ESSAYS ON PUBLIC POLICY AND CONSUMER WELFARE

A Dissertation

by

TRENT GAVIN MCNAMARA

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Steven L. Puller
Committee Members,	Fernando A. Luco
	Ragan A. Petrie
	Anastasia V. Shcherbakova
Head of Department,	Timothy J. Gronberg

May 2020

Major Subject: Economics

ABSTRACT

This research presents three essays on the effects of different institutions, technologies, and policies on informational outcomes, political support, and consumer welfare using experimental and structural research designs. Specifically, I consider the effects of social media, uniform gasoline taxes, and behavior changes in electricity markets.

In the first essay "The Economic Effects of Facebook", joint work with Roberto Mosquera, Mofioluwasademi Odunowo, Xiongfei Guo, and Ragan Petrie, we study the effects of Facebook on news awareness, subjective well-being, and daily activities. We use a large field experiment with a validated Facebook restriction to document the value of Facebook to users and its causal effect on news consumption and awareness, well-being, and daily activities. Those who are off Facebook for a week reduce news consumption, are less likely to recognize politically-skewed news stories, report being less depressed and engage in healthier activities. One week of Facebook is worth $67, and this increases by 19.6 percent after experiencing a Facebook restriction.

In the second essay "Who Supports Pigou? The Distributional Consequences of Pigouvian Taxes", joint work with Steve Puller, we study the distributional effects resulting from a uniform gasoline tax and its impact on political support. We first design an optimal gasoline tax that accounts for the external damages associated with local pollutants and show that there exists significant heterogeneity in both the costs and benefits borne from this uniform tax regime. Following this, we show that this distribution is an important determinant for an individual's level of support for a gasoline tax even after controlling for political identity, indicating that tax support is partially determined by self-interests rather than simply "tribal" politics. Then using this relationship, we design a revenue neutral tax policy and show that there exists a meaningful way in which revenue can be returned to individuals such that support for raising the gasoline tax increases to a median level of support of 5 on a scale of 10.

In the third essay "Price Leadership and Learning in Oligopoly: Evidence from Electricity Markets", I study the process by which firms transition between multiple equilbria. Specifically, in oligopoly markets where firms compete in supply functions, there exists a wide range of potential equilibria with significant differences in market outcomes. In this setting, transitioning between equilibria can be highly profitable. I use firm offer data into 15-minute electricity auctions to show the process by which firms transition from a low price supply function equilibrium (SFE) to a high price equilibrium. I first document a price leader's deviation from equilibrium play which serves as a signal for other firms to deviate as well. Firms forego short-term profits in a dynamic learning environment to transition to a high price equilibrium. This shift in equilibrium is associated with an average price increase of 5% but can be as large as 1,500% in some periods. This also generates profits significantly larger than those foregone by signaling. In order to speak to how learning occurs during the transition period, I integrate a fictitious play learning model into a model of dynamic profit maximization. In general, firms learn and respond to each other's more recent actions. From a market design perspective, this allows me to estimate how the timing and release of historical information impacts market outcomes. I show that with enough of a data release lag, firms would forego transitioning altogether.

DEDICATION

To Shelby, my wife, for her constant support in all of my endeavors. To my parents, Ron and Nancy, who from Heaven and Earth continue teaching me life's most valuable lessons.

ACKNOWLEDGMENTS

I am grateful to my advisor, Steve Puller, for continuous guidance, advice, and encouragement throughout my graduate studies. I am thankful for investing the time and energy to mentor me through each step of the research process and for teaching me to always think big.

I also want to thank the other members in my committee, Fernando Luco, Ragan Petrie, and Anastasia Shcherbakova, who provided valuable feedback on all my projects including those that are not a part of this dissertation. Additionally, I thank everyone in the applied microeconomics group at Texas A&M for comments throughout countless presentations and seminars.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supported by a dissertation committee consisting of Professors Steve Puller, Fernando Luco, and Ragan Petrie of the Department of Economics at Texas A&M and Professor Anastasia Shcherbakova of Department of Agricultural Economics at Texas A&M.

The analyses depicted in Section 2 were conducted in part by Roberto Mosquera, Mofioluwasademi Odunowo, Xiongfei Guo, and Ragan Petrie of the Department of Economics; the analyses depicted in Section 3 were conducted in part by Steve Puller of the Department of Economics at Texas A&M.

All other work conducted for the dissertation was completed by the student independently.

Funding Sources

Graduate study was supported by teaching and research assistantships from Texas A&M University.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE Page

LIST OF TABLES

TABLE Page

1. INTRODUCTION

The efficient design of social and economic policies requires knowledge and understanding of how institutional complexity, emerging technologies, and historical policies interact and ultimately impact outcomes of interest to society. In a lot of cases, policymakers, academics, and society may not be aware of how these institutions affect individual behaviors and firm decisions. By extension, it is also not well understood how fundamental changes to these institutions, whether through technological advancements or behavioral responses, impacts societal welfare. Following this, in this dissertation, I use experimental and structural research methods to analyze three issues that directly relate to policies affecting the welfare of consumers.

In Section 2, together with Roberto Mosquera, Mofioluwasademi Odunowo, Xiongfei Guo, and Ragan Petrie, we study the causal effect of Facebook on various aspects of daily life. We are interested in studying Facebook's direct impact on a suite of outcomes pertaining to news awareness, subjective wellbeing, daily activities related to mood, and Facebook's value to its users. The development and growth of social media has characterized the 21st century, and Facebook can arguably be labeled as the industry leader. Almost 2.2 billion people worldwide have an active Facebook account, and nearly 1.4 billion log on daily for an average of 50 minutes per day. While social media usage has increased substantially over the last 15 years, this technological advancement has fundamentally changed society and how individuals relate to each other as well as to other institutions. Facebook not only provides means to connect with friends and build social networks and capital, but it is also exposes users to a vast amount of information and news. Some of these impacts can be beneficial and others can be harmful. For example, as individuals rely more on social media as their primary source of information, they are connected to a much larger pool of news outlets. Some of these outlets are credible, others are less so. This exposure can manifest and influence in political markets through changes in voting behaviors. The consequences of social media in terms of news awareness and biases - highlighted by political investigations regarding Facebook's involvement in the 2016 U.S. Presidential Election - are largely unknown.

We do not know how people perceive the costs and benefits associated with social media. Hence, estimating Facebook's monetary value to users is a critical component in understanding the total utility an individual receives as well as the surplus generated to consumers from the platform. Furthermore, knowing the value individuals have for Facebook is an important measure for policymakers to address the potential unintended negative consequences and externalities of social media usage on various aspects of daily life. Therefore, Facebook's monetary value along with its effects on news awareness and well-being is an important but under-researched aspect of the 21st century.

In order to answer these questions, we designed and implemented a field experiment with a randomized and validated week-long Facebook restriction. This randomized intervention generates exogenous variation in Facebook usage, allowing us to identify how Facebook affects daily activities, news awareness and exposure, and subjective well-being. In order to quantify the monetary value users have for the platform, we embed an incentive-compatible auction mechanism into our experimental design. We find that that one week of Facebook is worth about $67 to users, with a median value of $40. Regarding Facebook's effects on news, we find that Facebook is a major source of news exposure. Importantly, individuals restricted from Facebook are less aware of news coming from politically-skewed sources. This effect is stronger for men than women. In order to explain the mechanism for this, we show evidence that a reduction in news consumption drives this result and that participants do not substitute towards other news sources or social media platforms when being off Facebook for a short period. Finally, regarding subjective well-being, we show no significant effect of using Facebook on overall life satisfaction. However, we do find a sizeable short-term reduction in feelings of depression when restricted from Facebook, especially for men. We build on existing research by studying the effect of Facebook on behaviors correlated with mood. We find that individuals restricted from using Facebook engage in healthier activities.

In Section 3, together with Steve Puller, we study the relationship between the distributional consequences of a uniform gasoline tax and political support for the tax itself. The externalities created by the personal transportation market for gasoline consumption in the United States impose a large cost on society as a whole. When considering the full spectrum of damages (e.g., local and global pollution, congestion and waste, health costs, crop and timber yields, and even more "minor" consequences such as noise pollution) total damage estimates incurred by driving can easily be upwards of $400 billion per year. In order to address this large societal cost, the United States has elected to regulate vehicle manufacturers through Corporate Average Fuel Economy (CAFE) Standards instead of through Pigouvian taxes despite a large and growing literature showing that the former is anywhere from 3 to 6 times less cost efficient than the latter. Given how large these costs are, it is important for social welfare to efficiently address the issue. If not, then there are substantial welfare consequences. One reason that this inefficiency remains status quo stems from the fact that American voters have never had a favorable opinion for gasoline taxes. This lack of support for a gasoline tax creates a complex paradigm for policy-makers and academics alike. If gasoline taxes are too politically unpopular, this means that for policy-makers to remain favorable with their respective constituency, they have strong incentives to pass sub-optimal policies (i.e., CAFE Standards) rather than increasing gasoline taxes. Furthermore, while uniform gasoline taxes can be welfare enhancing in the aggregate, there exists heterogeneity amongst who receives those benefits and who bears the costs.

In order to investigate the relationship between gasoline taxes and political support, we use confidential DMV data for the state of Texas augmented with data on local pollution damages (for NOX, VOC, PM25, and SO2) to design an optimal uniform gasoline tax that properly accounts for these local pollution damages. In practice, this gasoline tax amounts to $0.40 per gallon. We then counterfactually increase the price of gasoline by this uniform tax and estimate how this increases individual costs (gasoline expenditures) and county benefits (reduction in pollution damages). Comparing these two measures of costs and benefits, we find that there is substantial heterogeneity across the state that results from a uniform increase in the gasoline tax. However, individuals with small costs and large benefits tend to be located in urban and dense regions of the state, while those with large costs and small benefits tend to be located in the more rural areas. Following this result, we use data that comes from a stratified state-wide survey and show evidence that this distribution of costs and benefits is an important determinant in an individual's decision to support an increase in the gasoline tax. Even after controlling for an individual's political affiliation, increasing a counties' net benefits can explain up to 10% of an individual's support for a gasoline tax. This result has multiple important policy implications. First, it establishes a behavioral link between self-interests and tax support that goes beyond "tribal" politics. Second, this link enables policy-makers to design tax regimes that are "politically-sophisticated." Given this, we design a revenue neutral tax regime whereby individuals in counties with higher proportions of people who have costs that outweigh their benefits receive a larger tax dividend than those with smaller proportions. This "politically-sophisticated" tax regime results in strategically shifting the distribution of support for a Pigouvian gasoline tax from a median level of 4 (out of 10) to a 5, representing a significant and meaningful change.

In Section 4, I study the process by which firms transition between multiple equilbria in wholesale electricity markets. Multiple equilibria are pervasive in many markets and industries and their existence complicates both the design and the role of regulation. This is particularly the case for oligopoly markets where firms choose supply functions rather than just price or quantity. In these markets, there exists a wide range of potential equilibria resulting in an equally wide range of market outcomes in terms of prices and profits. Hence, since a significant number of prices above can be sustained in equilibrium, these differences generates the incentive for firms to transition away from low price equilibria towards high price equilibria. Despite this, surprisingly little is known about how an equilibrium is reached as well as how agents transition between them. Learning more about the transition process has important market design implications about how to optimally disclose and reveal information to the market and implications for how to regulate and monitor firms in a market.

In order to study these concepts, I first develop and make use of a model of static unilateral profit maximization. I show that a small firm in the market suddenly begins to deviate from a low-price supply function equilibrium (SFE) and foregoes static profits (≈$3,500 per offer) in favor of over-pricing a significant portion of its production. After multiple deviations by this price leader, the largest firm in the market begins to follow a similar behavior and begins to over-price its production (forgoing ≈\$1,200 per offer). As both firms iterate and learn about each other's actions, they are able to effectively reduce their production which results in prices associated with a high-price equilibrium. I show that this shift in equilibrium play is associated with an average price increase of 5%, but can importantly result in swings as large as 1,500%. I find preliminary evidence that these outcomes are consistent with a dynamic repeated game equilibrium. To explicitly analyze the process of learning during the transition period, I integrate a fictitious play learning model into a model of dynamic profit maximization. This characterizes a fixed point equilibrium regarding the beliefs that each firm must have about the other in order for it to be optimal to initially deviate. Since these belief parameters are used by firms to form expectations about each other's current and future actions, there is room for policy-makers to limit available information through the form of information disclosure policies. These types of policies define a procedure where market information is released with a lag. By preventing more recent information from being observable, firms would need to forgo an increased amount of static profits to arrive at the high price SFE. In fact, I estimate that revealing information with a 10-day lag prevents firms from transitioning to the high priced equilibrium altogether.

2. THE ECONOMIC EFFECTS OF FACEBOOK*

2.1 Introduction

Social media usage has increased dramatically over the past decade, and Facebook has dominated the market. Almost 2.2 billion individuals worldwide have an active Facebook account, and nearly 1.4 billion log on daily (Facebook, 2017) for an average of 50 minutes per day (Facebook, 2016). Facebook not only provides means to connect with friends and build social networks and capital (Bailey et al., 2018; Mayer and Puller, 2008; Shirley Cramer, 2017), but it is also exposes users to a vast amount of information and news. Despite the potential influence of Facebook on an individual's behavior via information and content provision, there is surprisingly little known about its direct and comprehensive effects on news exposure and awareness, subjective well-being and day-to-day activities.

Facebook's platform has several characteristics that lend well to investigating its effects on an individual's exposure to news content as well as its impact on well-being. The platform consolidates information from many sources, making it an important and compelling place to go on the internet to keep up with news. People tap into Facebook for local, national and international news. Indeed, roughly two-thirds of Americans get at least some of their news from social media sources (Pew Research Center, 2017). While there is a concern that news transmitted through social media could be fake or skewed and affect political outcomes (Allcott and Gentzkow, 2017), these type of platforms could also serve to uncover corruption (Enikolopov et al., 2016). As individuals rely more on social media and news aggregators as a primary source of information, segregation may increase (Gentzkow and Shapiro, 2011) and voting behavior can be affected (DellaVigna and Kaplan, 2007; Bond et al., 2012; Martin and Yurukoglu, 2017). The consequences of this in terms of news awareness and biases - highlighted by political investigations regarding Facebook's involvement in the 2016 U.S. Presidential Election - are largely unknown.

*Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

More broadly, there is little consensus on Facebook's impact on well-being, especially in the context of daily behaviors and activities. Facebook is often used to connect with friends and family, organize events and share information and photos (Laroche et al., 2012; De Vries et al., 2012; Ashley and Tuten, 2015; Lee and Ma, 2012; Bailey et al., 2017). Being able to seamlessly keep in touch with others might improve mood and happiness, but it might also induce negative emotions and habits from social comparison (Tromholt, 2016; Deters and Mehl, 2013). How Facebook directly affects well-being and mood in general and the correlation with daily activities is unclear.

Facebook's platform is provided for free to users and paid for by advertising, so the monetary value to users, as reflected in a market price, is untested. The platform facilitates building social networks and seamless access to relevant information. Usage rates, both in frequency and intensity, suggest this provides benefits to users. While the economic impact of Facebook on advertising has been estimated, the benefits to users and impact on behavior have been given more limited study.[1] Knowing the value of Facebook would inform an understanding of welfare effects and provide a monetary measure of the importance of Facebook to users.

We ran a field experiment in the Spring of 2017 with a randomized, and validated, Facebook restriction to investigate how Facebook may affect daily activities and news exposure and quantify how much users value access. In total, 1,769 individuals from a large U.S. university participated in the study. Using an incentive-compatible procedure (Becker et al., 1964), we asked participants how much they would need to be paid to not use Facebook for one week. Qualified participants were then randomly assigned to either a one-week Facebook restriction group or a control group that faced no restriction.

Our design has several important and unique features worth noting. First, we can exploit the rich data collected on the distribution of Facebook's value to check for possible selection effects in our results. Second, we enforced and validated the restriction by logging participants off Facebook and verified treatment compliance using an unobtrusive online monitoring procedure throughout the week. Our procedure was undetectable to the participant and did not involve direct contact

[1] It is estimated that the impact of Facebook through advertising is $77.6 billion in the U.S. (Deloitte, 2015). Evidence on the value of Facebook is given in Brynjolfsson et al. (2018).

which could potentially impact behavior. Finally, participants completed two surveys, the first prior to random assignment and a second survey one week later. These surveys were designed to provide a comprehensive view of behavior and measure the short-term effects of Facebook on news awareness and consumption, well-being, daily time allocation and daily activities.

We have several key results. First, our study reveals that one week of Facebook is worth about $67 to users, with a median value of $40. This value is in line with other studies (Brynjolfsson et al., 2018; Corrigan et al., 2018; Allcott et al., 2019; Sunstein, 2018; Herzog, 2018) and represents a significant portion of a typical university student's weekly budget and expenses (roughly 30% according to Flood et al. (2017)).[2] Individuals place a nontrivial value on Facebook usage, and the value increases 19.6% after not being able to use it for one week. This is consistent with addiction or the compounding loss of information, however, we note this is only suggestive as we are underpowered to detect a statistically significant effect.

Second, our data document that Facebook is an important source of news exposure. Individuals restricted from Facebook are less aware of politically-skewed sources, and this is stronger for men than women. Consistent with this result, the Facebook restriction reduces news consumption and participants do not substitute towards other news sources or social media platforms when being off Facebook for a short period of time. There is no effect on news awareness from mainstream sources. The causal estimates show that Facebook is an important conduit for news from non-mainstream outlets, and this echoes the findings of Allcott and Gentzkow (2017) who show that social media is correlated with the distribution of "fake news." Our results provide additional evidence that Facebook plays an important role in the acquisition of information by affecting what news is available to consume and thus an individual's ability to assess its veracity.

Third, our findings contribute to the literature that focuses on Facebook's effect on happiness and well-being. Early studies found mostly positive effects of social media on subjective well-being, perhaps through enhanced engagement, in cross-sectional studies (Ellison et al., 2007;

[2]We note that the BDM mechanism used in our study, and in other studies using the BDM or other mechanisms, involve hassle costs and some complexity that may affect values. Our participants face a one in two chance of experiencing a Facebook restriction, and this may reduce bias in value estimates when using elicitation mechanisms coupled with implementation uncertainty.

Valenzuela et al., 2009; Gonzales and Hancock, 2011; Kim and Lee, 2011) and laboratory experiments (Sagioglou and Greitemeyer, 2014; Vogel et al., 2015; Verduyn et al., 2015). More recent studies have found mixed results using panel data (Shakya and Christakis, 2017) and Facebook use limitations (Tromholt, 2016).[3] Cross-sectional evidence on the effect of Facebook on depression is mixed. Feinstein et al. (2013) finds depressive feelings are driven by negative outcomes from social comparison, but other studies find no relationship between Facebook and depression (Steers et al., 2014; Jelenchick et al., 2013; Tandoc et al., 2015). We contribute to this literature by using a randomized and verified Facebook restriction and show no significant effect of using Facebook on overall life satisfaction. However, we do find a large short-term reduction in feelings of depression when restricted from Facebook, especially for men.

Finally, we build on existing research by studying the effect of Facebook on behaviors largely found to be correlated with mood. We find suggestive evidence that individuals restricted from using Facebook engage in healthier activities. While our design does not allows us to recover the underlying mechanism, this finding is consistent with research in psychology (Salovey et al., 2000; Ostir et al., 2000; Fredrickson and Joiner, 2002; Blake et al., 2009; Kettunen et al., 2015; Newman et al., 2014; Sonnentag, 2001) that better mood is positively correlated with engagement in healthier behaviors.

Overall, the effects our study finds on news awareness, news consumption, feelings of depression and daily activities show that Facebook has significant effects on important aspects of life not directly related to building and supporting social networks. Furthermore, almost two years after our experiment, Allcott et al. (2019) find similar results for news awareness and subjective well-being for a different population, which supports our findings. The effects of Facebook are far reaching, and our results provide a more comprehensive documentation of these impacts on daily life. Users seem to understand this and place a substantial value on the experience that Facebook provides.

The paper is organized as follows. Section 2 describes the study design and implementation.

[3] Tromholt (2016) uses a one-week, self-enforced Facebook restriction and finds a positive effect on overall life satisfaction.

Section 3 reports results on the value of Facebook to users and the effect of the Facebook restriction on news awareness, subjective well-being and activities. Section 4 continues with robustness checks on our main findings. Section 5 concludes.

2.2 Study Design

A direct approach to analyze the causal effects of Facebook on daily life would be to take the population of Facebook users, randomly restrict usage for some and not others and then examine behavior across the restricted and not restricted groups. This is difficult to achieve, however, absent a random event that blocks some comparable users from accessing Facebook for period of time and not others and then identifying those users to examine behavior. As an alternative, we adopt an approach where we recruit volunteers and then randomize a Facebook restriction among them.[4] While feasible to implement, a challenge is the representativeness of the generated sample. Simply asking for volunteers willing to give up Facebook would likely result in a sample of low-value individuals. To address this issue, we collect additional information from our volunteers that allows us to account for this type of selection. Rather than merely asking for volunteers, we elicit an individual's value of Facebook for one week and then use the distribution of stated values to test if selection affects the results.

Our study occurs in three major phases, as outlined in Figure 2.1. In Phase 1, we elicit an individual's value of using Facebook for one week and recruit qualified participants into the Facebook restriction. In Phase 2, we administer a pre-treatment survey and then randomly assign participants into two groups – a group that experiences one week without Facebook and a group with no restriction. In Phase 3, participants return to complete a second survey and collect payments. In a surprise, we also re-elicit an individual's value of Facebook for one week. We ran this intervention between April and May 2017.

[4]The study is registered in the AEA Registry (AEARCTR-0003952)

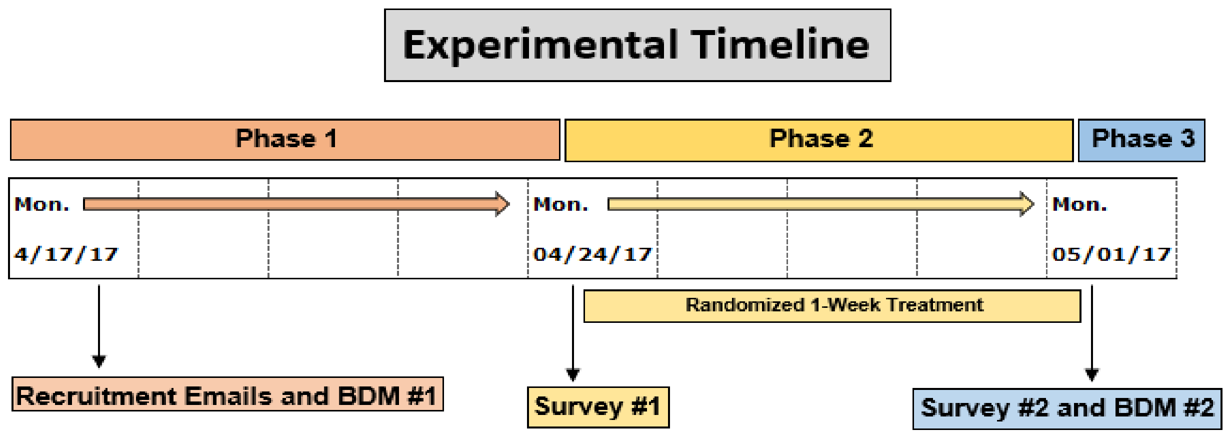

Figure 2.1: Timeline of Experimental Phases

Notes: Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

2.2.1 Phase 1 - Recruitment and value of Facebook

We sent an invitation email to recruit participants. The email contained a short description of the study and a link to an online survey that asked basic demographic information, determined if the participant had a Facebook account (95% did) and elicited the participant's value for not using Facebook for one week.[5]

An individual's value of Facebook is revealed with the Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964) and determines eligibility for participation in subsequent phases of the study. The participant is asked to submit her value of one week of Facebook usage. A random counter offer is drawn and shown to the participant. If the participant's value is less than the counter offer, then the participant is eligible for the next phases of the study and would be paid the counter offer upon study completion. If the participant's value is higher than the random offer, then she is not eligible to participate in any of the subsequent phases of the study and does not receive payment. Several examples of how the procedure works are included in the instructions to make sure that participants understand the procedure prior to submitting a value.[6] The examples

[5]The email text and online survey questions are in the Appendix, Sections A.1 & A.2.

[6]Our procedures made clear to participants that they would be paid the random offer upon study completion to

explicitly highlight that participants should optimally reveal their true value. To assure that reported values are not biased upwards, we follow the suggestion of Bohm et al. (1997) and Mazar et al. (2014) and leave the upper limit of the random offer unclear because that increases the validity of the BDM mechanism. This is implemented by informing participants that the minimum counter offer is $5 and the maximum is "our most reasonable estimate of the value of the time spent on Facebook."[7]

All eligible participants were invited by email to attend the next phase (Phase 2) on Monday of the following week.[8] The email explained that the next phase involves completing a comprehensive survey and being randomly assigned to log off Facebook for one week. In addition, the participants were informed that they would need to come back a second time (one week later) to complete another survey and receive cash payments of the counter offer they received. The time and location of the session is indicated in the email, and participants confirm their attendance.

2.2.2 Phase 2 - Pre-survey and Facebook restriction assignment

Participants were required to show up in person to complete a short survey that collects information on social media usage, news awareness, consumption behavior, time allocation, and subjective well being (Appendix Sections A.3 and A.4). The questions on social media usage included time spent, frequency of postings and emotions felt while using the platform.[9]

To capture news awareness, we tapped into a variety of news sources. In the week prior to the survey, we collected headlines from the front page of the eleven most popular newspapers as ranked by the Pew Research Center, including The New York Times, Washington Post, USA Today, Wall Street Journal, LA Times, New York Daily News, New York Post, Boston Globe, San Francisco

mitigate any uncertainty bias (Horowitz, 2006).

[7]For budgetary reasons and expected participation rates, the random counter offers were drawn with the following probabilities: (5, 15.14%; 7, 15.14%; 9, 11.14%; 10, 11.14%; 12, 11.14%; 14, 11.14%; 16, 7.14%; 18, 6.14%; 20, 5.14%; 21, 5.14%; 24, 0.64%; 25, 0.64%; 28, 0.14%; 30, 0.14%). The expected offer is $11.58.

[8]Those who are ineligible for subsequent phases are not contacted.

[9]Participants complete the survey in Phase 2, prior to random assignment to the Facebook restriction, and in Phase 3. One might be concerned that changes in outcomes are due to experimenter demand effects. First, participants are not aware they will complete the same survey questions a second time. Second, we find effects for some, but not most, of the outcomes, thus alleviating concerns of such an effect. Finally, in a later study with a similar design, Allcott et al. (2019) find similar results to ours while explicitly testing for demand effects.

Chronicle, The Chicago Tribune and The British Daily Mail. We used Breitbart as the source of skewed news.[10] There were no extraordinary news events during this period, like a mass shooting or major natural disaster, that might bias news knowledge. The participant is shown six headlines randomly chosen from the pool of mainstream sources and one randomly chosen from the skewed source and asked to identify if the event occurred or not. From the six mainstream sources, two headlines are changed slightly so as to make the headline false. All other headlines did appear on the front page of a newspaper or on Breitbart.[11]

Daily behavior is measured by presenting participants with a series of statements (e.g. "I save more money than I normally do", etc.) and asking them to identify on a scale of 1-5 whether they agree/disagree with the statement. Time allocation is measured with estimates for average time spent doing a variety of activities, such as working and exercising. Finally, our subjective well-being questions are constructed following the OECD Guidelines used to characterize the affective state of the respondent (OECD Better Life Initiative, 2013). These questions ask participants to respond on a 0-10 scale how frequently they feel a certain emotion (e.g. depression, happiness, etc.).

Upon completion of the survey, participants were randomly assigned to either a one-week Facebook restriction or no restriction based on the last digit of the participant's university-assigned ID number.[12] All participants complied with their assigned treatment and associated protocols.

The no restriction group is dismissed and asked to return the following Monday (one week later) to complete another survey and receive payment. The restriction group is required to log off of Facebook, and all its associated features, including Messenger, for one week. To validate compliance with the restriction, we created a Facebook account for the study and had treated participants become friends with our study account. As friends, we can monitor all access to their account through the "Last Active" feature in Facebook Messenger. This feature automatically up-

[10]We chose Breitbart given that its internet traffic as of March 2017 surpassed other major skewed news sources and was similar in magnitude to that of mainstream news sources such as The Washington Post according to data from alexa.com

[11]See the questionnaire in Appendix A.3

[12]The university randomly generates the last four digits of a student's ID number.

dates as soon as someone logs on to Facebook, thus we can validate if a participant complies or not with the restriction. A participant could go invisible, block or un-friend our Facebook account, but they would have to log in and we would observe this in our data. We saw no instances of this, and all participants complied with the restriction. After becoming friends with our Facebook account, participants logged off of all their active Facebook sessions on all their devices using Facebook's security settings. Finally, the restriction group was asked to return the following Monday (one week later) to complete another survey and receive payment.

2.2.3 Phase 3 - Post-survey and re-elicited value of Facebook

All participants returned one week later to complete another survey and receive payment. The survey is identical to the one given in Phase 2 and allows us to see how key indicators – social media use, news awareness and subjective well-being – have changed over the previous week.[13] After completing the survey, participants were instructed to go to a separate room for payment.

In the separate room, before receiving payment, we again elicited each participant's value for one week of Facebook usage. Up to this point, participants did not know they would again be asked their value of Facebook. This procedure gives us an unbiased measure of the change in Facebook's valuation following the restriction. We use the same BDM mechanism procedures as in Phase 1.[14] Afterwards, all participants receive a cash payment based on the counter offer from Phase 1 before leaving the session.

2.2.4 Implementation

Participants were recruited via email from a random sample of the undergraduate population at Texas A&M University during the Spring semester of 2017. Overall, 1,929 individuals initiated the Phase 1 online survey and 1,769 completed it, thus producing the distribution of stated values used to estimate the value of Facebook and to test if selection affects results. When we compare

[13] We updated the news pool to reflect headlines from the previous week.

[14] Participants are asked to write down their valuation and informed that their payment today is unaffected by their response. Eligible participants from this second BDM go through the same process as in Phase 2, return for a third and final survey in one week, and are paid their counteroffers from the second BDM. We do not include this third survey in our estimates.

the characteristics of the individuals who responded to the survey with the entire undergraduate population (based on year in school, home state and declared major), we find that our survey respondents are representative. Of those individuals who completed the Phase 1 survey, 562 were eligible for Phase 2 of the study, and eligibility does not depend on covariates.[15] Also, we find no evidence that participants who ended up being eligible or ineligible based on the randomly-drawn counter offer are different.[16]

All eligible participants were invited to Phase 2 of the study, and this session was held on main campus where participants came to complete the survey and be randomized into the Facebook restriction.[17] For the Phase 2 sessions, 167 participants showed up and completed the survey. Appendix Table A.1 shows the comparison between those who were eligible and showed up and those who did not. The only meaningful differences are that those who did not show up had a slightly lower value for Facebook and counteroffer. We test the robustness of our results to design-induced selection in Section 2.4.2 by re-weighting the sample.

Among the participants who completed the Phase 2 survey, fifty-four percent (n=90) were randomly assigned to the no restriction control group, and 46% (n=77) were assigned to the Facebook restriction treatment group. Comparing covariates of the control and treatment groups, we find there are significantly more women in the control group (71%) compared to the treatment group (57%), but otherwise, the two groups are balanced.[18] To address covariate differences by treatment assignment, our analysis controls for individual fixed effects so that treatment effects are identified through differences in changes in behavior before and after the one-week Facebook restriction across the treatment and control groups.

After one week of treatment, 90% (n=151) of the participants from Phase 2 returned to com-

[15]Eligibility for Phase 2 means that the submitted value was less than a randomly-selected counter-offer of no more than \$30. This is by the design of the elicitation mechanism – so all those with submitted values higher than \$30 were ineligible. Descriptive statistics for these groups are in Appendix Table A.1. In Section 2.4.2 of the paper, we test the robustness of the results to this design-induced selection.

[16]When we compare participants who submitted values less than or equal to \$30, so they could have been eligible to participate in Phase 2, there is no significant difference by age or gender between those who ended up being eligible or ineligible based on the counteroffer. See Appendix Table A.1

[17]Participants were aware of this procedure prior to submitting their value of Facebook in the Phase 1 online survey. Holding this session on main campus minimizes travel costs that might have affected valuations for Facebook.

[18]Appendix Table A.2 shows the balance of covariates across the treatment and control groups.

plete the Phase 3 survey. There is no significant difference in covariates between the participants who returned for Phase 3 and those who did not, and attrition is not correlated with treatment status. Our monitoring process validates compliance with the restriction.[19] Those in the treatment group reduced their use of Facebook by 1.7 hours per day. Given a baseline Facebook usage of 1.9 hours per day, this illustrates that the treatment group complied with the restriction.

All sessions were completed in April-May 2017. Time to complete the Phase 1 online survey was approximately five minutes, and each subsequent in-person survey took about 10-15 minutes. Average payment to participants was $16.79 (s.d. $5.22) at the completion of Phase 3.

2.3 Results

2.3.1 Description of the sample

In the baseline survey (Phase 2), participants report spending a mean of 1.9 hours per day on Facebook, including reading news feeds and news content (Figure 2.2, panel a). This is consistent with other surveys with college students that report an average of 2.6 hours spent on Facebook per day (EMarketeer, 2015), yet higher than the national average of 50 minutes per day (Neilsen Company, 2016). Engagement on Facebook is measured by how often participants post pictures and comment. This activity was rated on a scale of 1 (never) to 7 (several times per day). About 52% never or rarely post pictures, 28% once or twice a month and the remainder post once a week or more (Figure 2.2, panel b). In terms of posting comments, 48% never or rarely comment, 18% once or twice a month and the remainder post once a week or more (Figure 2.2, panel c).

Other social media platforms are also used. On a daily basis, participants report spending close to two hours on Facebook, Snapchat and YouTube, over one hour on Instagram, less than one hour on Twitter, and very little on Tumblr and Vimeo.[20] This is consistent with the number of friends and followers reported across platforms. On average, there are more friends and followers on Facebook (641) and Instagram (452) than on Tumblr (87) and Twitter (182).

Information is also collected on where participants get their news and time spent acquiring

[19] Participants did not interact with the study account in any way.
[20] Appendix Table A.3.

Figure 2.2: Time Spent on Facebook and Facebook Usage

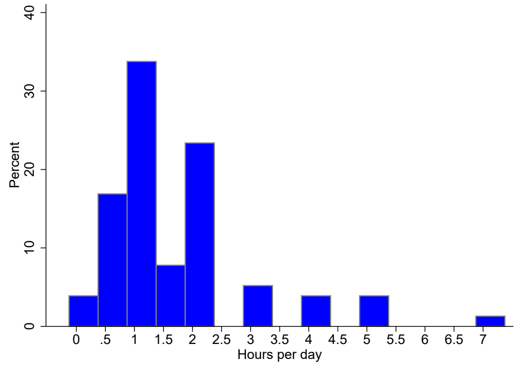

(a) Hours per day on Facebook

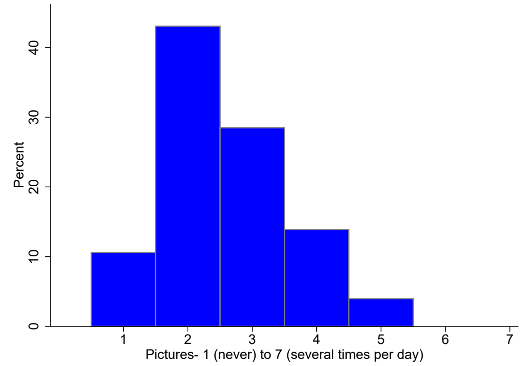

(b) Post pictures

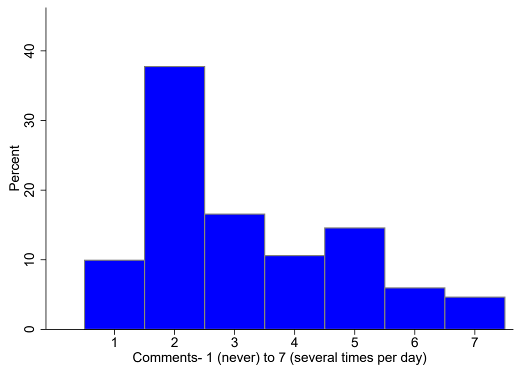

(c) Post comments

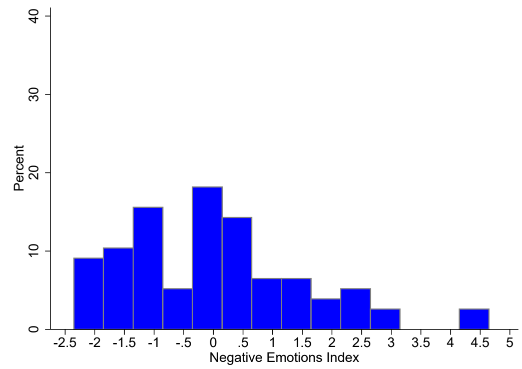

(d) Negative Emotions Index

Notes: This figure presents descriptive statistics on Facebook usage. The x-axis in panels (b) and (c) represents: 1 never, 2 rarely, 3 1-2 times per month, 4 once a week, 5 2-4 times per week, 6 once a day, 7 several times per day. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

news. Roughly, 15-30 minutes a day is spent reading or watching news, and most news is obtained from digital sources (e.g. online news, social media) as opposed to traditional outlets (e.g. cable tv, paper news, radio).[21] Participants reported their preferred news sources, and we rank each source's political bias on a scale of 1 (Left) to 5 (Right).[22] The average preferred news source has a political

[21] Appendix Table A.3. While we cannot say what proportion of news participants get from Facebook, 81% report opening up Facebook every day or several times a day to check their news feed.

[22] We use the rankings on www.allsides.com. If a participant lists a news outlet that is not reported on allsides.com, we treat their preferred news outlet as missing. The top five first choice sources are CNN (28.1%), FOX (12.6%), BBC (8.3%), NYT (4.7%), and ESPN (4.7%). Breitbart was not listed as a first choice, however, news from this source could appear on a Facebook news feed.

bias of 2.8 - slightly left of Center.

We further asked a variety of subjective wellbeing questions. On a scale of 0 (Never) to 10 (Very/Always), participants are generally satisfied with life (mean of 7.2) and responded with a mean of 3.4 to feelings of depression. These results are in line with the OECD's Better Life Initiative Survey for 2017 which reports an average overall life satisfaction score of 7.3.

Participants were asked to rate on a scale of 1 (never) to 6 (all the time) how often they felt certain negative emotions while using Facebook, such as envy/jealousy, loneliness, misery and annoyance. To generate a general measure of experiencing negative emotions while on Facebook, we take these four measures and combine them into a factor index that ranges from -2.35 to 4.37 using principal component analysis. A higher index indicates a participant feels more negative emotions (see Figure 2.2, panel d), and there is large variation in this index.[23]

2.3.2 Value of Facebook

Participant responses to the BDM lottery show that one week of Facebook usage is valued at \$24.84 on average ([23.02, 26.65] 95% confidence interval), and the median value is \$15 ([12.70, 17.30] 95% confidence interval).[24] We evaluate how sensitive the mean is to outliers by trimming the distribution at \$200, \$100 and \$50. With each cut, the mean BDM value changes to \$22, \$21 and \$18, respectively. The median BDM value remains fixed at \$15 with each cut of the distribution.[25] There is no bunching at \$5 which indicates that participants did not try to manipulate the BDM mechanism to be eligible for the next stage of the study.

Our experiment introduces a lottery in which an individual has a 50% chance of being restricted from Facebook. Given that the restriction is experienced half of the time, stated values could be dampened and the BDM would then produce an underestimate of Facebook's value. If we assume

[23]Appendix Figure A.1 shows the distribution of these emotions separately.

[24]We calculate the confidence intervals using bootstrap with 1000 replications.

[25]Our design also explored the willingness to pay (WTP) - willingness to accept (WTA) gap in the BDM mechanism (see Knetsch et al. (2001), Plott and Zeiler (2005), Horowitz (2006), and Brynjolfsson et al. (2018) for a discussion of this phenomenon). Half of the participants were asked the value in terms of selling participation in the study (WTA), "How much money would you need to be given to stop using Facebook for a week?" and half were asked in terms of purchasing participation (WTP), "What is the value of your weekly time on Facebook?" We find no significant difference in the reported value of Facebook from either solicitation method or by covariates across groups, so we pool the data in our analysis.

that stated valuations are half of the truth, then under risk neutrality, the mean value of Facebook would be $50 per week (median=$30) and $200 per month (median=$120). If individuals are risk averse and we assume a CRRA utility function with a risk aversion parameter within a reasonable range (0.1-0.3), then the mean value of Facebook would be $67 per week (median=$40) and $267 per month (median=$160). Throughout the remainder of the paper, we report values adjusted for risk aversion. However, results are qualitatively the same if we use the unadjusted reported values from the BDM mechanism. Figure 2.3 shows the distribution of the risk-adjusted values.[26] While our design does not separately consider hassle costs, other studies find similar values to ours, suggesting that hassle costs are minimal. Brynjolfsson et al. (2018) and Corrigan et al. (2018) find lower weekly median values ($3.92 and $15 respectively), and Allcott et al. (2019) find median monthly values ($100-$180) similar to ours.[27]

According to the Pew Research Center (2016), women are 8 percentage points more likely to use Facebook than men. Hence, we might expect to see differences in the value of Facebook across genders, however, we do not find a statistically significant difference. On average, one week of Facebook is worth $69.35 for men (median=$43.07) and $65.18 for women (median=$40.38). We also test for difference in the distributions of the value by gender and find no significant difference.[28]

There is a positive correlation between the value of Facebook and age in our data. For those aged 21 years or younger, one week of Facebook is worth $62.95 (median=$40.38), while for those older than 21 years, Facebook is worth $78.37 (median=$53.84). This could reflect differences in income or that younger participants are more likely to use other social media. Indeed, those 21

[26]The distribution is trimmed at $540 because of a few outliers in the data – the maximum value is $2,153. We use the nontrimmed, full sample in our analysis.

