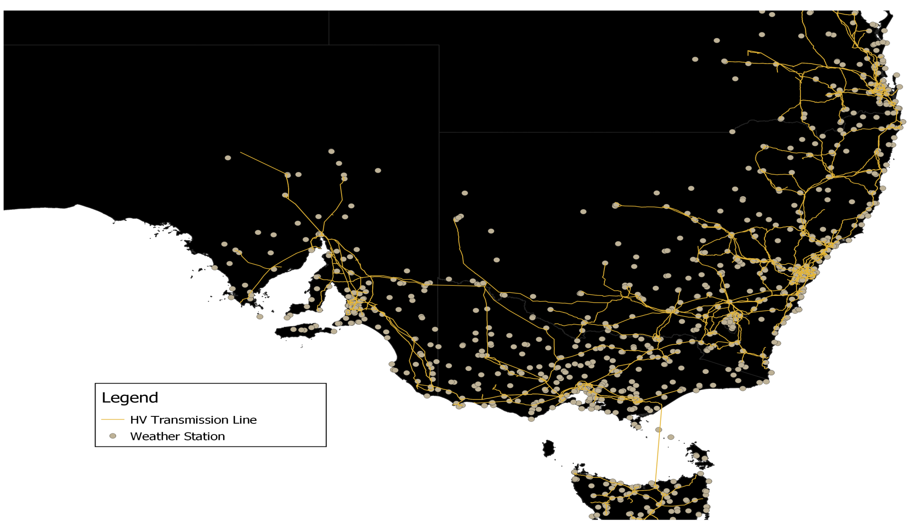

Note: This figure shows the location of weather stations and transmission lines for one part of the National Electricity Market in Australia after dropping weather stations more than 1 degree away from all transmission lines.

Figure A.7: Natural Gas Price History

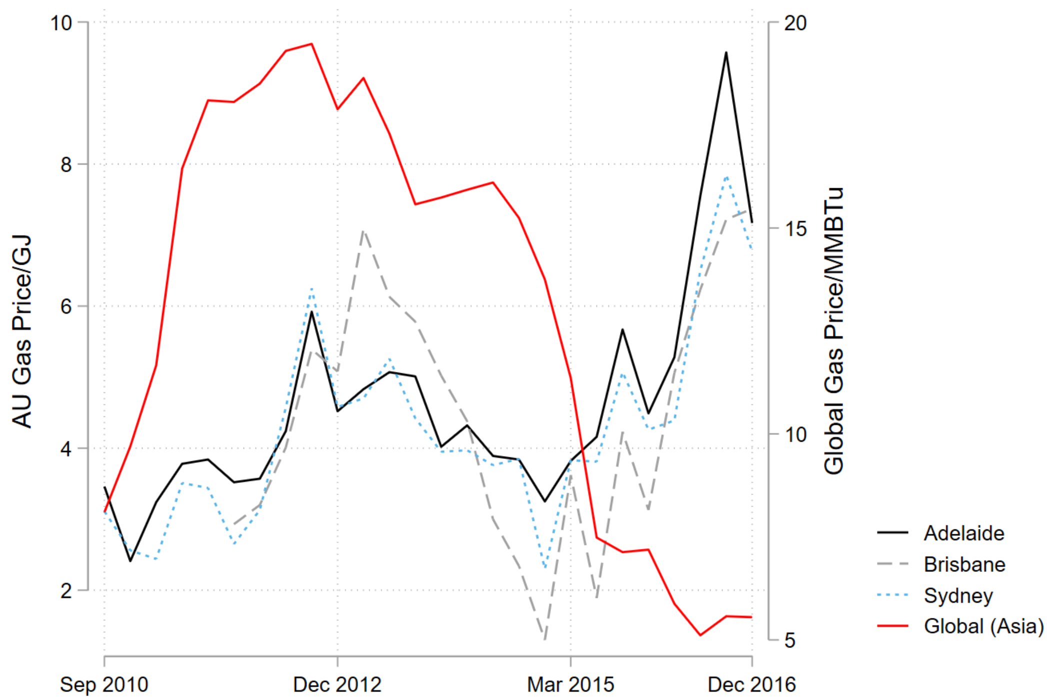

Note: This figure shows the natural gas price history over time. Starting in 2015 the two local natural gas prices in Australia are no longer correlated with the global natural gas price.

Figure A.8: Wholesale and Retail Prices over Time

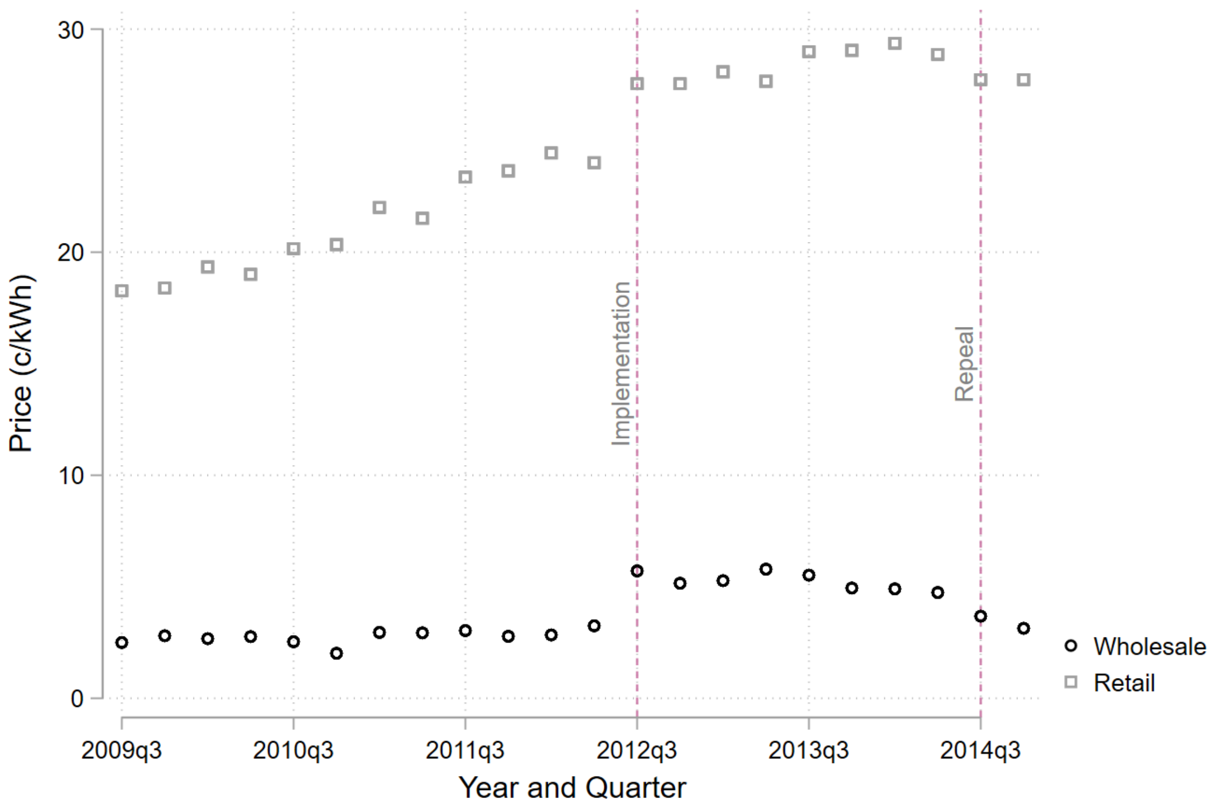

Note: This figure plots average wholesale and retail electricity prices in cents per kWh over time. The carbon price implementation coincides with the third quarter of 2012 and its repeal coincides with the third quarter of 2014. The effect of the carbon price is evident in the sharp increase in wholesale and retail electricity prices when it is implemented and the sharp decrease when it is repealed.

Figure A.9: Emissions Rates (tCO_2e/MWh) for Marginal Electricity Generation

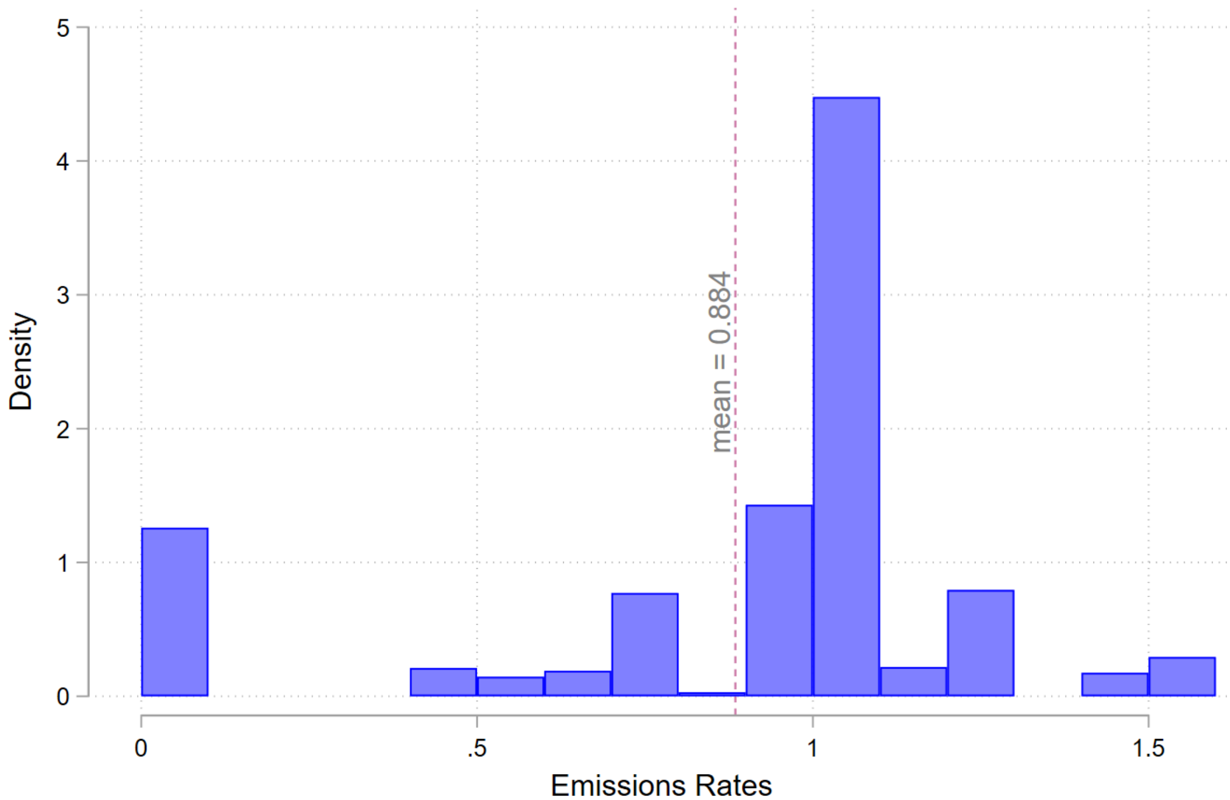

Note: This figure plots the distribution of emissions rates for electricity generated at the margin (price setters) throughout the National Electricity Market from 2012 to 2014. Approximately 60 percent of marginal plants use black coal, followed by 15 percent using brown coal, a little over 12 percent using hydro, and a little under 12 percent using natural gas.

