Environmental Regulation, Imperfect Competition and Market Spillovers: The Impact of the 1990 Clean Air Act Amendments on the US Oil Refining Industry

Richard L. Sweeney[*]

April 23, 2015

Abstract

The 1990 Clean Air Act Amendments imposed extensive restrictions on refined petroleum product markets, requiring select end users to purchase new cleaner versions of gasoline and diesel. I estimate the impact of this intervention on refining costs, product prices and consumer welfare. Isolating these effects is complicated by several challenges likely to appear in other regulatory settings, including overlap between regulated and non-regulated markets and deviations from perfect competition. Using a rich database of refinery operations, I estimate a structural model that incorporates each of these dimensions, and then use this cost structure to simulate policy counterfactuals. I find that the policies increased gasoline production costs by 7 cents per gallon and diesel costs by 3 cents per gallon on average, although these costs varied considerably across refineries. As a result of these restrictions, consumers in regulated markets experienced welfare losses on the order of $3.7 billion per year, but this welfare loss was partially offset by gains of $1.5 billion dollars per year among consumers in markets not subject to regulation. The results highlight the importance of accounting for imperfect competition and market spillovers when assessing the cost of environmental regulation.

JEL Codes: L11, L13, L71, Q48, Q52

[*]John F. Kennedy School of Government, Harvard University, rich_sweeney@hksphd.harvard.edu. I am grateful to Joe Aldy, Erich Muehlegger, Ariel Pakes and Rob Stavins for their continual guidance and encouragement. I thank Peter Berck, Thomas Covert, Evan Herrnstadt, Zhenyu Lai, Greg Lewis, Mark Shepard, Che-Lin Su, Thomas Wollmann, Ali Yurukoglu, Matthew Zaragoza-Watkins and participants of the Harvard Industrial Organization and Environmental Economics lunches and workshops for their helpful comments and discussions. Funding from a Resources for the Future Joseph L. Fisher Dissertation Fellowship, Vicki Norberg-Bohm Fellowship, Joseph Crump Fellowship, and Harvard University Dissertation Completion Fellowship is gratefully acknowledged. The data used in this paper were obtained under a researcher information access agreement with the Energy Information Administration, and executed with extensive support from Jacob Bournazian and Joseph Conklin of EIA. The views and analysis herein, including any errors, are mine alone.

1 Introduction

The US oil refining industry is by far the largest and most sophisticated in the world, producing over $750 billion worth of products in 2012. The 1990 Clean Air Act Amendments imposed a series of new restrictions on refineries, described at the time as the greatest operational challenge the industry had faced since the 1970s.[1] A major component of these regulations was the requirement that gasoline and diesel fuel meet higher environmental standards in certain markets. Areas of the country with severe ozone problems were required to use a new grade of gasoline, called reformulated gasoline (RFG), and highway diesel consumers were required to purchase a new low sulfur grade of distillate, called low sulfur diesel (LSD). These regulated products have since traded at a premium over their unregulated counterparts.

While price increases were expected, it was also widely acknowledged that the cost of producing these new fuels would vary across refineries (NPC 1993). The imposition of fuel content regulations balkanized previously integrated gasoline and distillate markets, with the extent of fragmentation varying considerably across regions of the country. This combination of heterogeneous compliance costs and differential access to newly created markets suggests that regulation aimed at reducing pollution externalities could have come at the expense of decreased allocative efficiency. The refining industry is characterized by a relatively small number of firms and enormous barriers to entry, making it a perennial concern for lawmakers and competition authorities (GAO 2004, FTC 2004, 2011). Initial evidence from the introduction of RFG found that the extent of price changes was correlated with the number of suppliers in a market (Brown et al., 2009). However, in-depth study mapping these price changes to refinery costs, profits, and consumer welfare has not been done, partly due to a lack of data availability. Government studies analyzing the impact of regulation on refineries use engineering models and run simulations that assume perfect competition.

In this paper I estimate the variable cost of producing the new fuels mandated under the 1990 Clean Air Act Amendments, as well as the the impact on refined product markets. In order to do this, I obtained access to a rich, previously unused confidential database of refinery-level production decisions. For every refinery in the United States, I observe detailed information on all product outputs, crude inputs, and installed technology, by month, beginning in 1986.

Estimating the impact of these regulations on refineries is complicated by a number of factors that are likely to affect other regulatory settings, particularly those related to energy and environmental policy. The main challenge is a lack of suitable control refineries or markets. Although content regulations varied geographically, supply and demand patterns overlap in way such that every refinery either served a market that became regulated or was linked strategically to one. Bulow, Geanakoplos, and Klemperer (1985) showed that, in oligopolistic settings, a firm's actions in one market can change a competitor's strategies in a second market. It is therefore likely that content regulations not only affected prices and production in directly regulated markets, but also spilled over into unregulated markets as well. The situation is

[1] Scherr, Richard, G. Allan Smalley Jr., and Michael E. Norman. 5/27/1991. "Refining in the '90s", *Oil & Gas Journal*. Accessed 10/14/2014.

complicated further by the fact that all refineries are multiproduct firms. Faced with increased costs or competition in one product market, such as gasoline, refineries could have responded by increasing relative production of other products, such as jet fuel. Accounting for this cross-product adjustment margin is important for understanding the net costs of the policy.

I overcome these challenges, as well as the more fundamental challenge of costs not being observed, by estimating a structural model of refinery decisions which directly incorporates each of these dimensions. Employing methods first introduced by Rosse (1970), the empirical approach is to estimate the costs of producing reformulated gasoline and low sulfur diesel by comparing refiners' willingness to supply gasoline and distillate across seemingly similarly profitable situations before and after the regulation. I assume constant elasticity of demand and Cournot competition, and develop a multiproduct marginal cost function for each refinery. After recovering the cost structure for the industry, I then simulate counterfactuals with the fuel content restrictions removed to calculate the price and welfare effects of the policy.

The main finding of the paper is that reformulated gasoline increased refinery costs by 7 cents per gallon and low sulfur diesel increased costs by 3 cents per gallon on average, although these costs varied considerably across refineries. I find that the demand for petroleum products at the wholesale level is more elastic than would be surmised from end-user consumption patterns alone. As a result, while refineries producing reformulated gasoline saw markups increase in those markets, average refinery profits from 1995 to 2003 were 8 percent lower than they would have been without content regulation. Contributing to this decline in profits was a reduction in margins in conventional product markets, as refineries which found it difficult to produce the new fuels reallocated output elsewhere. Thus, while consumers in regulated markets experienced welfare losses on the order of $34 billion during this period, this loss was partially offset by gains of $14 billion dollars among consumers in markets not subject to regulation.

Several papers have looked at the impact of gasoline content regulation on prices. Muehlegger (2006) estimates the relationship between regional content regulations and gasoline price spikes, finding that content regulation contributed to price spikes in California, Illinois and Wisconsin. Brown et al. (2008) use detailed weekly wholesale price data from 1994 to 1998, and compare prices across matched regulated and unregulated cities in the years immediately before and after gasoline content regulations went into effect. They find that regulation increased prices by 3 cents per gallon on average, but that this effect varied considerably across cities depending on their degree of market isolation. Chakravorty, Nauges, and Thomas (2008) use data from 1995 to 2002 to estimate the impact of "boutique" gasoline standards on state-level wholesale prices. They model differences in gasoline standards across states as being endogenous to the concentration of refineries in each state, and find that OLS estimates understate the effect of regulation on prices. Although no papers have estimated the cost of low sulfur diesel, Zhang (2011) estimates the cost of the subsequent switch from low sulfur to ultra-low sulfur diesel in 2006.[2] This paper extends this literature by incorporating refinery-level data

[2]In addition to these papers, several authors have empirically studied the oil refining industry. Berman and Bui (2001) find that productivity at refineries in Los Angeles increased during a period when they were subject to stringent point source emission regulations. Considine (2001) describes a structural model of markup pricing

and separately estimating costs and price effects. I show that costs vary significantly across refineries, and directly relate this cost heterogeneity to changes in market power and markups.

This paper also contributes to a growing empirical literature at the intersection of industrial organization and environmental economics (Millimet, Roy, and Sengupta 2009). Ryan (2012) estimates a dynamic structural model of the Portland cement industry to assess the impact of the 1990 Clean Air Act Amendments on that industry. He finds that focusing on static prices and profits alone generates negative cost estimates, but that the sign is reversed once changes in fixed costs are incorporated. In a subsequent paper, Fowlie, Reguant, and Ryan (2014) use this model to assess the relationship between alternative carbon pricing policies and product market distortions in this setting. They find that policies which fail to account for strategic firm behavior generate social welfare losses. In this paper, I show that allowing for imperfect competition and accounting for strategic interdependence between markets is important understanding the full effects of environmental regulation. Previous reduced form studies of the Clean Air Act have implicitly calculated the gross effect of regulation on regulated versus unregulated areas or firms, leaving the net impact ambiguous. In this paper, I account for intra-country shifting by modeling refinery decisions explicitly and then simulating policy counterfactuals to recover the net national effect of fuel content regulations. The results highlight the importance of allowing for market spillovers, and demonstrate how detailed firm-level data can be combined with assumptions about producer behavior to recover regulatory impacts in settings where the program evaluation framework is not applicable.

The remainder of the paper proceeds as follows. Section 2 provides a brief overview of the refining industry and describes the relevant environmental regulation. Section 3 introduces and summarizes the data. Section 4 develops a structural model of refinery behavior and describes the estimation procedure. Section 5 presents the main results of the paper, and Section 6 concludes.[3]

2 Institutional background

To provide intuition for the structural model specified below, in this section I give a brief overview of the refining industry and relevant environmental regulation. A key feature of the industry is that all refineries produce multiple products and production is an inherently joint process. The efficiency of this process varies across refineries, driven by differences in crude oil input quality and installed technology. Geography is also important, as pipelines are by far the cheapest way to transport products, and every refinery is not connected to every state by pipeline. This paper focuses on fuel content regulation, which mandated that certain drivers purchase reformulated gasoline or low sulfur diesel. The share of local gasoline and distillate

under joint production at refineries. Hendricks et al (2007) introduce a model of bilateral oligopoly to study mergers in this industry. Hastings and Gilbert (2007) find evidence of raising rivals costs after Tosco's acquisition of Unocal. Chesnes (2014) studies the impact of refinery outages on product prices and refinery investment. Further discussion of how this paper relates to prior structural models of the industry is delayed until Section 4, after more background on the industry is provided.

[3] Each section begins with a brief summary of the key points for the time constrained reader.

markets covered under the regulation varied across states, and, due to the incompleteness of the pipeline network, this meant that some refineries were more affected by the regulations than others.

2.1 The oil refining industry

Refineries lie at the middle of the US transportation fuel supply chain (Figure 1). Crude oil is extracted upstream, domestically or abroad, processed at a refinery, and then shipped out via pipeline or barge to wholesale terminals, where it is distributed by truck for local consumption. Crude oil as it comes out of the ground is a mixture of different length hydrocarbons, ranging from short hydrocarbons, which roughly correspond to butane and gasoline, to long hydrocarbons, which correspond to asphalt and tar. At the most basic level, oil refining consists of separating crude into streams of differing densities using heat and a complex series of catalytic processes. The "lighter" end products, which include gasoline, diesel and propane, are typically of much higher value. So, all else equal, a refiner tries to maximize the amount of light outputs produced from a given amount of crude oil.

The finished product mix obtainable from a barrel of crude oil is a function of the type of crude that is used and the type of processing capital the refinery has installed. Crude oil properties vary across oil fields, and are typically described by two characteristics, API gravity and sulfur content. API gravity, denominated in degrees, is a measure of crude density, which dictates the relative proportions of light and heavy oils that can be separated out during simple distillation. Heavier crudes contain a relatively larger share of long hydrocarbons, which translates into a larger share of heavy end products. Sulfur content, denominated in parts per million (ppm), was historically of interest because it causes corrosion in metals and other processing complications. More recently, environmental regulations have set caps on the amount of sulfur end products can contain. Thus, light "sweet" (i.e. low sulfur) crudes, like West Texas Intermediate (WTI), are typically the most valuable.

While API gravity and sulfur content map fairly directly into the proportions of refined outputs obtained from simple distillation, modern refineries in the United States are much more complex operations (Figure 2). It is therefore more useful to think of gravity and sulfur content as determining the amount of processing necessary to transform a given type of crude into a particular end product mix. The most basic refining technology is the distillation tower, which separates crude into different density streams by slowly boiling it. These streams can then be sent through any number of secondary processes, collectively referred to as "upgrading" capacity, which increase the yield of higher value light end-products. Finally, there are a host of technologies that can be employed at the end of the process to remove pollutants and impurities. From an environmental perspective, the most important of these processes is hydrotreating, which removes sulfur. Refineries in the US range from simple "topping" operations, which do not have any upgrading capacity and have light yields of less than 50 percent, to the most complex refineries in the world, where the amount of upgrading capacity exceeds distillation capacity and light yields routinely top 90 percent. However, even at the most sophisticated operations, it is impossible to transform an entire barrel of crude into gasoline or diesel, making refining an inherently joint production process.

The final important dimension of differentiation in the industry is geography. Figure 3 presents a map of refineries and refined product pipelines in the continental United States. Historically, refineries were set up where crude oil was easily available. As a result, US refining is relatively geographically concentrated and not particularly well correlated with the location of end-users. This map actually understates the extent of regional concentration, as the refineries in the Gulf Coast are much larger than other areas, giving this region almost 50 percent of total US capacity. In order to balance the location of supply with demand, an extensive pipeline distribution system evolved over the course of the 20th century to transport refined products to local markets. These pipeline routes are typically unidirectional, with product generally flowing north from refining centers to populations centers. There are two key points to note about this pattern. East Coast refining capacity is only about a third of its consumption. As result, it receives around 50 percent of its supply from the Gulf Coast and accounts for almost all of the refined product imports in the US. On the opposite extreme, the West Coast is shielded from European imports as well as domestic imports via pipeline from the rest of the country. It therefore relies on refineries in the region for almost all of its consumption.

2.2 Environmental regulation

The refining industry is one of the largest sources of air pollution in the United States, contributing both directly through emissions generated during the refining process and indirectly through the combustion of petroleum products at end sources. As a result, refineries have been subjected to considerable environmental regulation over the last half century, stemming primarily from the Clean Air Act of 1963 and subsequent amendments (Table 1). Direct emissions were covered under regulations which placed increased oversight on polluting facilities located in counties which did not meet newly established National Ambient Air Quality Standards (NAAQS) (Greenstone 2002). The 1970 Clean Air Act Amendments also targeted indirect refinery emissions by permitting regulation of the chemical composition of refined petroleum products. This authority was used to phase out lead in gasoline starting in 1975, and to reduce volatile organic compounds (VOCs) and other ozone precursors by imposing Reid vapor pressure (RVP) limits in summer months beginning in 1989.[4] The 1990 Clean Air Act Amendments (CAAA) marked the most significant regulation of indirect refinery emissions to date, requiring oxygenate be added to gasoline in some markets and a new reformulated version of gasoline be supplied in others, depending on the region and season.[5] The CAAA also imposed strict sulfur

[4]Reid vapor pressure is a measure of the propensity of gasoline to evaporate. RVP regulation was implemented in two phases, affecting summer months starting in 1989 and 1992. See Auffhammer and Kellogg (2011) for more information on RVP regulation.

[5]See Muehlegger (2004) for a thorough review of gasoline content regulation.

In addition to these refinery specific programs, many refineries were also covered under New Source Review. This program subjected large polluting facilities located in non-attainment areas to undergo additional regulatory review before making investments or plant alterations. Approximately 70% of refining capacity was located in regions that were out of attainment. Although this designation restricted investment opportunities, it did directly limits on highway diesel fuel.

In this paper, I focus on reformulated gasoline (RFG) and low sulfur diesel (LSD) from 1994 to 2003. Of the regulations introduced by the 1990 Clean Air Act Amendments, these were the two that called for direct alterations to the refining process.[6] In 1999, new sulfur limits were announced for gasoline and highway diesel, which phased in starting in 2004 and 2006 respectively. These new programs allowed for flexibility in compliance. Specifically, firms could average sulfur reductions across refineries, buy or sell excess reductions from other firms, and bank early compliance for future credit. Unfortunately, actual sulfur content levels and modes of compliance for each refinery are not publicly available.[7] Incorporating this phase-in period and flexibility requires a substantial extension to the model specified below, and is left for future work.

2.2.1 Reformulated Gasoline

The 1990 Clean Air Act Amendments mandated the adoption of reformulated gasoline in nine large metropolitan markets with severe ozone pollution. RFG is gasoline manufactured to reduce the amount of smog forming particles and toxic pollutants released into the air during combustion. The program was implemented in two phases, coming online in 1995 and 2000. Both phases set minimum oxygenate levels of 2 percent and capped benzene levels at 1 percent of volume. Phase I also required a 15 percent reduction in toxic air pollutants relative to conventional gasoline, and this reduction was increased to 25 percent in Phase II. In addition to these year round requirements, RFG imposed stricter VOC limits during summer ozone season (June 1 - September 15), again mandating reductions of 15 and 25 percent in Phases I and II. Finally, Phase II added a year round NOx requirement of 5.5 percent. While only nine areas were required to use RFG, other areas with moderate ozone pollution were allowed to opt into the program. Today, RFG is used in 17 states and the District of Columbia, making up 30 percent of US gasoline consumption (Figure 4). In March 1996, California and Arizona adopted a more stringent version of RFG, called CARB gasoline, which imposed tighter seasonal VOC limits and an 80 percent reduction in sulfur content.

