Industry Compliance Costs Under the Renewable Fuel Standard: Evidence from Compliance Credits*

Arthur R. Wardle†

Sherzod B. Akhundjanov‡

May 28, 2019

The Renewable Fuel Standard (RFS), which requires oil refineries to blend ethanol into domestic fuel supplies, is a market-based policy that implements tradable compliance credits to better equalize compliance costs across firms. We use variation in the prices of these credits to retrieve reduced-form estimates of how the RFS impacts the values of refining firms. In addition to fitting bivariate time series models to compliance credit and firm values, we use unanticipated mid-compliance-year regulatory announcements that changed mandated biofuel blending volumes to identify event study models. Previous evidence on cost pass-through suggests that the RFS should have little impact on refining firms. In contrast, we find a significant negative oil refinery stock price response to shocks in RFS compliance credit values. This negative effect is limited to refining firms with large market capitalizations and integrated downstream operations. This evidence discredits a widespread critique of the RFS claiming that integrated refiners are able to draw profits from merchant refiners that lack downstream blending and retail operations.

Key Words: Biofuels, Ethanol, Refinery, Renewable Fuel Standard, Renewable Identification Numbers

JEL Classification: H23; L71; Q35; Q41; Q42; Q48

---

*Early drafts of this paper benefited from comments by William Shughart, Ben Blau, and Megan Hansen. The Center for Growth and Opportunity at Utah State University provided financial support for data procurement and research. All errors are the author's own.

†Graduate Student, Department of Economics & Finance, Utah State University. arthur.wardle@aggiemail.usu.edu (corresponding author).

‡Assistant Professor, Department of Applied Economics, Utah State University. sherzod.akhundjanov@usu.edu

1 Introduction

The Renewable Fuel Standard (RFS), created under the Energy Policy Act of 2005 and greatly expanded by the Energy Independence and Security Act of 2007, mandates the use of various biofuels in domestic transportation fuel supplies.[1] The statute itself includes volumetric mandates for cellulosic, biomass-based biodiesel, "advanced," and total renewable fuels through 2022. Obligated parties (fuel refiners and importers) are required to submit specified numbers of Renewable Identification Numbers (RINs) to comply with the standard. RINs are created by biorefineries, which link them to barrels of renewable fuels, and are split from those barrels upon blending. RINs can be traded, allowing obligated parties to comply with RFS mandates either by blending renewable fuels themselves or buying excess RINs from other parties. The construction of the RIN market is functionally equivalent to other market-based environmental regulations, such as pollution permits.

Because regulated firms can comply with the RFS either by blending additional biofuels or by buying RINs, the basic fundamental value of a RIN (with some complications, described later in this paper) is the marginal cost to the refining sector of blending an additional unit of biofuel.

Even high RIN prices do not necessarily impact the bottom line of obligated parties if these costs are easy to pass through to consumers and demand for transportation fuels is inelastic. Indeed, a large literature establishes that refiners are able to fully (or even more than fully) pass RFS compliance costs onto consumers (Burkhardt, 2016; Knittel et al., 2017; Pouliot et al., 2017; Li and Stock, 2019; Lade and Bushnell, 2019). This, however, does not imply that the RFS has no financial impact US oil refiners whatsoever. Complete pass-through does not replace the profit margins on refined crude oil that counterfactually would have been sold without the mandate and may impose infrastructure costs to accommodate ethanol blends, negotiating costs with biorefineries, and myriad costs of doing business not fully captured in models of cost pass-through. Knittel and Smith (2015) give a fuller description of ethanol's impact on oil refining profitability. The goal of our analysis is to establish how changes in the prices of RFS compliance credits (RINs) impact the value of the policy's obligated parties (refiners). To accomplish this, we implement two different reduced-form methods, which allow us to avoid imposing any particular causal channel for how the policy might impact firms' value.

[1] A complete description of how the RFS works is available in Schnepf and Yacobucci (2013).

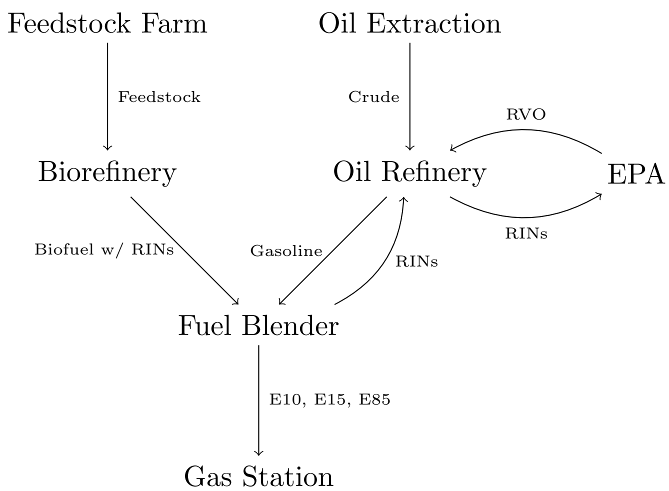

Figure 1: Obligated Parties (Oil Refiners) in a Simplified Gasoline Supply Chain

Outside of the academic pass-through literature, some stakeholders in RFS debates express concerns that larger refiners use the RFS to disadvantage smaller ones. Refiners in the United States can be split into "merchant" refiners, who do not blend their own fuel and are generally smaller, and "integrated" refiners, who do. Though refineries are the parties obligated by the RFS, RINs are not actually separated from biofuels until they are blended with gasoline. All this can be seen in Figure 1, with integrated refineries owning both "Oil Refinery" and "Fuel Blender" assets and merchant refiners owning only the former. Not owning blending assets leaves merchant refiners in the position of having to buy RINs on the market rather than being able to generate them themselves. Merchant refiners and their advocacy organizations often claim that being unable to generate RINs puts their operations at a competitive disadvantage and allows integrated refineries to sell excess RINs for windfall profits (see discussion and footnotes in Environmental Protection Agency, 2017, p. 21-31). The Environmental Protection Agency has dismissed such arguments, pointing primarily to economic research on RIN pass-through in doing so (Environmental Protection Agency, 2017). Research by Babcock et al. (2016) provides a theoretical explanation for why the RFS should not impact merchant and integrated refiners differentially.

We directly examine the impact of RIN price fluctuations on the stock prices of obligated refineries. First, we fit bivariate time series models for every firm $\times$ RIN combination. Modeling each firm separately allows us to investigate heterogeneity among firms. While these models provide a picture of how RINs and firm stock prices are associated over time, they are subject to endogeneity problems—both RINs and refining stocks are structurally related to commodity prices, fuel demand factors, and other variables. The intent of this paper is not to identify these structural relationships, and building them into the model is beyond the scope of this research. Instead, we take advantage of two large, exogenous shocks to RIN prices to identify a more plausibly causal estimate of how RIN price changes impact refining stocks. Following Lade et al. (2018), we use unanticipated regulatory announcements that drastically affected the price of RINs to identify the impact on every firm in our sample, as well as selected subgroups of firms guided by multivariate results.

We find that when RIN prices rise, the stock prices of refineries with large market capitalizations drop with a 3-5 day lag. The effect is statistically significant though economically small. Medium and small firms, however, exhibit no reaction to RIN price changes. These results are consistent across both the bivariate time series and event study analyses. Due to the reduced-form nature of our estimates, we are unable to identify any particular causal mechanism for these results, but we conclude with a list of potential hypotheses worth investigating in future research. In any case, our findings discredit claims that the RFS enables integrated refiners to take advantage of merchant refiners and raises doubts about the costs of so-called RIN "speculation."

2 Background on the Renewable Fuel Standard

The Renewable Fuel Standard is a nested mandate, meaning that blending higher-level biofuels also works to meet the mandate requirements at lower levels. RINs coming from corn ethanol generate D6 RINs, which can serve to fulfill only the lowest level of the mandate. Ethanol from more "advanced" sources such as sugarcane generates advanced ethanol RINs (D5) and constitutes a smaller, nested mandate. RINs from biomass-based diesel (D4), cellulosic biofuel (D3), or cellulosic diesel (D7) fulfill the mandate in their own categories, the advanced mandate, and the total mandate simultaneously. This nested structure is visualized in Figure 2.

