California Gasoline Demand Elasticity Estimated Using Refinery Outages*

Armando R. Colina — Banco de México
Bulat Gafarov — UC Davis
Jens Hilscher — UC Davis

May 2024

Abstract

This paper presents new gasoline price elasticity estimates for California. We use unique characteristics of the California gasoline market and a new set of proposed instruments that are strong and plausibly exogenous. As a first step, we take advantage of California's unique gasoline market, which is partially isolated from the rest of the U.S. due to environmental regulations. We control for persistent demand shocks and estimate a lower bound for the elasticity of demand of -0.23. In the second step, we use a new set of instruments to control for simultaneity. We use detailed information on refinery outages to capture short-run supply shocks. Our estimate of long-run demand elasticity is -0.57.

JEL codes: C22, C36 ,C51 ,D12 ,Q41.

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

*Corresponding Author: Armando R. Colina. Banco de México, Av. 5 de Mayo 2, Alc. Cuauhtémoc, México City, México. Email: armando.rangel@banxico.org.mx, Phone: +52 55-2563-6220. This paper is based on the second chapter of Armando's Ph.D. dissertation. The views and conclusions presented in this work are exclusively the responsibility of the authors and do not necessarily reflect those of Banco de México. We would like to thank Armando's advisers, Richard J. Sexton, James B. Bushnell, and Erich Muehlegger, for their helpful guidance and comments.

Electronic copy available at: https://ssrn.com/abstract=4629611

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 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 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 US, California has been at the forefront of regulating gasoline and energy markets to encourage the transition to a lower greenhouse gas (GHG) emission economy (e.g. cap and trade system, low carbon fuel standard). 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 in the US, given its large size and strict environmental regulations. To achieve the environmental standards set by the Cal-

[1] For 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.

Electronic copy available at: https://ssrn.com/abstract=4629611

ifornia 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 close to a separated gasoline market. Therefore, reductions in refining capacity cannot easily be compensated for 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. Including such lagged variables as controls, we estimate 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. There is still some attenuation bias from supply and demand simultaneity. These estimates can therefore be seen as a lower bound on the absolute size of the demand elasticity.

In the second step, we introduce refinery outages as instruments for supply shocks to address supply and demand simultaneity. 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 they are conditionally uncorrelated with gasoline demand.

Electronic copy available at: https://ssrn.com/abstract=4629611

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 occurred, and how much capacity was lost.

Based on these instruments, we can calculate two sets of price elasticity estimates for the California gasoline retail market. We estimate the one-month price elasticity of demand to be -0.2 and the long-term elasticity to be -0.57.

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 as compared to our estimate. The higher elasticity in California 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

Electronic copy available at: https://ssrn.com/abstract=4629611

discussion and concludes.

2 Previous gasoline elasticity estimation approaches

Gasoline consumption touches several aspects of our everyday lives. 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.

There are three broad areas of focus: (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.

Electronic copy available at: https://ssrn.com/abstract=4629611

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

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 prod. disruptions	Retail demand	-0.03	0.01
Li et al. (2014)	Gas taxes & oil prices	Retail and tax elast. of dem.	-0.07	0.02
Coglianese et al. (2017)	Gas taxes	Retail demand	-0.37	0.23
Kilian and Zhou (2024)	Gas taxes	Retail demand	-0.32	0.067

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 (Energy Information Administration (2022)). 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. This may be for two reasons: most of the US crude oil production disruptions that the authors consider are related to weather events in the U.S. Gulf Coast

Electronic copy available at: https://ssrn.com/abstract=4629611

(USGC). 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 appaoches 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.

Finally, Coglianese 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.

Using micro-level data, it can sometimes be possible to estimate demand elasticity precisely without the use of an IV approach. Levin et al. (2017) use disaggregated daily panel data for 243 US 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 because it is largely segmented and the largest state in the US.

Electronic copy available at: https://ssrn.com/abstract=4629611

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 US. 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.

