The Price of Biodiesel RINs and Economic Fundamentals

Scott H. Irwin, Kristen McCormack, and James H. Stock

The D4 RIN is the tradable compliance certificate for the biomass-based diesel (BBD) mandate in the renewable fuel standard (RFS). Understanding the price dynamics of the D4 RIN is important for understanding the RFS because its price sets a ceiling on the ethanol RIN (D6) and because some observers have suggested that RIN price fluctuations are too large to be explained by economic theory. We use option pricing theory to develop a model of the D4 RIN in terms of its economic fundamentals: the spread between the price of biodiesel and petroleum diesel and the status of the biodiesel blenders' tax credit. The resulting D4 fundamental price closely tracks actual D4 prices. We conclude that RIN price volatility arises because of the design of the RFS and intrinsic features of the U.S. fuel supply system.

Key words: Biodiesel, D4, fundamental, option, RIN, tax credit.

JEL codes: G13, Q4.

The renewable fuel standard (RFS) mandates the blending of biofuels into the surface transportation fuel supply, where the percentage blending rate is determined by an annual rulemaking from the U.S. Environmental Protection Agency (EPA). Refiners and importers of gasoline and diesel fuel ("obligated parties") demonstrate compliance with the RFS using the renewable identification number (RIN) system. A RIN is a unique electronic certificate that is generated when a gallon of biofuel is produced and is separated from the biofuel when it is blended with petroleum fuel. Once separated, the RIN can be traded. This enables obligated parties to purchase RINs, which they can then retire with the EPA to demonstrate compliance.

The total market value of RINs retired in 2017 was $14 billion.[1] Different categories of fuel generate different types of RINs. The two RINs that account for nearly all the market value are D6 RINs for conventional renewable fuels, which are mainly composed of corn starch ethanol, and D4 RINs for biomass-based diesel (BBD). As is evident in figure 1, RIN prices are highly volatile. This volatility creates compliance cost risk for obligated parties and undercuts the effectiveness of the RFS in stimulating investment in biofuels production and distribution infrastructure. The high volatility has also raised questions about how RIN prices are determined in practice and whether speculation and market manipulation could be part of the reason for RIN price volatility (e.g., Voegele 2013; Blewitt and Mider 2016).

This paper examines the extent to which D4 RIN prices are determined by economic fundamentals. D4 RINs are used to demonstrate compliance with the BBD requirement. However, they also can be used to demonstrate compliance with the conventional

Scott H. Irwin is the Laurence J. Norton Chair of Agricultural Marketing in the Department of Agricultural and Consumer Economics at the University of Illinois at Urbana-Champaign. Kristen McCormack is a PhD student in the Harvard Kennedy School at Harvard University. James H. Stock is the Harold Hitchings Burbank Professor of Political Economy in the Department of Economics and Harvard Kennedy School at Harvard University and the National Bureau of Economic Research. The authors thank Sandra Dunphy, Cynthia Lin Lawell, Eben Lazarus, Neelesh Nerukar, Aaron Smith, and the reviewers for helpful comments and/or guidance.
Correspondence may be sent to: sirwin@illinois.edu.

[1] Irwin, Scott H., and James H. Stock. 2018. The Renewable Fuel Standard. Unpublished manuscript.

Amer. J. Agr. Econ. 102(3): 734–752; doi:10.1002/ajae.12014
Published online February 3, 2020

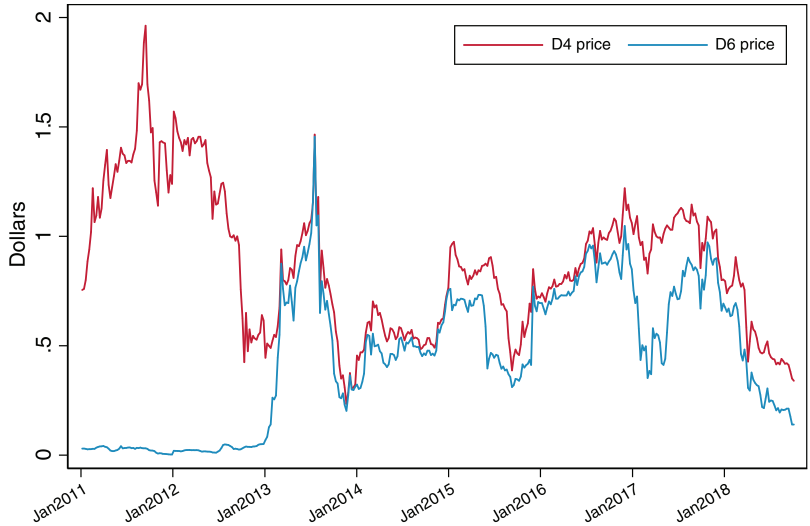

Figure 1. Weekly D4 biodiesel and D6 ethanol RIN prices, January 6, 2011–October 4, 2018 Weekly data (Thursday) are from the Oil Price Information Service (OPIS). Prices are of RINs generated in the current calendar year (current-year vintage RINs).