In order to qualify as reformulated, gasoline had to meet specific composition and emission performance criteria. Rather than an exact formula, there were many different ways to satisfy the criteria, allowing refineries the flexibility to make different tradeoffs based on the economics of their operations. As a result, producing RFG is not associated with any specific piece of refinery equipment or characteristic, although, at a minimum, refineries had to reconfigure their operations at the beginning of the program. For many refineries, the extent of reconfiguration required was too costly, and they opted not to participate in this market.

---

affect operations conditional on installed capital. In future work I plan to study the long run impacts of this program on refinery investment by nesting the static profit structure estimated in this paper within a dynamic refinery investment game.

[6]The other CAAA content regulation, oxygenated gasoline, simply involved blending in oxygenate at the end of the refining process. This was often done by a third party downstream.

[7]A Freedom of Information Act request to obtain realized trades and compliance methods was denied by the EPA.

2.2.2 Low Sulfur Diesel

In order to facilitate new particulate emissions standards on heavy duty diesel engines, the CAAA capped the sulfur content of on-highway diesel at 500 ppm starting in October 1993. Distillate fuel oil is a general classification of relatively heavier petroleum products used for domestic heating, industrial burners, and compression in ignition engines. Diesel fuel is distillate fuel burned in diesel engines. "On-highway" diesel fuel is diesel fuel used by trucks and passenger cars, whereas "off-highway" diesel is used in farms, construction, and marine vessels. Distillate fuel oil is thus primarily categorized by end use, rather than physical properties. This distinction is important, because the new distillate regulations imposed by the CAAA only applied to one of these categories. Home heating oil and other similar distillates were not required to meet the standard, which affected 46 percent of distillate demand, and around 8 percent of total petroleum demand at the time of enactment. Although the delineation was not as stark as RFG, there is still considerable heterogeneity in the fraction of distillate that was regulated under the program across states, driven by differences in relative demand by end-use. Figure 5 presents a map with the each state's share of distillate consumption that was low sulfur after the regulation began. In the northeast, substantial amounts of distillate are used for home heating oil, meaning that less than 50 percent of those markets were covered under the new regulation.[8] At the other extreme, in the southwest, over 90 percent of distillate sales after 1994 were highway diesel.

Producing this new low sulfur distillate was a significant achievment for refineries, as the national average sulfur content at the time was 3000 ppm (Lidderdale 1993). In contrast to RFG, LSD production was largely determined by the installed capacity of hydrotreating or "desulfurization" capacity. According to Lidderdale (1993), refineries with catalytic hydrocracking units may be able to reconfigure them to remove some sulfur, but, otherwise, LSD production would be largely limited by desulfurization capacity. Smaller, less sophisticated refineries were therefore particularly vulnerable to this regulation, and, in an effort to compensate for this, were given SO2 credits for sulfur removed from diesel. California again adopted a slightly more stringent version of the regulation, imposing a 10 percent aromatics limit on highway diesel in addition to the 500 ppm sulfur cap.

3 Data description and summary

In this section I describe and summarize the data obtained from EIA. Detailed input and production data is observed for every refinery in the United States. Sales are observed at the firm-product-state level, and the demand side of these transactions is primarily comprised of intermediaries which reallocate products to end consumers. The sample is restricted to 124 large, non-specialty refineries, 20 of which exit between 1986 and 2003. There is considerable variation in the level of participation in reformulated gasoline and low sulfur diesel markets across these refineries, driven largely by geographic location. Despite this exogenous variation in

[8] In other parts of the country, off-highway distillate is mainly used in farming equimpent or marine vessels.

regulatory exposure, reduced form regressions of refinery RFG and LSD production on gasoline and distillate productivity fail to identify any significant costs associated with these regulations.

3.1 Description of confidential EIA data

Through a confidential data request I obtained several data sets on refinery operations from the Energy Information Administration (EIA), each described in appendix Table A.1. The main data come from survey EIA-810, which contains monthly information on all inputs and outputs for every refinery in the United States from 1986 to 2012. Importantly, gasoline is reported as being conventional or reformulated, and distillate is broken out by sulfur level. The average API gravity and sulfur content of crude inputs is also reported, along with the amount of distillation capacity in operation at the start of each month. This monthly distillation capacity information is supplemented with annual data from survey EIA-820, which records the capacity of every refinery unit, including all upgrading and desulfurization capacity, at the start of each year.[9]

This refinery-level data is combined with several firm-level data sets. Survey EIA-782A is a census of monthly state-level sales by every firm which owns a refinery in the United States. Refiners report sales in the state where the transfer of title occurred, regardless of where that product is ultimately consumed. Both the volume sold and the price are reported, broken out by sales to end users (retail) and sales for resale (wholesale). Survey EIA-782C is a census of all "prime suppliers", which includes firms that own refineries as well as large importers and marketers. 782C asks respondents to only report sales for which they are the final supplier into a state and the fuel is going to be consumed within state. The 782C data does not contain price and does not break volume sold into retail and wholesale.

The distinction between the 782A and 782C data is important for understanding the demand system specified in this paper. Firms record 782A sales in the state where the transfer of product occurred and 782C sales in the state where the product is ultimately consumed. In this sense, the quantity reported on 782A reflects the total quantity demanded *from refineries* for transfer in each state, while 782C reflects the quantity demanded *by end-users* in a given state from all distributors.[10] In this paper, I assume the relevant demand curve facing refineries is reflected in the 782A data. There are several reasons for this. First, survey 782A records the location and price of transactions where they occurred, and it is the price at and shipping costs to this location, rather than where the purchaser ultimately transports the product, which presumably matters to refineries. Second, the 782A data breaks sales down by retail and wholesale channels, and I use this information in estimation below. Finally, as the purpose of the 782C data is to measure state-level consumption, many transactions are excluded from the data in order to avoid double counting.[11]

[9] EIA-820 was not collected in 1996 and 1998. For these years, I interpolate capacity for each refinery based on the reported values from the adjacent years

[10] In general, states with many refineries, located at key points in the pipeline network tend to have quantities in 782A which exceed prime supplier volumes in 782C, whereas states that are net importers of refined products the opposite is true. EIA (2009) discusses the differences between the surveys in great detail.

[11] For example, if a refinery sells product to a distributor who then resells it in another state, quantity from this first transaction would not appear in the 782C data.

One major limitation of the data is that distribution is observed at the firm level, as opposed to the refinery level. I attempt to overcome this by assuming that firms minimize transportation costs when serving end markets. I obtained GIS maps of the US refined product pipeline system and waterways suitable for petroleum transportation from EIA, along with GIS coordinates of each refinery. Following Muehlegger (2006), costs for transporting petroleum products by pipelines, barges and trucks of 2, 4.5 and 30 cents per gallon per thousand miles are taken from estimates presented before the Federal Trade Commission (Jacobs 2002). For each refinery, I calculate the least cost method of serving each state. I then minimize each firm's total cost of meeting its observed state-level sales from the 782A data, subject to the observed 810 output at each of its refineries.

The final confidential data set comes from survey EIA-14, which contains average crude oil prices, including cost of delivery, at the firm-PADD-month level starting in 2004. For earlier years I use predicted crude delivery price as a function of benchmark crude spot prices, region dummies, regional domestic crude first purchase prices, and crude prices binned by API gravity (Appendix B). This confidential data is supplemented with annual state-level population data from the Census Bureau, petroleum product taxes and vehicle registrations from the Federal Highway Administration, and monthly weather data from the National Oceanic and Atmospheric Administration.

3.2 Summary Statistics

Table 2 presents refinery summary statistics from the data. The sample is restricted to a pre period (1986-1992) and a post period (1995-2003). 1993 and 1994 are omitted because LSD and RFG, respectively, appear in production data for months before they are tracked in the consumption data, so price and destination are not observed. Refineries are grouped by Petroleum Administration Defense Districts (PADDs), which is a commonly used geographic aggregation dating back to World War II. Figure 6 provides a map of the regions, as well as a proportional representation of the refineries in the sample. 220 refineries appear at some point in the data. Small refineries, with less than 10,000 barrels per day operating capacity, and specialty refineries, with less than 50 percent light production, are excluded from the analysis.

Of the 124 refineries included in the sample, 9 exited between 1986 and 1995, and another 11 exited by 2003.[12] Most of these exits represent the tail end of an industry restructuring after deregulation. In 1981, the US removed domestic oil price controls, causing 101 relatively small or inefficient refineries to exit the industry between 1981 and 1985. This trend of exits continued, albeit at a much slower pace, through the 1990's. Figure 7 reports the number of operating refineries and total industry distillation capacity in each year. Although a large number of refineries exited the industry, they were relatively small operations, and the amount

[12] In four of these exits, the refinery was sold to a nearby refinery and integrated into that refinery's operations. Of the 11 refineries that exit post 1994, two in California and two Illinois produced RFG. The two in CA exited in the summer of 1995, citing inability to comply with the stricter CARB phase II specifications (FTC 2006). Premcor shut down its Illinois refineries in 2001 and 2002 citing high capital costs to meet upcoming sulfur regulations.

of operating capacity lost was more than offset by expansions at refineries that stayed. No new refineries have been built in the US since the 1970's. Despite this, market structure varies substantially over the course of this period, driven by a wave of mergers and acquisitions in the 1990s (Table 3).

Figure 8 displays national price and quantity trends for gasoline and distillate during the sample. Total gasoline volumes were relatively constant, while total distillate sales have been trending upward since the late 1980's due to increasing demand for highway diesel. Prices of all products varied substantially over this time period, driven mainly by global crude price movements (Choinard and Perloff 2008).

Turning back to the refinery summary statistics in Table 2, the average yields of gasoline, distillate and jet fuel are 45 percent, 23 percent, and 6 percent respectively, with much more variation across refineries within each region than across regions. The Gulf Coast has the largest refineries, with an average ability to process 150,000 barrels of crude oil per day. The West Coast has the most sophisticated refineries, with an average ratio of upgrading capacity to distillation capacity of over 67 percent. The West Coast also uses the heaviest crude oil during this period, while refineries in the Rocky Mountain region are both the smallest and the least sophisticated.

The next two rows in the table report the number of refineries in the sample that produce RFG and LSD by region. Only 54 of the 111 refineries operating after 1994 produce reformulated gasoline, where as 104 refineries produce low sulfur diesel. On the intensive margin, there is heterogeneity across refineries in the share of each regulated product produced (Figures 9 and 10). As was discussed in Section 2.2, geography is a major determinant of market access in this industry, and regions differed in how affected they were by content regulation. Table 4 formally tests the extent to which this influenced refinery production of these newly created products. The dependent variable in these regressions is the share of a refinery's gasoline that is RFG, the share of a refinery's distillate production that is LSD, and the change in desulfurization capacity between 1990 and 1996. The independent variable is the average share of RFG and LSD post-1994 in states that each refinery was serving prior to 1990, when the regulations were announced. In all three regressions the share of a refinery's pre-1990 markets which were subsequently subject to CAAA content regulation significantly determines its post-1994 production shares of the new fuels.

The bottom section of Table 2 reports average wholesale prices for each product and region. Conventional gasoline prices were higher on the coasts prior to the introduction of RFG. In the post 1994 period, RFG prices were around 10 cents per gallon higher than conventional prices. Low sulfur diesel was 6 cents per gallon more than high sulfur distillate in PADD 5, and only 1 cent per gallon more in PADD 4, which had the highest distillate prices to begin with. Appendix Table A.1 reports average prices by state, along with the average Herfindahl–Hirschman Index (HHI) for each market. The FTC generally considers markets with an HHI in excess of 0.15 to be moderately concentrated. At the state level, 11 gasoline and distillate markets met this criterion in the pre period, and 16 conventional markets, 10 RFG markets, and 19 LSD and HSD markets were above this threshold in the post period.

3.3 Reduced form results

Before turning to the structural model, it is useful to look at the reduced form relationship between RFG and LSD production and refinery productivity. There are two channels through which the regulation could have affected operations: by altering the amount of gasoline or diesel extracted from each barrel of crude, or by changing the amount of crude the refiner chose to process. One measure of refinery productivity is therefore the ratio of gasoline or distillate output to crude distillation capacity, which is the product of these two channels. Table 5 presents results from the following regression,

$$q_{ijt}/k_{it} = \beta X_{it} + \alpha_j \theta_{ij} + \nu_p + \mu_i + \gamma_t + \epsilon_{it} \tag{1}$$

Where q_{ijt} is the number of gallons of product j (i.e. gasoline or distillate) produced at refinery i in month t, and k is installed monthly distillation capacity, denominated in gallons of crude oil processable per month. θ_{ij} is the fraction of gasoline or distillate that is reformulated or low-sulfur respectively. ν_p is an indicator for the post-1994 period interacted with PADD dummies, and μ and γ are refinery and time period fixed effects. X includes the installed upgrading capacity and crude API gravity at each refinery each period.

Three separate specifications for θ_{ij} are run for both gasoline and distillate. In the first column, θ_{ij} is the average RFG and LSD share for each refinery from 1995 to 2003 (Figures 9 and 10). A 10 percent increase in RFG or LSD shares is associated with 0.8 and 0.5 percent higher yields of gasoline and diesel per unit of capacity. The second column uses monthly variation in θ_{ij}. Within-refinery deviations in RFG and LSD shares are again associated with higher gasoline and distillate productivity. Of course, θ_{ij} is chosen by the refinery, and could be correlated with other refinery unobservables or with residual demand shocks to relative RFG and LSD demand. The third column presents IV results where I instrument for θ_{ij} using the pre-1990 market share variable from Table 4 as well as gasoline and highway diesel taxes. The distillate regression also uses heating degrees days as an instrument, which should increase relative demand for high-sulfur distillate. After instrumenting, the RFG coefficient is slightly negative, but not significant, while the LSD coefficient is still large and significant. There does not appear to be strong evidence that producing RFG or LSD reduces refinery productivity, and, in the case of LSD, appears to increase the amount of distillate obtained from a given capital stock and crude type.[13]

4 A model of refinery operations

In this section I develop a structural model of the refining industry that extends the existing literature by simultaneously incorporating joint production, capacity constraints, and imperfect competition. In order to run policy counterfactuals, I need to recover all of the parameters gov-

[13] A log-log specification, rather than ratios, yields similar results, as does including PADD-year or PADD-time dummies.

erning refinery behavior, not just those related to content regulations. I specify a multiproduct cost function and use a logit transformation to incorporate the joint nature of the refining process. Costs are not directly observed, and are instead inferred from market clearing decisions under assumptions about producer behavior, following Rosse (1970). Firms compete simultaneously in quantities in each state, and estimation is based off first order conditions which equate marginal revenue in each end market to a centralized marginal cost of production for each refinery. Identification comes from an extensive set of seasonal, temporal and geographic demand shifters, as well as changes in refinery ownership and capacity shares, which vary the infra-marginal quantity each refinery internalizes when making production decisions.

4.1 Existing literature

Three authors have previously estimated structural models of the US refining industry using aggregated data on refinery operations. Muehlegger (2006) estimates a marginal cost function for refined products, but assumes separable production with perfect substitutability. With refinery-level production data, I am able to observe behavior that is much more consistent with joint production. Despite the fact the relative prices of end products vary by plus or minus 40 percent during this period, there is a not a single observation where a refinery produces only gasoline or diesel. Chesnes (2009) estimates a dynamic model of refinery investment, but assumes that product yields are fixed across refineries and over time. As was shown in Table 2 and Figure 8, product yields vary significantly across refineries and time periods. Since one potential response for a refinery facing content regulation was to alter its production mix, it is important in this paper to explicitly incorporate this margin.

Zhang (2011) estimates a multiproduct production function, but assumes perfect competition and treats refinery operating levels as exogenous and identical across refineries in a PADD. With refinery-level data I am able to observe that the average coefficient of variation of refinery utilization rates within a PADD-month is 0.15. Moreover, variation in utilization rates is correlated with market power. A regression of monthly refinery utilization on the controlling firm's share of total capacity in the PADD along with refinery and time fixed effects returns a coefficient of -0.327 (0.149), indicating that a 10 percent increase in regional market share is associated with a 3.2 percent reduction in marginal willingness to operate (see Appendix C for further details). Again, incorporating market power is particularly important in this paper, as it is possible that markups for the new products differed as well as costs. In what follows, I specify a model that includes a multiproduct cost function, incorporates endogenous refinery-level heterogeneity in yields and utilization, and allows for costs to differ from marginal prices based on market shares.

4.2 Structural model

Firms face a constant elasticity of demand curve for each product j and end market m:

$$\ln Q_{jm}(\alpha) = \alpha_{0jm} + \alpha_j \ln p_{jm} + \epsilon^D_{jm} \tag{2}$$

Competition is assumed to take place at the wholesale level. In the 782A data, 87 percent of gasoline sales and 83 percent of distillate sales are sales for resale. At this level, products are essentially homogeneous. Although branded gasoline often contains additives which garner a price premium at the pump, these are added at the wholesale terminal by the purchasing party. When shipped, refined products of a particular type commingle, with purchasers often unaware of which refinery the product was produced at.