[27]There are differences across studies. Brynjolfsson et al. (2018) use an online sample, one out of every 200 participants are randomized into the Facebook restriction, and respondents who do not use Facebook are not screened out for their weekly estimate. Corrigan et al. (2018) use a series of second-price auctions with different samples and compensation schemes. Allcott et al. (2019) also use a BDM mechanism but with an online sample.

[28]Women are typically found to be more risk averse than men. The risk-adjusted values of Facebook that we use assume that men and women have the same level of risk aversion. Women would need to be 37% more risk-averse than men for the difference to be significant at the 10% level, 41% more risk-averse than men for the difference to be significant at the 5% level, and 48% more risk-averse than men for the difference to be significant at the 1% level.

years and younger spend more time on Twitter and Snapchat and have more Instagram followers.[29]

The value of Facebook changes across user types, with those who are more active reporting higher values. Facebook is worth 20% more for participants who use it for more than one hour a day and for those who post at least once per month. There is a positive, but not significant, correlation between the value of Facebook and having a large number of friends on Facebook, however, there is a positive and significant correlation between the value and having a large number of friends on other social media platforms. Those with a large number of friends on other social media also have a lot of friends on Facebook, so this likely reflects the larger value that active Facebook users place on using the platform. There is a negative correlation between feeling depressed or experiencing negative emotions while on Facebook and the value of Facebook, but these correlations are not significant.[30]

To put some perspective on the magnitude of the stated values of Facebook in our sample, we compare its value with college students' mean income and some common expenses. The weekly average income of a college student is \$224.28 (Flood et al., 2017), so a week of Facebook usage is worth 30% of income.[31] In addition, university students spend roughly \$14 in clothing, \$14 in personal care and \$11.50 in technology (devices, plans and subscriptions) per week. Facebook is worth more than each of these and more than the average weekly expenditure of \$20 on coffee (Tuttle, 2012). Facebook has a large value for our participants relative to their income and other purchases.

2.3.3 Effects of the Facebook Restriction

We explore the effect of not using Facebook for one week on five outcomes: social media usage, news consumption, news awareness, subjective well-being, daily activities and the value of

[29]We did not ask questions on income but asked the zip code of where the participant lived at age 15. Using income data from this zip code, we find no significant difference in mean income for younger participants compared to older.

[30]Appendix Table A.4 presents the Pearson correlation coefficients between the value of Facebook and several measures that characterize Facebook users.

[31]In-state tuition at Texas A&M is \$11,200 per year, or \$350 per week, implying that participants value Facebook as much as 19% of the weekly cost of studying at the university. According to the College Board, the average university student in the U.S. spends \$225 per week (\$10,800 per year) on room and board. Facebook is then worth 30% of these expenses.

Figure 2.3: Distribution of the Value of Facebook (trimmed at $540)

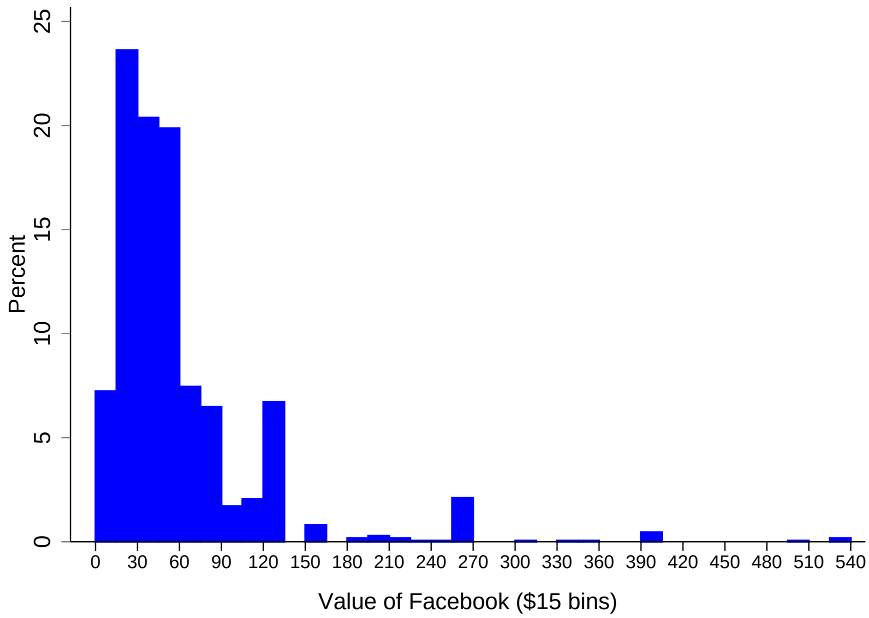

Notes: This figure presents the distribution of the Value of Facebook. We adjust user's reported BDM value using a CRRA utility function in order to address the fact that our experiment introduces a lottery in which there is a 50% chance of being restricted from Facebook. Distribution is trimmed at $540. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Facebook. Throughout the paper, indices are constructed using the procedure of Anderson (2008). We demean each variable using the mean of the control group in Phase 2 and convert it into an effect size by dividing it by the standard deviation of the control group in Phase 2. The index is the weighted average of the transformed outcomes, where the weights are derived from the inverse of the covariance matrix of the transformed outcomes. A key advantage of our design is that we can verify that participants assigned to the Facebook restriction remained logged off without having to directly contact participants with reminders and possibly affect their behavior. Our compliance rate is 95%, and throughout the paper we report intent-to-treat effects.[32]

To examine the effects of Facebook on behavior, we exploit the fact that we ask the same ques-

[32] All but three treated participants stayed off of Facebook for the entire week. The three who did log back into Facebook did so only once for less than an hour to communicate for a student organization via the organization's Facebook account. All three participants contacted the research team prior to logging in to inform us why they were logging back on. These participants are included in our intent-to-treat analysis. Instrumental variable estimates are 5% larger and slightly less precise.

tions in the pre and post-treatment surveys (administered in Phase 2 and Phase 3) and estimate the change in the outcome of interest and control for individual fixed effects. This approach identifies treatment effects based on changes in individual behavior and controls for any unbalancedness that might exist in covariates across the treatment and control groups. By relying on within-individual variation to identify effects, the only difference across individuals is random assignment to treatment and control.

Specifically, we estimate the following equation:

$$y_{it} = \beta_0 + \beta_1 PostSurvey_t + \beta_2 PostSurvey_t \cdot Treatment_i + \alpha_i + \varepsilon_{it} \qquad (2.1)$$

where $PostSurvey_t$ is a dummy variable for the survey given in Phase 3 after the one-week Facebook restriction and $Treatment_i$ indicates if individual i is randomly assigned to the Facebook restriction group. β_2 is our coefficient of interest. Individual fixed effects are included and thus control for treatment assignment and fixed individual covariates. Standard errors are clustered at the individual level. We estimate equation (2.1) for the full sample and explore heterogeneous effects by gender and different classifications of Facebook users.

In addition to testing differences in means, we test whether Facebook usage has an effect on the distribution of outcomes. We test for equality of the distributions, as well as first and second order stochastic dominance.[33]

Our analysis tests for effects on a large number of outcomes. To make sure that our results are not due to chance, we adjust the p-values to account for multiple comparisons and report these as our main findings.[34] We apply the procedure defined by Benjamini and Hochberg (1995) and Benjamini et al. (2006).[35]

[33]It would be important to test for effects at different quantiles, but we do not have enough power to estimate meaningful comparisons at the tails of the distribution. To test for distribution equality, let $F_{(1)}$ be the distribution of outcome y_{it} for the treated group and $F_{(0)}$ be the distribution of the control group. According to Abadie (2002), we define $F_{(1)}$ first order stochastic dominates $F_{(0)}$ if $\int_0^x dF_{(1)}(y) \leq \int_0^x dF_{(0)}(y)\ \forall x \geq 0$ and $F_{(1)}$ second order stochastic dominates $F_{(0)}$ if $\int_0^x \left(\int_0^z dF_{(1)}(y) \right) dz \leq \int_0^x \left(\int_0^z dF_{(0)}(y) \right) dz\ \forall x \geq 0$

[34]Doing this involves a trade off between a Type I error and the power of the test (Anderson, 2008). We control for the false discovery rate to adjust our p-values and achieve a balance between these two factors.

[35]For reference, both the unadjusted and adjusted p-values are reported in Table 2.1. All of our results remain

2.3.3.1 News Awareness

According to Gottfried and Shearer (2016), 64% of social media users access news from just one site, and on Facebook, 66% of users report getting at least some news while using the platform (Pew Research Center, 2016). This suggests that Facebook might play an important role in the distribution of news. If this is true, we should expect that logging individuals off Facebook for a week decreases awareness of current events. We use the news headlines quiz described in Section 2.2.2 to define three indicators that measure the effect of Facebook usage on news awareness: the proportion of news headlines participants correctly recognized as having occurred, the proportion they got wrong and the proportion for which they were uncertain (i.e. they answered "I don't know"). We calculate these measures for the questions from mainstream sources (six questions) and for the skewed news source.

Figure 2.4 shows the effect of the Facebook restriction on these three measures for mainstream and skewed sources. There is no significant effect of the restriction on news awareness for headlines from mainstream sources.[36] However, there is significant uncertainty of the veracity of headlines from skewed news. Those who experienced a week off of Facebook are 22.1 percentage points more likely to be uncertain about whether or not a politically-skewed news headline is true or not. And, they are 15.6 percentage points less likely to answer correctly if the event actually occurred.[37]

---

statistically significant at the 5% level or less, with the exception of the probability of answering "Don't Know" for skewed news, the healthy activities index and the change in the value of Facebook. We also do a more robust adjustment controlling for the family-wise error rate. When we use the free step-down method described by Anderson (2008), only the effects on Facebook use, news access through social media, news consumption and the correct answer of skewed news are statistically significant at conventional levels.

[36] We tested whether the Facebook restriction had different effects for true headlines and the false headlines we created (by changing a few words) in the news quiz. For both types of headlines, the point estimates are similar to the main results and statistically insignificant.

[37] Gender differences do emerge. While both men and women are less likely to be aware of the veracity of skewed news when off of Facebook, the effect is much stronger for men than women. This suggest that men, more than women, are exposed to politically skewed news when on Facebook.

Figure 2.4: Effects on News Awareness

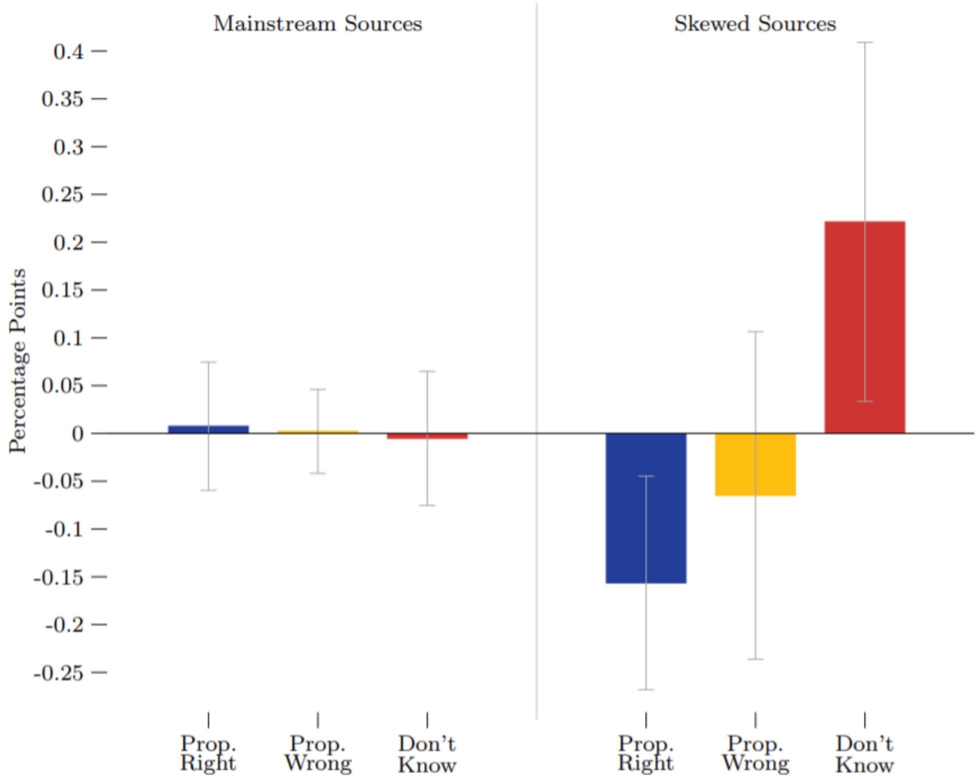

Notes: This figure presents the intent to treat effects of the Facebook restriction on news awareness. Prop. Right corresponds to the proportion of questions answered correctly on the news headlines quiz, Prop. Wrong corresponds to the proportions of questions answered incorrectly and Don't Know corresponds to the proportion of questions answered "I don't know". Estimates control for individual fixed effects. Each estimate corresponds to the change in the proportions of answers in each category. The figure displays the 95% confidence interval. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

2.3.3.2 Potential Mechanisms for the Reduction in News Awareness

The reduction in news awareness should be correlated with an overall decrease in access and consumption of news. We analyze how being logged off Facebook for a week affects the frequency with which individuals access different news media and whether consumption of different types of news changes. Participants reported their answers for news consumption and types of news using a Likert scale ranging from "not at all" (1) to "all the time" (7). Following the procedure described in

Figure 2.5: Effects on News Media Access

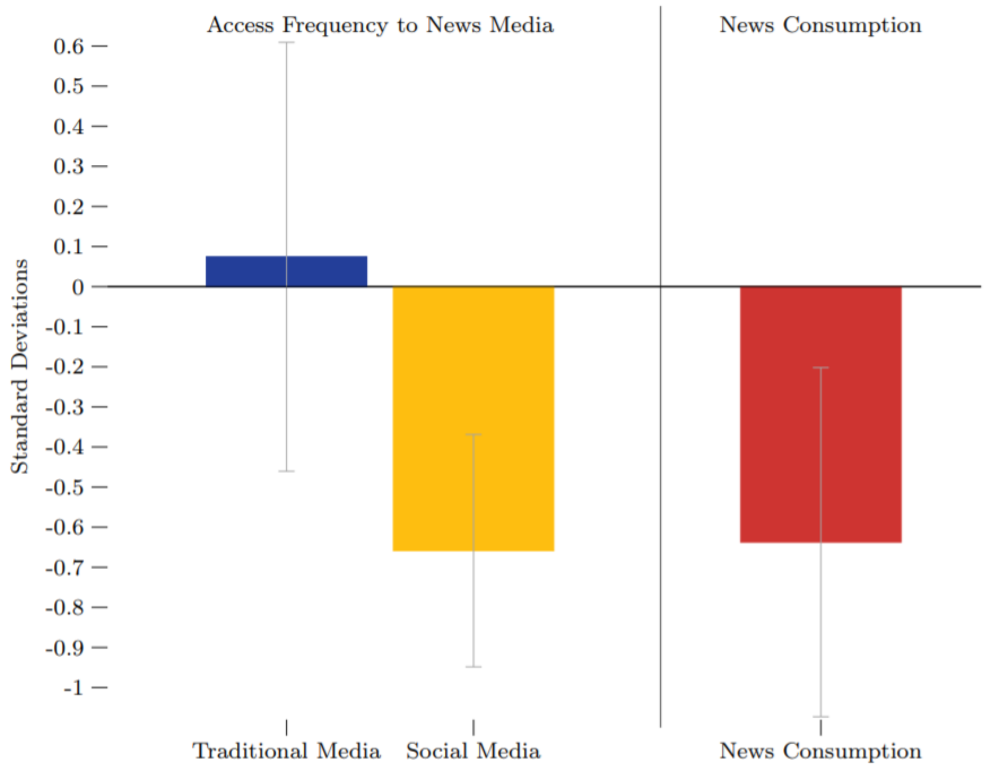

Notes: This figure presents the intent to treat effects of the Facebook restriction on access to two types of news media. Traditional media is an index that measures access to "traditional" news media (i.e. radio, newspapers, television and Internet sites). Social media is an index that measures access to news through social media and news feeds. Estimates control for individual fixed effects. Each estimate corresponds to the change in access frequency of a type of media. The figure displays the 95% confidence interval. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Section 2.3.3, we aggregate access to "traditional" news media (i.e. radio, newspapers, television and Internet sites) in one index (Traditional Media) and access to social media and news feeds into a second index (Social Media). We use the two indices to measure changes in access to news media.

The left panel in Figure 2.5 presents the effect of the Facebook restriction on access frequency to news media. On average, access to news through social media decreases by 0.66 standard deviations (significant at the 5% level), while there is no statistically significant change in access to "traditional" news media. These results are consistent with the fact that participants in the restriction group reduced their Facebook usage to zero but they do not substitute by increasing use of traditional media.[38] We also find that the distribution of the social media index for the restriction group first order stochastic dominates the distribution of the non-restricted group. This indicates that access to news through social media decreases not only at the mean, but throughout the distribution (see Appendix Table A.5). We find no distribution differences for access to "traditional" media. These results indicate that Facebook is an important source of news for our participants, and in the short term, they do not substitute with other news sources.

The right panel in Figure 2.5 presents the effect of the Facebook restriction on news consumption. We asked how frequently the participants read political, business, sports, international, culture, science, local and weather news, and we aggregate these measures into an index (News Consumption) to capture overall news consumption. On average, participants in the Facebook restriction group significantly decrease their consumption of news by 0.64 standard deviations with respect to the baseline (p-value < 0.05), and this effect is consistent across all news types. The reduction in consumption of news decreases not only at the mean but also across the entire distribution (see Appendix Table A.5.

In summary, these results indicate that Facebook is an important conduit for news awareness, specifically from skewed sources, for college students. News consumption decreases and there is no evidence of substitution to other news sources. In the next section we study the effects of Facebook on subjective well-being.

[38]Our research design restricted usage of Facebook for those in the treatment group, but participants were not restricted in their usage of other social media platforms. We validate that those in the treatment group did reduce their use of Facebook – by 1.7 hours per day. Given a baseline Facebook usage of 1.9 hours per day, this illustrates that the treatment group did comply with the restriction. While the treatment group refrained from using Facebook, we find that they did not increase their usage of other social media (e.g. Instagram, Snapchat, Tumblr, Twitter). This is consistent with studies finding low cross-platform usage for social media and a significant cost to switch to alternatives for one week (Pew Research Center, 2016). Only one-third of Facebook users are active on other social media platforms, yet about 90% of users of other platforms are active on Facebook (Pew Research Center, 2016).

2.3.3.3 Subjective Well-being

Previous studies have found mixed results on the effects of Facebook on happiness and well-being. We build on previous research by applying a validated Facebook restriction that does not interfere with participants during treatment, and by including a series of questions on daily habits and activities potentially correlated with well-being (Salovey et al., 2000; Ostir et al., 2000; Fredrickson and Joiner, 2002; Blake et al., 2009; Kettunen et al., 2015; Newman et al., 2014; Sonnentag, 2001).

We asked participants five subjective well-being questions (taken from the OECD Better Life Initiative) using a Likert scale (from 0-10). The questions assess overall life satisfaction, how worthwhile life is, happiness, level of worry, and depression.[39] Figure 2.6 presents the effects of the Facebook restriction on these measures. Estimates for overall life satisfaction, life is worthwhile, happiness and worry are small and statistically insignificant.[40] However, being off of Facebook does significantly reduce depression by 17% (0.57 points on the Likert scale). This result is consistent with findings from the social psychology literature using cross-sectional data that shows Facebook increases feelings of depression (Steers et al., 2014 and Feinstein et al., 2013).[41] We do not find evidence of distribution shifts (see Appendix Table A.5).

Our results suggest that using Facebook induces feelings of depression. While this could plausibly decrease an individual's well-being, our estimates reject significant changes in well-being. Evidence of a negative correlation between happiness and depression is weak (Rezaee et al., 2016), hence, a significant decrease in depression is not inconsistent with no change in well-being.

[39]The questions are: (i) Overall, how satisfied are you with life as a whole? (ii) Overall, to what extent do you feel that things you do in your life are worthwhile? (iii) How happy are you? (iv) How often do you worry? and (v) How often do you feel depressed? An alternative approach could have been to use the Day Reconstruction Method (Kahneman et al., 2004 and Kahneman and Krueger, 2006), however, to keep the survey short, we opted for the five OECD questions.

[40]Our results on life satisfaction are smaller than Tromholt (2016) who finds a significant effect of 0.26 standard deviations. The study's Danish sample is older (average age of 34 years) compared to our U.S. sample (average age of 20 years), and participants were contacted daily by the researcher team to follow their assigned treatment status.

[41]Subjective well-being measures can be sensitive to temporary events (e.g. the weather, long lines at a coffee shop, meeting somebody) (Krueger and Schkade, 2008), nonetheless, because our participants are randomly assigned to treatment, random shocks should be evenly distributed and our panel estimation allows us to directly control for events that affect both groups uniformly across time.

Figure 2.6: Effects on Subjective Wellbeing

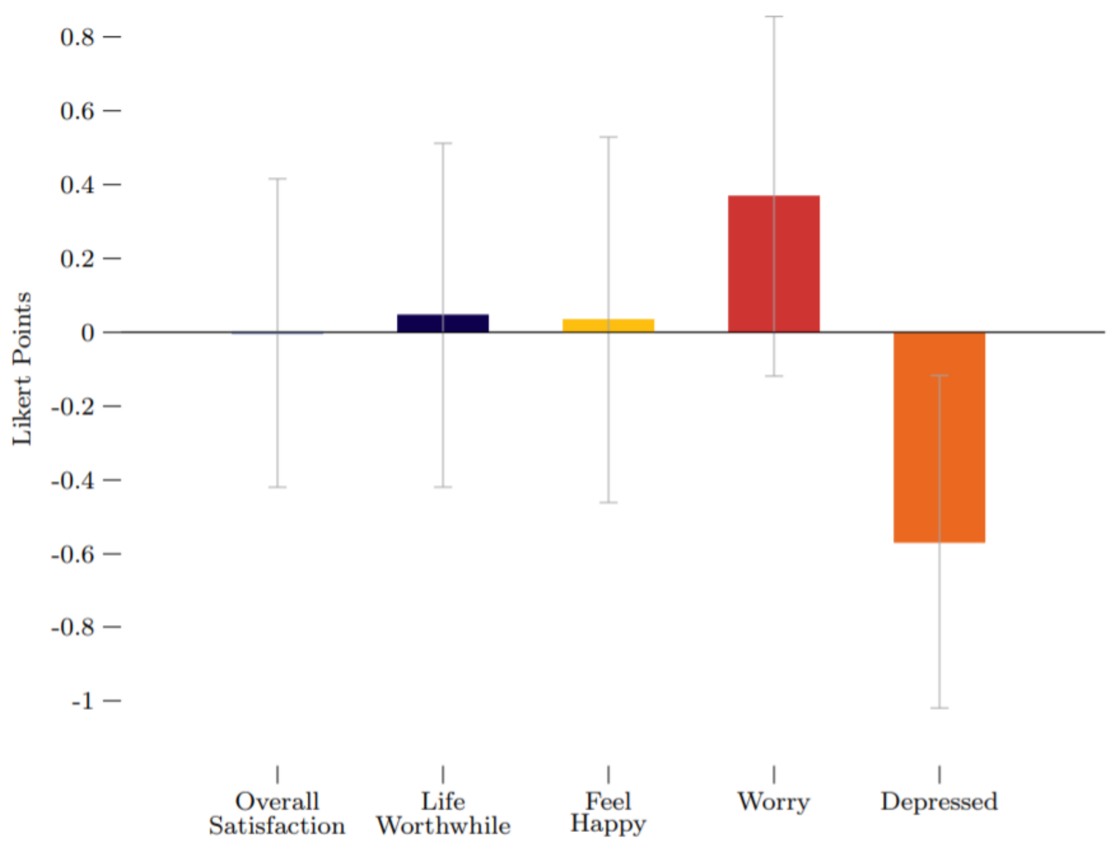

Notes: This figure presents the intent to treat effects of the Facebook restriction on five different measures of subjective well-being. Estimates control for individual fixed effects. The figure displays the 95% confidence interval. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

The reduction in feelings of depression from being logged off of Facebook could be driven by changes in behavior. To shed light on how people respond to losing Facebook access, we asked participants to report on a variety of activities the week prior to completing the pre and post-treatment surveys (Phases 2 and 3). Healthy behavior was measured by asking whether participants ate out less than usual, did less impulse buying, saved more money, ate healthier and exercised more.[42] We also asked what they expected their behavior would be the following week. Productive

[42]There is evidence that eating out is associated with excessive calorie intake (Urban et al., 2016), a less healthy diet (Wolfson and Bleich, 2015), increased hypertension (Seow et al., 2015) and a higher exposure to phthalates (Varshavsky et al., 2018), which have been linked to asthma, breast cancer, type 2 diabetes and fertility issues. Diet is

Figure 2.7: Effects on Activities and Time Use

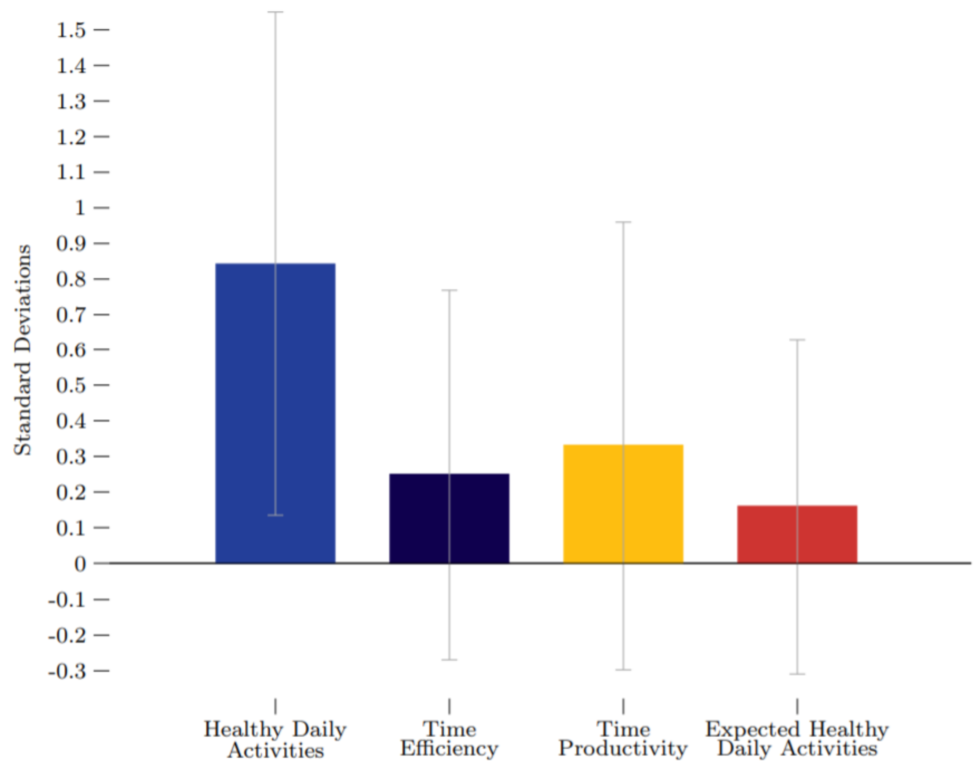

Notes: This figure presents the intent to treat effects of the Facebook restriction on four index measures of activities and time use. Healthy Daily Activities indexes engagement in "healthier" consumption/savings practices in the past week. Time Efficiency measures efficient time use. Time Productivity measure productive time use. Expected Healthy Daily Activities indexes the expected engagement in "healthier" consumption/savings practices the following week. Estimates control for individual fixed effects. The figure displays the 95% confidence interval. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

time use was measured by asking whether they spent more time studying, had time to relax and be with friends, and partied a lot. Time efficiency was measured by whether they wasted less time, achieved more than usual, were not late for class, were able to meet deadlines, were able to prevent distractions, discontinued wasteful activities, and procrastinated less.[43] Again, we use the procedure in Section 2.3.3 to aggregate these four categories of questions into four indices: healthy

correlated with an individual's mental health (O'Neil et al., 2014).

[43]Participants were asked on a scale 1-5 to what extent they agreed with a particular statement, where 1: Strongly Agree, 5: Strongly Disagree. We adjust the coding so a higher value indicates a "healthier" response.

daily activities, time efficiency, time productivity and expected healthy daily activities.

Figure 2.7 reports the effects of the one week Facebook restriction on these four measures. Overall, we find suggestive evidence that people behave in a healthier manner. Healthy daily activities increase by 0.86 standard deviations with respect to the baseline p-value=0.057). We find positive, but not statistically significant changes for the other indices. There are no significant effects on the distributions (see Appendix Table A.5).

In summary, a one-week Facebook restriction decreased feelings of depression and increased engagement in healthier activities. While we are not able to pinpoint the exact mechanism, these results suggest that Facebook can negatively affect components of daily life that go beyond any existing benefits of social media.

2.3.3.4 Change in the Value of Facebook

Being off Facebook for one week decreases news awareness and consumption, improves well-being by decreasing feelings of depression and promotes healthier behavior. If participants internalize these changes, we would expect a change in individuals' value of Facebook. Figure 2.8 shows the distribution of values for the restricted and unrestricted groups for those who completed the pre and post-surveys (Phases 2 and 3). Experiencing a week-long Facebook restriction increases the value of Facebook by 19.6% from $30.13 to $36.04, however, this effect should be interpreted cautiously given that we are not powered to detect significant results.[44] We find no significant distributional treatment effects (see Appendix Table A.5).

There are several potential explanations for this increase in value. First, the reduction in access to news may simply not be compensated by a better mood and healthier activities. Individuals would then need a higher payment to be willing to be off of Facebook for another week. Second, the increase in value is consistent with withdrawal effects of an addictive good.[45] If being on Facebook creates addiction, then the week-long restriction should increase the desire to be back

[44] The adjusted p-value is 0.125. Our sample size allows us to detect effects up to 0.182 percentage points at the 5% level with a power of 80%.

[45] A key characteristic of an addictive good is that its consumption exhibits "adjacent complementarity" (Becker and Murphy, 1988, and Gruber and Köszegi, 2001), which means that past consumption increases the marginal utility of present consumption.

Figure 2.8: Distribution of the Value of Facebook after Treatment

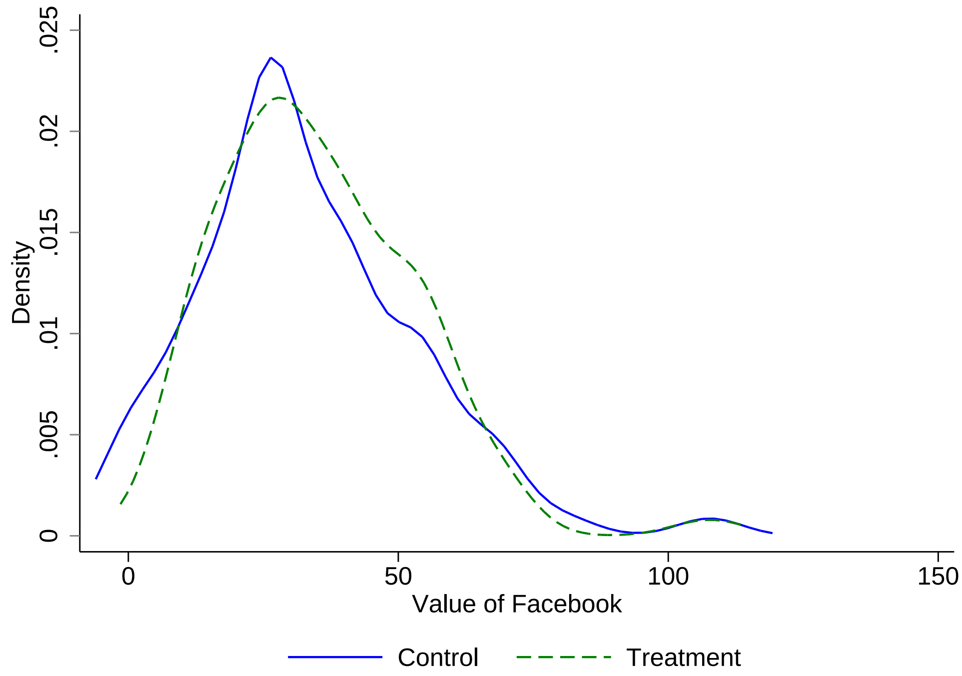

Notes: This figure compares the distribution of the value of Facebook after a one-week Facebook restriction for the participants who attended both Phase 2 and Phase 3. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

on Facebook. This would also explain the rise in value of Facebook. Third, Facebook further affects other dimensions of daily life that were not captured in our study. For instance, we do not measure the effects of losing access to Facebook's messenger service. These aspects along with their interactions may be utility increasing, which could explain the increase in value for an additional week off of Facebook.[46]

[46] Appendix Figure A.2 shows that while the level of depression in the treatment group has decreased relative to control group, there is no evidence that suggests that treated participants are internalizing this benefit by lowering their value for Facebook.

2.4 Robustness Checks

2.4.1 Multiple Comparisons

Our analysis thus far tests for effects on a large number of outcomes. To check that our main findings are not due to chance, we adjust the p-values to account for multiple comparisons.[47] We apply the procedure defined by Benjamini and Hochberg (1995) and Benjamini et al. (2006), and the results are shown in Table 2.1. We see that all of our results remain statistically significant at the 5 percent level or less, with the exception of the probability of answering "Don't Know" for skewed news, the healthy activities index and the change in the value of Facebook.[48]

2.4.2 Sample Selection

Our approach of recruiting volunteers to log off Facebook may induce selection by over-sampling low-value participants. To address this, we use the distribution of the stated BDM value of Facebook to re-weight the sample using the inverse probability of being eligible to participate in Phase 2 conditional on the stated value. Table 2.2 presents these results. Columns 1 and 2 show that the results pertaining to news awareness and news consumption remain and are robust to sample selection. The point estimates are robust to re-weighting the sample, although the weighted estimates are less precise, suggesting incremental power issues due to re-weighting. The point estimate of the effect on depression decreases from 0.57 (17% of baseline) to 0.39 (11% of baseline) Likert points and loses statistical significance. The same happens to the effect on daily activities. The point estimate decreases from 0.84 (17% of baseline) to 0.69 (11% of baseline) standard deviations.

This analysis suggests that the results on news consumption and awareness are robust to sample selection and representative of the broader population of college students. Conversely, the results on depression and daily activities speak to the population of college students who report having a

[47]To do this involves a trade off between a Type I error and the power of the test (Anderson, 2008). We control for the false discovery rate to adjust our p-values and achieve a balance between these two factors.

[48]We also do a more robust adjustment controlling for the family-wise error rate. When we use the free stepdown method described by Anderson (2008), only the effects on Facebook use, news access through social media, news consumption and the correct answer of skewed news are statistically significant at conventional levels.

Table 2.1: Adjustments for Multiple Comparisons

	Unadjusted P-value	FDR Adjusted P-value
Facebook Use	0.000***	0.000***
News Media Index - Traditional Media	0.785	1.000
News Media Index - Social Media	0.000***	0.000***
News Consumption Index	0.004***	0.027**
Probability Right Answer - Mainstream News	0.826	1.000
Probability Wrong Answer - Mainstream News	0.926	1.000
Probability Not Sure Answer - Mainstream News	0.885	1.000
Probability Right Answer - Skewed News	0.006***	0.030**
Probability Wrong Answer - Skewed News	0.458	0.723
Probability Not Sure Answer - Skewed News	0.022**	0.052*
Overall Satisfaction	0.993	1.000
Life is Worthwhile	0.845	1.000
Feel Happy	0.893	1.000
Worry	0.139	0.228
Depressed	0.014**	0.048**
Consumption Index	0.020**	0.057*
Productive Time Index	0.302	0.499
Efficient Time Index	0.346	0.530
Expected Consumption Index	0.504	0.743
Value of Facebook	0.068*	0.125

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

This table shows how the significance of the main results changes when we control for multiple comparisons. The table present the unadjusted p-values of our main estimates (Column 1) and their corresponding values adjusted for multiple comparisons (Column 2). We apply a false discovery rate control as described in Anderson (2008). Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

BDM value of Facebook up to $30 per week (84.4% of the student population who uses Facebook).

2.4.3 Gender differences

There is evidence to suggest that men and women use Facebook for different purposes and with different frequencies. According to the Pew Research Center (2018) report, more women (74%) use Facebook than men (62%). Women are more likely to use it daily (69%) than men (54%) (Statista, 2018), and they post more comments and pictures and send more messages (Muscanell and Guadagno, 2012). This is also evident in our sample.[49] These differences in Facebook usage may imply heterogeneous responses to the Facebook restriction.

Splitting our sample by gender, Table 2.2 shows that for men one week off Facebook decreases feelings of depression by 0.82 Likert points, which increases to 0.90 after re-weighting. Both are statistically significant at the 5% level. There are no significant effects for women. While the point estimate of the effect on healthy daily activities decreases from 0.84 to 0.69 standard deviations in the full sample, losing statistical significance, the effect remains large and significant at the 10% level for men. Both the weighted and unweighted results show that the group, in this case men, that is less depressed also engages in healthier activities, confirming the influence of Facebook on other aspects of daily life. This is also consistent with findings that men are more likely to feel depressed due to negative social comparisons (Steers et al., 2014).

Our finding on the reduction in awareness of skewed news is supported by the behavior of men. They are significantly less likely to be certain about the veracity of skewed news both in the weighted and unweighted samples, and women are unaffected. Women reduce their consumption of news via social media, as do men, but are otherwise not significantly affected by the Facebook restriction.

There is an increase in the value of Facebook after the restriction. This is driven by women. They significantly increase their value by 33%, which decreases to 19% after re-weighting.

2.5 Conclusions

Social media and Facebook have become entities of global proportions. However, we know little about their economic value to users, the effects on daily activities, consumption behavior and news awareness. Using a randomized, and validated, Facebook restriction in a large field experiment, we provide an estimate of an individual's value of Facebook. One week on Facebook is worth about $67 for our participants – a relatively large value considering that it represents 30% of average weekly income. We also examine the direct effect of being logged off Facebook for

[49] In our sample, about 43% of women post comments on Facebook at least once a week, compared to 21% of men. Also, 23% of women post pictures at least once a week compared to 8% of men.

one week on five outcomes: social media usage, news awareness, news consumption, subjective well-being, activities and the value of Facebook.

While individuals facing a Facebook restriction did refrain from using Facebook, they did not increase their usage of other social media. This is consistent with studies that find low usage across social media platforms and suggests that there is a significant switching cost between platforms.

In addition to not using other social media, participants did not look for news from other sources, even when the substitution cost for accessing news from other sources is low (i.e. turning on the television or radio or typing the web address of a news site instead of Facebook). Overall, awareness of mainstream news was not affected, but being off of Facebook resulted in more uncertainty about whether news from politically-skewed sources was fake or not. Those who experienced a week without Facebook were 22.1 percentage points more likely to be uncertain about a skewed news headline, and men's news awareness was most affected. These results imply that Facebook is an important source of news and may especially be a source of skewed news for men.

Our study has further implications. News aggregators that remove biases from news sources would better inform and educate the general public and could weaken the influence of skewed news. Facebook features (i.e. Instant Article, Trending News, etc.) suggest the company desired to serve as a news aggregation platform. However, recently Facebook eliminated these features out of concerns of propagating fake or skewed news, which goes in line with our finding on news consumption and awareness. While a news aggregator has the potential to provide an unbiased perspective of news and events (Mullainathan and Shleifer, 2005), our findings suggest that Facebook, as currently constructed, may not be well suited for this purpose.

Our results suggest that using Facebook induces feelings of depression, which plausibly decreases an individual's well-being. This effect is particularly pronounced for men, for active Facebook users and for those who experience negative emotions while on Facebook. Contrary to other studies (Tromholt, 2016; Valenzuela et al., 2009; Deters and Mehl, 2013), we find no effect with respect to reported overall life satisfaction. The reduction in depression we find from being off of Facebook might be explained by two mechanisms. First, being off Facebook could encourage individuals to engage in more positive, healthy activities, such as exercising and eating out less often, which could explain the improvement in mood. Second, Facebook itself might be a channel for decreasing subjective well-being, and changes in activities and consumption patterns could be a result of feeling better. Untangling the direction of causality would be an important area for future research.

