

Colina, Armando R.; Gafarov, Bulat; Hilscher, Jens

Working Paper

California gasoline demand elasticity estimated using refinery outages

Working Papers, No. 2025-04

Provided in Cooperation with:

Bank of Mexico, Mexico City

Suggested Citation: Colina, Armando R.; Gafarov, Bulat; Hilscher, Jens (2025) : California gasoline demand elasticity estimated using refinery outages, Working Papers, No. 2025-04, Banco de México, Ciudad de México

This Version is available at:

<https://hdl.handle.net/10419/322347>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Banco de México

Working Papers

N° 2025-04

**California Gasoline Demand Elasticity Estimated
Using Refinery Outages**

Armando R. Colina

Banco de México

Bulat Gafarov

University of California, Davis

Jens Hilscher

University of California, Davis

April 2025

La serie de Documentos de Investigación del Banco de México divulga resultados preliminares de trabajos de investigación económica realizados en el Banco de México con la finalidad de propiciar el intercambio y debate de ideas. El contenido de los Documentos de Investigación, así como las conclusiones que de ellos se derivan, son responsabilidad exclusiva de los autores y no reflejan necesariamente las del Banco de México.

The Working Papers series of Banco de México disseminates preliminary results of economic research conducted at Banco de México in order to promote the exchange and debate of ideas. The views and conclusions presented in the Working Papers are exclusively the responsibility of the authors and do not necessarily reflect those of Banco de México.

California Gasoline Demand Elasticity Estimated Using Refinery Outages^{*}

Armando R. Colina[†]
Banco de México

Bulat Gafarov[‡]
University of California, Davis

Jens Hilscher[§]
University of California, Davis

Abstract: This research paper presents new gasoline demand price elasticity estimates for California. The study leverages the unique characteristics of California's gasoline market and introduces a new set of proposed instruments. As a first step, the study takes advantage of California's partially isolated gasoline market, which is separated from much of the U.S. by environmental regulations. The estimating equation controls for persistent demand shocks, leading to a lower bound for the long-run elasticity of demand of -0.23. In the second step, a new set of instruments is employed to address supply and demand simultaneity, using detailed information on refinery outages to capture short-run supply shocks. The results show that the long-run demand elasticity is -0.60.

Keywords: Capacity outages, Gasoline demand price elasticity, Instrumental variable estimation, California gasoline market

JEL Classification: C22, C36, C51, D12, Q41

Resumen: Este trabajo de investigación presenta nuevas estimaciones de la elasticidad precio de la demanda de gasolina para California. El estudio aprovecha las características únicas del mercado de gasolina de California y propone un nuevo conjunto de instrumentos para estimar la elasticidad. Como primer paso, el estudio aprovecha el aislamiento parcial del mercado californiano respecto al resto de Estados Unidos debido a regulaciones ambientales. La ecuación de estimación controla por choques persistentes de demanda estimando un límite inferior para la elasticidad de la demanda a largo plazo de -0.23. En el segundo paso, se emplea un nuevo conjunto de instrumentos para abordar la simultaneidad entre oferta y demanda. La estrategia de estimación utiliza información detallada sobre interrupciones en refinerías para capturar choques de oferta de corto plazo. Los resultados muestran que la elasticidad de la demanda a largo plazo es de -0.60.

Palabras Clave: Fallas no programadas, Elasticidad precio de la demanda de gasolina, Estimación por variables instrumentales, Mercado de gasolina en California

^{*}This paper is based on the third chapter of Armando's Ph.D. dissertation. We want to thank Richard J. Sexton, James B. Bushnell, and Erich Muehlegger, for their helpful guidance and comments, as well as two anonymous reviewers at ENEECO and two anonymous reviewers at Banco de México. We are grateful to CONAHCYT, Banco de México, and the USMEX Institute at UC San Diego for their financial support.

[†] Dirección General de Asuntos del Sistema Financiero. Email: armando.rangel@banxico.org.mx

[‡] Department of Agricultural and Resource Economics. Email: bgafarov@ucdavis.edu

[§] Department of Agricultural and Resource Economics. Email: jhilscher@gmail.com

1 Introduction

The long-term price elasticity of demand for gasoline is an important parameter used in economic research, policy design, and business decisions.¹ Several previous studies estimate the average elasticity by aggregating data across many different regions. This strategy was largely due to the fact that the data was previously somewhat limited.

A recent contribution by Kilian and Zhou (2024) is one of the first papers to use more detailed state-level data to analyze heterogeneity in markets. They find that crude oil price pass-through into gasoline prices varies systematically across regions. Their findings imply that there may be important heterogeneity between states. Liu (2014) also finds evidence of heterogeneity in gasoline demand price elasticity. As the largest market by revenue in the United States, California has been at the forefront of regulating gasoline and energy markets to encourage the transition to a lower greenhouse gas emission economy (e.g. by introducing a cap and trade system, or low carbon fuel standards). Its fuel market is also largely segmented from the rest of the country. As a result of various policies and generally higher prices, California consumers may react differently to changes in gasoline prices.

In this paper, we focus on a recent sample period for California and estimate a California-specific price elasticity. Estimating the price elasticity using observational data has proven to be challenging. For effective estimation, researchers must control for demand shocks and require supply-side price variation, achieved through an instrumental variable (IV) approach. Our focus on the California gasoline market allows us to take advantage of its specific features, which in turn allows us to identify the price elasticity of demand.

The California gasoline market is unique within the United States, given its large size and strict environmental regulations.² To achieve the environmental standards set by the California Air Resources Board (CARB), refiners in California must build special and costly units that produce specific blending components that are not required in the rest of the country. As a result, CARB-grade gasoline is only consumed in California and is almost exclusively produced there. Consequently, California is considered to be nearly a separated gasoline market. Therefore, reductions in refining capacity in

¹For example, Holland et al. (2009) estimate the social costs of implementing the Low Carbon Fuel Standard (LCFS) and simulate different scenarios based on different supply and demand price elasticities. Parry et al. (2022) weigh the benefits and downsides of different carbon pricing policies.

²See Table 5 in the Annex for a detailed description of the specifications required of CARB-compliant gasoline.

California cannot be easily compensated by importing gasoline and can greatly impact prices. These market features allow us to control for demand shocks first and then use supply shocks that have strong associations with retail prices.

We start by using detailed controls for demand shocks. To the extent that demand is persistent, we can (at least partially) control for demand shocks by including lagged sales. The rich data in California allow us to also control for inventories, imports, and capacity utilization. For example, if a persistent demand shock hits, refiners will adjust their inventories and refinery utilization in anticipation of this prolonged shock.³ Including such lagged variables as controls, we estimate long-run demand elasticities between -0.24 and -0.20, slightly lower in magnitude but consistent with the findings in Coglianese et al. (2017). However, these estimates cannot perfectly control for all demand shocks; therefore there is still some attenuation bias from supply and demand simultaneity.⁴ Due to the attenuation bias, these estimates can be seen as a lower bound on the absolute size of the demand elasticity.

In the second step, we introduce unplanned refinery outages as instruments for supply shocks to address supply and demand simultaneity.⁵ During the sample period, California has seen large refinery outages. For example, due to an explosion at the Torrance refinery in 2015, seven percent of the refining capacity was unavailable for more than a year. In addition, there have been many other refinery outages that were smaller and more frequent. These outages are the result of power outages, operational accidents, or the need to replace parts. Importantly, they are plausibly exogenous and can therefore be used as instruments for supply shocks. Specifically, the instruments are relevant because they directly impact supply, consumers do not anticipate outages, and satisfy the exclusion restriction because they are conditionally uncorrelated with gasoline demand.

However, not all outages are the same. Outages rarely happen at a refinery-wide level. Instead, there are specific units within the refinery that stop working. Each unit has a different production capacity, and their outputs are important inputs to other refining processes. Losing a small refining unit that creates a critical component can have an outsized effect on the whole refining process. We use a detailed data set with information on which refining units stopped working, the dates when the outage

³The level of inventories in California is sufficient to cover 20 days of consumption, approximately (U.S. Energy Information Administration, 2024).

⁴See Section 4 for further discussion.

⁵As far as we can tell, this is the first research paper to use refinery outages to estimate the price elasticity of gasoline demand in California.

occurred, and how much capacity was lost.

The detailed data set allows us to differentiate between planned and unplanned outages. Planned outages are linked to routine maintenance procedures and can be timed to happen in low-demand periods. In contrast, unplanned outages happen following operational accidents like fires, or explosions or due to unforeseen circumstances like high-speed winds or power outages. Due to the haphazard nature of the unplanned outages, these provide exogenous variation on gasoline supply.

Based on these instruments, we calculate two sets of price elasticity estimates for the California gasoline retail market. We estimate the one-month price elasticity of demand to be -0.25 and the long-term elasticity to be -0.60.