Products are shipped to markets from refineries i at transportation costs τ, resulting in total post-shipping revenues:

$$Rev_i(\mathbf{q_i}) = \sum_j \sum_m (p_{jm}(Q_m) - \tau_{im}) q_{ijm} \tag{3}$$

These revenues are generated at the expense of a single centralized production cost for each refinery:

$$Cost_i(\mathbf{q_i}, c_i; \beta) = \sum_j (\beta_j + \beta_{jj} \frac{q_{ij}}{c}) q_{ij} + (\beta_c + p_c) c_i + f(\frac{c_i}{k_i}) \tag{4}$$

The cost function has two components, product specific processing costs and general operating costs. β_j reflects the constant marginal cost of producing each product, while β_{jj} represents the increasing difficulty of extracting higher yields for each product. In addition to these product specific costs, for each gallon of crude c processed, the refinery pays constant marginal operating costs β_c, crude oil costs p_c, and convex utilization costs $f(\frac{c}{k})$, where k is the amount of installed distillation capacity. Refineries report capacity as the number of gallons of crude processable under "normal" operating conditions, rather than the maximum processable amount. As a result, monthly utilization rates routinely exceed 100 percent, although the distribution drops off sharply after that point (Figure 11).

Content regulations RFG (r) and LSD (l) enter the cost function by shifting the product specific marginal processing costs of gas and distillate, mc_g and mc_d,

$$mc_{ir} = \beta_g + \beta_{gg} \frac{q_{ig}}{c_i} + \beta_r = mc_{ig} + \beta_r \tag{5}$$

$$mc_{il} = \beta_d + \beta_{dd} \frac{q_{id}}{c_i} + \beta_l = mc_{id} + \beta_l \tag{6}$$

Under this specification, β_r and β_l represent the dollar per gallon cost increase of RFG and LSD relative to conventional gasoline and distillate.

As was discussed in Section 2.1, refining is an inherently joint production process. Operations are centered around "production runs", where a refinery sets the amount of each end product to extract from each gallon of crude and then decides how much crude to process. Define the yield of gasoline and distillate from a given gallon of crude as $Y^g = \frac{q_g}{c}$ and $Y^d = \frac{q_d}{c}$, and the share of each that satisfies content restrictions as $\theta_r = \frac{q_r}{q_g}$ and $\theta_l = \frac{q_l}{q_d}$. The cost function becomes,

$$\begin{aligned} Cost_i(\mathbf{Y_i}, \theta_\mathbf{i}, c_i; \beta) &= c_i\left[(\beta_g + \beta_{gg}Y_i^g + \beta_r\theta_{ir})Y_i^g + (\beta_d + \beta_{dd}Y_i^d + \beta_l\theta_{il})Y_i^d + \beta_o Y_i^o\right] \\ &\quad + c_i(\beta_c + p_c) + f(\frac{c_i}{k_i}) \end{aligned} \tag{7}$$

Where $Y^o = (1 - Y^g - Y^d)$ is the refinery's outside option from increasing gasoline and distillate yields, which includes the all other refined petroleum products, such as jet fuel, residual fuel oil, asphalt and propane.

Firms are assumed to compete simultaneously in quantities. Combining equations (7) and (3), and summing over all refineries I_f owned by firm f yields the firm's optimization program:

$$\begin{aligned} Max_{\{q_{ijm}, Y_i, \theta_i, c_i\}}\pi &= \sum_{i \in I_f} Rev_i(\mathbf{q}_i) - Cost_i(\mathbf{Y_i}, \theta_\mathbf{i}, c_i; \beta) \\ s.t. \quad \sum_m q_{ijm} &\leq Y_i^j c_i \\ \sum_j Y_i^j &= 1 \end{aligned} \tag{8}$$

Despite relaxing many assumptions made in the previous literature, the model is still limited in several respects. Vertical integration in either direction is not captured. Upstream, this implies that firms are price takers in the crude oil market. Approximately 55 percent of refining capacity in the sample is owned by independent refiners with no upstream operations. Downstream, the model assumes that each refinery's incentive to supply a market is fully captured in the wholesale price. In reality, some wholesale sales are going to parties the refiner has a contractual obligation with, or that compete against its downstream arm.[14] However, information on these relationships is not available for all products or for the entire sample. Finally, demand is assumed to be static, ignoring the demand side's ability to smooth purchases across time through storage and refineries' consideration of inventories when setting quantities each period (Borenstein and Shepard 2002).

4.3 Estimation

Estimating equations are based on the first order conditions of (8). Two assumptions are made to simplify the problem, which contains hundreds of first order conditions. First, for all products, the Karush-Kuhn-Tucker conditions imply that marginal revenue, net of shipping costs, must be equal in expectation for all markets served by each refinery.

[14] Gilbert and Hastings (2005) find evidence that vertical integration had a significant impact on wholesale gasoline prices following Tosco's acquisition of Unocal's West Coast assets in 1997.

$$mr_{ijm} = p_{jm}(1 + \frac{s_{fjm}}{\alpha_j}) - \tau_{im} = \frac{\partial Cost_i(\mathbf{Y_i}, \theta_\mathbf{i}, c_i; \beta)}{\partial q_{ij}} \quad \forall q_{ijm} > 0 \tag{9}$$

Where s_{fjm} is firm f's market share and α_j is the demand elasticity parameter from (2). Let mr_{ij} be the expected marginal revenue of each product across all markets served by refinery i, assuming optimal allocation.

Second, rather than working in fractions with constraints, I use the logit transformation to convert yield choices into a continuous unbounded state space:

$$Y^g = \frac{e^{\delta_g}}{1 + e^{\delta_g} + e^{\delta_d}} \quad ; \quad Y^d = \frac{e^{\delta_d}}{1 + e^{\delta_g} + e^{\delta_d}} \tag{10}$$

Conceptually, δ_j can be thought of as encompassing all of the effort and resources a refinery allocates towards producing product j. This specification also incorporates the inherent multiproduct nature of the process, by imposing the logit assumption that the effect of increased effort towards product j is proportional to the yields of products j and k,

$$\frac{\partial Y^j}{\partial \delta_j} = Y^j(1 - Y^j) \quad ; \quad \frac{\partial Y^k}{\partial \delta_j} = -Y^k Y^j$$

Incorporating these assumptions reduces the refiner's problem to five choices variables for each month: how much refining effort to direct towards gasoline and diesel (δ_g, δ_d), the share of each to convert into RFG and LSD (θ_r, θ_l), and how many gallons of crude to process (c). Each decision variable is assumed to have an associated private cost shock ϵ^S known to the refinery at the time of production but unobserved to the econometrician. Under this formulation, the cost function becomes,

$$\begin{aligned} Cost_i(\delta_\mathbf{i}, \theta_\mathbf{i}, \mathbf{c_i}; \beta) \quad = \quad & c_i \Big[(\beta_g + \beta_{gg} Y_i^g + \beta_r \theta_{ir}) Y_i^g + (\beta_d + \beta_{dd} Y_i^d + \beta_l \theta_{il}) Y_i^d + \beta_o Y_i^o\Big] \\ & + c_i(\beta_c + P_c) + f(\frac{c_i}{k_i}) + c_i \epsilon_{ic}^S + k_i(\delta_{ig}\epsilon_{ig}^S + \delta_{id}\epsilon_{id}^S + \theta_{ir}\epsilon_{ir}^S + \theta_{il}\epsilon_{il}^S) \end{aligned} \tag{11}$$

Each of the product specific cost shocks is assumed to enter additively and scale with capacity across refineries. The interpretation here is that these shocks pertain to configuring the refinery to generate a specific output mix each production run, but do not directly affect marginal crude input decisions conditional on yields. Fixed costs are not identified in this model, and assumed to be zero.[15] Differentiating with respect to each choice variable yields five first order equations to be estimated simultaneously:

[15] Unlike electric power generators, refineries always operate except for scheduled maintenance (every 3-5 years) or an unscheduled disruption, such as a fire, which prohibits operation.

$$\text{(FOC1)} \quad \frac{d\pi}{d\delta_{ig}} = \frac{c_i}{k_i}\Big[[mr_{ig}(1-\theta_{ir}) + mr_{ir}\theta_{ir} - mr_{io}]Y^g_{ig} + [mr_{id}(1-\theta_{il}) + mr_{il}\theta_{il} - mr_{io}]Y^d_{ig}\Big]$$
$$-\frac{c_i}{k_i}\Big[(\beta_g + \beta_{gg}Y^g_i + \beta_r\theta_{ir})Y^g_{ig} + (\beta_d + \beta_{dd}Y^d_i + \beta_l\theta_{il})Y^d_{ig}\Big] - \epsilon^S_{ig} = 0$$
$$\text{(FOC2)} \quad \frac{d\pi}{d\delta_{id}} = \frac{c_i}{k_i}\Big[[mr_{ig}(1-\theta_{ir}) + mr_{ir}\theta_{ir} - mr_{io}]Y^g_{id} + [mr_{id}(1-\theta_{il}) + mr_{il}\theta_{il} - mr_{o}]Y^d_{id}\Big]$$
$$-\frac{c_i}{k_i}\Big[(\beta_g + \beta_{gg}Y^g_i + \beta_r\theta_{ir})Y^g_{id} + (\beta_d + \beta_{dd}Y^d_i + \beta_l\theta_{il})Y^d_{id}\Big] - \epsilon^S_{id} = 0$$
$$\text{(FOC3)} \quad \frac{d\pi}{dc_i} = [mr_{ig}(1-\theta_{ir}) + mr_{ir}\theta_{ir} - mr_{io}]Y^g_i + [mr_{id}(1-\theta_{il}) + mr_{il}\theta_{il} - mr_{io}]Y^d_i + mr_{io}$$
$$-(\beta_g + \beta_{gg}Y^g_i + \beta_r\theta_{ir})Y^g_i - (\beta_d + \beta_{dd}Y^d_i + \beta_l\theta_{il})Y^d_i$$
$$-\beta_c - p_c - f'(\frac{c_i}{k_i}) - \epsilon^S_{ic} = 0$$
$$\text{(FOC4)} \quad \frac{d\pi}{d\theta_{ir}} = \frac{c_i}{k_i}Y^g_i\left[(mr_{ir} - mr_{ig}) - \beta_r\right] - \epsilon^S_{ir} = 0$$
$$\text{(FOC5)} \quad \frac{d\pi}{d\theta_{il}} = \frac{c_i}{k_i}Y^d_i\left[(mr_{il} - mr_{id}) - \beta_l\right] - \epsilon^S_{il} = 0$$

Where $Y^j_k = \frac{\partial Y^j}{\partial \delta_k}$, and $f'(\frac{c}{k})$ is the marginal utilization cost per gallon of crude processed.

Figure 12 provides a graphical representation of the firm's problem. The top panel represents the yield choice facing the refiner. When choosing δ_g and δ_d, the marginal revenue from increasing yields at a given crude level is equated to the marginal revenue of the outside option. Content regulations shift the intercept of the marginal net revenue curves for gasoline and distillate by $(P_r - P_g - \beta_r)$ and $(P_l - P_d - \beta_l)$ respectively. The slope of the curves will also change to the extent that market shares in regulated markets differ from unregulated markets. The shaded gray area reflects the total marginal revenue gained at optimal yields Y^* for a given level of crude inputs c. The second panel represents the refiner's decision of how much crude to process, where $mr(Y^*, c)$ indexes this optimal marginal revenue for every possible level of operation. Refineries face increasing utilization costs as they approach capacity, while the price of crude oil and other constant marginal costs shift the point at which these utilization costs intersect marginal operating revenue.

All of the marginal cost intercepts are modeled as linear functions of refinery characteristics, $\beta_j = \sum_i \beta_{jk}X_{jki}$. For both gas and diesel, X_j includes a constant, API gravity and API gravity squared, crude oil sulfur content, the ratio of upgrading capacity to distillation capacity, and the interaction between API gravity and upgrading capacity. X_g also includes indicators for summer months (May - September) post-1989 and 1992 to capture summer gasoline RVP restrictions. β_o includes the share of the outside option that is jet fuel, which is the most valuable product not explicitly modeled here. All other products in the outside option are assumed to trade at the price of residual fuel oil, which is a benchmark bottom of the barrel petroleum product tracked in the EIA data. Operating costs β_c are assumed to be zero, beyond refinery and time fixed effects, and marginal utilization costs are modeled as a cubic B-spline with knots placed at quartiles of the utilization distribution.

LSD costs are modeled as follows,

$$\begin{aligned}\beta_{li} &= \beta_{l0} + \beta_{l1} PctUpgrading + \beta_{l2} PctDesulf + \beta_{l3} PctDesulf^2 + \beta_{l4} Sulfur \\ &\quad + \beta_{l5} PctDesulf * Sulfur + \beta_{l6} API + \beta_{l7} PctDesulf * API + \beta_{l8} CA + \beta_{l9} SmallRefinery\end{aligned} \tag{12}$$

Where $PctUpgrading$ and $PctDesulf$ are the ratios of installed upgrading and desulfurization capacity to total distillation capacity, $Sulfur$ is the sulfur content of crude oil, CA is an indicator for the state of California, which imposed stricter diesel limits, and $SmallRefinery$ is an indicator for refineries eligible to receive SO2 credits for producing LSD.

As was discussed above, unlike LSD, RFG capability was not a function of any observable refinery technology. I therefore estimate refinery-specific RFG costs,

$$\begin{aligned}\beta_{ri} &= \beta_{r0i} + \beta_{r1} RFG1Summer + \beta_{r2} RFG2 + \beta_{r3} RFG2Summer \\ &\quad + \beta_{r4} CARB + \beta_{r5} CARBSummer + \beta_{r6} MTBE\end{aligned} \tag{13}$$

Where β_{ri} is a refinery dummy for all refineries with positive RFG production, $RFG1$ and $RFG2$ indicate the phases of the RFG program, and $Summer$ indicates summer months, which involved added restrictions. $CARB$ indicates reformulated gasoline sold in California or Arizona after March 1996, and $MTBE$ is an indicator for whether the refinery was able to use MTBE as an oxygenate in making RFG. For the early years of the program, MTBE was the preferred mode of satisfying the oxygenate requirements of RFG. Beginning in the late 1990s, there was increasing public concern that MTBE was in fact carcinogenic. In response to this, a number of states banned MTBE between 2000 and 2006, at which point a federal ban was enacted (Anderson and Elzinga 2012).

I estimate four demand functions from equation (2): gasoline demand, pre-1993 distillate, post-1994 high sulfur distillate and low sulfur diesel.[16] Each demand equation contains state-month dummies to account for variations in seasonality and time dummies. These four equations are estimated jointly with the supply side first order conditions, resulting nine equations and errors $\epsilon = (\epsilon^D, \epsilon^S)$.

Estimation proceeds via 2-stage GMM by jointly minimizing $E(\epsilon' Z)$, where $Z = \{Z_D, Z_S\}$ is a set of instruments that are assumed to be uncorrelated with demand and supply errors respectively (Hansen 1982). Z_D includes regional crude prices, regional refinery capacity concentration, and pipeline outages, which should be correlated with prices but unrelated to demand shocks. Z_S includes end-market population, weather, and fuel taxes, all of which shift demand but should not alter refinery production costs. Z_S also includes firm level capacity share, the number of refineries operating in each region, and regional capacity concentration, which vary considerably during the sample and shift the residual demand curve facing each refinery. Finally,

[16] Although distillate sales are broken out into diesel and fuel oil prior to 1993, Marion and Muehlegger (2008) show demand for these two products was jointly determined, as diesel consumers sought to evade taxes by purchasing untaxed distillate intended for off-highway use. Concurrent with the introduction of low sulfur diesel, the government mandated that non-highway distillate be marked with a dye to prevent illegal sales.

Z_S also includes month dummies. Both gasoline and distillate exhibit considerable seasonality. Figure 13 shows that while gasoline demand increases in the summer, distillate demand is relatively higher in the winter. Figure 14 breaks distillate sales down further to reveal that highway diesel demand is actually slightly higher in the summer, while demand for high sulfur distillate, which is largely used for heating oil, is over twice as high in the winter compared to the summer. This seasonal variation in relative demand helps pin down the convexity associated increasing gasoline and distillate yields, as well as the costs of converting distillate into low-sulfur diesel.

All parameters enter the supply equations linearly, and $\epsilon_g^S, \epsilon_d^S$, and ϵ_c^S each contain refinery and year dummies. These equations therefore use two sources of variation to identify the costs of content regulation. First, with refinery and year fixed effects, (FOC1-FOC3) compare refineries' willingness to supply gas and distillate, and to operate, before and after content regulation. As was discussed in section 2.2, refineries were differentially exposed to content regulation based on their proximity to regulated markets. In addition to this cross-period comparison, equations (FOC4-FOC5) compare refineries' willingness to tradeoff between regulated and unregulated products within a given month.

5 Results

This section presents the main results from the paper. Cost function estimation returns intuitive coefficients on crude quality and refinery technology, and significant heterogeneity in the marginal cost of producing RFG and LSD across refineries. Wholesale demand estimates are more elastic than commonly studied end-use petroleum product demand elasticities, and the precision of these estimates increases substantially when the supply side and demand side are estimated simultaneously. In order to recover the industry equilibrium price and quantity effects of content regulation, I simulate counterfactual market outcomes for the entire United States with fuel content restrictions removed. Under this counterfactual, prices in RFG and LSD markets are 6 and 3 percent lower on average, and prices in conventional gasoline and distillate markets are 2 and 1 percent higher. Although markups in RFG markets are higher with the restrictions, refineries experience decreased profits on average due to decreased markups and volumes on other products.