Table 2.2: Weighting Adjustments

	Full Sample		Men		Women	
	Unweighted	Weighted	Unweighted	Weighted	Unweighted	Weighted
Facebook Use	-1.73***	-1.88***	-1.27***	-1.34***	-2.09***	-2.22***
News Media Index - Traditional Media	0.07	-0.05	0.47	0.77	-0.17	-0.48
News Media Index - Social Media	-0.66***	-0.61***	-0.81***	-0.68**	-0.53***	-0.55**
News Consumption Index	-0.64***	-0.59**	-1.01**	-0.59	-0.46*	-0.62**
Probability Right Answer - Mainstream News	0.01	0.04	-0.09	-0.03	0.03	0.05
Probability Wrong Answer - Mainstream News	0.002	0.005	-0.02	-0.04	0.02	0.02
Probability Not Sure Answer - Mainstream News	-0.01	-0.04	0.10*	0.06	-0.05	-0.07
Probability Right Answer - Skewed News	-0.16***	-0.14**	-0.32***	-0.33**	-0.09	-0.09
Probability Wrong Answer - Skewed News	-0.06	-0.03	-0.27**	-0.22**	0.06	0.06
Probability Not Sure Answer - Skewed News	0.22**	0.17*	0.59***	0.55***	0.03	0.03
Overall Satisfaction	-0.002	-0.02	0.12	0.21	-0.07	-0.10
Life is Worthwhile	0.05	-0.12	0.51	0.80**	-0.17	-0.42
Feel Happy	0.03	0.16	0.05	0.21	0.06	0.23
Worry	0.37	0.48*	0.42	0.28	0.37	0.57*
Depressed	-0.57**	-0.39	-0.82**	-0.90**	-0.44	-0.24
Healthy Daily Activities Index	0.84**	0.69	1.47**	1.16*	0.52	0.67
Productive Time Index	0.25	0.04	0.40	0.10	0.19	0.12
Efficient Time Index	0.33	-0.13	0.41	0.12	0.31	-0.18
Expected Healthy Daily Activities Consumption Index	0.16	0.22	0.28	0.12	0.14	0.32
Value of Facebook (% change)	0.20*	0.19**	-0.16	0.10	0.33**	0.19*

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

This table compares the main results with the weight-adjusted estimates. We use the inverse probability of being eligible as weights. All the p-values are adjusted for multiple comparisons. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

3. WHO SUPPORTS PIGOU? THE DISTRIBUTIONAL CONSEQUENCES OF PIGOUVIAN TAXES

3.1 Introduction

Governments and policy-makers often have the choice between simple uniform tax regimes and differentiated tax regimes. A uniform tax, at least in theory, allows for easier implementation on many differing dimensions. However, in the presence of heterogeneous firms and consumers, a uniform tax does a poor job at maximizing consumer welfare. This has been shown across many different industries including alcohol and liquor (Miravete et al., 2017), the insurance market (Finkelstein et al., 2009), and more generally, in markets where differentiated goods face identical prices (Orbach and Einav, 2007). The impact of uniform prices and taxes have large distributional consequences when applied in heterogeneous environments.

The focus of this paper is on uniform Pigouvian taxes in the personal transportation market for gasoline consumption in the United States. The externalities that result from driving impose a large cost on society as a whole. These costs include both local and global pollution damages (over $20 billion per year),[1] congestion and waste (over $78 billion per year),[2] and health costs resulting from accidents (estimated yearly damages of over $220 billion).[3] These examples account for only a few of the major costs borne from this industry. When considering the full spectrum, total damage estimates incurred by driving could easily push into the range of $400 billion per year. In order to address this large societal cost, the United States has primarily elected to regulate vehicle manufacturers through Corporate Average Fuel Economy (CAFE) Standards instead of through Pigouvian taxes.[4] While the current federal gasoline tax is only $0.184 per gallon, there has been

[1] From https://www.nytimes.com/2008/04/20/magazine/20wwln-freakonomics-t.html

[2] As reported by the Texas A&M Transportation Institute's 2007 mobility report.

[3] From Edlin and Karaca-Mandic (2006)

[4] As described by the U.S. Department of Transportation, "the purpose of Corporate Average Fuel Economy (CAFE) Standards is to reduce energy consumption by increasing the fuel economy of cars and light trucks. The CAFE standards are fleet-wide averages that must be achieved by each automaker for its car and truck fleet, each year, since 1978. When these standards are raised, automakers respond by creating a more fuel-efficient fleet, which improves (the) nation's energy security and saves consumers money at the pump, while also reducing greenhouse gasoline (GHG) emissions."

a growing interest which has spurred research focused on determining what the optimal Pigouvian gasoline tax should be. Most estimates suggest a tax of around $1.00 - $2.00 per gallon to address the full range of damages, but estimates as large as $3.00 per gallon are not uncommon as we learn more about the many pervasive effects (Parry and Small, 2005; Lin and Prince, 2009; Mankiw, 2009).

Given how large these costs are, it is important for social welfare to efficiently address the issue. If not, then there are substantial welfare consequences. A large body of literature is devoted to evaluating the efficiency of both CAFE Standards and gasoline taxes. While there is evidence documenting that CAFE has increased the fuel efficiency of the vehicle fleet over the last few decades, this has come at a substantial loss of efficiency due to the fact that CAFE standards are not as efficient as Pigouvian taxes.[5] From a conceptual standpoint, CAFE standards only correct for the extensive margin of driving - that is, a consumer's vehicle purchase decision. These regulations however do not correct for the intensive margin of driving - the decision of how much to drive and how much gasoline to purchase. On the other hand, an optimal Pigouvian tax accounts for both margins of driving. By using CAFE standards and not accounting for both margins, an inefficiency is created. Quantifying this inefficiency, it has been estimated that in terms of per gallon of gasoline, CAFE Standards are regressive policies that cost 3 to 6 times as much as a gasoline tax (Jacobsen, 2013; Austin and Dinan, 2005). These facts ultimately create an interesting paradox in trying to understand why the United States continues to rely on CAFE standards instead of implementing optimal Pigouvian gasoline taxes. One reason why this remains status quo is due to the general population's well documented dislike for taxes, and in particular, gasoline taxes. Survey evidence shows that this sentiment is true currently and has largely been true throughout the history of the United States (Agrawal and Nixon, 2014; Knittel, 2014).

In this paper, we address three main questions. First, what are the distributional consequences of a uniform gasoline tax from an equity and fairness standpoint?[6] Second, can this distribution

[5]Before CAFE standards were implemented in 1975, the average American car got about 13.5 miles per gallon. By 2016, the average American car gets about 25 miles per gallon.

[6]Previous papers have estimated the inefficiencies of a uniform gasoline tax, most notably Knittel and Sandler (fcm) and Bento et al. (2009). These papers find that a uniform gasoline tax does a poor job at minimizing dead

help explain the support (or lack thereof) for an increase in the gasoline tax? And finally, do revenue neutral tax policies exist that are capable of shifting the distributional consequences in a way such that it increases political support? We show three primary things.

First, by using confidential DMV data for the state of Texas augmented with data on local pollution damages (for NOX, VOC, PM25, and SO2), we calculate a uniform gasoline tax that properly accounts for these local pollution damages and find that there is substantial heterogeneity across the state when it comes to both the costs and benefits that arise from a uniform gasoline tax. A uniform tax imposes the same per gallon tax rate on individuals regardless of where one lives. We then counterfactually increase the price of gasoline by this uniform tax. We estimate how this impacts both individual expenditures and pollution damages. Overall, most people end up with higher expenditures. However, in response to higher gasoline prices, individual consumers reduce their vehicle miles traveled (VMT). This reduction leads to a decrease in the exhaust of local pollution and their associated damages within each county. On top of this, an individual benefits not only from their own reduction in driving and pollution, but also from everyone in their county reducing pollution as well. In most counties where pollution does very little harm, there are very few benefits generated. By contrast, in counties where pollution causes large damages, a reduction in VMT corresponds with a large benefit. There is substantial heterogeneity across Texas in regards to these measures.

Second, using data that comes from a stratified state-wide survey, we show suggestive evidence that the distribution of costs and benefits is an important determinant in an individual's decision to support an increase in the gasoline tax. Even after controlling for an individual's political affiliation, increasing a counties' net benefits can explain up to 10% of an individual's support for a gasoline tax. This result has important policy implications. There are generally two schools of thought when it comes to understanding support for a gasoline tax. One school tells us that support is determined by party lines (i.e., Republican vs. Democrat). Another school tells us that support is

weight loss and that there are gains to be had if the tax can be implemented at finer levels. This paper differs from these previous two in that we will focus mainly on the equity consequences of a uniform gasoline tax instead of the efficiency consequences.

determined by individual self-interests. Untangling these two from each other is important because of their implications for potential policy remedies. In a world where support is entirely determined by "tribal" politics, then there are not any policy levers available to influence support. However, if support is in part determined by self-interests, then policy makers can use this to strategically return the revenue generated under a Pigouvian tax in order to increase support for the tax itself. By doing so, raising support for the tax is socially beneficial if it can replace CAFE Standards.

Third, we show that through revenue neutral tax systems there is room for policy-makers to strategically shift the distribution of costs and benefits so that more counties end up with more people who have benefits that outweigh their costs. By doing so, this would increase support for a Pigouvian gasoline tax from a median level of 4 (out of 10) to a 5. There is growing interest in revenue neutral tax policies, but for policy-makers interested in moving away from CAFE standards in favor of Pigouvian taxes, we show a meaningful and significant shift in the distribution of political support.

The paper proceeds as follows. In Section 3.2, we discuss some of the background for the federal gasoline tax along with the longstanding public opinion for the tax. In Section 3.3, we reference and describe the three main sources of data used in this paper. Section 3.4 discusses the methods and estimation techniques used to analyze the main questions. Results are presented in Section 3.5. Finally, Section 3.6 discusses the policy implications as well as concludes the paper.

3.2 Background

The federal gasoline tax is currently set to \$0.184 per gallon. It was last increased by \$0.043 in 1993 by the Clinton administration through the Omnibus Budget Reconciliation Act. While there is some heterogeneity among state gasoline taxes, the average gasoline tax (federal and state combined) across the nation is about \$0.40/gallon.[7] Since its' inception in 1932, the federal gasoline tax has only been increased 9 times.[8] For the most part, the gasoline tax has only been

[7]See https://taxfoundation.org/state-gas-tax-rates-2019/ for a state-by-state summary of the state level taxes as of 2019. The main text reports the average as of 2008.

[8]Appendix Table B.2.1 shows the year of each increase as well as the reason for what the additional funds would be allocated towards.

raised in times of war with the intention to help reduce deficit spending or as a revenue generating instrument for transportation and infrastructure spending.

Furthermore, public opinion has never been favorable towards gasoline taxes. Survey evidence shows that this sentiment is currently true and has also largely been true throughout the history of the United States. Today, support for raising the gasoline tax is about 25% - 30% (Agrawal and Nixon, 2014; Knittel, 2014). Despite growing research showing the efficiency gains resulting from a gasoline tax, this lack of support has remained very static throughout American history. Even during times when people felt material consequences from not raising the gasoline tax, survey data showed that people strongly preferred rationing gasoline - after experiencing gasoline lines firsthand - over increasing the gasoline tax.[9] This sentiment can be contrasted with the fact that 92% of economic experts support the gasoline tax over CAFE Standards compared to the 22% of average Americans (Sapienza and Zingales, 2013).

The lack of support for a gasoline tax creates an interesting paradigm for policy-makers and academics alike. A policy-maker's primary objective is to either get elected or, if already elected, to remain in office. As mentioned in Section 3.1, this means that policy-makers can have strong incentives to pass sub-optimal policies (i.e., CAFE Standards) if a gasoline tax is too politically unpopular with their respective constituency. As for academics, it remains an open question why the general public does not support a gasoline tax. It was once thought that there was not any support because the tax was regressive, but this has been shown to not be the case (Poterba, 2011). Whether this fact is salient to the public is not clear however. Ironically, CAFE Standards have been shown to be regressive (Jacobsen, 2013). Again, it is not evident whether this fact is also salient to the public.

On top of showing the distributional consequences from an equity standpoint, one of the main contributions of this paper is to help answer why the general public does not support a gasoline tax. Specifically, we show that an individual is more likely to support a gasoline tax if their own benefits exceed their costs. There are large policy implications from being able to answer this ques-

[9]Knittel (2014) details multiple surveys throughout US history showing the lack of support a gasoline tax has received.

tion. First, it establishes a behavioral link between self-interest and tax support that goes beyond "tribal" politics. Second, this link enables policy-makers to design tax policies that are "politically-sophisticated." Finally, by designing a tax with an increased level of support, this can pave the way for the transition towards a more efficient policy solution that addresses the externalities imposed by the gasoline market and away from the inefficient CAFE Standards. Ultimately, this transition represents the potential for large societal welfare gains.

Currently, local governments are experimenting with alternative tax regimes in order to establish either gasoline taxes or carbon taxes. The state of Washington is leading the way. In 2016, Ballot Initiative 732 was put up for vote that would have formally established a carbon emissions tax.[10] Ballot Initiative 732 lost by a 4-6 margin. Two years later in 2018, Washington Senate Bill 6203 was introduced to try and again establish a state-wide price on carbon. What was different about this attempt is that it featured revenue neutrality. By definition of revenue neutrality, all revenue collected through the tax is returned back to the state through various investments and tax cuts.[11] This was not the first time policy-makers started thinking about revenue neutrality. In February of 2017, George P. Shultz and James A. Baker III put together a carbon tax plan with the Climate Leadership Council (CLC) that would start a carbon tax at about $40 per ton, escalating by 2 to 5 percent annually, and reaching as high as $65 per ton by 2030. Importantly, The Baker-Shultz proposal would return all of the tax revenue to American households in the form of monthly rebate checks. It is estimated that for most households, the amount rebated would be greater than the increased costs they would pay for fuel (Lavelle, 2019). These examples and recent events show the increased interest in designing tax plans that focus on equity through various revenue neutral policies. If people are motivated by their own self interests, then there exists the potential for these revenue return systems to help raise support for the tax and push the transition towards

[10]In summary, "(t)his measure would impose a carbon emission tax on the sale or use of certain fossil fuels and fossil-fuel-generated electricity, at $15 per metric ton of carbon dioxide in 2017, and increasing gradually to $100 per metric ton (2016 dollars adjusted for inflation), with more gradual phase-in for some users. It would reduce the sales tax rate by one percentage point over two years, increase a low-income sales tax exemption, and reduce certain manufacturing taxes." (Ballotpedia)

[11]Senate Bill 6203 features a 100% return of all revenues generated under the tax. 50% would be invested in carbon reductions, 20% to water and forest funds, 15% to low income programs, and 15% to rural economic development.

gasoline taxes in the future.

3.3 Data

The empirical setting of this paper is the Texas personal transportation market. Data for this project comes from three primary sources. First, we utilize confidential data from the Department of Motor Vehicles for the state of Texas for the years between 2004 and 2010. Second, we use marginal damage data that dollarizes the damages associated with an additional pollutant in each specific county. Finally, we use rich survey data that comes from the Texas A&M Transportation Institute in order to estimate individual behaviors and support for gasoline taxes.

3.3.1 Texas Department of Motor Vehicles Data

The primary source of data is confidential vehicle registration data from the Texas Department of Motor Vehicles (DMV) between the years of 2004 and 2010. This data provides a census of all vehicles in the Texas state fleet and allows us to identify all vehicles owned by a household, each vehicle's respective annual VMT, and any changes to the ownership status of a vehicle.[12] Not only does this data allow us to measure a household fleet, but these records also contain information on the unique vehicle identification number (VIN). This information allows us to build on the DMV data by combining it with very rich engineering data specific to each class/make of vehicle. We merge each VIN with DataOne Software to obtain characteristics for each registered vehicle including miles-per-gallon (MPG), horsepower (HP), size, etc.[13] This data provides the benchmark from which we will work with. Table 3.1 provides summary statistics for this data.

According to Table 3.1, we have a census of over 5 million households in Texas containing approximately 8 million unique vehicles indicating that households on average own 1.5 vehicles. Importantly, most vehicles in the state fleet are driven for about 12,000 miles per year. However, since the standard deviation on this measure is large, this will later speak to the heterogeneity in

[12]To calculate annual VMT (AVMT), we average a vehicles odometer readings over the time between readings, then multiply this by 365. If a vehicle has an associated safety inspection from TCEQ, we augment our primary readings with these additional odometer readings too.

[13]DataOne is a software product that offers VIN decoding services and provides an automotive database for vehicle characteristics.

Table 3.1: DMV Summary Statistics

	Median	Mean	Std. Dev
Yearly Avg HH Count		5,209,694	
Yearly Avg Vehicle Count		8,038,925	
Panel A) Vehicle Characteristics			
MPG	19	19.29	4.67
MSRP	23,485	25,338	10,691
Curb Weight	3,494	3,536	1,122
Size (HxWxL ft^3)	487	526	144
HP	181	174	86
Panel B) Driving Characteristics			
Avg Annual VMT	10,750	12,034	17,970
Avg Annual Gallons	561	666	1117
Panel C) Census Data			
Median Income		\$54,727	
Bachelors Degree +		28.1%	
People per Household		2.84	
% White		79.4%	
% Black		12.6%	
% Hispanic		39.1%	

This table provides summary statistics of confidential DMV data for the state of Texas for the years 2004 - 2010. Panel A describes various technical characteristics for vehicles. Panel B describes driving characteristics for households, and Panel C gives census level information for the state of Texas. Medians and standard deviations shown when able to calculate.

individual costs resulting from a change in the price of gasoline. Furthermore, most of the vehicles in the fleet have a MPG rating of about 19 and tend to be mid-sized based on the curb weight and volume size.

3.3.2 Pollutant Marginal Damages

This paper relies on an accurate assessment of local pollutants and the damages they impose at granular levels. To do so, we rely on the AP2 model - an integrated assessment air pollution model. The AP2 model maps emissions from different local pollutants (PM25, VOC, NOX, SO2) into ambient air concentrations of O2, SO2, and PM25 at the county level. Furthermore, Muller and Mendelsohn extend AP2 to estimate welfare effects imposed by these pollutants. Using data from the U.S. Census and the U.S. Department of Agriculture, Muller and Mendelsohn consider and incorporate in their model how pollution impacts human health, damages crop and timber yields, degrades buildings, infrastructure, and materials, and causes reductions in visibility and recreation.

These exposures are then translated into physical effects using concentration-response functions reported in related literature. Perhaps most importantly, Muller and Mendelsohn estimate the impact of exposure to a specific pollutant into a monetary value for adult mortality rates. In order to calculate these marginal damages, a baseline level of pollution is established from which associated damages are calculated. Then from a polluting source, AP2 adds 1 ton of a specific pollutant to that county and damages are recalculated. The difference between the baseline level of damages and the recalculated level is referred to as the marginal damage of pollutant p in county c. Table 3.2 provides summary statistics for the marginal damages of local pollutants across the counties in Texas.

Table 3.2 shows the variation in pollutant damages across the 254 counties in Texas. The damages associated with PM25 and SO2 have larger magnitudes relative to VOC and NOX, but the variation across county for each of these pollutants is large. Because this paper focuses on personal transportation and gasoline consumption, we also need to determine specific vehicle emission rates for each pollutant at a granular level in order to assess how much of each pollutant a vehicle is emitting when being used. We use several sources to determine these emissions rate for gasoline vehicles.[14] Primarily, we use EPA emission standard tiers for various vehicle bins indicating a

[14]We only consider vehicles that use gasoline and diesel as fuel. This represents the majority of the state fleet.

Table 3.2: Pollutant Marginal Damages

	Mean	Std. Dev.	Min	Max
PM25	42,078	35,025	5,436	369,609
VOC	1,860	1,560	434	15,801
SO2	13,403	7,904	7,202	98,762
NOX	4,115	1,707	1,135	9,860
N	254	254	254	254

This table provides summary statistics for the marginal damages associated with various local pollutants imposed on the counties in Texas. Values are as reported by Muller and Mendelsohn's AP2 model, an integrated assessment air pollution model. Values are in terms of $2,000/Ton.

gram/mile emission estimate for PM25, VOC, and NOX.[15] SO2 emissions are directly proportional to gasoline consumption which allows us to use standard conversion factors provided by GREET. The combination of county level marginal damages with vehicle specific emissions rates gives rise to significant heterogeneity in the costs and damages across the state. For example, two identical vehicles can have very different welfare implications depending on which county they are driven (polluting) in. Similarly, high fuel efficient vehicles may impose large costs on society if they are driven in susceptible counties whereas older, less efficient ones may not cause any damage if it is being driven in other regions. To summarize, our detailed sets of data allows us to identify the fuel efficiency of vehicles, where they are being driven, by how much they are being used, the quantity of pollutants they are emitting, and what this means in terms of dollarized pollution damages with respect to the county the vehicle is registered in.

3.3.3 Texas A&M Transportation Institute Poll

The final source of data for this paper comes from the Texas A&M Transportation Institute (TTI). TTI launched a Transportation Policy Research Center in 2013 that focuses on many different issues related to transportation including finance, freight, congestion, public engagement, and

[15]References can be found in the EPA's emission standards reference guide for light duty vehicles and trucks.

Table 3.3: Summary Statistics of Poll Respondents

	Obs.	Mean	Std. Dev.
Age	4,363	51.60	14.04
Income ($'000s)	3,783	76.97	54.91
Fraction with a Bachelor's Degree	4,363	0.43	0.5
Fraction White	4,363	0.80	0.4
Fraction Male	4,359	0.50	0.5
Number of Household Vehicles	4,363	2.28	1.38
Household Size	4,363	2.72	1.81
Fraction Republican	4,363	0.35	0.48

This table provides summary statistics of basic demographic information for respondents in the Texas A&M Transportation poll. Sample is stratified across the state. Appendix Section B.1 contains a full depiction of the stratified regions.

technology. The Institute brings together experts from engineering, finance, economics, technology, policy and public engagements fields, from both academics and the private sector. To start identifying important problems and potential solutions, TTI has acquired data from a variety of sources. For this paper, we use data that comes from the Texas Transportation Poll conducted by TTI.

The TTI poll has been administered twice - once in 2014 and another time in 2016.[16] For each wave, TTI randomly samples approximately 5,000 participants to answer travel behavior questions as well as to get opinions on daily transportation choices, challenges, funding, and potential solutions.[17] The poll was conducted through stratified sampling across 12 different regions of Texas.[18] Appendix Section B.1 contains a full depiction of the 12 regions. The poll also gathers demographic information on the respondents. Table 3.3 reports this below.

The average respondent is about 52 years old with a household income of $77,000 per year.

[16]We only use data coming from the 2016 wave. Data in 2016 has information on a respondent's zip-code of residence, whereas the 2014 poll only gives information on an individual's region of residence.

[17]The random sample poll responses were collected by mail, by phone, or online. The sample size of more than 5,000 provides for a statewide margin of error of plus or minus 1.5 percent.

[18]The 12 regions are defined as Region 1 - Houston; Region 2 - Dallas; Region 3 - Fort Worth; Region 4 - San Antonio; Region 5 - Austin; Region 6 - Laredo, Pharr; Region 7 - Corpus Christi, Yoakum; Region 8 - Bryan, Waco; Region 9 - Atlanta, Beaumont, Lufkin, Paris, Tyler; Region 10 - Amarillo, Childress, Lubbock, Wichita Falls; Region 11 - Abilene, Brownwood, Odessa, San Angelo; Region 12 - El Paso.

Both of these are slightly higher than the average Texas-wide population (45 years old, with a household income of about $56,700 per year).[19] Respondents are predominantly white with an equal proportion of responses coming from men and women. Lastly, there is similar representation between republicans and democrats.[20]

The poll asks individuals a variety of questions about daily transportation choices, challenges, funding, and potential solutions. For the purposes of this paper, we will be using responses from questions that ask individuals about their level of support for increasing the state gasoline tax,[21] as well as responses from questions about gasoline tax opposition.[22] We also use some questions as control variables. One of the major benefits of this poll is the rich set of questions and responses that is provided at an individual level. Namely, we use responses that capture an individual's opinion on whether they support federal/state/local government involvement, and to what degree,[23] as well as responses about the degree to which congestion is a local problem.[24] Our measure of gasoline tax support asks individuals whether they support it as an instrument to fund transportation infrastructure, not necessarily whether they support it for Pigouvian tax purposes. We are not concerned that this alternative wording would change results. Conventionally, academic economists refer to the gasoline tax as a Pigouvian tax. But to the general public and policy-makers alike, the gasoline tax has been used as a funding instrument for its entire existence.[25] Hence, it does not seem that one would respond differently - especially under revenue neutral systems in which

[19]As reported by the US Census Bureau. Since the poll does not include individuals less than 18 years old, we have adjusted the Census Bureau's age estimates to just those of an adult age to get the 45 reported in the paper.

[20]There are about 1,500 Republicans and 1,100 Democrats. The remaining are either Independent, "Other", or preferred not to respond.

[21]Individuals are first prompted "On a scale from 0 (Oppose) to 10 (Support), how strongly do you oppose or support these potential ways to fund transportation in Texas?" Following this, two questions are "Increasing the state fuel tax by 5 cents per gallon" and "Increasing the state fuel tax by 10 cents per gallon."

[22]Individuals are first prompted "Which of the following statements best reflects your opposition to increasing the state fuel tax to generate additional transportation funding?" Following this, they are then asked "I oppose any type of tax increase."

[23]Individuals are first prompted "On a scale from 0 (Disagree) to 10 (Agree), how strongly do you agree with the following statements?" They are then given three following statements "Local (State, Federal) government should take a more significant role in addressing transportation issues in my region."

[24]Individuals are asked "On a scale of 0 (Not bad at all) to 10 (Extremely bad), how would you rate congestion in your community?

[25]Appendix Table B.2.1 shows how the federal gasoline tax has been used and to what specific purpose each cent in revenue has been designated towards.

Table 3.4: Summary Statistics of Respondent Answers

	Obs.	Mean	Std. Dev.	Min	Max
Support 5 Cent Increase (0-10)	4,147	3.95	3.47	0	10
Support 10 Cent Increase (0-10)	4,145	2.91	3.20	0	10
Oppose any Tax Increase (0-1)	4,363	0.14	0.34	0	1
Local Gov. Support (0-10)	4,171	7.57	2.55	0	10
State Gov. Support (0-10)	4,164	7.40	2.60	0	10
Federal Gov. Support (0-10)	4,152	5.39	3.55	0	10
Congestion Rating (0-10)	3,024	6.47	2.38	0	10

This table provides summary statistics of various selected questions and answers for respondents in the Texas A&M Transportation poll.

revenue from the tax is reinvested (or given back) directly to the individual. Table 3.4 reports summary statistics for the primary questions that will be used going forwards.

The first two rows of Table 3.4 provide summary statistics for answers about an individual's level of support for raising the gas tax by either 5 or 10 cents. These answers will be used as our primary outcome variables in later analysis. Options for these answers range from 0 to 10, with 0 - 4 indicating decreasing opposition, 5 indicating a neutral position, and 6-10 indicating increasing support.[26] Overall, survey respondents on average lean slightly towards opposing a 5 cent increase in the gasoline tax ($mean = 3.95$), but become increasingly less supportive for a 10 cent increase ($mean = 2.91$). A small proportion, about 14% of the population, oppose a tax increase of any kind. Furthermore, responses indicate that individuals are more inclined to support finer levels of government involvement when it comes to transportation policy-making. Support for local government involvement ($mean = 7.57$) is more favorable than involvement by the state government ($mean = 7.40$), but both are significantly more preferred to federal government involvement ($mean = 5.39$).

The data from this survey provides one of the largest within state samples that gathers information on individual support for a gasoline tax while also collecting very rich demographic data

[26]We also bin responses into a binary variable where responses with values of 0 - 4 indicate opposition and responses with a value of 6 - 10 indicate support. For this variable, we drop neutral (5) responses. Our binary variable, $fivecents_i$, takes a value of 0 if individual i's survey metric of support is between 0 and 4, and takes a value of 1 if between 6 and 10.

for each respondent as well. While support for a gasoline tax has been frequently surveyed, it has either been done at the national level (Gallup, PEW, etc) with smaller sample sizes (∼1,000), or they have been conducted by local news stations that only survey individuals from their local cities. Neither of these sources provides a similarly rich set of information on the respondents and because of this, does not allow for a rigorous analysis of the determinants of gasoline tax support.

3.4 Methodology

3.4.1 Front-End Calculations

The goal of this first step can be broken into three main steps. First, we design an optimal uniform Pigouvian tax that accounts for the damages incurred from the local pollutants emitted by the state vehicle fleet. Second, we counterfactually impose this tax by raising the price of gasoline. Third, we determine how households would respond to this price increase and compare the additional expenditures with reduced pollution damages in this counterfactual scenario to the actual in order to analyze the distributional consequences of a uniform tax.

3.4.1.1 Design of an Optimal Uniform Tax

To design an optimal uniform Pigouvian tax, we follow a procedure commonly used throughout the literature. The Pigouvian tax per gallon of gasoline is calculated as the average externality per gallon across all vehicles in Texas for a specific year. Local externalities (namely PM25, VOC, SO2, and NOX) are valued using the county level marginal damages as reported in Muller and Mendelsohn (2009).[27] We merge these damages to a vehicle's county of registration.[28] Then, we define the marginal damage per gram of pollutant p in county c to be θ_c^p. Similarly, we define a vehicle's emission rate of pollutant p in grams per mile by vehicle i to be ϵ_i^p. With these two measures, θ_c^p and ϵ_i^p, we are able to define the externality per mile of vehicle i, denoted by E_i, as:

[27] We do not consider the impact of Carbon. While there are global estimates for the social cost of carbon (SCC), we need data on the local benefits from carbon reduction.

[28] By doing so, we assume that an individual drives solely in their county of residence. While this may be a strong assumption, this will bias our results downwards. Due to the nature of urban sprawl, most people commute to more urban/dense cities but these also tend to be the areas where pollution damages are higher.

$$E_i = \sum_p \epsilon_i^p * \theta_c^p. \tag{3.1}$$

The externality per mile, E_i, sums across pollutants the dollarized damages imposed by a vehicle for each mile driven. Following this, we use the externality per mile for every vehicle in the state fleet to calculate a uniform Pigouvian tax per gallon for year y according to the following equation:

$$\tau^y = \frac{1}{N^y} \sum_{i=1}^{N^y} \frac{E_i}{MPG_i}. \tag{3.2}$$

The tax per gallon, τ^y, is defined to be the average across every vehicle's externality per mile scaled by its fuel efficiency, MPG_i. τ^y accounts for the damages incurred from the local pollutants emitted by the state vehicle fleet.

3.4.1.2 Counterfactual Response to a Uniform Tax Increase

With a tax per gallon now in place, our next step is to counterfacually raise the price of gasoline and determine how households would respond to these higher prices. To do this, we use credible estimates from the literature for the VMT elasticity and the pass-through rate.[29] For the VMT elasticity, γ, we use estimates from Levin et al. (2017) and Wang and Chen (2014). In practice, we use an estimate of $\gamma = -0.30$. There may a wide range of heterogeneity in the elasticity by income and region, but our uniform response will provide a lower bound for our results. For the pass-through rate, there is a difference between a federal tax and a state tax. A federal tax is passed on entirely to the consumer while a state tax is split almost equally between consumers and suppliers (Marion and Muehlegger, 2011). Furthermore, the incidence is shown to take effect almost immediately. For simplicity, we assume that consumers are responsible for the full tax incidence. This mimics a consumers response to price increases as if they were imposed under the

[29]The VMT elasticity can also be estimated using the following equation: $ln(VMT_{ijt}) = \beta ln(DPM_{ijt}) + \nu_j + \nu_t + \epsilon_{ijt}$, where i indexes vehicles, j indexes geographic location, and t indexes time. VMT_{ijt} is a measure of the vehicle miles traveled by vehicle i in region j at time t. DPM_{ijt} is a measure of the daily price of gasoline for vehicle i in region j at time t. Regression also controls for region and time fixed effect in ν_j and ν_t respectively.

federal gasoline tax. In order to determine how individuals would response to a uniform tax, we calculate the following equation:

$$VMT_i^{\tau^y} = VMT_i^y[1 - \gamma(\frac{\tau^y}{P_g^y})]. \tag{3.3}$$

Equation 3.3 represents a household's miles-driven response after a tax τ^y is implemented.[30] This takes into consideration the VMT elasticity, γ, the pass-through rate, and the price of gasoline, P_g^y. We use the average daily retail gasoline price in Texas as reported by the EPA for our measure of P_g^y.[31]

3.4.1.3 Assessment of Distributional Consequences

Following this increase in gasoline prices, our last step is to analyze the distributional consequences of a uniform tax. We do this by comparing an individual's additional expenditures with the reduction in pollution damages between actual driving behaviors and their counterfactual response to a uniform tax. Intuitively, as prices increase, the quantity of miles driven should decrease according to the VMT elasticity, γ, and the pass-through rate. We can then compare the additional expenditures a household faces between the two states, $VMT_i^{\tau^y}$ and VMT_i^y. The additional costs to a household imposed under τ^y are determined by:

$$C_i^y = \sum_n [P_g^y \frac{VMT_i^y}{MPG_i} - (P_g^y + \tau^y)\frac{VMT_i^{\tau^y}}{MPG_i}]. \tag{3.4}$$

Equation 3.4 assesses a household's additional expenditures on gasoline purchases over a year. For each household, we sum across the total number of vehicles owned where n represents the number of cars owned by household i.

Furthermore, since $VMT_i^{\tau^y} < VMT_i^y$, the incurred damages associated with local pollutants will decrease due to the reduction in gallons of gasoline being used and burned. Because of this, each individual in the county benefits from every one else decreasing their individual gasoline

[30]We assume that a household does not substitute miles between their household fleet. Similarly, at least in the short-run, we also assume that a household does not respond to the gasoline tax by scrapping a car in their fleet.

[31]See Appendix Figure B.3.2 for a time series of this price.

consumption. Therefore, we create a measure for County Benefit that captures the benefits that result from the decrease in damages under a uniform tax by using the marginal damages for each pollutant p in each county c. This measure is defined as:

$$B_c^y = \sum_i VMT_i^y(\epsilon_i^p \theta_c^p) - VMT_i^{\tau^y}(\epsilon_i^p \theta_c^p). \tag{3.5}$$

In Equation 3.5, i indexes all cars in county c. We assess the difference in pollution damages between the two states, $VMT_i^{\tau^y}$ and VMT_i^y.

Using the measure of individual cost, C_i^y, and the measure of county benefits, B_c^y, that result from a uniform tax τ^y, we compare the two in order to summarize the distributional consequences of a uniform gasoline tax. Intuitively, for households that experience larger benefits than they do costs, $C_i^y < B_c^y$, we define as "Winners" ($Winner_i$). On the other hand, for households that experience larger costs than they do benefits, $C_i^y > B_c^y$, we define as "Losers" ($Loser_i$). This comparison provides the primary metric that we use to evaluate the distributional properties of a uniform gasoline tax.

3.4.2 Back-End Regressions

The overall objective of this second section is to establish a behavioral link between the distributional consequences imposed under a uniform tax and an individual's level of support for a gasoline tax. In an ideal world, researchers would be able to randomly assign variation in the distributional consequences of an individual and track their support for the gasoline tax over time. This process would yield an unbiased estimate of the impact that heterogeneity in the distributional burden has on tax support. Unfortunately, this idealized experiment is not feasible in practice due to the complicated nature of any potential intervention(s). In order to bypass this complication, we exploit the richness of our data.

3.4.2.1 Linking the Distributional Consequences to Political Support

Our primary specification will be to estimate effects within an ordinary least squares (OLS) framework. The richness of our data allows us to control for and acknowledge a significant number of potentially confounding stories that would bias the relationship between the distributional consequences and political support. Formally, our main specification is as follows:

$$y_i = \beta_0 + \beta_1 Win\%_i + \beta_k X_k + \beta_n X_n + \epsilon_i. \tag{3.6}$$

In Equation 3.6, y_i represents an indicator for supporting an increase to the gas tax, $Win\%_i$ is the County Win Percentage for individual i, X_k is a vector of county level covariates including the 2008 Election GOP Vote Percentage as well as county-level ambient air quality, and X_n is a vector of individual covariates. Importantly, X_n includes a suite of individual political characteristics like party affiliation and whether the individual is an active voter or not.

The results from the estimation linking distributional consequences to political support have significant policy implications when it comes to the potential solutions for addressing the externalities imposed on society by this market. While it may be socially beneficial to transition away from CAFE Standards and utilize Pigouvian taxes instead, the lack of political support maintains the status quo. There are generally two schools of though when it comes to the determinants of support for a gasoline tax. One school tells us that support is dictated by political party affiliation (i.e., Republican vs. Democrat). There is a growing body of literature looks at the strong effects of social identity on political behavior (Klar, 2013; McLeish and Oxoby, 2011).

Another school tells us that support is determined by self-interest. Understanding which of these two is responsible (and by how much) for determining political support is important because of its implications for potential remedies. For example, in a world where support is only based on "tribal" politics, then there really are not any policies that can address this. However, if support is in part determined by self-interests (i.e., the distributional consequences), then policy makers can potentially use this information to strategically return the revenue generated under a tax increase in order to increase support for the tax itself. In effect, by increasing support for the tax, this would be socially beneficial if it can replace CAFE Standards.

The richness of our data allows us to consider both of these explanations as well as their respective magnitudes. Furthermore, we are also able to control for a vast amount of potentially confounding factors. These features allow us to garner insight into the behavioral determinants of political support for gasoline taxes.

3.5 Results

3.5.1 Optimal Uniform Tax and its Distributional Consequences

In this section, we follow the procedures as described in Section 3.4.1. Recall, the overall objective is to design an optimal uniform Pigouvian tax that accounts for the damages incurred from the local pollutants emitted by the state vehicle fleet. We then counterfactually impose this tax by raising the price of gasoline. Finally, we determine how drivers respond to this price increase and compare the additional expenditures with the reduction in pollution damages under a hypothetical tax increase in order to analyze the distributional consequences of a uniform tax.

One of the main component's in this section is the distribution of pollutant marginal damages for each county in Texas. Figure 3.1 shows this distribution across the state for each of the primary local pollutants. According to these distributions, PM25 and VOC have significantly higher damages in more urban areas, like Dallas and Houston, whereas NOX and SO2 are more region specific. NOX has much larger impacts on the northeast region and SO2 impacts the southern band more than it does other areas. Furthermore, Table 3.2 shows that on average, PM25 and SO2 have qualitatively larger damages compared to VOC and NOX. Hence, areas experiencing reductions in these two pollutants have room for substantial benefits.

We further combine this data with EPA emissions rates per mile as well as VMT data from the DMV to determine the dollarized rate of each vehicle's externality per mile. Following this, we then uniformly distribute this damage rate across the state fleet to obtain a tax per gallon of gasoline. In practice, this amounts to a uniform gasoline tax of about \$0.40 per gallon. Hence, this result implies that in order to optimally address the externalities from local pollutants, the gasoline tax should be \$0.40 per gallon. While at first glance, this tax is lower than the numbers reported in Section 3.1, it is important to note that estimates ranging between \$1.00 - \$2.00 include the full spectrum of externalities, like accidents, congestion, and carbon emissions, and even noise

Figure 3.1: Distribution of Pollutant Marginal Damages across Texas

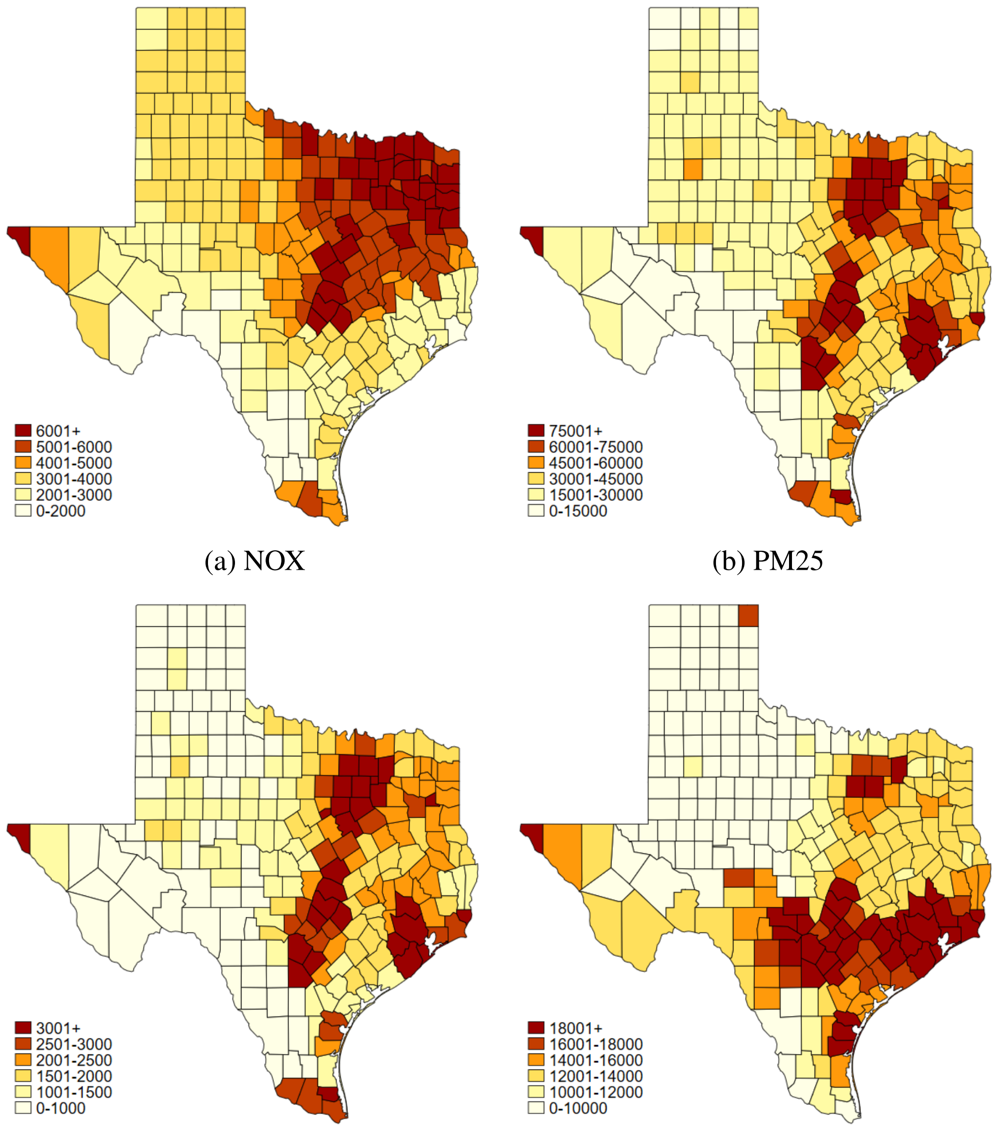

(a) NOX (b) PM25

(c) VOC (d) SO2

This figure reports the marginal damages by county for the primary local pollutants NOX, PM25, VOC, and SO2 in Texas. Values are as reported by Muller and Mendelsohn's AP2 model, an integrated assessment air pollution model. Values are in terms of $2,000/Ton.

pollution. When decomposing those estimates by each of the respective externalities, a $0.40 per gallon tax is in line, albeit slightly higher, than what the current research has established.