requirement; that is, a D4 RIN can be used instead of a D6 RIN but not vice versa. Thus, D4 RIN prices provide a cap on D6 RIN prices. D6 RIN prices were low prior to 2013 because blending ethanol into gasoline at ratios below 10% was chosen by the market as the most cost-effective available octane booster. The increase in D6 RIN prices at the start of 2013 reflected concern that conventional (ethanol) mandates could soon exceed 10%, breeching the so-called E10 blend wall; if so, satisfying the mandate would require either selling ethanol blends exceeding 10% (which require a subsidy) or using D4 RINs to satisfy the D6 requirement (Irwin and Good 2013). Subsequently, the cap on the D6 price provided by the D4 price has been binding much of the time since 2013. We focus on D4 RIN prices because we are able to observe the key fuel prices that economic theory suggests are the economic fundamentals of D4 RIN pricing, whereas this is not possible for D6 RINs. Because the D4 price is a binding cap on the D6 price for much of this period, if economic fundamentals explain D4 prices, then they explain much of the variation in D6 prices as well.[2]

A RIN can be retired in the year it is generated ("current-year") or it can be held and used to satisfy obligations incurred in the next year. This is similar to the banking provisions used in some pollution permit trading schemes (e.g., Cronshaw and Kruse 1996; Kling and Rubin 1997). Because it can be retired any time during this window, the D4 RIN is in effect an American call option. Economic theory indicates that the price of the underlying asset depends on (*a*) the price spread between biodiesel and its petroleum substitute, ultra low-sulfur diesel (ULSD); and (*b*) whether the \$1 per gallon biodiesel blenders' tax credit is in effect contemporaneously. We propose a simple model for these two fundamentals—the spread is a random walk, and the biodiesel tax credit follows a Markov process—that is consistent with their time-series properties. Using option pricing theory, we derive two models for the D4 price. The first allows for the possibility that the biodiesel-ULSD spread might be negative and yields a closed-form expression for the option price derived under

[2] Absent uncertainty and in a static model, if the conventional mandate exceeds 10% and biodiesel is the marginal gallon used to fill the conventional mandate, then the D4 and D6 RIN prices should be equal. In practice, the D4 RIN price always exceeds the D6 price, even if only by a small amount. This arises because the D4 RIN can satisfy both the D4 and D6 mandate, which makes D4 RINs more valuable because of the possibility that the D4 and D6 prices could separate in the future, for example, because the conventional mandate reverts to being below the E10 blend wall.

the additional assumption that future values of the spread have a Gaussian distribution, conditional on currently available information. The resulting predicted D4 price is a nonlinear function of the current spread and the status of the tax credit. The second expression does not require normality but assumes that the probability of a negative fundamental is negligible, and it provides a convenient and intuitive expression for the D4 price as a linear function of the current spread and tax credit status. The use of option pricing theory to understand the behavior of energy markets is hardly new (e.g., see Litzenberger and Rabinowitz 1995); however, to the best of our knowledge, we are the first to apply it to prices of bankable compliance permits.

It turns out that the two models yield similar predictions for the D4 price, although the nonlinear model outperforms the linear model when the spread is low. Figure 2 shows the D4 price and its predicted value based on the nonlinear economic fundamentals model, averaged across the predictions for the three markets for which we have data on the biodiesel-ULSD price spread (Chicago, the Gulf, and New York Harbor [NYH]). Evidently, the economic fundamentals do a good job explaining the variation in RIN prices at the monthly frequency and longer. There are short-term (one- or two-week) departures from the fundamentals, which we take to represent unmodeled transitory developments in the fuels market such as weather-related supply disruptions, reactions to news, or rumors regarding U.S. congressional changes to the RFS, or adjustments to changes in the implementation of the annual standards. There are also some longer departures from fundamentals, such as in the first half of 2016; however, those departures are relatively small (the average prediction error from January to June 2016 is $0.13; over all of 2016 it is $0.02). The departures of prices from our fundamentals-based price also could arise from market participants using more information to update their beliefs about whether the biodiesel tax credit will be in effect in the next year than is used in our simple Markov model. In fact, these departures can be used to infer time-varying market beliefs, an extension we take up later in the article.

This paper contributes to the literature on RIN pricing. The most closely related contributions are Irwin and Good (2017) and Lade, Lin Lawell, and Smith (2018). Irwin and Good (2017) price D4 RINs using the contemporaneous economic fundamentals and do not incorporate the option value or the uncertainty surrounding the biodiesel tax credit. Lade, Lin Lawell, and Smith (2018) develop a dynamic model with uncertainty for jointly pricing multiple RINs, including the nesting cap. Relative to their paper, by focusing solely on the D4 RIN, we are able to obtain a closed-form solution for the option price; in addition, we use a more

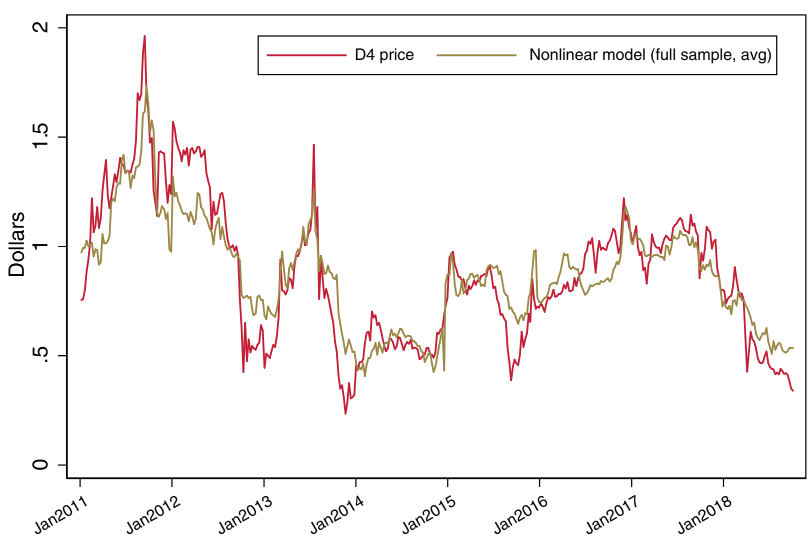

Figure 2. Weekly D4 biodiesel RIN price and its predicted value based on the nonlinear economic fundamentals model, averaged over three markets (Chicago, Gulf, New York Harbor), January 6, 2011–October 4, 2018