Figure A.10: Emissions Rates (tCO_2e/MWh) for Marginal Electricity Generation by Region

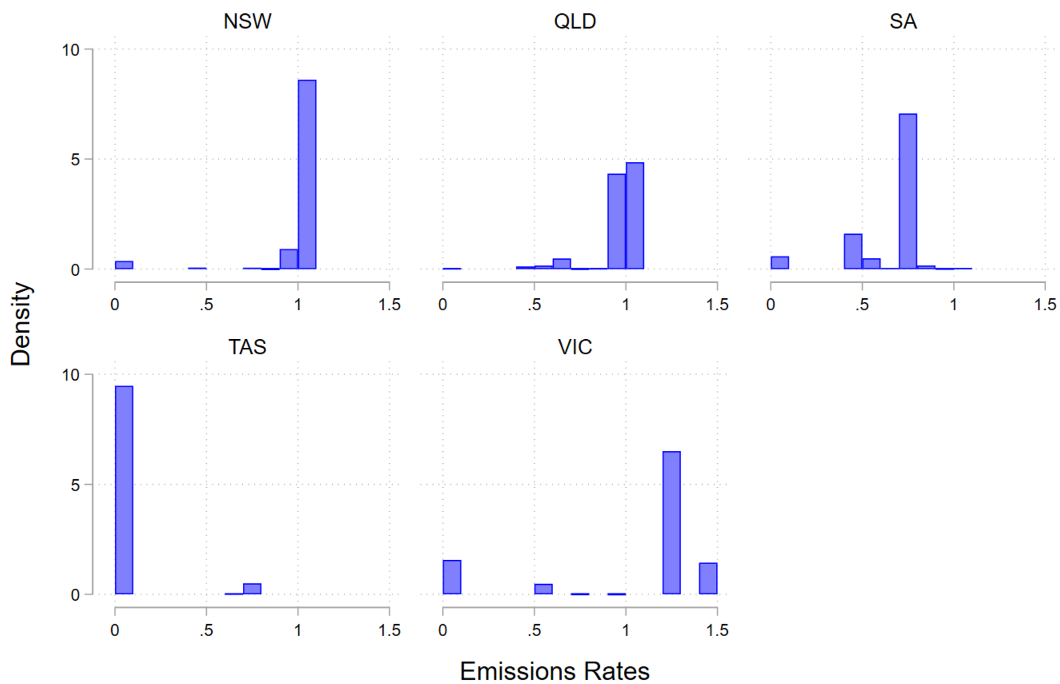

Note: This figure plots the distribution of emissions rates from 2012 to 2014 for electricity generated at the margin (price setters) broken down by region. Despite heterogeneity in the energy source in the aggregate, marginal electricity generators in each region use predominantly one fuel. In Victoria, 80 percent of the time the marginal plant used brown coal. In South Australia 94 percent of the time the marginal plant used natural gas. In both Queensland and New South Wales, over 90 percent of marginal plants used primarily black coal. In contrast, marginal plants in Tasmania used hydro almost 95 percent of the time.

Figure A.11: Plot of Lerner Index and Firm-Specific Changes in Quantity per dollar of the Carbon Tax

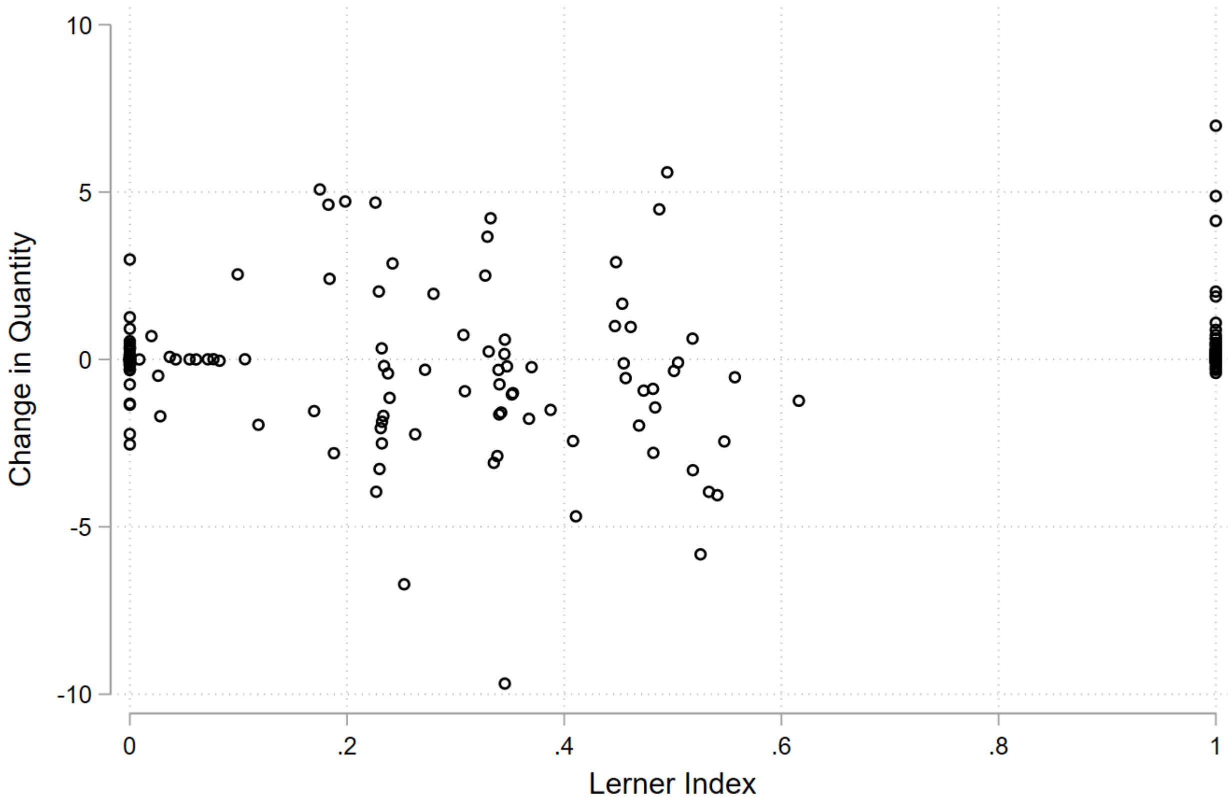

Note: This figure plots each firm's estimated Lerner index ($(P_m - MC_i)/P_m$) and changes in quantity per dollar of the tax ($dq_i/d\tau$). This figure uses the changes in quantity estimated with the most restrictive specification shown in Figure 1.6, specification 5. The correlation between the Lerner Index and changes in quantity is 0.0975.

Figure A.12: Percent Difference in Estimates of Consumer Burden: Constant Decrease for All Firms

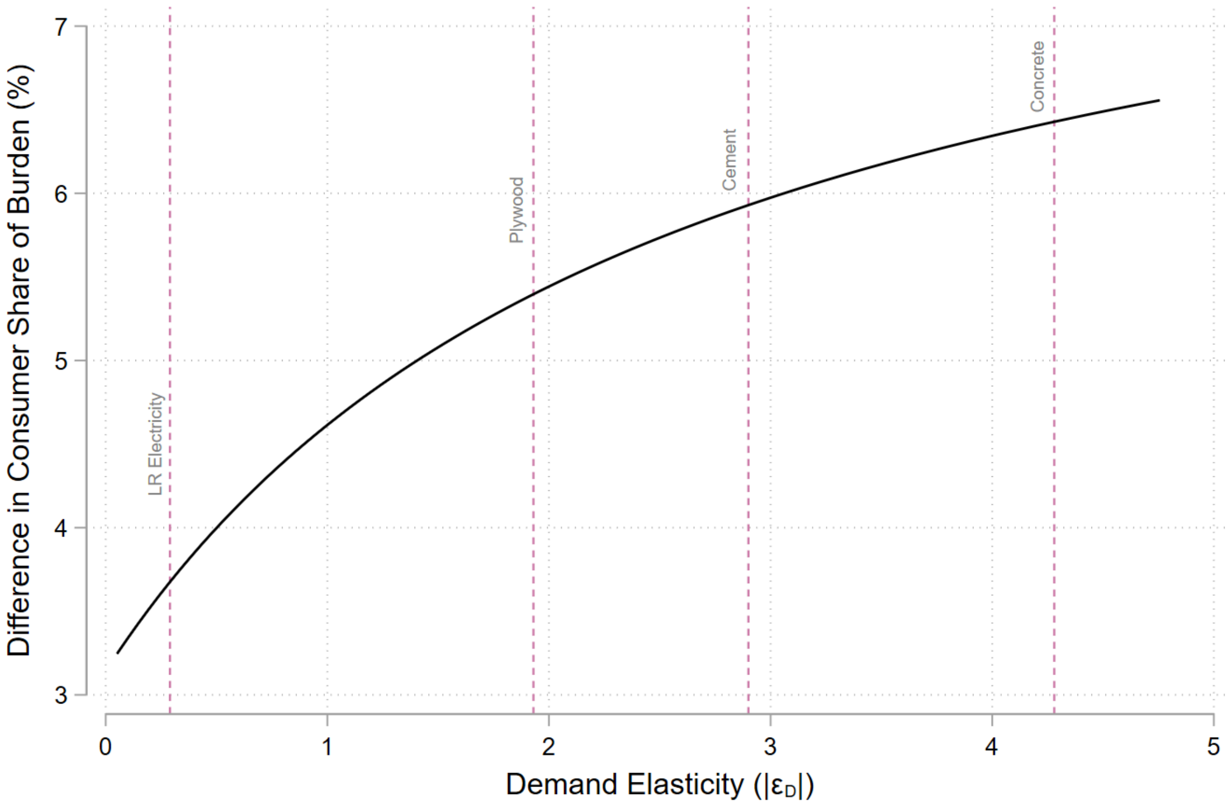

Note: This figure shows the percent difference between estimates of the consumer share of the burden when producers are assumed to be symmetric compared to when producers are allowed to be asymmetric. All estimates where the elasticity is larger than 0.03 are the result of simulations where I increase the demand elasticity from the baseline by simulating larger decrease in production in response to the carbon tax ($dQ/d\tau$). This is done by assuming that all electricity generators now experience the same constant decrease in production. In other words, I assume that $\frac{dq_i^{new}}{d\tau} = \frac{dq_i}{d\tau} - a$ where a is a different positive constant at each point along the x-axis. I then estimate the values of lost production when firms are assumed to be symmetric (α_{hom}) compared to asymmetric (α_{het}) and incidence in both cases. All estimates of the changes in quantities for each firm use as a baseline the preferred specification, Figure 1.6 specification 5, and the preferred pass-through estimates with the raw price data, Table 1.2 column 1. All estimates in the figure assume the same market conditions estimated in the context of the Australian electricity market, but with larger quantity changes in response to changes in the carbon tax. Gray lines act only as reference points for estimates of demand elasticity in various markets. The LR electricity demand elasticity is -0.36, which lies in the middle of estimates from Deryugina et al. (2020). Plywood, cement, and concrete demand elasticity estimates are from Ganapati et al. (2020).

Figure A.13: Percent Difference in Estimates of Consumer Burden: Larger Decrease for Firms that Decreased Production and are Likely Marginal

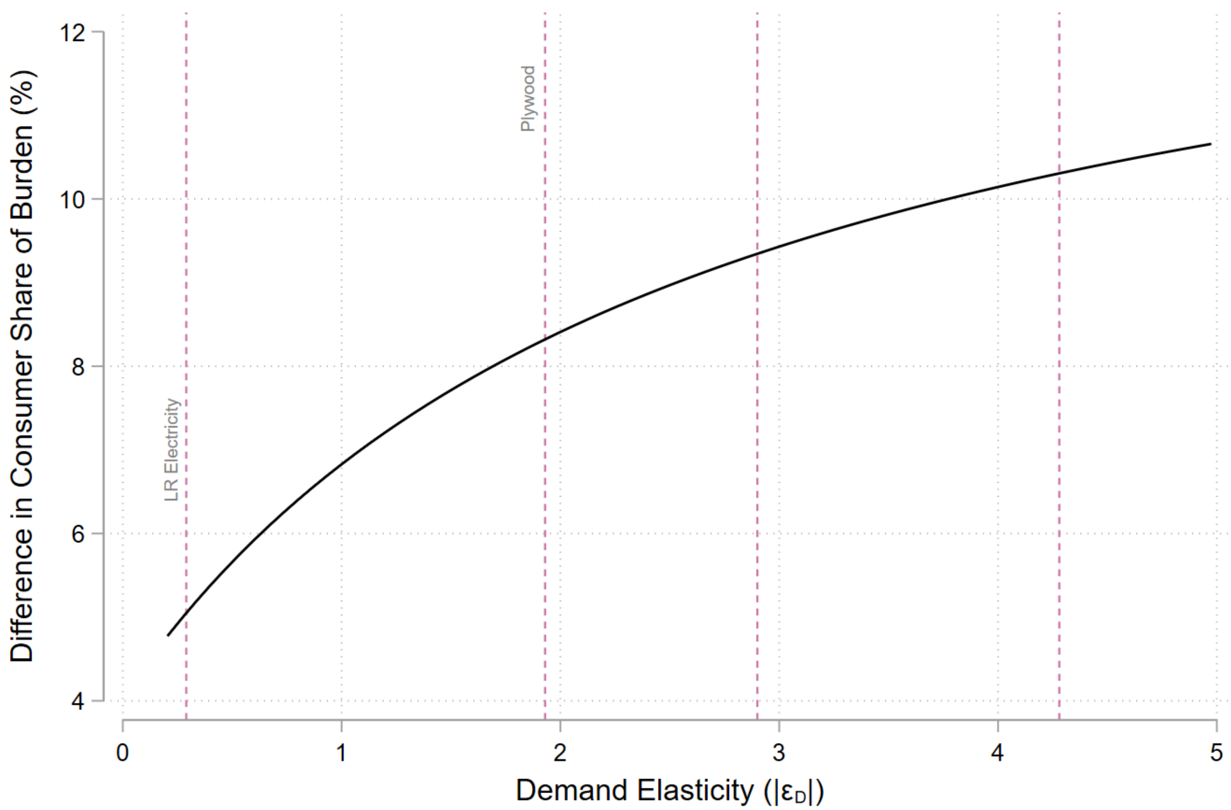

Note: This figure shows the percent difference between estimates of the consumer share of the burden when producers are assumed to be symmetric compared to when producers are allowed to be asymmetric. All estimates where the elasticity is larger than 0.03 are the result of simulations where I increase the demand elasticity from the baseline by simulating larger decrease in production in response to the carbon tax ($dQ/d\tau$). This is done by assuming that all electricity generators now experience the same constant decrease in production. In other words, I assume that $\frac{dq_i^{new}}{d\tau} = (c - L_i) \cdot \frac{dq_i}{d\tau}$ where c is a different constant greater than 2 at each point along the x-axis. I then estimate the values of lost production when firms are assumed to be symmetric (α_{hom}) compared to asymmetric (α_{het}) and incidence in both cases. All estimates of the changes in quantities for each firm use as a baseline the preferred specification, Figure 1.6 specification 5, and the preferred pass-through estimates with the raw price data, Table 1.2 column 1. All estimates in the figure assume the same market conditions estimated in the context of the Australian electricity market, but with larger quantity changes in response to changes in the carbon tax. Gray lines act only as reference points for estimates of demand elasticity in various markets. The LR electricity demand elasticity is -0.36, which lies in the middle of estimates from Deryugina et al. (2020). Plywood, cement, and concrete demand elasticity estimates are from Ganapati et al. (2020).