5.1 Cost function estimates

Cost function estimates are presented in Table 6. Model 1 uses wholesale prices in mr_{ij} for all sales, and Model 2 uses a weighted average of wholesale prices and sales through company owned retail outlets for each firm. As expected, using a higher quality crude (one with a high API gravity and low sulfur content) is associated with lower costs of gasoline and distillate production. Costs are also reduced by having more upgrading capacity, but the interaction of these two is positive, reflecting the fact that higher quality crude needs less processing. The slope coefficient on gasoline is slightly higher than for distillate, with both positive and significant, reflecting the fact that it is costly to increase product yields conditional on capital and crude quality. Jet fuel is estimated to cost 37 cents per gallon more than the outside option. Utilization costs are estimated to be essentially flat over the most of the distribution, increasing sharply once a refinery's crude inputs exceeds 100 percent of installed capacity. In Model 2, retail sales of gasoline and distillate have an effect equivalent to increasing marginal costs by 16 and 4 cents per gallon respectively, although this includes both retailing costs and the average marginal profit from retailing.

Table 7 presents the estimated costs of RFG and LSD in detail, along with projected costs from the EPA and the National Petroleum Council.[17] The average intercept β_{r0i} across the 54 refineries which produced RFG during the sample is 12.3 and 9.3 cents per gallon in the models excluding and including retail prices. Phase I summer restrictions are not estimated to significantly increase RFG costs beyond nationwide summer RVP limits. Costs are actually estimated to be slightly lower during Phase II, although the summer component of Phase II, which was the most stringent addition, is large and significant. Summer restrictions in California increase costs by a similar amount. The inability to use MTBE is found to substantially increase RFG costs by 3.6 to 3.8 cents per gallon. In combination, the results indicate that RFG was 7.1 to 9.6 cents per gallon more expensive than conventional gasoline on average. When weighted by quantity of RFG produced, the implied increase in refining costs is 6.2 to 8.0 cents per gallon. All but the unweighted Model 1 estimates fall within the range predicted by EPA, and below that predicted by the NPC. Figure 15 plots the average RFG cost for each refinery with positive RFG production in the sample. Each point is a refinery, and they are sorted by cost. The cost estimates vary considerably across refineries, ranging from -2.7 cents per gallon all the way up to 23 cents per gallon, 16 cents higher than the mean.

Turning to diesel, for a large refinery with no desulfurization capacity, LSD is estimated to cost 7.2 and 6.5 cents per gallon more than high-sulfur distillate depending on whether or not retail is included. At average desulfurization capacity levels, costs are reduced considerably to 2.9 and 3.3 cents per gallon respectively. While the intercept on desulfurization capacity is positive, the interaction with crude sulfur content and API gravity is negative, indicating that the technology reduced costs primarily by allowing refineries to process light high-sulfur crudes into low sulfur diesel. California's aromatics restriction was estimated to add an additional 2.4 cents per gallon to LSD costs in both models. Although small refineries gained valuable SO2 permits for making LSD, their costs are still estimated to be 1.1 to 1.5 cents higher than larger refineries. Presumably actual costs at these facilities are much higher. When weighted by quantity of LSD produced, the average compliance cost during the sample is estimated to be a modest 2.2 to 2.6 cents per gallon, slightly lower than estimated by EPA, and less than half what was predicted by the NPC. One reason variable costs were so low was because refineries invested heavily in desulfurization capacity after the 1990 Clean Air Act Amendments were passed. The costs of these investments are estimated in Sweeney (2015). . Figure 16 plots the average LSD cost for each refinery, sorted by cost. Similar to RFG, costs vary considerably

[17] The NPC is a federal advisory committee to the Secretary of Energy made up of petroleum industry executives. The purpose of the NPC is to advise and make recommendations to the Secretary of Energy (FTC 2006).

across refineries, ranging from from 0 to 10 cents per gallon, or 7 cents higher than the mean.

5.2 Demand estimates

Demand elasticity results are presented in Table 8. For comparison purposes, Equation (2) was estimated independent of the supply side (i.e. "offline") as well as jointly. Offline OLS results are presented first, and reveal wholesale demand to be fairly inelastic, particularly for gasoline. There is reason to believe OLS price coefficients will be biased towards zero for petroleum products (Davis and Killian 2011). However, as is well documented in the literature, it is difficult to come up with valid instruments for price at the state level that have sufficient power. The third row of the table presents offline IV estimates, where regional domestic crude prices, refinery concentration, and pipeline outages were used as instruments. Only the gasoline regressions have sufficient power in the first stage, possibly due to the larger sample. For three of the four products, the IV point estimates are more elastic than OLS, but the standard errors increase dramatically, meaning that only the gasoline elasticity is statistically distinguishable from zero. The final row of the table presents the results from jointly estimating the demand and the supply sides. Demand is estimated to be much more elastic than the offline results would suggest, although the joint estimates all lie within the 95 percent confidence intervals of the IV estimates.[18] High sulfur distillate demand post regulation is estimated to be the most elastic and gasoline the least.

For each product, the estimated wholesale price elasticities are substantially larger than more commonly reported end-use demand elasticities. For example, in a recent paper, Li, Linn, and Muehlegger (2014) estimate state-level end-use gasoline demand elasticities ranging from 0.109 to 0.365. However, if that were the relevant demand curve facing refineries, it would be difficult to reconcile such inelastic demand with observed markups. In the EIA-782 data, state-level market shares of 30 percent are not uncommon. At that level, even a relatively high end-use elasticity of 0.3 would imply Lerner markups of 100 percent above marginal costs. Yet wholesale gasoline prices are only 40 percent higher than WTI crude spot prices on average during this period. Thus, even ignoring processing costs and the fact that gasoline yields per gallon of crude are less than one, markups, at least in a quantity setting model, appear much too low to rationalize commonly reported end-use elasticities. Instead, the estimates in Table 8 suggest that the demand curve facing refineries, which is comprised of logistics firms and large intermediaries capable of storage, is fairly elastic.

5.3 Counterfactuals

Having recovered the cost structure of the industry and the additional costs of producing RFG and LSD, I simulated counterfactual market outcomes with these policies removed.[19] Simula-

[18] Other industrial organization papers have found supply side moments helpful in pinning down demand elasticities. For example, Berry, Levinsohn, and Pakes (1995).

[19] Simulations carried out in AMPL using the PATH complementarity problem solver of Steven Dirske, Michael C. Ferris and Todd Munson, available at http://pages.cs.wisc.edu/~ferris/path.html.

tion of the full model would involve solving simultaneous equilibrium in every state for every product, which was not computationally feasible. I therefore aggregated end-markets into 9 regional markets broadly reflective of pipeline patterns and refinery concentration. A map of the simulation regions is provided in Figure 15. Moving from state-level to region-level demand results in a minor loss of fidelity between the baseline simulated prices and observed prices in sample. The difference between the two series is less than 2 percent on average, with an R^2 of 0.983. All counterfactual results therefore compare simulated counterfactual results against simulated baseline results.[20]

Table 9 presents the price results. All differences presented are the baseline outcomes under the regulation minus the outcomes with the policies removed. Gasoline prices in RFG areas are 6 to 8.2 cents per gallon higher than they would have been if those regions could have purchased conventional gasoline. However, from a national perspective, these price increases are partially offset by decreased prices in conventional gasoline markets, as refineries that found it costly to produce RFG reallocated supply towards unregulated markets. This reallocation had the largest price effect on the West Coast, where RFG made up over 80 percent of total gasoline demand and shipment to other conventional markets was not economical. Overall, the net effect on US gasoline consumers is $14 billion over the 9 year sample.

Turning to distillate markets, price changes in highway diesel markets are pretty uniform across the country, increasing by about 2.3 to 3.4 cents per gallon. The largest prices increases were on the West Coast, where California required a more costly form of highway diesel. Similar to RFG, LSD restrictions drive down the price of high sulfur distillate, partially offsetting the total cost of the policy. In combination, the net effect on distillate consumers was $6 billion in lost consumption.

The last column reports the average increase in marginal costs at refineries supplying each RFG and LSD market. In each region, RFG market prices increase by more than costs, with average markup increases ranging from 0.3 to 2.5 cents per gallon. For LSD, costs increase by more than prices for 6 of the 9 markets, indicating that LSD reduced distillate margins on average. Table 10 presents the combined impact on refineries by PADD. Profits decline by 0.45 cents per gallon on average. Refineries in the Gulf Coast are least affected, having relatively low costs of producing RFG and LSD and access to the most end markets. PADD 4 refineries experience the biggest decline in profits, as these refineries are the least sophisticated and also serve the Northwest, which saw relatively larger declines in conventional gasoline prices. In sum, profits were $10 billion lower than they would have been without content regulation, an 8 percent reduction.

For RFG, the estimated price increases can be compared to previous estimates from the literature.[21] Brown et al. estimate RFG price effects using an unweighted average of unbranded weekly city-level wholesale gasoline prices from 1994 to 1998. Their estimates range from -1.7 to

[20]Multiple equilibria may be possible. In order to check for this, I randomly selected 50 baseline periods and simulated equilibria for each from 10 separate starting points. In every case all 10 runs converged to the same point, suggesting that multiplicity of equilibria is not a concern in this setting.

[21]As far as I am aware, no other studies have estimated the price impact of low sulfur diesel.

10.1 cents per gallon across cities, with a mean of 4.1.[22] The estimated price increase for RFG regions in this paper during the same period ranges from 5.7 to 7.6 cents per gallon. Although these estimates falls within the range estimated by Brown, the mean is slightly higher.[23] Using state level data and a three stage least squares approach that instruments for the size of each state's RFG market, Chakravorty, Nauges, and Thomas (2007) find that moving from zero to 100 percent RFG would increase a state's wholesale gasoline prices by 15.1 percent (16 cents per gallon on average).

5.4 Discussion

In Table 5, reduced form productivity regressions showed that, if anything, RFG and LSD were associated with increases in gasoline and distillate output per unit of capacity during the sample. However, by explicitly accounting for changes in operating incentives across time, structural estimation recovered large statistically significant marginal cost increases associated with producing these products. Simulation sheds additional light on what's going on here, revealing that refineries which found it costly to produce RFG and LSD were not only excluded from these markets, but also experienced lower margins in conventional markets.

The policy counterfactuals also highlight the importance of accounting for interactions between regulated and unregulated markets when estimating the costs of environmental regulation. Previous reduced-form studies of the Clean Air Act have implicitly calculated the gross effect of regulation on regulated versus unregulated areas or firms, leaving the net impact ambiguous. One such paper, Greenstone (2002), summarizes the importance of this distinction,

> It would be informative if the estimated regulation effects could be used to determine how much production (and employment) was shifted abroad as a result of the non-attainment designations. This would provide one measure of the national costs of these regulations. Unfortunately, such a calculation is not possible because it cannot be determined whether the lost activity in non-attainment counties moved to foreign countries or attainment counties. Since it is likely that the regulation effects partially reflect some shifting of manufacturing activity within the United States, they probably overstate the national loss of activity due to the non-attainment designations. Moreover, the possibility of intra-country shifting means that the regulation effects are also likely to overstate losses in non-attainment counties. The reason is that the identification strategy relies on comparisons between non-attainment and

[22]Estimates, taken from Table 8 and average from Table 5 in Brown et al. 2008. Results in that paper were reported in nominal dollars, and are converted here to 2013 dollars to match the estimates in this paper.

[23]There are several possible reasons for this. With city level data, Brown et al. are able to pick up within-state variation that is not captured in the EIA data. To the extent that RFG areas within a state had higher prices to begin with, state-level data could overstate the price change associated with the shift to RFG. Another difference is that their price data is one step further downstream from the demand curve estimated in this paper. Their data includes resellers and marketers, as well as refineries. Finally, Brown et al. do not have quantity information, and therefore use a straight average. If sellers with low volumes also changed their prices less, then this stickiness would bias estimates of the average price effects downwards.

attainment counties, which leads to 'double counting' when production is moved from a non-attainment county to an attainment one.

In this paper, I find that if reformulated gasoline and low sulfur diesel restrictions were removed, approximately one third of the consumption surplus gained by consumers in those markets would be offset by decreased surplus in conventional gasoline and distillate markets.

Although the focus of this paper has been on the costs of fuel content regulation, it is important to relate them to the benefits of these programs. EPA claims that RFG Phase I and II reduced smog-causing pollution and toxics by 17 percent and 22 percent respectively compared to conventional gasoline, resulting in emissions reductions of 64,000 and 105,000 tons per year (EPA 1999). However, subsequent study has called into question whether any ozone benefits actually materialized. Auffhammer and Kellogg (2011) use detailed pollution monitor data to study changes in ozone concentrations around the programmatic and seasonal fuel restriction changes. They find no evidence that RVP regulations or federal reformulated gasoline improved ground level ozone. They do, however, find economically and statistically significant improvements in California, where stricter limits were placed on how refineries could comply with the regulation.

As far as I am aware there has been no retrospective empirical study of the benefits of the low sulfur diesel program specifically. Assuming that the pre-1990 diesel sulfur content levels would have persisted in the absence of this policy, the move to LSD represents an 80 percent reduction in sulfur emissions from diesel, which would correspond to a 1 to 2 percent reduction in national SO2 emissions during this period. It is estimated that approximately 12 percent of urban sulfur dioxide emissions are converted to particulate matter (PM) in the atmosphere. Retrospective reviews of the costs and benefits of the Clean Air Act by EPA found that benefits exceeded the costs by an order of magnitude, with most of the benefits attributed to reductions in premature mortality due to reductions in ambient PM (EPA 2011). Low sulfur diesel was implemented concurrent with several other programs targeting PM, as well as regulations on heavy duty truck engines, and a full calculation of the additional gains from LSD specifically is beyond the scope of this paper. However, the relatively modest costs of low sulfur diesel estimated here suggest that if even a small percentage of realized PM reductions during this time frame are attributable to this program, then those benefits alone would far outweigh the costs for this policy.

6 Conclusion

I estimate the impact of fuel content regulations imposed by the 1990 Clean Air Act Amendments on the US oil refining industry. In doing so, I account for the existence of spillovers between regulated and unregulated markets and imperfect competition. I find that reformulated gasoline increased production costs by 7 cents per gallon on average, and that low sulfur diesel increased costs by 3 cents. These costs varied significantly across refineries, resulting in gains for some refineries, particularly those able to cheaply produce RFG. However, I find that the demand curve facing refineries is significantly more elastic than would be gathered from looking at end-use consumer price responsiveness alone. As a result, operating profits were 8 percent lower from 1995 to 2003 than they would have been absent the policy[24]. These cost increases translated to consumer surplus losses of $2.85 billion and $884 million per year in RFG and LSD markets. However, these losses were partially offset by consumer surplus gains of $1.28 billion and $240 million in non-regulated markets as refineries reallocated production.

This paper has only considered reformulated gasoline and low sulfur diesel standards as they were actually implemented. However, the wide heterogeneity in refinery productivity and compliance costs estimated here suggest that, when feasible, incentive-based regulation could substantially increase cost effectiveness in this industry. Although RFG is probably not conducive to a market based approach, such as cap and trade, given the steep damages associated with local ozone concentrations, subsequent sulfur regulations and pending carbon dioxide regulations appear ideal candidates.

Finally, this paper has primarily focused on the static costs of content regulation, taking capacity investments, mergers and acquisitions as given. Yet the relatively large profit impacts estimated above are likely to have also affected investment decisions and potentially even expedited closures. Understanding the dynamic implications of fuel content regulation, such as its impact on long run industry concentration and point source refinery emissions, is an interesting topic for future research.

[24]This decline in in profits only includes operating profit changes. The fixed costs associated with the low sulfur diesel program are estimated in Sweeney (2015).