immediate measure of fundamentals that we expect should improve fit (biodiesel prices and ULSD, whereas they use soybean oil and crude oil prices), and we incorporate the biodiesel tax credit, which turns out to be empirically important. There is also a growing literature on the pass-through of RIN prices through the fuel supply chain (for references, see Knittel, Meiselman, and Stock 2017; Lade and Bushnell 2019); however, that literature focuses on the consequences of a movement in RIN prices, not on the economic reasons for RIN price variation in the first place.

The RFS: A Brief Summary

The renewable fuel standard (RFS) was established by the Energy Policy Act of 2005 and was substantially expanded as part of the Energy Independence and Security Act (EISA) of 2007. The RFS divides renewable fuels into four categories: total renewable, advanced, biomass-based diesel (BBD), and cellulosic. All fuels that qualify for the RFS have at least a 20% reduction in life-cycle greenhouse gas emissions compared to their petroleum counterparts. To qualify as an advanced biofuel, the reduction in life-cycle emissions must be at least 50%. As shown in figure 3, the renewable fuel categories in the RFS are nested based on the level of emission reductions. Since biofuels also differ by their energy content per gallon, the standards adjust for these differences by assigning an energy-equivalence value (EV) to each biofuel relative to a gallon of ethanol. For example, ethanol has EV = 1, biodiesel has EV = 1.5, and renewable diesel has EV = 1.7.

The U.S. Environmental Protection Agency (EPA) is responsible for implementing the RFS. Each year, the EPA determines a volume (a renewable volume obligation [RVO]) of each of the four categories of renewable fuels for use in the U.S. fuel supply for the coming year. The EPA's determination starts with a table of RVOs in the 2007 EISA, modified using various discretionary authorities specified in the statute. The difference between the total RVO and the total advanced RVO is capped in the EISA at 15 billion gallons. This volume is referred to as the conventional RVO and has been met primarily by corn ethanol.

The EPA translates the annual RVOs into percentage standards, expressed as the renewable percentage requirement per gallon of petroleum fuel blended. Obligated parties under the RFS are petroleum refiners and importers, who must demonstrate that this fraction of renewable fuels has been blended into the total volume of petroleum gasoline

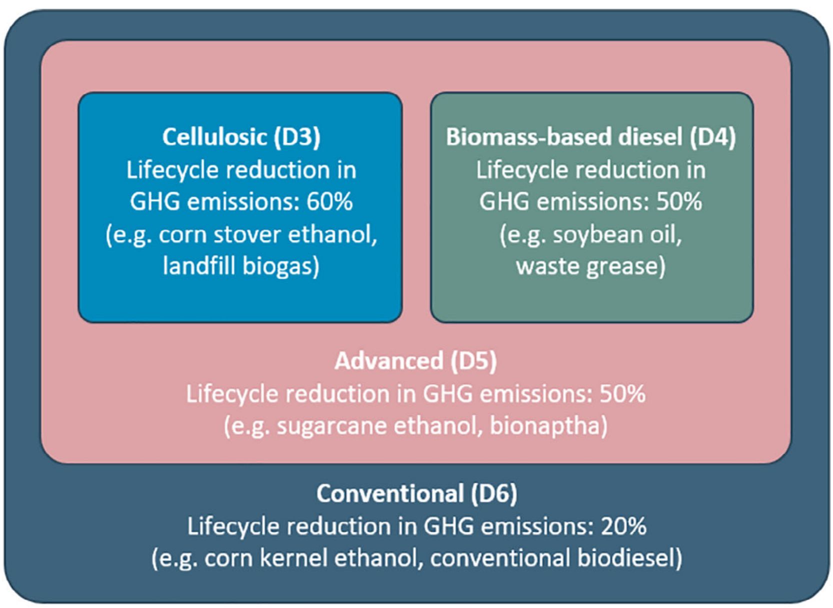

Source: Environmental Protection Agency

Figure 3. The renewable fuel standard nesting scheme for biofuels

or diesel that they sell into the surface transportation pool. Obligated parties demonstrate compliance using the renewable identification number (RIN) system. A RIN is an electronic certificate that is generated when a gallon of biofuel is produced, and it is separated when that biofuel is blended into finished gasoline or diesel for retail sale. As shown in figure 3, each category of biofuel has its own RIN number: D6 (conventional), D4 (biomass-based diesel), D7 (cellulosic diesel), D3 (non-diesel cellulosic), and D5 (all other advanced). A separated RIN can be retired with the EPA to demonstrate compliance, sold to an obligated party that needs additional RINs to demonstrate compliance, or banked for future use. The nested structure of the RFS biofuel categories means that RINs also have a nested structure, whereby a RIN for a higher-ordered category can be used to satisfy the RVO of a lower-ordered category. In particular, a D4 RIN, which is generated by producing advanced, non-cellulosic biomass-based diesel, can be used instead of a D6 RIN to satisfy the conventional RVO.[3]

D4 RIN Pricing Model

At the low blend ratios currently in use, pure petroleum diesel (ULSD) and blends of biodiesel and ULSD are effectively perfect substitutes, after adjusting for biodiesel having 92.7% of the energy content of ULSD. Because biodiesel is more expensive than ULSD, it would not enter the market were it not for subsidies (e.g., Irwin 2014). The two national-level subsidies for biodiesel are through the RFS, in the form of the D4 RIN and the \$1 per gallon biodiesel blenders' tax credit.