Figure A.14: Percent Difference in Estimates of Consumer Burden: Larger Decrease for Firms that are Likely Marginal

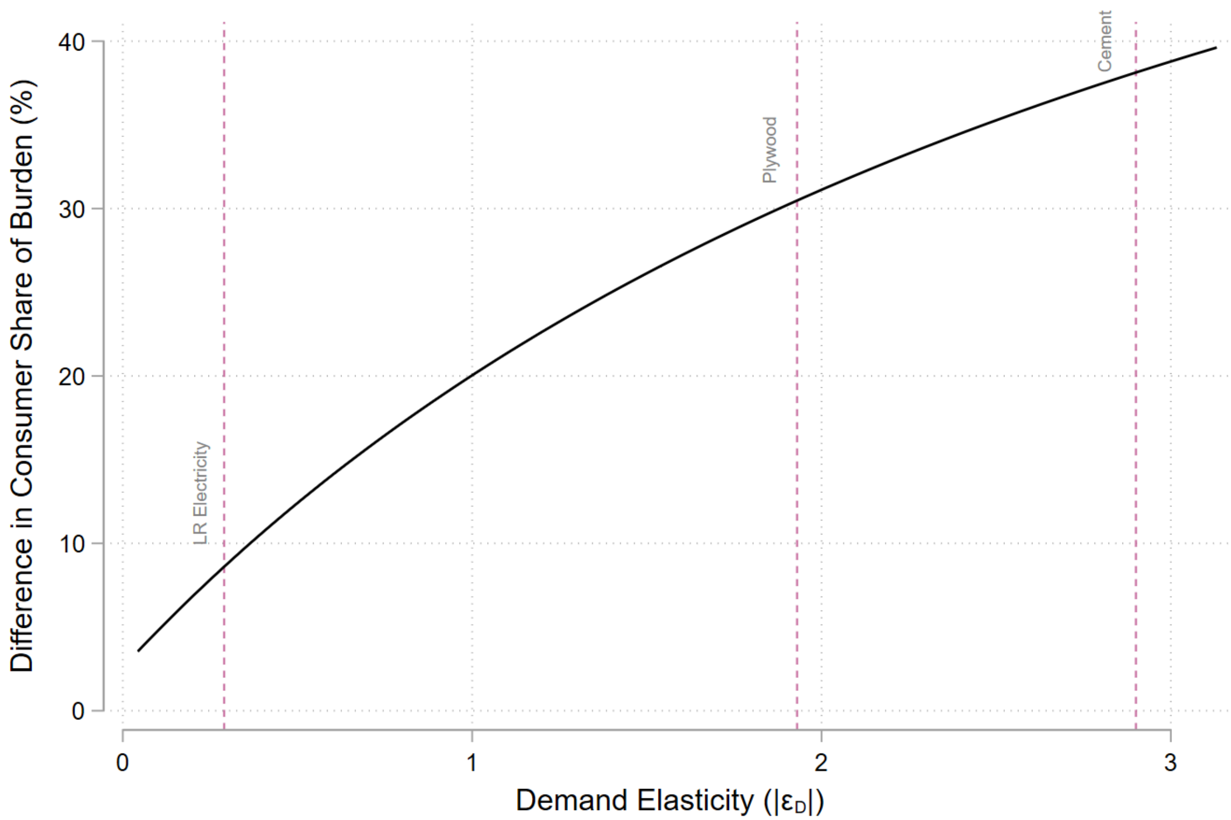

Note: This figure shows the percent difference between estimates of the consumer share of the burden when producers are assumed to be symmetric compared to when producers are allowed to be asymmetric. All estimates where the elasticity is larger than 0.03 are the result of simulations where I increase the demand elasticity from the baseline by simulating larger decrease in production in response to the carbon tax ($dQ/d\tau$). This is done by assuming that all electricity generators now experience the same constant decrease in production. In other words, I assume that $\frac{dq_i^{new}}{d\tau} = \frac{dq_i}{d\tau} - a - (1 - L_i)$ where a is a different positive constant at each point along the x-axis. I then estimate the values of lost production when firms are assumed to be symmetric (α_{hom}) compared to asymmetric (α_{het}) and incidence in both cases. All estimates of the changes in quantities for each firm use as a baseline the preferred specification, Figure 1.6 specification 5, and the preferred pass-through estimates with the raw price data, Table 1.2 column 1. All estimates in the figure assume the same market conditions estimated in the context of the Australian electricity market, but with larger quantity changes in response to changes in the carbon tax. Gray lines act only as reference points for estimates of demand elasticity in various markets. The LR electricity demand elasticity is -0.36, which lies in the middle of estimates from Deryugina et al. (2020). Plywood, cement, and concrete demand elasticity estimates are from Ganapati et al. (2020).

Appendix B

Appendix: Paying for Integers

B.1 Data Refinement Procedure

The refinement procedure follows Haggag and Paci (2014).

1. Drew 2,000 random taxi driver and car pairs for each year
2. Dropped duplicate observations.
3. Drop-off time occurs before pick-up time.
4. Drop-off time occurs after subsequent trip pick-up time.
5. Ride duration was zero or longer than 3 hours.
6. Trip distance was zero or greater than 100 miles.
7. Surcharge amount was greater than \$1.00.
8. Fare was less than \$2.50 or negative fare amounts.
9. MTA tax was larger than \$0.50.

10. Driver drove fewer than 100 rides for a given year.

11. Multiple cars were associated with the same driver during the same shift.

12. Driver's shift was longer than 20 hours.

13. Driver's shift was shorter than 30 minutes.

14. Either the pickup or drop-off location could not be mapped to census tract in New York, New Jersey, Connecticut or Pennsylvania

15. Dropped fares were categorized as "Dispute" or "No Charge"

16. Switched variable names between "Tip Amount" and "Tolls Amount" for Dec2011 fare.

17. Dropped rides with cash transactions.[1]

[1]See Table B1 for a comparison between cash and credit transactions.

B.2 Simulating Impact of Integer Default Tip Suggestions

The impact of integer default tip suggestions on the utility from the menu option is relatively straightforward. A small change in the fare that leads to an integer tip suggestion can impact the utility of this option depending on the value of b_i, as is shown in Figure B1. When $b_i = 0$ there is no change in the utility of the default tip option based on whether or not the tip suggestion, $t_i^D F_i$, is an integer. However, when $b_i > 0$ a passenger's utility from the default option exhibits discontinuously higher utility when the tip suggestion is an integer.

Figure B1: Individual's Utility from Default Tip Suggestion in Response to Fare amount, by Different b_i

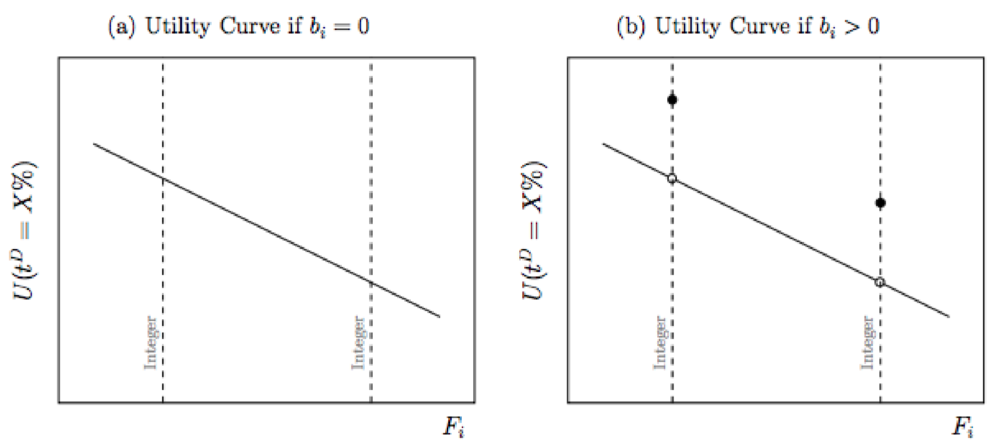

Notes: Figures presents the relationship between utility of taking default tip suggestion and fare amounts under the extended model by different b_i, for a given default tip rate $t_i^D = X\%$. Panel (a) presents the relationship when we set $b_i = 0$. Under this case, increases in F_i *smoothly* decreases one's utility. Panel (b) presents the same relationship when we set $b_i > 0$. Under this case, we observe discontinuous sharp increases in $U(t_i^D)$ when $t_i^D F_i \in \mathbb{Z}$.

The impact of integer default tip suggestions on the utility from custom tips is less clear as the preferred custom tip depends on comparing the tip rate that satisfies equation (2.2) with alternative tip rates that lead to integer tips, as shown in the right panel from Figure 2.4. Although it is unlikely, one could imagine that tip rates that satisfy equation (2.2) tend to lead to integer tip suggestions when default tip rates are integers. To explore this, we parameterize the utility function and plot the utility of the preferred custom and default tips based on distance to the integer tip suggestion.

The primary piece of customer's utility that we need to put structure to in order to simulate utility is the norm-deviation cost $v(T_i, t_i)$. Following Donkor (2020), we define the norm-deviation cost as $\theta(T_i - t_i)^2$. We can then write a generic passenger's utility function as:

$$Max_{t_i} U = -t_i F_i - \underbrace{\theta(T_i - t_i)^2}_{v(T_i,t_i)} - \mathbb{1}\{t_i \neq t_i^D\}[c_i^{non} - \alpha_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}] + b_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\} \quad \text{(B1)}$$

where θ "scales" the impact of deviating from what the passenger perceives as the socially accepted tip.

We are primarily interested in investigating whether, under reasonable parameters, utility from custom tips exhibit a discontinuity when default tip suggestions are integers. For this exercise, we will thus make the following parameter assumptions:

- $T_i = 0.15$ or $T_i = 0.18$
- $t_i^D = 0.2$
- $\theta = 1000$
- $c_i^{non} = 0.6$
- $\alpha = 0.1$
- $b_i = 0.1$

For fares ranging from 0 to 100, we then calculate the utility for the tip rate that satisfies equation (2.2) and the closest tip rates that lead to integer tips. We then calculate $U(t_i^C)$

as the custom tip, integer or not, that gives the highest utility to the passenger for that fare. Alternatively, for default tips, we calculate the utility for a single default tip rate of 0.2 for all the fares from 0 to 100.

Given the default tip rate of 0.2, integer default tip suggestions will occur at fares of 5, 10, 15, etc. To highlight any discontinuities in utility around these values, we calculate the average utility for default and custom tip rates at values around the integer default tip suggestion. Specifically, we calculate the distance between the fare and the closest fare that leads to an integer default tip suggestion. In practice, this means that fares of 4.5 and 9.5 would be treated similarly since their decimal distance is -0.5 (-50 cents), while fares of 5.5 and 10.5 would have decimal distance equal +0.5. We then calculate the average default tip option and custom tip option utility based on the decimal distance. If there is a discontinuity, on average, then this would be shown in a spike at the value of 0. Figure B2 shows that this is evident for default tip suggestions, but not custom tips. Importantly, the lack of a discontinuity for custom tip rates does not appear to be a result of the choice of T_i as the results are robust to alternative T_i besides those shown here. In addition, in all alternative specifications for the other parameters (θ, c_i^{non}, α, and b_i) that we have simulated, the conclusions are similar although the utility levels and magnitudes of the spikes for the default option can vary.