Our California-specific estimate is larger in magnitude than some other recent estimates. Kilian and Zhou (2024) find that elasticities can vary substantially across subsamples, reaching very similar magnitudes to our estimate. The higher elasticity in California, when compared to the aggregate estimates for the United States, may be driven by a greater awareness of energy markets and environmental effects, as well as a higher awareness of price changes. Indeed, Kilian and Zhou (2024) find that elasticities are lower if the income is higher. California has above median income, suggesting that our higher elasticity estimate is reflecting unique California-specific effects.

Our paper contributes to the existing literature on demand elasticity estimation. Many previous papers used instruments that either take information from the crude oil market or from changes in taxes. As we learn more about demand estimation, a growing body of literature has found significant limitations in the instruments used. These studies have identified three main challenges: the instruments' relevance, the instruments' conditional correlation with consumer expectations, and the conditional correlation between the instruments and aggregate economic activity (Houthakker et al., 1974; Ramsey et al., 1975; Li et al., 2014; Coglianese et al., 2017; Kilian and Zhou, 2024). In Section 2, we discuss these challenges in further detail. Section 3 discusses how local California outages are robust to these challenges. Section 4 presents our results from the OLS and IV estimates. Section 5 adds an additional discussion and concludes.

2 Previous Gasoline Elasticity Estimation Approaches

Gasoline consumption touches several aspects of the everyday lives of consumers. Therefore, it is not surprising that the price elasticity of demand is used to inform public

policy decisions, business decisions, and economic research as a whole (Hastings, 2004; Yeh and Sperling, 2010; Carter et al., 2011; Knittel and Tanaka, 2021). However, there are challenges when identifying the parameter using observational data.

The main challenge for parameter identification is the simultaneity of supply and demand when using observational data from market outcomes. To address the parameter identification challenge, two components are necessary: exogenous variation that is uncorrelated with unobserved demand components but strongly correlated with gasoline prices (Gandhi and Nevo, 2021); second, a set of controls for demand shifters (Berry and Compiani, 2021). The first set of components has not always been easy to find.

A good instrument for estimating the price elasticity of demand for gasoline should satisfy three key criteria: (i) relevance of the instruments; (ii) conditional independence of the instrument from consumer expectations; (iii) conditional independence of the instrument from aggregate economic activity. In the following subsections, we describe the instruments that have previously successfully addressed these difficulties using various data sets from specific sample periods and geographic locations. Before introducing our California-specific instruments, we provide a short overview.

Table 1 provides an overview of selected studies that have estimated the price elasticity of demand by addressing simultaneity bias; most of them use information from the crude oil market or from tax changes.

2.1 Instruments Based on the Crude Oil Market

One of the first studies to tackle this problem was Houthakker et al. (1974) who use lagged price as an instrument. Another approach is to instrument for price using information from the crude oil market (Ramsey et al., 1975). Crude oil is the primary input into gasoline production (Gary et al., 2007c). Therefore, crude oil prices strongly correlate with gasoline prices through a cost channel. However, crude oil prices are not uncorrelated to gasoline demand. After distilling crude oil, close to 50% of its output is gasoline blending components (U.S. Energy Information Administration, 2022b). This makes crude oil and gasoline prices interconnected through consumers' income and their expectations of future economic activity.

Another set of instruments from the crude oil market is disruptions to crude oil production. Hughes et al. (2008) find that these are not strong predictors of gasoline prices. One possible reason for this finding is that most of the U.S. crude oil production

Table 1: Selected works on IV estimation of the price elasticity of demand for gasoline

Authors	Instrument	Estimated price elasticity	Estimate	St. Errors
Houthakker et al. (1974)	Lagged prices of gasoline	Private demand	-0.24	Not reported
Ramsey et al. (1975)	Relative prices of other distillates	Private demand	-0.65	0.36
Hughes et al. (2008)	Weather-related oil production disruptions	Retail demand	-0.03	0.01
Li et al. (2014)	Gas taxes & oil prices	Retail and tax elasticity of demand	-0.07	0.02
Coglianesi et al. (2017)	Gas taxes	Retail demand	-0.37	0.23
Kilian and Zhou (2024)	Gas taxes	Retail demand	-0.32	0.067
Colina (2024)	BLP instruments	Retail demand	-0.64	0.088

Notes: Estimation procedures that do not take into consideration the simultaneity of prices and quantities risk having elasticity estimates biased towards zero. Previous studies have addressed simultaneity through instruments based on data from the crude oil market, changes in taxes, highly disaggregated observations across time, or quasi-natural experiments.

disruptions that the authors consider are related to weather events in the Gulf of Mexico. The general occurrence of these events follows a seasonal pattern and each specific weather event can be forecasted with more than two weeks of anticipation. The seasonal pattern and the ability to forecast weather events allow refiners to adjust their purchase levels to compensate for disruptions to their supply of inputs. Supply disruptions may therefore have a weak first-stage regression due to anticipatory behavior from refiners. In contrast, our approach of narrowing the geographic area to California allows us to improve on this approach. As we will discuss in more detail, our detailed data on specific types of outages also allow us to strengthen the first stage.

2.2 Other Approaches to Address Endogeneity

Next, we provide a brief overview of other approaches. Using a panel data approach, one strategy is to use tax changes as an instrument for price changes. This strategy has been successfully used and refined in several studies. Davis and Kilian (2011) show that state tax changes may have a noticeable effect on prices at the state level.

Li et al. (2014) show that the salience of the tax implementation may generate endogeneity. For example, two tax changes of the same magnitude publicized differently may produce different consumer reactions through the expectation channel. The authors therefore control for news coverage in anticipation of a tax change.

Coglianesi et al. (2017) account for the anticipatory behavior of consumers to a tax that is being implemented. They include leads and lags of retail prices to control for anticipatory and forward-looking behavior.

Finally, Colina (2024), uses a natural quasi-experiment to address simultaneity concerns and uses gas station level characteristics to control for local demand shifts; he instruments station characteristics with BLP instruments to address endogeneity between observed and unobserved station attributes and prices as in Berry et al. (1995) and Davis (2006).

There are alternative approaches that do not use an instrumental variable. Levin et al. (2017) use disaggregated daily panel data for 243 U.S. cities from 2006 to 2009 and include city and day-of-sample fixed effects to control for supply and demand simultaneity.

Our focus is on specifically studying the California market. This market is relevant since it is the largest state in the United States in terms of gasoline revenue, and economic activity (U.S. Energy Information Administration, 2022a; State of California,

2024). Perhaps more importantly for our setting, California is largely segmented from the rest of the United States due to its environmental regulations.

3 Using Refinery Outages as Exogenous Price Shocks

To address endogeneity, we follow an approach similar to Hughes et al. (2008), who use supply disruptions in Venezuela, Iraq, and the United States. Focusing on California, we propose refinery outages as a new set of instruments to estimate the price elasticity of demand. This set of instruments solves the three main documented issues mentioned in Section 2. Due to the institutional arrangement of gasoline production, a refinery outage reduces available installed capacity and increases costs for producers.

Gasoline production follows a multistep process. Refineries produce gasoline blending components, and these blending components are mixed to achieve specific performance properties (Gary et al., 2007c). The blend is transported to a city terminal, mostly by pipelines, and then mixed with ethanol to produce finished gasoline. The finished gasoline is then distributed within the city to the gasoline stations (Borenstein et al., 1992).

Different markets need different performance properties of their finished gasoline. Refineries achieve these properties and optimize their configuration by choosing, amongst other things, which refinery units to install and how to connect them together. In this process, the output of one refining unit is used as input to another refining unit.

This combination of products is meant to maximize the refiner’s profits, conditional on achieving performance requirements. This configuration results in refining units connected in a complex multistage process (also see additional discussion in Section 3.3.).

Heat, pressure, catalysts, and other chemicals are used throughout the refining stages. Because of the nature of these processes, every so often, refineries need to stop operating one of the refining units for repair. Sometimes, a specific stage of the refining process suffers from an accident or a malfunction. These incidents lead to unplanned stops in the operation of a refining unit. We refer to the loss in refining capacity in a specific refining unit as an outage.

Several circumstances lead to unplanned outages. Examples include an unplanned power outage; an unplanned flaring event;⁶ a malfunction in the refining unit caused

⁶A flaring event happens when excess hydrocarbons are burned rather than released straight into

by a leak, a crack, or loss in pressure or a fire; unexpected high winds; and unexpected malfunction of the crude pipeline that supplies the refinery.

3.1 Data and Summary Statistics

We use a detailed data set from Bloomberg, where we observe the output of each individual refinery unit within a refinery. We need this granularity in the data in order to ensure a strong first stage for the IV estimation and also more precise overall estimates. For explanatory purposes, we group the relevant refinery units into Primary Conversion Units and Secondary Conversion units. The former are involved in the initial processing and breakdown of crude oil into more useful components and produce a sizeable share of the finished output and sulfur-reducing units that produce specific distillates. The latter improve the quality of the components by reducing pollutants or increasing octane levels; they produce smaller amounts of throughput but are essential to achieve regulatory requirements; therefore, a small outage has an outsized effect. We discuss the refining process and how it relates to our instrument in more detail in Section 3.3.