Electronic copy available at: https://ssrn.com/abstract=4629611

Several circumstances lead to outages. Examples include an unplanned power outage; an unplanned flaring event;[2] a malfunction in the refining unit caused by a leak, a crack, or loss in pressure or a fire; unexpected high winds; unexpected malfunction of the crude pipeline that supplies the refinery; replacing an old unit with a newer one; or replacing a part due to wear and tear.

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 middle-stage units that produce a sizeable share of the finished output and sulfur-reducing units that produce specific distillates. The latter produce smaller amounts of distillates but are essential to achieve the regulatory requirement; 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 US Energy Information Administration. The outage data set comes from Bloomberg; these data are available starting in 2011. Our sample period therefore runs from January 2011 until March 2023.

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 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

[2]A flaring event happens when excess hydrocarbons are burned rather than released straight into 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).

Electronic copy available at: https://ssrn.com/abstract=4629611

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. 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 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 :					
:: Alkylation unit (MMgals/month)	139	21.1	41.6	0.0	233.8
:: Coker unit (MMgals/month)	139	23.2	33.1	0.0	139.8
:: Hydrocracking unit (MMgals/month)	139	29.4	41.6	0.0	233.2
:: FCC unit (MMgals/month)	139	57.2	73.8	0.0	316.3

Notes: MM denotes millions

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

Electronic copy available at: https://ssrn.com/abstract=4629611

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

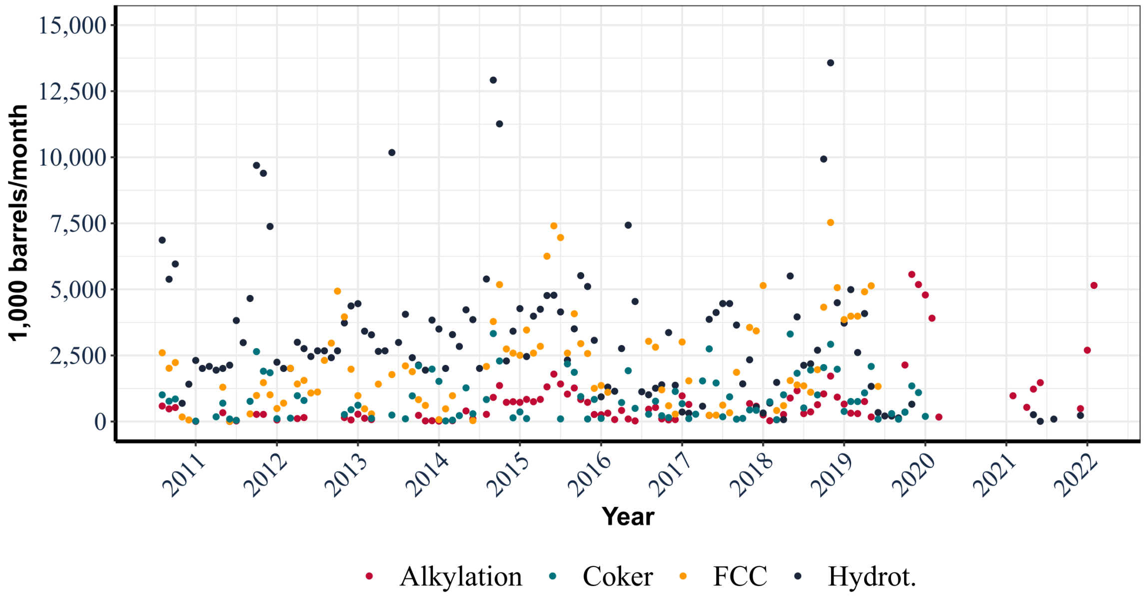

supplier, which is costly (American Petroleum Institute (2013)); reduce total output while maintaining the optimal blend, which increases inventory costs due to unused components (Energy Information Administration (2007)); or produce a suboptimal blend subject to achieving performance requirements (Valentine 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 U.S. are reported to the Occupational Safety and Health Administration (OSHA). Large, unplanned outages

Electronic copy available at: https://ssrn.com/abstract=4629611

are reported in specialized news outlets. To account for possible salience effects, we will include lagged retail sales.