In summary, the simulation shown in Figure B2 highlights that custom tip utility appears to be continuous when presented with default tip suggestions. Intuitively, this is because the primary concern was that custom tip rates that satisfy equation (2.2) lead to integer tips more frequently when the default tip suggestion is also an integer. There is no reason ex-ante to think that this would be the case, which is supported by Figure B2.[2]

[2]Intuitively, one could think that customer's prefer a tip rate of 0.1, which would also frequently have integer tip suggestions when $t_i^D = 0.2$. Our theory, however, would suggest that even if passenger's believe the socially accepted tip rate is 0.1, they would "shade downwards" their preferred custom tip.

Figure B2: Custom and Default Tip Utility by Distance to Integer Default Tip Suggestion

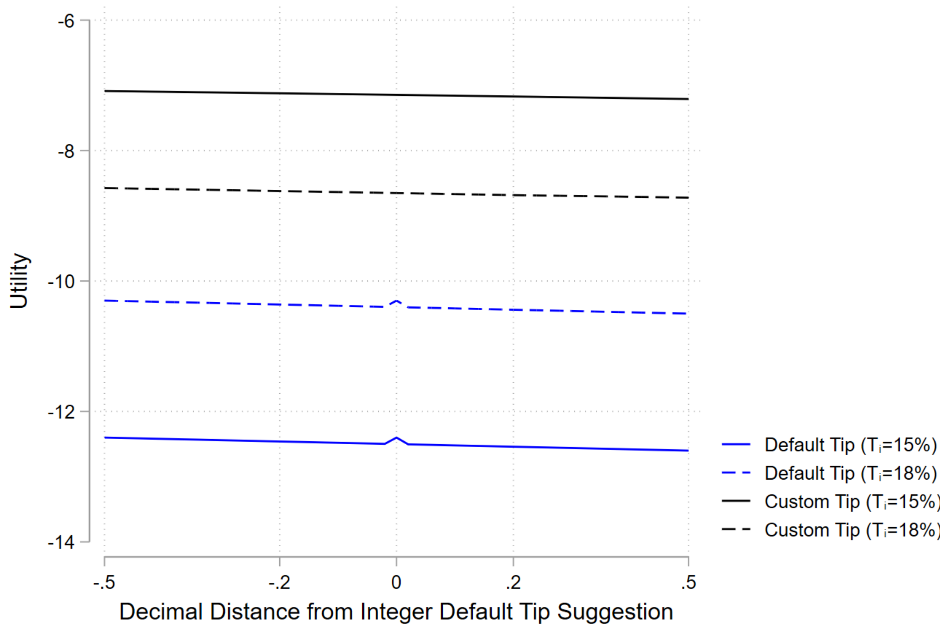

Notes: This figure plots the utility of choosing a custom tip compared to a default tip option based on the distance of the fare from the closest fare that leads to a default tip suggestion that is an integer. The range of fares used to create this figure is from 0 to 100. For the default tip rate of 0.2 used here, this means that the utility shown at 0 corresponds to the average utility at fares of 5, 10, 15, etc, while -0.5 represents 4.5, 9.5, 14.5, etc. The utility function used for this figure is:

$$U = -t_i F_i - \underbrace{\theta(T_i - t_i)^2}_{v(T_i, t_i)} - \mathbb{1}\{t_i \neq t_i^D\}[c_i^{non} - \alpha_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}] + b_i \cdot \mathbb{1}\{t_i F_i \in \mathbb{Z}\}$$

where we set $\theta = 1000$, $c_i^{non} = 0.6$, $\alpha = 0.1$, and $b_i = 0.1$. Solid lines show when $T_i = 0.15$ and dashed lines show when we set $T_i = 0.18$. To calculate the default tip utility for each fare, we change F_i leaving all else constant. To calculate the custom tip utility for each fare, we change F_i and find the custom tip rate that maximizes utility, ignoring the default option, at that point.

B.3 Additional Figures and Tables

Figure B1: Passenger Display for CMT in 2012

Notes: This figure shows the screen for a CMT outfitted vehicle in 2012. The source is the online appendix to Haggag and Paci (2014), Figure A.1, which was a photo taken by the authors.

Figure B2: Distribution of Tip Suggestion: Feb – Aug 2012

(a) Low Suggestion (20%)

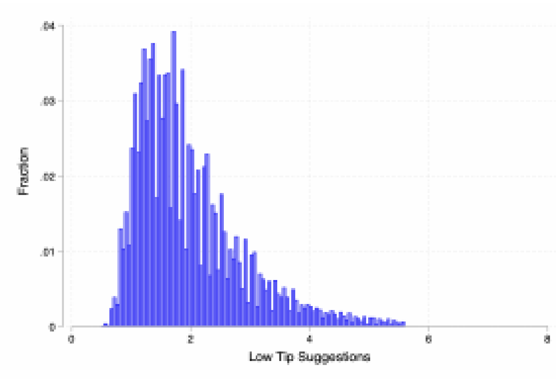

(b) Middle Suggestion (25%)

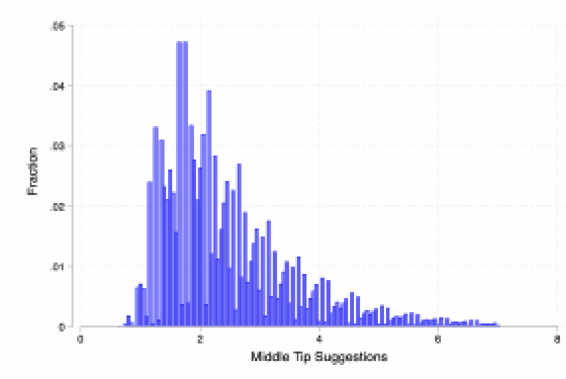

(c) High Suggestion (30%)

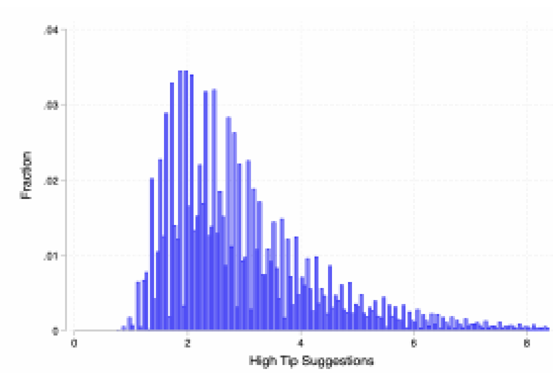

Notes: Panels (a) (b) and (c) shows the distributions of tip suggestions for the low, middle and high options. Extreme tip suggestion ($> 99^{th}$ percentile) are excluded from the figure. During Feb–Aug 2012, the % tip suggestion options (20-25-30) were identical to CMT and VTS taxis.

Figure B3: Distribution of Custom Tip Amount: Feb – Aug 2012

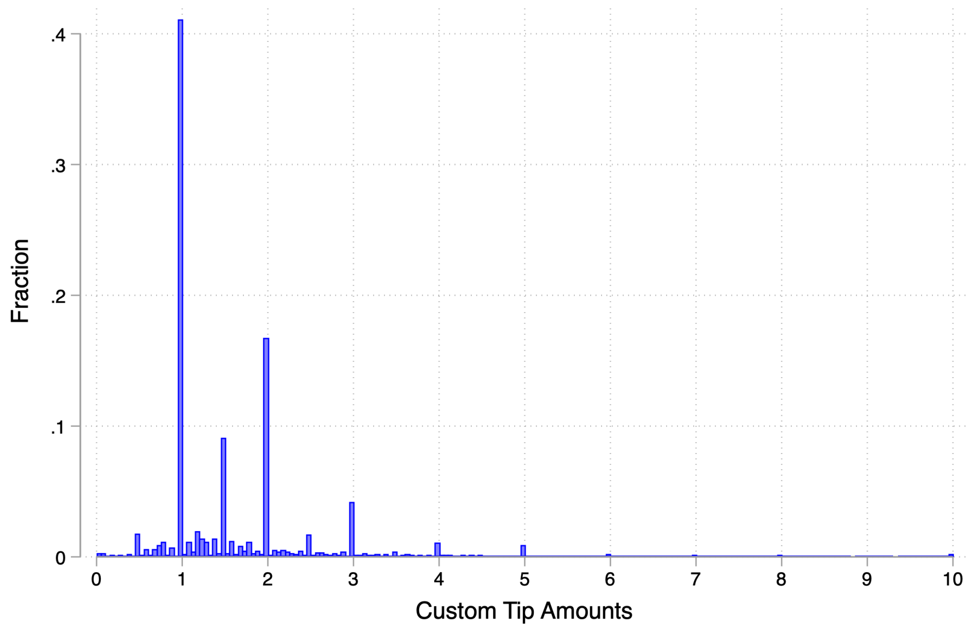

Notes: This figure shows the distribution of Custom tip amounts for all non-airport trips that were paid by credit card. Custom tips includes all non-zero tips that are not equal to any of the tip suggestions. Extreme custom tip amounts ($> 99^{th}$ percentile) are excluded from the figure. All tip amounts are in nominal dollar value.

Figure B4: The Decision Process for Agents under the Extended Model

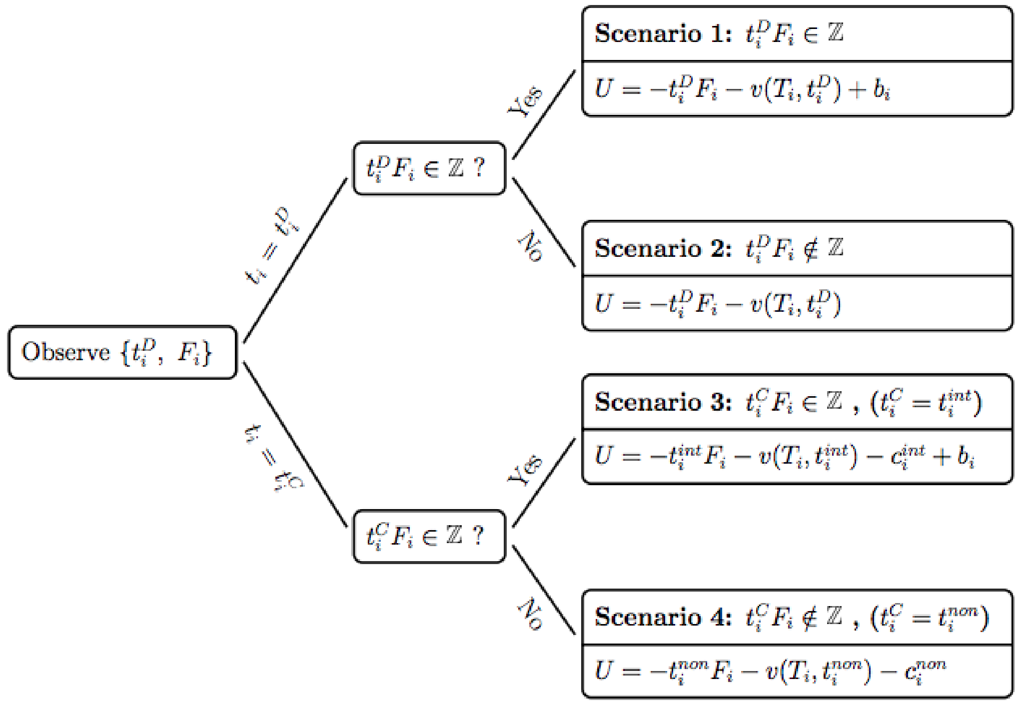

Notes: This figure presents the decision process that is described by the extended model (equation. 2.4). After observing t_i^D and F_i, a customer decide whether or not to the choose custom tipping option. If she chooses to enter a custom tip, she will incur a cognitive cost of either c_i^{non} or c_i^{int} depending on whether the chosen tip amount is an integer. We assume the cognitive cost is lower if she chooses to tip an integer amount $c_i^{int} < c_i^{non}$. In addition, she experiences utility gains (b_i) from choosing an integer tip, i.e, $t_i \in Z, t_i \in \{t_i^C, t_i^D\}$.