Fuel sales data are from the California Department of Tax and Fee Administration; we source retail prices, inventories, imports, capacity utilization, and WTI prices from the U.S. Energy Information Administration. The outage data set comes from Bloomberg; these data are available starting in 2011 as part of the OIL <GO> market monitor. Bloomberg aggregates data from national public sources like the Occupational Safety and Health Administration (OSHA) and the Environmental Protection Agency (EPA);⁷ regional agencies like the South Coast Air Quality Management District or the Bay Area Air Quality Management District;⁸ and private market research firms that specialize in fuel markets.

Our sample period runs from January 2011 until March 2023 with observations at the state-monthly level. Table 2 reports summary statistics of the variables used in the estimation. The upper part of the table reports statistics for control variables, while

the atmosphere. Plants usually inform local authorities about planned flaring. But sometimes pressure builds up to dangerous levels, resulting in an unplanned flaring event. For more information, see Gary et al. (2007b).

⁷Refineries are mandated to report work place injuries and events that result in exceeding emission limits (Occupational Safety and Health Administration, 1970; Environmental Protection Agency, 1984).

⁸Refineries are required to report flaring events to air quality management districts within a specified time as well as their root causes (California State Legislature, 1976).

the bottom part of the table shows statistics for outages, specifically the change in capacity due to an outage. We note that sales are quite stable, suggesting that the size of shocks to demand is not large when measured as a percentage of sales. Importantly, the standard deviations of capacity lost due to outages are similar in magnitude, in some units, to the standard deviation of sales, suggesting a potentially large impact of outages on supply. We also note that the mean of inventories is quite large so that sales can be smoothed. This link means that inventories may be a relevant control variable when we estimate elasticities. As we discuss more later, we also see that capacity utilization has an average of 84.5% and a maximum of 97.7%, implying potentially large impacts of outages.

Figure 1 plots the dates and sizes of various types of outages over the sample period. We note that with the exception of 2020 (probably related to the effects of the COVID 19 pandemic), outages occur frequently and throughout the sample. There are no visible trends in outages and there also do not appear to be large and prolonged differences in variability across different subsamples.

Table 2: Summary statistics

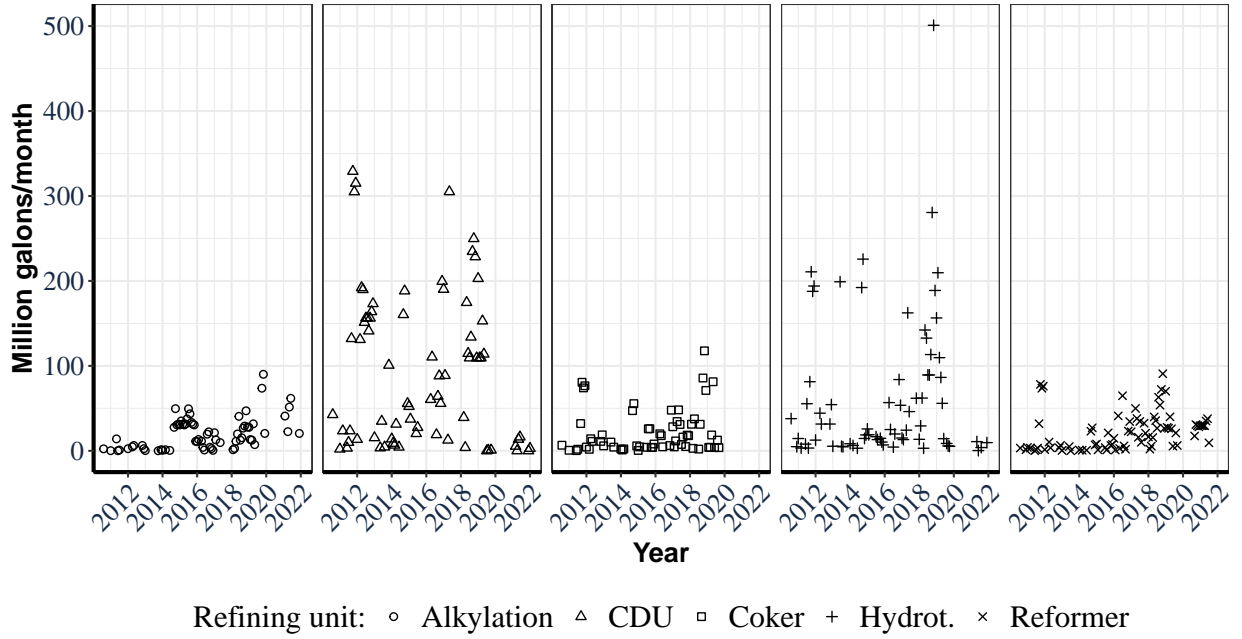
Variable	N	Mean	SD	Min	Max
Date (monthly)	139	.	.	Jan 2011	Aug 2022
Sales (MMgals/month)	139	1,220.7	98.4	713.6	1,514.2
Retail price (USD/gal)	139	3.6	0.7	2.4	6.2
Inventories (MMgals)	139	753.9	68.7	617.4	952.6
Imports (MMgals/month)	139	946.8	209.0	419.4	1,641.4
Capacity util. (%)	139	84.5	6.7	60.5	97.7
WTI (USD/barrel)	139	70.0	23.7	16.5	114.8
Outages (MMgals/month)					
:: Alkylation unit	139	9.7	16.6	0.0	90.1
:: Crude distilling unit	139	49.6	80.4	0.0	329.0
:: Coker unit	139	9.9	20.8	0.0	117.7
:: Hydrotreating unit	139	33.5	70.5	0.0	500.9
:: Reformer unit	139	29.4	41.6	0.0	233.2

Observations are at the state-month level. MM denotes millions

Sources: California Department of Tax and Fee Administration, US Energy Information Administration, and Bloomberg L.P. (2022)

Notes: Note that sales are quite stable, suggesting that the size of shocks to demand is not large when measured as a percentage of sales. Additionally, the standard deviations of capacity lost due to outages are similar in magnitude to the standard deviation of sales, suggesting a potentially large impact of outages on supply.

Figure 1: Refining capacity lost for unplanned outages in different refining units



Source: Authors' calculations with data from Bloomberg L.P. (2022)

Notes: Unplanned outages are frequent, and their magnitude varies across time. Due to the interconnections of the units, a small outage in the secondary conversion units like the Alkylation, Hydrotreating, or Reformer units can have an outsized impact on the ability of refineries to produce CARB-grade gasoline. Therefore, measuring outages at the refining-unit level, as opposed to aggregate terms, is important due to the potential attenuation.

3.2 Characteristics of Outages

In Section 2 we summarized three documented concerns about instruments for the price of gasoline. We will now relate those concerns to gasoline outages.

The first concern is the relevance of the instruments. Outages in specific refining units result in increased costs for the refiner. As one of the components of the optimized blend is missing, refiners either: source the missing component from an outside supplier, which is costly (American Petroleum Institute, 2013); reduce total output while maintaining the optimal blend, which increases inventory costs due to unused components (U.S. Energy Information Administration, 2007); or produce a suboptimal blend subject to achieving performance requirements (Valentine and Josefson, 2017).

In each case, there is an increase in operational costs and the possibility of a reduced output, which will impact market prices.

A second concern is the conditional independence of the instrument from consumer expectations. For example, if consumers expect higher prices in the future, they can buy gasoline before a price increase happens. However, due to the unexpected nature of accidents, consumers cannot engage in anticipatory buying before a specific operational problem occurs when dealing with refining accidents.

Third, the salience of an event affects consumer expectations of the magnitude of the impact. Regarding refinery outages, news agencies cannot cover an accident in anticipation of it happening. Therefore, the level of coverage cannot affect expectations before the incident. However, once the unplanned outage occurs, it may be covered by the news. Most unplanned outages in the United States are reported to the Occupational Safety and Health Administration (OSHA). Large, unplanned outages are reported in specialized news outlets. To account for possible salience effects, we will include lagged retail sales.