We begin by describing the date-t fundamental value of the D4 RIN, first without the biodiesel tax credit in place and then with the tax credit in place. Because the biorefiner produces both the "wet" (physical) biodiesel and the D4 RINs attached to the biodiesel, the value received by the biorefiner is the sum of the wet fuel value, which is the energy-adjusted ULSD price, and the price of the D4 RIN. Economic theory suggests that, absent the biodiesel tax credit, the D4 RIN price will adjust so that the supply of biodiesel equals the demand for biodiesel, where the demand is determined by the EPA annual RFS rulemaking. Because each gallon of wet biodiesel generates 1.5 D4 RINs, absent the tax credit the price based on contemporaneous economic fundamentals at date t is $P_t^* = \max\left((P_t^{Biodiesel} - 0.927 P_t^{ULSD})/1.5, 0\right)$, where $P_t^{Biodiesel}$ is the biodiesel price, P_t^{ULSD} is the ULSD price, and 0.927 is the energy content adjustment for biodiesel.[4] The fundamental price is truncated at zero because if the biodiesel price is less than the energy-adjusted ULSD price, the BBD mandate will not be binding and the D4 RIN price will be zero.

The biodiesel blenders' tax credit provides a tax credit of \$1 for each gallon of biodiesel that is blended with ULSD. Because ULSD and biodiesel are perfect substitutes (energy-adjusted), under perfect competition the blenders' tax credit will accrue to the biorefiner (and thus to the feedstock producer). Thus, in our base model, the fundamental value of the D4 RIN at date t is $P_t^* = \max\left((P_t^{Biodiesel} - 0.927 P_t^{ULSD} - B_t)/1.5, 0\right)$, where $B_t = 1$ if the biodiesel tax credit is in effect on date t and $B_t = 0$ otherwise.

A D4 RIN can be used to demonstrate compliance for an obligation incurred in the year it is generated or in the year thereafter. Thus, it is an American option with no dividend and an expiration date of December 31 in the year after it is generated. As a result, the D4 RIN can be priced as a European option, where its fundamental value on the expiration date is the price of the underlying asset. For a risk-neutral firm,[5] this gives the pricing formula

$$(1) \quad P_t^{D4} = e^{-r(T-t)} E_t \max[(S_T - B_T)/1.5, 0],$$

where $S_T = P_T^{Biodiesel} - 0.927 P_T^{ULSD}$, T is December 31 of the year after the RIN was

[3] For additional information on the RFS and its history, see Schnepf and Yacobucci (2013) and Bracmort (2019).

[4] The fundamental price here is expressed for biodiesel, which, as noted in the previous section, generates 1.5 RINs per gallon, and for which the main feedstock in the United States is soybean oil. The BBD mandate in the RFS also can be met using renewable diesel, which is produced by hydrotreatment, is fully compatible with petroleum diesel, and generates 1.7 RINs per gallon because of its higher energy content. In equilibrium, there would also be a D4 fundamentals equation relating the price of renewable diesel to ULSD. We focus on biodiesel because its volumes are larger than renewable diesel, because we have biodiesel prices but not renewable diesel prices, and because soy biodiesel is generally considered the marginal fuel in the industry.

[5] Risk neutrality is not needed to obtain equation (1). From the fundamental theorem of asset pricing, the D4 price is $P_t^{D4} = E_t(m_{t,T} \times \max[(S_T - B_T)/1.5, 0])$, where $m_{t,T}$ is the stochastic discount factor. Equation (1) follows if the stochastic discount factor is uncorrelated with the fundamental price.

generated, E_t denotes the expectation conditional on information at time t, and the factor of 1/1.5 adjusts for the fact that one gallon of biodiesel generates 1.5 RINs. The maximum in equation (1) imposes the condition that the price of the D4 RIN cannot be negative. Equation (1) is the same as equation (2) in Lade, Lin Lawell, and Smith (2018), extended to include the biodiesel tax credit.

We complete the model by assuming that S_T follows a random walk and that B_T follows a Markov process:

$$(2) \quad E(S_{t+\tau}|S_t, B_t) = S_t, \tau \geq 1, \text{and}$$

$$(3) \quad \Pr(B_T = 1|S_t, B_t) = E(B_T|S_t, B_t) = pB_t + (1-q)(1-B_t),$$

of S_T given $S_t = s_t$ and $B_t = b_t$ is Gaussian, specifically $N[s_t, \sigma^2(T-t)]$, where σ^2 is the variance of ΔS_t (we treat this variance as constant here for simplicity but in the empirical work allow it to vary over time). This Gaussian assumption does not require the change in the spread to be Gaussian but rather can be motivated by a Central Limit Theorem, in which the average change in the spread is normally distributed; this appeal to the Central Limit Theorem allows for heavy-tailed daily or weekly changes that get averaged out over the relatively long life of the RIN, which extends to December 31 of the year after its generation. Under these assumptions, a calculation yields

$$(4) \quad P_t^{D4} = e^{-r(T-t)} \sum\nolimits_{b=0,1} E[\max(S_T - B_T, 0)|S_t = s_t, B_T = b] \Pr(B_T = b|B_t)/1.5$$
$$= e^{-r(T-t)} \{ f_{0,t} - (f_{0,t} - f_{1,t}) [pB_t + (1-q)(1-B_t)] \} / 1.5$$

where p and q are the probabilities of staying in states 1 and 0, respectively; that is, $p = \Pr[B_T = 1|B_t = 1]$ and $q = \Pr[B_T = 0|B_t = 0]$.