The nested relationship gives rise to a price hierarchy for RINs, which binds empirically at almost all times (Whistance and Thompson, 2014):

$$P_{D6} \leq P_{D5} \leq \min\{P_{D4}, P_{D3,D7}\}$$

The absence of strictness to that inequality is not merely theoretical. Federal regulation

Figure 2: Structure of the Nested RFS Mandate

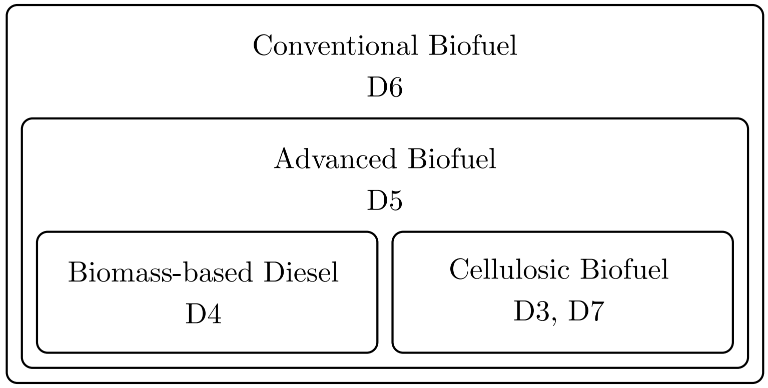

prevents most consumer fuels (excepting E85 and E15, which contain up to 85% and 15$ ethanol respectively, are available at a limited number of fueling stations, and can be used only in certain vehicles) from containing more than 10% ethanol. When national gasoline stocks are saturated with 10% ethanol, sales of additional ethanol can occur only through comparatively miniscule E85 and E15 channels. Thus, at mandate levels beyond 10% of nationwide gasoline sales, refiners must take advantage of the nested mandate structure and sell additional biodiesel to meet their requirements (Korting et al., 2019). In those market conditions, D6 prices track closely to D4 prices (Irwin, 2014).

Congress intentionally set statutory RFS volume mandates optimistically high. To prevent undue financial pressures on the transportation fuels industry, Congress explicitly allows the Environmental Protection Agency (EPA) to review the statutory volumetric standards and reduce them if compliance would be infeasible. The EPA has found it necessary to invoke that power numerous times. The following statement, released along with a proposed adjustment of the 2014-2016 mandates, illuminates the EPA's role in tempering the statutory requirements:

> Due to constraints in the fuel market to accommodate increasing volumes of ethanol, along with limits on the availability of non-ethanol renewable fuels, the volume targets specified by Congress in the Clean Air Act for 2014, 2015 and 2016 cannot be achieved. However, EPA recognizes that the statutory volume targets were intended to be ambitious; Congress set targets that envisioned growth at a pace that far exceeded historical growth rates. Congress clearly intended the RFS program to incentivize changes that would be unlikely to occur absent the RFS program. Thus

> while EPA is proposing to use the tools provided by Congress to waive the annual volumes below the statutory levels, we are proposing standards that are directionally consistent with Congress' clear goal of increasing renewable fuel production and use over time (Environmental Protection Agency, 2015).

In reviewing and adjusting the yearly mandates, the EPA issues a proposed rule, gathers public comments on that proposal, and then issues a final rule. The final rule is supposed to be complete by November 30 of the preceding year (e.g., 2015's final rule should be issued by November 30, 2014). In the lifespan of the RFS, the EPA has repeatedly missed that deadline (Bracmort, 2015). Final rules are often made partway through the compliance year and in one case a final rule was set almost a full year after the compliance year had passed. Research by Lade et al. (2018) demonstrates that such announcements shock RIN values as well as some commodity markets and biorefinery firm values, but no currently published research uses the shocks to identify the RFS's impact on the industry it actually regulates.

2.1 Industry Impact

While the RFS mandate's point of compliance generally rests with refiners, the associated costs can be pushed up or downstream if the industrial organization of consumer fuels markets allows it. As mentioned in the introduction, a developed literature discusses this very question. Most papers characterize the RFS as a subsidy to high-ethanol fuels and measure how changes in the value of the RIN 'subsidy' percolate to prices of E10 and E85 fuels. The general thrust of the literature concludes that pass-through is complete (or even more than complete) at the wholesale level (Burkhardt, 2016; Knittel et al., 2017; Lade and Bushnell, 2019), complete for E10 at the retail level (Pouliot et al., 2017; Lade and Bushnell, 2019; Li and Stock, 2019), and less than complete for most E85 (Lade and Bushnell, 2019; Li and Stock, 2019). The general interpretation of these results is that the RFS is irrelevant to the refiner, costly to most consumers, and beneficial to consumers of high-ethanol fuel and some E85 retailers.

Understanding pass-through does get us most of the way to understanding how the RFS affects refiners financially, but pass-through is not the only avenue by which the RFS could impose compliance costs. Costs associated with biofuel procurement and profit margins on gasoline sales lost to ethanol are just two additional potential ways the RFS could impose costs on refiners, even with complete pass-through. The methods undertaken in this paper impose no particular cost channel on examining the RFS's impact, instead electing to allow efficient financial markets price these anticipated costs. Examining stock price responses to variations in RIN prices gives a highly reduced-form but generalized and unstructured understanding of how the RFS impacts refineries.

We also analyze heterogeneities among refinery responses, rather than taking refiners as a homogeneous group. Existing research on whether the impact of the RFS is heterogeneous across obligated firms is scant. There are at least a few reasons differences might exist. For example, firms which own refining capacity in close proximity to ethanol production may react less negatively to RFS cost hikes. The RIN market is designed much like pollution permits in market-based environmental policies, so that the theoretical market equilibrium in RINs should equalize the shadow prices of additional ethanol blending across firms (Montgomery, 1972). Whether the RIN system effectively accomplishes that is an open empirical question. Findings by LaRiviere et al. (2017) indicate that the RFS imposed heavier burdens on retail gasoline consumers distant from ethanol production centers because ethanol's physical properties make it more expensive to transport, but this is further downstream than refining and blending.

We also investigate firm size heterogeneities by asking whether larger firms have an easier time passing-through or otherwise dealing with RFS costs. Built-in exemptions for small refiners facing economic hardship seem to indicate that the drafters of the RFS policy must have worried about differential impacts. Beyond minimizing impacts on smaller refineries, exemptions can also shift the burden of meeting the volumetric mandates to larger refineries if they are issued before the final rule (Coppess and Irwin, 2017). Exemptions hit a low point in 2015, with only 7 petitions granted representing 3 billion gallons of gasoline, as compared to 29 petitions granted representing over 13 billion gallons just two years later (Environmental Protection Agency, 2018). Of course, political connections, easier access to credit, and other factors benefiting large firms may mean that smaller firms still bear the brunt of the regulation.

Large firms also tend to have integrated downstream operations, allowing them to blend their own ethanol to generate RINs rather than being required to buy them from separated downstream blenders. As discussed in the introduction, this is the basis for a perennial complaint by small refineries, who argue that integrated operations profit from their ability to generate and sell their excess RINs.

All sorts of heterogeneities in firm operations are of interest to regulatory agencies because differences in cost structures could counter-intuitively incite lower-cost firms to support the law as a new, artificial source of comparative advantage (Salop and Scheffman, 1983). My reduced-form estimates of how firms respond to positive price changes in RFS compliance costs can lend credence to or discredit some of these theories.

2.2 2015 RIN Shocks

As mentioned previously, the EPA sometimes misses regulatory announcement deadlines, resulting in major regulatory announcements that occur mid-compliance year. In 2013, the EPA released the final rule on August 6, leaked a draft proposal for 2014 on October 11, and released 2014's official proposed rule on November 15. Lade et al. (2018) measures how 2013's mid-year announcements influenced RIN prices themselves, related commodity markets, and the stock prices of biorefining firms. They find that those shocks, which drove the prices of all RIN varieties downward, mainly impacted commodities and firms tied to advanced biofuels, which they considered to be the marginal compliance fuel. The year 2013 was a volatile time for RIN markets, and visual inspection of RIN price histories throughout 2013 in Figure 3 reveals that the announcement dates analyzed by Lade et al. (2018), while significant, are unremarkable compared to baseline volatility.

Unlike 2013, 2015's proposed and final rules, which were released on May 29 and November 30, are clearly and unambiguously apparent by visual inspection of Figure 4.[2] Surrounding a proposed rulemaking on May 29, 2015, the price of a D6 RIN dropped from \$0.69 one week before to \$0.3775 one week afterwards, a drop of more than 45%. A jump of similar magnitude surrounded a final rulemaking on November 30, 2015.

In this study, we focus on 2015 because the EPA made two major policy announcements that we will exploit as structural breaks following Lade et al. (2018). We choose 2015 over 2013, the primary year analyzed by Lade et al. (2018), for two reasons: First, the announcement breaks are much clearer in the data; that can be seen by comparing Figures 3 and 4, but note the change in the extent of the price axis. Second, RIN prices are sufficiently high throughout 2015 to guarantee that the mandate was binding, whereas early 2013 prices were so low that the RFS

[2] In an online supplement, Lade et al. (2018) repeat the portion of their analysis detailing how policy announcements affect RIN prices for 2015 event dates. Unlike their 2013 analysis, they do not examine how 2015 announcements affected biofuel stocks or commodity markets.