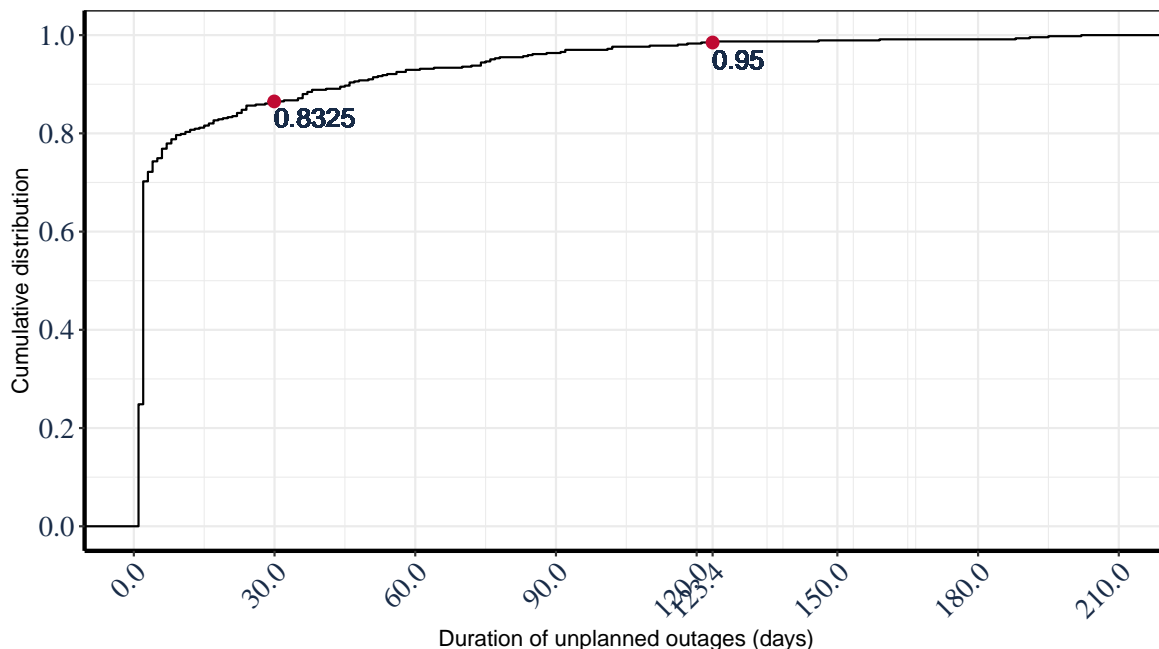
Another possibility is that consumers may expect the duration of the outage to be long-lived, and they react differently from a regular price change. However, the vast majority of outages are resolved in a short period of time; 83% of them are solved in less than a month. Therefore, outages are unlikely to elicit changes in long-term consumer behavior, such as buying a more fuel-efficient vehicle, moving closer to their job to reduce commute distances, or choosing alternative modes of transportation. Figure 2 provides additional detail on the duration of outages.

Another concern is the conditional independence of the instrument from aggregate economic activity. It is possible that accidents in a refinery are more likely to occur when the units are running at full capacity. This would violate the assumption that accidents happen randomly. To account for the possible systematic variation in accidents, we control for the percentage of operating capacity at which the refineries operate.

California's gasoline market faces a set of regulatory and infrastructure constraints that make it a partially isolated market from the rest of the 47 contiguous states and the District of Columbia. This unique setting causes the proposed instrument to have a strong first-stage regression because it is relatively difficult to substitute for the loss of local production capacity.

Due to environmental regulations, Californians consume the cleanest gasoline in the United States. However, within the United States, only California refineries produce this blend for the majority of our sample period (Pyziur, 2016). Therefore, when there

Figure 2: Cumulative Distribution of the duration of outages



Source: Authors' calculations with data from Bloomberg L.P. (2022)

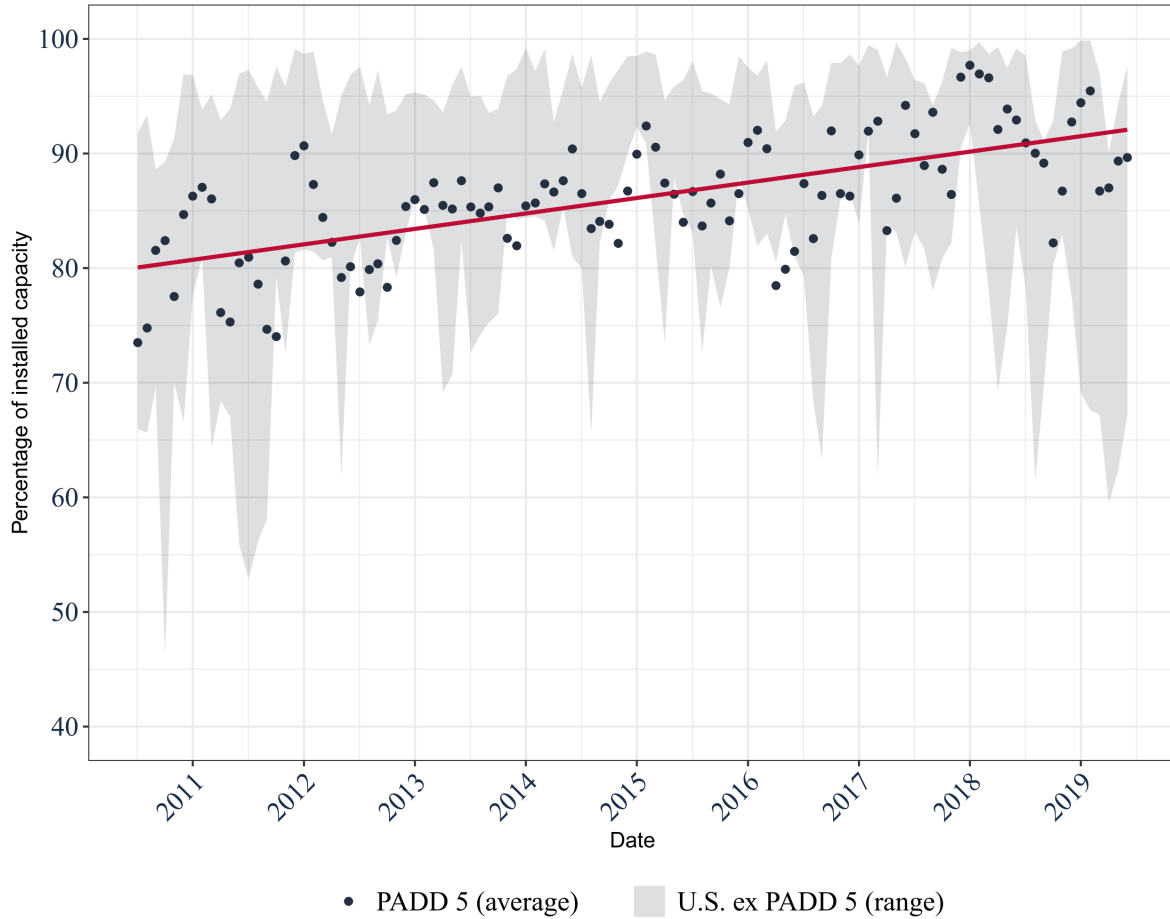
Notes: Most unplanned outages are short-lived. The red dots highlight that 83.25% of the outages last less than a month, while 95% of outages last less than four months.

is an outage and refining capacity is lost, wholesalers cannot substitute local production with refined products produced elsewhere in the United States

An alternative to sourcing refined products from outside the state is to substitute with products made within the state. However, refineries in California have been operating increasingly close to their installed capacity, making it harder to increase regional production in response to a local outage. Figure 4 shows the throughput of refineries in PADD 5 as a percentage of the installed available capacity.⁹ PADD 5, encompasses California and other states along the West of the United States, of which California is by far the largest. Additionally, the pipeline transportation infrastructure within California is limited, and the major refining centers in the north and south are not connected to each other by pipeline. This lack of connectivity makes it challenging to move fuels between

⁹A Petroleum Administration for Defense Districts (PADD) is a geographic division of crude and fuels markets established during World War 2 to ration gasoline consumption. Today, market participants use the division to analyze regional trends (U.S. Energy Information Administration, 2012).

Figure 4: Utilization of refining capacity



Source: Authors' calculations with data from the U.S. Energy Information Administration

Notes: Refineries in PADD 5 are running ever closer to the limit of their installed capacity as time progresses. The gray area shows the range of the throughput of refineries in PADD 5 as a percentage of the installed available capacity. The dark dots indicate the average level at which refineries operated during a given month, while the red line shows a time trend.

regions, especially to and from inland areas (U.S. Energy Information Administration, 2015).

A second alternative to substitute for the loss of local production capacity is to import refined products from outside the United States into California. Two other countries produce the refined products needed to achieve the required blend for California: Singapore and South Korea. According to California Energy Commission (2020a), the minimum number of days needed for a vessel to reach California and be fully unloaded

is 19 and 13 days, respectively. However, weather conditions across the Pacific Ocean and local logistics constraints at California’s ports can extend this timeframe. The lag between an outage and when imported products may arrive creates a temporary contraction in supply.

One of the possible local logistics constraints when importing distillates is scheduling their transportation once the vessel is unloaded. Schremp (2015) explains that only two sets of pipelines transport products in California. The first one starts in San Francisco and finishes in Reno passing through several refineries along the way. The second set of pipelines starts in Los Angeles and forks to Las Vegas and Phoenix. The limited number of pipelines leads to a strict scheduling system in which users need to buy space and time in the pipeline in advance. Then, it is likely that an importer would have to buy the already reserved pipeline capacity to move imported products (California Energy Commission, 2020a). This would increase transportation costs, which may ultimately be passed on to the retail price, resulting in a strong first-stage regression.

3.3 Different Types of Outages

Gasoline production is a multistage process. Refinery units are involved in different stages in this process, each with a different installed capacity. Therefore, an outage of 10 million gallons per month will have a very different impact depending on the stage of the process in which it occurs. To improve the power of the instruments, we differentiate by the source of the outage at the refinery unit level.