Figure 2: Distribution of the duration of outages

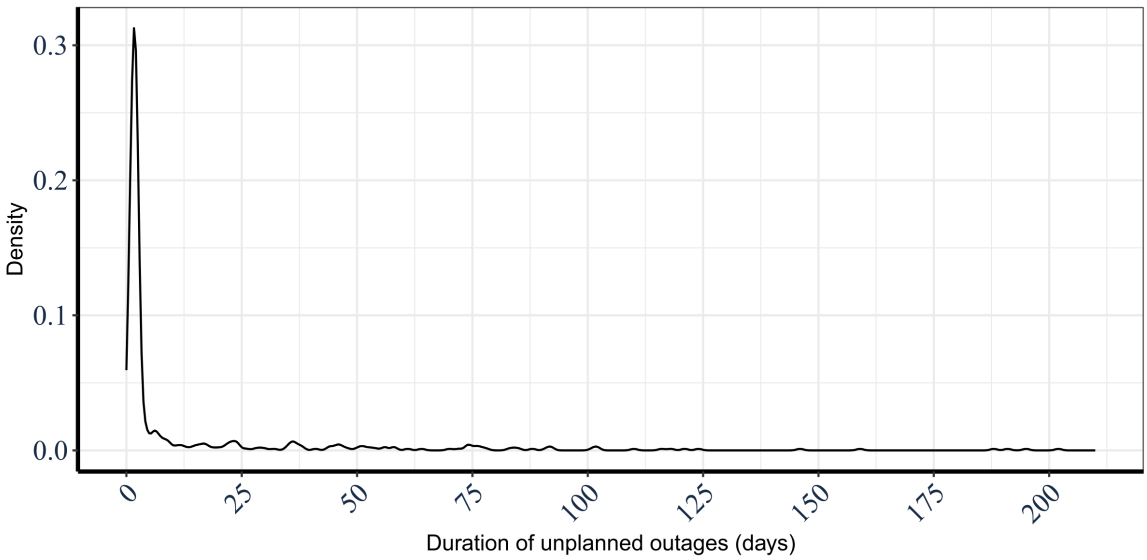

Another possibility is that consumers may expect the duration of the outage to be long-lived, and they react differently from a regular price change. However, the vast majority of outages are resolved in a short period of time; 83% of them are solved in less than a month. Therefore, outages are unlikely 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

Electronic copy available at: https://ssrn.com/abstract=4629611

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 US. However, within the US, only California refineries produce this blend for the majority of our sample period (Pyziur (2016)). Therefore, when there is an outage and refining capacity is lost, wholesalers cannot substitute local production with refined products produced elsewhere in the U.S.