7 Tables and Figures

Figure 1: Petroleum Product Supply Chain

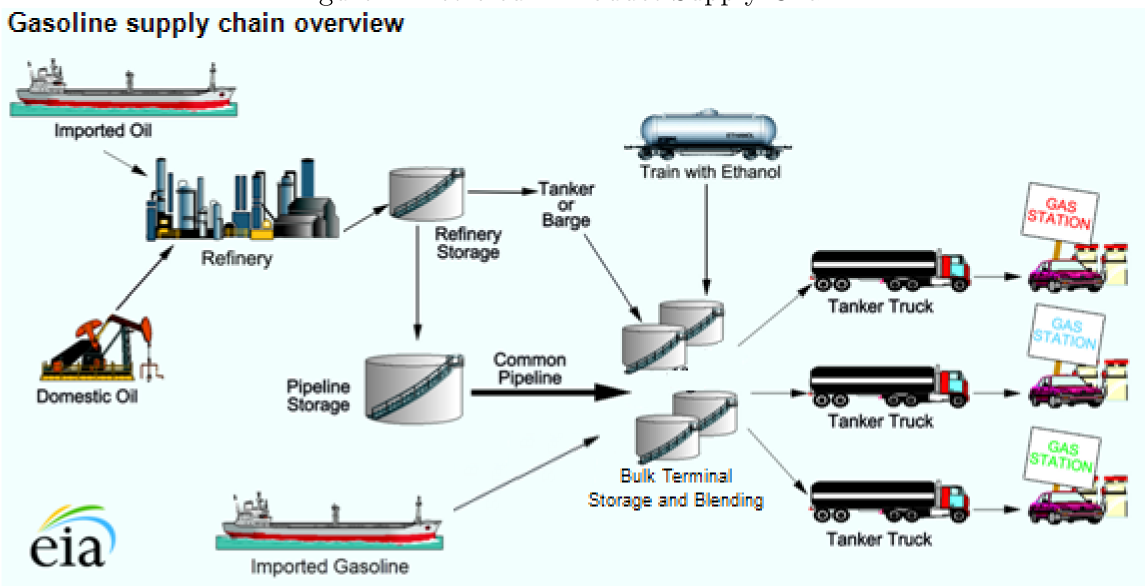

Figure 2: Modern Refinery Configuration

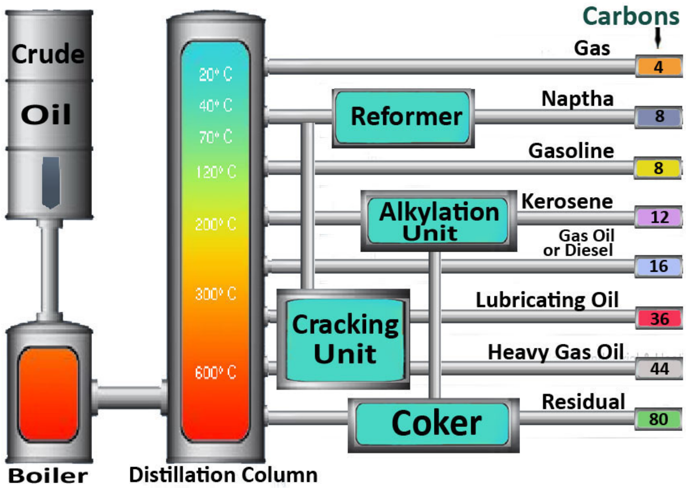

Figure 3: Refinery and Pipeline Locations

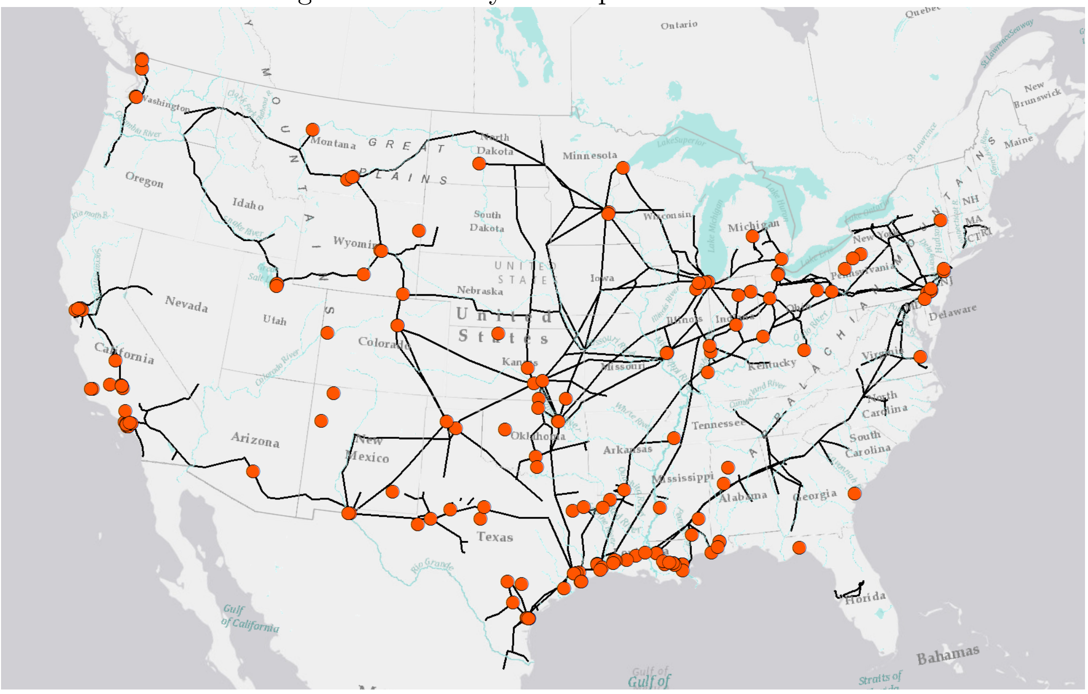

Table 1: **Summary of Environmental Regulation**

Regulation	**Description**	**Dates**
RVP Limits	Limits on gasoline in summer months (May 1 - September 15)	Phase 1: 1989 - 1991 Phase 2: 1992 - present
Oxygenated Gasoline	Required oxygenate be blended into gasoline in severe CO non-attainment counties (November - February)	1992 - present
Federal RFG	Content and performance limits on gasoline in severe ozone non-attainment counties	Phase 1: 1995 - 2000 Phase 2: 2000 - present
CARB Gas	RFG with additional restrictions for CA	Phase 1: 1992 Phase 2: 3/96 Phase 3: 4/03
Low Sulfur Diesel	Highway diesel capped at 500 (ppm)	10/93 - 6/06
Tier 2 Low Sulfur Gas	Average gasoline sulfur content set to 30 ppm	Phased in: 2004 - 2006
Ultra Low Sulfur Diesel	Highway diesel capped at 15 (ppm)	Phase-in: 6/06 - 5/10 Binding 6/10

Figure 4: Counties Requiring Reformulated Gasoline

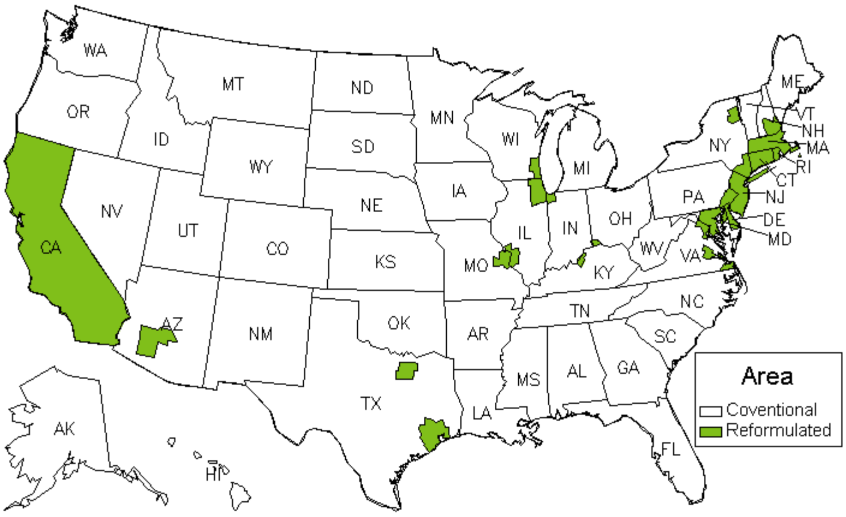

Figure 5: Share of Post-94 Distillate Sales that are Low-Sulfur

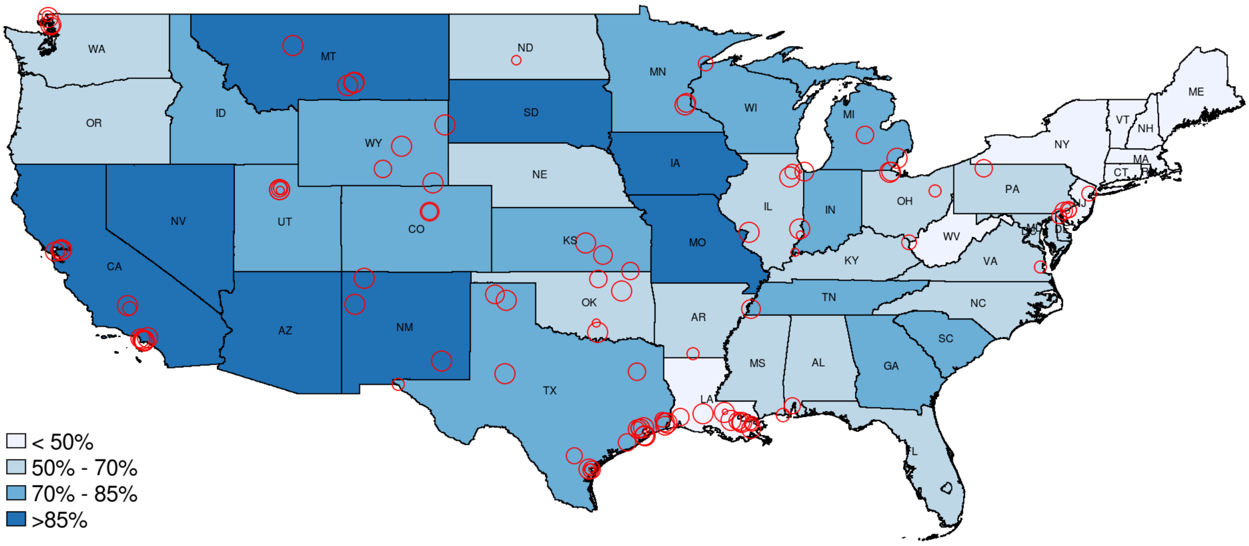

NOTE: Circles proportional to the average fraction of distillate that is low-sulfur for each refinery

Table 2: Summary Stats By Region

	East Coast (PADD 1)		Midwest (PADD 2)		Gulf Coast (PADD 3)		Rockies (PADD 4)		West Coast (PADD 5)	
# Refineries										
1986	11		30		47		16		20	
1995	9		29		44		15		18	
2003	9		24		40		15		16	
Average Yields (%)										
Gasoline	42.4	(7.9)	51.1	(6.)	45.0	(7.8)	44.5	(8.1)	43.0	(12)
Distillate	24.5	(4.7)	23.6	(6.9)	22.8	(6.8)	26.2	(4.9)	17.9	(7.1)
Jet Fuel	5.2	(4.6)	4.6	(4.5)	7.3	(5.5)	4.7	(4.2)	9.4	(7.2)
Capacity (KBbl/cd)	133	(69)	113	(83)	151	(119)	35	(16)	121	(76)
Upgrading Capacity (%)	50.8	(25)	49.7	(12)	58.0	(40)	42.6	(15)	67.2	(35)
API Gravity (degrees)	33.7	(5.1)	35.3	(4.6)	34.4	(6.)	35.5	(5.6)	26.2	(5.)
Crude Sulfur (%)	0.8	(0.7)	0.9	(0.6)	1.1	(0.8)	0.9	(0.8)	1.1	(0.3)
# Producing RFG	8		9		24		0		13	
# Producing LSD	9		25		39		15		16	
Wholesale Gasoline Prices										
Pre-1993 Conventional	1.25	(.15)	1.18	(.16)	1.13	(.15)	1.19	(.17)	1.26	(.16)
Post-1994 Conventional	1.01	(.2)	1.02	(.22)	0.96	(.2)	1.11	(.2)	1.15	(.21)
Post-1994 Reformulated	1.13	(.21)	1.17	(.22)	1.03	(.2)	-	-	1.23	(.23)
Wholesale Distillate Prices										
Pre-1993 High-Sulfur	1.06	(.18)	1.06	(.17)	1.01	(.17)	1.12	(.19)	1.05	(.2)
Post-1994 High-Sulfur	0.90	(.21)	0.93	(.2)	0.85	(.2)	1.03	(.21)	0.95	(.22)
Post-1994 Low-Sulfur	0.92	(.21)	0.95	(.21)	0.89	(.2)	1.03	(.21)	1.02	(.21)

Notes: Sample restricted to refineries with more than 10KBbl/cd and at least 50% light yield. All prices in real (2013) dollars per gallon.

Figure 6: Refinery Size and Exits by PADD

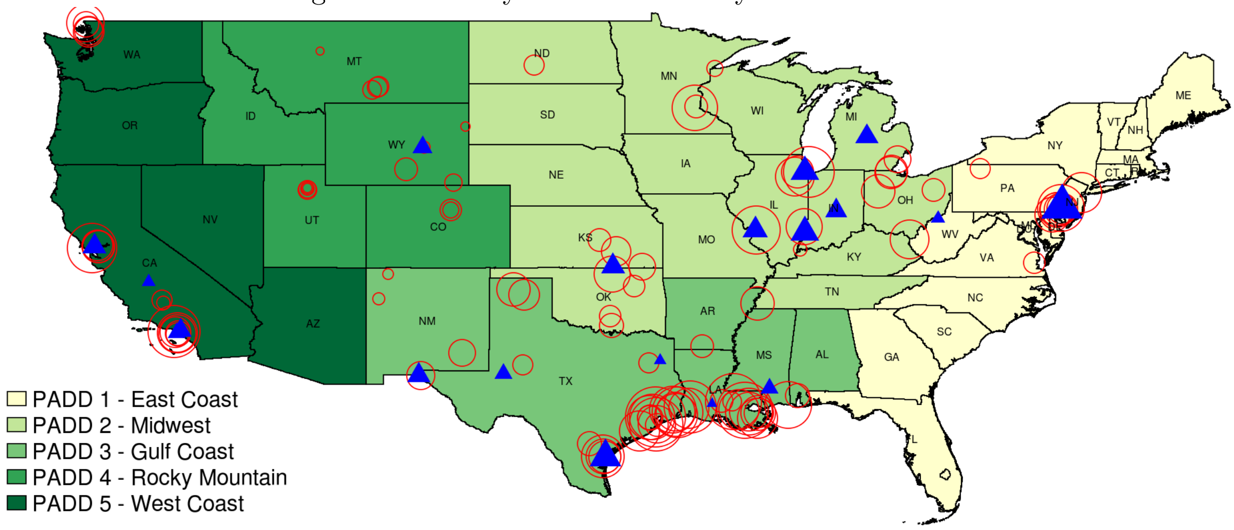

NOTE: Blue triangles indicate refineries which exited by 2003. Points are proportional to ending distillation capacity.

Figure 7: Distillation Capacity Changes

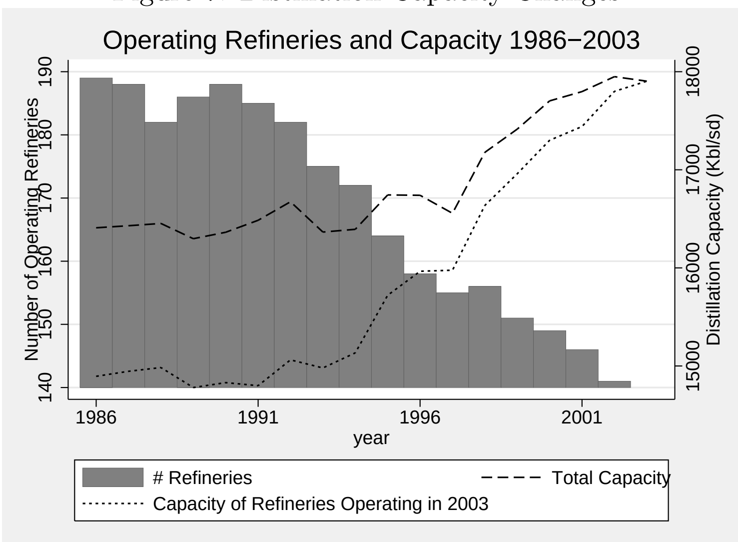

Table 3: Operable Refineries and Sales 1986 - 2003

Year	Operable Refineries	Refinery Sales
1986	216	15
1987	219	9
1988	213	21
1989	204	5
1990	205	4
1991	202	9
1992	199	2
1993	187	2
1994	179	4
1995	175	6
1996	169	7
1997	164	19
1998	160	40
1999	159	11
2000	158	14
2001	155	29
2002	153	6
2003	149	8

Notes: A detailed genealogy of US refiners can be found at `http://www.eia.gov/finance/genealogy/`.