Not all outages are the same; simply aggregating capacity loss would result in noisier estimates. For example, the crude distilling unit (CDU) is usually the largest refining unit within a refinery, and all refineries have one. This unit is the first one to receive the crude oil at the refinery; it then applies heat and produces the first batch of distillates. A loss of 10 million gallons per month will not have a big impact since its output would not be a limiting factor in producing California-compliant gasoline.¹⁰ In contrast, the alkylation unit has a smaller capacity as it is used mainly at the end of the refining process and produces specific distillates that make the refining blend California-compliant. The alkylation unit is expensive to build and install; only select refiners have one. Losing 10 million gallons per month would create a large disruption in the market. Appendix Section A discusses the differences in refinery units in more

¹⁰The installed refining capacity in California is approximately 51.3 million barrels per month (California Energy Commission, 2023).

detail.

3.4 Statistical Diagnostics for IV

We are now ready to start our estimation. The first step is to estimate the first stage, where we regress the log of the monthly retail gasoline price on various measures of outages and other explanatory variables. We also control for seasonal monthly fixed effects. A necessary condition for our estimation to be valid is a choice of instruments that do not suffer from the weak-instrument problem. We choose three different specifications, including different sets of control variables corresponding to our main results, which we report in the next section.

Table 3 reports the results. We find that the instruments are strong jointly; the conventional instrument threshold for the F-statistic is 10 and this is cleared in all three specifications by a wide margin. As expected, we find that the coefficients on outages are, in most cases, positive and jointly statistically significant. We also find that the statistically significant coefficients are positive, supporting the interpretation that these identify supply shocks. Specifications 1 and 2, in particular, provide evidence that alkylation and reformer unit outages are strongly correlated with prices. As we mentioned earlier, past capacity utilization can be correlated with future outages. Indeed, in specification 3, inclusion of this variable marginally reduces statistical significance of individual outages, but the overall F-stat remains high. Overall, we prefer specification 3 since it results in a strong first stage regression and controls for the possibility of persistent demand shocks affecting capacity utilization.

One of the advantages of using more detailed data on outages is our ability to account for the different effects of the various types of outages. As discussed, the volume and unit cost of production across components used in the blending process vary. This approach, therefore results in stronger overall instruments.

Finally, one might imagine that unplanned outages are, in fact, related to market conditions in a way that would decrease their validity. We have performed an additional test to address this concern. Specifically, in a partial adjustment model, which is what we are estimating, we need to check that instruments are lag-exogenous. We check this by running Granger causality tests of the unplanned outages on lagged prices. The results are reported in Appendix Table 6. We find that the instruments are not predictable by past market conditions and are, therefore, lag-exogeneous, supporting our estimation approach.

Table 3: First-stage estimation results.

<i>Dep. Variable: log Retail Prices</i>	First stage regression		
	<i>Specification 1</i>	<i>Specification 2</i>	<i>Specification 3</i>
	(1)	(2)	(3)
A) Coefficient estimates			
Alkylation unit outage	1.335**	1.218**	0.861
:: s.e.	(0.583)	(0.557)	(0.594)
CDU outage	-0.189	-0.169	-0.118
:: s.e.	(0.138)	(0.142)	(0.153)
Coker unit outage	0.330	0.135	0.310
:: s.e.	(0.616)	(0.585)	(0.680)
Hydrotreating unit outage	-0.130	-0.104	-0.144
:: s.e.	(0.205)	(0.202)	(0.219)
Reformer unit outage	1.480***	1.432***	1.192*
:: s.e.	(0.505)	(0.514)	(0.636)
Constant	11.681***	14.391***	19.806***
:: s.e.	(2.223)	(3.242)	(3.711)
Seasonal Month fixed effects	Yes	Yes	Yes
log Retail sales, lagged	Yes	Yes	Yes
log WTI, lagged	Yes	Yes	Yes
log Inventories, lagged		Yes	Yes
log Imports, lagged		Yes	Yes
log Capacity util., lagged			Yes
B) Model stats.			
Observations	138	138	138
R-squared	0.745	0.752	0.761
F stat.	13.92	16.29	20.81

Newey-West standard errors in parentheses.

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

Outage units are in billion barrels per month to improve legibility.

Notes: We find that the instruments are strong jointly; clearing the conventional instrument threshold of 10, for the F-statistic, by a wide margin. We also find that the statistically significant coefficients are positive, supporting the interpretation that these identify supply shocks.

4 Estimating Short and Long Run Elasticity

This chapter examines the estimation of short- and long-run gasoline demand elasticities using a dynamic partial adjustment model. By leveraging refinery outages as exogenous instruments, the analysis addresses simultaneity bias and demonstrates that estimation procedures that fail to account for simultaneity tend to underestimate elasticities. The results reveal that California’s gasoline consumers exhibit a short-run elasticity of -0.25 and a long-run elasticity of -0.60, reflecting both immediate and persistent responses to price changes.

4.1 Partial Adjustment Model

Following Balestra and Nerlove (1966); Houthakker et al. (1974); Paul et al. (2009); Taylor and Houthakker (2009) we use a traditional dynamic log-log specification that relates the quantities consumed to the prices observed on the market. We use equation (1) to estimate the parameters of interest: the one-month price elasticity of demand, β_1 , and the long-run elasticity of demand implied by the partial adjustment model.

We use additional lagged covariates to control for demand shocks. In different specifications, we control for aggregate and speculative demand shocks. We include the level of inventories, the level of imports, the refinery utilization rate, and the price of crude oil.¹¹ To control for dynamic adjustments, we estimate different model specifications given by lagged covariates in the following model,

$$q_t = \beta_0 + \beta_1 p_t + \sum_{s=1}^{11} m_t^s + \sum_{\ell=1}^L \left(\rho_{\ell} q_{t-\ell} + \lambda_{\ell} wti_{t-\ell} + \gamma_{\ell} inv_{t-\ell} + \delta_{\ell} imp_{t-\ell} + \theta_{\ell} util_{t-\ell} \right) + \varepsilon_t, \quad (1)$$

where $\beta, \gamma, \delta, \rho, \lambda, \theta$ are parameters, and m_t^s are monthly seasonal fixed effects. The variable q_t is the log of the level of gallons of gasoline sold in month t in millions. The variable p_t is the log of the monthly retail price of regular gasoline; inv_t is the log of the level of blending components in inventories; imp_t is the log of imports into California of motor gas blending components; wti_t is the log spot price of a barrel of West Texas Intermediate crude oil (WTI). The model in equation (1) includes the lagged dependent variable $q_{t-\ell}$ to control for partial adjustments, and the monthly fixed effects m_t^s to account for demand seasonality. ε_t is a shock of unobserved time-varying

¹¹As has been emphasized by Kilian (2009) and Kilian and Zhou (2024), the raw WTI price is likely correlated with aggregate demand shocks.

factors with mean zero. We use the Bayesian information criterion (BIC) to choose specifications with $L = 1$.

An unexpected shock to demand may take time to propagate. In theory, there are two reasons why there can be partial adjustments. There may be autocorrelation in the errors; alternatively, it is possible that the shock comes from independent and identically distributed distribution but there is a partial adjustment in the quantity, e.g., due to market frictions or for institutional reasons. From an econometric perspective, these two are observationally equivalent. Our estimation approach can be used in either case. We also account for any remaining unmodeled autocorrelation using Newey-West standard errors.

Having estimated the short-term elasticity, we can then calculate the long-term price elasticity of demand, i.e. the response in quantity to a permanent increase in price, assuming that all other variables remain fixed. The response of the quantity after a sufficient amount of time has passed - the long-run elasticity - is then given by $\beta^* \equiv \beta_1 / (1 - \rho_1)$.¹²

4.2 Results

The model parameter estimates described in equation (1) are presented in Table 4. There are three model specifications, each with different lagged control variables. The estimates of each model specification are presented side by side, comparing two estimation techniques, ordinary least squares (OLS) and instrumental variables (IV).

We first present OLS estimates with Newey and West (1987) standard errors; we label these columns as *OLS*. However, this estimation procedure does not address the simultaneity of supply and demand and will, therefore, suffer from simultaneity bias, which results in estimated coefficients that are biased towards zero.¹³ Thus, the OLS estimates provide a lower bound for the elasticity estimates. In addition, if we find that OLS coefficients are lower in absolute value than IV-estimates this represents evidence in support of the IV strategy that we employ. The columns labeled ‘*IV*’ present

¹²Consider the model in Equation 1. Under the long-run assumptions mentioned in the main text, the variables in the model take the same value in period t and $t - 1$. We proceed to calculate the total derivative with respect to price, and then solve for $\Delta q / \Delta P$ to obtain β^* .

¹³Let ε^d be unobserved demand shocks. From a structural supply and demand model, we know that increases in demand, all else constant, cause the equilibrium price of gasoline, P^* , to increase, then $cov(P^*, \varepsilon^d) > 0$. From economic theory, we expect β_1 in Equation 1 to be ≤ 0 . From a univariate OLS regression that fails to control for econometric endogeneity, the estimated coefficient will be $\hat{\beta}_1 = \beta_1 + cov(P^*, \varepsilon^d) / var(\varepsilon^d)$, biasing the estimate towards zero.

results from the second estimation technique which uses an instrumental variable estimation procedure to control for simultaneity. We implement a two-stage least squares estimating procedure.