Figure 3: Utilization of refining capacity

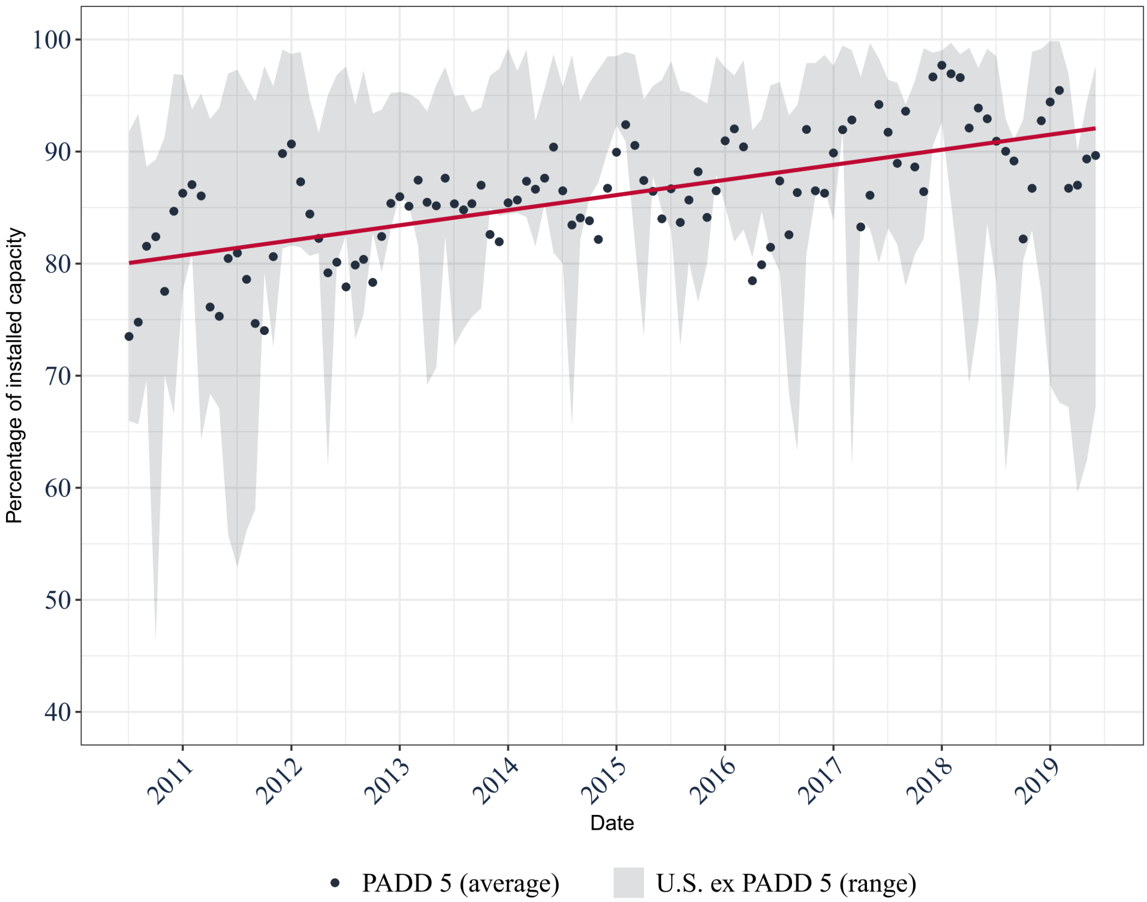

An alternative to sourcing refined products from outside the state is to substitute with products made within the state. However, refineries in California have been operating increasingly close to their installed capacity, making it harder to increase regional production in response to a local outage. Figure 3 shows the throughput of refineries in PADD 5 as a percentage of the installed available capacity.[3] PADD

[3]A Petroleum Administration for Defense Districts (PADD) is a geographic division of crude

Electronic copy available at: https://ssrn.com/abstract=4629611

5, encompasses California and other states along the West of the U.S., of which California is by far the largest.

A second alternative to substitute for the loss of local production capacity is to import refined products into California from outside the US. There are two other countries that 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 required 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

and fuels markets established during World War 2 to ration gasoline consumption. Today, market participants use the division to analyze regional trends (Energy Information Administration (2012)).

Electronic copy available at: https://ssrn.com/abstract=4629611

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. 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.[4] In contrast, the alkylation unit has a smaller capacity as it is used mostly 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 refniery units in more 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; the conventional instrument threshold for the F-statistic is 10 and this is cleared in all three specifications. As expected, we find that the coefficients on outages are positive and in most cases statistically significant.[5]

[4]The installed refining capacity in California is approximately 2.1 billion gallons per month

[5]However, we note that the outages of the hydrotreating unit have a significant and, at first

Electronic copy available at: https://ssrn.com/abstract=4629611

One of the advantages of using more detailed data on outages is our ability to account for different effects of the various outage types. 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.

Table 3: First-stage estimation results.