Figure 3: Price Histories for 2013 Vintage RINs

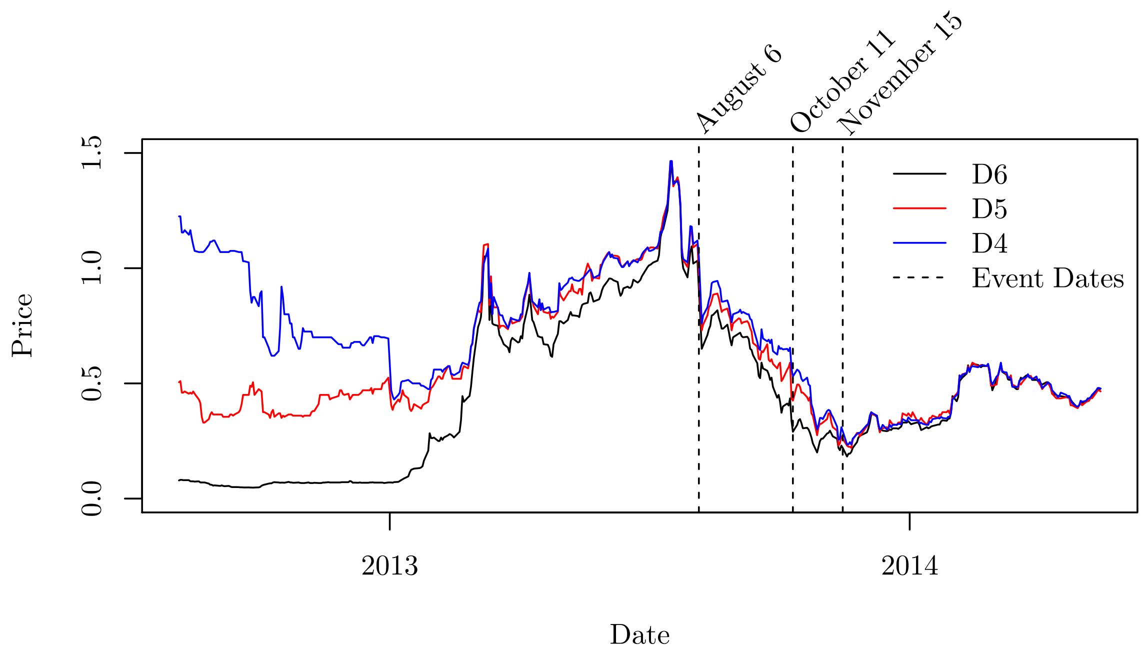

may not have actually been stringent enough to alter refiner behavior (Whistance et al., 2016).

2.3 Identification of RFS Industry Impact

This paper employs two methodologies to study the impact of the Renewable Fuel Standard on the refining industry. First, we use multivariate time series methods to quantify the response of refining firm values to shocks in RIN price series. In particular, for every Firm × RIN combination, we estimate impulse response functions using vector auto-regressions and vector error correction models. These allow us to determine the dynamic response of the value of refining firms to shocks in the value of RIN prices. This portion of our analysis builds on past literature using time series methods to unravel how RINs transmit to wholesale and consumer fuels (Knittel et al., 2017) and biofuel-related commodities (Whistance and Thompson, 2014; Whistance et al., 2016). Second, we implement an event study methodology similar to that used by Lade et al. (2018). Even if endogeneity complicates measurement of direct industry impact using conventional multivariate time series methods, we anticipate that the large shocks constituting nearly half of RIN prices should induce measurable impacts on refining firms if indeed there is an impact. The event study methodology also side-steps nonlinearities in the

Figure 4: Price History for 2015 Vintage RINs

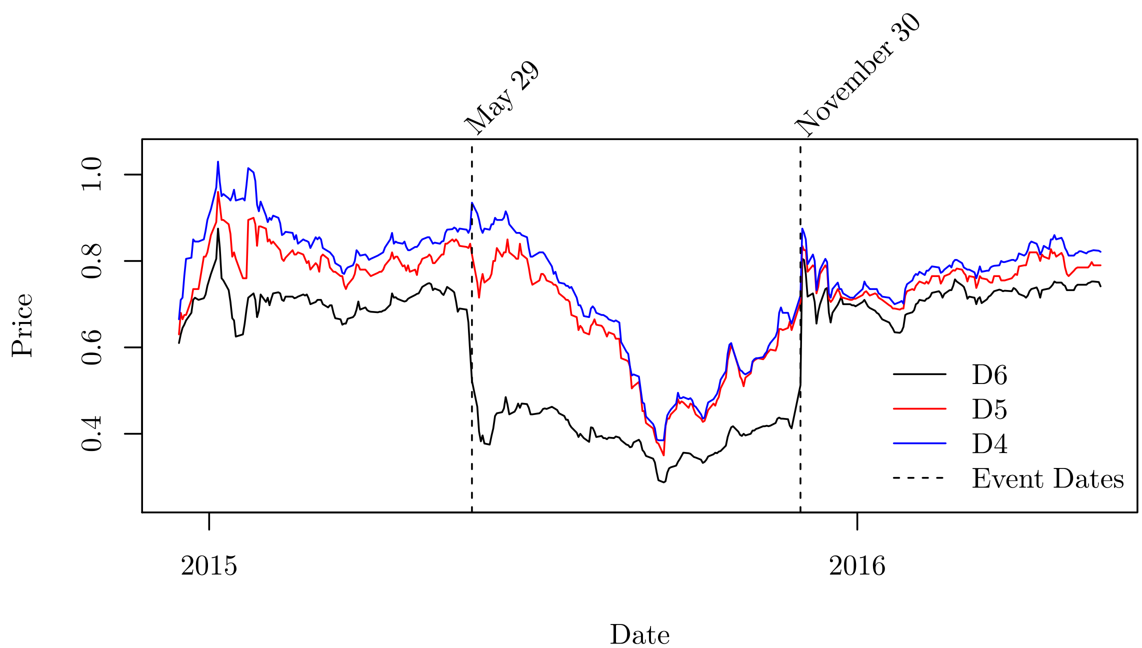

RIN-stock system that may complicate the bivariate models (Serra et al., 2011).

Both methodologies are reduced-form. By examining stock prices rather than specific product prices, we should be able to identify the impact of *any* channel whereby a larger RFS mandate influences the net values of refining firms in the short-run. The downside of this approach is that we cannot validate any particular channel of causality. Nevertheless, establishing the existence or non-existence of an effect in and of itself offers insights as to whether cost pass-through is an adequate stopping point for research seeking to understand how regulated industries respond to the RFS and other fuel blend regulations.

3 Data

As described in the introduction, RINs are compliance credits that refiners obligated by the RFS can use to meet their biofuel utilization mandates. Because multiple, nested categories exist within the RFS mandate, there are four commonly traded types of RINs, as shown in Figure 2.

The mandate for D3 RINs is consistently minuscule, and markets for those RINs are commensurately quite thin. Therefore, we restrict our analysis to D6, D5, and D4 RINs. RIN series also differ by their year of creation; apart from some flexibility allowing limited inter-year banking and borrowing of RIN stocks, RINs are mostly used to comply with the mandate for the year in which they were generated. Thus, we limit our analysis to RINs of a single compliance year.[3] As discussed in Section 2.2, we analyze data from 2015. Because of some pre-year trading and the fact that RINs are not actually submitted for compliance until a few months after year end, our data does extend a little beyond a single calendar year.

All data is daily, excluding weekend and other financial market closures such as holidays. Data for RIN series come from the Oil Price Information Service (OPIS), a company that provides pricing for numerous petroleum products and is a frequent supplier of RIN data for economic researchers. Stock prices for nearly all firms are adjusted prices from Yahoo! Finance (accessed via the quantmod package in R). Stocks delisted since 2015 are not available on Yahoo! Finance and are instead sourced from Bloomberg. While these prices are not adjusted, none of those firms underwent stock splits or reverse-splits in the study period, so this should be of little consequence.

3.1 Firm Characteristics

This study encompasses all publicly traded firms with at least 200,000 barrels per day of refining capacity as of January 1, 2015, plus Western Refining (Energy Information Administration, 2015). We sort firms along two dimensions: size and exposure to RFS regulatory costs as of the beginning of 2015. We use market capitalization as a relevant measure of firm size (Fama and French, 1992) and the percentage of a firm's refining capacity in PADD 2 and Alaska as our measure of exposure (or, rather, non-exposure).[4]

PADDs are collections of states by which many government-issued petroleum data sources are aggregated. The Midwest, encompassed by PADD 2, contains the vast majority of ethanol refining capacity because refinery location decisions are driven primarily by access to feedstocks (Lambert et al., 2008). Because ethanol cannot be transported using existing pipeline infras-

[3]The majority of the literature examining RIN prices do not restrict their analyses to a single compliance year. It is not always clear how this literature stitches together price series for RINs belonging to difference compliance years, especially when RINs for multiple compliance years are sold contemporaneously.