Our next step is raise the price of gasoline and estimate how consumers would respond to these higher prices.[32] From the literature, we assume both a VMT elasticity as well as a pass-through rate. For the elasticity, we use a value of -0.30 (Levin et al., 2017; Wang and Chen, 2014). The pass through for a federal gas tax is estimated to be very close to 1, so we use this value (Marion and Muehlegger, 2011). We then estimate Equation 3.3 using the previous values. This allows us to estimate how individuals would respond and change their driving behavior (VMT) for a uniform price increase across the state.

A household's driving response has two primary consequences. Since $VMT_i^{\tau^y} < VMT_i^y$, we calculate the increase in a household's expenditures on gasoline purchases over a year as in Equation 3.4. For each household, we sum across the total number of vehicles owned by the household. Furthermore, the incurred damages associated with local pollutants will decrease due to the reduction in gallons of gasoline being used and burned. Because of this, each individual in the county benefits from every one else decreasing their individual gasoline consumption. This decrease in VMT for individual results in a county wide benefit that everyone experiences, as defined by Equation 3.5. Finally, we compare individual costs to county benefits in order to determine who "wins" (benefits $>$ costs) and who "loses" (costs $>$ benefits). The proportion of "Winners" within each county is reported in Figure 3.2.[33]

Figure 3.2 presents a strong result. Most of the individuals who benefit from a uniform tax are located in more populated regions, that is, the "Texas Triangle." The Texas Triangle is formed by four of the major cities in Texas totaling a population of over 13 million; Houston, Dallas-Fort Worth, San Antonio, and Austin, connected by Interstate 45, Interstate 10, and Interstate 35.[34] As documented earlier, Muller and Mendelsohn's marginal damage model places significant value on mortality and health factors. Similarly, the distribution of damages associated with local pollutants has large impacts on this area (Figure 3.1). However, there are also a handful of counties outside

[32]Note that the current gasoline tax in the state of Texas is about \$0.38 (state and federal combined). Because of this, we only increase the price of gasoline by the difference between τ^y and the current gasoline tax. Depending on the actual price of gasoline (Appendix Figure B.3.2), this implies price increases ranging anywhere from 0.5% to 1.4%. As more externalities are considered, this magnitude will increase accordingly. As currently constructed, these increases provide enough variation in the distributional consequences for us to glean insight from.

[33]Appendix Figure B.3.3 reports a histogram of the percent of counties across various "Win" proportions.

[34]In 2004, there were 13.8 million people living in the Texas Triangle. In 2015, this increased to 18.1 million.

Figure 3.2: County "Win" Proportions

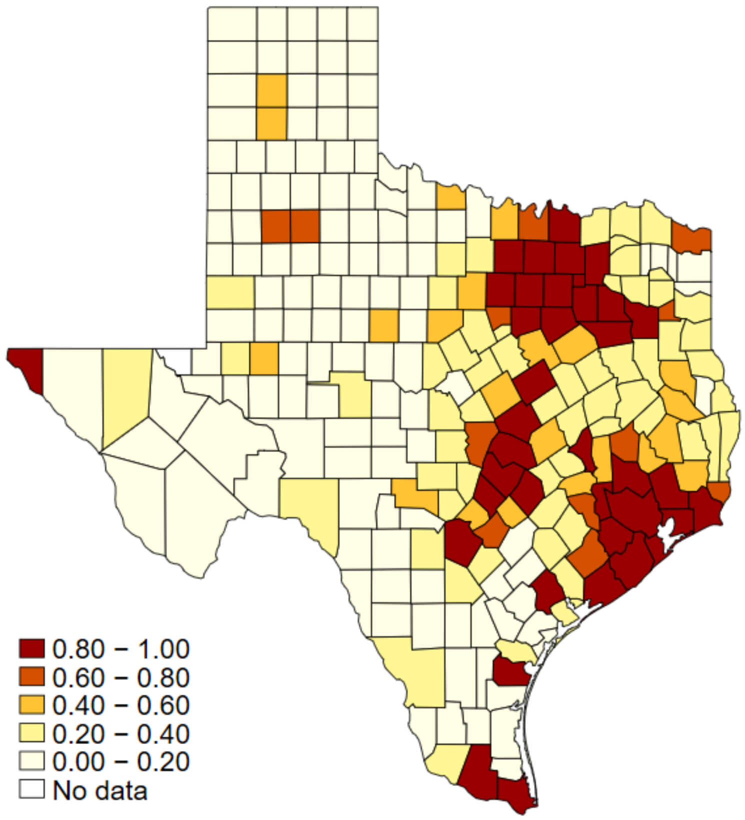

This figure presents results for the front-end calculations, and depicts the proportion of "Winners" (benefits > costs) in each county. "Winners" tend to be concentrated in the more urban regions of the state, but can also be driven in rural counties where pollution damages are high.

of the Texas Triangle that end up "Winning" depending on the specific distribution of costs and benefits. On average, most counties have more "Losers" than they do "Winners" and these tend to be the more rural areas.[35] It further points out that the more rural counties have very little to gain from a uniform increase to the gasoline tax. These counties tend to be relatively "clean," meaning that local pollutants impose very little harm on them. Hence, reductions in these pollutants offers very little benefit but large increases in their expenditures (costs).[36]

[35]See Appendix Figure B.3.3 for a histogram of this distribution.

[36]This may partially be captured by a heterogenous VMT elasticity. However, rural areas tend to be less elastic than their urban counterparts indicating that our estimates are most likely lower bounds for the costs and expenditures that these regions would face.

3.5.2 Distributional Consequences and Political Support

The overall objective of this section is to establish a behavioral link between the distributional consequences imposed under a uniform tax as reported in Section 3.5.1 and an individual's level of support for a gasoline tax.[37] In order to bypass potential confounders between these two, we exploit the richness of our data. To do this, we estimate equation 3.6 using each county's win percent as our main variable of interest to predict an individual's support for raising the gasoline tax by 5 cents. Table 3.5 shows these results across various specifications.[38]

Overall, column (1) of Table 3.5 shows a significant and positive relationship between the distributional consequences and support for an increase in the gasoline tax. However, there are various confounders that might bias this relationship. In order to circumvent this, we exploit the richness of our data. In column (2) we add covariates for an individual's party affiliation and an indicator for being a high-mileage driver. In column (3), we control for various county level characteristics like the number of daily ambient air quality violations and population. In columns (4) through (6), we continue adding a full spectrum of potentially related variables, like education, minority status, employment status, a dummy for registered voters, age, gender, support for local, state, and federal government involvement, number of vehicles owned, and household size. Across each of these specifications, Table 3.5 shows a consistent and robust positive estimate for the impact of county win percentage on support for an increase in the gasoline tax. This result provides evidence that the distribution of costs and benefits is an important determinant in an individual's decision to support an increase in the gasoline tax.

Important to this paper, we show that even after controlling for an individual's political affiliation, increasing a counties' net benefits explains more than 10% of an individual's support for a gasoline tax. This result has important policy implications. In a world where support is entirely de-

[37] Appendix Figure B.3.1 show the distribution of support for a 5 and 10 cent increase in the gasoline tax across the state of Texas. Panel (c) also shows opposition to increasing a tax.

[38] Appendix Table B.4.1 reports results for the full specification. Results using support for a 10 cent increase as our outcome variable are consistent with results for a 5 cents increase. Appendix Table B.4.3 shows full results for an individual's support for a 10 cent increase in the gasoline tax. Furthermore, results using the raw scale of support are similar to those presented to those in 3.5. See Appendix Table B.4.2 for results where the raw scale of support is the outcome variable of interest.

Table 3.5: Relationship Between County "Win" Percent and 5 Cent Support

	1	2	3	4	5	6
County Win Percent	0.18***	0.17***	0.14***	0.12***	0.11**	0.10**
	(0.03)	(0.03)	(0.05)	(0.05)	(0.05)	(0.04)
Republican		-0.09***	-0.09***	-0.11***	-0.12***	-0.11***
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
High Mileage		-0.07***	-0.07***	-0.04**	-0.02	-0.02
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
County GOP Voting Percentage			-0.04	-0.05	-0.05	-0.03
			(0.11)	(0.09)	(0.09)	(0.09)
Independent/Other		X	X	X	X	X
County Air Quality Violations			X	X	X	X
High Education				X	X	X
Minority				X	X	X
Employed					X	X
Registered Voter					X	X
Old Age					X	X
High Income					X	X
Female					X	X
Local Gov. Support						X
State Gov. Support						X
Federal Gov. Support						X
HH Vehicle #						X
HH Size						X
R-sq	0.016	0.022	0.023	0.044	0.053	0.063
N	3631	3631	3631	3631	3631	3631

Standard errors in parentheses: * p<0.1, ** p<0.05, *** p<0.01. Standard errors are clustered at the county level. Regression (1) reports baseline results, and (2) adds covariates related to political party affiliation. Regression (3) covariates include a county level control for the number of daily ambient air quality violations, and population. Regression (4) covariates add controls for education and minority status. Regression (5) covariates add employment status, a dummy for registered voters, age, and gender. Regression (6) covariates also include an individual's support for local and state government, number of vehicles owned, and household size.

termined by "tribal" politics, then there are not any policy levers available.[39] Our results suggests that people do not base their support for a gasoline tax entirely on party-lines. Instead, an equally important determinant for individuals is how the tax directly impacts them, that is, their own self interests. Because of this relationship, policy makers can use this to strategically return the revenue

[39] A growing body of literature shows strong effects of social identity on political behavior (Klar, 2013; McLeish and Oxoby, 2011).

generated under the tax to increase support for the tax itself. In effect, increasing support for the tax is socially beneficial if it can replace CAFE Standards.

3.5.3 Revenue Neutral Tax Policies

Using the relationship between self-interests and tax support, the objective of this section is to design a tax policy that strategically returns the revenue generated by the tax to increase support for the tax itself. This link enables policy-makers to design tax regimes that are "politically-sophisticated." That is, tax regimes that are beneficial to welfare but also have increased levels of political support. We turn to revenue neutrality as our primary method of thinking about alternative tax regimes.

As explained in Section 3.2, the concept of revenue neutrality has garnered significant interest and attention from both political parties in the United States. By definition of revenue neutrality, all of the revenue collected through a tax is returned through various investments and tax cuts. In this section, we use county level tax dividends as a method of "returning" revenue back to individuals.[40] In order to design a revenue neutral tax policy and issue county dividends, we begin by calculating the additional revenue generated by a uniform increase in the price of gasoline. That is,

$$Revenue = [P_g^y + \tau^y] * Q(P_g^y, \tau^y) - P_g^y * Q(P_g^y). \tag{3.7}$$

In Equation 3.7, P_g^y represents the price of gasoline in year y, τ^y is the uniform gasoline tax as determined by Equation 3.2, and $Q(\cdot)$ is the total quantity of gasoline used by all vehicles in the state fleet. Using this increase in revenue, we design a revenue neutral tax regime and analyze its impact on support for a gas tax. Specifically, we create tax dividends and return revenue differentially to counties based on their distribution of costs and benefits. For example, counties with high proportions of "Losers" are assigned large tax dividends. We then include this county-specific return in an individual's benefit calculation, and use our estimates from Section 3.5.2 to

[40] In practice, this can be implemented by local county tax offices. Every year, vehicle owners are required to register their vehicle with their local tax office as well as pay the associated fees. Similarly, a tax dividend can be issued during this process.

Figure 3.3: Revenue Neutrality and Gasoline Tax Support

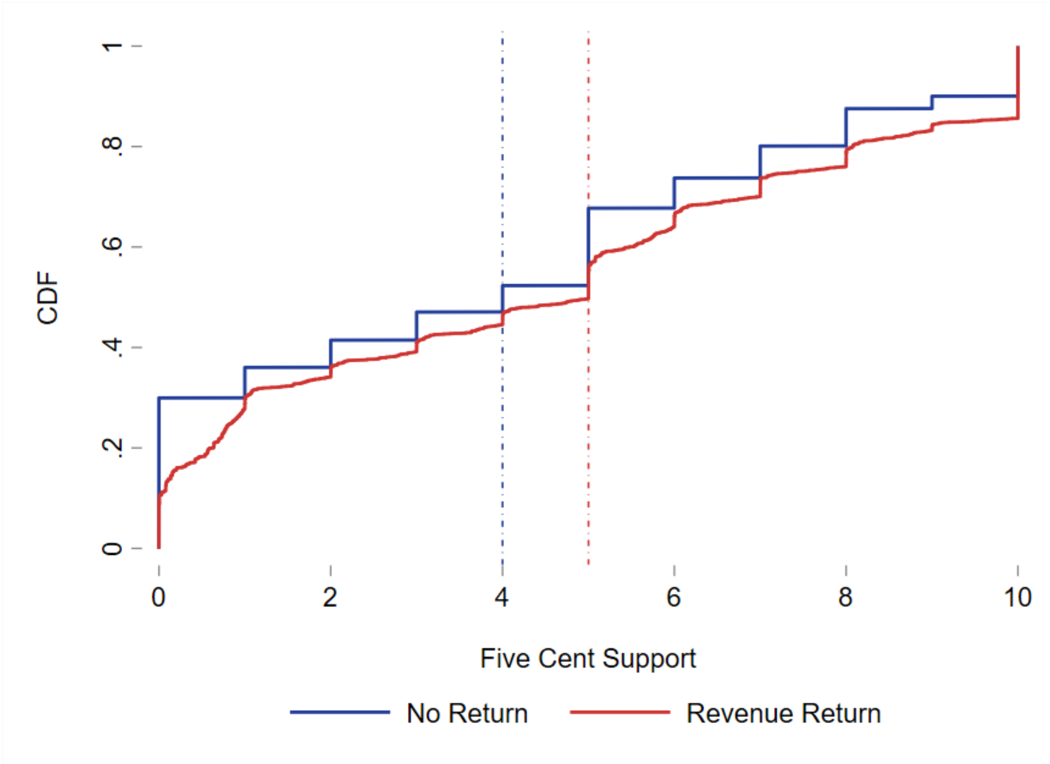

This figure presents results describing the impact of revenue neutrality on support for raising the gasoline tax. Revenue generated from the tax increase is differentially returned to counties with a higher proportion of "Losers". Using estimates obtained in Section 3.5.2, we counterfactually estimate how this would impact an individual's support for the tax itself. Results are shown as CDFs across all levels of support (0-10). Answers between 0 and 4 indicate opposition, 5 indicates indifference, and those between 6 and 10 indicate support.

determine how this would change support for a gas tax. Similarly, counties with small proportions of "Losers" are assigned smaller tax dividends. We determine how this would impact support for a gas tax as well. Figure 3.3 shows the change in tax support under the revenue neutral tax regime previously described.

Figure 3.3 shows the distribution across the raw scale of support for an increase of 5 cents to the gasoline tax. Options for this answer range from 0 to 10, with 0 - 4 indicating decreasing opposition, 5 indicating a neutral position, and 6-10 indicating increasing support. The blue distribution ("No Return") shows the actual distribution. On average, individuals are slightly in opposition of increasing the tax $(mean = 3.95)$, with a median voter at 4 (as indicated by the blue dashed lined).

The red distribution ("Revenue Return") shows the distribution of support after implementing the revenue neutral tax policy as described in this section. This second distribution has shifted significantly to the right relative to the first, indicating that people become more supportive of a tax increase under this policy. Importantly, this tax regime is also responsible for shifting the median voter to a 5 (as indicated by the red dashed lined).

There is a large literature showing the importance of the median voter in political systems. The Hotelling-Downs model of median voters predicts and shows that political candidates take on the positions of the median voter of the entire population even if the median voter of their respective party changes (Hotelling, 1929; Downs, 1957; J. Osborne, 1995).[41] With this in mind, shifting the median voter in our population from a 4 (someone in opposition to a gasoline tax) to a 5 (someone who is indifferent) represents a significant and meaningful change in the level of support for a Pigouvian tax. Given the large efficiency losses of CAFE Standards relative to Pigouvian taxes, garnering the political support to transition towards gasoline taxes represents the potential for significant gains in societal welfare. Our results indicate that this is potentially feasible with "politically-sophisticated" tax regimes designed to strategically increase support.

3.6 Conclusion

Given the large cost of driving and the externalities imposed by this market on society, it is important to be able to optimally address them. The United States has opted to address this through CAFE Standards rather than Pigouvian taxes. CAFE Standards do not correct for the intensive margin of driving - the decision of how much to drive and how much gasoline to purchase. On the other hand, an optimal Pigouvian tax accounts for both margins of driving. By using CAFE standards and not accounting for both margins, an inefficiency is created. However, the lack of political support a gasoline tax receives makes it difficult for policy-makers and government officials to act otherwise. Public opinion has never been favorable towards gasoline taxes. Survey evidence shows

[41]The Hotelling-Downs model predicts that in a two-candidate election, each candidate should take the positions of the median voter to increase their respective probability of winning (Downs, 1957; Hotelling, 1929). This result is robust under a few conditions, namely, that voter preferences are single-peaked and that the number of candidates does not exceed two (J. Osborne, 1995). This is the case for primary elections, after each candidate secures the vote from their respective party. For party elections, candidates adopt the positions of the party's median voter first.

that this sentiment is true currently and has largely been true throughout the history of the United States.

This paper answers three main questions. First, what are the distributional consequences of a uniform gasoline tax? We combine data from the Texas DMV with estimates on marginal damages to create a uniform tax. We impose this tax in the state of Texas, estimate changes in driving decisions, and document the distribution of costs and benefits. Overall, there is significant heterogeneity across the state, but most of the benefits ("Winners") are consolidated in more urban regions. Smaller and more rural counties have very little to gain from a uniform gasoline tax.

Second, does this distribution explain support (or the lack thereof) for a Pigouvian gasoline tax? Using a stratified survey across the state we show that this distribution is a strong predictor for an individual's support for a gasoline tax. Importantly, this fact remains true even after controlling for an individual's political orientation along with many other potential confounders.

Third, can revenue neutral tax regimes positiely shift the distribution of support for gasoline taxes. We show that through revenue neutrality, policy-makers have the potential to shift the distribution of "Winners" in a positive direction. This result shows that there is the potential to design "politically-sophisticated" tax regimes with revenue neutrality that have increased levels of support for a Pigouvian gasoline tax. By designing tax regimes with revenue neutral components, there are ways to motivate individuals based on their self-interests which in turn can be translated into support for a gas tax.

4. PRICE LEADERSHIP AND LEARNING IN OLIGOPOLY: EVIDENCE FROM ELECTRICITY MARKETS

4.1 Introduction

Multiple equilibria are pervasive in many markets and their existence complicates both the design of these markets and the role of regulation within them. This is particularly the case for oligopoly markets where firms choose supply functions rather than either just price or quantity. In these markets, there exists a wide range of potential equilibria resulting in a similarly wide range of market outcomes. For example, any outcome ranging between Bertrand and Cournot competition can be sustained in equilibrium (Klemperer and Meyer, 1989). This implies that there are a significant number of prices above marginal cost that can be sustained as well. Differences in prices, and ultimately profits, generates the incentive for firms to transition away from low price equilibria towards high price equilibria. Despite this, surprisingly little is known about how an equilibrium is reached as well as how agents transition between them.

Learning more about the transition process has market design implications about how to optimally disclose and reveal information to the market. In some markets, information about prices, auction results, and the actions firms take is not made available in real-time. It is often the case that there is a lag until a firm can observe these historical results (this process is referred to as information disclosure). Being one of the few levers available to policy-makers, setting optimal information disclosure policies can help regulators take an ex-ante approach to preventing the formation of some equilibria. This is potentially very important for one large class of equilibria - namely, those falling under the umbrella of collusion, coordination, and cooperation - which have received a large share of policy interest. In order to protect consumers from the direct impacts of these equilibria, antitrust policy primarily relies on theories about the market conditions that facilitate these equilibria. Furthermore, case studies are largely used to study the behaviors of economic agents. Given the importance of the topic, little is known about how the transition between various equilibrium is facilitated.

In this paper, I study the role of price leadership and learning in the establishment of a new equilibrium and the transition towards it in the Texas wholesale electricity market during the summer of 2013. The Texas electricity market is useful for this study because it is comprised of a few private firms in a highly restricted market that compete over a homogeneous product with inelastic demand, who share similar cost structures.[1] In addition to this, the existence of rich and high-frequency data on both firm behavior and marginal cost provides a unique opportunity to study how firms in an oligopoly learn and transition between equilibrium.

I first develop and use a model of static unilateral profit maximization and show that a small firm in the market deviates from a low-price supply function equilibrium (SFE) and foregoes static profits (≈$3,500 per offer) in favor of over-pricing a significant portion of its production. After multiple deviations by the price leader, the largest firm in the market begins to follow a similar behavior and over-prices its production (forgoing ≈$1,200 per offer). As both firms iterate and learn about each other's actions, they are able to reduce their production and set prices associated with a high-price equilibrium. I show that this shift in equilibrium play corresponds to an average price increase of 5%, but can importantly result in swings as large as 1,500%. I find preliminary evidence that these outcomes are consistent with a dynamic repeated game equilibrium.

To explicitly analyze the process of learning during the transition period, I integrate a fictitious play learning model into a model of dynamic profit maximization. This characterizes a fixed point equilibrium regarding the beliefs that each firm must have about the other in order for it to be optimal to initially deviate. Since these belief parameters are used by firms to form expectations about each other's current and future actions, there is room for policy-makers to limit available information through the form of information disclosure policies. These types of policies define a procedure where market information is released with a lag. By preventing more recent information from being observable, firms would need to forgo an increased amount of static profits to transition to the high price equilibrium. In fact, I estimate that revealing information with a 10-day lag

[1]These are shown to be key characteristics that facilitate oligopoly behavior, namely that of price leadership (Markham, 1951)

prevents firms from transitioning to the high priced equilibrium altogether.

This paper adds to several branches of the economics literature. To the best of my knowledge, there is only one other paper that studies the explicit initiation of a new equilibrium and the process by which firms transition towards it. Byrne and de Roos (2019) study the retail gasoline market in Perth, Australia in which they document initial behavior changes by a price leader. Specifically, they document how a price leader experiments with different focal point prices as a coordinating device to reveal its collusive intent to other firms. Over time, other retailers in the market learn to recognize these focal point price "signals" and adopt a similar pricing behavior as the price leader. Differing from this study, I study the role of price leadership and learning in the context of firms competing in supply functions rather than choosing a single price. This arguably reveals a higher level of strategic interaction given the added complexity of the environment for firms to behave in.

I also contribute to the growing literature on the effects of price leadership and coordination in general. In the U.S. telecommunications industry, Busse (2000) finds that firms with multi-market contact coordinate price increases between 7%-10% by both enhancing the ability to punish as well as increasing the scope of leadership. Retail gasoline is a market saturated with research on this topic. While estimates vary significantly depending on the environment, the presence of a price leader increases margins anywhere from 0.6% to 75% (Byrne and de Roos, 2019; Lemus and Luco, 2018; Lewis, 2012). Chilet (2018) studies the Chilean pharmaceutical market and identifies coordinated price increases ranging between 1-132 percentage points. Similar to the setting of my paper, Chilet finds that successful price leadership is facilitated by the smallest firm. Finally, Miller et al. (2018) develop a model of oligopolistic price leadership in order to study the role of price leadership in the U.S. beer industry in the context of mergers. In comparison to these studies, I am one of the first to study price leadership in electricity markets while also providing some of the largest estimates that this behavior can potentially have on prices (1,500%).

My research also builds on the study of oligopoly behavior in wholesale electricity markets. Notably, Hortacsu and Puller (2008) develop a model of unilateral best response to study firm behavior during the early phases of the market in 2002. Their research rejects equilibrium behavior and concludes that firms are over-pricing their production (bidding too-steep) relative to static profit maximization. In order to explain this, Hortacsu et al. (2019) incorporate a model of cognitive hierarchy and show that this discrepancy is in part due to heterogeneity in strategic behavior. My research studies the market more than 10 years after these early findings and shows that another explanation for firms over-pricing their production can be to transition between equilibria. This result is similar is vain to some of the work that studies market power in the California electricity market (Borenstein et al., 2002). Doraszelski et al. (2018) study the frequency response market in the UK electricity system and show evidence that firms learn about demand as well as the behaviors and actions of other firms. I also show evidence that firms learn about each other's behaviors, but can use this to transition between multiple equilibria.

4.2 Background on the Texas Wholesale Electricity Market

The Electric Reliability Council of Texas (ERCOT) manages electricity for more than 25 million people in Texas, directing over 46,500 miles of transmission and more than 600 individual generators. Its real-time spot market operates as an auction in which participants bid to buy and offer to sell energy. Every 5 minutes, ERCOT clears these bids and offers at the lowest cost of production conditional on transmission constraints and lines loss. In the absence of these two constraints, the market clearing price is uniform. I analyze behavior in this spot market.

The basic mechanics of this market can be summarized as follows. Every 15 minutes, participants selling electricity submit offer curves specifying the quantity they are willing to produce at every price.[2] The market operator, ERCOT, aggregates these offers and clears the market subject to minimizing the cost of production every 5 minutes. ERCOT thus establishes a market clearing price that equates supply and demand. When it comes time to settle however, not every firm is paid this market price for every megawatt it produces. A firm may have other pre-determined quantities and prices that it gets paid. Primarily, these are bilateral contracts and financial positions. Bilateral contracts are pre-specified contracts held between an electric supplier and a buyer. Financial positions can vary significantly in terms of length. The most common financial position a firm agrees

[2] In practice, supply offers are bounded between prices of -\$250 and \$5,000.

to comes in the Day-Ahead Market. In the day-ahead market, firms choose to submit offers to sell electricity for the following day. ERCOT follows a similar procedure explained above and establishes a day-ahead price for which a firm is guaranteed the following day. The rationale behind a day-ahead market is to provide market participants a mechanism to voluntarily hedge against the volatility of prices in the real-time market.[3]

4.2.1 Setting Descriptives

The focus of this paper revolves around the summer of 2013 in the Texas wholesale electricity market. Namely, the observation of several key features leads itself towards the study of equilibrium transition and changes in firm behavior during this sample period. First, as mentioned before, this is a market with a few number of private firms competing in a restricted market over a homogeneous product with inelastic demand, who all share similar cost structures. In order to participate in the market, a firm needs to either meet significant credit constraints or have ownership of a power plant itself, which also requires significant credit. The inelasticity of demand is due to the fact that almost all consumers pay a fixed price for electricity, so consumer behavior is unresponsive to wholesale prices.[4] Lastly, the similar cost structures are due to the fact that generator technology is shared between firms. The literature shows that these are all key characteristics that facilitate successful price leadership (Markham, 1951). Summary statistics for the largest ten firms in the market are shown in Table 4.1.

As indicated by Table 4.1, 57% of total generation is on average met by the largest 5 firms, and may even be as large as 80% in a given auction. The largest two firms in the market (Luminant and NRG) own over 27,000 MWs of combined capacity across 82 generating units, representing about 35% of total output and owning over 40% of ERCOT's installed capacity. Smaller and medium-

[3]According to ERCOT, "(t)he Day-Ahead Market (DAM) is a voluntary, financially-binding forward energy market. The DAM matches willing buyers and sellers, subject to network security and other constraints, whereby energy is co-optimized with Ancillary Services and certain Congestion Revenue Rights. It provides a platform to hedge congestion costs in the day-ahead of the Operating Day, and instruments to mitigate the risk of price volatility in Real-Time."

[4]It should be noted that there are emerging services (i.e., Griddy Energy) that enables a consumer to purchase electricity directly at the wholesale price. However, these services began in 2017 and have had slow take-up relative to the total market share.

Table 4.1: Firm Summary Statistics

Firm	Resources	Capacity (MWs)	Avg. Mkt. Share (%)	Std. Dev	Min	Max
Luminant	42	16601.5	24.2	3	17.6	34.2
NRG Energy	40	10307.82	10	1.5	5.6	13.9
Calpine	15	7205	8.3	2.5	2.3	13
CPS Energy	26	6404	7.5	0.9	3.4	10.3
Nextera Energy	26	5145.7	4.9	1.4	0	8.5
GDF Suez	14	4127	6.7	0.8	4	8.7
Exelon	11	3255	2.8	0.6	0.8	4.9
LCRA	23	2524	2.1	0.3	1.3	3.3
City of Austin	20	2424.7	1.3	0.5	0	3
Direct Energy	8	2030	2.3	0.5	0.6	4
Constellation Energy	5	1276	0.4	0.5	0	2.1

This table presents summary statistics for the largest firms in the market. Capacity is measured by the maximum quantity offered at any price in any auction over the sample period. Average Market Share is calculated from each firm's telemetered net output as a fraction of the total market output. This is calculated for each auction. The final two columns provide min/max statistics for the market share calculations.

sized private firms (Calpine, Exelon, GDF, Nextera) produce about 23% of total generation while owning about 30% of installed capacity. Lastly, municipalities and co-ops (City of Austin, CPS, LCRA) represent a smaller average market share of about 11% despite owning close to 18% of installed capacity. This highly dense market structure suggests a large role any one or handful of firms can have on the market as a whole at potentially any point in time.

Furthermore, this time period also observes several significant price events. Figure 4.1 plots a time series over the sample period of the average market price. In this paper, the average market price, P_t, is defined as the load-weighted average price across all nodes in the market.[5] The average market-wide price, P_t, is \$31 across all hours of the day. However, there are also times when prices spike to be anywhere between \$100 and \$800.

High prices by themselves are not by all means an indication of behavior changes or strategic interactions. Especially in electricity markets, there are a few potential explanations for price spikes that might be the product of supply and demand in a particular auction. These can come from demand-side shocks, any shocks to supply (generator outages, transmission failures, etc.), as

[5]This is common practice when studying nodal electricity markets.

Figure 4.1: Market Price

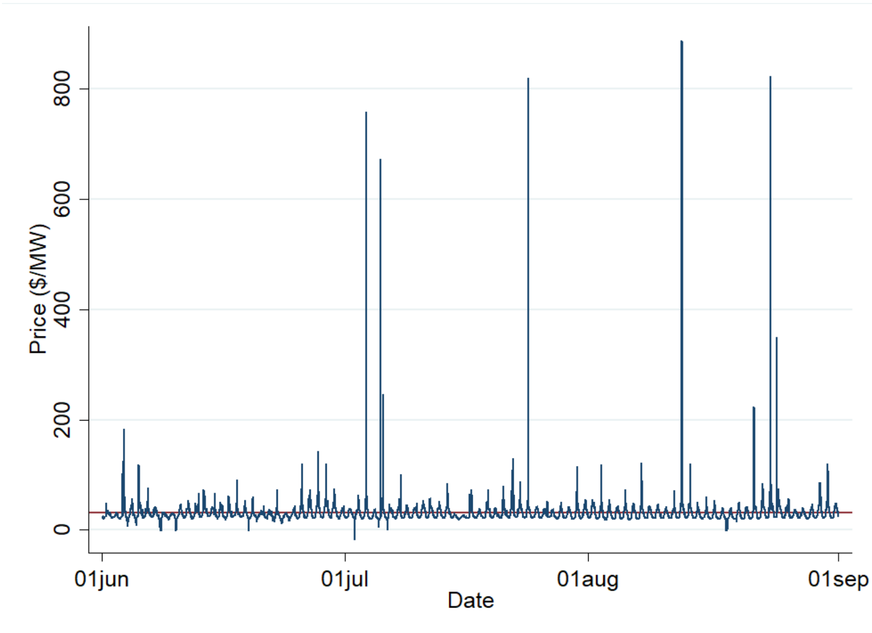

Notes: This figure plots the average market price, P_t, for the summer months of 2013. The horizontal red line refers to the average summer price ($31).

well as more technical engineering constraints that electricity grids are subject to. Allowing prices to properly reflect these conditions is an important component of a well-designed and functioning market. However, despite this, demand and supply during this sample period do not predict significant price events.[6] Figure 4.2 shows a similar time series for these market characteristics.

Panel a) of Figure 4.2 shows the capacity available to the market operator in the real-time spot market. Panel b) shows the realized demand in each auction. And lastly, panel c) depicts a time series of both coal and gas prices.[7] These summary figures indicate that there are not any unexpected shocks to supply or demand that could potentially explain the significant price events.[8]

Looking at this time period more closely, I turn towards the study of firm behavior and the

[6]Appendix Figure C.2.1 shows a comparison between actual prices and predicted prices.

[7]Appendix Figure C.2.2 shows hourly summary figures for a) the market price, b) available capacity, and c) aggregate demand. As demand increases throughout the course of a day, so does online capacity and market prices accordingly.

[8]Similarly, I also test for congestion as an explanation for these price events and do not find any significant correlation. Using the variation of nodal prices as a proxy for congestion, I screen out any potentially congested time period. The significant price events remain even after screening out auctions where the quantity-weighted standard deviation is more than 5-10%.

Figure 4.2: Market Summary

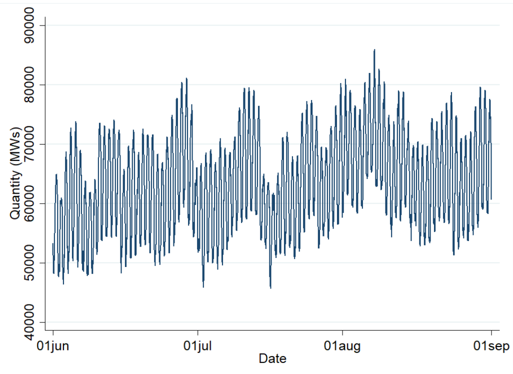

(a) Capacity

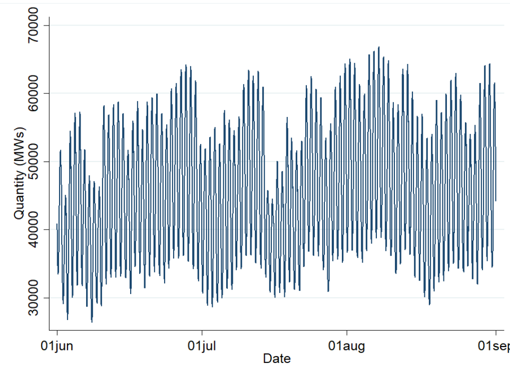

(b) Aggregate Demand

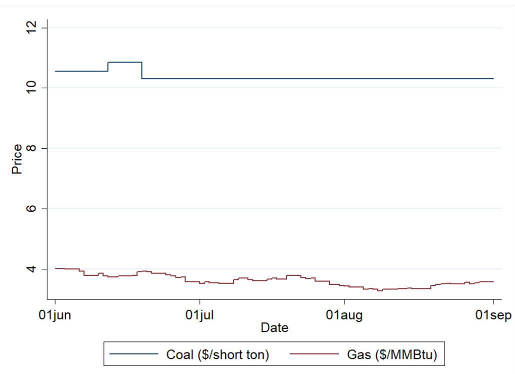

(c) Fuel Costs

Notes: These figures plot time series of important market components. Panel a) shows the total capacity being offered into each auction, while Panel b) shows aggregate demand. Panel c) shows fuel prices for coal and natural gas.

interactions between firms. As mentioned before, this is in the context of firms submitting supply functions. In highly dense oligopoly markets such as electricity markets, the way in which firms offer their production, either individually or in the aggregate, has a direct impact on the market clearing price. I first use a model of static unilateral profit maximization to characterize equilibrium behavior. This model takes into account both supply and demand from the perspective of the firm and allows me to analyze changes in behavior, specifically when it comes to price leadership and learning in the Texas wholesale electricity market during this time period.

4.3 Methodology and Data

4.3.1 Methodology

I start by developing an equilibrium model of bidding in the ERCOT spot market that incorporates the uncertainty faced by each firm at the time of making an offer into the market. This process follows the uniform price share auction setup of Wilson (1979), and has been built upon frequently in the literature (Hortacsu and Puller, 2008; Hortacsu et al., 2019; Mercadal, 2018).

4.3.1.1 Demand

I begin by modeling demand as

$$D_t(p) = d_t(p) + \epsilon_t \tag{4.1}$$

where $d_t(p)$ is a deterministic component and ϵ_t is a stochastic shock. This is a natural specification due to the nature of demand in wholesale electricity markets. Most consumption pay a fixed rate per MWh and are thus price insensitive to the short-term market clearing process. On top of this, there are no demand-side bids in the spot market and ERCOT clears according to demand forecasts.

4.3.1.2 Firm Offers

For every 15 minute auction t, firms must make decisions about how much to produce. Each firm submits an offer $S_{it}(p)$ into the spot market auction.[9] These offers specify the quantity the firm is willing to produce at every price.[10] After clearing the market and establishing a price, p_t^*, a firm is given dispatch orders specifying the quantity it has been cleared to produce, $S_{it}(p_t^*)$.

Furthermore, when offering into the spot market, each firm must take into account any financial or bilateral contracts that it may have. These contracts specify a price, PC_{it}, and quantity, QC_{it}, at which a firm has locked in place. Accounting for these contract positions is important due to the impact it has on the incentives for a firm to offer into the spot market (Wolak, 2000,

[9] I generate firm-level offer curves by aggregating over every generating resource that the firm owns.

[10] In practice, ERCOT allows each generator to submit an offer curve with 20 price-quantity "tranches." However, after aggregating across a firm's generating fleet, a firm has available many steps available to them in their supply function.

2003; Green, 1999). For example, if the market clearing price, p_t^*, is greater than a firm's contract price, PC_{it}, then the firm is required to pay the difference $(p_t^* - PC_{it})QC_{it}$ to the contract buyer (i.e., the firm is a net buyer). If on the other hand p_t^* is less than PC_{it}, then the firm is paid the difference $(p_t^* - PC_{it})QC_{it}$ (i.e., the firm is a net seller). These contract positions are settled as contracts-for-differences. Hence, a firm's offer into the spot market depends on this contract position, $S_{it}(p, QC_{it})$.

Lastly, each firm has a cost function $C_{it}(q)$ for every level of output q. I assume that each firm knows the cost structure of all firms in the market.[11] This assumption is not strong since firms interact between each other frequently and have been doing so for years on end. On top of this, most of the cost is generated by the technological characteristics of the generating unit itself. This information is largely publicly available.

4.3.1.3 Market Clearing

In the absence of transmission constraints, the market clearing price in a nodal market is uniform across the grid. Market operators collect every firm's offers into the auction and equate this to demand in a least-cost manner,

$$\sum_{i=1}^{N} S_{it}(p_t^*, QC_{it}) = D_t(p_t^*) \tag{4.2}$$

where p_t^* is the equilibrium price that each firm is paid.

4.3.1.4 The Firm's Problem

Upon realization of p_t^*, each firm is paid $S_{it}(p_t^*, QC_{it})p_t^*$ due to the uniform pricing rule. This means that firm i's ex-post profit in auction t is

$$\pi_{it} = S_{it}(p_t^*, QC_{it})p_t^* - C_{it}(S_{it}(p_t^*)) - (p_t^* - PC_{it})QC_{it} \tag{4.3}$$

where, as indicated earlier, $(p_t^* - PC_{it})QC_{it}$ takes into consideration a firm's contract position.

[11] This is a standard assumption made frequently in the literature (Hortacsu and Puller, 2008; Hortacsu et al., 2019; Mercadal, 2018)

The most important source of uncertainty in the profit equation comes from the market clearing price p_t^*. This uncertainty has two components, but can ultimately be summarized by uncertainty over the residual demand the firm faces in auction t.[12] First, a firm has uncertainty about the demand shock ϵ_t. Second, each firm faces uncertainty about all other firm's contract positions.

Following Hortacsu and Puller (2008), I define a probability measure over the realization over the clearing price from the perspective of firm i. This measure is conditional on i's private information about it's own contract position, i's submission of offer $S_{it}(p, QC_{it})$, and other firms playing their equilibrium strategies $\{S_{jt}(p, QC_{jt}), j \neq i\}$. Here,

$$H(p, S_{it}(p); QC_{it}) \equiv Pr(p_t^* \leq p | QC_{it}, S_{it}) \tag{4.4}$$

represents the uncertainty that firm i faces. Given i's contract position, this is the probability that firm i will be paid price p when selling quantity $S_{it}(p)$ and all other firms submit their equilibrium offers. Hortacsu and Puller (2008) show that this can be substituted back into the firm's profit problem and solved. Furthermore, if bids are additively separable in price and contract positions, then a firm's optimality condition can be summarized by

$$p - C_{it}^{'}(S_{it}) = \frac{S_{it} - QC_{it}}{-RD_{it}^{'}(p)}. \tag{4.5}$$

This resulting condition essentially corresponds to an optimal markup rule, $p - C_{it}^{'}(S_{it})$. In practice, this rule depends on whether or not a firm is a buyer or seller (and by how much), $S_{it} - QC_{it}$, scaled by the firm's market power, $RD_{it}^{'}(p)$, where $RD_{it}^{'}(p)$ is the slope of firm i's residual demand curve. Observe that in order to solve for this condition, one would need to observe each firm's contract position. In order to do this, I rely on a common assumption made in this literature. That is, if $C_{it}^{'}(S)$ is observed, a firm's contract position can be determined by the quantity where a firm's offer curve intersects its marginal cost curve.[13]

[12] Firm i's residual demand curve, $RD_{it}(p)$, is determined as $RD_{it}(p) = D_t(p) - \sum_j^N S_{jt}(p, QC_{jt})$, such that $j \neq i$.