	First stage regression		
Dep. Variable: log Retail Prices	*Specification 1* (1)	*Specification 2* (2)	*Specification 3* (3)
A) Coefficient estimates			
Alkylation unit outage	0.473	0.576*	0.458
:: s.e.	(0.313)	(0.324)	(0.297)
Hydrotreating unit outage	-0.626***	-0.606**	-0.576***
:: s.e.	(0.121)	(0.139)	(0.112)
FCC unit outage	0.625**	0.585***	0.553**
:: s.e.	(0.217)	(0.221)	(0.200)
Hydrocracking unit outage	0.702*	0.668	0.617
:: s.e.	(0.421)	(0.432)	(0.431)
Coker unit outage	1.461**	1.281***	1.305**
:: s.e.	(0.469)	(0.482)	(0.471)
log Retail sales, lagged	x	x	x
log WTI, lagged	x	x	x
log Inventories, lagged		x	x
log Imports, lagged		x	x
log Capacity util., lagged			x
Constant	7.821***	10.321***	15.603***
:: s.e.	(2.105)	(3.242)	(3.838)
Seasonal Month fixed effects	Yes	Yes	Yes
B) Model stats.			
R-squared	0.790	0.804	0.811
F stat.	18.71	20.55	26.80

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.

4 Model and results

Following 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

glance, counterintuitive negative coefficient. A possible explanation is that some maintenance outages may have occurred during the months of the year with low demand.

Electronic copy available at: https://ssrn.com/abstract=4629611

of interest: the contemporaneous price elasticity of demand, β 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.[6] To control for dynamic adjustment, we estimate different model specifications given by lagged covariates in the following model,

$$q_t = \beta_0 + \sum_{s=1}^{11} m_t^s + \sum_{\ell=1}^{L} \Big(\beta_\ell p_{t-\ell+1} + \rho_\ell q_{t-\ell} + \lambda_\ell wti_{t-\ell} + \gamma_\ell inv_{t-\ell} + \delta_\ell imp_{t-\ell} + \theta_\ell util_{t-\ell}\Big) + \varepsilon_t \tag{1}$$

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 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. It is possible that there may be autocorrelation in the errors; alternatively, it is possible that the shock is i.i.d. but there is partial adjustment in the quantity, e.g. due to 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.

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

Electronic copy available at: https://ssrn.com/abstract=4629611

Having estimated the 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 $\beta^\star \equiv \beta_1/(1-\rho_1)$.

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. Therefore, the OLS estimates provide a lower bound for the elasticity estimates. The second estimation technique uses an instrumental variable estimation procedure to control for simultaneity. We label these columns as *IV*.

Table 4: Elasticity estimates under various specifications

Dep. Var.: log Retail Sales	*Specification 1*		*Specification 2*		*Specification 3*	
	OLS (1)	*IV* (2)	*OLS* (3)	*IV* (4)	*OLS* (5)	*IV* (6)
A) Long term elasticity (β^*)	-0.204*	-0.613**	-0.235**	-0.713**	-0.230**	-0.569*
:: s.e.	(0.118)	(0.299)	(0.118)	(0.325)	(0.112)	(0.300)
B) Coefficient estimates						
log Retail price (β_1)	-0.041	-0.160*	-0.044	-0.199	-0.047	-0.198*
:: s.e.	(0.029)	(0.095)	(0.027)	(0.122)	(0.040)	(0.106)
log Retail sales, lagged (ρ_1)	0.797***	0.737***	0.809***	0.720***	0.794**	0.651***
:: s.e.	(0.084)	(0.070)	(0.078)	(0.093)	(0.172)	(0.177)
log WTI, lagged	x	x	x	x	x	x
log Inventories, lagged			x	x	x	x
log Imports, lagged			x	x	x	x
log Capacity util., lagged					x	x
Constant (β_0)	4.182**	5.376***	4.790**	6.901***	5.103*	8.296**
:: s.e.	(1.817)	(1.460)	(1.602)	(2.345)	(3.011)	(2.330)
Seasonal month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	138	138	138	138	138	138
R-squared	0.717	0.700	0.723	0.696	0.723	0.699
Bayesian Information Criteria		-373.6		-362.0		-358.4

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

Note: Newey-West standard errors in parenthesis.