Table 4: Elasticity estimates under various specifications

Dep. Var.: log Retail Sales	<i>Specification 1</i>		<i>Specification 2</i>		<i>Specification 3</i>	
	<i>OLS</i>	<i>IV</i>	<i>OLS</i>	<i>IV</i>	<i>OLS</i>	<i>IV</i>
	(1)	(2)	(3)	(4)	(5)	(6)
A) Long term elasticity (β^*)	-0.204*	-0.824**	-0.235**	-0.797***	-0.230**	-0.603**
:: s.e.	(0.118)	(0.330)	(0.118)	(0.325)	(0.112)	(0.270)
B) Coefficient estimates						
log Retail price (β_1)	-0.041	-0.254	-0.044	-0.242	-0.047	-0.251
:: s.e.	(0.029)	(0.309)	(0.027)	(0.324)	(0.040)	(0.308)
log Retail sales, lagged (ρ_1)	0.797***	0.690***	0.809***	0.695***	0.794**	0.601***
:: s.e.	(0.084)	(0.158)	(0.078)	(0.186)	(0.172)	(0.223)
Constant (β_0)	4.182**	6.315**	4.790**	7.496	5.103*	9.401
:: s.e.	(1.817)	(3.142)	(1.602)	(5.155)	(3.011)	(5.730)
Seasonal month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
log WTI, lagged	Yes	Yes	Yes	Yes	Yes	Yes
log Inventories, lagged			Yes	Yes	Yes	Yes
log Imports, lagged			Yes	Yes	Yes	Yes
log Capacity util., lagged					Yes	Yes
Observations	138	138	138	138	138	138
R-squared	0.717	0.664	0.723	0.680	0.723	0.681

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

Note: HAC standard errors in parenthesis. IV estimation is done through 2SLS

Note: Long term elasticity $\beta^* \equiv \beta_1 / (1 - \rho_1)$.

Notes: Our parameters of interest are β_1 and β^* . Our preferred specification is # 3 since it estimates the parameters of interest with the most precision.

Table 4 Panel B reports estimates of the model parameters. The price coefficients in all specifications are negative but not statistically significant in the case of OLS. The corresponding IV estimates are approximately six times as large; the somewhat low level of statistical significance is most likely driven by the relatively small sample size. The reason for this is that the data on outages start only in 2011.¹⁴

The fact that the OLS estimates are much smaller in magnitude compared to their IV counterparts is consistent with the well-documented case that estimates of demand elasticity that control poorly for simultaneity tend to be biased towards zero (see Davis and Kilian (2011); Coglianese et al. (2017); Kilian and Zhou (2024) for further details). The change in magnitude in the elasticity estimates validates the assumption that the

¹⁴In an earlier version of the paper, using a shorter sample period, power and statistical significance were indeed lower.

instruments control for simultaneity bias.

Panel A reports long-run elasticities (see definition above). We find that all IV estimates for the long-term price elasticity of demand are statistically significant at the 5% level. Our preferred specification is #3 – it is likely that the additional controls in that specification are effective at reducing a hypothetical bias from persistent demand shocks in the estimate. We note in particular that, when additional controls are added relative to specification #2, the point estimate drops quite a bit and the standard error is reduced by almost 20%.

One may be also interested in the length of time until prices reach the long-run equilibrium in response to a supply shock. This can be done, for example, using impulse response analysis in the local projection framework of Jordà (2005). A careful analysis using this methodology would require a longer time series to get sufficient power compared to the partial adjustment model we use. Nevertheless, in Section C of the Appendix, we report preliminary results of such an LP approach. These results suggest that the long-run equilibrium is reached eleven months after a shock.

4.3 Discussion and Implications

We find that the one-month price elasticity is -0.25 and that the long-term elasticity is -0.60. The one-month price elasticity estimates, for our sample period, suggest that California consumers could be more price-sensitive than the average U.S. consumer as reported by Kilian and Zhou (2024), since our central estimate falls well outside the 95% confidence interval reported by the authors. The magnitude of the long-term elasticity is consistent with cross-sectional heterogeneity documented by Kilian and Zhou (2024) and the aggregate, long-term elasticity estimates of Colina (2024). Given the large interest in energy and environmental policies in California, our estimates can inform optimal policy. Similarly, knowing that consumers are quite price-sensitive could open the door to carbon tax policies that are more palatable to elected officials (Parry et al., 2022).

Gasoline consumption touches on several aspects of consumers’ everyday lives, and knowing consumers’ price sensitivity is essential to the design of public policies and informing business decisions. Our estimates show substantial consumer price-sensitivity. Based on this evidence, policymakers, investors, and researchers may re-evaluate the implications of consumers’ reactions to price changes.

For example, Holland et al. (2009) estimate the welfare costs of implementing the

Low Carbon Fuel Standard (LCFS) and simulate different scenarios based on different supply and demand price elasticities values. Our estimates exceed the range of values for which they simulate welfare outcomes. Still, based on their argument, the conclusion follows that welfare costs of adjusting to the new standards would be lower since consumers are more responsive to the implied subsidies of the LCFS.

5 Conclusion

This study introduces refinery outages as novel instruments to estimate gasoline demand elasticity in California, a market distinguished by its unique regulatory environment and infrastructure constraints. The granularity of the outage data enables precise identification of short- and long-term elasticities, with estimates of -0.25 and -0.60, respectively. These findings indicate that California consumers are more price-sensitive than previously believed in commonly cited papers. By controlling for simultaneity and leveraging detailed information on unplanned outages, this analysis addresses key econometric challenges in demand estimation.

The robustness of the proposed methodology lies in its focus on exogenous supply shocks and the specific characteristics of the California gasoline market. We differentiate between types of outages and their impact on the production of California-compliant gasoline, ensuring strong first-stage regressions and credible instruments. These findings underscore the potential of the approach to be extended to other geographic regions or industries that face similar regulatory environments. Such an approach could help identify market-specific elasticities and further explore heterogeneity across markets, contributing to our understanding of drivers of variation in elasticities, also see Gafarov et al. (2023) and Gaarder et al. (2024).

These results have important policy implications. The demonstrated sensitivity of consumers to gasoline price changes suggests opportunities for designing more effective environmental and taxation policies. For instance, the higher elasticities imply that carbon pricing or subsidies for cleaner alternatives could induce significant behavioral changes, aiding in the transition to lower greenhouse gas emissions. This analysis thus contributes not only to refining elasticity estimates but also to informing the broader discourse on sustainable energy policy and consumer welfare as in Colina (2024).

References

- AMERICAN PETROLEUM INSTITUTE (2013): “Economic and Supply Impacts of a Reduced Cap on Gasoline Sulfur Content,” Tech. rep., Turner, Mason and Company.
- ANGRIST, J. D. AND J.-S. PISCHKE (2009): “IV and Causality,” in *Mostly Harmless Econometrics: An Empiricist’s Companion*, 115–137.
- BALESTRA, P. AND M. NERLOVE (1966): “Pooling Cross Section and Time Series Data in the Estimation of a Dynamic Model: The Demand for Natural Gas,” *Econometrica*, 34, 585–612.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.
- BERRY, S. T. AND G. COMPIANI (2021): “Empirical Models of Industry Dynamics with Endogenous Market Structure,” *Annual Review of Economics*, 13, 309–334.
- BLOOMBERG L.P. (2022): “Refinery outage data from Jan 2011 to Aug 2022. Retrieved from Bloomberg database,” .
- BORENSTEIN, S., A. C. CAMERON, AND R. GILBERT (1992): “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?” *The Quarterly Journal of Economics*, 112, 305–339.
- CALIFORNIA AIR RESOURCES BOARD (2014): “Brief Summary: CARB Phase 3 Gasoline Specifications and Test Methods,” Series: Fuels Enforcement Program. Fuels Specifications and Test Methods. Gasoline Specifications and Test Methods URL: <https://ww2.arb.ca.gov/sites/default/files/2020-03/gasspecs.pdf>.
- CALIFORNIA ENERGY COMMISSION (2020a): “How Petroleum Products Move,” Tech. Rep. March 2020, California Energy Commission, California, series: Petroleum Watch.
- (2020b): “Refining Operations Under Decreased Demand,” Tech. Rep. May 2020, California Energy Commission, California, series: Petroleum Watch. URL: https://www.energy.ca.gov/sites/default/files/2020-05/2020-05_Petroleum_Watch.pdf.