[13] See Hortacsu and Puller (2008); Hortacsu et al. (2019); Mercadal (2018).

4.3.2 Data

This paper uses data from various sources that can be summarized into two main categories - real-time spot market data and marginal cost data.

Real-Time Spot Market data is publicly available from ERCOT. I first obtain each firm's offer curve made into ERCOT's Security Constrained Economic Dispatch software (SCED) from ERCOT's 60-Day SCED Disclosure Reports. These reports contain detailed information on every firms's 15 minute offer, including the price-quantity values that make up an actual supply offer curve, ramp-rates, generator start-up costs, facility status, cleared quantity orders, maximum and minimum sustained energy production capabilities, etc. I also make use of ERCOT's 2-Day Load Reports, which provide aggregate measures of realized demand for each 15-minute interval.

In order to obtain price data, I rely on ERCOT's Locational Marginal Price (LMP) reports. As of 2010, ERCOT operates as a nodal market meaning that each node on the grid has its own price (i.e., the LMP).[14] LMP's are constructed from three components; the energy cost, congestion cost, and line loss. Specifically,

$$LMP_{nt} = MEC_t + MCC_{nt} + MLC_{nt}, \tag{4.6}$$

where LMP_{nt} is the locational marginal price of node n in period t, MEC_t is the marginal energy cost of adding one more MW to the grid period in period t, MCC_{nt} is the congestion cost at node n in period t, and MLC_{nt} is the marginal line loss at node n in period t. Observe that in the absence of congestion and loss, the nodal price is uniform across the grid, that is $LMP_{nt} = MEC_t$. Following this, I construct the market price as the load-weighted average across all nodes,[15]

$$P_t = \sum_{n}^{N} \frac{q_{nt}}{Q_t} LMP_{nt}, \tag{4.7}$$

[14] See Appendix Figure C.1.1 for an example of prices with no grid congestion and Appendix Figure C.1.2 for an example with congestion.

[15] This practice is common when analyzing nodal markets (Bushnell et al., 2008).

where q_{nt} represents that quantity produced at node n in period t, Q_t is the total quantity of electricity in period t, and LMP_{nt} is defined as in Equation 4.6. Since I do not explicitly model the transmission grid, and thus potential congestion constraints, I use the variation of the nodal prices to screen out potentially congested periods.[16,17]

As previously mentioned, the relevant marginal costs that firms must consider when deciding how much to produce can be obtained directly from data. In order to calculate marginal cost curves for each generating resource, I combine data from three main sources. First, I use data from the U.S. Energy Information Administration (EIA) to collect data on fuel prices. Coal prices ($/short ton) are constructed from the Powder River Basin daily spot price and natural gas prices ($/MMBtu) are constructed from the Henry Hub daily spot price.[18] I also account for transportation costs by using the EIA's transportation cost estimates.

Furthermore, I obtain data on a generator's heat rate from the U.S. Environmental Protection Agency's (EPA) Emissions and Generation Resource Integrated Database (eGRID) as well as data from ERCOT directly. Sometimes this information is incomplete in which case I augment the heat rate data with the EIA's Form 923. I also obtain data from the Texas Commission on Environmental Quality (TCEQ) in order to get cost estimates on pollution permits. Lastly, I use data from the Federal Energy Regulatory Commission (FERC) Form 1 that gives operation and maintenance cost estimates. These combined costs represent the bulk of the share of a firm's total cost of production.

4.4 Price Leadership

Price leadership has been shown to have significant impacts on profits, margins, and consumer welfare across a variety of industries (Busse, 2000; Byrne and de Roos, 2019; Lemus and Luco, 2018; Lewis, 2012; Miller et al., 2018). Given this, it is important to investigate the role of strategic

[16] The literature suggests two primary ways to deal with congestion without having to model the transmission grid. The first method requires researchers to choose a specific hour of the day when congestion is likely low (Hortacsu and Puller, 2008; Hortacsu et al., 2019), while the other method relies on sample periods in which there is little geographic variation of LMPs (Bushnell et al., 2008; Mansur, 2008). I follow and modify the latter of these two methods. In practice, line loss averages 5% but can be as large as 15%. I screen out periods in which P_t has a standard deviation greater than this range.

[17] At times, I also make use of ERCOT's data on shadow prices and binding transmission constraints.

[18] These hubs are commonly used when studying ERCOT.

behaviors in light of significant price events.

In many markets, regulatory emphasis tends to be placed on larger firms rather than smaller firms (Gates and Leuschner, 2007). However, a recent literature shows the importance of smaller firms in oligopoly behaviors specifically in regards to price leadership (Chilet, 2018). Due to the nature of supply and demand, market power in wholesale electricity markets may exist for any firm depending on the specific market conditions. Despite this fact, ERCOT's regulator (the PUCT) excludes firms with less than 5% of installed capacity from having market-wide market power.[19] Even with this protocol in place, it is not clear from a firm's perspective whether this exclusion applies to them. In particular, defining capacity is difficult and has warranted its fair share of administrative meetings. In order to bypass this uncertainty, a small firm - GDF Suez - enters into a voluntary mitigation plan (VMP) with the PUCT that explicitly defines the firm's capacity as 4.99% of installed capacity, thereby ensuring the firm exclusion from any regulation of market-wide market power.[20]

Shortly after the acceptance of GDF's VMP, there is a notable change in the firm's offering behavior.[21] On ∼June 13th, GDF begins to offer its production between $400 and $1,200 - a noticeable increase from $20-$30 as the firm had historically been doing. Panel a) of Figure 4.3 shows this transition for one of GDF's generators.[22]

Following this significant change in the offering behavior by a GDF generator, Panel b) of Figure 4.3 documents a change in behavior by one owned by Luminant about three days after the fact (∼ June 16th). While maybe not as drastic, this still represents a change from historically being willing to sell electricity for about $20 to sporadically increasing its offer anywhere from $100

[19]Firms with generation portfolios less than 5% of installed capacity are considered "Small Fish" and are regulated according to the "Small Fish Swim Free Rule" (hereafter SFSF) rule. SFSF is an Electric Substantive Rule implemented by the PUCT on August 13th, 2006. Small Fish are deemed not to have ERCOT-wide market power, and therefore have "an absolute defense against an allegation of an abuse of market power through economic withholding with respect to real-time energy offers up to and including the system-wide offer cap." See Subst. R. §25.504(c).

[20]A voluntary mitigation plan is an agreement between a market participant and the PUCT explicitly stating actions, conditions, and/or behavior as not being an abuse of market power. See Subst. R. §25.504(d).

[21]A firm with 5% installed capacity could potentially have ERCOT-wide market power beginning in mid-June. There is an abundance of excess capacity in lower-demand months meaning one firm has a minimal impact. But in hotter periods, almost every firm could be a marginal supplier. See Appendix Table C.2.3.

[22]Appendix Table C.2.4 shows a longer summary figure of GDF's historical offering behavior.

Figure 4.3: Departure from Historical Behavior

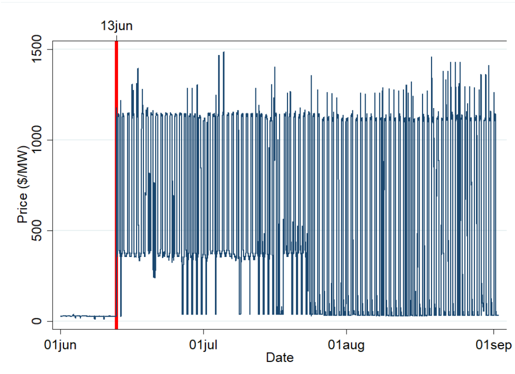

(a) GDF Average Price

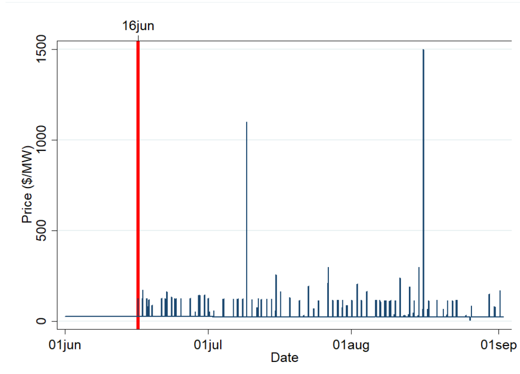

(b) LUM Average Price

Notes: This figure plots the historical average offer price, for the summer months of 2013 by two generators. Panel a) plots a GDF generator whereas Panel b) plots a Luminant generator. The average offer price is calculated for each offer by quantity weighting each price-tranche in the generator's offer curve. By doing so, each offer curve is collapsed into a single observation and the time series of these is plotted above.

to $1,000. In order to rule out changes in cost or market conditions, I use the model outlined in Section 4.3.1 to characterize each firm's unilateral static profit maximizing offer for each auction. Importantly, this model takes into account both realized demand and supply as well as the uncertainty facing each firm allowing for the construction of optimal static supply offers. I then use this model to compare the model predictions for a firm's ex-post optimal offer with their actual offer. Unlike many other settings where firms choose a single price or quantity, firms submit multiple prices and quantities in order to complete an entire supply function. Because of this, it is important to analyze the full range of prices and not any single one. It could be the case that different behaviors exist at multiple points along the supply function. To the best of my knowledge, this is the first paper that studies price leadership in the context of firms competing in supply functions rather than setting a single price or quantity. Figure 4.4 shows the results of this comparison.

To summarize Figure 4.4, I first construct a firm's ex-post optimal offer and compare this to the firm's actual offer in every auction.[23,24] In order to control for differences in contract positions

[23]These offers, as well as the market clearing process more generally, take into consideration important economic and engineering components, like price-quantity values, generator start-up costs and facility status, maximum/minimum sustained energy production capabilities, etc.

[24]Appendix Figure C.3.1 shows Figure 4.4 examples in price-quantity offer space. Figure 4.4 shows the first $100

over time, I scale the contract position to be zero. Hence, each plotted offer only captures net-selling behavioral differences. After doing this, I take the difference between a firm's optimal offer and their actual offer, $S_{it}^{ExPO} - S_{it}$, and then scale this by their available capacity upon fulfilling their contract positions, $Cap_{it} - QC_{it}$, where Cap_{it} is firm i's total capacity in auction t. Hence, each respective line in Figure 4.4 corresponds to the scaled difference between a single offer and the benchmark offer. Note that for simplicity, I only show the first \$100 for each offer.[25] Therefore, the multiplicity of lines captures the sequential time-series element. By definition, any offer above 0 on the y-axis corresponds to offers that are "too-steep" relative to static best response profit maximization (i.e., the firm is over-pricing) whereas offers that are below 0 on the y-axis correspond to offers that are "too-flat" relative to static profit maximization (i.e., the firm is under-pricing). In summary, any line that is not horizontal at 0 represents an offer that is not profit maximizing for the firm from a static perspective.

Looking at representative samples of firm behavior in Figure 4.4 through the lens of the model described in Section 4.3.1, Panel (a) corresponds to the time period before GDF Suez departs from its historical behavior. In general, these offers can be classified as profit-maximizing from a static ex-post perspective. Similarly, Panel (b) indicates that Luminant also behaves in a profit-maximizing manner.[26] Panel (c) shows that once GDF departs from its historical behavior, it is also deviating from profit maximization and is actively "leaving money on the table" by over-pricing its production. Given that this knowledge does not immediately become publicly available, we expect a natural lag in responses from other firms.[27] Panel (d) shows that in GDF's initial deviating auctions, Luminant does not systematically change its behavior. Panel (e) shows that GDF continues to deviate, and is finally followed in Panel (f) with similar deviations from Luminant once knowledge of GDF's deviations are publicly available to the firm.

In the context of this paper, I define the initial equilibrium (corresponding to Panels (a) and

of the smoothed difference between "Ex-Post Optimal" and "Actual Offer" starting at the firm's contract position.

[25] In practice, a firm's supply function can range between -\$250 and \$5,000 (the offer cap).

[26] There are some instances where Luminant "under-prices" which may be due to risk aversion of being the largest firm in the market.

[27] There is a two day lag in the disclosure of aggregate market information. However, there are a variety of other reports that provide firms with more current information about its rivals production (i.e., the Current Operating Plan).

Figure 4.4: Deviation from Ex-Post Best Response Bidding

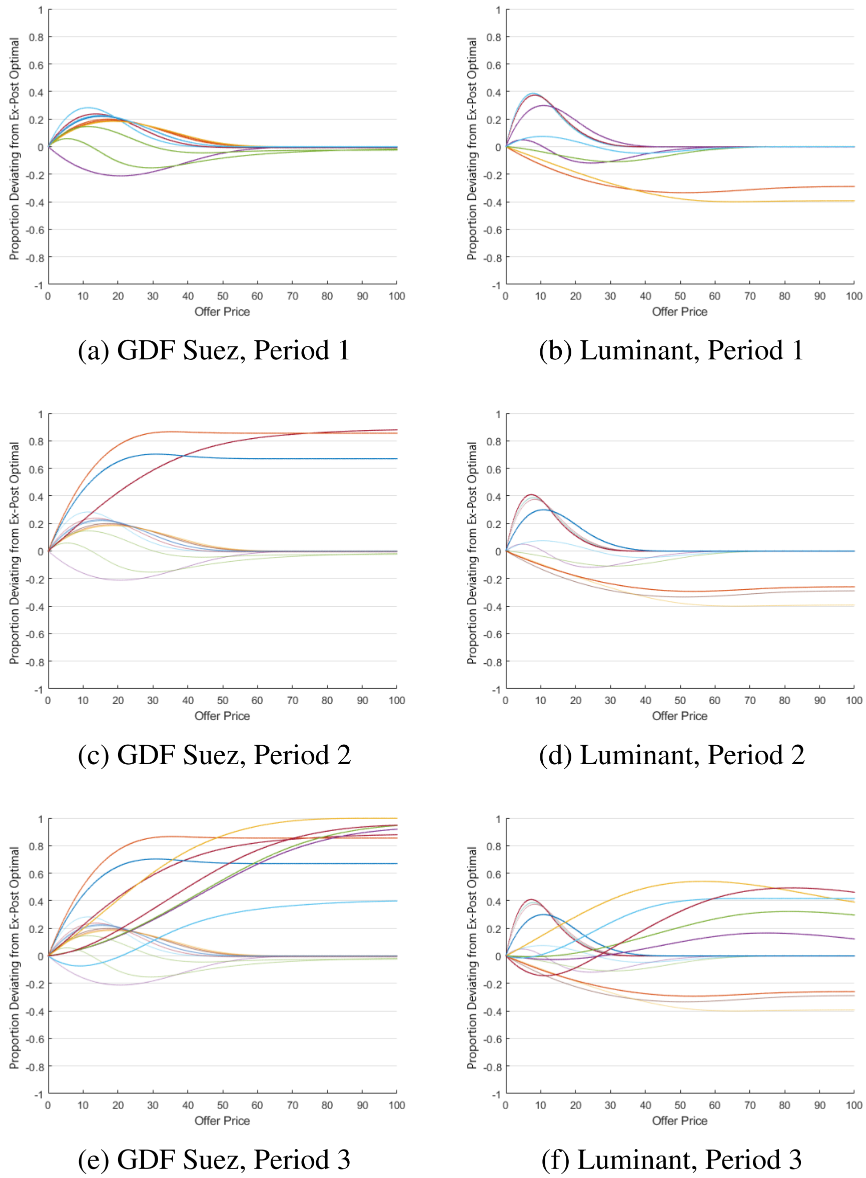

(a) GDF Suez, Period 1

(b) Luminant, Period 1

(c) GDF Suez, Period 2

(d) Luminant, Period 2

(e) GDF Suez, Period 3

(f) Luminant, Period 3

Notes: Each line corresponds to a unique offer and represents the smoothed deviation from ex-post best response. Depicted offers are selected from a representative sample of firm behavior. This is then scaled by the firms available capacity. By definition, anything above 0 is interpreted as an offer that is "over-priced" whereas a line below 0 is interpreted as being "under-priced". Panels (a) and (b) show that initially, both firms are behaving according to ex-post best response. Panel (c) indicates that GDF Suez deviates from this equilibrium play in period 2. Panel (d) shows that while GDF deviates, Luminant continues to behave as it did in Period 1. Panel (e) shows that GDF continues to deviate while in panel (f), Luminant begins to deviate from ex-post best response in period 3.

Figure 4.5: Equilibrium Price Events

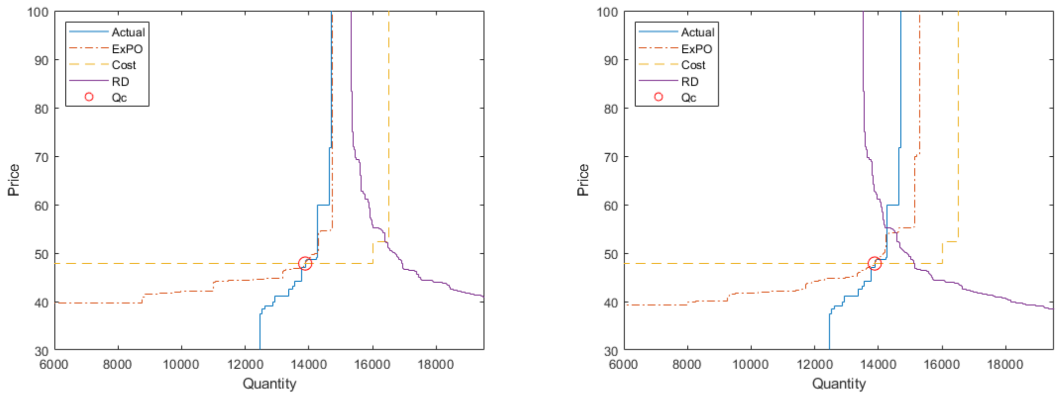

(a) Ex-Post Optimal (b) Counterfactual

Notes: This figure summarizes Luminant's behavior in two different settings. Panel (a) shows realized market data, whereas in Panel (b) I counterfactually replace GDF's offer with their ex-post optimal offer from period 1. Panel (a) indicates that Luminant's actual offer and their ex-post offer align upon fulfillment of their contract position, whereas Panel (b) shows that Luminant's actual offer is "too-steep" relative to their ex-post optimal offer. The results of this figure are mirrored for the same counterfactual but from GDF's perspective.

(b) in Figure 4.4) as a low price supply function equilibrium (SFE). The sequential deviation from static best response behavior by multiple firms can be interpreted as a period of learning by firms towards another more profitable equilibrium. After repeated iterations of this, the firms arrive at a second equilibrium point that more closely resembles the Cournot outcome (vertical offers).[28,29] I define this second equilibrium as a high price equilibrium. One important feature of this second equilibrium is that depending on a firm's residual demand curve, there is a higher probability of price events which generates significant profits to the firms. For example, consider Figure 4.5.

Figure 4.5 provides an example of equilibrium behavior where it is actually in each firm's unilateral interests to offer "vertically" at lower quantities. Panel a) shows that Luminant's actual offer, which is very vertical in nature, lines up with its ex-post optimal offer. However, this would not have been optimal for the firm had GDF never changed behavior in the first place. Panel b)

[28]As shown by Klemperer and Meyer (1989), in the context of equilibrium in supply functions, almost anything between Bertrand and Cournot competition can be sustained as a static equilibrium outcome.

[29]There is evidence that in some auctions, behavior between the two firms is greater than the Cournot outcome. This result is consistent with a dynamic equilibrium where there are repeated game effects.

shows that if GDF offered as it had historically, Luminant would have been better off offering much closer to its marginal cost.

Furthermore, this example shows an important difference between the two equilibrium outcomes. Depending on market conditions, that is tight demand and supply, it may be the case that the two firms are able to set significantly high prices. If their residual demand curves are relatively steep, the two firms can choose prices along these steeper portions. This is the case in Figure 4.5 where the market price is close to $800. This feature of supply functions in electricity markets is one reason for the existence of a multiple equilibria.

4.5 Results

4.5.1 Differences in Equilibrium Price

To summarize so far, there exist multiple equilibria which may be optimal for a firm to play depending on the actions of other firms. In the setting of the Texas wholesale electricity, firms begin playing a low price SFE but depart from this in favor of learning to play a high price equilibrium. In order to show differences between the low price SFE and high price equilibrium, I create an indicator, $I(Post)$, to capture all of the post transition periods and use this indicator to estimate the difference in market prices. Table 4.2 presents the difference in the market clearing price between the two equilibria.

There are a handful of factors that can impact market prices. In order to show that these are not the driving force behind the price differences, I control for the most important features of the electricity market. Namely, Table 4.2 adds controls for fuel costs, demand shocks, and available capacity in columns 2 and 3 respectively. Furthermore, there are certain times of the day as well as days of the week where price may vary in response to consumption patterns. In order to compare prices only between periods of the same hour of the same day, column 4 controls for hour and day fixed effects. Column 4 also controls for various weather elements that can impact electricity consumption. Lastly, one might worry that the market clearing price picks up congestion pricing in the later Summer periods. To address this, I screen for congestion and delete

Table 4.2: Equilibrium Price Differences

	1	2	3	4	5
Post	1.23**	1.38*	1.61**	1.57**	1.17*
	(0.49)	(0.76)	(0.69)	(0.68)	(0.70)
Fuel Cost		X	X	X	X
Demand/Supply			X	X	X
Day Fixed Effects				X	X
Hour Fixed Effects				X	X
Weather				X	X
Transmission					X
N	8729	8729	8729	8729	6368

Standard errors in parentheses. * p<0.1 ** p<0.05 *** p<0.01. Column 1 shows the regression of $I(Post)$ on P_t. Columns 2 and 3 add controls for fuel costs, load, and available capacity respectively. Column 4 includes fixed effects for hour of the day and day of the week and controls for various weather components (temperature, wind speed, rain). In column 5, I screen out potentially congested periods.

any of these potentially congested periods in column 5. Across all of these specifications, a high price equilibrium is associated with an average price increases of close to \$1.50, or a 5% increase relative to the mean (\$31.26). As depicted in Figure 4.5, this type of play can potentially result in significant price increases as large as 1,500%.[30]

4.5.2 Opportunity Cost of Equilibrium Transitions

Despite the resulting price increases, recall Figure 4.4. In order to learn and facilitate this second equilibrium, it required two firms - a leader and a follower - to deviate from a low price SFE. Put another way, both firms had to initially forego profits in order to learn and experiment how to transition. In order to calculate the foregone profits for both the leader and follower, I use two different counterfactuals to obtain appropriate upper and lower bounds. In a deviating auction, I obtain an upper bound on foregone profits by replacing a firm's actual "too-steep" offer with their

[30]For sake of clarity, it should be noted that prices could have been larger during these events had either firm offered at higher prices. There are other factors (e.g. regulatory) that might play an important role in why firms do not choose prices as high as they could.

Table 4.3: Foregone Profits of Transition

	Actual Price (1)	Estimated Price (2)	Best-Response (3)	Period 1 Offer (4)
Panel A: Leader				
Foregone Profits			4332.67	3432.85
			(5267.71)	(5387.083)
Scenario Price	45.67	45.40	43.21	43.92
	(15.13)	(10.03)	(6.43)	(7.63)
Panel B: Follower				
Foregone Profits			1503.44	1205.83
			(2975.41)	(2607.14)
Scenario Price	46.79	47.88	45.82	47.49
	(19.21)	(14.25)	(10.59)	(16.38)

Standard deviations in parentheses. Panel A reports the foregone profits for the Leader during the afternoon offers for the latter half of June, 2013. Panel B reports the foregone profits for the Follower during the afternoon offers for the latter half of June, 2013. Column 1 reports the average market clearing price P_t. Column 2 shows the estimated market clearing price. Column 3 reports the counterfactual impact had each firm offered ex-post optimal supply curves in Periods 2 and 3 whereas Column 4 reports the counterfactual impact had either firm offered a randomly selected offer from Period 1 in Periods 2 and 3. In this light, Column 3 can be thought of as an upper-bound and Column 4 can be thought of as a lower bound for the impact of each firm's deviations.

ex-post optimal offer and calculate the change in profits. However, this may be an overestimate due to the implicit assumption of perfect foresight. To obtain a lower bound, I replace a firm's actual "too-steep" offer with a randomly selected offer from the period prior to deviations. Following this, I calculate the difference in profits. Columns 3 and 4 of Panel (A) and (B) of Table 4.3 show the upper and lower bounds respectively for the leader (GDF) and follower (Luminant) respectively.

Column 3 of Panel (A) in Table 4.3 indicates that an upper bound on the leader's foregone profits is about $4,300 whereas column 4 indicates a lower bound of $3,400.[31] Furthermore, by

[31] Appendix Figure C.4.1 Panel (a) shows the density of foregone profits for GDF. I screen out potentially congested periods and report similar estimates in Panel (b) of Appendix Table C.4.1.

over-pricing its generation, the leader's actions actually result in higher prices for all other market participants. Had GDF not deviated, I estimate that prices would have been \$1.48 - \$2.19 lower than actual. These differences are displayed in Panel (a) of Table 4.3 under the "Scenario Price" row.[32]

Column 3 of Panel (b) in Table 4.3 indicates that an upper bound on the follower's foregone profits is about \$1,500 whereas column 4 indicates a lower bound of \$1,200.[33] Furthermore, by over-pricing its generation, the follower is also generating higher prices for all other market participants. Had Luminant not deviated, I estimate that prices would have been \$0.39 - \$2.06 lower than predicted. These differences are displayed in Table 4.3 under the "Scenario Price" row.[34]

When comparing the difference in the magnitude of foregone profits between the leader and follower, it is worth pointing out that the larger firm (the follower) actually foregoes fewer profits than the smaller firm (the leader). This is a feature of the larger firm submitting offer curves that typically start deviating from profit maximization further along the schedule. By doing so, the firm is able to retain profits upon market clearing while also deviating at other points in its offer. This is a unique feature of price leadership and learning in supply-function space relative to a price setting environment. Highly sophisticated firms are able to employ a diverse set of strategies at many points in its supply function.[35]

4.5.3 Consumer Surplus

Given the nature of the electricity market, changes in price correspond to a transfer between consumers and producers. In order to quantify the size of this transfer, I calculate the change in total market profits had the two firms not changed their behaviors.[36] However, for simplicity, I only

[32] Appendix Figure C.4.2 Panel (a) shows the density of how GDF's behavior impacted the market clearing price.

[33] Appendix Figure C.4.1 Panel (b) shows the density of foregone profits for Luminant. Estimates are conditional on Luminant being a net-seller in the spot market. I screen out potentially congested periods and report similar estimates in Panel (b) of Appendix Table C.4.2

[34] Appendix Figure C.4.2 Panel (b) shows the density of how Luminant's behavior impacted the market clearing price.

[35] This is consistent with the literature showing that larger firms in this market are more sophisticated strategically (Hortacsu et al., 2019).

[36] It should be noted that this is a back-of-the-envelope characterization. I do not model pass-through or competition at the retail level. The complete function relating prices from the wholesale market to retail provision is a very intricate process. However, there are emerging services (i.e., Griddy Energy) that actually enables a consumer to purchase

Table 4.4: Changes in Profits

	Mean	Std. Dev.	Min	Max
GDF	82.2	26.1	61.5	120.4
LUM	459.0	85.6	376.5	546.3
Market	2,285.8	227.9	1,945.3	2,427.3

This table reports summary statistics (presented in thousands, $‘000) for the profits generated under the price events as described in Figure 4.5. A single price event, which is more probable to occur when both firm play according to the high price equilibrium, generates significant profits for both the leader and follower making up for those foregone as described in Table 4.3.

consider transfers resulting from the "price events" as described in Figure 4.5. The effect of these price events represents a large share of this transfer. Table 4.4 provides summary statistics for the price event transfers.

Assuming that 5% of the market quantity is actually exposed to the market clearing price, this difference is then equally dispersed across the total number of households in Texas.[37] Table 4.5 shows this calculation.

Panel A shows that for each price event, if the profit difference is accounted for entirely in the following month, each household's average electricity bill would be $0.24 higher, or about 0.2%. Over time, repeated occurrences of price events can lead to noticeable increases in a consumer's bill. Panel B shows that after firms continually behave in this manner over the summer period, this in total increases electricity bills by 3.2%.

electricity directly at the wholesale price.

[37]It is important to note that most of the market quantity is hedged through bilateral/financial contracts. Hence, the quantity actually exposed to the spot price can vary by the hour. However, most estimates show that this is on average about 5-10%.

Table 4.5: Impact on Consumers

	Dollars	Percent of Bill
Panel A: Per Event		
Next Month	0.24	0.2%
12-Month	0.02	0.0175%
Panel B: Total		
Next Month	1.92	3.2%
12-Month	0.16	0.26%

Changes in market profits resulting from price events are passed along directly to a consumer's electricity bill. Panel A shows this for each price event and Panel B shows the total across all price events. The differences are dispersed across all households directly either on their "Next Month's" average electricity bill, or spread across a typical "12-Month" contract.

4.6 Learning and Information Disclosure

Given the impacts that one equilibrium can have over another, it is important to understand the process of how firms transition between multiple equilibria. This can potentially allow for policy measures in the form of information disclosure that focus on taking an ex-ante approach to monitoring equilibrium selection, rather than the current ex-post approach. One of the primary levers that policy-makers currently have available - but have invested little research in - is through information disclosure. Information disclosure is the process by which historical information about market outcomes and firm actions is made public to participants. Hence, market participants do not have the most immediate information available and must wait a certain period of time for it to be disclosed.

In dynamic settings, firms must maximize profits conditional on their expectations about what their rivals will do. These expectations depend on a lot of factors including the ability to observe the past play of other market participants. In essence, information disclosure defines a window where a firm does not have the information available to them about how its rivals and other firms most recently behaved. Hence, in this setting where firms learn about each other, the availability of recent information plays an important role in how each firm forms their expectations about how another firm will behave in the future.

In order to capture these elements, I embed a fictitious play learning model (whereby a firm places a weight on its rivals historical actions in order to form an expectation about what they will do today) into a model of dynamic profit maximization. These weights are then estimated dynamically and indicates how both firms form expectations about the probability of future actions by their rival. This model setup culminates in a fixed point equilibrium on each firm's probability of jointly maximizing profits and the other firm's belief about this probability.

4.6.1 Model Outline

Belief-based learning starts with the premise that players keep track of their rival's history of play and form beliefs about what they will do in the future based on this historical play. While there may be a few adaptive learning models that fall within the category of belief-based learning, fictitious play is largely considered to be the leading candidate model.[38] I assume that in auction t, firm i must choose an offer S_{it}. The firm also forms a belief about its rivals' offer $S_{-i,t}$ by sampling from the empirical distribution of past offers.[39] Furthermore, I allow for geometrically declining sampling weights to capture the fact that firms may believe recent offers are more similar to more recent observations. In practice, this amounts to assigning a sampling weight $\theta^{t-\tau-1}$ to rival offers $S_{-i,\tau}$ in auctions $\tau \leq t-1$, such that $\theta \in [0,1]$. I normalize θ such that the weights sum to one. Hence, in the event that $\theta = 1$, a firm believes that its rivals actions will be sampled equally from all past observations. On the other hand, $\theta = 0$ corresponds to a firm believing that its rivals actions today will be identical to its actions in $t-1$.

In order to model the transition period, I look at the behaviors of only GDF Suez (the leader) and Luminant (the follower) and treat all other firms as part of a fringe supply.[40] In auction t, each

[38] See Aguirregabiria and Jeon (2018) for further discussion on this.

[39] There is some discussion about the bounds of this distribution. For computational reasons, I sample from the most recent week of offers following a bounded approach. See Sela and Herreiner (1999) for a more detailed remark on this.

[40] For sake of clarity, I do not assume that all other firms offer at marginal cost. I assume that their actual offer

firm takes one of two actions, $a_{it} \in \{a^S, a^J\}$, where a^S corresponds to the firm maximizing its static profits and a^J corresponds to the firm maximizing dynamic profits and producing the jointly withheld quantity.[41] Each player's probability of playing a^J, $Pr_{it}(a_{i,t} = a^J)$, is determined by its expected profits and by their belief about the probability of their rival playing a^J, $Pr_{it}(a_{-i,t} = a^J) + \epsilon_i$, too. Each firm knows the weight they assign to the historical play of its rivals, while the other only has a belief about what this private value is.

I further model the probability that the leader plays a^J as $Pr_{it}(a_{i,t} = a^J) = Pr(\pi_i(a^J) > \pi_i(a^S))$, where:

$$\pi_i(a^J) = \pi_t(a_i^J, a_{-i}^S) + \sum_{t+n}^{T} \delta^{t+n}[Pr_{it}(a_{-i,t} = a^J) * \pi_{t+n}(a_i^J, a_{-i}^J) + (1 - Pr_{it}(a_{-i,t} = a^J)) * \pi_{t+n}(a_i^J, a_{-i}^S)] + \epsilon_i. \quad (4.8)$$

To summarize, the leader must send an initial signal, $\pi_t(a_i^J, a_{-i}^S)$ in the form of foregone static profits to the follower. However, the leader has a belief about when the follower will respond to this signal and similarly play a^J. Once both firms play a^J, the leader receives the high-price profits, $\pi_{t+n}(a_i^J, a_{-i}^J)$. On the other hand, the leader could never send the initial signal and instead receive static profits, $\pi_i(a^S)$. Formally,

$$\pi_i(a^S) = \pi_t(a_i^S, a_{-i}^S) + \sum_{t+n}^{T} \delta^{t+n}[\pi_{t+n}(a_i^S, a_{-i}^S)] + \epsilon_i. \quad (4.9)$$

In order to model the follower's behavior, the follower must first observe a signal from the leader. That is, in period $t - n$, the follower observes $a_{-i,t-n} = a^J$. The follower then decides whether to play a^J or a^S. I model this probability again as $Pr_{it}(a_{i,t} = a^J) = Pr(\pi_i(a^J) > \pi_i(a^S))$,

would not change in response to the actions of GDF and Luminant.

[41] I find that prices in the third period are on average slightly higher than the Cournot outcome. Hence, for the sake of this section, I choose to model this second period as corresponding to an outcome from a repeated game.

but here the profit functions are now instead

$$\pi_i(a^J) = Pr_{it}(a_{-i,t} = a^J) * \pi_t(a_i^J, a_{-i}^J) + \sum_{t+n}^{T} \delta^{t+n}[Pr_{it}(a_{-i,t+n} = a^J)* \\ \pi_{t+n}(a_i^J, a_{-i}^J)] + \epsilon_i. \quad (4.10)$$

The probability of the follower playing a^J depends on its belief about the leader continuing to play a^J going forwards too. Alternatively, the follower could obtain static profits $\pi_i(a^S)$ denoted by

$$\pi_i(a^S) = Pr_{it}(a_{-i,t} = a^J) * \pi_t(a_i^S, a_{-i}^J) + \sum_{t+n}^{T} \delta^{t+n}[Pr_{it}(a_{-i,t} = a^J) * \pi_{t+n}(a_i^S, a_{-i}^J)+ \\ (1 - Pr_{it}(a_{-i,t} = a^J)) * \pi_{t+n}(a_i^S, a_{-i}^S)] + \epsilon_i. \quad (4.11)$$

The follower has the option of capitalizing on the leader playing a^J. By doing so however, the leader will revert back to a^S at some point.

4.6.2 Estimation Strategy

This model represents a fixed point equilibrium on probabilities. In equilibrium, the probability that the leader plays a^J corresponds to the follower's belief about the probability of the leader playing a^J too. Similarly, the probability that the follower plays a^J corresponds to the leader's belief about the probability of the follower playing a^J. Using this setup, I estimate the parameters of the model by maximizing the associated likelihood function:

$$L = \sum Pr^F(a_{Lt} = a^J; \theta_F)^{I(a_{Lt}=a^J)} \cdot Pr^L(a_{Ft} = a^J; \theta_L)^{I(a_{Ft}=a^J)}, \quad (4.12)$$

such that the above conditions hold true.

To summarize, this model estimates values of the weighting parameters, θ_F and θ_L, that best explain the observed probabilities of both firms playing a^J. In practice, these estimate amount to values of $\theta_F^* = 0.5$ and $\theta_L^* = 0.3$.[42] This indicates that both the leader and the follower place higher weight on their rival's more recent actions. However, it should be noted that these two estimates are not qualitatively different from each other, so the exact magnitudes should be interpreted with a grain of salt.

4.6.3 Counterfactual Information Disclosure Windows

Given that each firm observes and puts more weight on the other's more recent actions, the availability of information plays an important role in this equilibrium selection process. For example, without any disclosure window, if the leader sends out a signal in period t, then the follower can decide how to respond in period $t + 1$. However, in the event that there is a one day window, the leader must send out a signal in period t and $t + 1$ before the follower is able to respond in period $t + 2$. Intuitively, for longer windows, the leader will have to sacrifice more profits before being able to observe how the follower will respond.

To estimate the impact of different information disclosure windows, I allow firms to only weight observations explicitly outside of this window. For various window lengths w and parameter values θ_F^* and θ_L^*, I can determine how both firms incorporate and respond to their rival's historical actions $\{a_{-i,t-w}, ..., a_{-i,t-w-n}\}$. Figure 4.6 shows the impact that different window lengths w have on market prices. In summary, longer window lengths are associated with a lower average market clearing price. A two and five-day window are associated with average prices of \$52.00 and \$51.90 respectively, whereas a 10 and 20-day window are associated with average prices of \$51.65 and \$50.98.

Important to the content of this paper, we are interested in information disclosure as a potential policy lever used to prevent the formation of the high priced equilibrium. With this in mind, Figure 4.6 also shows that the frequency of high priced occurrences (>\$100) does not change when going from two to five days, but falls significantly at 10 days. Under these information lags, firms must sacrifice more static profits in the short-term to signal their intent to the other firm. As these foregone profits accumulate, it no longer becomes profitable for firms to even try to reach the

[42]For computational reasons, these results are based on only a subset of afternoon hours.

Figure 4.6: Impact of Information "Window" Lags

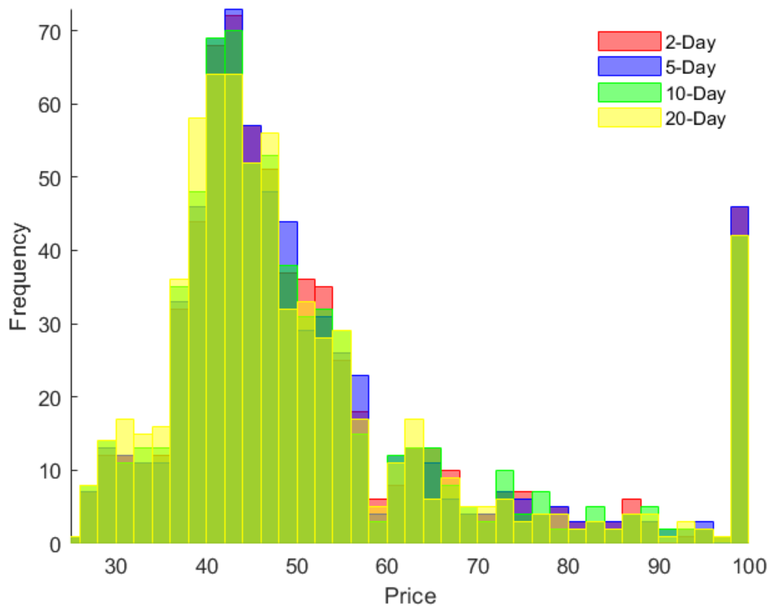

Notes: This figure plots the distribution of market prices, P_t, for various counterfactual information disclosure "window" lag lengths. Prices above $100 are pooled together in the last bin.

second higher priced equilibrium.

Information disclosure policies are important components in equilibria transition and firm learning. For the sake of completeness, it should be cautioned that there may be benefits of shorter window lengths for the market as a whole. Specifically, with more recent information, firms are potentially able to start-up low/medium cost generators if they have information that they will be called upon and used in the near future. Generators can ramp up their generation over time. Without this notice, lower cost generators may not have the opportunity to be on (for example, maintenance) leading to the market operator dispatching higher cost generators instead. Given the importance that information disclosure has on the market, future work should comprehensively look at the full spectrum of costs and benefits resulting from shorter or longer information disclosure windows. Furthermore, this analysis assumes that any and all information is not accessible to firms (i.e., an information "black out") within the defined window length.[43] Determining the

[43]It should also be noted that it may not be practical or feasible to "black out" all sources of information as firms

heterogeneous benefits and costs of each information source is also important to consider.

4.7 Conclusion

The multiplicity of equilibria makes designing and regulating markets a complex problem. In oligopoly markets where firms compete in supply functions, this is even more the case. Most research to date focuses on studying the specific market characteristics that facilitate one equilibrium over another, rather than the process by which firms can transition between many equilibria. However, this does not aide in monitoring markets from an ex-ante perspective. Differences in prices (and ultimately profits) can generate an incentive for firms to transition away from lower priced equilibria towards higher priced ones.