If the biodiesel tax credit is in place at date t, it is in place for the current calendar year. Thus, equation (3) applies to terminal date T in the calendar year subsequent to the current date t. This is the appropriate timing for evaluating the price of current-year RINs. Hence, $\Pr[B_T = 1|S_t, B_t = 0]$ represents the probability that the biodiesel tax credit will be in effect at date T, conditional on the (nonretroactive) contemporaneous status of the tax credit at time t. We examine these assumptions empirically in the next section and show that they are consistent with the spread and tax credit data, with the exception that there is some evidence that the level of the spread depends on the value of the tax credit. We generalize the model to allow for this possibility below, after first solving equations (1)–(3).

Two closed-form solutions for the D4 RIN price are provided. The first further assumes that the conditional distribution

where

$$(5) \quad f_{b,t} = \sigma\sqrt{T-t}\,\phi\left(\frac{s_t - b}{\sigma\sqrt{T-t}}\right) + (s_t - b)\Phi\left(\frac{s_t - b}{\sigma\sqrt{T-t}}\right),$$

and $\phi(.)$ is the normal density and $\Phi(.)$ is the cumulative normal distribution.[6]

The second solution assumes that the probability of the fundamental price going below zero is negligible, in which case $E_t\max[(S_T - B_T)/1.5, 0] \approx E_t(S_T - B_T)/1.5$ and

$$(6) \quad P_t^{D4} = e^{-r(T-t)} E_t(S_T - B_T)/1.5$$
$$= e^{-r(T-t)} \{ S_t - [pB_t - (1-q)(1-B_t)] \} / 1.5.$$

Equation (6) also obtains as a limiting approximation to equations (4) and (5) for small σ. The linear pricing expression in equation (6) has the intuitive interpretation that the D4 RIN price is the expected value of the

[6] Write $S_T - b = (S_T - s_t) + (s_t - b) = z\tau + m$, where $m = s_t - b$, $\tau = \sigma\sqrt{T-t}$, and $z = (S_T - s_t)/\tau$. Conditional on $B_T = b$ and $S_t = s_t$, $z \sim$ N(0,1). Thus, $E[\max(S_T - B_T, 0)|S_T = s_t, B_T = b] = E\max[z\tau + m, 0] = E[(z\tau + m)1(z > -m/\tau)] = \int_{-m/\tau}^{\infty}(z\tau + m)\phi(z)dz = \tau\int_{-m/\tau}^{\infty} z\phi(z)dz + m\int_{-m/\tau}^{\infty}\phi(z)dz = \tau(2\pi)^{-1/2}\int_{-m/\tau}^{\infty} ze^{-z^2/2}dz + m[1 - \Phi(-m/\tau)] = \tau\phi(m/\tau) + m\Phi(m/\tau)$. Substituting the expressions for m and τ into this final expression and collecting terms yields equations (4) and (5).

spread at the end of its lifetime, minus the current expectation of whether the tax credit will then be in effect, discounted to the present and adjusted for the energy value factor of 1.5.

It is tempting to try to extend this approach to the D6 RIN, the RIN generated by corn ethanol. This is not readily done, however, because ethanol is not a direct substitute for gasoline after energy adjustment. Ethanol has typically had a higher octane value than petroleum gasoline feedstock, so at blends less than 10% it is used as an octane booster (Irwin 2019). At blends greater than 10%, it faces the E10 blend wall and consumers need an incentive to blend ethanol. Thus, although the ethanol supply price (the price of bulk ethanol with a RIN) is observed in commodity markets, the ethanol demand price depends on the blend ratio and is not observed. In addition, the fact that D4 RINs can be used to meet the conventional mandate further complicates the analysis. For additional discussion, see Lade, Lin Lawell, and Smith (2018).

Data and Empirical Results

We use weekly Oil Price Information Service (OPIS) data (Thursday) of national average prices for D4 RINs that expire in the current year and of spot prices for wholesale ULSD and biodiesel at Chicago, the Gulf, and the New York Harbor. The D4 weekly price data and the Chicago and Gulf fuel price data span September 3, 2009–October 4, 2018. The New York Harbor data span October 19, 2012–October 4, 2018.

The RFS underwent a transition in 2010 with new volumes and regulations. The first year of the new regime ("RFS2") in which the required volumes were known in real time was 2011. We therefore begin our estimation in the first week of January 2011. We use earlier data on the spread to estimate the variance of the change in the spread, as discussed below. We also collected data on when the biodiesel tax credit was in effect contemporaneously and, on each date, when it was set to expire if it was in place contemporaneously.[7]

For the interest rate, we use the six-month Treasury rate.