[4]PADD stands for Petroleum Administration for Defense Districts. Using refining capacity in PADDs 2, 3, and Alaska or, inversely, using refining capacity in PADDs 1 and 5 minus Alaska, hardly changes the groupings.

tructure, it is plausible to think that the RFS puts refineries located near biorefineries at a comparative advantage. Geographic disparity in the impacts of the RFS have already been documented for retail gasoline prices (LaRiviere et al., 2017) and pass-through of RIN prices by rack sellers at fuel terminals (Pouliot et al., 2017). Alaskan refining capacity is also included in this measure because Alaska is exempted from the RFS. We sort firms into one of nine bins, one for each element of the cross product of three market capitalization and three RFS exposure bins, as shown in Table 1. The cutoffs were chosen somewhat arbitrarily to result in balanced bins, but, as the next paragraph further explains, the cutoffs do run parallel to important qualitative differences between firms.

Firms characterized as large (market capitalizations in excess of $100B) are also the US firms generally considered to be large, integrated refineries: BP, Shell, Chevron, and ExxonMobil. In mid-2015, an analysis released by Valero Energy Corporation argued that those firms were already separating RINs in excess of their own RFS obligations, yielding them “windfall profits,” because their sales of branded gasoline made up large percentages of their refinery production (Valero Energy Corporation, 2015). Patterns of price response among these energy firms could lend evidence for or against those conjectures. The line between small and medium refining firms also separates relatively smaller nationwide operations from merely regional companies.

Table 1: Firm Characteristics with Ticker Symbols

	100% Exposed	<100% & >70% Exposed	<70% Exposed
Large: >$100B	Shell (RDS.A) Chevron (CVX) Total (TOT)	Exxon Mobil (XOM)	British Petroleum (BP)
Medium: < $100B, > $10B		Valero (VLO) Phillips 66 (PSX)	Marathon (MPC)
Small: <$10B	Carlyle Group (CG) Western Refining (WNR)	Andeavor (ANDV)	HollyFrontier (HFC)

3.2 Stationarity Tests

Both methods of our analysis require stationary data inputs, so we begin by testing for stationarity in our data. Tests for stationarity are known to be biased towards conclusions of non-stationarity in the presence of structural breaks (Perron, 1989). However, despite evidence indicating the presence of structural breaks in RIN price series (Mason and Wilmot, 2016; Lade et al., 2018), no prior research to our knowledge considers structural breaks in tests of RIN price stationarity, including RFS research specifically related to large breaks, such as Lade et al. (2018).

Most popular structural break tests allow for endogenous breakpoint selection, reflecting that most structural break research requires first identifying the exact location of the breakage. We have no need for that—we know the precise dates of the breaks ex ante—and running those tests would sacrifice power unnecessarily. Lee and Strazicich (2003) depart from that norm and offer a stationarity test (hereafter the LS test) that allows for up to two exogenously defined structural breaks.[5]

The LS test allows for breaks in both mean and trend, using a vector of structural break variables for breaks occurring at $t = T_1$ and $t = T_2$: $Z_t = [DU_{1,t}, DU_{2,t}, DT_{1,t}, DT_{2,t}]'$, where $DU_{i,t} = 1$ for $t \geq T_i + 1$ and zero otherwise and $DT_{i,t} = t - T_i$ for $t \geq T_i + 1$ and zero otherwise for $i = 1, 2$. The LS two-break unit root test is estimated by the equation

$$\Delta y_t = \delta' \Delta Z_t + \phi \tilde{S}_{t-1} + u_t, \tag{1}$$

where y_t is the variable whose stationarity is being tested; Δ is the difference operator; $\tilde{S}_t = y_t - \tilde{\psi}_x - Z_t\tilde{\delta}$ for $t \geq 2$; $\tilde{\delta}$ are coefficients recovered from a regression of Δy_t on ΔZ_t (i.e., $\Delta y_t = \delta' \Delta Z_t + e_t$); $\tilde{\psi}_x = y_1 - Z_1\tilde{\delta}$; and u_t satisfies the normality conditions defined by Phillips and Perron (1988, p. 336). The test statistic for the unit root null hypothesis is the t-statistic on the parameter estimate for ϕ, and significance is checked using critical values for an exogenous 2-break unit root test reported in Lee and Strazicich (2003).[6]

For testing stationarity of our non-RIN series, the LS test conveniently exhibits nice size and power properties even when the actual data generating process contains no breaks. Results of

[5]An endogenous version of the same test is also described by Lee and Strazicich (2003) and is the test most frequently associated with the paper.

[6]Lee and Strazicich (2003) only report critical values for tests with a sample size of 100, whereas our sample size is 357. This will result in the test being underpowered.

the LS test are reported in Table 2 alongside the typical Augmented Dickey-Fuller (ADF) and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) tests, which do not account for structural breaks. The three tests conclude unanimously that every RIN and stock price series is non-stationary.

4 Multivariate Time Series Analysis

To understand the relationship between RIN prices and refining firm values, we fit bivariate time series models for each RIN $\times$ Firm pair in our data. For cointegrated series, we estimate vector error correction models (VECM), otherwise we estimate vector autoregressive models (VAR) in differences to account for the aforementioned non-stationarity of our data series. Cointegration offers a systematic methodology to deal with non-stationary data in a multivariate context. Specifically, if a pair of non-stationary series are cointegrated, then their linear combination produces a stationary series (Engle and Granger, 1987). After fitting the VAR and VECM models, we analyze the significance of lagged values of RIN prices on the stock values of refining firms. Using those results, we visualize the effect of exogenous shocks to RIN prices on refinery values using impulse response functions.

4.1 Cointegration Tests

To select an appropriate bivariate time series model for a pair of non-stationary series, we first need to establish whether they are cointegrated (Engle and Granger, 1987). While cointegration tests are affected by structural breaks, structural breaks have little impact on the size or power of Johansen's cointegration test (Campos et al., 1996). For that reason, we elect to use Johansen's maximum eigenvalue and trace tests without modification (Johansen, 1991) to evaluate the relationship between each pairwise combination of RIN series and firms.[7] Table 3 reports the results of these tests.

The tests for cointegration indicate that there is some cointegrating relationships between RIN prices and refinery stock prices. The only two significant cointegrating relationships between both the maximum eigenvalue and trace tests are biomass-based biodiesel RINs and Exxon Mobil (XOM) and Chevron (CVX) stock prices. Both Chevron and Exxon Mobil are large companies

[7]Johansen et al. (2000) introduce a modification to the canonical cointegration test that explicitly allows for structural breaks. Future refinements to this section may utilize that methodology.

Table 2: Summary Statistics and Stationarity Tests

	Summary Statistics			Stationarity Tests		
	Mean	St. Dev	Obs.	ADF	KPSS	LS
Null				Non-stationary	Stationary	Non-stationary
D6	0.599	0.155	357	-1.557	1.425***	-1.931
D5	0.721	0.119	357	-1.630	1.420***	-1.019
D4	0.759	0.134	357	-1.491	1.777***	-0.255
VLO	53.780	5.633	357	-2.421	3.164***	-1.230
MPC	42.283	5.931	357	-2.395	2.173***	-1.117
XOM	72.530	4.362	357	-1.004	1.674***	-3.204
PSX	71.721	5.792	357	-2.940	3.688***	-1.163
CVX	82.563	8.145	357	-1.046	2.526***	-2.219
BP	28.244	3.072	357	-2.212	4.593***	-1.033
HFC	35.932	5.337	357	-1.429	1.438***	-2.250
RDS.A	43.017	4.934	357	-2.308	5.253***	-2.790
CG	16.077	3.596	357	-1.646	5.500***	-0.494
TOT	39.943	2.310	357	-2.905	1.598***	-2.730
ANDV	90.745	11.316	357	-2.119	1.396***	-2.042
WNR	39.370	7.515	357	-1.793	3.760***	-1.784

Note: Stationarity test column headers are acronyms for the Augmented Dickey-Fuller, Kwiatkowski-Phillips-Schmidt-Shin, and Lee-Strazicich stationarity tests. The Lee-Strazicich test allows for the presence of two structural breaks without being biased towards non-stationarity (Lee and Strazicich, 2003). D6, D5, and D4 are RIN varieties, all other variables are stock tickers corresponding to firms as in Table 1. Significance at alpha levels of 0.1, 0.05, and 0.01 are reported with *, **, and ***, respectively.

Table 3: Cointegration Tests

	Maximum Eigenvalue			Trace		
Firm	D6	D5	D4	D6	D5	D4
VLO	8.892	8.961	9.549	12.524	11.781	12.005
MPC	7.674	4.508	3.984	9.265	6.330	5.844
XOM	6.709	14.762*	19.877**	8.378	16.020	22.400**
PSX	9.901	10.607	10.719	14.258	15.047	14.014
CVX	6.628	14.909*	18.066**	8.311	16.990	20.659**
BP	3.477	3.665	4.821	5.910	5.551	6.971
HFC	15.759**	7.838	6.408	16.815	8.646	7.243
RDS.A	4.424	5.410	4.929	6.779	7.661	6.891
CG	8.016	4.898	4.721	9.941	6.399	5.901
TOT	9.546	12.623	13.839*	13.099	15.174	16.523
ANDV	10.986	8.514	8.182	13.629	10.540	9.960
WNR	6.037	3.999	3.261	6.677	4.685	3.916

Note: Maximum eigenvalue and trace tests are Johansen's (1991) two methodologies of testing for cointegration. D6, D5, and D4 are RIN varieties, row names are stock tickers corresponding to firms as in Table 1. Significance at alpha levels of 0.1, 0.05, and 0.01 are reported with *, **, and ***, respectively.

with market capitalizations above \$100 billion and relatively high "exposure" to the RFS, as defined in Section 3.1. Chevron lacks any refining capacity whatsoever in Alaska and PADD 2, and only 12% of Chevron's capacity is located there.