In this paper, I study the role of price leadership and learning in the establishment of a new equilibrium and the transition towards it in the Texas wholesale electricity market during the summer of 2013. I use firm offer data into 15-minute electricity auctions to document the systematic deviations from static unilateral best response behavior of a price leader and a follower. These deviations impose a cost on the firms in the form of foregone profits, but in the long run, transitioning to the high price equilibrium results in price increases of 5% relative to the average market price during the study time period, and the potential for price events as large as 1,500%.

Given these impacts, there are multiple important policy implications. First, while most regulatory emphasis is placed on large firms, smaller firms are also important to consider especially in regards to the entire market (Gates and Leuschner, 2007). This result generalizes to many concentrated industries like technology and search engines (e.g., Google, Yahoo, Microsoft), wireless telecommunications (e.g., Verizon Wireless, AT&T Inc., Sprint Nextel Corporation), airlines (e.g., Southwest, Delta, American, United), among many others. While the impacts of a single firm in highly concentrated markets can be large, the actions of smaller firms are also important to consider.

Second, the way in which information is released and made available to the market matters. To explicitly analyze the process of learning during the transition period, I integrate a fictitious

may have their own private sources that they use.

play learning model in a model of dynamic profit maximization. I show that with a 10-day lag in information disclosure, these firms would not have found it profitable to transition altogether. Hence, one way to prevent this transition may be through the use of information disclosure policies. In markets where information is disclosed to participants (financial and stock, etc.), this result shows that the timing of information release is an important factor to consider when designing a market. Given this, further research should focus on documenting the full range of costs and benefits that are associated with various information disclosure policies.

5. SUMMARY AND CONCLUSIONS

The three essays comprising this dissertation use experimental and structural research methods to study how institutions, emerging technologies, and historical policies interact and ultimately impact outcomes of interest to consumers and society.

In Section 2, together with Roberto Mosquera, Mofioluwasademi Odunowo, Xiongfei Guo, and Ragan Petrie, we study the causal effect of Facebook on various aspects of daily life as well as the monetary value that users have for the platform. Using an incentive-compatible auction mechanism, we find that one week of Facebook is estimated to be worth $67 for participants in our experiment. This is a relatively large value considering that it represents approximately 30 percent of the average weekly income for university students. Second, we find that Facebook is a primary source of news for participants as well. When restricted from its use, individuals did not search for news from other sources, even from sources with low substitution costs (e.g., radio, television, internet). This decrease in news access is also associated with increased difficulty in assessing the veracity of news. Being off of Facebook resulted in more uncertainty about whether news from politically-skewed sources was fake or not. Finally, we find that using Facebook induces negative feelings, such as depression, and that participants switch to healthier activities when they do have access to Facebook.

In Section 3, with Steve Puller, we study the relationship between the distributional consequences of a uniform gasoline tax and political support for the tax itself. In order to address the externalities imposed on society from gasoline consumption in the personal transportation market, the United States has elected to regulate vehicle manufacturers through Corporate Average Fuel Economy (CAFE) Standards instead of through Pigouvian taxes despite a large and growing literature showing that the former is anywhere from 3 to 6 times less cost efficient than the latter. This status quo stems from the fact that gasoline taxes are politically unpopular among American voters. While uniform gasoline taxes can be welfare enhancing in the aggregate, there exists heterogeneity amongst who receives the benefits and who bears the costs. In this paper, we first design an optimal uniform gasoline tax that properly accounts for local pollution damages (in practice, this amounts to \$0.40 per gallon), counterfactually increase the price of gasoline, and find that there is substantial heterogeneity across the state in terms of the distribution of costs and benefits. We then show that, even after controlling for an individual's political affiliation, this distribution is an important determinant in an individual's decision to support a gasoline tax and can explain up to 10% of an individual's decision. This result establishes a behavioral link between self-interests and tax support that goes beyond "tribal" politics allowing us to design tax regimes that are "politically-sophisticated." We design a revenue neutral tax regime whereby individuals in counties with higher proportions of people who have costs that outweigh their benefits receive a larger tax dividend than those with smaller proportions. This tax regime shifts the distribution of support for a gasoline tax from a median level of 4 (out of 10) to a 5.

In Section 4, I study the process by which firms transition between multiple equilbria in wholesale electricity markets. In oligopoly markets where firms choose supply functions rather than just price or quantity, there exists a wide range of potential equilibria resulting in an equally wide range of market outcomes in terms of prices and profits. Hence, these differences generates the incentive for firms to transition away from low price equilibria towards high price equilibria. Despite this, surprisingly little is known about how an equilibrium is reached as well as how agents transition between them. I show that a small firm in the market suddenly begins to deviate from a static low-price supply function equilibrium (SFE) using a model of unilateral profit maximization. This firm foregoes static profits (≈\$3,500 per offer) in favor of over-pricing a significant portion of its production. Shortly after this behavior, the largest firm in the market begins to behave similarly and over-prices its production (forgoing ≈\$1,200 per offer). I show that this shift in play is associated with an average price increase of 5%, but can also result in swings as large as 1,500%. To explicitly analyze the process of learning during the transition period, I integrate a fictitious play learning model into a model of dynamic profit maximization. This model allows me to estimate and paramaterize the beliefs needed to be held by both firms in order for them to behave in this manner. I further show that by preventing the most recent information from being observable, firms would need to forgo an increased amount of static profits to arrive at the high priced equilibrium. Revealing information with a 10-day lag prevents firms from transitioning to the high priced equilibrium altogether.

REFERENCES

Abadie, A. (2002). Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models. *Journal of the American Statistical Association*, 97(457):284–292.

Agrawal, A. and Nixon, H. (2014). What do americans think about federal tax options to support public transit, highways, and local streets and roads? results from year five of a national survey what do americans think about federal transportation tax options?

Aguirregabiria, V. and Jeon, J. (2018). Firms' Beliefs and Learning: Models, Identification, and Empirical Evidence. Working Papers tecipa-620, University of Toronto, Department of Economics.

Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2019). The welfare effects of social media. Technical report, National Bureau of Economic Research.

Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2):211–36.

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association*, 103(484):1481–1495.

Ashley, C. and Tuten, T. (2015). Creative Strategies in Social Media Marketing: An Exploratory Study of Branded Social Content and Consumer Engagement. *Psychology & Marketing*, 32(1):15–27.

Austin, D. and Dinan, T. (2005). Clearing the air: The costs and consequences of higher cafe standards and increased gasoline taxes. *Journal of Environmental Economics and Management*, 50(3):562 – 582.

Bailey, M., Cao, R., Kuchler, T., and Stroebel, J. (2017). The economic effects of social networks: Evidence from the housing market.

Bailey, M., Cao, R., Kuchler, T., Stroebel, J., and Wong, A. (2018). Social connectedness: Measurement, determinants, and effects. *Journal of Economic Perspectives*, 32(3):259–80.

Becker, G. M., Degroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3):226–232.

Becker, G. S. and Murphy, K. M. (1988). A Theory of Rational Addiction. *Journal of Political Economy*, 96(4):675–700.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 289–300.

Benjamini, Y., Krieger, A. M., and Yekutieli, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3):491–507.

Bento, A. M., Goulder, L. H., Jacobsen, M. R., and von Haefen, R. H. (2009). Distributional and efficiency impacts of increased us gasoline taxes. *American Economic Review*, 99(3):667–99.

Blake, H., Mo, P., Malik, S., and Thomas, S. (2009). How effective are physical activity interventions for alleviating depressive symptoms in older people? a systematic review. *Clinical rehabilitation*, 23(10):873–887.

Bohm, P., LindÃľn, J., and SonnegÃěrd, J. (1997). Eliciting reservation prices: Becker-degroot-marschak mechanisms vs. markets. *The Economic Journal*, 107(443):1079–1089.

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., and Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295.

Borenstein, S., Bushnell, J. B., and Wolak, F. A. (2002). Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market. *American Economic Review*, 92(5):1376–1405.

Brynjolfsson, E., Eggers, F., and Gannamaneni, A. (2018). Using massive online chohice experiments to measure changes in well-being. NBER Working Paper 24514.

Bushnell, J. B., Mansur, E. T., and Saravia, C. (2008). Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets. *The American Economic Review*, 98(1):237–266.

Busse, M. R. (2000). Multimarket Contact and Price Coordination in the Cellular Telephone Industry. *Journal of Economics & Management Strategy*, 9(3):287–320.

Byrne, D. P. and de Roos, N. (2019). Learning to Coordinate: A Study in Retail Gasoline. *American Economic Review*, 109(2):591–619.

Chilet, J. A. (2018). Collusive Price Leadership in Retail Pharmacies in Chile. Working paper, Editorial Express.

Corrigan, J. R., Alhabash, S., Rousu, M., and Cash, S. B. (2018). How much is social media worth? Estimating the value of Facebook by paying users to stop using it. *PloS One*, 13(12):e0207101.

De Vries, L., Gensler, S., and Leeflang, P. S. (2012). Popularity of Brand Posts on Brand Fan Pages: An Investigation of the Effects of Social Media Marketing. *Journal of Interactive Marketing*, 26(2):83–91.

DellaVigna, S. and Kaplan, E. (2007). The fox news effect: Media bias and voting*. *The Quarterly Journal of Economics*, 122(3):1187–1234.

Deloitte (2015). Facebook's global economic impact: A report for facebook.

Deters, F. g. and Mehl, M. R. (2013). Does posting facebook status updates increase or decrease loneliness? an online social networking experiment. *Social psychological and personality science*, 4(5):579–586.

Doraszelski, U., Lewis, G., and Pakes, A. (2018). Just Starting Out: Learning and Equilibrium in a New Market. *American Economic Review*, 108(3):565–615.

Downs, A. (1957). An economic theory of political action in a democracy. *Journal of Political Economy*, 65(2):135–150.

Edlin, A. S. and Karaca-Mandic, P. (2006). The accident externality from driving. *Journal of Political Economy*, 114(5):931–955.

Ellison, N. B., Steinfield, C., and Lampe, C. (2007). The benefits of facebook "friends:" social capital and college studentsâĂŹ use of online social network sites. *Journal of Computer-Mediated Communication*, 12(4):1143–1168.

EMarketeer (2015). College students still spend most social time with facebook.

https://www.emarketer.com/Article/College-Students-Still-Spend-Most-Social-Time-with-Facebook/1012955.

Enikolopov, R., Makarin, A., and Petrova, M. (2016). Social media and protest participation: Evidence from Russia. *Universitat Pompeu Fabra*.

Facebook (2016). Facebook reports first quarter 2016 results.

Facebook (2017). Facebook reports third quarter 2017 results.

Feinstein, B. A., Hershenberg, R., Bhatia, V., Latack, J. A., Meuwly, N., and Davila, J. (2013). Negative Social Comparison on Facebook and Depressive Symptoms: Rumination as a Mechanism. *Psychology of Popular Media Culture*, 2(3):161.

Finkelstein, A., Poterba, J., and Rothschild, C. (2009). Redistribution by insurance market regulation: Analyzing a ban on gender-based retirement annuities. *Journal of Financial Economics*, 91(1):38–58.

Flood, S., King, M., Ruggles, S., and Warren, J. R. (2017). Integrated public use microdata series, current population survey: Version 5.0. dataset, University of Minnesota, Minneapolis, MN. https://doi.org/10.18128/D030.V5.0.

Fredrickson, B. L. and Joiner, T. (2002). Positive emotions trigger upward spirals toward emotional well-being. *Psychological science*, 13(2):172–175.

Gates, S. M. and Leuschner, K. J. (2007). Is special regulatory treatment for small businesses working as intended? Technical report.

Gentzkow, M. and Shapiro, J. M. (2011). Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839.

Gonzales, A. L. and Hancock, J. T. (2011). Mirror, mirror on my facebook wall: Effects of exposure to facebook on self-esteem. *Cyberpsychology, Behavior, and Social Networking*, 14(1-2):79–83.

Gottfried, J. and Shearer, E. (2016). News use across social media platforms 2016.

Green, R. (1999). The Electricity Contract Market in England and Wales. *The Journal of Industrial Economics*, 47(1):107–124.

Gruber, J. and Köszegi, B. (2001). Is Addiction "Rational"? Theory and Evidence. *The Quarterly Journal of Economics*, 116(4):1261–1303.

Herzog, B. (2018). Valuation of digital platforms: experimental evidence for google and facebook. *International Journal of Financial Studies*, 6(4):87.

Horowitz, J. K. (2006). The Becker-DeGroot-Marschak mechanism is not necessarily incentive compatible, even for non-random goods. *Economics Letters*, 93(1):6–11.

Hortacsu, A., Luco, F., Puller, S. L., and Zhu, D. (2019). Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets. forthcoming, American Economic Review.

Hortacsu, A. and Puller, S. L. (2008). Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Spot Market. *The RAND Journal of Economics*, 39(1):86–114.

Hotelling, H. (1929). Stability in competition. *The Economic Journal*, 39(153):41–57.

J. Osborne, M. (1995). Spatial models of political competition under plurality rule: A survey of some explanations of the number of candidates and the positions they take. *Canadian Journal of Economics*, 28:261–301.

Jacobsen, M. R. (2013). Evaluating us fuel economy standards in a model with producer and household heterogeneity. *American Economic Journal: Economic Policy*, 5(2):148–87.

Jelenchick, L. A., Eickho, J. C., and Moreno, M. A. (2013). Facebook depression? social networking site use and depression in older adolescents. *Journal of Adolescent Health*, 52(1):128–130.

Kahneman, D. and Krueger, A. B. (2006). Developments in the Measurement of Subjective Well-Being. *The Journal of Economic Perspectives*, 20(1):3–24.

Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., and Stone, A. A. (2004). A Survey Method for Characterizing Daily Life Experience: The Day Reconstruction Method. *Science*, 306(5702):1776–1780.

Kettunen, O. et al. (2015). Effects of physical activity and fitness on the psychological wellbeing of young men and working adults: associations with stress, mental resources, overweight and workability.

Kim, J. and Lee, J. (2011). The facebook paths to happiness: Effects of the number of facebook friends and self-presentation on subjective well-being. *CyberPsychology, Behavior, and Social Networking*, 14(6):359–364.

Klar, S. (2013). The influence of competing identity primes on political preferences. *The Journal of Politics*, 75(4):1108–1124.

Klemperer, P. D. and Meyer, M. A. (1989). Supply Function Equilibria in Oligopoly under Uncertainty. *Econometrica*, 57(6):1243–1277.

Knetsch, J. L., Tang, F.-F., and Thaler, R. H. (2001). The endowment effect and repeated market trials: Is the vickrey auction demand revealing? *Experimental Economics*, 4(3):257–269.

Knittel, C. and Sandler, R. (fcm). The welfare impact of indirect pigouvian taxation: Evidence from transportation. *American Economic Journal: Economic Policy*.

Knittel, C. R. (2014). The political economy of gasoline taxes: Lessons from the oil embargo. *Tax Policy and the Economy*, 28(1):97–131.

Krueger, A. B. and Schkade, D. A. (2008). The Reliability of Subjective Well-Being Measures. *Journal of Public Economics*, 92(8):1833–1845.

Laroche, M., Habibi, M. R., Richard, M.-O., and Sankaranarayanan, R. (2012). The Effects of Social Media Based Brand Communities on Brand Community Markers, Value Creation Practices, Brand Trust and Brand Loyalty. *Computers in Human Behavior*, 28(5):1755–1767.

Lavelle, M. (2019). Carbon tax plans: How they compare and why oil giants support one of them.

Lee, C. S. and Ma, L. (2012). News Sharing in Social Media: The Effect of Gratifications and Prior Experience. *Computers in Human Behavior*, 28(2):331–339.

Lemus, J. and Luco, F. (2018). Price Leadership and Uncertainty about Future Costs. *SSRN Electronic Journal*.

Levin, L., Lewis, M. S., and Wolak, F. A. (2017). High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, 9(3):314–47.

Lewis, M. (2012). Price Leadership and Coordination in Retail Gasoline Markets with Price Cycles. *International Journal of Industrial Organization*, 30(4):342–351.

Lin, C.-Y. C. and Prince, L. (2009). The optimal gas tax for California. *Energy Policy*, 37(12):5173–5183.

Mankiw, N. G. (2009). Smart taxes: An open invitation to join the pigou club. *Eastern Economic Journal*, 35(1):14–23.

Mansur, E. T. (2008). Measuring Welfare in Restructured Electricity Markets. *The Review of Economics and Statistics*, 90(2):369–386.

Marion, J. and Muehlegger, E. (2011). Fuel tax incidence and supply conditions. *Journal of Public Economics*, 95(9):1202 – 1212. Special Issue: The Role of Firms in Tax Systems.

Markham, J. W. (1951). The Nature and Significance of Price Leadership. *The American Economic Review*, 41(5):891–905.

Martin, G. J. and Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–99.

Mayer, A. and Puller, S. L. (2008). The Old Boy (and Girl) Network: Social Network Formation on University Campuses. *Journal of Public Economics*, 92(1):329–347.

Mazar, N., Koszegi, B., and Ariely, D. (2014). True context-dependent preferences? the causes of market-dependent valuations. *Journal of Behavioral Decision Making*, 27(3):200–208.

McLeish, K. N. and Oxoby, R. J. (2011). Social interactions and the salience of social identity. *Journal of Economic Psychology*, 32(1):172 – 178.

Mercadal, I. (2018). Dynamic Competition and Arbitrage in Electricity Markets: The Role of Financial Traders. *SSRN Electronic Journal*.

Miller, N., Sheu, G., and Weinberg, M. (2018). Oligopolistic Price Leadership: An Empirical Model of the U.S. Beer Industry. *SSRN Electronic Journal*.

Miravete, E. J., Seim, K., and Thurk, J. (2017). One markup to rule them all: Taxation by liquor pricing regulation. Working Paper 24124, National Bureau of Economic Research.

Mullainathan, S. and Shleifer, A. (2005). The market for news. *American Economic Review*, 95(4):1031–1053.

Muller, N. Z. and Mendelsohn, R. (2009). Efficient pollution regulation: Getting the prices right.

American Economic Review, 99(5):1714–39.

Muscanell, N. L. and Guadagno, R. E. (2012). Make new friends or keep the old: Gender and personality differences in social networking use. *Computers in Human Behavior*, 28(1):107–112.

Neilsen Company, T. (2016). 2016 neilsen social media report. The Neilsen Company.

Newman, D. B., Tay, L., and Diener, E. (2014). Leisure and subjective well-being: A model of psychological mechanisms as mediating factors. *Journal of Happiness Studies*, 15(3):555–578.

OECD Better Life Initiative (2013). OECD Guidelines on Measuring Subjective Well-Being.

O'Neil, A., Quirk, S. E., Housden, S., Brennan, S. L., Williams, L. J., Pasco, J. A., Berk, M., and Jacka, F. N. (2014). Relationship between diet and mental health in children and adolescents: A systematic review. *American Journal of Public Health*, 104(10):e31–e42. PMID: 25208008.

Orbach, B. and Einav, L. (2007). Uniform prices for differentiated goods: The case of the movie-theater industry. 27:129–153.

Ostir, G. V., Markides, K. S., Black, S. A., and Goodwin, J. S. (2000). Emotional well-being predicts subsequent functional independence and survival. *Journal of the American Geriatrics Society*, 48(5):473–478.

Parry, I. W. H. and Small, K. A. (2005). Does britain or the united states have the right gasoline tax? *American Economic Review*, 95(4):1276–1289.

Pew Research Center (2016). Social media update 2016.

Pew Research Center (2018). Social media use 2018.

Plott, C. R. and Zeiler, K. (2005). The willingness to pay-willingness to accept gap, the "endowment effect," subject misconceptions, and experimental procedures for eliciting valuations. *The American Economic Review*, 95(3):530–545.

Poterba, J. (2011). Is the gasoline tax regressive? *Tax Policy and the Economy*, 5.

Rezaee, M., Hedayati, A., Naghizadeh, M. M., Farjam, M., Sabet, H. R., and Paknahad, M. (2016). Correlation between happiness and depression according to beck depression and oxford happiness inventory among university students. *Galen Medical Journal*, 5(2):75–81.

Sagioglou, C. and Greitemeyer, T. (2014). FacebookâĂŹs emotional consequences: Why facebook causes a decrease in mood and why people still use it. *Computers in Human Behavior*, 35:359–363.

Salovey, P., Rothman, A. J., Detweiler, J. B., and Steward, W. T. (2000). Emotional states and physical health. *American Psychologist*, 55(1):110.

Sapienza, P. and Zingales, L. (2013). Economic experts vs. average americans. *SSRN Electronic Journal*.

Sela, A. and Herreiner, D. (1999). Fictitious Play in Coordination Games. *International Journal of Game Theory*, 28:189–197.

Seow, D. Y., Haaland, B., and Jafar, T. H. (2015). The association of prehypertension with meals eaten away from home in young adults in singapore. *American journal of hypertension*, 28(10):1197–1200.

Shakya, H. B. and Christakis, N. A. (2017). Association of facebook use with compromised well-being: a longitudinal study. *American journal of epidemiology*, 185(3):203–211.

Shirley Cramer, B. I. (2017). State of mind social media and young people's mental health and wellbeing. pages 13–16.

Sonnentag, S. (2001). Work, recovery activities, and individual well-being: a diary study. *Journal of occupational health psychology*, 6(3):196.

Statista (2018). Frequency of facebook use in the united states as of february 2018 by gender. Retrieved June 23, 2018 from https://www.statista.com/statistics/653014/frequency-usage-facebook-usa-gender/.

Steers, M.-L. N., Wickham, R. E., and Acitelli, L. K. (2014). Seeing Everyone Else's Highlight Reels: How Facebook Usage is Linked to Depressive Symptoms. *Journal of Social and Clinical Psychology*, 33(8):701–731.

Sunstein, C. R. (2018). Valuing Facebook. *Behavioural Public Policy, Forthcoming*.

Tandoc, E. C., Ferrucci, P., and Duy, M. (2015). Facebook use, envy, and depression among college students: Is facebooking depressing? *Computers in Human Behavior*, 43:139–146.

Texas A&M Transportation Institute (2016). Texas Transportation Poll.

Tromholt, M. (2016). The facebook Experiment: Quitting Facebook Leads to Higher Levels of Well-Being. *Cyberpsychology, Behavior, and Social Networking*, 19(11):661–666.

Tuttle, B. (2012). How Much You Spend Each Year on Coffee, Gas, Christmas, Pets, Beer, and More. TIME Magazine. http://business.time.com/2012/01/23/how-much-you-spend-each-year-on-coffee-gas-christmas-pets-beer-and-more/.

Urban, L. E., Weber, J. L., Heyman, M. B., Schichtl, R. L., Verstraete, S., Lowery, N. S., Das, S. K., Schleicher, M. M., Rogers, G., Economos, C., et al. (2016). Energy contents of frequently ordered restaurant meals and comparison with human energy requirements and us department of agriculture database information: a multisite randomized study. *Journal of the Academy of Nutrition and Dietetics*, 116(4):590–598.

Valenzuela, S., Park, N., and Kee, K. F. (2009). Is there social capital in a social network site?: Facebook use and college students' life satisfaction, trust, and participation. *Journal of Computer-Mediated Communication*, 14(4):875–901.

Varshavsky, J. R., Morello-Frosch, R., Woodruff, T. J., and Zota, A. R. (2018). Dietary sources of cumulative phthalates exposure among the us general population in nhanes 2005–2014. *Environment international*, 115:417–429.

Verduyn, P., Lee, D. S., Park, J., Shablack, H., Orvell, A., Bayer, J., Ybarra, O., Jonides, J., and Kross, E. (2015). Passive facebook usage undermines affective well-being: Experimental and longitudinal evidence. *Journal of Experimental Psychology: General*, 144(2):480.

Vogel, E. A., Rose, J. P., Okdie, B. M., Eckles, K., and Franz, B. (2015). Who compares and despairs? the effect of social comparison orientation on social media use and its outcomes. *Personality and Individual Differences*, 86:249–256.

Wang, T. and Chen, C. (2014). Impact of fuel price on vehicle miles traveled (vmt): do the poor respond in the same way as the rich? *Transportation*, 41(1):91–105.

Wilson, R. (1979). Auctions of Shares. *The Quarterly Journal of Economics*, 93(4):675–689.

Wolak, F. A. (2000). An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market. *International Economic Journal*, 14(2):1–39.

Wolak, F. A. (2003). *Identification and Estimation of Cost Functions Using Observed Bid Data: An Application to Electricity Markets*, volume 2 of *Econometric Society Monographs*, page 133âĂŞ169. Cambridge University Press.

Wolfson, J. A. and Bleich, S. N. (2015). Is cooking at home associated with better diet quality or weight-loss intention? *Public Health Nutrition*, 18(8):1397–1406.

APPENDIX A

THE ECONOMIC EFFECTS OF FACEBOOK

A.1 Phase 1 Recruitment Email:

Howdy "Student's Name",

Did you know that the average person spends about 50 minutes a day on Facebook? Over an individual's lifetime, this will amount to 5 or more years.

To date, the impact of this usage is unclear. Texas A&M's Department of Economics is seeking current TAMU students who are Facebook users to participate in a research study. You are receiving this email because you are on A&M's email list. Our team is examining the effects of Facebook on everyday life, and we are looking for students to help us out.

If you have an active Facebook account, you may be eligible to participate in this paid research study. In an unusual turn of events, we are asking **you** to tell us how much money you would need to be paid to stay off Facebook for a week. Please note that if selected for this study, staying off Facebook for one week will be a part of the protocol.

Participation in this study involves:

- Cash payouts based on an auction
- Coming to the Evans Library on main campus to complete two surveys
- The potential to be without Facebook for a week

If you are interested in participating in this study, please click the link below for more information.
Take the survey

If you have any questions or would like more information about this study, please contact the research team by email at rpetrie@tamu.edu.

Thank you,

Prof. Ragan Petrie
TAMU Department of Economics
3035 Allen Building, College Station, TX 77845

Study Title: The Behavioral Effects of Social Media
IRB2017-0189M

Follow the link to opt out of future emails:
Click here to unsubscribe

Notes: Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

A.2 Phase 1 Survey:

TEXAS A&M UNIVERSITY

We want to analyze the effects of social media on individual behavior including the value individuals place on social media, its' effects on social awareness, consumption behavior, expectations about the future, and general well-being.

This survey, which will take no more than 5 minutes, has two objectives. First, you will help us learn about the value of social media. Second, this survey will help us determine participation for an experiment in which we will randomly require that half of the group not use their Facebook accounts for one week.

In this survey we will collect basic demographic information and pre-select participants for this experiment.

We will collect email addresses of those pre-selected to participate. This information will be used to contact chosen participants and to match the data collected in this short survey with two additional surveys that will be conducted during the experimental sessions. After the match, email addresses will be destroyed.

If you have any questions about this study, please contact Dr. Ragan Petrie and the research team at rpetrie@tamu.edu.

IRB2017-0189M

Would you like to be considered as a participant in this experiment?

- Yes
- No

>>

If no, the next screen shows

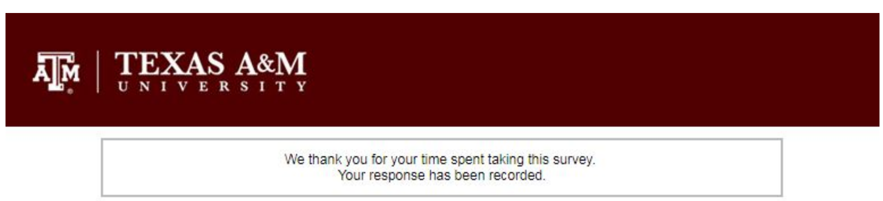

TEXAS A&M UNIVERSITY

We thank you for your time spent taking this survey.
Your response has been recorded.

If yes, the next screen shows

TEXAS A&M UNIVERSITY

Thank you for your interest in helping us learn more about the effects of social media. Before continuing, please tell us a little bit about yourself.

What is your gender?

- Male
- Female

What is your age in years?
Please round up.

Which college are you currently enrolled in?

What is your major?

Did you live in the United States when you were 15 years old?

- Yes
- No

>>

TEXAS A&M UNIVERSITY

What was the zip code of your address when you were 15 years old?

>>

TEXAS A&M UNIVERSITY

What country did you live in when you were 15 years old?

>>

TEXAS A&M UNIVERSITY

Do you have an active Facebook account?

- Yes
- No

>>

TEXAS A&M UNIVERSITY

Let's play a game!

Please think carefully about your value of the time you spend on Facebook over a week. You will be asked to enter this value later.

Afterwards, we will present a counter-offer! This counter-offer will be randomly drawn from an interval of $5 to a maximum that is our most reasonable estimate of Facebook's value over a week.

If our counter-offer is greater than or equal to your valuation YOU WILL BE CONSIDERED TO PARTICIPATE IN THE EXPERIMENT. We will randomly select the final participants from this group. FINAL PARTICIPANTS will be paid the value of our random counter-offer.

If our counter-offer is lower than your valuation we will not be able to compensate you fairly. YOU WILL NOT BE CONSIDERED TO PARTICIPATE IN THE EXPERIMENT.

The next screen provides examples.

>>

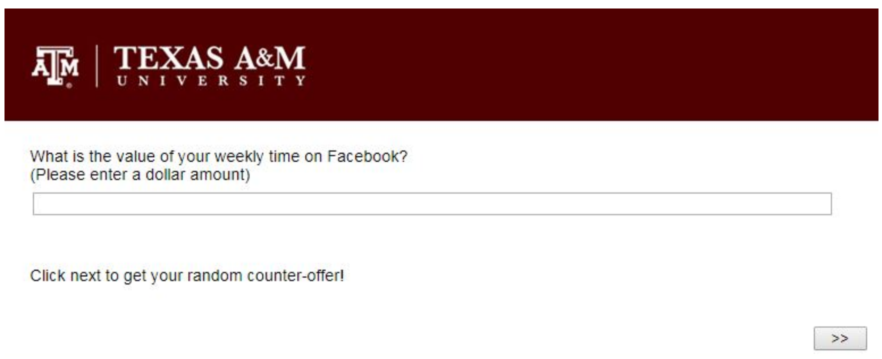

† The screen above represents the WTP setting. Half of the subjects received this wording while the other half were asked "How much money would you need to be given to stop using Facebook for a week?", which reflects the WTA setting.

TEXAS A&M UNIVERSITY

Our counter offer is $10.

Congratulations! You have been pre-selected to participate in this experiment.

On Thursday, April 20, we will send you an email if you are randomly selected as a final participant.

If selected, we will further explain the details of the experiment on Monday, April 24, at Evans Library. We will email you the exact time and room number. There will be a second session on Monday, May 1, at Evans Library. At most you will spend 1 hour between the two sessions. You will be paid $10 - our counter-offer - at the end of the experiment.

Remember, we will randomly require that half of the final participants do not use their Facebook accounts by any means for one week.

Please enter your preferred email address below:

Thank you for your time!

>>

For the case where the counter offer is less than the valuation:

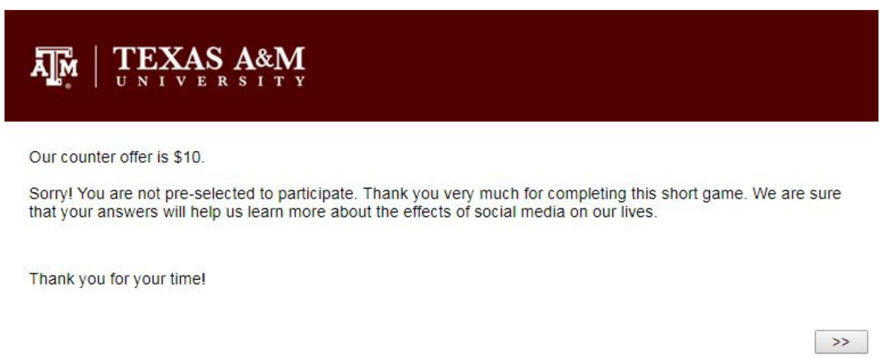

TEXAS A&M UNIVERSITY

Our counter offer is $10.

Sorry! You are not pre-selected to participate. Thank you very much for completing this short game. We are sure that your answers will help us learn more about the effects of social media on our lives.

Thank you for your time!

>>

Notes: Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

A.3 News Quiz

News Quiz in phase 2 (before treatment)

A1	Read the following list of events. Did these events happen in the previous week?			
		Definitely happened	I do not know	Definitely did not happened
A11	Serena Williams, the best women's tennis player, is expecting her first child and will not play again until next year.	❑	❑	❑
A12	Thousands of people gathered in the rain Saturday on the soggy grounds of the Washington Monument to turn Earth Day into an homage to science.	❑	❑	❑
A13	Facebook killer, Steve Stephens, was arrested in Ohio.	❑	❑	❑
A14	Vice President Mike Pence visited the demilitarized zone as the U.S. kept its options open on North Korea.	❑	❑	❑
A15	Stanford University, said that it would permit the conservative author Ann Coulter to speak on campus in early May, just one day after it canceled her appearance.	❑	❑	❑
A16	MSNBC analyst calls for ISIS to bomb Trump property.	❑	❑	❑
A17	General Motors has become the latest multinational company to pull out of Venezuela after it says government authorities illegally seized its plant there.	❑	❑	❑

News Quiz in phase 3 (after treatment)

A1	Read the following list of events. Did these events happen in the previous week?			
		Definitely happened	I do not know	Definitely did not happened
A11	Bulls bow out of playoffs with blowout loss to Celtics in Game 6.	❑	❑	❑
A12	Federal agencies take actions to implement President Trump's order to strip fund from municipal governments that refuse to cooperate fully with immigration agents.	❑	❑	❑
A13	Obama begins new phase of public life with Chicago visit.	❑	❑	❑
A14	Tens of thousands of people protested the president's rollback of rules protecting the environment.	❑	❑	❑
A15	President Trump has instructed his advisers to keep the corporate tax rate close to 30 percent.	❑	❑	❑
A16	In France's most consequential election in recent history, voters on Sunday chose Emmanuel Macron and Marine Le Pen to go to a runoff to determine the next president.	❑	❑	❑
A17	Trump wants to send astronauts to Mars during his presidency.	❑	❑	❑

Notes: Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

A.4 Survey Questionnaire

Date: May 1, 2017 Time: ☐☐ : ☐☐ PM

UIN	Please enter your TAMU UIN:	☐☐☐☐☐☐☐☐☐	Please enter your TAMU Email:

N1	How much time did you spend reading or watching the news per day last week?
	❑ Less than 15 min ❑ More than 15 minutes but less than 30 minutes ❑ More than 30 minutes but less than 1 hours ❑ More than 1 hour but less than 2 hours ❑ More than 2 hours

N2 Please indicate how frequently you used the following types of news media last week. Please answer on a scale of 1 to 7, where 7 is a type of media that you used frequently, and 1 is a type of media you used infrequently.

		◄ Not at all						All of the time ►
		1	2	3	4	5	6	7
N21	Cable TV	❑	❑	❑	❑	❑	❑	❑
N22	Paper news	❑	❑	❑	❑	❑	❑	❑
N23	Radio	❑	❑	❑	❑	❑	❑	❑
N24	Online news	❑	❑	❑	❑	❑	❑	❑
N25	Social media	❑	❑	❑	❑	❑	❑	❑
N26	News feed	❑	❑	❑	❑	❑	❑	❑
N27	Other 1: ______________	❑	❑	❑	❑	❑	❑	❑
N28	Other 2: ______________	❑	❑	❑	❑	❑	❑	❑

N3 Please indicate how frequently you used the following methods to obtain news last week. Please answer on a scale of 1 to 7, where 7 is a method you used frequently, and 1 is a method you used infrequently.

		◄ Not at all						All of the time ►
		1	2	3	4	5	6	7
N31	Watch	❑	❑	❑	❑	❑	❑	❑
N32	Read	❑	❑	❑	❑	❑	❑	❑
N33	Listen	❑	❑	❑	❑	❑	❑	❑

N4	List the top 3 news outlets/sources you got your news from last week.		
	1st Choice: ______________________	2nd Choice: ______________________	3rd Choice: ______________________

N5 What type of news did you frequently read, watch or listen to last week? Please answer on a scale of 1 to 7, where 7 is a type you used frequently, and 1 is a type you used infrequently.

		◄ Not at all						All of the time ►
		1	2	3	4	5	6	7
N51	Political	❑	❑	❑	❑	❑	❑	❑
N52	Sports	❑	❑	❑	❑	❑	❑	❑
N53	Business	❑	❑	❑	❑	❑	❑	❑
N54	International	❑	❑	❑	❑	❑	❑	❑
N55	Local news	❑	❑	❑	❑	❑	❑	❑
N56	Culture	❑	❑	❑	❑	❑	❑	❑
N57	Science	❑	❑	❑	❑	❑	❑	❑
N58	Weather	❑	❑	❑	❑	❑	❑	❑

N6 For the following sources, indicate how frequently you used each last week. Please answer on a scale of 1 to 7 where 7 is a source you used frequently, and 1 is a source you used infrequently.

		◄ Not at all						All of the time ►
		1	2	3	4	5	6	7
N61	Battalion	❑	❑	❑	❑	❑	❑	❑
N62	KBTX	❑	❑	❑	❑	❑	❑	❑
N63	MSC website	❑	❑	❑	❑	❑	❑	❑
N64	Local radio	❑	❑	❑	❑	❑	❑	❑
N65	Local newspaper	❑	❑	❑	❑	❑	❑	❑
N66	National newspaper	❑	❑	❑	❑	❑	❑	❑
N67	Online news	❑	❑	❑	❑	❑	❑	❑
N68	Online social network	❑	❑	❑	❑	❑	❑	❑
N69	Friends	❑	❑	❑	❑	❑	❑	❑

M1	Last week, how much time did you spend **each day** doing the following activities?				
M11	Sitting in a library on campus	_______ Hours	M17	Attending class	_______ Hours
M12	Studying	_______ Hours	M18	Sleeping (average number of hours per night)	_______ Hours
M13	Working for pay	_______ Hours	M19	Attending a party or social event (fill in for time you spent in total last week)	_______ Hours
M14	Exercising	_______ Hours	M110	At what time do you typically go to bed?	☐☐ : ☐☐ ❑ AM / ❑ PM
M15	Hanging out with friends	_______ Hours	M111	At what time do you typically wake up?	☐☐ : ☐☐ ❑ AM / ❑ PM
M16	Reading news	_______ Hours			
M2	Last week, how much time did you spend **each day** on the following types of social media?				
M21	Facebook	_______ Hours	M27	Vimeo	_______ Hours
M22	Instagram	_______ Hours	M28	YouTube	_______ Hours
M23	Twitter	_______ Hours	M29	Other 1: _______________	_______ Hours
M24	Tumblr	_______ Hours	M210	Other2 : _______________	_______ Hours
M25	Snapchat	_______ Hours			
M3	How many friends do you have on Facebook? (Feel free to open your FB account to check)				
M4	How many followers do you have on Instagram?				
M5	How many followers do you have on Tumblr?				
M6	How many followers do you have on Twitter?				

F1	How often do you do the following on Facebook?							
		Never	Rarely	1-2 times per month	Once a week	2-4 times per week	Once a day	Several times per day
F11	Open up FB to check your news feed	❑	❑	❑	❑	❑	❑	❑
F12	Read news feed content	❑	❑	❑	❑	❑	❑	❑
F13	Post pictures	❑	❑	❑	❑	❑	❑	❑
F14	Post comments	❑	❑	❑	❑	❑	❑	❑

F4	When you are on Facebook, how often do you feel the following?						
		Never	Rarely	Sometimes	Frequently	Often	All the time
F41	Envy/jealousy	❑	❑	❑	❑	❑	❑
F42	Happiness	❑	❑	❑	❑	❑	❑
F43	Misery	❑	❑	❑	❑	❑	❑
F44	Satisfaction	❑	❑	❑	❑	❑	❑
F45	Connected with friends	❑	❑	❑	❑	❑	❑
F46	Up to date on my friends' activities	❑	❑	❑	❑	❑	❑
F47	Lonely	❑	❑	❑	❑	❑	❑
F48	Annoyed	❑	❑	❑	❑	❑	❑
F49	Inspired	❑	❑	❑	❑	❑	❑

Please think about what you did __last week__ as you answer the following questions.

C1	(1-Strongly agree, 2-Agree, 3-Neither agree nor disagree, 4-Disagree, 5-Strongly disagree)					
		1	2	3	4	5
C11	I ate out less than I normally do	❑	❑	❑	❑	❑
C12	I did less impulse buying than usual	❑	❑	❑	❑	❑
C13	I saved more money than I usually do	❑	❑	❑	❑	❑
C14	I ate healthier than usual	❑	❑	❑	❑	❑
C15	I exercised more than usual	❑	❑	❑	❑	❑
C2	(1-Strongly agree, 2-Agree, 3-Neither agree nor disagree, 4-Disagree, 5-Strongly disagree)					
		1	2	3	4	5
C21	I wasted less time than I normally do	❑	❑	❑	❑	❑
C22	I achieved more than I normally do	❑	❑	❑	❑	❑
C23	I spent more time studying and doing school related work	❑	❑	❑	❑	❑
C24	I was not late for classes, meetings or work	❑	❑	❑	❑	❑
C25	I was able to meet deadlines without rushing at the last minute.	❑	❑	❑	❑	❑
C26	I was able to prevent distractions from achieving high priority tasks.	❑	❑	❑	❑	❑
C27	I discontinued any wasteful or unprofitable activities or routines.	❑	❑	❑	❑	❑
C28	I had time to relax and be with friends	❑	❑	❑	❑	❑
C29	I procrastinated less than I normally do	❑	❑	❑	❑	❑
C210	I partied a lot	❑	❑	❑	❑	❑

Please think about what you are going to do __this coming week__ as you answer the following questions.