The Spread and the Biodiesel Tax Credit

Table 1 presents statistics describing the stochastic process followed by the energy-adjusted spread, S_t, between biodiesel and ULSD prices. Column (1) presents a levels autoregression and Dickey-Fuller test for a unit root, column (2) presents the same regression imposing a unit root (i.e., first-differences regression), and column (3) examines whether the coefficients in the first-differences regression depend on the status of the biodiesel tax credit. For all three spreads, the results are consistent with the base model assumption that S_t follows a random walk and that its coefficients do not depend on the biodiesel tax credit.[8]

Columns (4) and (5) in table 1 examine the possibility that the level of the spread depends on whether the biodiesel tax credit is in place. The evidence suggests that (*a*) the Chicago spread averages \$0.76 higher if the \$1 tax credit is in effect (\$0.80 for the Gulf and \$0.69 for the NYH spread, which are over a shorter sample), and (*b*) the residual from regression in column (4) follows a random walk. This latter finding is consistent with S_t following a random walk with jumps on the dates that the tax credit comes into effect. Figure 4 shows the Chicago spread and its predicted value from regression (4); this predicted value is a step function that depends on the status of the tax credit. This variation in the spread related to the biodiesel tax credit is large economically as well as statistically (the R^2 of regression (4) for Chicago is 0.47); however, because the spread is integrated of order 1 and there are only a few times that the tax credit turns off and on contemporaneously, the coefficient on the tax credit is estimated

[7] The biodiesel blenders' tax credit was in place contemporaneously for the full calendar years of 2007–9, 2011, 2013, and 2016. For calendar years 2010, 2012, 2014, 2015, and 2017, the tax credit was restored retroactively, with the exception of the last three weeks of December 2014 when the tax credit was reinstated contemporaneously. For 2018, the tax credit had expired, and, as of the end of the data set, it was not known whether it would be restored retroactively. Thus, for 2010, 2012, 2014–15, 2017, and 2018, the tax credit was not in place, but market participants did not know whether the tax credit would be restored retroactively or whether it would be reinstituted for the subsequent year. For 2011, 2013, and 2016, it was in place but was scheduled to expire at the end of the year, and it was unknown whether it would be extended into the subsequent year.

[8] The Dickey-Fuller tests do not reject a unit root for all three spreads. For Chicago and New York Harbor, lags of the spread beyond the first do not enter the spread regression at the 10% significance level, consistent with the random walk model. For the Gulf, however, they are significant at the 5% level. The sum of the coefficients on lagged first differences for the Gulf is small (0.03), so forecasts including those lags are consistent with random walk forecasts at horizons beyond a week. Given the long horizon for the forecasts because of the RIN retirement date, we therefore use the random walk approximation for all three spreads.

imprecisely.[9] Because this dependence of the spread on B_t is perfectly collinear with the included regressors $(1 - B_t)$ and B_t in equation (6), it does not change the D4 predicted price; however, as discussed below, it changes the interpretation of the coefficients in the D4 pricing model and has an interesting substantive interpretation of its own.

It is more difficult to check the assumptions of the biodiesel tax credit Markov model in equation (3) because of the history of the tax credit. Historically since 2012, the tax credit was on for at most the current year, never for future years, and until 2018 it was regularly reinstated retroactively after it expired. Thus, with the benefit of the full data set, it looks as though the probability of the tax credit being on in the future was nearly 1 regardless of whether it was currently in effect. In real time, however, there was always uncertainty as to whether the U.S. Congress would in fact enact the credit in the next year or restore it retroactively.[10]

D4 Pricing with Nonlinear and Linear Models

For the Chicago and Gulf spreads, the weekly linear models (equation (6)) were estimated over January 6, 2011–October 4, 2018. For the NYH spread, the linear model was estimated over the full span of the available weekly data, October 19, 2012–October 4, 2018. Constructing the terms $f_{0,t}$ and $f_{1,t}$ in the nonlinear model (equations (4) and (5)) requires an additional parameter, the variance of ΔS_t (its square root is σ in equation (5)). We estimated this variance using a rolling fifty-two-week retrospective equal-weighted moving average of $(\Delta S_t)^2$, not including the current week. For the Chicago and Gulf spreads, we have more than a year of presample data available to estimate the initial variance, so the model estimation sample is January 6, 2011–October 4, 2018, and all observations use the fifty-two-week retrospective rolling variance. For the NYH spread, the first observation for our NYH data is October 19, 2012. To maximize the estimation span for the NYH nonlinear model, we used a recursive estimator of the variance of ΔS_t for the first fifty-two weeks (so for these initial periods, the equal-weighted average runs from October 19, 2012, to the week before date t), and thereafter a fifty-two-week rolling estimator. This allows us to estimate the NYH nonlinear model over the span October 25, 2012–May 31, 2018. The remaining two free parameters in the nonlinear model, q and p in equation (4), are estimated by OLS estimation of equation (4).[11]

The full-sample estimates for the nonlinear model are given in columns (1), (3), and (5) of table 2 for the Chicago, Gulf, and NYH spreads, respectively, and the average of the predicted values from these regressions is shown in figure 2 (from January 6, 2011, to October 18, 2012, the average is of the Chicago and Gulf predicted values; thereafter, all three predicted values are averaged). Taken literally, the estimated values of q for the Gulf spread indicate that, if the tax credit is not in effect, the market believes there is an approximately 18% chance that it will be in effect at the RIN expiration date next year. If the tax credit is currently in effect, the estimated value of p indicates that the markets believe there is a 70% chance that it will be in effect next year. The estimates of p and q from the Chicago and NYH spreads are within one standard error of the estimates for the Gulf.