However, many other firms categorized as highly exposed to the RFS still lack cointegrating relationships with RINs. HollyFrontier Corporation (HFC), whose stock is cointegrated with D6 RINs according to the maximum eigenvalue test, is a small company with relatively little RFS exposure.

4.2 VAR and VECM Modeling

With cointegration results in hand, we model every Firm × RIN pair using either a reduced-form bivariate VAR model in differences (for non-cointegrated series) or a bivariate VECM model (for series that are cointegrated at a significance level of 0.05 in either test). Following Sims (1980), the VARs are specified in differences and have the following form:

$$\begin{aligned}\Delta FIRM_t &= c_1 + \sum_{l=1}^{m}(\phi_{1,1}^{l}\Delta FIRM_{t-l} + \phi_{1,2}^{l}\Delta RIN_{t-l}) + e_{1,t}\\ \Delta RIN_t &= c_2 + \sum_{l=1}^{m}(\phi_{2,1}^{l}\Delta FIRM_{t-l} + \phi_{2,2}^{l}\Delta RIN_{t-l}) + e_{2,t},\end{aligned} \tag{2}$$

where $\Delta FIRM_t$ and ΔRIN_t are differenced values for a RIN and firm stock price at time t respectively; c_i is a regression constant for regression $i \in \{1,2\}$; m is the VAR lag length determined based on the Akaike Information Criterion (AIC); each $\phi_{i,j}^{l}$ is an autoregressive coefficient for the lth lag of either $\Delta FIRM_t$ or ΔRIN_t; and $e_{i,t}$ is the VAR error term for $i \in \{1,2\}$. In particular, we are interested in the statistical significance of $\phi_{1,2}^{l}$, the conditional effect of ΔRIN_{t-l} on $\Delta FIRM_t$, at all lags against the null hypothesis of $\phi_{1,2}^{l} = 0$. Significance of one or more lags would indicate that there is a feedback relationship between the two series, and hence RIN price changes are associated with subsequent firm stock price changes.

Following Johansen (1991), we use VECM models for the cointegrated series, which are expressed as:

$$\begin{aligned}\Delta FIRM_t &= c_1 + \pi_{1,1}FIRM_{t-1} + \pi_{1,2}RIN_{t-1} + \sum_{l=1}^{m}(\phi_{1,1}^{l}\Delta FIRM_{t-l} + \phi_{1,2}^{l}\Delta RIN_{t-l}) + e_{1,t}\\ \Delta RIN_t &= c_2 + \pi_{2,1}FIRM_{t-1} + \pi_{2,2}RIN_{t-1} + \sum_{l=1}^{m}(\phi_{2,1}^{l}\Delta FIRM_{t-l} + \phi_{2,2}^{l}\Delta RIN_{t-l}) + e_{2,t}\end{aligned} \tag{3}$$

The VECM specification is similar to the above VAR model except for the inclusion of the error correction terms (the *undifferenced* single-lag values of the component series, i.e. $\Delta FIRM_t$ and ΔRIN_t), which captures the long-run equilibrium relationship between the two series. The error correction term allows multivariate estimation of two series with a common stochastic trend that otherwise would cause omitted variable bias. Significance of the error correction terms merely validates that a simple VAR model that did not account for the series' cointegration would be flawed. Again, for the purposes of this research, the values of $\phi^l_{1,2}$ at all lags are the parameters whose significance is of interest.

These models illuminate interactions between RIN markets and firm stock prices. Because the relationship of interest is the impact of the RFS on refining firms, the analysis and results reported below will focus on the differenced RIN lag parameters and error correction terms in the equation modeling movements in firm stock prices, i.e. the first equations in Equations 2 and 3. Tables 4, 5, and 6 report parameter estimates and significance levels for the autoregressive coefficients $\phi^l_{1,2}$, regression constants, and, when applicable, an error correction term coefficient for 36 total regressions covering every Firm $\times$ RIN pair.

Also reported in the aforementioned tables are observation counts (varying slightly between models depending on lag lengths) and test statistics for three tests of model assumptions. Included among these three are the Ljung-Box test for autocorrelation among the residuals and the Jarque-Bera and Shapiro-Wilk tests for normality of the error term.

4.3 Results

The results from bivariate time series analysis suggest that the impacts of RIN price movements are not homogeneous across firms. In fact, most small- and medium-size firms values' are not impacted by the lagged values of RIN prices. By contrast, large firms show negative reactions to RIN price movements after a few lags. This is especially evident in results from D5 bivariate modeling in Table 5, where Chevron (CVX), British Petroleum (BP), Shell (RDS.A), and Total (TOT) all have significant negative fourth lag terms for D5 RINs. Exxon Mobil (XOM), the only remaining large market cap firm in the data, is also negatively associated with D5 RINs, but at the fifth lag. This firm size pattern is not an artifact of selecting lag length by AIC; forcing all bivariate models to include six lags does not result in any more significant fourth lag terms in the D5 models.

Table 4: Bivariate Time Series Model with D6 RINs

	VLO	MPC	XOM	PSX	CVX	BP	HFC	RDS.A	CG	TOT	ANDV	WNR
Constant (c_1)	0.025	-0.008	0.007	0.028	-0.001	-0.005	2.485***	-0.021	-0.011	0.001	0.010	-0.039
EC Term ($\pi_{1,2}$)							-1.558***					
Lag 1 ($\phi^1_{1,2}$)	-2.336	-0.662	-1.615	-1.821	-1.188	-0.864	1.681	-0.587	-0.299	-1.093	5.508	-2.566
Lag 2 ($\phi^2_{1,2}$)	2.349	-1.883	0.421	1.365	1.136	0.946		4.088*	1.122	0.682	-5.718	1.529
Lag 3 ($\phi^3_{1,2}$)			0.884		-2.156	-2.469*		-4.902**				
Lag 4 ($\phi^4_{1,2}$)			-4.842*									
Observations	354	354	352	354	353	353	355	353	354	354	354	354
Ljung-Box	0.002	0.002	0.064	0.000	0.111	0.108	0.001	0.148	0.012	0.015	0.008	0.000
Jarque-Bera	129.692***	36.575***	23.626***	23.241***	2.719	26.951***	26.656***	3.127	16.531***	3.084	175.717***	5.096*
Shapiro-Wilk	0.962***	0.981***	0.984***	0.986***	0.990**	0.980***	0.981***	0.996	0.986***	0.995	0.964***	0.995

Note: Column headers are stock tickers corresponding to firms as in Table 1. Each column reports results from either a VAR or VECM model as described in Equations 2 and 3. VECM models include an error correction (EC) term ($\pi_{1,2}$). Optimal lag order for either model is determined based on AIC. Significance at alpha levels of 0.1, 0.05, and 0.01 are reported with *, **, and ***, respectively.

Table 5: Bivariate Time Series Model Results with D5 RINs

	VLO	MPC	XOM	PSX	CVX	BP	HFC	RDS.A	CG	TOT	ANDV	WNR
Constant (c_1)	0.025	-0.009	0.004	0.029	-0.005	-0.005	-0.019	-0.024	-0.013	0.003	0.010	-0.038
EC Term ($\pi_{1,2}$)												
Lag 1 ($\phi^1_{1,2}$)	-1.417	-1.627	1.508	-2.779	1.233	0.243	0.125	1.432	0.937	-0.839	2.825	-1.970
Lag 2 ($\phi^2_{1,2}$)	0.790	-0.520	-2.788	1.340	-2.178	-2.102	-1.583	-1.205	0.713	-0.437	−2.083	0.385
Lag 3 ($\phi^3_{1,2}$)			-0.411		-3.352	-0.620		-1.936	-0.766	-2.909		
Lag 4 ($\phi^4_{1,2}$)			-4.302		-9.040**	-4.126***		-7.595***		-5.047***		
Lag 5 ($\phi^5_{1,2}$)			-4.705*									
Lag 6 ($\phi^6_{1,2}$)			3.624									
Observations	354	354	350	354	352	352	354	352	353	352	354	354
Ljung-Box	0.001	0	0	0	0.053	0.004	0.001	0.005	0.025	0	0.003	0.001
Jarque-Bera	129.568***	42.073***	24.033***	23.747***	1.571	20.588***	28.044***	0.818	14.22***	0.475	170.158***	5.075*
Shapiro-Wilk	0.963***	0.98***	0.984***	0.986***	0.992**	0.984***	0.981***	0.997	0.988***	0.997	0.965***	0.994

Note: Column headers are stock tickers corresponding to firms as in Table 1. Each column reports results from either a VAR or VECM model as described in Equations 2 and 3. VECM models include an error correction (EC) term ($\pi_{1,2}$). Optimal lag order for either model is determined based on AIC. Significance at alpha levels of 0.1, 0.05, and 0.01 are reported with *, **, and ***, respectively.