Figure 8: Price and Quantity Trends

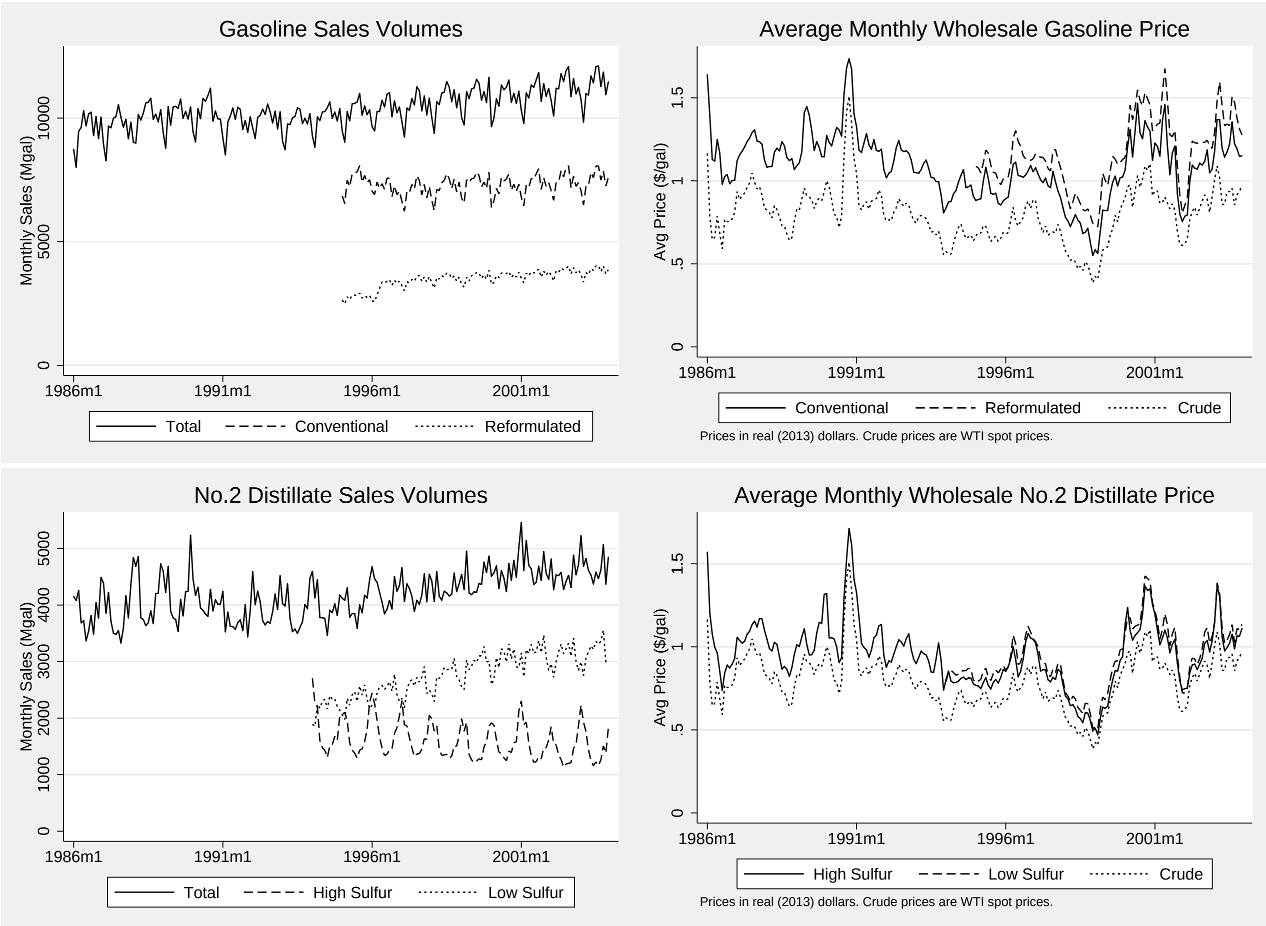

Figure 9: Average RFG Share by Refinery

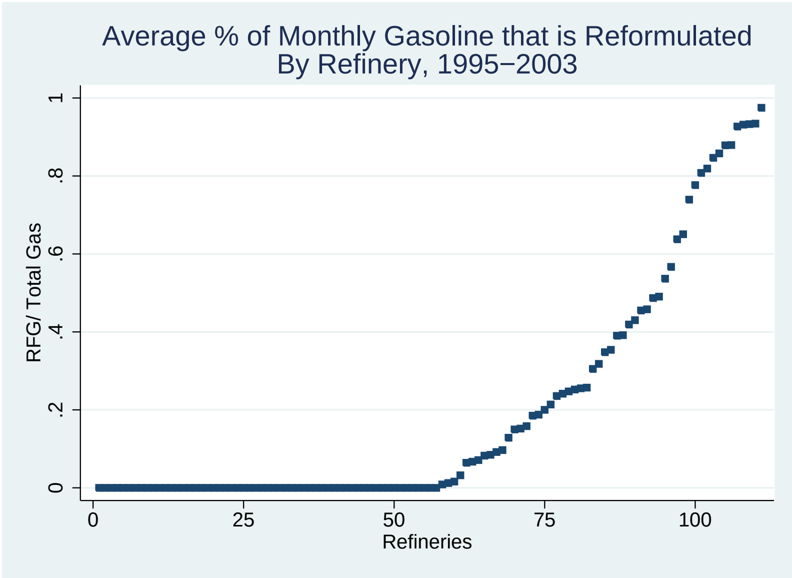

Figure 10: Average LSD Share by Refinery

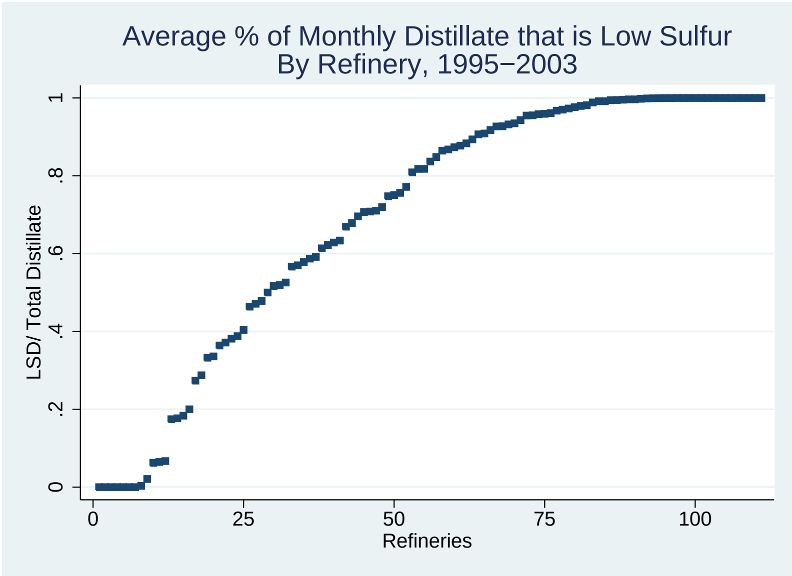

Table 4: Geographic Determinants of RFG and LSD

	Post 94 % RFG	Post 94 % LSD	Δ Desulf Cap 90-96
% States Served Pre 1990	0.934***	1.341***	0.900***
	(0.0673)	(0.421)	(0.235)
N	122	122	121
F	192.5	10.14	14.70

Notes: The independent variable in each regression is the share of each refinery's pre-1990 markets which subsequently became regulated post-1994. The dependent variable in the first two models is the share of each refinery's gasoline and distillate production which is regulated in the post period. The dependent variable in third column is the change in desulfurization capacity per unit of distillation capacity.

Table 5: Impact of Content Regulation on Output

	Average RFG %	RFG %	RFG % - IV
Gasoline Output / Capacity	0.0852***	0.0527*	-0.0139
	(0.0323)	(0.0271)	(0.0528)
N	18595	18595	18560

	Average LSD %	LSD %	LSD % - IV
Distillate Output/ Capacity	0.0528***	0.0340***	0.0445***
	(0.0126)	(0.00964)	(0.0109)
N	18595	18595	18560

Notes: All regressions include API gravity, upgrading capacity for major downstream units, indicators for the post-1994 period interacted with PADD dummies, and refinery and time dummies. The first stage F-stat for the IV models is 13.48 for RFG and 74.11 for LSD. Standard errors are clustered at the refinery level.

Figure 11: Observed Utilization Rates

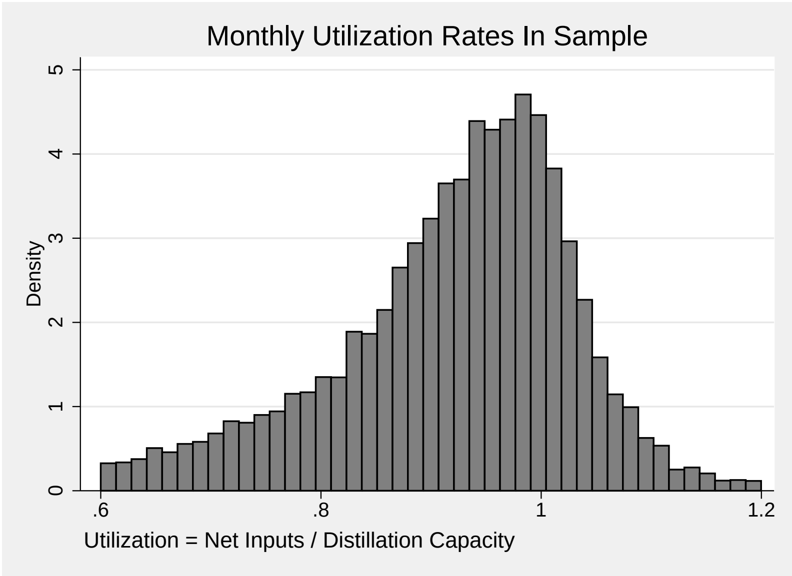

Figure 12: Refiner's Problem

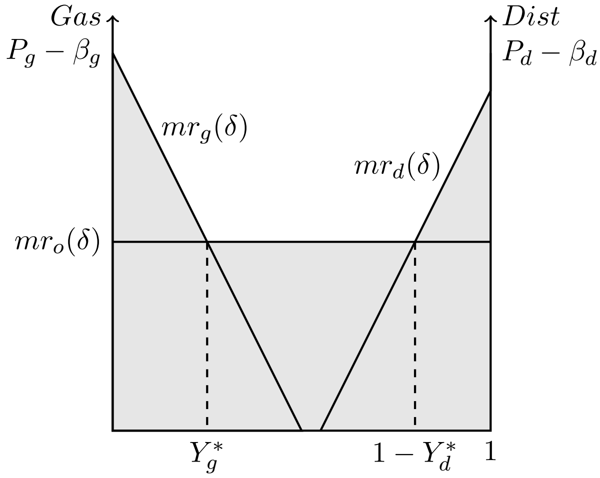

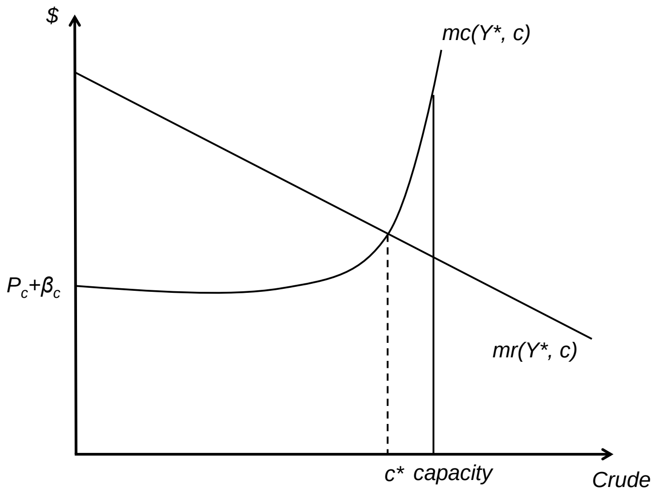

Figure 13: Seasonality in Gasoline and Distillate Demand

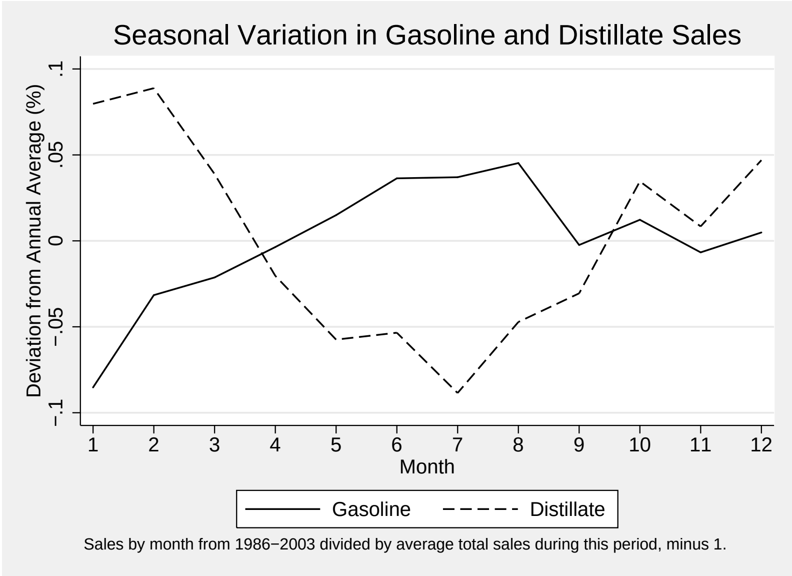

Figure 14: Seasonality in Distillate Demand by Sulfur Content

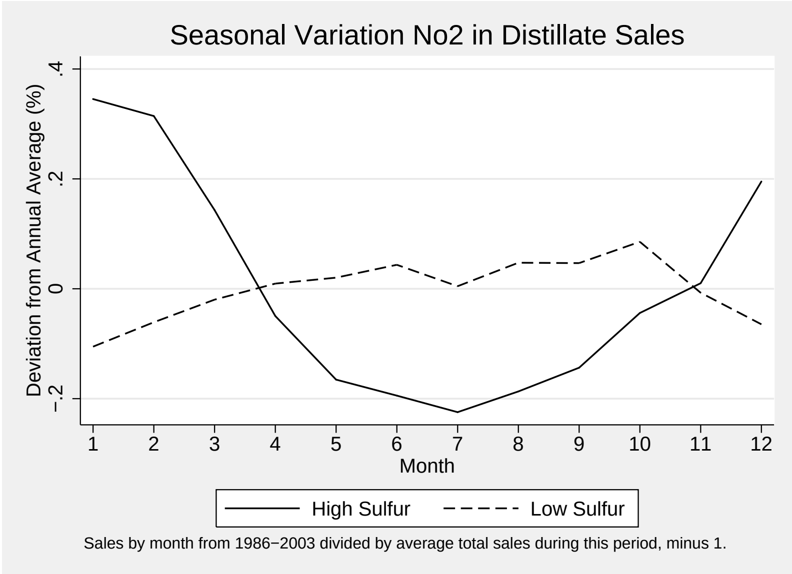

Table 6: Cost Function Estimates

	Model 1		Model 2	
Gasoline	**Est**	**SE**	**Est**	**SE**
Constant	0.379	(0.051)	0.582	(0.059)
API Gravity	0.176	(0.234)	-1.026	(0.273)
API Gravity^2	-1.243	(0.310)	0.233	(0.352)
% Upgrading	-0.181	(0.039)	-0.311	(0.044)
% Upgrading * API	0.480	(0.120)	0.913	(0.132)
Crude Sulfur	-0.003	(0.002)	-0.003	(0.002)
% Retail			0.161	(0.004)
RVP 1989	0.084	(0.004)	0.087	(0.004)
RVP 1992	-0.005	(0.003)	-0.002	(0.003)
Yield	0.147	(0.017)	0.137	(0.017)
Distillate				
Constant	0.241	(0.053)	0.415	(0.054)
API Gravity	0.706	(0.233)	-0.303	(0.240)
API Gravity^2	-1.536	(0.305)	-0.137	(0.314)
% Upgrading	-0.048	(0.037)	-0.132	(0.038)
% Upgrading * API	0.064	(0.112)	0.303	(0.114)
Crude Sulfur	0.007	(0.002)	0.001	(0.002)
% Retail			0.043	(0.002)
Yield	0.092	(0.021)	0.087	(0.022)
Other				
Jet fuel	0.376	(0.009)	0.373	(0.009)
Utilization = .85	-0.209	(0.040)	-0.083	(0.036)
Utilization = .94	-0.166	(0.036)	-0.060	(0.032)
Utilization = .99	-0.165	(0.037)	-0.055	(0.033)
Utilization = 1.2	0.106	(0.049)	0.164	(0.047)
Regulation Costs				
Average RFG Cost	0.096	(0.008)	0.071	(0.009)
Average LSD Cost	0.029	(0.004)	0.033	(0.004)
Elasticities				
Gasoline	1.61	(0.040)	1.32	(0.032)
Pre-94 Dist	2.00	(0.082)	2.05	(0.092)
HSD	3.18	(0.100)	4.34	(0.208)
LSD	2.09	(0.074)	2.22	(0.089)
Psuedo R2				
Gasoline Moment	0.66		0.67	
Distillate Moment	0.62		0.61	
Crude Moment	0.63		0.67	
RFG Moment	-0.03		0.15	
LSD Moment	-0.06		-0.06	
Gasoline Demand	0.87		0.88	
HSD Demand Pre-94	0.92		0.92	
HSD Demand Post-94	0.83		0.82	
LSD Demand	0.94		0.94	

Notes: All models contain 20,227 observations. The first three moments (FOC1-FOC3) contain refinery dummies and year dummies. All demand equations contain state-month dummies and time dummies. Robust standard errors presented.

Table 7: Content Regulation Costs

	Model 1		Model 2	
Reformulated Gasoline	**Est**	**SE**	**Est**	**SE**
Constant	12.3	(0.48)	9.3	(0.56)
RFG 1 Summer	-0.7	(0.33)	0.0	(0.36)
RFG 2	-1.7	(0.12)	-0.7	(0.15)
RFG 2 Summer	3.5	(0.24)	3.1	(0.29)
CARB	-0.6	(0.59)	0.7	(0.68)
CARB Summer	4.6	(0.36)	3.4	(0.41)
MTBE	-3.6	(0.25)	-3.8	(0.29)
Average RFG Cost	9.6	(0.80)	7.1	(0.93)
Wgt. Avg. RFG cost	8.0		6.2	
EPA Estimates				
Phase 1	4.8-7.8			
Phase 2	8.6			
NPC Estimates				
Phase 1	8.6			
Phase 2	11.5			
Low Sulfur Diesel				
Constant	2.3	(1.27)	0.7	(1.43)
Upgrading Capacity	-0.4	(0.30)	-2.8	(0.35)
% Desulfurization	3.8	(2.03)	7.2	(2.29)
% Desulfurization ^ 2	1.3	(0.60)	1.1	(0.63)
Crude Sulfur	0.7	(0.24)	1.1	(0.27)
% Desulfurization * Sulfur	-1.9	(0.33)	-2.1	(0.38)
API Gravity	11.9	(2.93)	17.2	(3.38)
% Desulfurization * API	-28.9	(4.05)	-32.8	(4.67)
CA	2.4	(0.26)	2.4	(0.28)
Small Refinery	1.1	(0.15)	1.5	(0.17)
Average LSD Cost	2.9	(0.36)	3.3	(0.42)
Wgt. Avg. LSD cost	2.2		2.6	
EPA Estimate	4.3			
NPC Estimate	6.8			

Notes: Sources: EPA (1990, 1993), NPC (1990). All costs in real (2013) cents per gallon.