Figure B5: Placebo Effects

(a) Default Take-up

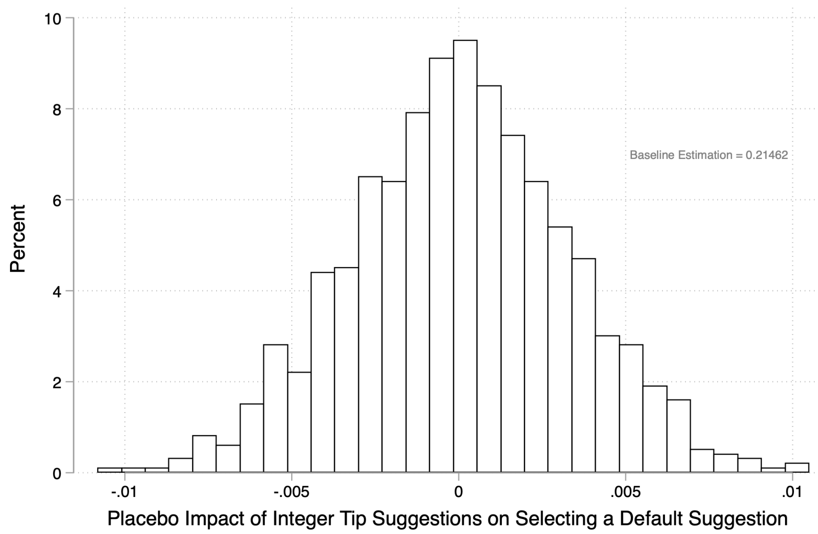

(b) Tip Rate

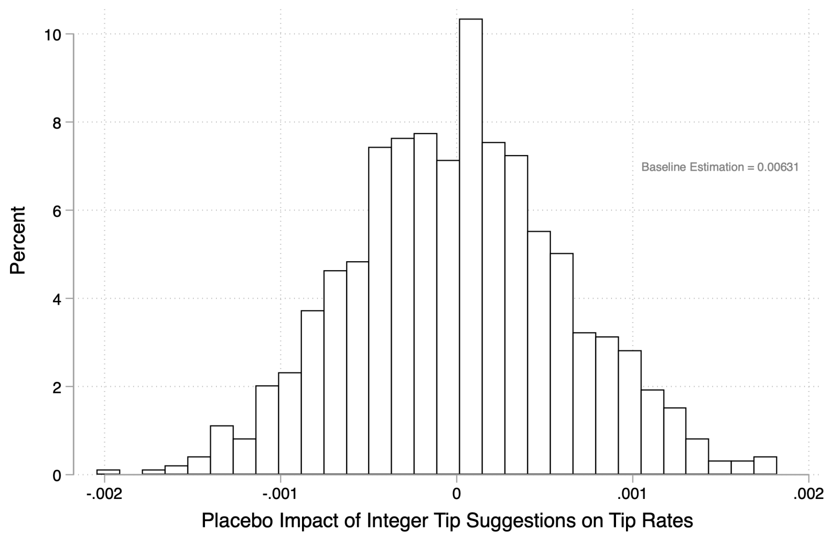

Notes: Figures shows the empirical distribution of estimated placebo treatment effects from 1,000 random treatment (trip with integer tip suggestion) assignments. The actual treatment effects are estimated from Table 2.2 Column (3) and Table 2.3 Column (3). p-values under the placebo tests are both < 0.001.

Figure B6: Impact of the VTS Menu Change on Tip Rates: RD in Time

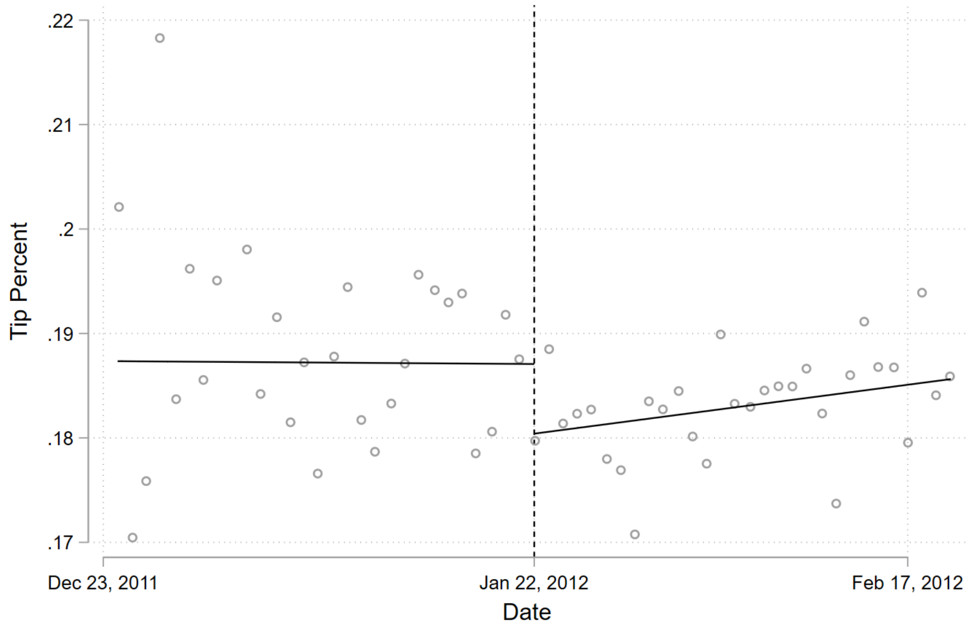

Notes: This figure presents the plot of regression discontinuity in time. Prior to the event, the tip suggestion menu offered by VTS had 2, 3, and 4 dollars suggestions for rate fare (fare + surcharge) below \$15. After the week of January 22, 2012, VTS removed the dollar tip suggestion and replaced it with the 20, 25, 30 percent tip suggestions.

Figure B7: Effect of VTS Menu Change in 2012

(a) Selecting Options from the Menu

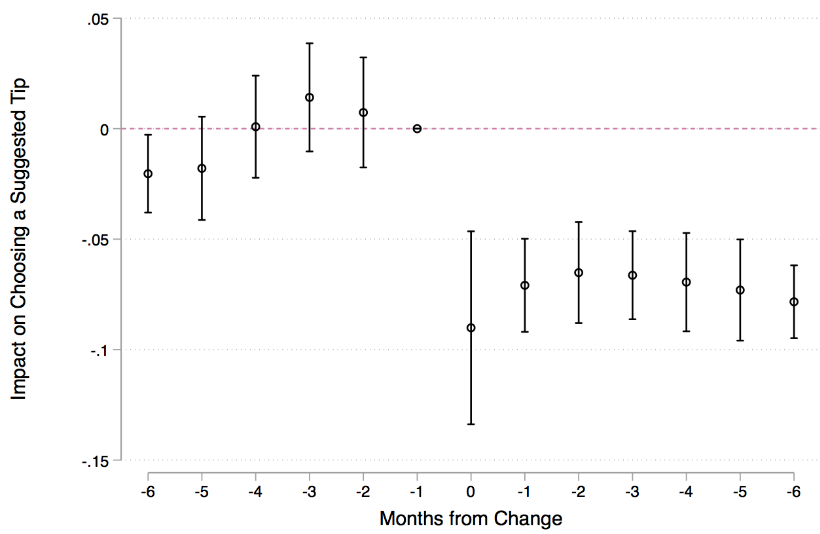

(b) Tip Rate

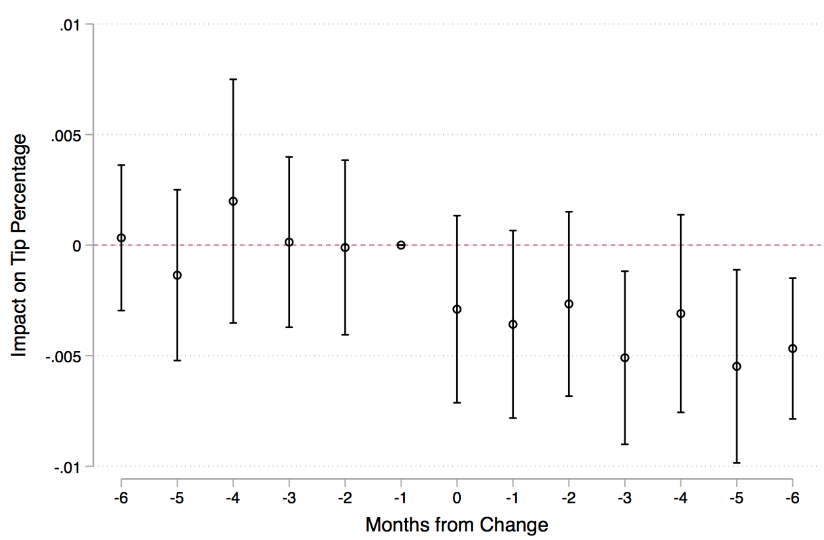

Notes: Figures show the event study plots for VTS menu change in Jan.2012. We control for tip or fare policy variations, and we include pick-up date fixed effects and vendor fixed effects. In addition, we include $x(d, mph)$ control. We include samples with non-zero tips. We cluster the standard errors at the pick-up date level.

Table B1: Summary of Cash and Credit Differences: Feb-Aug 2012

	(1) Cash	(2) Credit	(3) Difference
Fare Amount	8.58	9.48	-0.90***
	(4.76)	(4.93)	(0.01)
Trip Length (in minutes)	10.78	12.10	-1.33***
	(7.25)	(7.48)	(0.01)
Trip Distance (in miles)	2.20	2.54	-0.34***
	(2.03)	(2.11)	(0.00)
Fraction VTS	0.50	0.50	-0.00***
	(0.50)	(0.50)	(0.00)
Pickup Location Median Income	95,948.18	95,919.98	28.20
	(38,495.11)	(36,906.74)	(62.21)
Fraction Low Option Integer	0.03	0.02	0.01***
	(0.17)	(0.16)	(0.00)
Fraction Mid or High Option Integer	0.01	0.01	-0.00
	(0.10)	(0.10)	(0.00)
Observations	785,300	710,059	1,495,359

Notes: This table presents the summary statistics for the random sample of 2,000 taxi drivers during the time period of our main study: February to August 2012. Standard deviations are in parenthesis.

Appendix C

Appendix: Asymmetric Pass-through from Carbon Taxes to Electricity Prices

C.1 Mineral Resources Tax and Federal Income Tax Cuts

The implementation date of the carbon price was shared with implementing two other policies: mineral resources tax and federal income tax cuts. Using the repeal of the mineral resources tax, I find that this had no effect on wholesale electricity prices. At first it is puzzling that the mineral resources tax has minimal impact on electricity prices. It was a 30 percent tax rate on profits over 75 million from the mining of non-renewable resources. Intuitively this seems like it should have some impact on the wholesale electricity market. The lack of any effect, however, makes sense when one takes into account the fact that the tax was incredibly ineffective. In the first year the mineral resources tax raised less than 7 percent of its projected 3 billion dollars. The lack of response in the wholesale electricity market is thus not surprising since the burden for mining companies was small.

The second policy change was cuts to federal income taxes. Specifically, they decreased income taxes for households earning less than 80,000 AUD including an increase in the tax free threshold. That these income taxes did not have an immediate impact on electricity demand is not surprising given that the benefits from the income tax cuts would be experienced by consumers in the following year. Moreover, there is an extensive literature documenting that consumers do not respond to marginal electricity prices (Deryugina et al. 2020). In other words, in order for the tax cuts to increase electricity prices, households with incomes less than 80,000 would need to increase electricity consumption in response to income tax cuts that they have not yet experienced.

C.2 Additional Figures

Figure C1: Raw Electricity Price Distribution

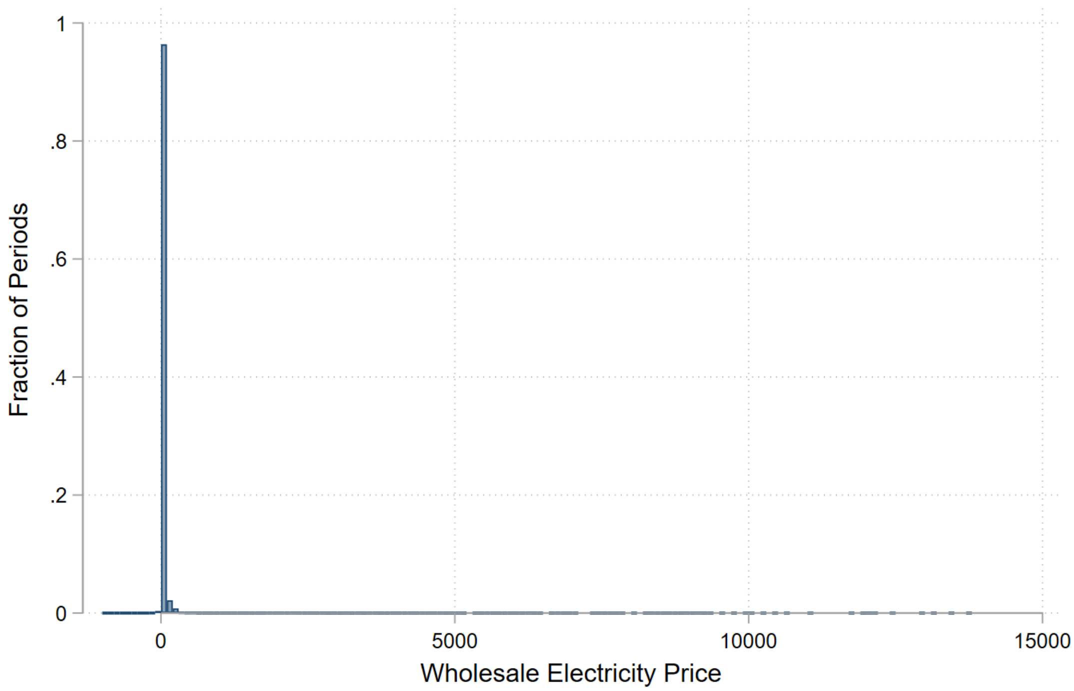

Note: This figure plots the distribution of electricity prices, which highlights the large outlier wholesale electricity prices. Bins are width 100 with well over 90 percent of observations between 0 and 100 AUD/MWh.