Table 6: Bivariate Time Series Model Results with D4 RINs

	VLO	MPC	XOM	PSX	CVX	BP	HFC	RDS.A	CG	TOT	ANDV	WNR
Constant (c_1)	0.025	-0.008	1.246	0.029	1.163	-0.004	-0.019	-0.024	-0.010	0.001	0.012	-0.038
EC Term ($\pi_{1,2}$)			-0.029		-0.275							
Lag 1 ($\phi^1_{1,2}$)	-4.469	-2.723	-1.065	-5.874	-1.833	-0.363	-2.590	-0.580	-1.096	-1.885	2.311	-3.943
Lag 2 ($\phi^2_{1,2}$)	1.894	-1.110	3.122	3.556	-0.141	0.113	-0.444	3.335	0.688	3.720	-5.413	-0.298
Lag 3 ($\phi^3_{1,2}$)			-1.157		-5.226			-3.746		-5.41**		
Lag 4 ($\phi^4_{1,2}$)			-3.751		-6.417			-6.528***				
Lag 5 ($\phi^5_{1,2}$)					-6.766							
Lag 6 ($\phi^6_{1,2}$)					-0.547							
Lag 7 ($\phi^7_{1,2}$)					8.044*							
Lag 8 ($\phi^8_{1,2}$)					-7.370							
Lag 9 ($\phi^9_{1,2}$)					4.175							
Lag 10 ($\phi^{10}_{1,2}$)					10.561**							
Observations	354	354	352	354	346	354	354	352	354	353	354	354
Ljung-Box	0	0	0.029	0	0.045	0.008	0.001	0	0.008	0.127	0.004	0.001
Jarque-Bera	131.37***	41.336***	23.55***	24.843***	1.039	32.193***	28.844***	0.382	22.297***	1.773	170.472***	5.553*
Shapiro-Wilk	0.962***	0.98***	0.984***	0.986***	0.994	0.979***	0.981***	0.999	0.985***	0.996	0.965***	0.994

Note: Column headers are stock tickers corresponding to firms as in Table 1. Each column reports results from either a VAR or VECM model as described in Equations 2 and 3. VECM models include an error correction (EC) term ($\pi_{1,2}$). Optimal lag order for either model is determined based on AIC. Significance at alpha levels of 0.1, 0.05, and 0.01 are reported with *, **, and ***, respectively.

Results across the D6 and D4 models are largely consistent with the D5 results, with no significant lags among small and medium firms, and significant parameters among most of the large firms after a few lags. In the D6 results, Shell (RDS.A), British Petroleum (BP), and Exxon Mobil (XOM) all have negative third or fourth lag terms, though Exxon Mobil and British Petroleum's are only marginally significant. Total (TOT) and Shell experience significant negative responses in the D4 models, and Chevron (CVX) and Exxon Mobil have negative responses of similar magnitudes at the same lag lengths, though these are not statistically significant. Chevron's D4 model exhibits significant responses in the other direction at even later lags, but these are erratic and likely spurious. Though the overlap between model results is not perfect, it is striking that across all RIN types, there are significant negative responses at lags three and four among large firms only. The consistency of sign and timing among these responses suggest an underlying pattern.

Figure 5 reports a subset of the orthogonal impulse response functions (IRFs) for refinery stock responses to exogenous shocks in D5 RINs. Results from the shocks applied are in terms of the differenced stock values and are accompanied by their respective error band. Based on the IRFs, it is clear that a shock to D5 RINs generally produces a significant decrease in the value large refineries' stocks (left panel of Figure 5). Specifically, a unit increase in the cost of D5 RINs is associated with negative stock price movements, hitting a minimum between 4-5 days after the shock with single-day losses between 5 and 20 cents at that time for Chevron, Exxon Mobil, and British Petroleum. After these significant lags, point estimates vary tightly around zero before converging towards a long-run equilibrium. Though the economic significance of these estimates is small (cumulative losses below a dollar for stocks worth over $40), patterns of statistical significance are consistent. For small and medium size firms (right panel of Figure 5), there is no real pattern of responses to RIN prices, convergence happens more quickly, and error bands preclude statistical significance by a wide margin.

The fact that our most significant results come from the D5 models comport with results reported by Lade et al. (2018), who find that advanced biofuels were the marginal compliance fuel and that advanced biofuel firms and commodities reacted most severely to 2013's shocks. Use of biodiesel and advanced ethanol as marginal compliance fuels is driven by the fact that the RFS is a nested mandate (allowing more advanced fuels to meet conventional mandates as well) and that conventional ethanol blends are limited by regulations on maximum ethanol concentrations,

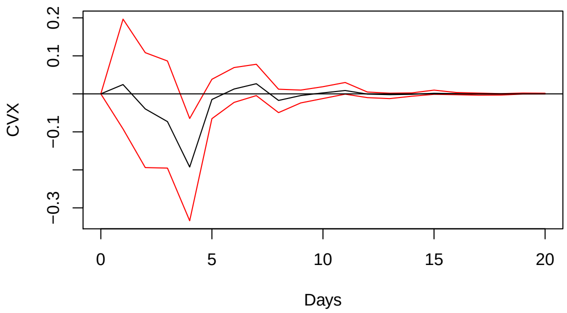

(a) Chevron

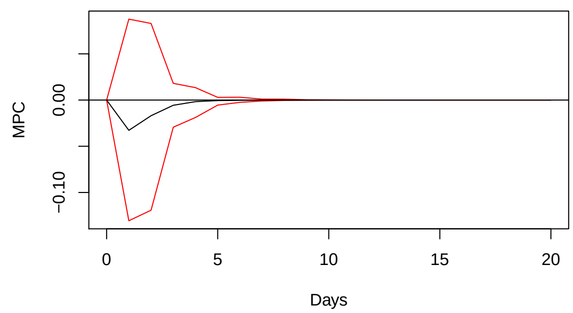

(b) Marathon Petroleum

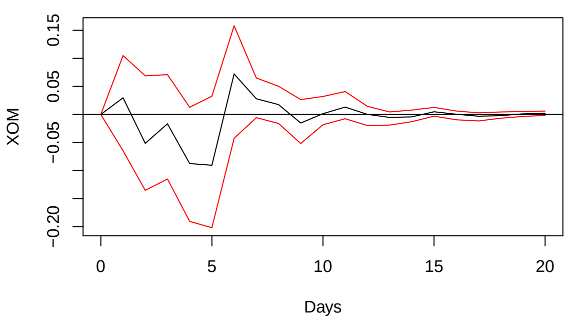

(c) Exxon Mobil

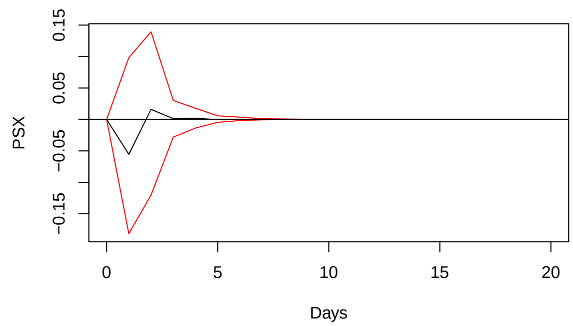

(d) Phillips 66

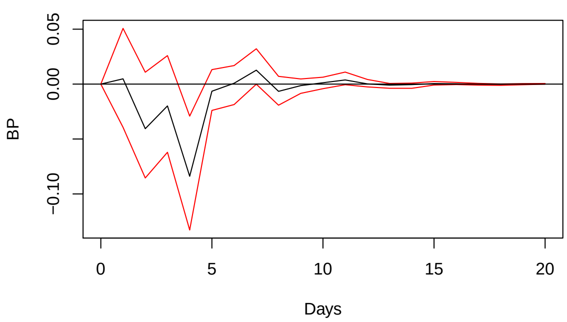

(e) British Petroleum

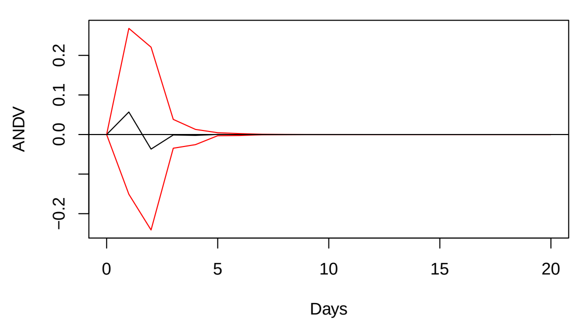

(f) Andeavor (Tesoro)

Figure 5: Refinery Stock Response to a D5 RIN Shock, Selected Large Firms (Left) and Small/Medium Firms (Right)

which creates the "blend wall," the point past which refiners cannot mix additional ethanol into most consumer fuel blends.