---

[9] The regression in column (4) is $S_t = \alpha + \beta_B B_t + u_t$, where (under the assumptions of the text) u_t follows a random walk with var$(\Delta u_t) = \sigma^2_{\Delta u}$. The persistence of the tax credit and the random walk assumption for the error term leads to a nonstandard sampling distribution for $\hat{\beta}_B$, the OLS estimator of β_B. Eliminate the intercept from this regression by subtracting off the mean of S_t and B_t, and define $\delta(\tau) = B_{[\tau T]} - \bar{B}$, where [.] is the least greater integer function and $\bar{B}$ is the sample mean of B_t. Then the coefficient on B_t in regression (4), $\hat{\beta}_B$, has the limiting representation, $T^{-1/2}\left(\hat{\beta}_B - \beta\right) \Rightarrow \sigma_{\Delta u}\int_0^1 \delta(\tau)W^{\mu}(\tau)d\tau / \int_0^1 \delta^2(\tau)d\tau$, where W^{μ} is demeaned Brownian motion. This has a limiting normal distribution, so a 95% confidence interval for β_B can be computed as ± 1.96 standard errors of $\hat{\beta}_B$. From the limiting expression, it follows that $\text{var}\left(\hat{\beta}_B\right) = T\sigma^2_{\Delta u}\int_0^1\int_0^1 d(\tau)d(r)\min(\tau, r)d\tau dr / \left(\int_0^1 \delta^2(\tau)d\tau\right)^2$ (the simplification of the covariance kernel of demeaned Brownian motion arises because $\int_0^1 \delta(\tau)d\tau = 0$). The standard error is computed from this expression using the standard deviation of Δu_t as an estimate of $\sigma^2_{\Delta u}$ and by numerical evaluation of the double integral.

[10] The nonlinear model also assumes that the conditional distribution of S_T given $S_t = s_t$ and $B_t = b_t$ is $N[s_t, \sigma^2(T - t)]$. Empirically, this distribution does in fact appear to be reasonably approximated as Gaussian. For example, at the $T - t$ = twenty-four-week horizon, the skewness and kurtosis of the spread are, respectively, −0.12 and 2.75 for Chicago, −0.16 and 2.83 for the Gulf, and −0.32 and 3.51 for NYH (these are 0 and 3, respectively, for a normal distribution). Longer horizons are also consistent with normality but have fewer nonoverlapping observations.

[11] At first glance, it may appear that the OLS estimates of q and p cannot be estimated very precisely because there are so few instances of the biodiesel tax credit turning on or off. However, this is not the case because the estimates are derived from weekly observations that incorporate expectations for the credit being in place or not. Consequently, the probabilities are estimated as the means of weekly residuals during the relevant (on or off) period, and since the residuals (after de-meaning) are stationary, the probabilities can in principle be well estimated.

Table 1. Autoregressive Models of the Spread (S_t) between Ultra Low Sulfur Diesel and Biodiesel Prices, January 6, 2011–October 4, 2018

Model	(1)	(2)	(3)	(4)	(5)
Dependent Variable	S_t	ΔS_t	ΔS_t	S_t	$\hat{u}_t$
Regressors	$S_{t-1},\ldots, S_{t-6}$	$\Delta S_{t-1},\ldots, \Delta S_{t-5}$	$\Delta S_{t-1},\ldots, \Delta S_{t-5}$ $B_{t-1} \times \Delta S_{t-1},\ldots, B_{t-5} \times \Delta S_{t-5}, B_t$	B_t	$\hat{u}_{t-1},\ldots, \hat{u}_{t-6}$
Panel A. Chicago, January 6, 2011–October 4, 2018 (*n* = 405)					
Intercept	0.033 (0.021)	–0.003 (0.005)	–0.013 (0.007)	1.624	–0.001 (0.006)
Coefficient on B_t	–	–	0.018 (0.011)	0.764 (0.570)	
ADF test	–2.57	–	–	–	–
Sum of coefficients on lagged levels	0.978 (0.013)	–	–		0.964 (0.021)
F-test, all lags (except S_{t-1}) and interactions (p-value)	1.43 (0.212)	1.18 (0.320)	1.20 (0.288)	–	0.75 (0.589)
F-test, all interactions (p-value)	–	–	1.33 (0.252)	–	
Panel B. U.S. Gulf at New Orleans, January 6, 2011–October 4, 2018 (*n* = 405)					
Intercept	0.030 (0.020)	–0.003 (0.006)	–0.015 (0.007)	1.650	–0.002 (0.006)
Coefficient on B_t	–	–	0.020 (0.011)	0.796 (0.570)	
ADF test	–2.25	–	–	–	–
Sum of coefficients on lagged levels	0.980 (0.011)	–	–		0.963 (0.018)
F-test, all lags (except S_{t-1}) and interactions (p-value)	2.44 (0.034)	2.17 (0.057)	1.48 (0.144)	–	0.74 (0.596)
F-test, all interactions (p-value)	–	–	1.03 (0.400)	–	
Panel C. New York Harbor, October 19, 2012–October 4, 2018 (*n* = 307)					
Intercept	0.044 (0.028)	–0.001 (0.006)	–0.012 (0.007)	1.477	–0.001 (0.006)
Coefficient on B_t	–	–	0.026 (0.014)	0.691 (0.422)	
ADF test	–1.85	–	–	–	–
Sum of coefficients on lagged levels	0.970 (0.018)	–	–		0.960 (0.026)
F-test, all lags (except S_{t-1}) and interactions (p-value)	0.81 (0.546)	0.68 (0.639)	1.28 (0.240)	–	0.70 (0.623)
F-test, all interactions (p-value)	–	–	1.27 (0.277)	–	

Note: Standard errors are in parentheses below coefficients; p-values are in parentheses below F-statistics. All regressions include an intercept. The ADF test includes a linear trend and six lags and rejects the null of a unit root at the **1% and *5% level. In column (4), the standard errors for the coefficients on B_t in the levels regression (4) are computed using the Gaussian functional limit described in the text; standard error for the intercept is not substantively relevant and is not computed. In column (5), $\hat{u}_t$ are the residuals from regression (4) of S_t on B_t.