Figure C2: Electricity Price Distribution Without Outliers

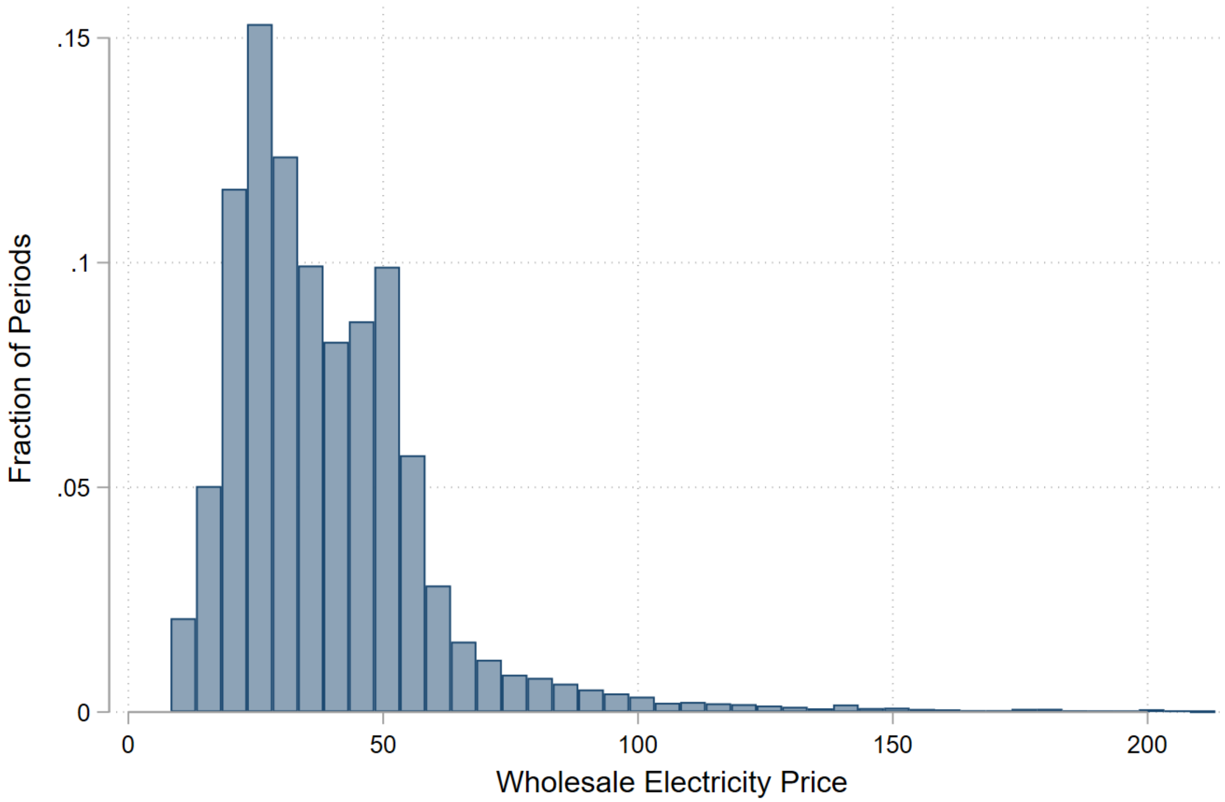

Note: This figure plots the middle 98 percent of the distribution of electricity prices. Bins are width 5.

Figure C3: Placebo RD in 2010 and 2011

(a) Placebo RD (July 1, 2010)

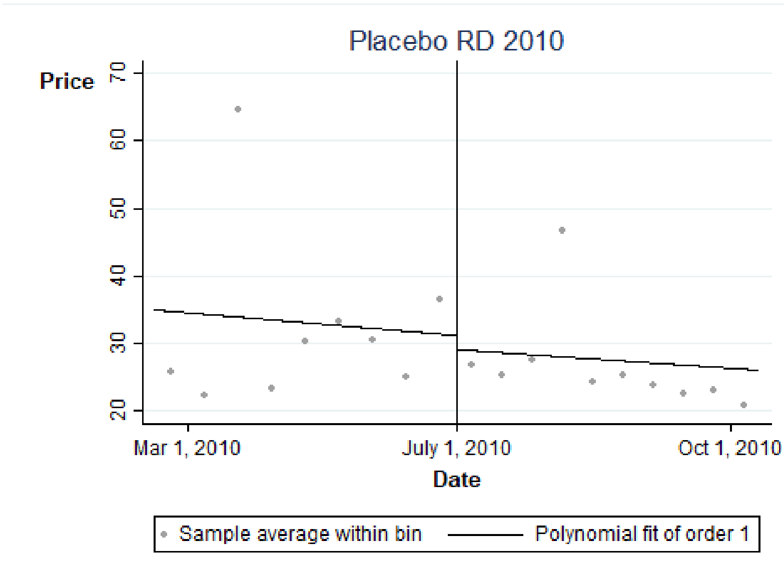

(b) Placebo RD (July 1, 2011)

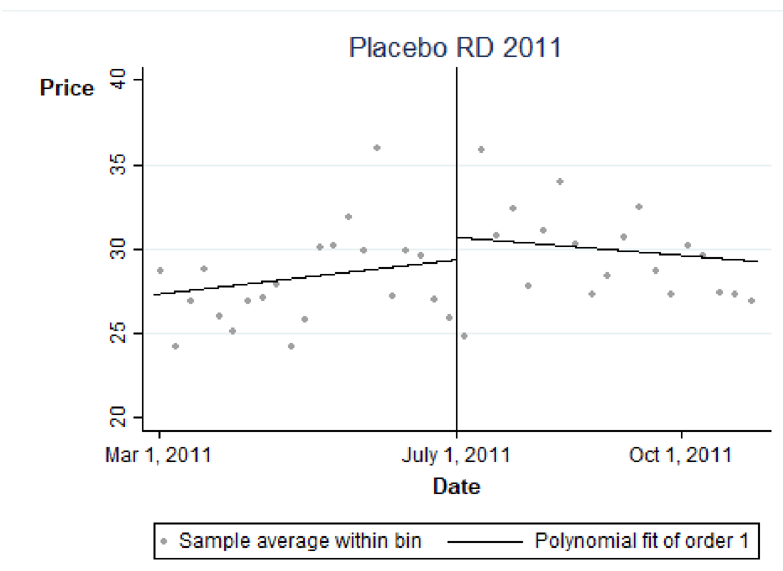

Note: This figure shows statistically insignificant estimates of placebo regression discontinuities on July 1 of the two years before the implementation date. Both figures are local linear regressions with triangular kernels, no controls, and data-driven optimal bandwidths (Calonico et al. 2015).

Figure C4: Equilibrium Quantity Around Implementation Date

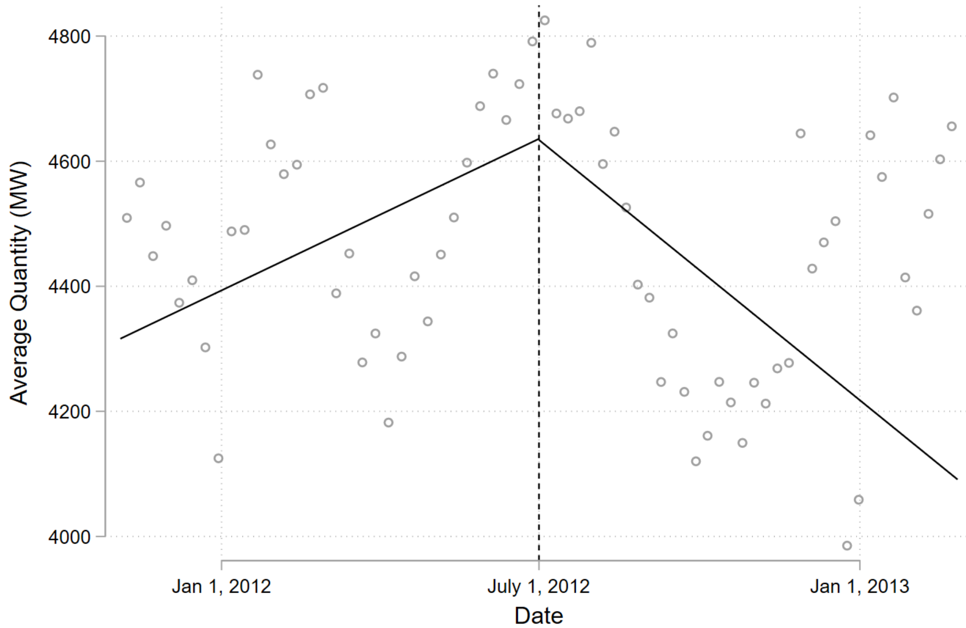

Note: This figure shows statistically insignificant estimates of the change in quantity when the carbon tax is implemented. This figure estimates a local linear regressions with a triangular kernel, no controls, and data-driven optimal bandwidth (Calonico et al. 2015).

Figure C5: Equilibrium Price around Mining Tax Repeal

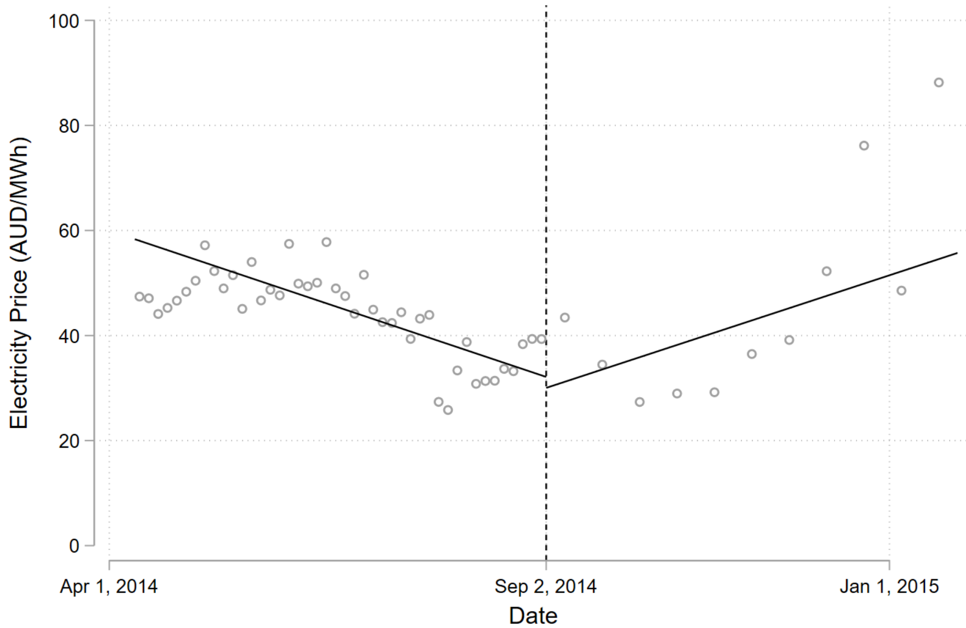

Note: This figure shows statistically insignificant estimates of the change in price when the mining tax is repealed. This figure estimates a local linear regressions with a triangular kernel, no controls, and data-driven optimal bandwidth (Calonico et al. 2015).