Some readers may question the choice of a bivariate model, given that numerous other financial series influence the fundamental value of both RINs and refining stocks. We offer a few brief responses to this objection. First, our data is limited, and each additional covariate eliminates disproportionately more degrees of freedom in VAR and VECM models. Second, the non-linear relationship between commodities and intermediate products and RINs or firm values further complicates modeling decisions and eats away at degrees of freedom. Third, whatever limitations the exclusion of these covariates imposes on the interpretation of our model results are much less of a concern for the event study modeling, which largely validates the result of the bivariate models. Nonetheless, we do repeat our analysis including a continuous front-month future price for ethanol as a third variable. The addition of this third variable penalizes the AIC enough that lag selection generally does not extend to the fourth lag, but when we coerce the lag length to match the optimum chosen in the bivariate case, almost the exact same parameters are significant and their signs are the same.

Our Ljung-Box test statistics fail to reject the null of serial independence for every bivariate model. The Jarque-Bera and Shapiro-Wilk test statistics, however, suggest substantial non-normality among some of our error terms. We consider this a minor issue, given the lack of serial dependance, reasonably large sample size, and the fact that many of our most notable results (such as the D5 and D4 models of Shell and Total) come from models that do not exhibit normality issues.

Before discussing the results of these models further, we conduct the event study analysis to see whether results from an alternative specification corroborate these findings.

5 Event Study Analysis

While multivariate time series models are suitable for characterizing everyday interactions between variables, event study models effectively characterize short-run, one-off responses to large market shocks. Regulatory changes that precipitate massive swings in compliance credit values offer such a market shock. Event studies of compliance credit shocks can be decomposed to study regulation's distributional effects, as in Bushnell et al. (2013), who counterintuitively show that stock prices of carbon-intensive firms actually fell in response to a slump in EU carbon prices.

In the context of the RFS, Lade et al. (2018) apply event studies to 2013 shocks to examine the policy's effect on commodity markets and biorefining firms.

Similar to Lade et al. (2018), the specification of our event study model takes the following form:

$$\Delta \ln(Y_{i,t}) = \beta_0 + \sum_{i=1}^{p} \beta_i t^i + \sum_{i=1}^{s} \sum_{i=0}^{m} \gamma_{s,m} 1(t \in \{\mathbb{T} + m\}) + \boldsymbol{\Theta}' \boldsymbol{\Delta} \ln(\mathbf{X_t}) + \lambda_{\text{MoY}} + \lambda_{\text{DoW}} + e_{i,t}, \quad (4)$$

where $\Delta \ln(Y_{i,t})$ is the log-differenced value of a refiner i's stock at time t; β_0 is a regression constant; the remaining β_i's are parameters on time controls of polynomial order p; $\gamma_{s,m}$ are parameters on dummies $1(t \in \{\mathbb{T} + m\})$ for the mth lag after each event in $\mathbb{T}$; s is the number of distinct events in $\mathbb{T}$; $\boldsymbol{\Theta}$ is a vector of parameters on log-differenced "normal return" controls $\boldsymbol{\Delta} \ln(\mathbf{X_t})$; λ's are month-of-year (MoY) and day-of-week (DoW) fixed effects; and $e_{i,t}$ is the disturbance.

Subjective modeling decisions were made to mirror Lade et al. (2018) as closely as possible. Specifically, we use the same polynomial of time controls, the same fixed effects, and the same normal return control: the RUS3000 index. Our one deviation is the length of lags we consider; given that our bivariate time series models suggested there may be some significant responses after the fourth lag, we set $m = 6$. The model is estimated on the entire sample of firm values jointly, then separately for each bin of firms by market capitalization.

The results from event study analysis are reported in Table 7. As in Lade et al. (2018), p-values are based on sample quantile tests for significance designed to accommodate event studies with few firms, as discussed in Gelbach et al. (2013).[8] The basic idea behind this test is to use the empirical cdf of non-event date residuals as the distribution for hypothesis testing, since event date parameters are identified off of variation within a single day.

Patterns of statistical significance in the event studies closely mirror results from the bivariate time series models. It is apparent that the first event does not register for any firms in the event studies. From Figure 4, the first event date seems to affect D6 RINs most dramatically. The effects on D5 and D4 RINs, however, are smaller and seemingly opposite to one another. If biodiesel and advanced ethanol really are the marginal compliance fuels under the RFS, as Lade et al. (2018) reveals and our bivariate time series results confirm, the patterns in the RINs following the first event should not be expected to generate significant responses in firm value.

[8] More information on the test in the context of difference-in-differences modeling is available in Conley and Taber (2011).

Table 7: Results from Event Studies

	Large Firms	Medium Firms	Small Firms	All Firms
Event 1				
Lag 0 ($\gamma_{1,0}$)	0.004	0.013	0.009	0.008
Lag 1 ($\gamma_{1,1}$)	-0.002	-0.009	-0.010	-0.006
Lag 2 ($\gamma_{1,2}$)	0.013	-0.012	-0.020	-0.004
Lag 3 ($\gamma_{1,3}$)	-0.003	-0.020	-0.016	-0.012
Lag 4 ($\gamma_{1,4}$)	0.000	-0.011	-0.004	-0.004
Lag 5 ($\gamma_{1,5}$)	0.003	0.003	0.006	0.004
Lag 6 ($\gamma_{1,6}$)	0.011	0.000	0.005	0.006
Event 2				
Lag 0 ($\gamma_{2,0}$)	0.011	-0.006	-0.007	0.001
Lag 1 ($\gamma_{2,1}$)	-0.012	-0.001	0.006	-0.003
Lag 2 ($\gamma_{2,2}$)	-0.016	-0.003	0.000	-0.008
Lag 3 ($\gamma_{2,3}$)	0.000	-0.017	-0.008	-0.007
Lag 4 ($\gamma_{2,4}$)	-0.032**	-0.005	-0.017	-0.020
Lag 5 ($\gamma_{2,5}$)	-0.025*	-0.013	0.002	-0.013
Lag 6 ($\gamma_{2,6}$)	-0.009	-0.001	-0.007	-0.006
Number of Firms	5	3	4	12
Observations	1780	1068	1424	4272

Note: Each column reports selected parameter estimates and significance levels for an estimation of Equation 4 on a subset of firms. See Table 1 to examine the firms in each size bin. Hypothesis testing is conducted using the sample quantile test described in Gelbach et al. (2013). Significance at alpha levels of 0.1, 0.05, and 0.01 are reported with *, **, and ***, respectively.

In contrast, Figure 4 shows that the second event has a much clearer and consistently positive price impact on all RINs. As in the bivariate time series models, firms with small and medium market capitalizations do not significantly react to the event, and that lack of response is reflected in the pooled regression as well. But firms with large market capitalizations do experience losses within the same time frame described by the bivariate time series.

The general finding of both the bivariate time series models and the event study is that large, integrated refineries lose value 3-5 days after an increase in RIN prices. Small and medium firms exhibit no such losses, and there is no pattern of geographical heterogeneity in price responses whatsoever.

6 Discussion

Our empirical findings do not perfectly support any prevailing narrative about the effects of the RFS. While the lack of price response among small- and medium-size firms is in line with the conclusions of the pass-through literature, our discovery of a negative price response among large, integrated refineries is novel and not predicted by existing research on the RFS. Though this finding is interesting in and of itself, the reduced-form nature of our analysis does not allow us to pinpoint any one theoretical explanation for this result, which is the admitted limitation of this style of analysis. It is important to note that our results do not necessarily mean that large firms are not fully passing through RIN costs. In fact, a unique inability of larger firms to pass through RIN costs is theoretically unfounded and seems unlikely. Below, we provide a number of other candidate explanations.

First, it is notable that each of the large firms in our sample are also the firms with substantial investments in assets downstream of refining. If the RFS hurts the value of downstream assets without affecting the value of refining assets, compliance cost hikes would only affect the firms in our sample that own both. My results could also be a consequence of inefficient internal transfer pricing (Hirshleifer, 1956). If blending and refining occur under two separate profit-maximizing divisions and internal transfer prices are set inappropriately, a firm could lose profit in moving RINs from its blending to its refining division (see Figure 1).