- (2023): “California’s Oil Refineries,” Series: Energy Almanac. URL: <https://www.energy.ca.gov/data-reports/energy-almanac/californias-petroleum-market/californias-oil-refineries>.
- CALIFORNIA STATE LEGISLATURE (1976): “California Health and Safety Code.Division 26, Air Resources,” Tech. rep., California.
- CARTER, C. A., G. C. RAUSSER, AND A. SMITH (2011): “Commodity Booms and Busts,” *Annual Review of Resource Economics*, 3, 87–118.
- COGLIANESE, J., L. W. DAVIS, L. KILIAN, AND J. H. STOCK (2017): “Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand,” *Journal of Applied Econometrics*, 15, 1–15.
- COLINA, A. R. (2024): “Mexico’s Gasoline Markets: A Study of the Consumer Welfare Impacts of Price Liberalization and Permit Curtailment,” Repository: Social Science Research Network.
- DAVIS, L. W. AND L. KILIAN (2011): “Estimating the Effect of a Gasoline Tax on Carbon Emissions,” *Journal of Applied Econometrics*, 26, 1187–1214.
- DAVIS, P. (2006): “Spatial competition in retail markets: Movie theaters,” *RAND Journal of Economics*, 37, 964–982.
- DEATON, A. (2010): “Instruments, Randomization, and Learning about Development,” *Journal of Economic Literature*, 48, 424–455.
- ENVIRONMENTAL PROTECTION AGENCY (1984): “Emergency Planning and Community Right-to-Know Act (EPCRA),” Series: Laws & Regularions. URL: <https://www.epa.gov/epcra>.
- GAARDER, I., L. H. D. FRAHAN, M. MOGSTAD, A. TORGOVITSKY, AND O. VOLPE (2024): “Linear Supply and Demand in Heterogeneous Markets,” Repository: Yale Department of Economics.
- GAFAROV, B., T. GONG, AND J. HILSCHER (2023): “Secular Rise and Pro-Cyclical Variation in Markups: Evidence from US Grocery Stores,” Repository: Social Science Research Network.

- GANDHI, A. AND A. NEVO (2021): “Empirical Models of Demand and Supply in Differentiated Products Industries,” Tech. rep., National Bureau of Economic Research, volume: 4 Issue: 1.
- GARY, J. H., G. E. HANDWERK, AND M. J. KAISER (2007a): “Alkylation and Polymerization,” in *Petroleum Refining. Technology and Economics*, Boca Raton, FL: Taylor & Francis Group, LLC, 231–255, 5 ed.
- (2007b): “Control of Atmospheric Pollution,” in *Petroleum Refining. Technology and Economics*, Boca Raton, FL: Taylor & Francis Group, LLC, 293, 5 ed.
- (2007c): “Industry Characteristics,” in *Petroleum Refining. Technology and Economics*, Boca Raton, FL: Taylor & Francis Group, LLC, 18–19, 5 ed.
- (2007d): “Introduction,” in *Petroleum Refining. Technology and Economics*, Boca Raton, FL: Taylor & Francis Group, LLC, 293, 5 ed.
- HASTINGS, J. S. (2004): “Vertical Relationships and Competition in Retail Gasoline Markets : Empirical Evidence from Contract Changes in Southern California,” *American Economic Review*, 94, 317–328.
- HOLLAND, S. P., J. E. HUGHES, AND C. R. KNITTEL (2009): “Greenhouse Gas Reductions under Low Carbon Fuel Standards?” *American Economic Journal: Economic Policy*, 1, 106–46.
- HOUTHAKKER, H. S., P. K. VERLEGER, AND D. P. SHEEHAN (1974): “Dynamic Demand Analyses for Gasoline and Residential Electricity,” *American Journal of Agricultural Economics*, 56, 412–418.
- HUGHES, J. E., C. R. KNITTEL, AND D. SPERLING (2008): “Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand.” *Energy Journal*, 29, 113–134.
- JORDÀ, (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- KILIAN, L. (2009): “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 99, 1053–1069.

- KILIAN, L. AND X. ZHOU (2024): “Heterogeneity in the Pass-through from Oil to Gasoline prices: A New Instrument for Estimating the Price Elasticity of Gasoline Demand,” *Journal of Public Economics*, 232, 105099.
- KNITTEL, C. R. AND S. TANAKA (2021): “Fuel Economy and the Price of Gasoline: Evidence from Fueling-level Micro Data,” *Journal of Public Economics*, 202, 104496.
- LARSON, B. (2018): “U.S. Gasoline Requirements Map,” Series: American Petroleum Institute Policy Resources. URL: <https://www.api.org/~media/files/policy/fuels-and-renewables/2016-oct-rfs/us-fuel-requirements/us-gasoline-requirements-map.pdf>.
- LEAMER, E. E. (1985): “Vector autoregressions for causal inference?” *Carnegie-Rochester Conference Series on Public Policy*, 22, 255–304.
- LEVIN, L., M. S. LEWIS, AND F. A. WOLAK (2017): “High Frequency Evidence on the Demand for Gasoline,” *American Economic Journal: Economic Policy*, 9, 314–347.
- LI, S., J. LINN, AND E. MUEHLEGGGER (2014): “Gasoline Taxes and Consumer Behavior,” *American Economic Journal: Economic Policy*, 6, 302–342.
- LIU, W. (2014): “Modeling gasoline demand in the United States: A flexible semiparametric approach,” *Energy Economics*, 45, 244–253.
- NEWKEY, W. K. AND K. D. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- OCCUPATIONAL SAFETY AND HEALTH ADMINISTRATION (1970): “Occupational Safety and Health Act of 1970,” URL: <https://www.osha.gov/laws-regs/oshact/completeoshact>.
- PARRY, I. W., S. BLACK, AND K. ZHUNUSSOVA (2022): “Carbon Taxes or Emissions Trading Systems? Instrument Choice and Design,” Tech. Rep. Climate Note 2022/006, International Monetary Fund.
- PAUL, A. C., E. C. MYERS, AND K. L. PALMER (2009): “A Partial Adjustment Model of U.S. Electricity Demand by Region, Season, and Sector,” .

- PETERSON, J. R. (1996): “Alkylate is Key for Cleaner Burning Gasoline,” Tech. rep., STRATCO.
- PYZIUR, M. (2016): “Understanding California’s High Transportation Fuel Prices,” Tech. rep., Energy Policy Research Foundation, issue: November.
- RAMSEY, J., R. RASCHE, AND B. ALLEN (1975): “An Analysis of the Private and Commercial Demand for Gasoline,” *The Review of Economics and Statistics*, 57, 502–507.
- SCHREMP, G. (2015): “California Transportation of Petroleum,” Tech. rep., California Energy Commission, Crockett, California.
- STATE OF CALIFORNIA (2024): “California Remains the World’s 5th Largest Economy,” URL: <https://www.gov.ca.gov/2024/04/16/california-remains-the-worlds-5th-largest-economy/>.
- TAYLOR, L. D. AND H. S. HOUTHAKKER (2009): *Consumer Demand in the United States: Prices, Income, and Consumption Behavior*, Springer Science & Business Media.
- TRANSPORTPOLICY.NET (2017): “U.S. Fuels: Diesel and Gasoline,” URL: <https://www.transportpolicy.net/standard/us-fuels-diesel-and-gasoline/>.
- U.S. ENERGY INFORMATION ADMINISTRATION (2007): “Refinery Outages: Description and Potential Impact On Petroleum Product Prices,” Tech. Rep. SR/OOG/2007-01, Office of Oil and Gas.
- (2012): “PADD Regions Enable Regional Analysis of Petroleum Product Supply and Movements,” URL: <https://www.eia.gov/todayinenergy/detail.php?id=4890>.
- (2015): “West Coast Transportation Fuels Markets,” Tech. Rep. 20585, U.S. Energy Information Administration, Washington, DC.
- (2022a): “Table E20. Motor Gasoline Price and Expenditure Estimates, Ranked By State, 2022,” .
- (2022b): “What is Crude Oil and What are Petroleum Products,” Series: Energy Explained URL: <https://www.eia.gov/energyexplained/oil-and-petroleum-products/>.

——— (2024): “California Law and Refinery Closure Reflect Ongoing Challenges for the State’s Fuel Market,” .

VALENTINE, J. AND A. JOSEFSON (2017): “Best Practices for Solving Crude and Fuel Blending Challenges Webinar,” Series: Emerson Webinars URL: https://www.youtube.com/watch?v=1S8dsiR_IOW&ab_channel=MicroMotionCoriolisFlowandDensityMeasurement.

YEH, S. AND D. SPERLING (2010): “Low Carbon Fuel Standards: Implementation Scenarios and Challenges,” *Energy Policy*, 38, 6955–6965.

A Types of Refining Units

There are several ways to classify refinery units: the stage of the refining process at which they are located, the physical and chemical reactions they use, or the outputs they generate. We make a broad classification, for explanatory purposes, into Primary Conversion Units and Secondary Conversion Units mostly based on the stage of the refining process in which the units are located.

A.1 Primary Conversion Units

The crude distilling unit (CDU) and coking units are part of the Primary Conversion Units. These units are shown at the beginning of the refining process in Figure 6 highlighted in orange. The main task of these units is to transform heavy and medium hydrocarbons into lighter distillates by “cracking” them.¹⁵ These lighter distillates have higher octane levels.