Figure C6: Comparison of Treatment Effects with Half-Donut and Sharp RDs

(a) Implementation RD without Controls

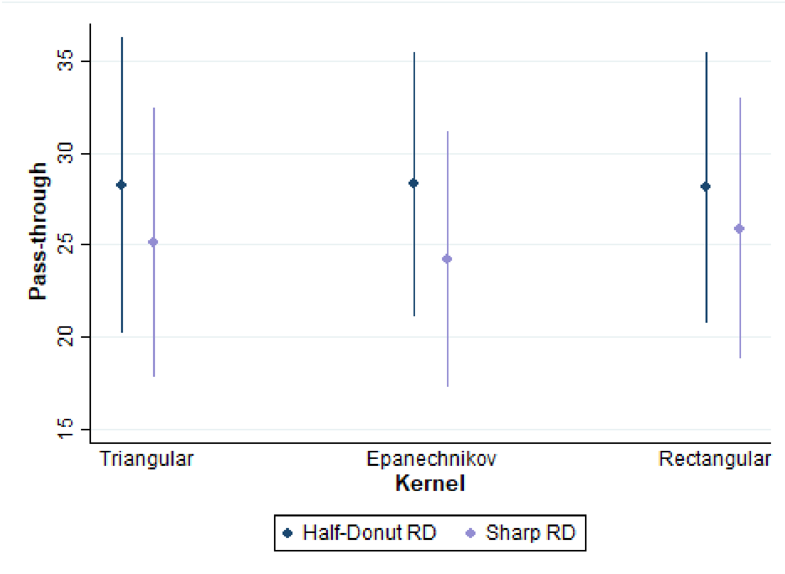

(b) Implementation RD with Controls

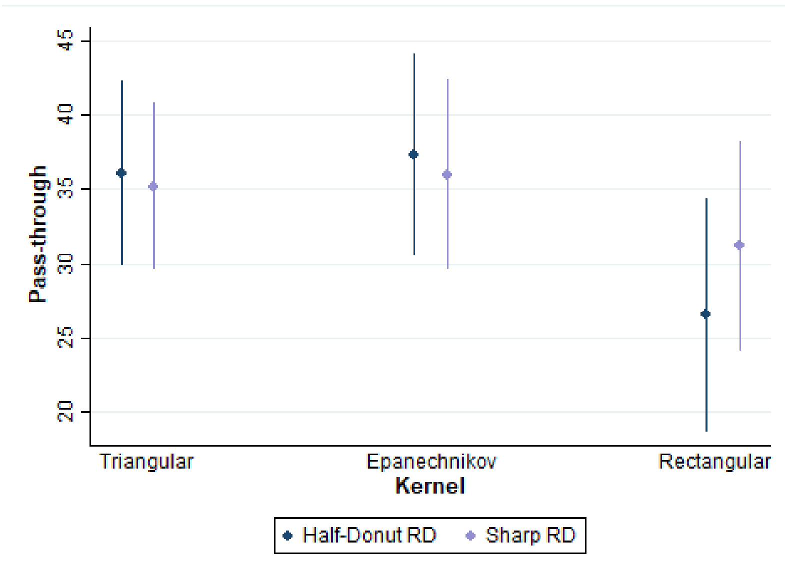

Note: This figure shows the results from estimating half-donut (dropping observations from July 1 to July 17, 2012) and sharp regression discontinuities with the implementation date. The top figure does not include controls while the bottom figure does. The x-axis displays the kernel used and the y-axis the pass-through estimate. Each dot represents a separate estimate and the bars are the 95% confidence interval. The sharp RD is shown in light blue, while the half-donut is shown in dark blue. All half-donut RD estimates lie within the 95% confidence interval for the sharp RD.

Figure C7: Equilibrium Price around Proposed Carbon Tax Repeal being Blocked

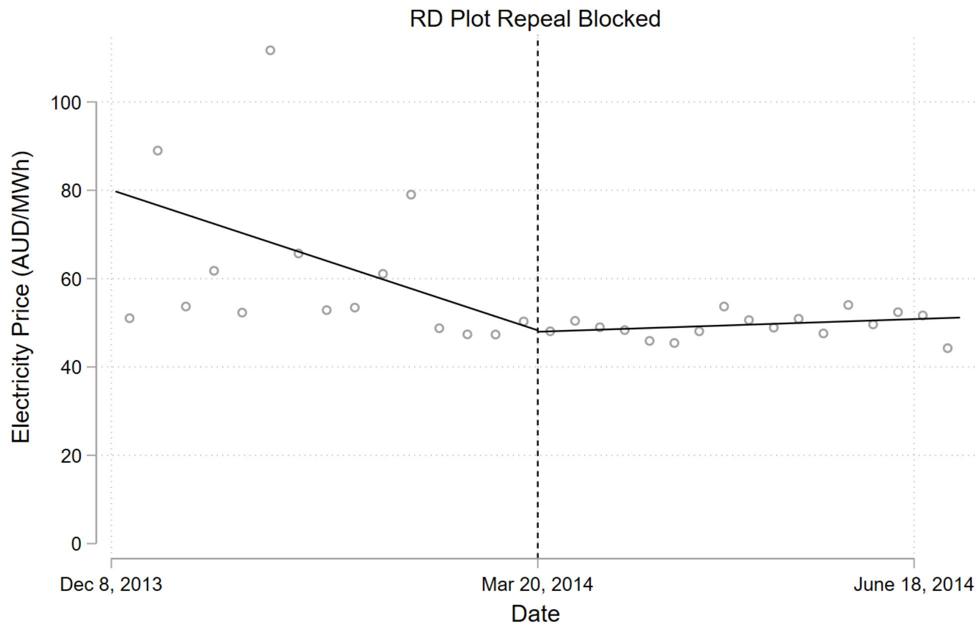

Note: This figure shows statistically insignificant estimates of the change in price when a proposed carbon tax repeal bill was blocked early in 2014. This figure estimates a local linear regressions with a triangular kernel, no controls, and a bandwidth of approximately 100 days.

Figure C8: Average Willingness-to-Accept using Contemporaneous Demand Projections

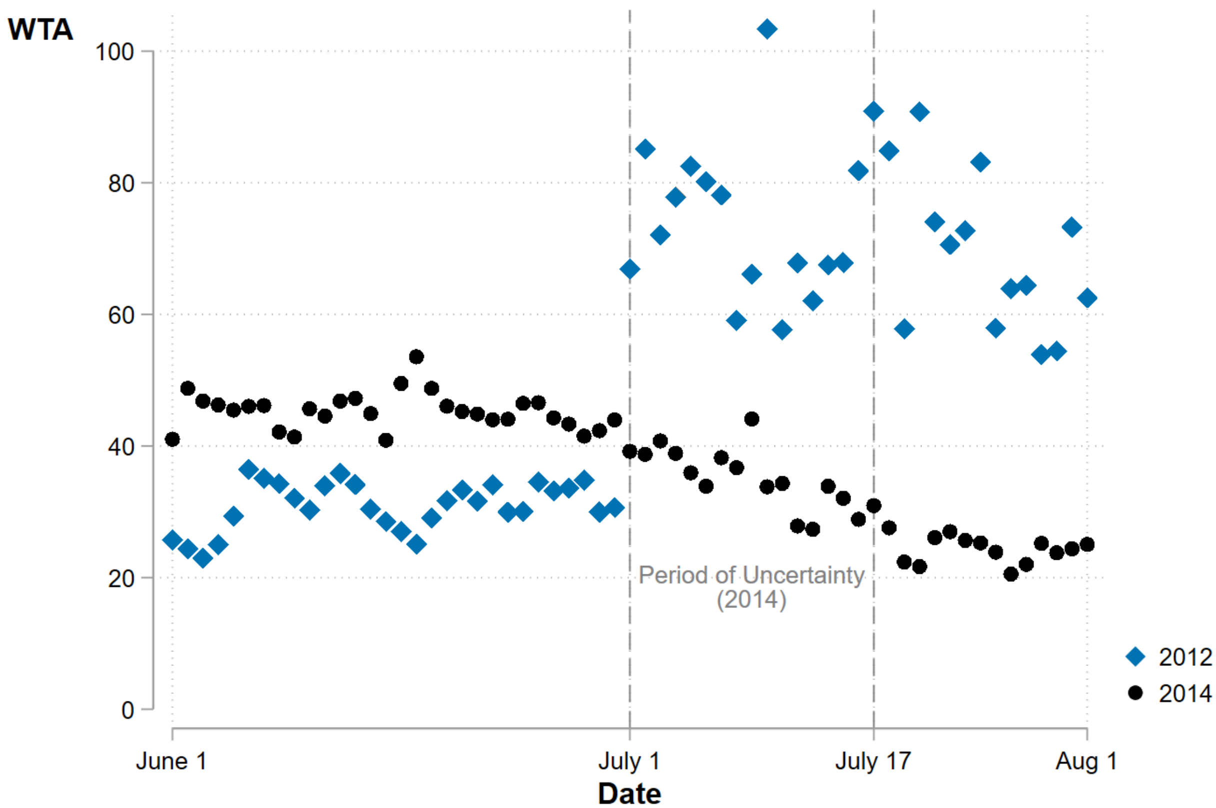

Note: This figure plots the daily average willingness-to-accept in the region of the supply curve that follows within the middle 80% of demand projections supplied for all periods of that day. Blue diamonds represent 2012 averages, while black circles are 2014 averages. The gray dashed lines represent the two key policy dates of July 1 and July 17. For 2012, July 1 is the implementation date. For 2014, July 1 is the legal repeal date while July 17 is the actual date the repeal bill passed. The area in-between is the period of uncertainty for electricity generators where the marginal emissions costs were unknown.

Figure C9: Average Willingness-to-Accept using 2012 Demand Projections

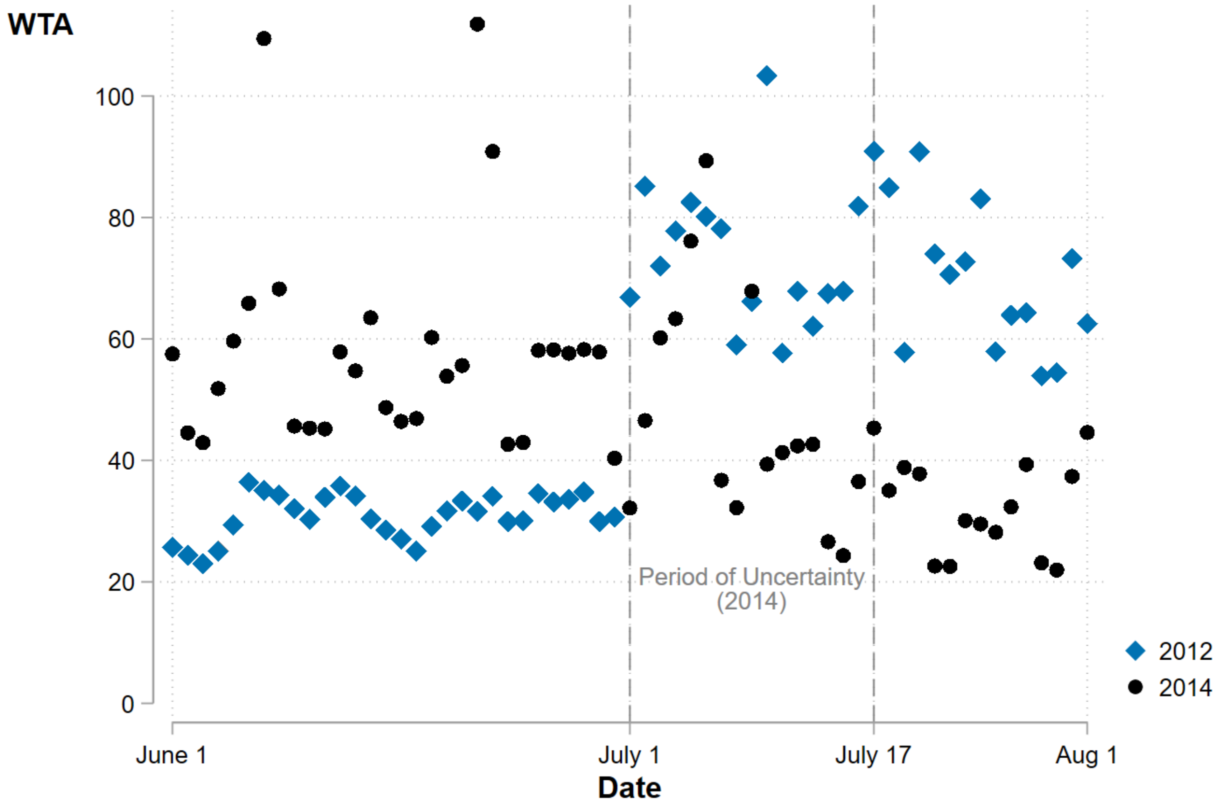

Note: This figure plots the daily average willingness-to-accept in the region of the supply curve that follows within the middle 80% of demand projections supplied for that period in 2012. Blue diamonds represent 2012 averages, while black circles are 2014 averages. The gray dashed lines represent the two key policy dates of July 1 and July 17. For 2012, July 1 is the implementation date. For 2014, July 1 is the legal repeal date while July 17 is the actual date the repeal bill passed. The area in-between is the period of uncertainty for electricity generators where the marginal emissions costs were unknown.