Figure 15: RFG Cost Distribution

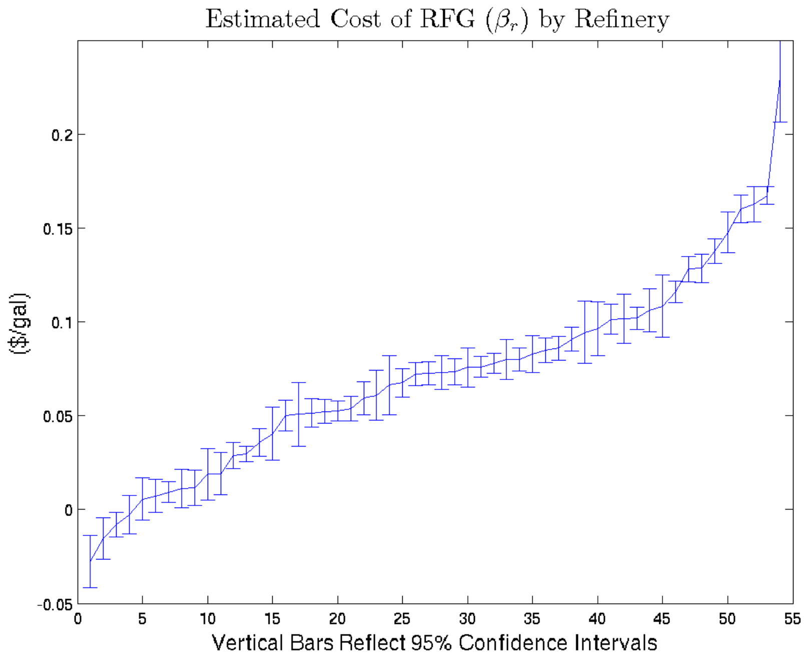

Figure 16: LSD Cost Distribution

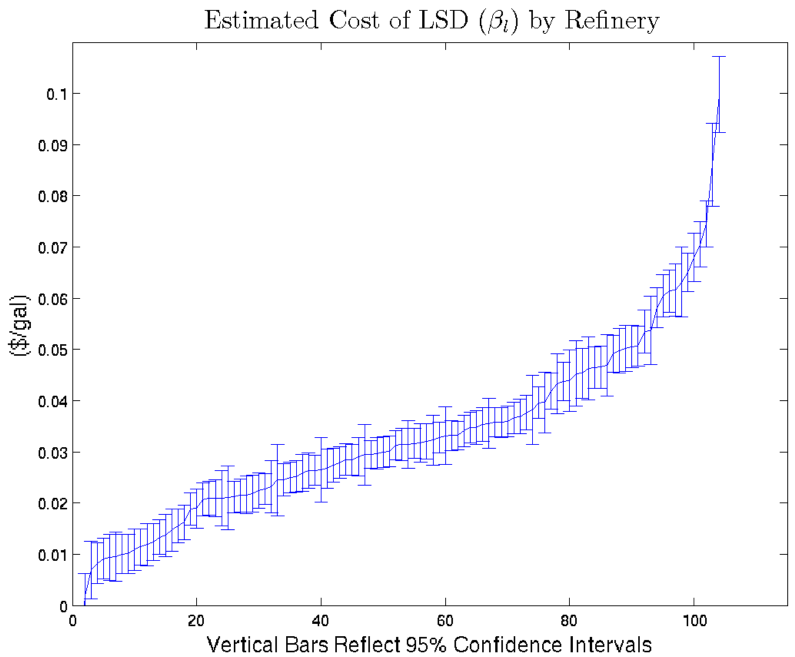

Table 8: Demand Estimates

	Gasoline	All Distillate	High Sulfur Distillate	Low Sulfur Diesel
Years	1986-2003	1986-1993	1994-2003	1994-2003
Uninstrumented	0.569	0.790	0.852	0.756
	(0.163)	(0.177)	(0.149)	(0.248)
Instrumented	1.119	0.425	1.157	1.268
	(0.371)	(1.326)	(2.430)	(0.910)
First Stage F-stat	25.35	4.35	11.56	8.5
Joint Estimates	1.324	2.046	4.335	2.219
	(0.032)	(0.092)	(0.208)	(0.089)

Notes: All regressions contain state-month and time dummies. Standard errors in parentheses. Offline regressions clustered at state level. Joint standard errors are robust.

Figure 17: Simulation Regions

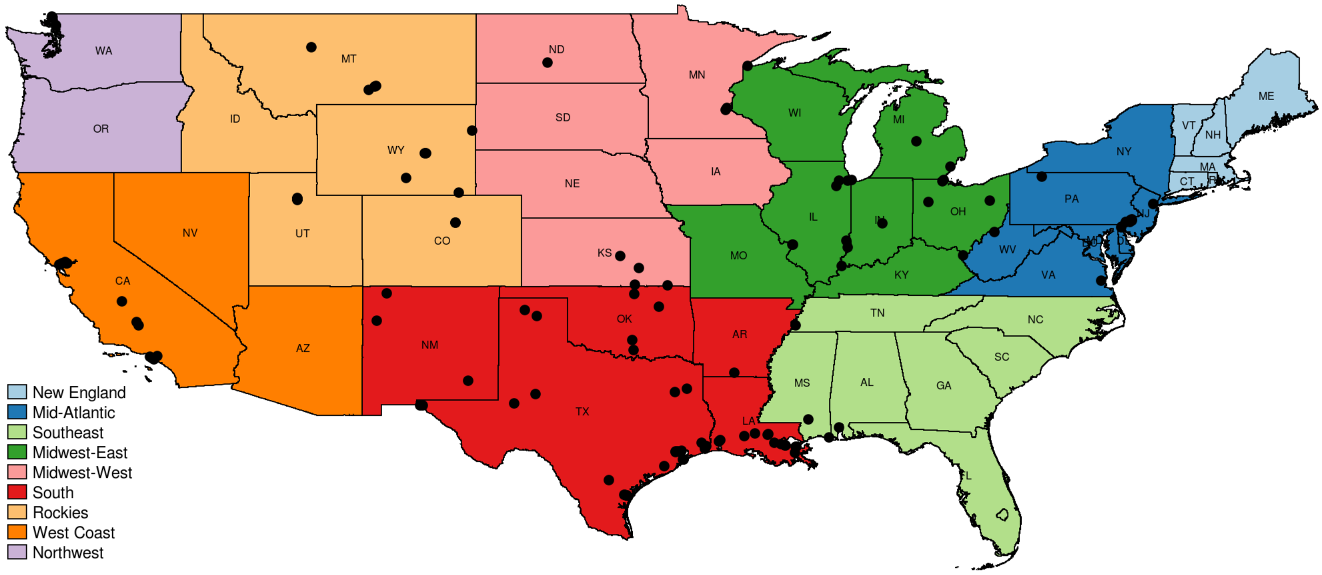

Table 9: Counterfactual Results

	Gasoline Price Changes						
	Reformulated			Conventional			Average Cost
Region	**c/gal**	**%**	**Surplus ($M)**	**c/gal**	**%**	**Surplus ($M)**	**Increase (c/gal)**
New England	6.0	5.3%	-2,210	-2.0	-1.8%	85	5.7
Mid-Atlantic	6.5	5.9%	-7,772	-1.5	-1.5%	1,115	5.4
Southeast	-	-	-	-1.1	-1.1%	2,293	
Midwest - East	8.2	7.2%	-3,811	-1.4	-1.4%	2,510	5.7
Midwest - West	-	-	-	-1.2	-1.1%	720	
South	6.3	5.8%	-2,617	-1.2	-1.2%	1,914	4.5
Rocky Mountain	-	-	-	-1.5	-1.4%	590	
West Coast	7.0	5.4%	-9,241	-3.7	-3.4%	973	5.9
Northwest	-	-	-	-3.4	-3.0%	1,286	
Total			**-25,651**			**11,486**	

	Distillate Price Changes						
	Highway Diesel			Non-Highway Distillate			Average Cost
Region	**c/gal**	**%**	**Surplus ($M)**	**c/gal**	**%**	**Surplus ($M)**	**Increase (c/gal)**
New England	2.6	2.6%	-126	-0.9	-1.0%	98	2.3
Mid-Atlantic	2.3	2.4%	-855	-1.3	-1.4%	619	2.4
Southeast	2.5	2.6%	-1,164	-1.3	-1.4%	258	2.7
Midwest - East	2.5	2.6%	-1,410	-1.2	-1.4%	337	2.5
Midwest - West	2.7	2.8%	-704	-1.1	-1.2%	87	1.8
South	2.4	2.5%	-1,716	-1.4	-1.6%	478	2.7
Rocky Mountain	2.5	2.5%	-439	-2.8	-2.9%	121	4.0
West Coast	3.4	3.2%	-1,267	-5.7	-6.4%	132	4.2
Northwest	2.7	2.7%	-276	-0.7	-0.8%	44	3.3
Total			**-7,957**			**2,172**	

Notes: All numbers presented are changes relative to the counterfactual where content restrictions are removed (for example, refinery profits per gallon with RFG and LSD in place minus profits in a world with no content restrictions). The last column reports the average incremental cost of producing RFG and LSD relative to conventional gasoline and distillate at refineries serving each region.

Table 10: Counterfactual Results - Refinery Profits

Padd	**Region**	**c/gal**	**($M)**	**%**
1	East Coast	-0.79	-1,646	-12.1
2	Midwest	-0.60	-3,070	-9.4
3	Gulf Coast	-0.24	-2,399	-5.1
4	Rocky Mountain	-1.48	-1,107	-24.8
5	West Coast	-0.45	-1,475	-7.4
	Total	**-0.45**	**-9,696**	**-8.2**

Notes: All numbers presented are changes relative to the counterfactual where content restrictions are removed (for example, refinery profits per gallon with RFG and LSD in place minus profits in a world with no content restrictions).

References

[1] Andersen, Soren and Andrew Elzinga (2014). "A Ban on One is a Boon for the Other: Strict Gasoline Content Rules and Implicit Ethanol Blending Mandates." *Journal of Environmental Economics and Management* 67 (3): pages 258-273.

[2] Auffhammer, Maximilian and Ryan Kellogg (2011). "Clearing the Air? The Effects of Gasoline Content Regulation on Air Quality." *American Economic Review*, 101(4), pages 2687-2722.

[3] Berman, Eli and Linda Bui (2001). "Environmental Regulation and Productivity: Evidence from Oil Refineries." *Review of Economics and Statistics*, 83(3): 498–510.

[4] Berry, Steven, James Levinsohn, and Ariel Pakes (1995). "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

[5] Brown, Jennifer, Justine Hastings, Erin T. Mansur, and Sofia B. Villas-Boas (2008). "Reformulating Competition? Gasoline Content Regulation and Wholesale Gasoline Prices." *Journal of Environmental Economics and Management*, Vol. 55, No. 1, pages 1-19.

[6] Borenstein, Severin and Andrea Shepard (2002). "Sticky Prices, Inventories, and Market Power in Wholesale Gasoline Markets." *RAND Journal of Economics*, Vol. 33, No 1.

[7] Bulow, Jeremy, John Geanakoplos, and Paul Klemperer (1985). "Multimarket Oligopoly: Strategic Substitutes and Complements." *Journal of Political Economy*, Vol. 93, no. 3.

[8] Chakravorty, Ujjayant, Celine Nauges, Alban Thomas (2008). "Clean Air Regulation and Heterogeneity in US Gasoline Prices." *Journal of Environmental Economics and Management*, Vol. 55, pages 106-122.

[9] Chesnes, Matthew (2014). "The Impact of Outages on Prices and Investment in the US Oil Refining Industry." *(working paper)*

[10] Chouinard, Haley H. and Jeffrey M. Perloff, "Gasoline Price Differences: Taxes, Pollution Regulations, Mergers, Market Power and Market Conditions", *The B.E. Journal of Economic Analysis and Policy* , *7*(1), article 8.

[11] Davis, Lucas and Lutz Killian (2001). "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics*, 16(7), 1187-1214.

[12] Environmental Protection Agency (1990). "Regulatory Impact Analysis: Control of Sulfur and Aromatics Contents of On-Highway Diesel Fule." Report EPA-420-R-90-103.

[13] Environmental Protection Agency (1993). "Final Regulatory Impact Analysis for Reformulated Gasoline." Report EPA-420-R-93-017.

[14] Environmental Protection Agency (2011). "The Benefits and Costs of the Clean Air Act from 1990 to 2020." Report.

[15] The Federal Trade Commission (2004). "The Petroleum Industry: Mergers, Structural Change and Antitrust Enforcement." Report. August 2004.

[16] The Federal Trade Commission (2006). "Investigation of Gasoline Price Manipulation and Post-Katrina Gasoline Price Increases." Report.

[17] The Federal Trade Commission (2006). "Gasoline Price Changes and the Petroleum Industry: An Update." Report. September 2011.

[18] The Federal Trade Commission (2010). *Horizontal Merger Guidelines*, Issued August 19, 2010.

[19] Fowlie, Meredith, Mar Reguant, and Stephen Ryan (2014). "Market-based Environmental Regulation and the Evolution of Market Structure", forthcoming *Journal of Political Economy*

[20] General Accounting Office (2004). "Effects of Mergers and Market Concentration in the U.S. Petroleum Industry." Report to the Ranking Minority Member, Permanent Subcommittee on Investigations, Committee on Governmental Affairs, U.S. Senate. GAO-04-96.

[21] Gilbert, Richard and Justine Hastings (2005). "Market Power, Vertical Integration and the Wholesale Price of Gasoline." *The Journal of Industrial Economics*, Vol. LIII (4).

[22] Greenstone, Michael (2002). "The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufacturers." *Journal of Political Economy*, 110(6), pages 1175-1219.

[23] Hansen, L.P. (1982), "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, 50, 1029-1054,

[24] Hendricks, Ken, Preston McAfee, Michael Williams (2007), "Evaluating the Likely Competitive Effects of Horizontal and Vertical Mergers: A New Approach", *Antitrust Report*, (2007) Issue 2, pp. 33-40.

[25] Li, Shanjun, Joshua Linn, and Erich Muehlegger (2014), "Gasoline Taxes and Consumer Behavior.", American Economic Journal: Economic Policy, *forthcoming.*

[26] Lidderdale, Tancred, "Demand, Supply, and Price Outlook for Low-Sulfur Diesel Fuel." Energy Information Administration, Monthly Energy Review, August 1993.

[27] Marion, Justin and Erich Muehlegger (2008), "Measuring Illegal Activity and the Effects of Regulatory Innovation: Tax Evasion and the Dyeing of Untaxed Diesel." *Journal of Political Economy*, 114:6, pages 633-666.

[28] Millimet, Daniel L., Santanu Roy, and Aditi Sengupta (2009). "Environmental Regulations and Economic Activity: Influence on Market Structure." *Annual Review of Resource Economics*, Vol. 1, pages 99-117.

[29] Muehlegger, Erich (2004). "Product Differentiation in Gasoline Markets: A Discussion of Regional Content Regulations" *(working paper)*

[30] Muehlegger, Erich (2006). "Gasoline Price Spikes and Regional Gasoline Content Regulation: A Structural Approach." *(working paper)*

[31] National Petroleum Council (1993). "U.S. Petroleum Refining: Meeting Requirements for Cleaner Fuels and Refineries." Report, August 1993.

[32] Reiss, P. C., & Wolak, F. A. (2007). "Structural Econometric Modeling: Rationales and Examples from Industrial Organization." *Handbook of Econometrics*, 6, 4277-4415.

[33] Rosse, J.N. (1970). "Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology". Econometrica 38 (2), 256–275

[34] Ryan, Stephen P. (2012). "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica*, Vol. 80, No. 3 (May), pages 1019-1062.

[35] Sweeney, Richard L. (2015). "Estimating the Sunk Costs of Environmental Regulation on a Large Concentrated Industry", *(working paper)* .