Table 4 Panel B reports estimates of the model parameters. The price coefficients

Electronic copy available at: https://ssrn.com/abstract=4629611

in all specifications are negative but not statistically significant in the case of OLS. The corresponding IV estimates are approximately four times as large and become significant at the 10% level for specifications 1 and 3. 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.[7]

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), and Kilian and Zhou (2024) for further details). The change in magnitude in the elasticity estimates validates the assumption that the 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% or 10% level. Our preferred specification is #3 – it is likely that the additional controls in that specification are effective at reducing potential bias in the estimate. We note in particular that, when additional controls are added relative to specification #2, the point estimate drops quite a bit. Another reason for this preference is that specification #2 does not have a statistically significant short-term (one month) price elasticity estimate. Finally, the long-term elasticity has a slightly lower standard error.

5 Discussion and conclusion

In this paper, we propose refinery outages as a new set of instruments to estimate the price elasticity of demand for retail gasoline in California. Our estimation approach applies a new set of instruments to the California context and differentiates between outages based on the refinery unit level. We focus on coking, HCU and FCC units since they produce naphtha, a significant blending component by volume of finished

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

Electronic copy available at: https://ssrn.com/abstract=4629611

gasoline. We also use outages in the alkylation and hydrotreating units. While these units contribute less to the finished product when measured by volume, the chemical properties of their output are essential to make CARB-compliant gasoline, making their outage have an outsized effect. Due to the granularity of the outage data set and the conditions in the California gasoline market, the proposed instruments have a strong first stage. Similarly, due to the largely unpredictable nature of the outages, these are conditionally uncorrelated with unobserved demand components.

We find that the one-month price elasticity is -0.2 and that the long-term elasticity is -0.57. These estimates imply that California consumers are more price sensitive than the average, perhaps because of a heightened awareness to the consistently high prices in California. The magnitude is consistent with cross-sectional heterogeneity documented by Kilian and Zhou (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).

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, but 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.

Gasoline consumption touches on several aspects of our everyday lives, and knowing the price sensitivity of consumers 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' reaction to price changes.

Electronic copy available at: https://ssrn.com/abstract=4629611

References

American Petroleum Institute (2013): "Economic and Supply Impacts of a Reduced Cap on Gasoline Sulfur Content," Tech. rep., Turner, Mason and Company.

Berry, S. T. and G. Compiani (2021): "Empirical Models of Industry Dynamics with Endogenous Market Structure," *Annual Review of Economics*, 13, 309–334.

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," .

California Energy Commission (2020a): "How petroleum products move," Tech. Rep. March 2020, California Energy Commission, California.

——— (2020b): "Refining operations under decreased demand," Tech. Rep. May 2020, California Energy Commission, 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.

Davis, L. W. and L. Kilian (2011): "Estimating the Effect of a Gasoline Tax on Carbon Emissions," *Journal of Applied Econometrics*, 26, 1187–1214.

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.

Electronic copy available at: https://ssrn.com/abstract=4629611

——— (2012): "PADD regions enable regional analysis of petroleum product supply and movements," .

——— (2022): "What is crude oil and what are petroleum products?" .

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 Author (s): Justine S . Hastings Source : The American Economic Review , Vol. 94, No . 1 (Mar ., 2004), pp . 317-328," *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.

Electronic copy available at: https://ssrn.com/abstract=4629611

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.

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," .

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. Muehlegger (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 semi-parametric approach," *Energy Economics*, 45, 244–253.

Newey, W. K. and K. D. West (1987): "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 55, 703–708.

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 US electricity demand by region, season, and sector," .

Electronic copy available at: https://ssrn.com/abstract=4629611

Peterson, J. R. (1996): "Alkylate is Key for Cleaner Burning Gasoline," Tech. rep., STRATCO.