C3	(1-Strongly agree, 2-Agree, 3-Neither agree nor disagree, 4-Disagree, 5-Strongly disagree)					
		1	2	3	4	5
C33	I expect to spend less on eating out and hanging out with friends	❑	❑	❑	❑	❑
C34	I expect to save more money	❑	❑	❑	❑	❑
C35	I will cut down on my impulse buying	❑	❑	❑	❑	❑
C36	I will spend more time studying	❑	❑	❑	❑	❑
C37	I will eat more healthy food	❑	❑	❑	❑	❑
C38	I will exercise more than I normally do	❑	❑	❑	❑	❑

S1	**Overall, how satisfied are you with life as a whole?**										
	◄ Not at all satisfied										Completely satisfied ►
	0	1	2	3	4	5	6	7	8	9	10
	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑
S2	**Overall, to what extent do you feel the things you do in your life are worthwhile?**										
	◄ Not at all worthwhile										Very worthwhile ►
	0	1	2	3	4	5	6	7	8	9	10
	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑
S3	**How happy are you?**										
	◄Very unhappy										Very happy ►
	0	1	2	3	4	5	6	7	8	9	10
	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑
S4	**How often do you worry?**										
	◄ Never										All of the time ►
	0	1	2	3	4	5	6	7	8	9	10
	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑
S5	**How often do you feel depressed?**										
	◄ Never										All of the time ►
	0	1	2	3	4	5	6	7	8	9	10
	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑	❑

D1	**What is your race?**		
	❑ White ❑ Black/ African American ❑ American Indian/ Alaskan Native	❑ Asian ❑ Native Hawaiian/ Other Pacific Islander ❑ Other: ___________________________	
D2	**What is your ethnicity?**	❑ Hispanic/ Latino	❑ Not Hispanic/ Latino

O1	**Is there anything else you would like to tell the research team?**

Notes: Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

A.5 Additional Results

Figure A.1: Facebook Negative Emotions

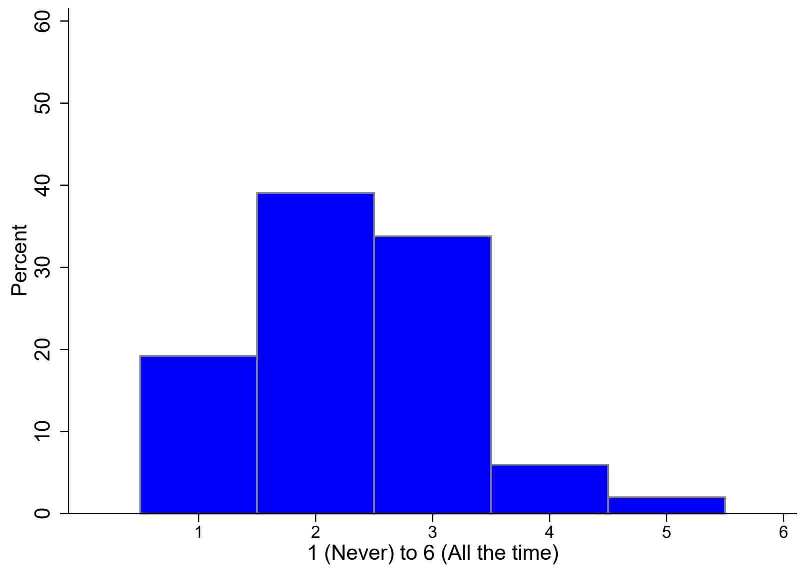

(a) Envy/Jealousy

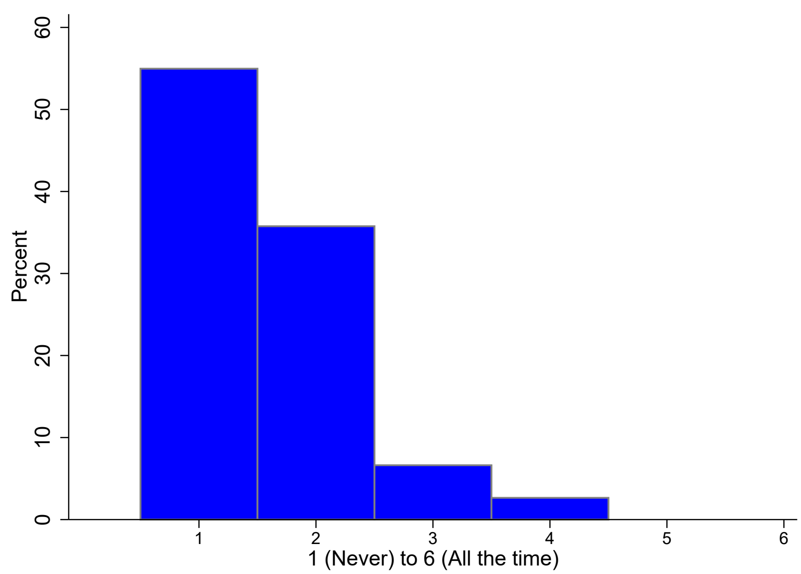

(b) Misery

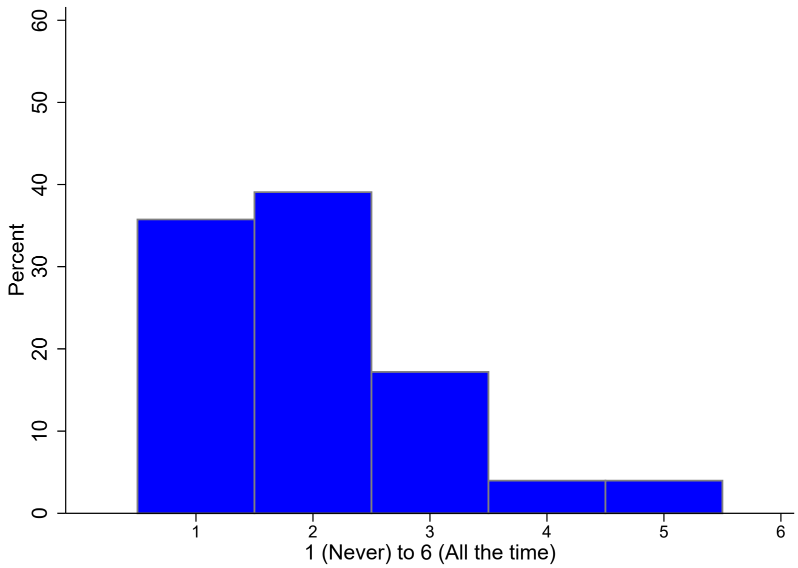

(c) Lonely

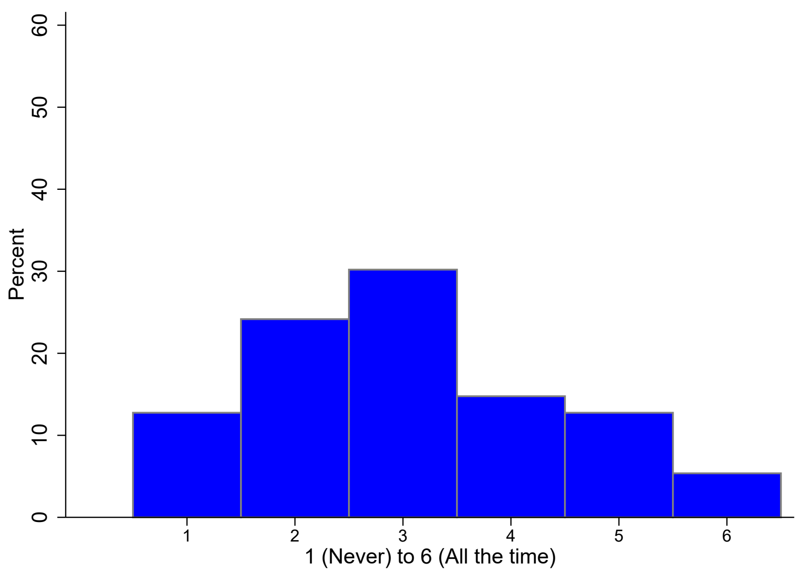

(d) Annoyed

Notes: Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Figure A.2: Change in Reported Depression and Change in the Value of Facebook

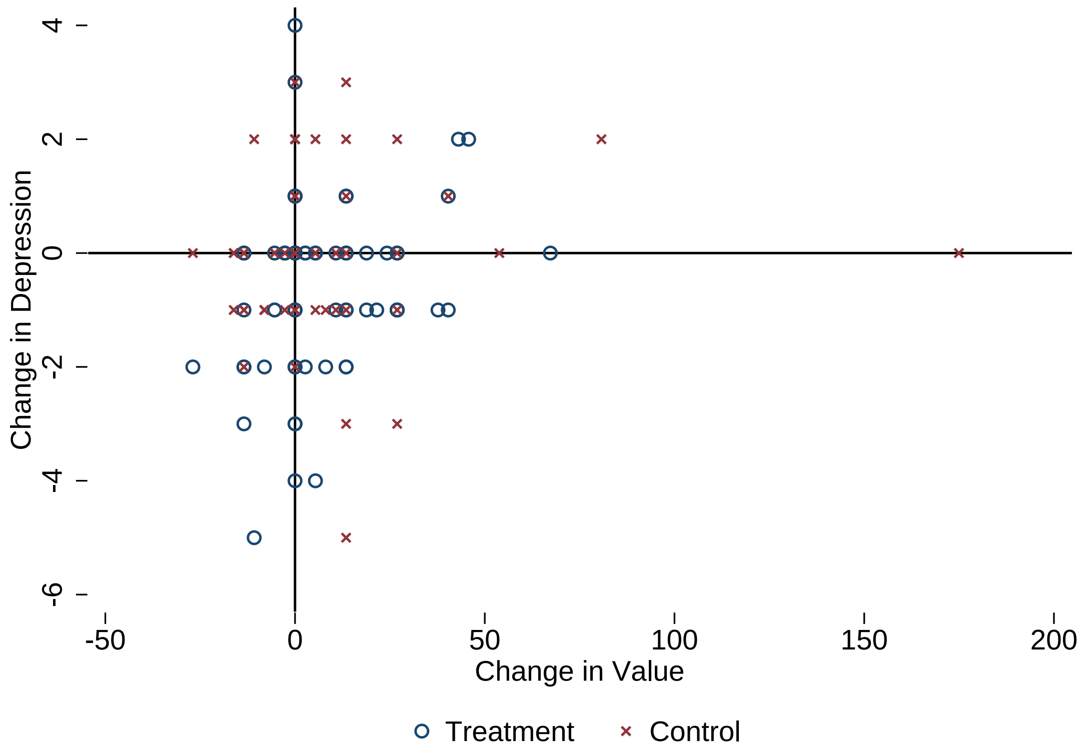

Notes: Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Table A.1: Descriptive Statistics by Survey Phases

	Ineligible	Eligible	P-value	Eligible (Show)	Eligible (No-Show)	P-value
Value of Facebook	85.35	27.11	0.000	28.97	26.33	0.025
	(119.88)	(12.72)		(12.98)	(12.55)	
Woman	0.60	0.59	0.720	0.65	0.57	0.089
	(0.49)	(0.49)		(0.48)	(0.50)	
Age	20.77	20.55	0.009	20.59	20.53	0.693
	(1.65)	(1.68)		(1.99)	(1.53)	
Income ($)	67,204	71,761	0.109	69,509	72,286	0.512
	(55,192)	(68,778)		(63,207)	(71,032)	
N	1,207	562		167	395	

This table presents the means for eligible and ineligible participants from the Phase 1 survey and for the eligible participants that showed up to complete the Phase 2 survey and those that were eligible but did not show up for phase 2. The p-values represents the difference of means for each group. Standard deviations are in parentheses. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Table A.2: Facebook Restriction - Balance of Covariates

	Treatment	Control	P-value
Value of Facebook	28.42	29.43	0.618
	(11.27)	(14.33)	
Woman	0.57	0.711	0.060
	(0.50)	(0.46)	
Age	20.69	20.51	0.569
	(2.41)	(1.56)	
Income($)	67,900	75,986	0.482
	(55,988)	(68,904)	
N	77	90	

The first two columns present the means of different observables characteristics for the Facebook restriction treatment group and the no restriction control group. Columns 3 presents the p-values of the difference of means between these groups. Standard deviations are in parentheses. Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Table A.3: Phase 2 Survey - Summary Statistics

	Mean	Median	Std. Dev.
Daily Time Reading or Watching News (1-5)[1]	2.15	2	1.19
Frequency of Use (1-7)[2]			
Cable TV	1.93	1	1.49
Paper News	1.31	1	0.67
Radio	2.46	2	1.66
Online News	4.55	5	1.73
Social Media	5.60	6	1.56
News Feed	4.14	4	1.99
Political Nature of Preferred News (1-5)[3]	2.81	3	0.97
Daily Social Media Usage (hours)[4]			
Facebook	1.87	1	2.21
Instagram	1.28	1	1.60
Twitter	0.86	0	2.06
Tumblr	0.35	0	1.57
Snapchat	1.95	1	3.02
Vimeo	0.03	0	0.16
YouTube	1.85	1	2.65
Social Media Friends and Followers (number)[5]			
Facebook	640.99	538	442.04
Instagram	452.36	350	511.77
Tumblr	87.32	0	571.74
Twitter	182.12	0	333.80
Subjective Well-Being (0-10)[6]			
Satisfied with life	7.15	8	1.92
Things in life are worthwhile	7.37	8	1.88
How happy are you	7.17	8	2.12
How often do you worry	6.79	7	2.33
How often do you feel depressed	3.40	3	2.63

Notes: [1]Responses to the question "How much time did you spend reading or watching the news per day last week?" Response options: 1) Less than 15 min, 2) More than 15 minutes but less than 30 minutes, 3) More than 30 minutes but less than 1 hour, 4) More than 1 hour but less than 2 hours, and 5) More than 2 hours. N=167 obs. [2]Responses to the question "Please indicate how frequently you used the following types of news media last week." Scale was from 1 to 7 where 1 indicates "Not at all" and 7 indicates "All of the time." N=167 obs. [3]List top news outlets/sources from the previous week. We categorized each 1st choice as either being 1) Left, 2) Left-Center, 3) Center, 4) Right-Center, or 5) Right based on www.allsides.com. N=57 obs. [4]Time spent each say on various social media platforms. [5]How many friends and followers on various social media platforms. [6]Subjective well-being questions, with 0 indicating "never and" 10 "very/always." Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Table A.4: Correlations between the Value of Facebook and User's Characteristics

	Value of Facebook	High Time	High Engage	Depressed	High Negative	High Friends in Facebook	High Friends in other Social Media
Value of Facebook	1.00						
High Time	0.23***	1.00					
High Engage	0.20**	0.32***	1.00				
Depressed	-0.11	0.23***	0.05	1.00			
High Negative	-0.06	0.17**	0.09	0.32***	1.00		
High Friends on Facebook	0.06	0.10	0.21***	-0.02	0.01	1.00	
High Friends on other Social Media	0.17**	0.18**	0.38***	-0.10	0.01	0.42***	1.00

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

This table presents the Pearson correlation coefficients between the stated value of Facebook and characteristics of its users based on Phase 2 survey responses. High Time refers to individuals who on average use Facebook for more than one hour per day; High Engage refers to individuals who post pictures and comments on Facebook at least once or twice per month; Depressed refers to individuals who reported feeling depressed above the reported median value; High Negative refers to individuals who are above the median of the factor index that combines measures of feeling envy, misery, lonely and annoyed while on Facebook; High Friends in Facebook refers to individuals who have more than 564 friends in Facebook (median number of friends); and High Friends in other Social Media refers to individuals who have more than 529 friends in Facebook (median number of friends in other social media). Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

Table A.5: Distribution Shift Tests

	Equality	FSD C-T	SSD C-T	FSD T-C	SSD T-C
Facebook Use	0.00***	0.00***	0.00***	0.93	1.00
News Media Index -Traditional Media	0.41	0.22	0.10*	0.60	0.57
News Media Index -Social Media	0.00***	0.00***	0.00***	0.94	1.00
News Consumption Index	0.07*	0.04**	0.00***	0.95	0.75
Probability Right Answer - Mainstream News	0.37	0.77	0.80	0.18	0.23
Probability Wrong Answer - Mainstream News	0.61	0.70	0.55	0.33	0.34
Probability Not Sure Answer - Mainstream News	0.55	0.29	0.21	0.58	0.51
Probability Right Answer - Skewed News	0.01***	0.01***	0.01***	0.54	0.99
Probability Wrong Answer - Skewed News	0.46	0.50	0.75	0.23	0.23
Probability Not Sure Answer - Skewed News	0.37	0.51	0.81	0.19	0.19
Overall Satisfaction	0.25	0.11	0.20	0.58	0.76
Life is Worthwhile	0.28	0.14	0.09*	0.62	0.79
Feel Happy	0.17	0.09*	0.11	0.93	0.82
Worry	0.21	0.90	0.79	0.10*	0.11
Feel Depressed	0.32	0.16	0.22	0.91	0.98
Consumption Index	0.03**	0.97	0.89	0.01**	0.00***
Productive Time Index	0.10	0.97	0.90	0.05*	0.02**
Efficient Time Index	0.10*	0.98	0.90	0.05*	0.01***
Expected Consumption Index	0.07*	0.79	0.63	0.03**	0.01***
Value of Facebook	0.47	0.70	0.99	0.25	0.14

* $p < 0.1$ ** $p < 0.05$ *** $p < 0.01$

This table presents the bootstrap p-values of Kolmogorov-Smirnov statistics that test for equality of distributions, first order stochastic dominance and second order stochastic dominance between treatment and control after a one week Facebook restriction. In column 1 the null hypothesis is that the distributions are the same, in column 2 the null hypothesis is that the treatment group first order stochastic dominates the control group, in column 3 the null hypothesis is that the treatment group second order stochastic dominates the control group, in column 4 the null hypothesis is that the control group first order stochastic dominates the treatment group, and in column 5 the null hypothesis is that the control group first order stochastic dominates the treatment group. First order stochastic dominance and second order stochastic dominance are defined as in Abadie (2002). Reprinted with permission from Roberto Mosquera, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, The Economic Effects of Facebook, published 2019, Experimental Economics.

APPENDIX B

WHO SUPPORTS PIGOU? THE DISTRIBUTIONAL CONSEQUENCES OF PIGOUVIAN TAXES

B.1 Transportation Poll Stratified Regions

Figure B.1.1: TTI Regions

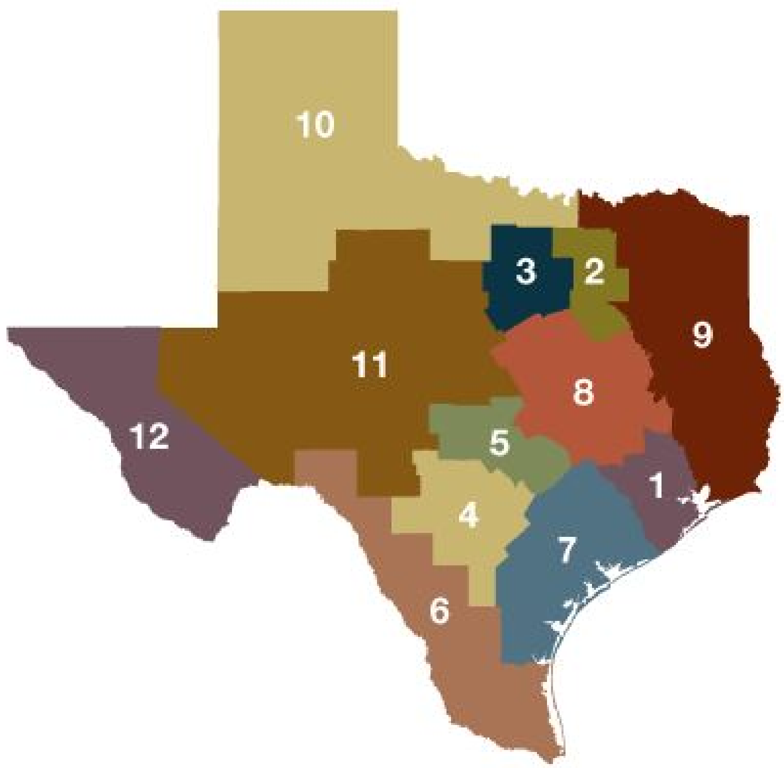

Region Designations:
Region 1 - Houston; Region 2 - Dallas; Region 3 - Fort Worth; Region 4 - San Antonio; Region 5 - Austin; Region 6 - Laredo, Pharr; Region 7 - Corpus Christi, Yoakum; Region 8 - Bryan, Waco; Region 9 - Atlanta, Beaumont, Lufkin, Paris, Tyler; Region 10 - Amarillo, Childress, Lubbock, Wichita Falls; Region 11 - Abilene, Brownwood, Odessa, San Angelo; Region 12 - El Paso

Reprinted from Texas A&M Transportation Institute's Texas Transportation Poll (2016).

B.2 Federal and State Gasoline Tax

Figure B.2.1: Federal Gas Tax Spending Allocation

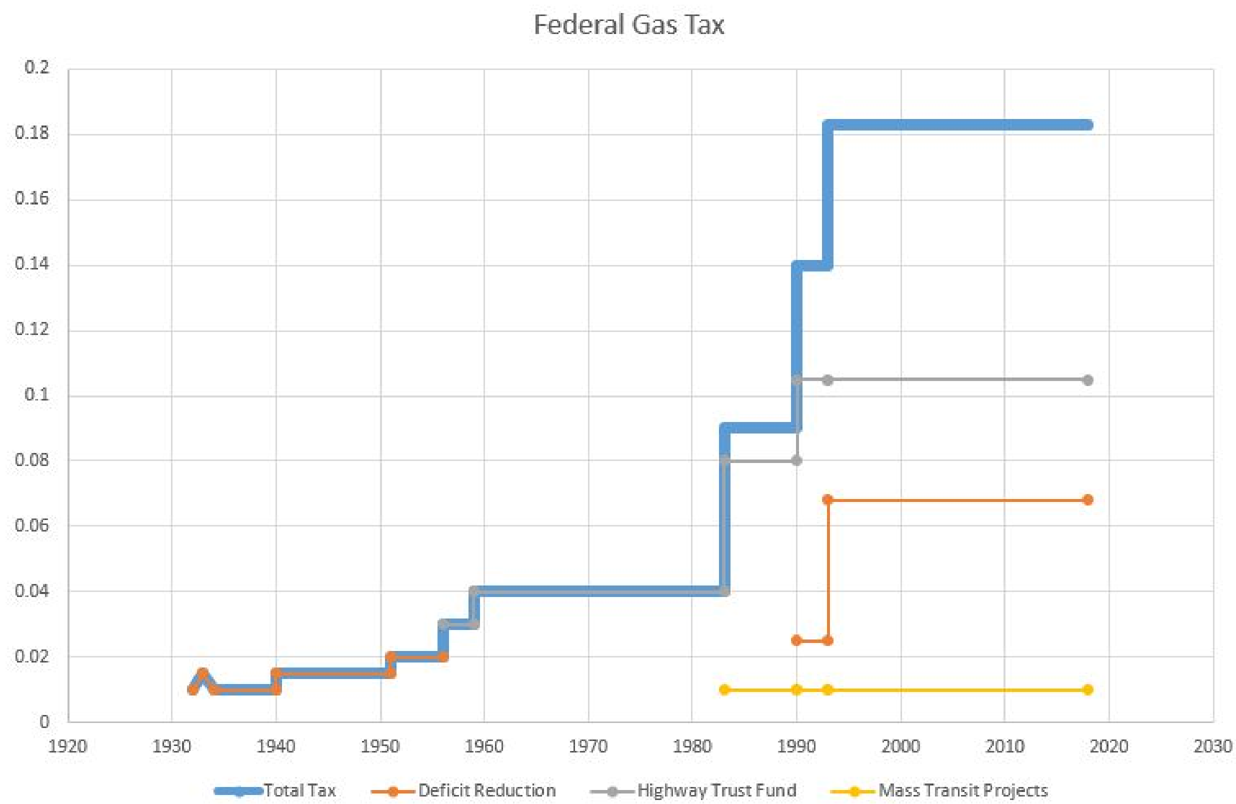

Gas Tax Uses:
This figure shows that historically, the federal gas tax has been used for two main purposes: 1) Deficit Reduction and 2) Infrastructure Spending via the Highway Trust Fund or Mass Transit Projects. Every time the gas was increased, each cent has been allocated to either one of these purposes.

Table B.2.1: State Gasoline Taxes

	Gasoline			Diesel		
	Excise	LUST Fees	Total	Excise	LUST Fee	Total
Federal	0.183	0.001	0.184	0.243	0.001	0.244

	Gasoline			Diesel		
	State	Taxes & Fees	State Total	State	Taxes & Fees	State Total
Alabama	0.240	0.010	0.250	0.250	0.018	0.268
Alaska	0.080	0.010	0.090	0.080	0.010	0.090
Arizona	0.180	0.010	0.190	0.180	0.010	0.190
Arkansas	0.245	0.003	0.248	0.285	0.003	0.288
California	0.473	0.082	0.555	0.360	0.342	0.702
Colorado	0.220	0.013	0.233	0.205	0.013	0.218
Connecticut	0.250	0.000	0.250	0.290	0.175	0.465
Delaware	0.230	0.000	0.230	0.220	0.000	0.220
District of Columbia	0.235	0.000	0.235	0.235	0.000	0.235
Florida	0.040	0.304	0.344	0.040	0.313	0.353
Georgia	0.279	0.008	0.287	0.313	0.008	0.321
Hawaii	0.160	0.025	0.185	0.160	0.025	0.185
Idaho	0.320	0.010	0.330	0.320	0.010	0.330
Illinois	0.380	0.151	0.531	0.455	0.151	0.606
Indiana	0.300	0.144	0.444	0.490	0.010	0.500
Iowa	0.305	0.000	0.305	0.325	0.000	0.325
Kansas	0.240	0.000	0.240	0.260	0.000	0.260
Kentucky	0.246	0.014	0.260	0.216	0.014	0.230
Louisiana	0.200	0.009	0.209	0.200	0.009	0.209
Maine	0.300	0.014	0.314	0.312	0.007	0.319
Maryland	0.262	0.107	0.369	0.270	0.107	0.376
Massachusetts	0.240	0.029	0.269	0.240	0.029	0.269
Michigan	0.263	0.136	0.399	0.263	0.165	0.428
Minnesota	0.285	0.001	0.286	0.285	0.001	0.286
Mississippi	0.180	0.004	0.184	0.180	0.004	0.184
Missouri	0.170	0.004	0.174	0.170	0.004	0.174
Montana	0.320	0.008	0.328	0.295	0.008	0.302
Nebraska	0.293	0.009	0.302	0.293	0.003	0.296
Nevada	0.230	0.008	0.238	0.270	0.008	0.278
New Hampshire	0.222	0.016	0.238	0.222	0.016	0.238
New Jersey	0.105	0.310	0.415	0.135	0.351	0.486
New Mexico	0.170	0.019	0.189	0.210	0.019	0.229
New York	0.080	0.258	0.338	0.080	0.240	0.320
North Carolina	0.361	0.003	0.364	0.361	0.003	0.364
North Dakota	0.230	0.000	0.230	0.230	0.000	0.230
Ohio	0.385	0.000	0.385	0.470	0.000	0.470
Oklahoma	0.190	0.010	0.200	0.190	0.010	0.200
Oregon	0.360	0.000	0.360	0.360	0.000	0.360
Pennsylvania	0.576	0.011	0.587	0.741	0.011	0.752
Rhode Island	0.340	0.011	0.351	0.340	0.011	0.351
South Carolina	0.220	0.008	0.228	0.220	0.008	0.228
South Dakota	0.280	0.020	0.300	0.280	0.020	0.300
Tennessee	0.260	0.014	0.274	0.270	0.014	0.284
Texas	0.200	0.000	0.200	0.200	0.000	0.200
Utah	0.311	0.007	0.318	0.311	0.007	0.318
Vermont	0.121	0.187	0.308	0.280	0.040	0.320
Virginia	0.162	0.006	0.168	0.202	0.006	0.208
Washington	0.494	0.027	0.521	0.494	0.027	0.521
West Virginia	0.205	0.152	0.357	0.205	0.152	0.357
Wisconsin	0.309	0.020	0.329	0.309	0.020	0.329
Wyoming	0.230	0.010	0.240	0.230	0.010	0.240

Values are as reported from the EIA's State-by-State Fuel taxes. Last updated on February 2020.

B.3 Gasoline Tax Support Maps

Figure B.3.1: Political Support Maps

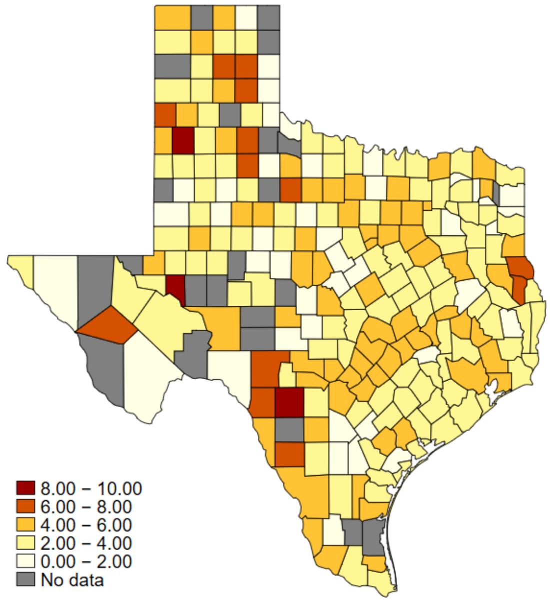

(a) 5 Cent Support

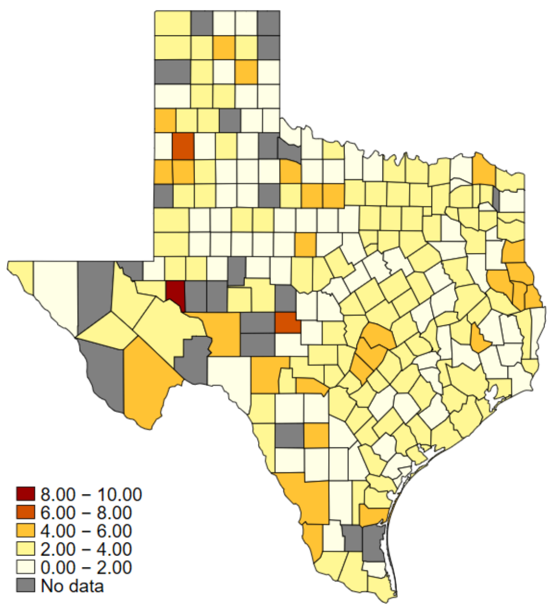

(b) 10 Cent Support

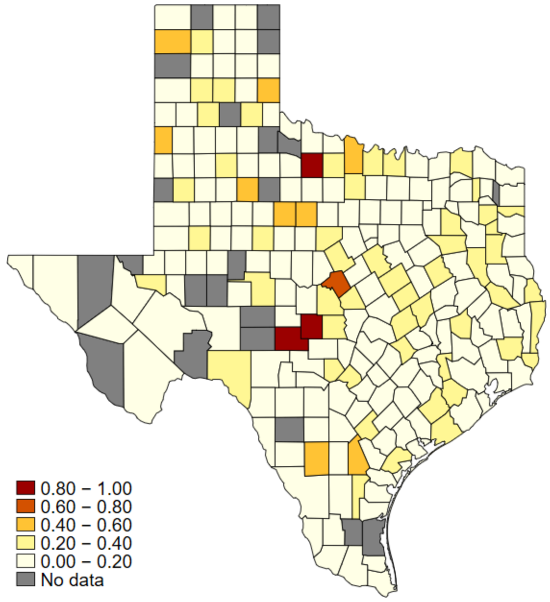

(c) Oppose Increase

Figure B.3.2: Historical Retail Gasoline Prices in Texas (2004-2010)

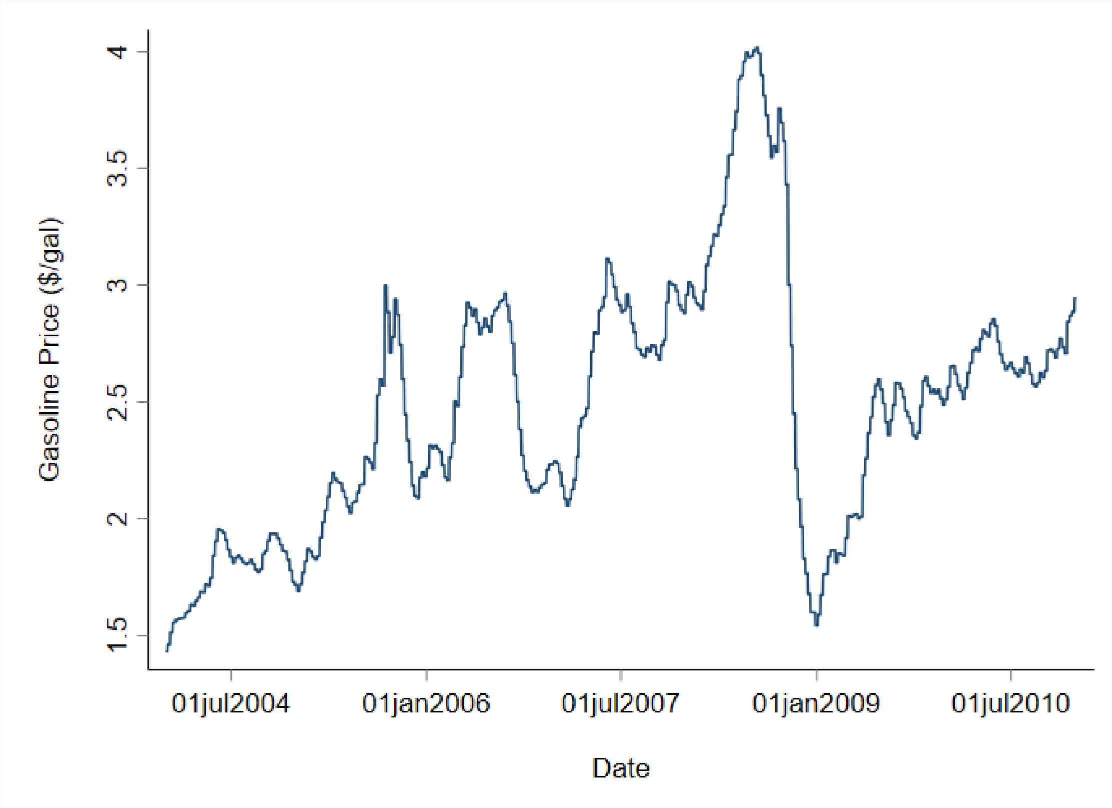

This figure presents a time series of the average daily retail gasoline price in Texas as reported by the EPA. Prices are in terms of dollars per gallon.

Figure B.3.3: County "Win" Proportions

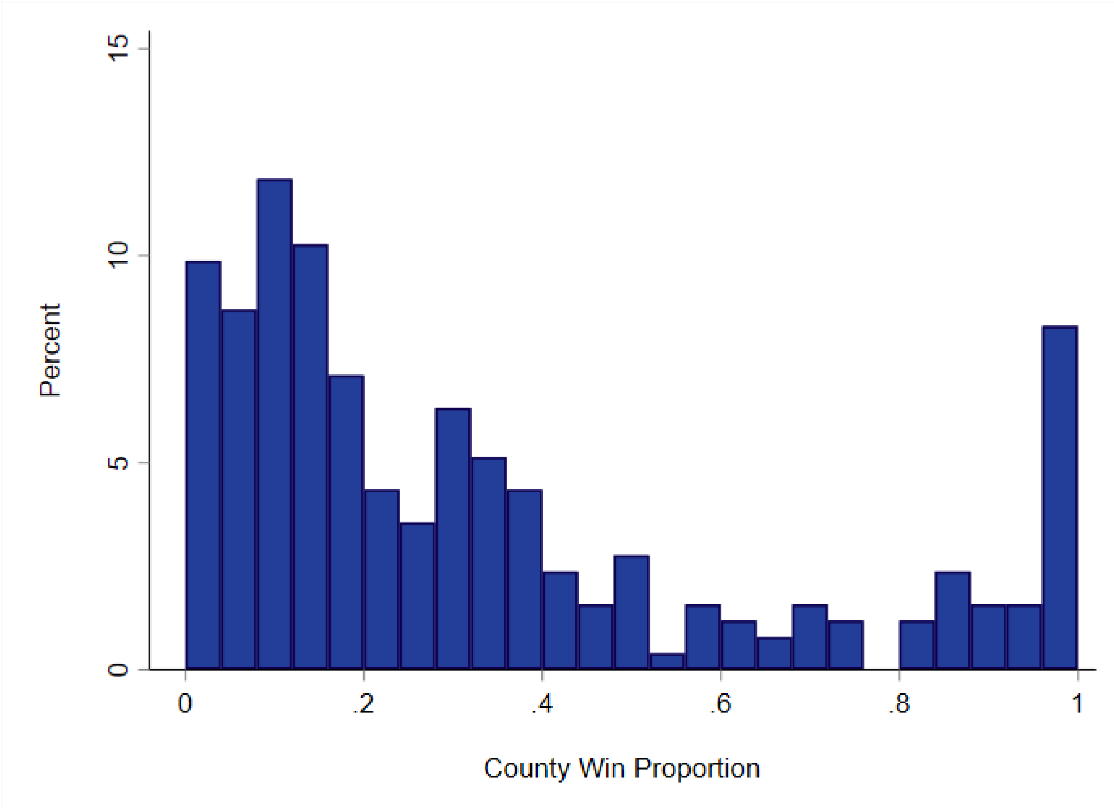

This figure presents results for the front-end calculations, and depicts a histogram of the proportion of "Winners" (benefits > costs) in each county. "Winners" tend to be concentrated in the more urban regions of the state, but can also be driven in rural counties where pollution damages are high.

B.4 Full Regressions

Table B.4.1: County Win Percent on 5 Cent Support

	1	2	3	4	5	6
County Win Percent	0.18***	0.17***	0.14***	0.12***	0.11**	0.10**
	(0.03)	(0.03)	(0.05)	(0.05)	(0.05)	(0.04)
Republican		-0.09***	-0.09***	-0.11***	-0.12***	-0.11***
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Independent/Other		-0.06***	-0.06***	-0.07***	-0.09***	-0.07***
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
High Mileage		-0.07***	-0.07***	-0.04**	-0.02	-0.02
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Air Quality Violations			<-0.01	<-0.01	<-0.01	<-0.01
			<(0.01)	<(0.01)	<(0.01)	<(0.01)
County GOP Voting Percentage			-0.04	-0.05	-0.05	-0.03
			(0.11)	(0.09)	(0.09)	(0.09)
High Education				0.13***	0.12***	0.12***
				(0.02)	(0.02)	(0.02)
Minority				-0.08***	-0.08***	-0.08***
				(0.02)	(0.02)	(0.02)
Employed					<-0.01	<-0.01
					(0.02)	(0.02)
Registered Voter					-0.02	<-0.01
					(0.03)	(0.03)
Old Age					-0.01	-0.02
					(0.02)	(0.02)
High Income					0.04***	0.06***
					(0.02)	(0.02)
Female					-0.08***	-0.09***
					(0.01)	(0.01)
Local Gov. Support						0.06**
						(0.03)
State Gov. Support						0.08***
						(0.02)
Federal Gov. Support						0.04**
						(0.02)
Household Vehicle #						-0.01
						<(0.01)
Household Size						<-0.01
						<(0.01)
Constant	0.28***	0.36***	0.40***	0.38***	0.44***	0.31***
	(0.02)	(0.02)	(0.09)	(0.08)	(0.08)	(0.09)
R-sq	0.016	0.022	0.023	0.044	0.053	0.063
N	3631	3631	3631	3631	3631	3631

Standard errors in parentheses: * p<0.1, ** p<0.05, *** p<0.01. Standard errors are clustered at the county level. Regression (1) reports baseline results, and (2) adds covariates related to political party affiliation. Regression (3) covariates include a county level control for the number of daily ambient air quality violations, and population. Regression (4) covariates add controls for education and minority status. Regression (5) covariates add employment status, a dummy for registered voters, age, and gender. Regression (6) covariates also include an individual's support for local and state government, number of vehicles owned, and household size.