Another possibility is that larger firms simply have a harder time complying with the RFS. Small refinery exemptions could play a role in this; the EPA is allowed to exempt refineries for whom complying with RFS mandates would cause serious economic hardship. Who receives these waivers is not public information (except when firms voluntarily elect to announce it), but exemptions as a whole were at a low point in 2015. Moreover, because these exemptions are issued at the refinery rather than the firm level, it is not clear that they should shield only small and medium firms from price impacts. Larger firms could also have a harder time procuring adequate quantities of biofuels. Biorefineries are known to be affected by diseconomies of scale in feedstock procurement because of transportation costs (Nguyen and Prince, 1996), and notoriously expensive biofuel transportation may generate a similar problem. At extremes, regional blend walls could act as constraints on firm's ability to generate RINs, creating nonlinearities in compliance costs for firms with otherwise low marginal blending costs.

Because there are only a small number of refining firms in the US and fewer still that we characterize as large, there is the lingering possibility that our results could be driven by peculiarities of just a few firms. If large firms generally held the belief that the EPA was going to reduce the mandate by more than they actually did, these firms could have been in a RIN-short position at the time of the regulatory announcements.

Untangling these possibilities will require structural models and, in many cases, proprietary data. Though these explanations each have their flaws, the effect that they must jointly describe, while significant, is economically small. Even without a full theoretical understanding of why large, integrated firms lose value while others don't, our results can still inform policy insofar as they thoroughly discredit the claim that integrated refiners take advantage of merchant refiners in the RIN market. This claim, which motivated policy proposals like transferring the point of RFS obligation from refiners to blenders, simply does not hold empirical water.

7 Conclusion

The Renewable Fuel Standard requires massive amount of biofuels to be blended into domestic fuel supplies, and the point of obligation for that blending rests mainly on oil refineries. A developed literature concludes that refineries are capable of passing through the costs of purchasing RINs, but pass-through is only one potential channel by which the RFS may affect the stock prices of obligated firms.

We use bivariate time series modeling and an established event study methodology to measure the price response of obligated firms to RIN price movements. We find a complete lack of price response among all firms except large companies with integrated downstream operations. This result is both counterintuitive and unexplained by the current RFS literature but is robust between both of our estimation techniques. The result discredits the idea that the Renewable Fuel Standard allows integrated refiners to reap substantial profits at the expense of merchant refiners that lack in-house blending and retail operations. The reduced-form nature of our analyses precludes us from being able to specify a causal mechanism, but we propose a number of potential hypotheses that could explain these results as a starting point for future research.

References

Babcock, B. A., Lade, G. E., and Pouliot, S. (2016). Impact on Merchant Refiners and Blenders from Changing the RFS Point of Obligation. *CARD Policy Briefs*.

Bracmort, K. (2015). The Renewable Fuel Standard (RFS): Waiver Authority and Modification of Volumes. Technical report, Congressional Research Service.

Burkhardt, J. H. (2016). Incomplete Regulation in an Imperfectly Competitive Market : The Impact of the Renewable Fuel Standard on U . S . Oil Refineries. *Working Paper*.

Bushnell, J. B., Chong, H., and Mansur, E. T. (2013). Profiting from Regulation: Evidence from the European Carbon Market. *American Economic Journal: Economic Policy*, 5(4):78–106.

Campos, J., Ericsson, N. R., and Hendry, D. F. (1996). Cointegration tests in the presence of structural breaks. *Journal of Econometrics*, 70(1):187–220.

Conley, T. G. and Taber, C. R. (2011). Inference with Difference in Differences with a Small Number of Policy Changes. *Review of Economics and Statistics*, 93(1):113–125.

Coppess, J. and Irwin, S. (2017). Another Wrinkle in the RFS: The Small Refinery Exemption. *farmdoc daily*, (7):224.

Energy Information Administration (2015). Refiners' Total Operable Atmospheric Crude Oil Distillation Capacity as of January 1, 2015. In *Annual Refinery Report*, pages 38–43.

Engle, R. F. and Granger, C. W. J. (1987). Co-Integration and Error Correction: Representation, Estimation, and Testing. *Econometrica*, 55(2):251–276.

Environmental Protection Agency (2015). EPA Proposes Renewable Fuel Standards for 2014, 2015, and 2016, and the Biomass-Based Diesel Volume for 2017.

Environmental Protection Agency (2017). Denial of Petitions for Rulemaking to Change the RFS Point of Obligation.

Environmental Protection Agency (2018). RFS Small Refinery Exemptions.

Fama, E. F. and French, K. R. (1992). The Cross-Section of Expected Stock Returns. *The Journal of Finance*, 47(2):427–465.

Gelbach, J. B., Helland, E., and Klick, J. (2013). Valid Inference in Single-Firm, Single-Event Studies. *American Law and Economics Review*, 15(2):495–541.

Hirshleifer, J. (1956). On the Economics of Transfer Pricing. *The Journal of Business*, 29(3):172–184.

Irwin, S. (2014). Rolling Back the Write Down of the Renewable Mandate for 2014: The RINs Market Rings the Bell Again. *farmdoc daily*, (4):148.

Johansen, S. (1991). Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models. *Econometrica*, 59(6):1551–1580.

Johansen, S., Mosconi, R., and Nielsen, B. (2000). Cointegration analysis in the presence of structural breaks in the deterministic trend. *The Econometrics Journal*, 3(2):216–249.

Knittel, C. R., Meiselman, B. S., and Stock, J. H. (2017). The Pass-Through of RIN Prices to Wholesale and Retail Fuels under the Renewable Fuel Standard. *Journal of the Association of Environmental and Resource Economists*, 4(4):1081–1119.

Knittel, C. R. and Smith, A. (2015). Ethanol Production and Gasoline Prices: A Spurious Correlation. *The Energy Journal*, 36(1):73–113.

Korting, C., de Gorter, H., and Just, D. R. (2019). Who Will Pay for Increasing Biofuel Mandates? Incidence of the Renewable Fuel Standard Given a Binding Blend Wall. *American Journal of Agricultural Economics*, 101(2):492–506.

Lade, G. E. and Bushnell, J. (2019). Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard. *Journal of the Association of Environmental and Resource Economists*, 6(3):563–592.

Lade, G. E., Lin Lawell, C.-Y. C., and Smith, A. (2018). Policy Shocks and Market-Based Regulations: Evidence from the Renewable Fuel Standard. *American Journal of Agricultural Economics*, 100(3):707–731.

Lambert, D. M., Wilcox, M. D., English, A., and Stewart, L. A. (2008). Ethanol Plant Location Determinants and County Comparative Advantage. *Journal of Agricultural and Applied Economics*, 40(1):117–135.

LaRiviere, J., Lima, L., and Musinov, S. (2017). Spatial Incidence of National Policies and Fixed Infrastructure: An Application to the Ethanol Mandate. *Working Paper*.

Lee, J. and Strazicich, M. C. (2003). Minimum Lagrange Multiplier Unit Root Test with Two Structural Breaks. *Review of Economics and Statistics*, 85(4):1082–1089.

Li, J. and Stock, J. H. (2019). Cost pass-through to higher ethanol blends at the pump: Evidence from Minnesota gas station data. *Journal of Environmental Economics and Management*, 93:1–19.

Mason, C. F. and Wilmot, N. A. (2016). Price discontinuities in the market for RINs. *Journal of Economic Behavior & Organization*, 132(B):79–97.

Montgomery, W. D. (1972). Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3):395–418.

Nguyen, M. H. and Prince, R. G. H. (1996). A Simple Rule for Bioenergy Conversion Plant Size Optimisation: Bioethanol from Sugar Cane and Sweet Sorghum. *Biomass and Bioenergy*, 10(5/6):361–365.

Perron, P. (1989). The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis. *Econometrica*, 57(6):1361–1401.

Phillips, P. C. B. and Perron, P. (1988). Testing for a Unit Root in Time Series Regression. *Biometrika*, 75(2):335–346.

Pouliot, S., Smith, A., and Stock, J. (2017). RIN Pass-Through at Gasoline Terminals. *Iowa State University Economics Working Papers*.

Salop, S. C. and Scheffman, D. T. (1983). Raising Rivals' Costs. *The American Economic Review*, 73(2):267–271.

Schnepf, R. and Yacobucci, B. D. (2013). Renewable Fuel Standard (RFS): Overview and Issues. Technical report, Congressional Research Service.

Serra, T., Zilberman, D., Gil, J. M., and Goodwin, B. K. (2011). Nonlinearities in the U.S. corn-ethanol-oil-gasoline price system. *Agricultural Economics*, 42(1):35–45.

Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica*, 48(1):1–48.

Valero Energy Corporation (2015). RE: Supplement to Valero Comments on Proposed Renewable Fuel Standards for 2014, 2015, and 2016 and Biomass-Based Diesel Volume.

Whistance, J., Ripplinger, D., and Thompson, W. (2016). Biofuel-related price transmission using Renewable Identification Number prices to signal mandate regime. *Energy Economics*, 55:19–29.

Whistance, J. and Thompson, W. (2014). A Critical Assessment of RIN Price Behavior and the Implications for Corn, Ethanol, and Gasoline Price Relationships. *Applied Economic Perspectives and Policy*, 36(4):623–642.