Pugliaresi, L. and M. Pyziur (2015): "Gasoline blending an EPRINC primer," Primer, Energy Policy Research Foundation, Washington D.C.

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.

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," .

Valentine, J. and A. Josefson (2017): "Best Practices for Solving Crude and Fuel Blending Challenges Webinar," .

Yeh, S. and D. Sperling (2010): "Low carbon fuel standards: Implementation scenarios and challenges," *Energy Policy*, 38, 6955–6965.

Electronic copy available at: https://ssrn.com/abstract=4629611

A Types of refining units

A.1 Middle-stage refining units

The coking, fluid catalytic cracking unit (FCC) and hydrocracking unit (HCU) are part of the middle stages of the refining process. Units are shown in Figure 4 highlighted in orange. The main task of these units is to transform heavy and medium hydrocarbons into lighter distillates by "cracking" them.[8] These lighter distillates have higher octane levels.

Coking unit: 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.

FCC unit: transforms heavy hydrocarbons into lighter products like gasoline and naphtha by applying heat and catalysts.

HCU unit: transforms blends with high sulfur content into naphthas and gasoline by using high pressure, hydrogen, and catalysts.

The output of the FCC and HCU units contribute approximately 40% of the volume of the finished gasoline blend (Pugliaresi and Pyziur (2015)). Therefore, an outage in these units constrains overall output by volume.

A.2 Sulfur-reducing refining units

One thing that makes California's gasoline market unique is its environmental regulations. 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 Lar-

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

Electronic copy available at: https://ssrn.com/abstract=4629611

son (2018), most geographies in the US use a blend known as Conventional Blendstock for Oxygenate Blending (CBOB), while some highly populated areas along the North-east coast of the U.S. 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	*CBOB* (1)	*RBOB* (2)	*CARBOB* (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 reach the desired environmental regulations. There are two refining units that are instrumental in reaching these dual objectives: the alkylation unit and the hydrotreating units.

While the volume of output from these units is low compared to the FCC, HCU, and coking units, these units are essential to meeting environmental regulations. Therefore, their outages strongly affect the ability to produce CARB-compliant gasoline. Figure 4 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

Electronic copy available at: https://ssrn.com/abstract=4629611

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 polimerization 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 U.S. 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 US (California Energy Commission (2020b)). Therefore, an outage in these units noticeably affects the ability of refiners to meet CARB's standards.

Electronic copy available at: https://ssrn.com/abstract=4629611

Figure 4: Refinery flowchart from Gary et al. (2007d).