Figure C10: Average Willingness-to-Accept using 2014 Demand Projections

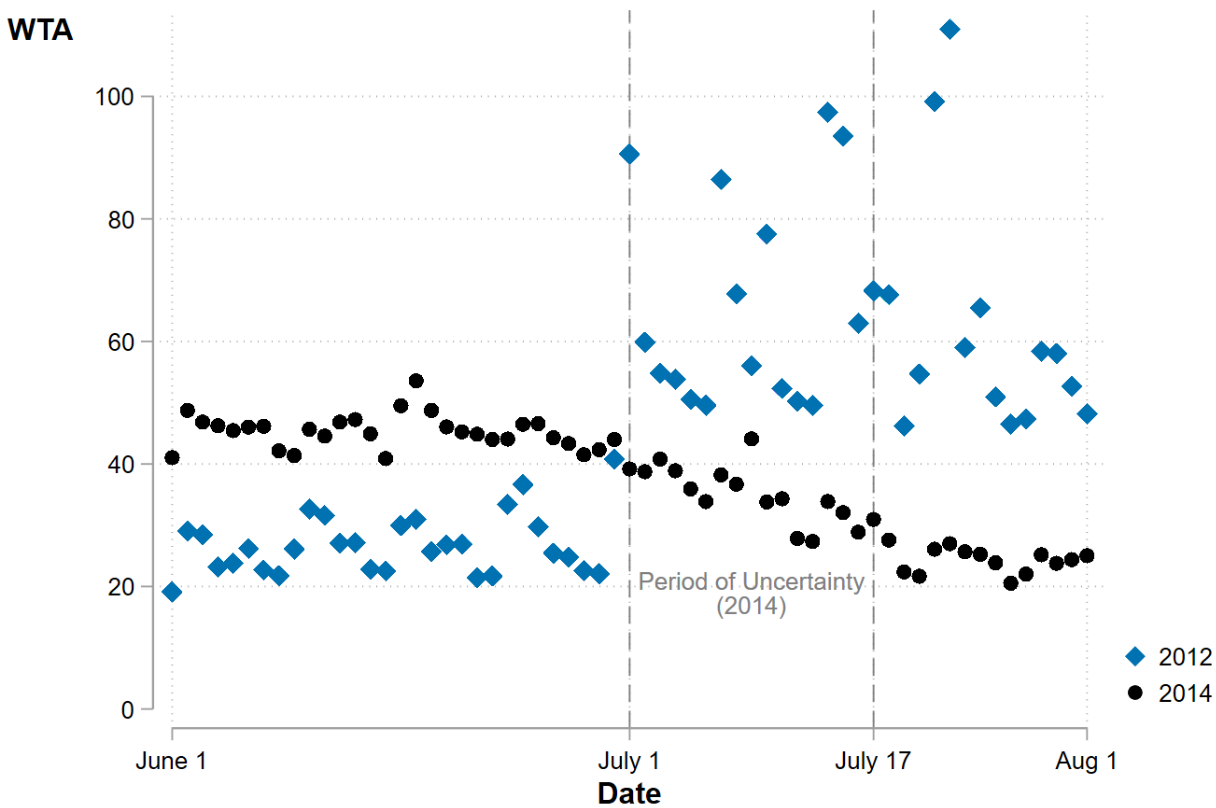

Note: This figure plots the daily average willingness-to-accept in the region of the supply curve that follows within the middle 80% of demand projections supplied for that period in 2014. Blue diamonds represent 2012 averages, while black circles are 2014 averages. The gray dashed lines represent the two key policy dates of July 1 and July 17. For 2012, July 1 is the implementation date. For 2014, July 1 is the legal repeal date while July 17 is the actual date the repeal bill passed. The area in-between is the period of uncertainty for electricity generators where the marginal emissions costs were unknown.

References

Alexander, D., Boone, C., & Lynn, M. (2021). The effects of tip recommendations on customer tipping, satisfaction, repatronage, and spending. *Management Science*, *67*(1), 146–165.

Allen, A. (2014). Review of emission factors for use in the cdeii. *Report to Australian Energy Market Operator*.

Allen, E. J., Dechow, P. M., Pope, D. G., & Wu, G. (2017). Reference-dependent preferences: Evidence from marathon runners. *Management Science*, *63*(6), 1657–1672.

Andersson, J. J. (2019). Carbon taxes and co 2 emissions: Sweden as a case study. *American Economic Journal: Economic Policy*, *11*(4), 1–30.

Australian Energy Regulator. (2018, dec). *Annual report on compliance and performance of the retail energy market 2017-2018* (Tech. Rep.).

Azar, O. H. (2007). The social norm of tipping: A review. *Journal of Applied Social Psychology*, *37*(2), 380–402.

Azar, O. H. (2008). Strategic behavior and social norms in tipped service industries. *The BE Journal of Economic Analysis & Policy*, *8*(1).

Ball, L., & Mankiw, N. G. (1994). Asymmetric price adjustment and economic fluctuations. *The Economic Journal*, *104*(423), 247–261.

Barreca, A. I., Guldi, M., Lindo, J. M., & Waddell, G. R. (2011). Saving babies? revisiting the effect of very low birth weight classification. *The Quarterly Journal of Economics*, *126*(4), 2117–2123.

Benabou, R., & Gertner, R. (1993). Search with learning from prices: does increased inflationary uncertainty lead to higher markups? *The Review of Economic Studies*, *60*(1), 69–93.

Benzarti, Y., Carloni, D., Harju, J., & Kosonen, T. (2020). What goes up may not come down: asymmetric incidence of value-added taxes. *Journal of Political Economy*, *128*(12), 4438–4474.

Borenstein, S., Cameron, A. C., & Gilbert, R. (1997). Do gasoline prices respond asymmetrically to crude oil price changes? *The Quarterly journal of economics*, *112*(1), 305–339.

Bushnell, J. B., Chong, H., & Mansur, E. T. (2013). Profiting from regulation: Evidence from the european carbon market. *American Economic Journal: Economic Policy*, *5*(4), 78–106.

Cabral, L., & Fishman, A. (2012). Business as usual: A consumer search theory of sticky prices and asymmetric price adjustment. *International Journal of Industrial Organization*, *30*(4), 371–376.

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, *82*(6), 2295–2326.

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2015). Optimal data-driven regression discontinuity plots. *Journal of the American Statistical Association*, *110*(512), 1753–1769.

Chandar, B., Gneezy, U., List, J. A., & Muir, I. (2019). *The drivers of social preferences: Evidence from a nationwide tipping field experiment* (Tech. Rep.). National Bureau of Economic Research.

Deryugina, T., MacKay, A., & Reif, J. (2020). The long-run dynamics of electricity demand: Evidence from municipal aggregation. *American Economic Journal: Applied Economics*, *12*(1), 86–114.

Donkor, K. B. (2020). How social norms and menus affect choices: Evidence from tipping.

Doyle, J. J., & Samphantharak, K. (2008). $2.00 gas! studying the effects of a gas tax moratorium. *Journal of public economics*, *92*(3-4), 869–884.

Evans, W. N., Ringel, J. S., & Stech, D. (1999). Tobacco taxes and public policy to discourage smoking. *Tax policy and the economy*, *13*, 1–55.

Eyster, E., Madarász, K., & Michaillat, P. (2020). Pricing under fairness concerns. *Journal of the European Economic Association*.

Fabra, N., & Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, *104*(9), 2872–99.

Farber, H. S. (2015). Why you can't find a taxi in the rain and other labor supply lessons from cab drivers. *The Quarterly Journal of Economics*, *130*(4), 1975–2026.

Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy cost pass-through in us manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, *12*(2), 303–42.

Genesove, D., & Mullin, W. P. (1998). Testing static oligopoly models: conduct and cost in the sugar industry, 1890-1914. *The RAND Journal of Economics*, 355–377.

Goulder, L. H., Hafstead, M. A., Kim, G., & Long, X. (2019). Impacts of a carbon tax across us household income groups: What are the equity-efficiency trade-offs? *Journal of Public Economics*, *175*, 44–64.

Grainger, C. A., & Kolstad, C. D. (2010). Who pays a price on carbon? *Environmental and Resource Economics*, *46*(3), 359–376.

Haggag, K., & Paci, G. (2014). Default tips. *American Economic Journal: Applied Economics*, *6*(3), 1–19.

Harasztosi, P., & Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, *109*(8), 2693–2727.

Harris, L. (1991). Stock price clustering and discreteness. *The Review of Financial Studies*, *4*(3), 389–415.

Hausman, C., & Rapson, D. S. (2018). Regression discontinuity in time: Considerations for empirical applications. *Annual Review of Resource Economics*, *10*, 533–552.

Hoover, H. (2019). Default tip suggestions in nyc taxi cabs. *Available at SSRN 3333460*.

Hortaçsu, A., Luco, F., Puller, S. L., & Zhu, D. (2019). Does strategic ability affect efficiency? evidence from electricity markets. *American Economic Review*, *109*(12), 4302–42.

Isaac, M. S., Wang, Y., & Schindler, R. M. (2020). The round-number advantage in consumer debt payoff. *Journal of Consumer Psychology*.

Ito, K. (2014). Do consumers respond to marginal or average price? evidence from nonlinear electricity pricing. *American Economic Review*, *104*(2), 537–63.

Kim, H. (2021). Heterogeneous impacts of cost shocks, strategic bidding and pass-through: Evidence from the new england electricity market. *American Economic Journal: Microeconomics*, *Forthcoming*.

Lee, D. S., & Card, D. (2008). Regression discontinuity inference with specification error. *Journal of Econometrics*, *142*(2), 655–674.

Lerner, A. (1934). The concept of monopoly and the measurement of monopoly. *Review of Economic Studies*, *1*(3), 157-75.

Loris, N., & Jolevski, F. (2014). Epa's climate regulations will harm american manufacturing. *Heritage Foundation Issue Brief*, *4158*.

Loy, J.-P., Weiss, C. R., & Glauben, T. (2016). Asymmetric cost pass-through? empirical evidence on the role of market power, search and menu costs. *Journal of Economic Behavior & Organization*, *123*, 184–192.

Lynn, M., Flynn, S. M., & Helion, C. (2013). Do consumers prefer round prices? evidence from pay-what-you-want decisions and self-pumped gasoline purchases. *Journal of Economic Psychology*, *36*, 96–102.

Marshall, A. (1890). Principles of economics macmillan. *London (8th ed. Published in 1920)*.

Murnane, R. J., & Willett, J. B. (2010). *Methods matter: Improving causal inference in educational and social science research.* Oxford University Press.

Nazifi, F., Trück, S., & Zhu, L. (2017). Carbon pass-through rates on electricity spot prices in australia.

Parente, P. M., & Silva, J. M. S. (2016). Quantile regression with clustered data. *Journal of Econometric Methods*, *5*(1), 1–15.

Peltzman, S. (2000). Prices rise faster than they fall. *Journal of political economy*, *108*(3), 466–502.

Ritz, R. (2015). The simple economics of asymmetric cost pass-through.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., & Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *The Stata Journal*, *19*(1), 4–60.

Schindler, R. M., & Wiman, A. R. (1989). Effects of odd pricing on price recall. *Journal of Business Research*, *19*(3), 165–177.

Sijm, J., Neuhoff, K., & Chen, Y. (2006). Co2 cost pass-through and windfall profits in the power sector. *Climate policy*, *6*(1), 49–72.

Tappata, M. (2009). Rockets and feathers: Understanding asymmetric pricing. *The RAND Journal of Economics*, *40*(4), 673–687.

Teeter, P., & Sandberg, J. (2017). Constraining or enabling green capability development? how policy uncertainty affects organizational responses to flexible environmental regulations. *British Journal of Management*, *28*(4), 649–665.

Thakral, N., & Tô, L. (2019). *Tipping and the dynamics of social norms* (Tech. Rep.). Mimeo.

Vesterberg, M. (2016). The hourly income elasticity of electricity. *Energy Economics*, *59*, 188–197.

Weyl, E. G., & Fabinger, M. (2013). Pass-through as an economic tool: Principles of incidence under imperfect competition. *Journal of Political Economy*, *121*(3), 528–583.

Woerman, M. (2019). Market size and market power: Evidence from the texas electricity market. *Work. Pap.*.

Zachmann, G., & Von Hirschhausen, C. (2008). First evidence of asymmetric cost pass-through of eu emissions allowances: Examining wholesale electricity prices in germany. *Economics Letters*, *99*(3), 465–469.