Specifically, the CDU separates crude oil into different components based on boiling points. While the coking unit processes heavy residual components from the CDU into hydrocarbons with lower boiling points and solid petroleum coke. The lighter hydrocarbons with further serve as inputs into other units. After crude oil passes through the CDU, some heavy residual fuels remain. Coking units convert these heavy residuals into hydrocarbons with lower boiling points that can be used as inputs to other units.

A.2 Secondary Conversion Units

These units modify and improve the quality or properties of the products derived from the primary units by further reducing the level of polluting elements or increasing the octane level. For example, the alkylation, hydrotreating, and reformer units are essential to reach CARB-grade standards. Specifically, California Air Resources Board (CARB) mandates specifications that reduce the pollution from gasoline consumption compared to conventional gasoline. These specifications are reached largely thanks to the alkylation and hydrotreating units.

Finished gasoline is a blend of petroleum distillates and ethanol. According to Larson (2018), most geographies in the United States use a blend known as Conventional

¹⁵Cracking is the process where heavy hydrocarbon molecules are broken up into lighter molecules, usually with higher octane levels.

Blendstock for Oxygenate Blending (CBOB), while some highly populated areas along the Northeast coast of the United States use a blend with more stringent environmental requirements known as Reformulated Blendstock for Oxygenate Blending (RBOB).

However, in California, a different blend known as California Reformulated Gasoline Blendstock for Oxygenate Blending (CARBOB) is mandated. Compared to RBOB, this blend has even lower volatility or tendency to vaporize, as measured by Reid Vapor Pressure (RVP); it also has lower levels of toxic pollutants like sulfur and benzene than conventional blends.

Table 5: Specifications for different types of gasoline

Gasoline parameter	<i>CBOB</i>	<i>RBOB</i>	<i>CARBOB</i>	
	(1)	(2)	(3)	Units
Benzene content	.	1.3	1.22	% of volume
Reid Vapor Pressure	7.9	7	5.99	psi
Sulfur content	80	80	21	ppm

Source: California Air Resources Board (2014) and TransportPolicy.net (2017).

To achieve CARB’s standards, California refiners need to include components into the blend that will accomplish two opposing objectives: reaching the desired octane level and the desired environmental regulations. Three refining units are instrumental in reaching these dual objectives: alkylation, reformer, and hydrotreating units.

Although the volume of output from the secondary conversion units is low compared to the CDU and coking units, these units are essential for meeting environmental regulations. Therefore, their outages strongly affect the ability to produce CARB-compliant gasoline. Figure 6 shows a flow chart of the refinery process and highlights in red where the alkylation and hydrotreating units participate in creating the gasoline blend.

Alkylation units produce alkylates, a distillate with low RVP, low sulfur, and high octane levels (Peterson, 1996). However, there are two downsides to alkylation units: the first is that they have high fixed costs; the second one is that one of the main inputs into the unit is isobutane.

Isobutane is a chemical produced during one of the refining stages and has limited availability. Isobutane is also a main input to another process called polymerization, which produces a high-octane distillate, yet this distillate is more polluting than the alkylates. Installing a polymerization unit is substantially cheaper than an alkylation unit and yields a high-octane product, making it more attractive for refiners (Gary

et al., 2007a). Therefore, alkylation units tend to be used only in regions with very strict RVP requirements, such as California (Peterson, 1996).

Hydrotreating units are the second set of units that help refiners achieve CARB's standards. The main objective of these units is to reduce a product's sulfur content. These units are common throughout the United States and not only in California, as opposed to alkylation units. However, California refineries tend to rely more on hydrotreating units than their counterparts in the contiguous United States (California Energy Commission, 2020b). Therefore, an outage in these units noticeably affects the ability of refiners to meet CARB's standards.

Notes: Refining petroleum products is a multi-step process in which the output of some refining units serves as input to other refining units. Due to environmental regulations, the output of some refining units is critical to meet environmental standards; therefore, an outage in one of these units may have an outsized effect on the capacity of a refinery to produce gasoline.

B Testing Exogeneity of Unplanned Outages

As an additional check, we perform Granger causality tests to investigate if past prices can predict unplanned outages. Our estimation strategy assumes that unplanned outages are unrelated to past prices. To test this assumption, we regress the five unplanned outages, which we include in the IV estimation strategy (see Table 3), on past log prices and monthly controls to allow for within-year seasonality and the lagged log WTI price to control for aggregate energy demand changes.

Table 6: Granger causality test

Dep. Variable	Coefficient estimate	s.e.	<i>p</i> -value
Alkylation unit outage	0.86	0.56	0.13
CDU outage	-0.13	0.14	0.36
Coker unit outage	-0.66	0.44	0.14
Hydrotreating unit outage	-0.14	0.13	0.30
Reformer unit outage	0.15	0.47	0.76

Notes: Following Angrist and Pischke (2009); Deaton (2010) we argue that exogeneity cannot be directly tested for. However, we show evidence consistent with conditional mean independence by showing that past prices cannot predict future unplanned outages.

Table 6 reports the results of the Granger causality tests. We find that past prices do not predict unplanned outages, consistent with our assumption of exogeneity of the instrument. For all of the unplanned outages, we cannot reject the null hypothesis of no relationship with past prices. The *p*-values on past prices lie above 10% for unplanned outages. Furthermore, the sign has no consistent pattern: two out of five coefficients are positive, while three are negative. This evidence confirms *no precedence* of prices on outages in the sense of Leamer (1985).

C Statistical Association Between Quantities Consumed Across Time and an Initial Price Shock

It is interesting to examine the market dynamics after an initial price shock to have a better understanding between short-term and long-term effects. To learn about how a price shock passes onto market outcome variables, we estimate a model identical to Equation (1) but change the dependent variable to q_{t+h} for $h = 0, 1, 2, \dots, 13$ using the local projection approach of Jordà (2005).

This exercise measures the statistical association between a price shock at time t and subsequent market-equilibrium quantity outcomes. Since this exercise uses observational variables q_{t+h} , $h = 0, 1, 2, \dots, 13$, but it no longer controls for demand nor supply changes h periods after the initial shock, we cannot isolate consumer behavior from the observed market outcome.¹⁶

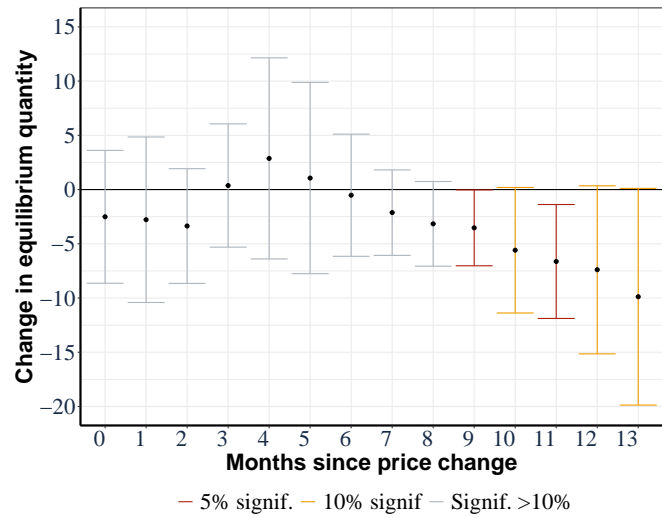
Table 7: Parameter estimates of the dynamic association[†]

Event time	Equilibrium quantity	s.e.
0	-2.51%	3.08
1	-2.78%	3.84
2	-3.36%	2.66
3	0.37%	2.86
4	2.87%	4.67
5	1.06%	4.44
6	-0.52%	2.83
7	-2.12%	1.98
8	-3.16%	1.96
9	-3.53% **	1.76
10	-5.59% *	2.91
11	-6.63% **	2.64
12	-7.39% *	3.90
13	-9.88% *	5.03

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

[†] between market-level quantities and a price increase of 10%.

Figure 9: Event study of a 10% price increase on the market equilibrium quantities



Notes: The figure shows the dynamic association between market-equilibrium quantities and an initial, and sustained, price increase of 10% at $t = 0$.

Table 7 reports the local projection coefficients based on model Specification 3. We find that a 10% price increase is associated with drop in equilibrium quantities during the current month and the next two months as well; however, the estimates are noisy, likely following from our limited sampling period. The results show that nine months after an initial price change, the associated change in quantity is larger in magnitude and significant ranging from -3.5% to -6.6%, in line with the long-term elasticity estimates

¹⁶Due to the somewhat short sample period and the fact that, when implementing this local projection methodology we lose observations as we are projecting forward, the power of our estimate is reduced. If we were to include further demand controls, the power will be further reduced as we lose degrees of freedom. In the future, as more data become available and the resulting power of our empirical analysis increases, we will estimate the impulse response function.

reported in table 4. Associated with this table, we present Figure 9, which plots the estimated correlation coefficients along their 95% confidence intervals.

Due to the somewhat short sample period and the fact that, when implementing this local projection methodology, we lose observations as we are projecting forward, the power of our estimate is reduced when compared with the partial adjustment model estimates in the main text of text paper. In the future, as more data become available and the resulting power of our empirical analysis increases, it will be possible to estimate the impulse response function with more precision.