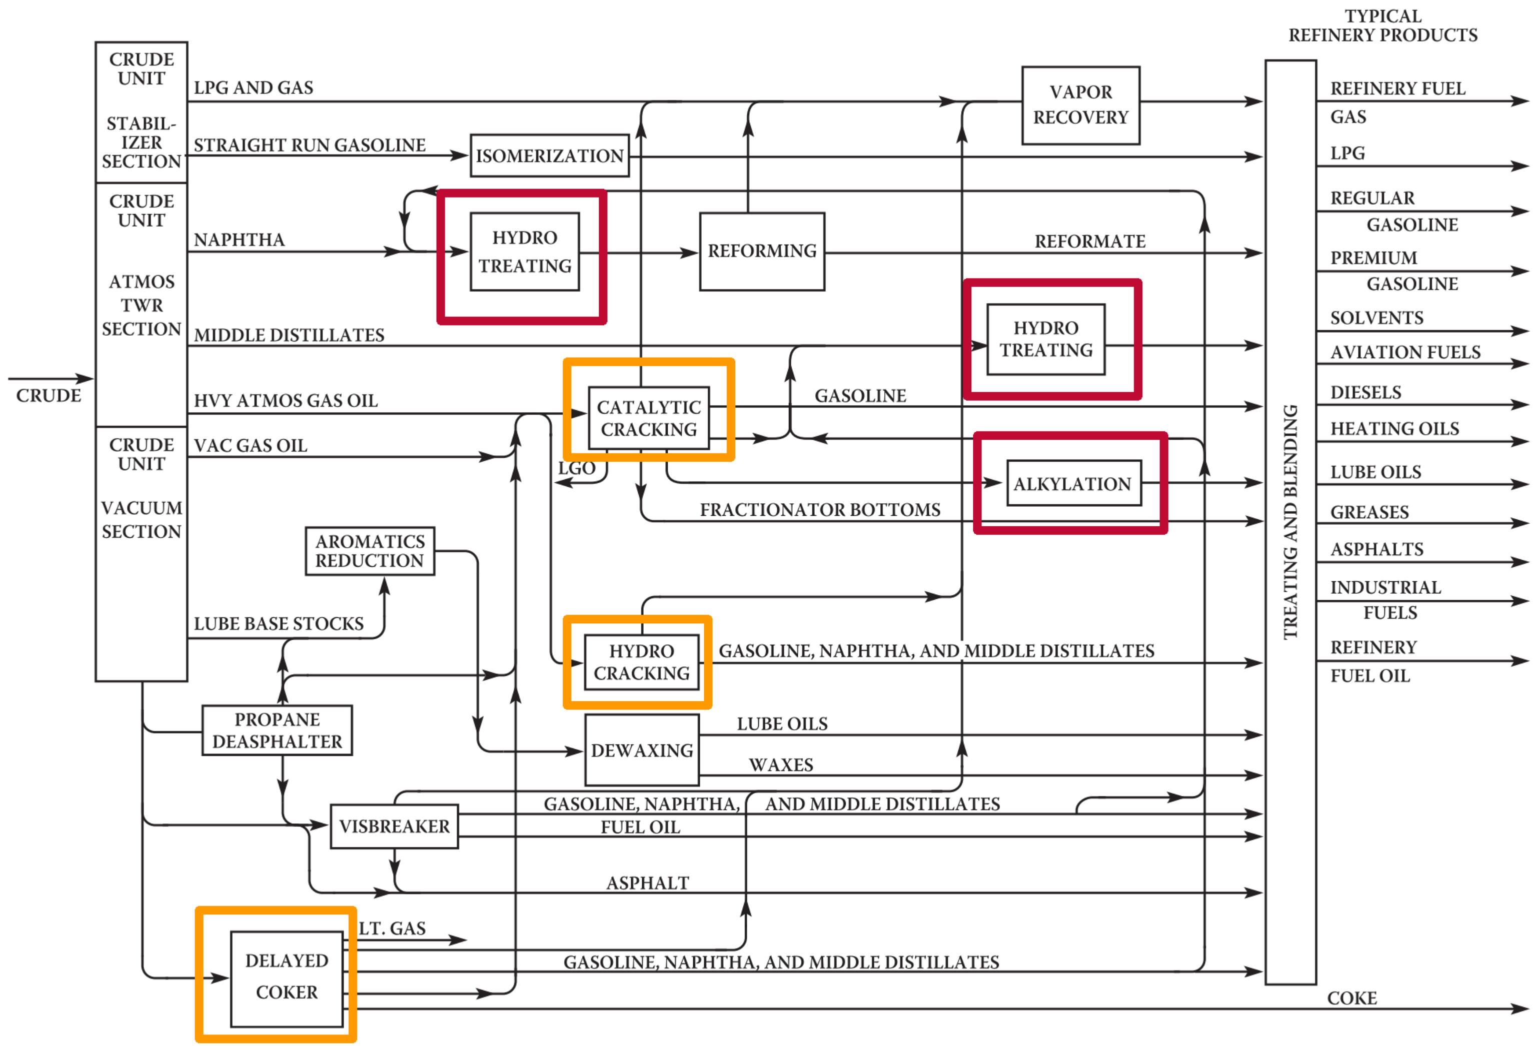

Electronic copy available at: https://ssrn.com/abstract=4629611