[36] Zhang, Ker (2011). "The Effect of Sulfur Content Regulation on the U.S. Diesel Market." *(working paper)*

A Appendix A: Data Appendix

Table A.1: Description of EIA Data

Survey	Dates	Description
Monthly Refinery Report (EIA-810)	1986-2012	Collects information regarding the balance between the supply (beginning stocks, receipts, and production) and disposition (inputs, shipments, fuel use and losses, and ending stocks) of crude oil and refined products located at refineries.
Annual Refinery Report (EIA-820)	1986-1995 1997 1999-2012	Collects data on: fuel, electricity, and steam purchased for consumption at the refinery; refinery receipts of crude oil by method of transportation; current and projected capacities for atmospheric crude oil distillation, downstream charge, and production capacities.
Refiners' Monthly Cost Report (EIA-14)	2002-2012	Collects data on the weighted cost of crude oil at the regional Petroleum for Administration Defense District (PADD) level at which the crude oil is booked into a refinery.
Refiners'/Gas Plant Operators' Monthly Petroleum Product Sales Report (EIA-782A)	1986-2012	Price and volume data at the State level for 14 petroleum products for various retail and wholesale marketing categories are reported by the universe of refiners and gas plant operators
Monthly Report of Prime Supplier Sales of Petroleum Products Sold for Local Consumption (EIA-782C)	1986-1990 1992-2012	Prime supplier sales of selected petroleum products into the local markets of ultimate consumption are reported by refiners, gas plant operators, importers, petroleum product resellers, and petroleum product retailers that produce, import, or transport product across State boundaries and local marketing areas and sell the product to local distributors, local retailers, or end users.

Notes: Additional information as well as the survey forms for each dataset available at http://www.eia.gov/survey/.

Figure A.1: Wholesale Prices and Concentration by State

	Gasoline								Distillate							
	Pre-1993		1995-2003						Pre-1993		1995-2003					
	Conventional		Conventional		Reformulated		RFG	Price	High Sulfur		High Sulfur		Low Sulfur		LSD	Price
State	**Price**	**HHI**	**Price**	**HHI**	**Price**	**HHI**	**Share**	**Diff.**	**Price**	**HHI**	**Price**	**HHI**	**Price**	**HHI**	**Share**	**Diff.**
Padd 1 - East Coast																
CT	1.35	0.14			0.92	1.00			1.09	0.20	0.91	0.28	0.94	0.21	0.40	0.03
DC	1.42	0.31			0.00	0.00			1.07	0.46	0.88	0.81	1.03	0.38	0.48	0.11
DE	1.25	0.21			0.00	0.00			1.04	0.79	0.91	0.77	0.93	0.63	0.66	0.02
FL	1.24	0.07	1.04	0.11					1.06	0.10	0.90	0.14	0.93	0.11	0.70	0.03
GA	1.18	0.07	1.00	0.12					1.04	0.08	0.89	0.15	0.91	0.11	0.76	0.02
MA	1.34	0.15			0.80	0.93			1.09	0.22	0.90	0.28	0.95	0.23	0.37	0.05
MD	1.34	0.13	1.00	0.14	1.00	0.14	0.90	0.14	1.08	0.12	0.89	0.21	0.94	0.15	0.56	0.04
ME	1.23	0.20	1.02	0.25	1.02	0.25	0.24	-0.01	1.09	0.23	0.91	0.31	0.96	0.30	0.36	0.04
NC	1.17	0.08	0.99	0.12					1.04	0.09	0.89	0.11	0.91	0.13	0.66	0.02
NH	1.34	0.22	1.04	0.45	1.04	0.45	0.87	0.11	1.10	0.31	0.91	0.61	0.98	0.43	0.45	0.06
NJ	1.27	0.10	0.92	0.31	0.92	0.31	0.90	0.19	1.05	0.11	0.88	0.17	0.91	0.17	0.37	0.03
NY	1.35	0.14	1.03	0.20	1.03	0.20	0.51	0.18	1.09	0.12	0.92	0.16	0.96	0.16	0.34	0.04
PA	1.21	0.11	1.01	0.12	1.01	0.12	0.26	0.11	1.07	0.10	0.90	0.26	0.93	0.20	0.53	0.02
RI	1.31	0.16			0.89	1.00			1.08	0.34	0.91	0.32	0.95	0.38	0.35	0.04
SC	1.17	0.08	0.99	0.11					1.05	0.10	0.90	0.15	0.91	0.12	0.73	0.01
VA	1.23	0.08	0.98	0.12	0.98	0.12	0.57	0.14	1.05	0.08	0.89	0.12	0.92	0.13	0.63	0.03
VT	1.34	0.57	1.08	0.54					1.13	0.32	0.92	0.61	0.96	0.50	0.39	0.04
WV	1.23	0.23	1.02	0.19					1.09	0.23	0.95	0.52	0.95	0.26	0.47	0.00
Padd 2 - Midwest																
IA	1.18	0.08	1.03	0.11					1.08	0.12	0.95	0.21	0.97	0.14	0.86	0.02
IL	1.21	0.12	1.00	0.12	1.00	0.12	0.54	0.20	1.04	0.14	0.90	0.21	0.92	0.14	0.64	0.02
IN	1.17	0.12	1.02	0.18	1.02	0.18	0.14	0.12	1.06	0.13	0.93	0.19	0.93	0.15	0.71	0.00
KS	1.12	0.09	1.00	0.12					1.04	0.11	0.92	0.22	0.95	0.15	0.80	0.03
KY	1.21	0.16	1.01	0.20	1.01	0.20	0.26	0.10	1.08	0.17	0.93	0.30	0.94	0.31	0.57	0.02
MI	1.21	0.11	1.05	0.16					1.08	0.15	0.93	0.28	0.96	0.16	0.74	0.03
MN	1.20	0.11	1.08	0.13					1.09	0.12	0.96	0.23	0.98	0.19	0.70	0.02
MO	1.17	0.08	1.02	0.12	1.02	0.12	0.15	0.10	1.05	0.09	0.91	0.20	0.94	0.12	0.87	0.03
ND	1.21	0.15	1.05	0.21					1.11	0.21	1.00	0.69	1.01	0.31	0.67	0.01
NE	1.16	0.08	1.03	0.12					1.07	0.13	0.94	0.24	0.97	0.16	0.79	0.03
OH	1.20	0.17	1.05	0.19					1.10	0.17	0.95	0.31	0.96	0.24	0.66	0.00
OK	1.11	0.09	0.97	0.11					1.02	0.09	0.90	0.30	0.92	0.14	0.69	0.02
SD	1.19	0.08	1.05	0.11					1.10	0.11	0.96	0.27	1.00	0.14	0.91	0.03
TN	1.17	0.07	1.00	0.11					1.04	0.08	0.90	0.15	0.92	0.12	0.71	0.02
WI	1.18	0.09	1.02	0.13	1.02	0.13	0.29	0.10	1.07	0.13	0.93	0.18	0.95	0.21	0.73	0.02
Padd 3 - South																
AL	1.17	0.07	0.99	0.11					1.03	0.08	0.88	0.14	0.90	0.11	0.69	0.02
AR	1.15	0.07	0.99	0.10					1.04	0.09	0.88	0.23	0.92	0.12	0.65	0.04
LA	1.14	0.08	0.97	0.10	0.97	0.10	0.02	-0.02	1.00	0.12	0.84	0.18	0.88	0.12	0.31	0.04
MS	1.13	0.08	0.97	0.11					1.01	0.15	0.86	0.20	0.89	0.11	0.61	0.03
NM	1.23	0.09	1.07	0.12					1.12	0.11	0.96	0.43	1.00	0.17	0.92	0.04
TX	1.12	0.06	0.94	0.08	0.94	0.08	0.29	0.09	1.00	0.06	0.85	0.14	0.89	0.09	0.72	0.03
Padd 4 - Plains																
CO	1.19	0.10	1.07	0.13					1.08	0.12	0.96	0.48	0.99	0.15	0.84	0.03
ID	1.21	0.14	1.14	0.15					1.17	0.15	1.05	0.29	1.07	0.19	0.70	0.03
MT	1.20	0.17	1.13	0.23					1.12	0.20	1.00	0.76	1.05	0.24	0.98	0.07
UT	1.19	0.12	1.13	0.15					1.14	0.16	1.04	0.31	1.07	0.22	0.81	0.04
WY	1.21	0.12	1.10	0.14					1.12	0.20	1.02	0.84	1.03	0.15	0.79	0.00
Padd 5 - West Coast																
AZ	1.25	0.12	1.16	0.13	1.16	0.13	0.43	0.08	1.12	0.13	0.96	0.45	1.00	0.13	0.88	0.04
CA	1.27	0.10	1.07	0.31	1.07	0.31	0.93	0.16	1.04	0.11	0.98	0.46	1.03	0.17	0.96	0.05
NV	1.22	0.14	1.18	0.17					1.10	0.15	1.01	0.45	1.01	0.13	0.90	0.00
OR	1.23	0.14	1.14	0.18					1.06	0.21	0.97	0.44	0.99	0.20	0.64	0.02
WA	1.21	0.12	1.16	0.15					1.05	0.18	0.94	0.38	0.99	0.15	0.60	0.05

Notes: Prices are average monthly prices for resale ($2013 / gal). Markets with sales in less than 50% of the sample are excluded.

Appendix B: Crude price variable construction

Crude oil makes up over 80 percent of a refinery's variable costs. While crude oil quality and volumes are observed at the refinery level in the EIA data, crude oil prices are not. Prices are observed at the firm-PADD level beginning in 2004. In order to construct a proxy for refinery-level crude prices during the sample, I estimate a regression relating this firm-PADD level data to publicly available crude price series that are available for the entirety of the sample.[25] I then predict crude prices for each refinery in each month and use this variable when estimating the structural model.

Beginning in 2004, survey EIA-14 records the monthly total cost and volume of domestic (D) and imported (I) crude oil acquired across all refineries owned by a firm in each Petroleum Administration Defense District.[26] These costs are assumed to be a function of benchmark crude prices, PADD-level domestic crude prices and cost shifters, and a quality premium on API gravity. West Texas Intermediate and Brent spot prices serve as domestic and imported crude benchmark prices. Regional variation in domestic crude prices is captured by the EIA's cost of first purchase price (p^{cofp}) series, which reports the average price received by domestic oil producers in each PADD. The EIA reports the average landed cost of imported crude (p^{land}) by API gravity bin (b) beginning 1986.[27] Using the midpoint of each bin, I construct a price per API gravity degree premium variable for each month, $\zeta_b^I = \frac{p_b^{land} - p^{brent}}{API_b - API^{brent}}$. Domestic prices by API gravity bin are not available before 1994. Therefore the domestic API gravity premium is proxied with the average price premium of two domestic heavy crudes, Alaska North Slope an Gulf Heavy.

Given the available price data, firm-PADD crude price are modeled as follows,

$$\begin{aligned} p_{fr} \quad = \quad & \alpha_0 + s_{fr}(\alpha_{D0r} + \alpha_{D1r}p^{wti} + \alpha_{D2}\zeta^D\Delta_{fr}^D + \alpha_{D3}(p_r^{cofp} - p^{wti})) \\ & +(1 - s_{fr})(\alpha_{I0r} + \alpha_{I1r}p^{brent} + \alpha_{I2}\zeta_{fr}^I\Delta_{fr}^I) + \alpha_2 s_{fr}\zeta^D\Delta_{fr}^D(1 - s_{fr})\zeta_{fr}^I\Delta_{fr}^I \end{aligned}$$

Where s_{rf} is the fraction of crude processed by firm f in PADD r that is domestic, and Δ is the difference in API gravity between this crude and the benchmark crude (i.e. $\Delta_{fr}^I = API_{fr} - API^{brent}$). The final term is an interaction between domestic and imported API gravity premiums to account for the fact than refineries only report average API gravity each month, rather than separate figures for domestic and imported crude streams.

[25] All of the crude price series discussed here can be downloaded at http://www.eia.gov/petroleum/data.cfm#prices

[26] The survey actually begins in 2002, but data was only collected at the national level until 2004.

[27] For example, the average cost of imported crude with API gravity between 20 and 25 degrees.

Table A.2: Crude Price Estimates

	Average Crude Price	
Constant	-0.159	(2.090)
Domestic - PADD 1	7.520***	(2.766)
Domestic - PADD 2	-2.461	(2.273)
Domestic - PADD 3	2.756	(2.284)
Domestic - PADD 4	4.453*	(2.429)
Domestic - PADD 5	1.207	(2.290)
COFP - PADD 1	1.021***	(0.207)
COFP - PADD 2	0.499***	(0.102)
COFP - PADD 3	1.139***	(0.0646)
COFP - PADD 4	0.962***	(0.115)
COFP - PADD 5	0.827***	(0.0417)
WTI - PADD 1	0.948***	(0.0207)
WTI - PADD 2	1.057***	(0.0104)
WTI - PADD 3	1.025***	(0.0101)
WTI - PADD 4	1.008***	(0.0148)
WTI - PADD 5	1.004***	(0.0103)
API Premium - Domestic	0.0554***	(0.0161)
Imported - PADD 1	3.071	(2.270)
Imported - PADD 2	12.11***	(2.343)
Imported - PADD 3	-2.336	(2.203)
Imported - PADD 4	12.87***	(2.355)
Imported - PADD 5	7.499***	(2.811)
Brent - PADD 1	0.920***	(0.0108)
Brent - PADD 2	0.717***	(0.0122)
Brent - PADD 3	0.979***	(0.00863)
Brent - PADD 4	0.687***	(0.0125)
Brent - PADD 5	0.924***	(0.0204)
Brent - PADD 6	0.999***	(0.0240)
API Premium - Imported	0.137***	(0.0106)
API Premium Interaction	-0.0319***	(0.00405)
N	8258	
r2	0.929	

Standard errors in parentheses

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

Appendix B: Market power and refinery utilization rates

In order to demonstrate the importance of market power in this setting, I run several versions of the following regression,

$$Utilization_{it} = \alpha_i + \gamma Capshare_{fi} + \epsilon_{it}$$

$Capshare_{ft}$ is the sum of distillation capacity across all refineries owned by firm f in the same PADD as refinery i divided by the total distillation capacity in that PADD during that month t. Regressions are run on the full set of data available, beginning in 1986 and ending in 2012. Ownership concentrations within a PADD vary considerably during this time period, as can be seen in Figure A.2. The primary driver of these changes was a wave of mergers and acquisitions seen in Table 3. In the IV regressions, I used the prior prior year's capacity at each refinery aggregated using the current year's ownership status. This instrument intuitively takes advantage of variation from ownership changes alone rather than increases in investment capacity within a given year.

Figure A.2: Changes in Capacity Concentration by PADD

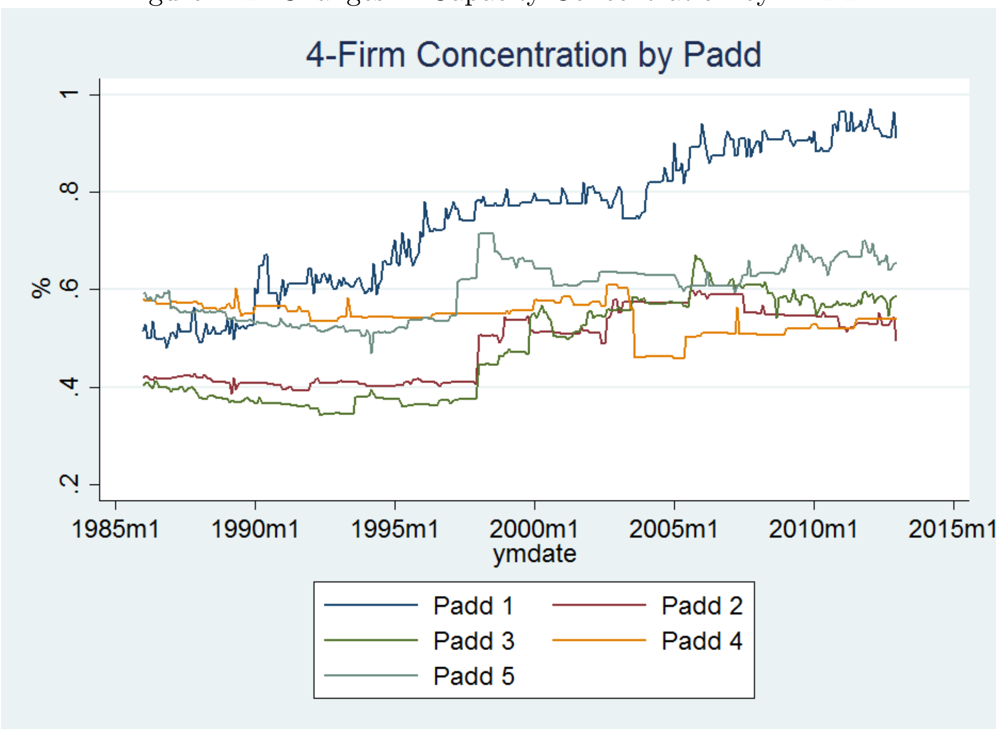

Table A.3: Refinery Utilization

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
share_totcap_padd	-0.398***	-0.362**	-0.399***	-0.363**	-0.327**	-0.294**	-0.327**	-0.295**
	(0.142)	(0.143)	(0.142)	(0.143)	(0.148)	(0.150)	(0.149)	(0.150)
IV		X		X		X		X
Refinery FE	X	X	X	X	X	X	X	X
Year FE	X	X			X	X		
Month FE	X	X			X	X		
Time FE			X	X			X	X
Padd*Year FE					X	X	X	X
Observations	41608	39916	41608	39916	41608	39916	41608	39916
R-Squared	0.358	0.340	0.363	0.344	0.367	0.349	0.372	0.354

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors reported, clustered by refinery. All specifications include refinery fixed effects. The dependent variable in all specifications monthly utilization (inputs divided by distillation capacity).