Table B.4.2: County Win Percent on 5 Cent Support (Raw)

	1	2	3	4	5	6
County Win Percent	1.252***	1.152***	0.931***	0.728**	0.690**	0.657**
	(0.236)	(0.227)	(0.357)	(0.334)	(0.325)	(0.320)
Republican		-0.650***	-0.615***	-0.755***	-0.852***	-0.747***
		(0.160)	(0.142)	(0.147)	(0.147)	(0.154)
Independent/Other		-0.461***	-0.446***	-0.473***	-0.618***	-0.545***
		(0.158)	(0.146)	(0.149)	(0.158)	(0.156)
High Mileage		-0.513***	-0.521***	-0.278***	-0.127	-0.136
		(0.109)	(0.110)	(0.106)	(0.108)	(0.112)
Air Quality Violations			-0.008	-0.007	-0.008	-0.008
			(0.015)	(0.013)	(0.013)	(0.013)
County GOP Voting Percentage			-0.443	-0.573	-0.521	-0.385
			(0.636)	(0.556)	(0.551)	(0.555)
High Education				1.018***	0.927***	0.915***
				(0.106)	(0.113)	(0.113)
Minority				-0.549***	-0.618***	-0.610***
				(0.110)	(0.117)	(0.116)
Employed					0.112	0.146
					(0.112)	(0.114)
Registered Voter					-0.150	-0.083
					(0.213)	(0.209)
Old Age					-0.205	-0.277*
					(0.138)	(0.142)
High Income					0.313***	0.394***
					(0.109)	(0.114)
Female					-0.579***	-0.615***
					(0.097)	(0.094)
Local Gov. Support						0.348**
						(0.169)
State Gov. Support						0.549***
						(0.150)
Federal Gov. Support						0.297***
						(0.111)
Household Vehicle #						-0.063
						(0.042)
Household Size						-0.064**
						(0.029)
Constant	3.025***	3.613***	3.971***	3.878***	4.350***	3.512***
	(0.154)	(0.174)	(0.536)	(0.488)	(0.559)	(0.596)
R-sq	0.013	0.021	0.023	0.048	0.059	0.067
N	4059	4059	4059	4059	4059	4059

Standard errors in parentheses: * p<0.1, ** p<0.05, *** p<0.01. Standard errors are clustered at the county level. Regression (1) reports baseline results, and (2) adds covariates related to political party affiliation. Regression (3) covariates include a county level control for the number of daily ambient air quality violations, and population. Regression (4) covariates add controls for education and minority status. Regression (5) covariates add employment status, a dummy for registered voters, age, and gender. Regression (6) covariates also include an individual's support for local and state government, number of vehicles owned, and household size.

Table B.4.3: County Win Percent on 10 Cent Support

	1	2	3	4	5	6
County Win Percentage	0.16***	0.14***	0.09**	0.07*	0.07*	0.07
	(0.04)	(0.03)	(0.04)	(0.04)	(0.04)	(0.04)
Republican		-0.12***	-0.11***	-0.13***	-0.14***	-0.11***
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Independent/Other		-0.08***	-0.08***	-0.08***	-0.10***	-0.08***
		(0.03)	(0.02)	(0.03)	(0.03)	(0.03)
High Mileage		-0.03	-0.03	-0.01	0.01	0.01
		(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Air Quality Violations			<0.01	<0.01	<0.01	<0.01
			<(0.01)	<(0.01)	<(0.01)	<(0.01)
County GOP Voting Percentage			-0.10	-0.10	-0.09	-0.06
			(0.10)	(0.09)	(0.09)	(0.08)
High Education				0.09***	0.08***	0.08***
				(0.02)	(0.02)	(0.02)
Minority				-0.04**	-0.05***	-0.05***
				(0.02)	(0.02)	(0.02)
Employed					-0.02	-0.01
					(0.01)	(0.01)
Registered Voter					-0.02	-0.01
					(0.02)	(0.02)
Old Age					-0.04*	-0.04*
					(0.02)	(0.02)
High Income					0.03**	0.04***
					(0.01)	(0.01)
Female					-0.09***	-0.10***
					(0.02)	(0.01)
Local Gov. Support						0.01
						(0.03)
State Gov. Support						0.05**
						(0.02)
Federal Gov. Support						0.07***
						(0.01)
Household Vehicle #						-0.02***
						<(0.01)
Household Size						0.0000
						<(0.01)
Constant	0.16***	0.25***	0.33***	0.32***	0.41***	0.31***
	(0.02)	(0.02)	(0.08)	(0.07)	(0.08)	(0.08)
N	3314	3314	3314	3314	3314	3314

Standard errors in parentheses: * p<0.1, ** p<0.05, *** p<0.01. Standard errors are clustered at the county level. Regression (1) reports baseline results, and (2) adds covariates related to political party affiliation. Regression (3) covariates include a county level control for the number of daily ambient air quality violations, and population. Regression (4) covariates add controls for education and minority status. Regression (5) covariates add employment status, a dummy for registered voters, age, and gender. Regression (6) covariates also include an individual's support for local and state government, number of vehicles owned, and household size.

APPENDIX C

PRICE LEADERSHIP AND LEARNING IN OLIGOPOLY: EVIDENCE FROM ELECTRICITY MARKETS

C.1 Nodal Electricity Grid

Figure C.1.1: Nodal Prices without Congestion

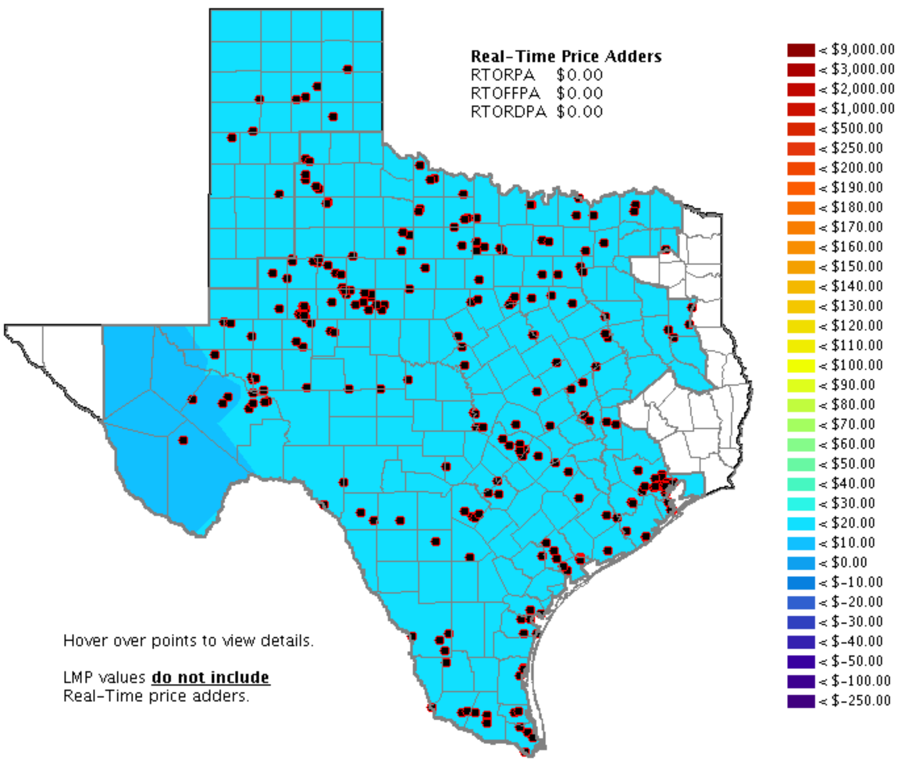

Figure C.1.2: Nodal Prices with Congestion

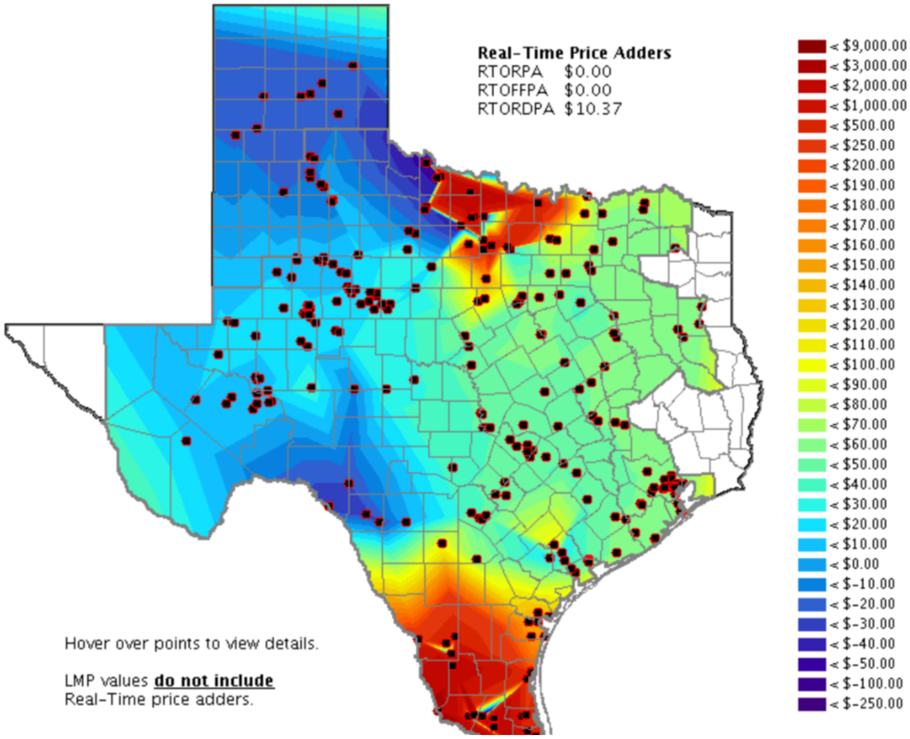

C.2 Market Summary

Figure C.2.1: Actual Vs. Predicted Market Price

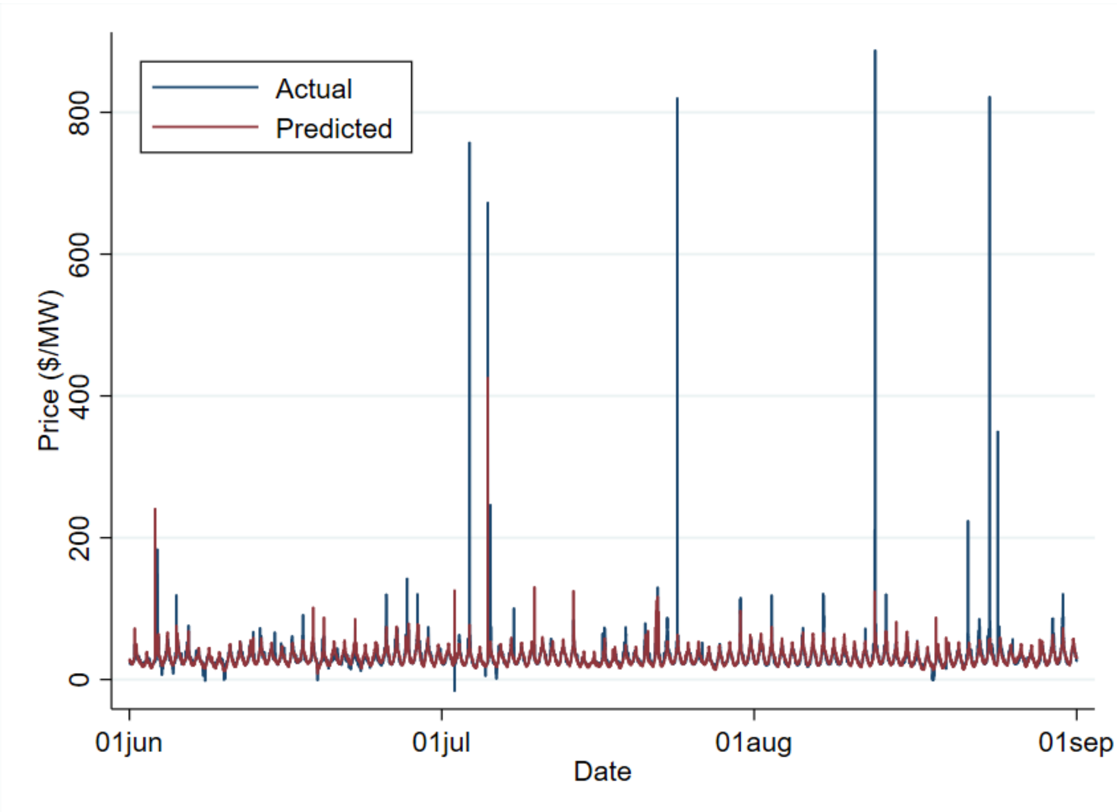

Figure C.2.2: Hourly Market Summary

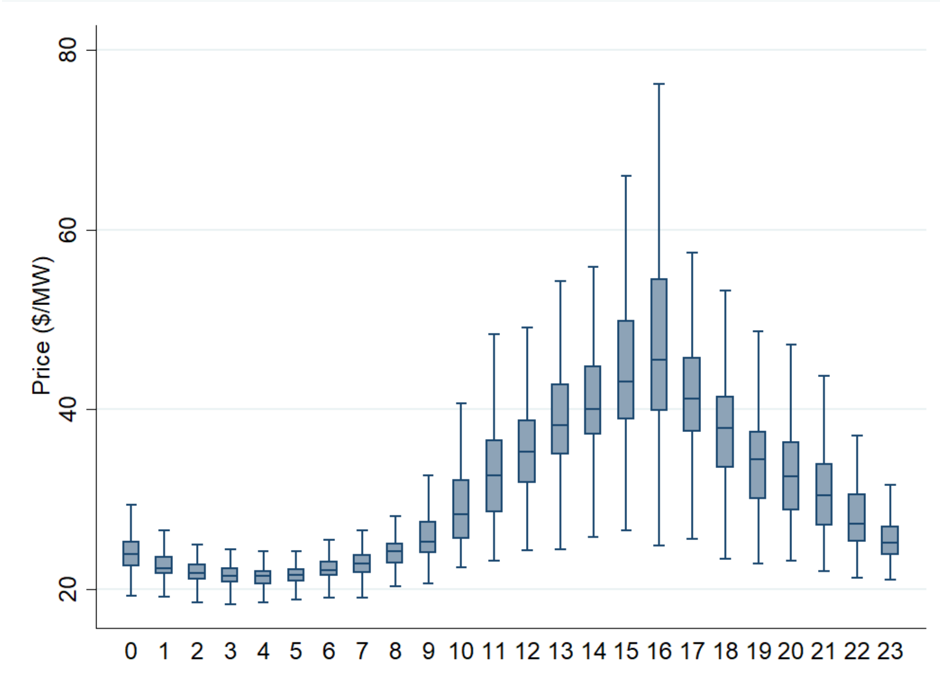

(a) Market Price

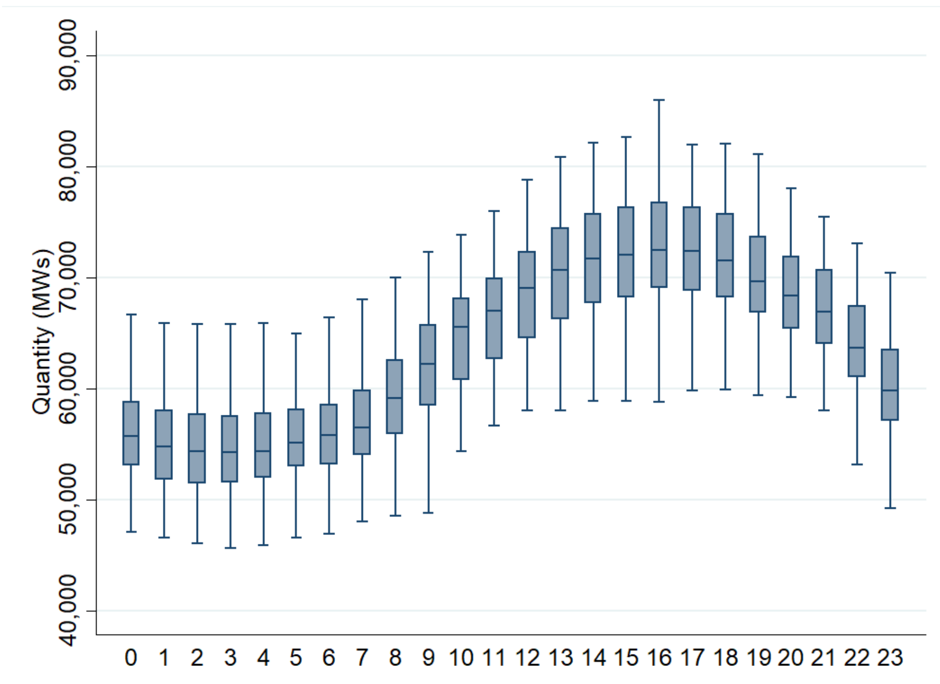

(b) Capacity

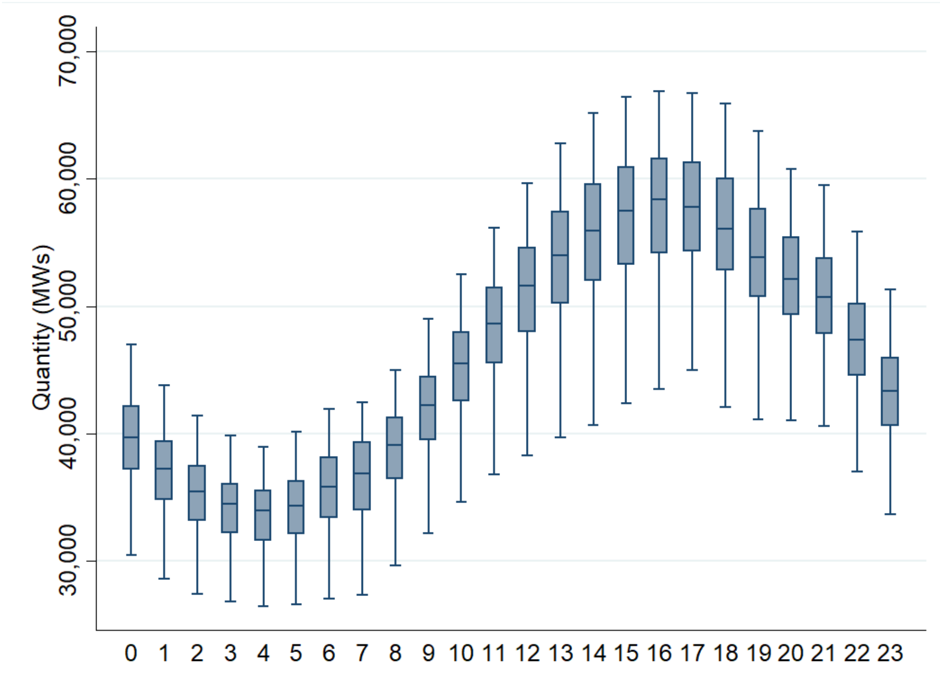

(c) Aggregate Load

Figure C.2.3: Excess Capacity by Month

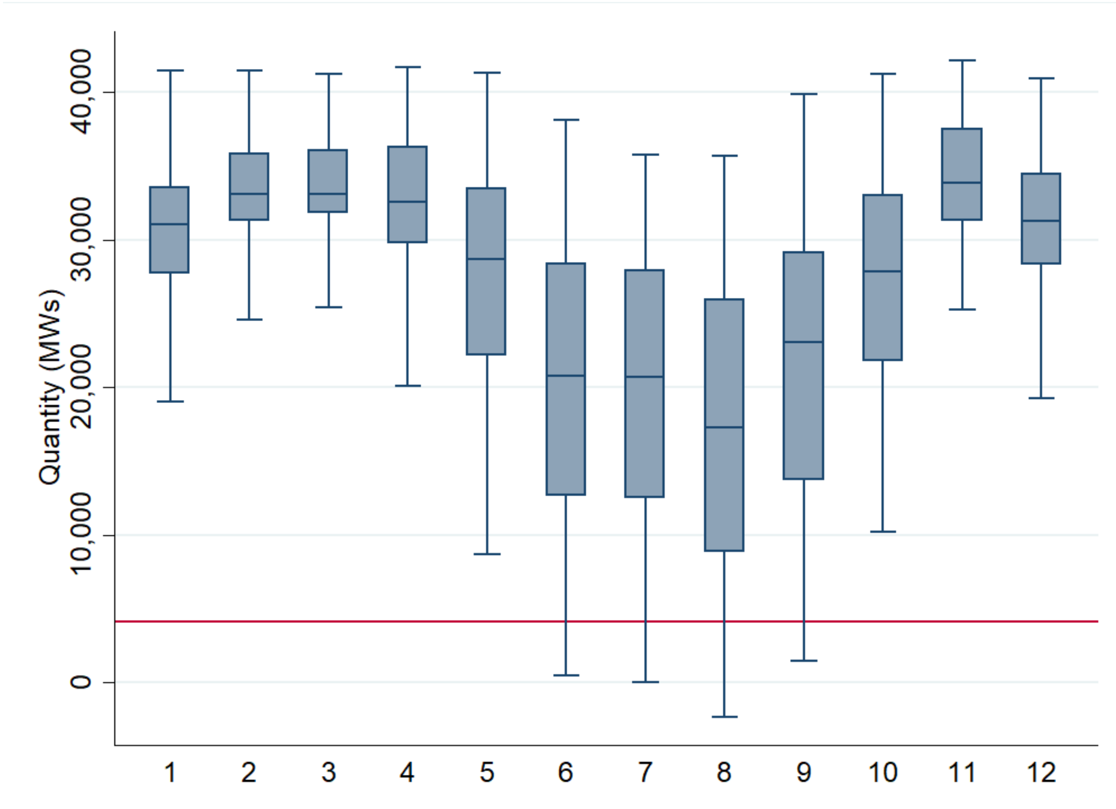

Figure C.2.4: GDF Historical Pricing

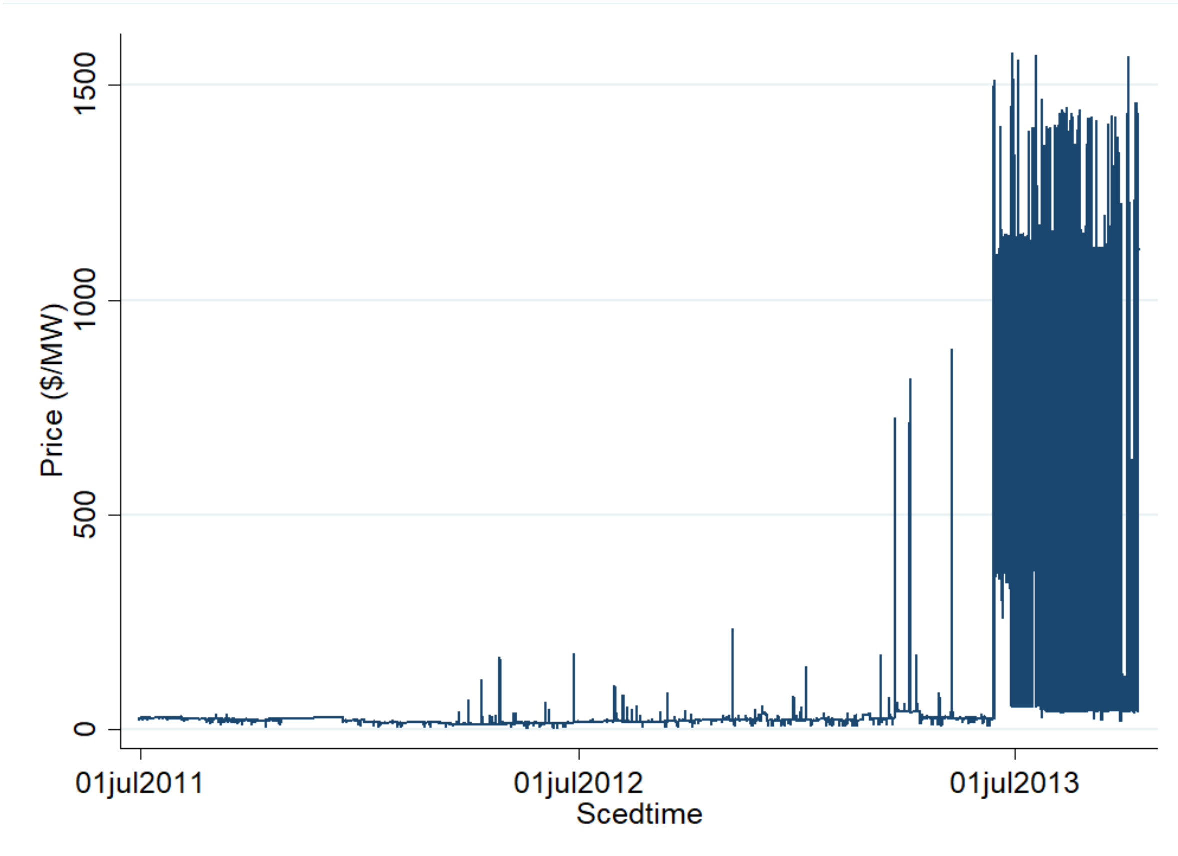

C.3 Deviations in "Offer" Space

Figure C.3.1: Deviation from Ex-Post Best Response Bidding

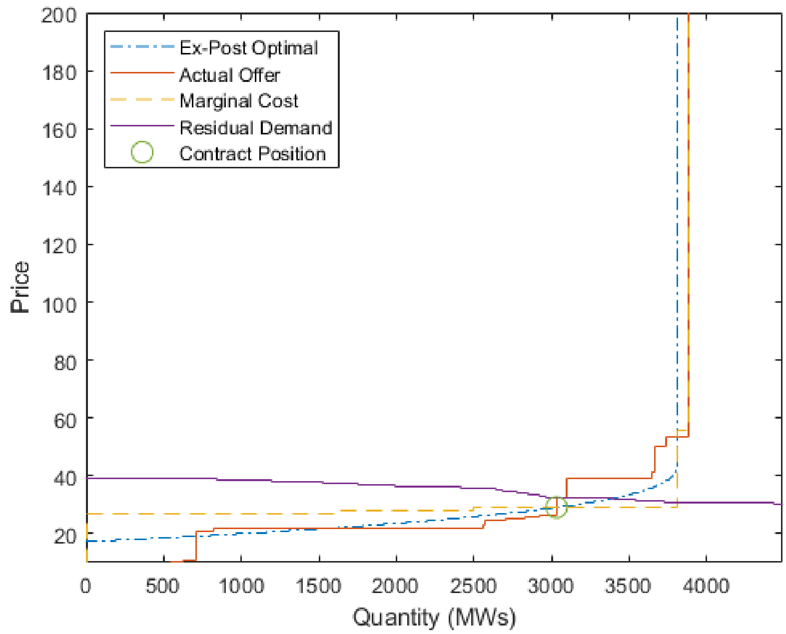

(a) GDF Period 1

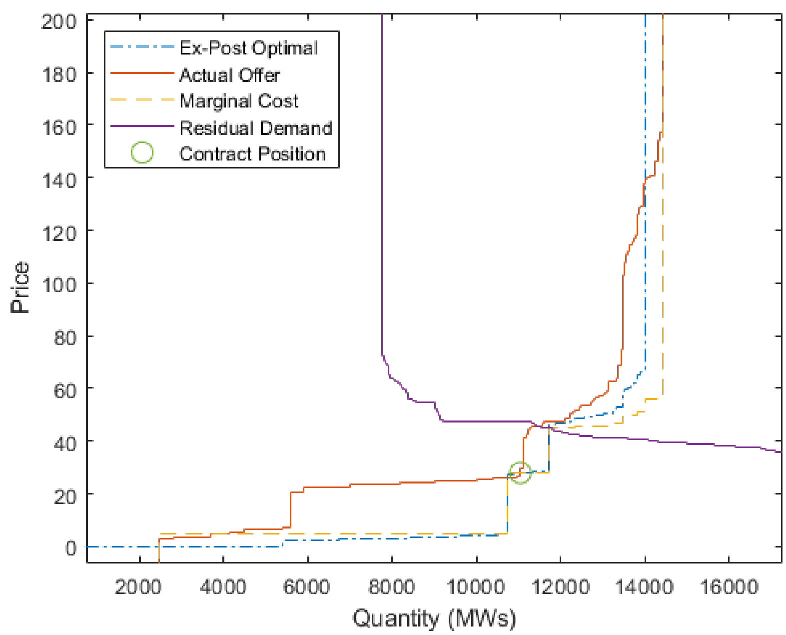

(b) Luminant Period 1

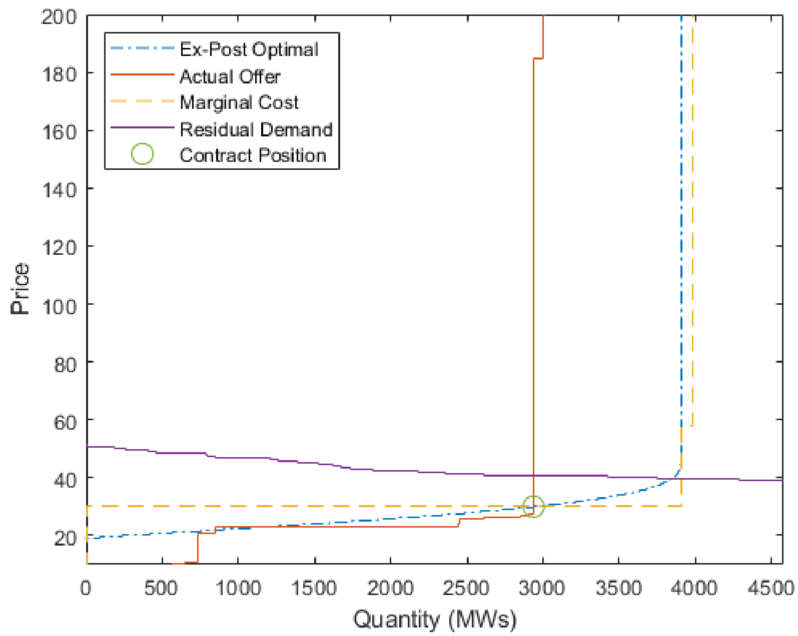

(c) GDF Period 3

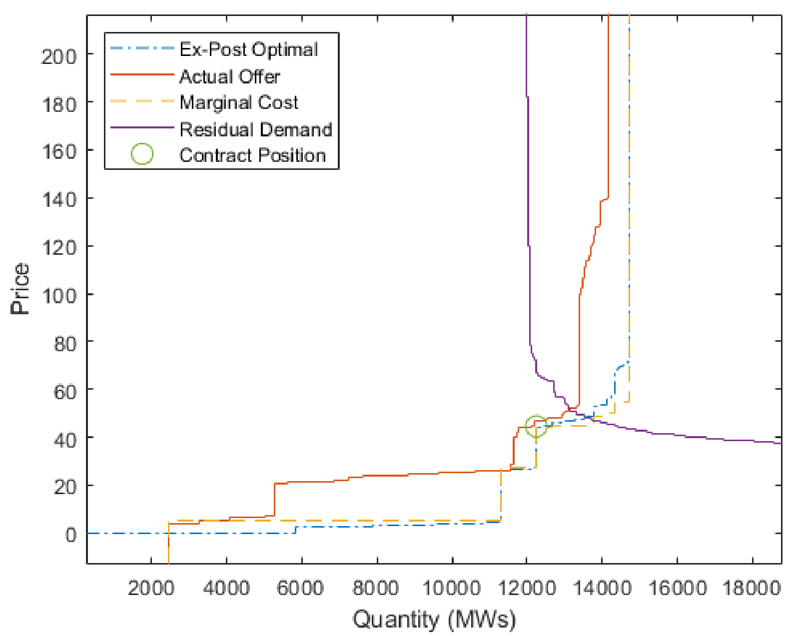

(d) Luminant Period 3

C.4 Additional Results

Table C.4.1: Foregone Profits by Leader

	Actual Price	Estimated Price	GDF-BR	GDF-OLD
	(1)	(2)	(3)	(4)
Panel A: Not Adjusted for Tx				
Foregone Profits			4332.67	3432.85
			(5267.71)	(5387.083)
Scenario Price	45.67	45.40	43.21	43.92
	(15.13)	(10.03)	(6.43)	(7.63)
Panel B: Adjusted for Tx				
Foregone Profits			4213.13	3469.26
			(3571.18)	(4528.19)
Scenario Price	44.76	45.29	43.18	43.71
	(12.11)	(8.86)	(5.44)	(5.35)

Standard deviations in parentheses. Panel A reports the foregone profits for the Leader during the afternoon offers for the latter half of June, 2013. Panel B removed potentially congested periods. Column 1 reports the average market clearing price P_t. Column 2 shows the estimated market clearing price. Column 3 reports the counterfactual impact had GDF offered ex-post optimal bids in Periods 2 and 3 whereas Column 4 reports the counterfactual impact had GDF offered a randomly selected offer from Period 1 in Periods 2 and 3. In this light, Column 3 can be thought of as an upper-bound and Column 4 can be thought of as a lower bound for the impact of GDF's deviations.

Table C.4.2: Foregone Profits by Follower

	Actual Price	Estimated Price	LUM-BR	LUM-OLD
	(1)	(2)	(3)	(4)
Panel A: Not Adjusted				
Foregone Profits			1503.44	1205.83
			(2975.41)	(2607.14)
Scenario Price	46.79	47.88	45.82	47.49
	(19.21)	(14.25)	(10.59)	(16.38)
Panel B: Adjusted for Tx				
Foregone Profits			1453.51	1212.14
			(2866.01)	(2633.48)
Scenario Price	45.27	46.73	44.81	46.53
	(17.64)	(13.56)	(9.80)	(15.56)

Standard deviations in parentheses. Panel A reports the foregone profits for the Follower during the afternoon offers for the latter half of June, 2013. Panel B removed potentially congested periods. Column 1 reports the average market clearing price P_t. Column 2 shows the estimated market clearing price. Column 3 reports the counterfactual impact had Luminant offered ex-post optimal bids in Periods 2 and 3 whereas Column 4 reports the counterfactual impact had Luminant offered a randomly selected offer from Period 1 in Periods 2 and 3. In this light, Column 3 can be thought of as an upper-bound and Column 4 can be thought of as a lower bound for the impact of Luminant's deviations.

Figure C.4.1: Foregone Profits for Leader and Follower

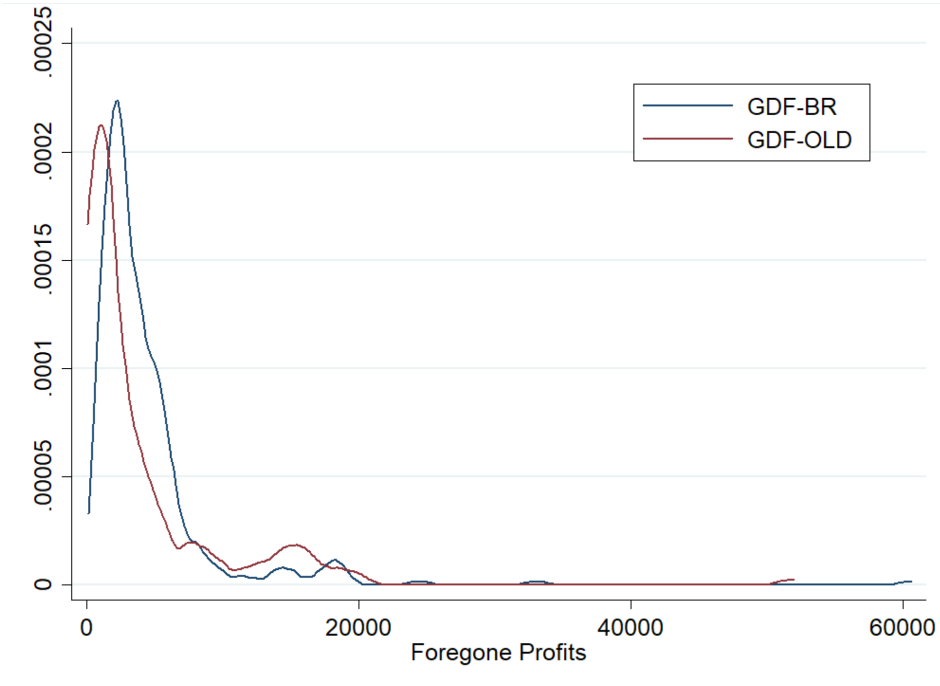

(a) Leader's Foregone Profits

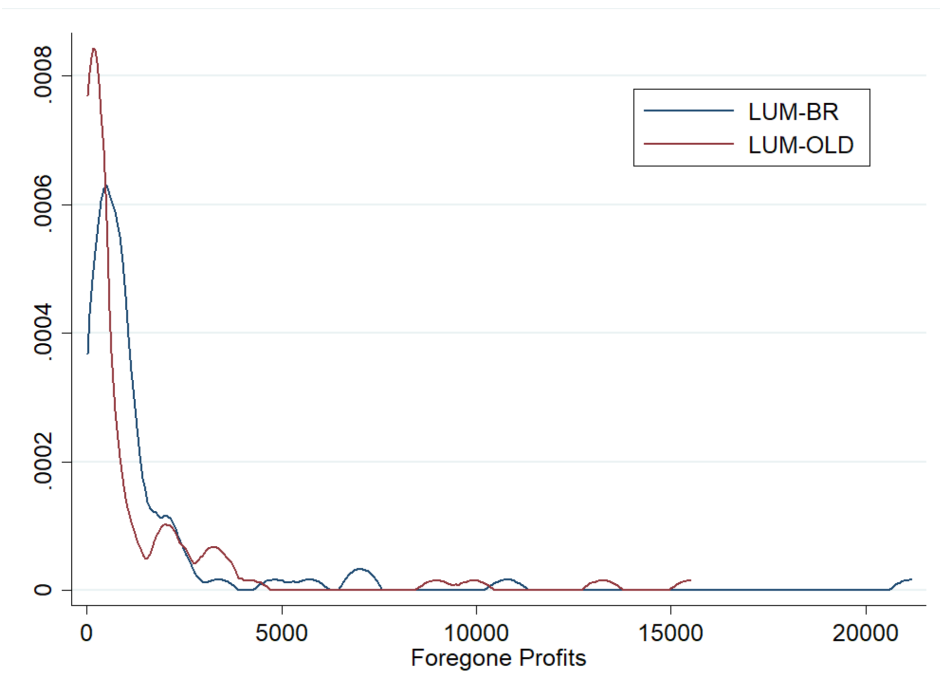

(b) Follower's Foregone Profits

Figure C.4.2: Price Impacts by Leader and Follower

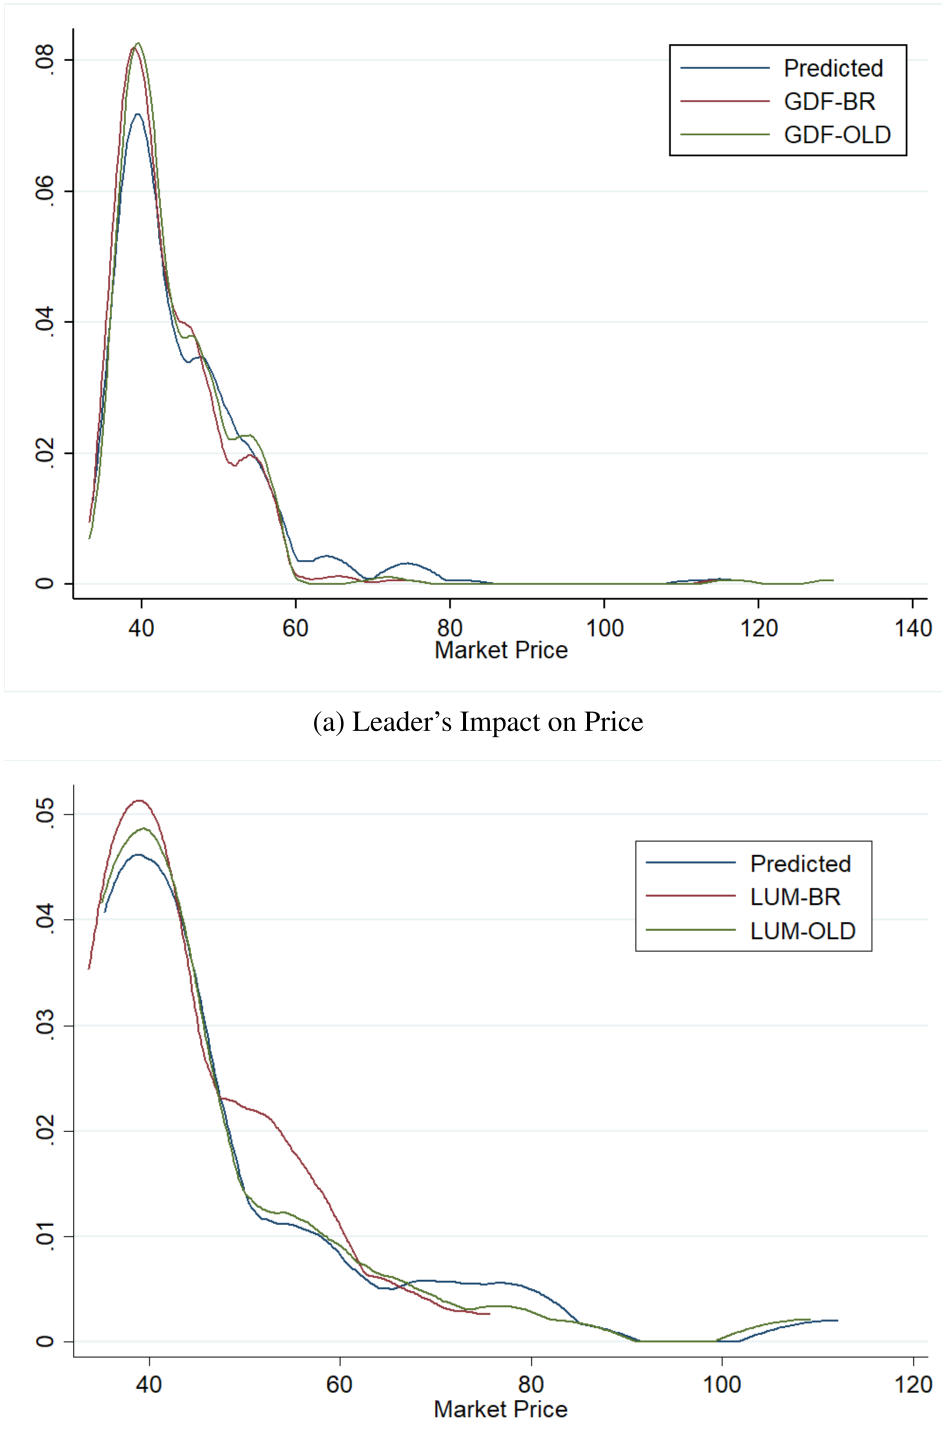

(a) Leader's Impact on Price

(b) Follower's Impact on Price