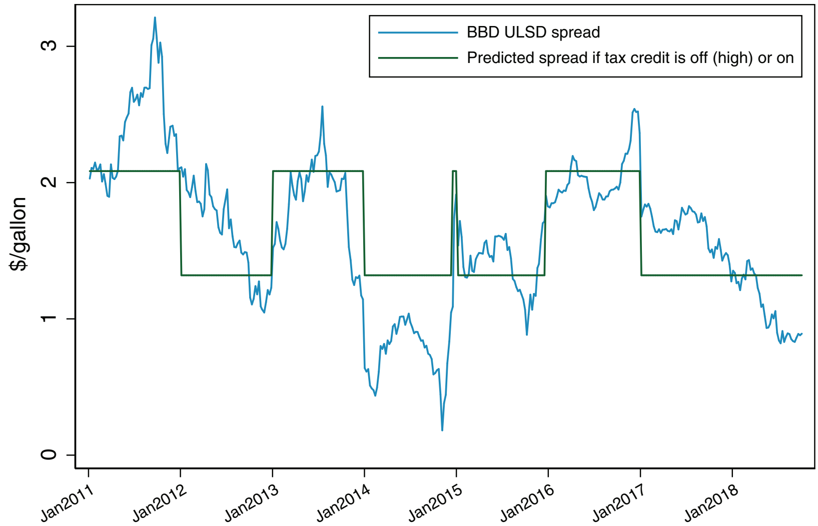

Figure 4. Weekly Chicago biodiesel and ultra-low sulfur diesel spread and its predicted value based only on whether the biodiesel tax credit is in effect contemporaneously, January 6, 2011–October 4, 2018

Columns (2), (4), and (6) of table 2 report estimates of the linearized model in equation (6). The estimated probabilities of the tax credit being in effect in the next year are somewhat smaller for the linear model than for the nonlinear model. Notably, the fit of the nonlinear model is on average slightly better than that of the linear model.

The models discussed so far use the full sample of RIN prices to estimate the transition probabilities q and p, so the resulting prices would not have been available in real time. To provide real-time prices, we therefore estimated the nonlinear model over a rolling 104-week window; for instance, $t - 105, \ldots, t - 1$ and substituting the resulting rolling estimates of q and p along with the values of B_t and S_t at date t into equation (4) yields a real-time price (recall that the volatility is estimated over a fifty-two-week retrospective window ending in $t - 1$).[12] Because all our data are unrevised asset price data available in real time, the rolling predicted prices therefore are feasible real-time prices. We refer to this model as the real-time nonlinear model.

Figure 5 presents the D4 price and the average predicted value from (*a*) the nonlinear models in table 2, (*b*) the linear models in table 2, and (*c*) the real-time rolling nonlinear models. There are three salient features of this chart. First, at the monthly frequency, the models generally track each other closely. Second, the full-sample nonlinear and linear models tend to differ the most when the RIN price is low. This corresponds to dates at which the term $S_t - B_t$ is close to zero, so that the probability of hitting zero is non-negligible and the nonlinear terms—that is, the option value component—come into play. In these cases, the nonlinear terms improve the fit (see the episodes in early and late 2014). In contrast, when $S_t - B_t$ is far from zero, the predicted values for the linear and nonlinear models are quite close. Third, the only time that the real-time nonlinear model has different prices than the full-sample nonlinear model for an extended period is the first half of 2017, when the fit of the full-sample model is better. The first half of 2017 was a period of evolving expectations during the early months of the Trump administration, so one interpretation of this discrepancy is that probabilities estimated using data from the final two years of the Obama administration are inappropriate estimates of the actual market probabilities of reinstatement of the tax credit during this period.

Figure 6 plots the real-time predicted values from the nonlinear rolling models for the three spreads separately. Evidently there are high-

[12] The results are insensitive to using a twenty-six-week window or thirty-nine-week window for the rolling volatility estimate instead of a fifty-two-week window. The coefficients in columns (1), (3), and (5) of table 2 all change by less than one standard error of the original estimates when alternative windows are used, and of those tested, the R^2 values in table 2 are greatest when using a fifty-two-week window.

Table 2. Estimation Results for D4 Biodiesel RIN Pricing Models, January 6, 2011–October 4, 2018

Model	(1) Chicago Nonlinear	(2) Chicago Linear	(3) Gulf Nonlinear	(4) Gulf Linear	(5) NYH Nonlinear	(6) NYH Linear
Estimate of $1-q$	0.186 (0.064)	0.098 (0.048)	0.183 (0.065)	0.111 (0.045)	0.253 (0.046)	0.146 (0.038)
Estimate of p	0.664 (0.065)	0.615 (0.054)	0.697 (0.059)	0.658 (0.048)	0.751 (0.060)	0.696 (0.043)
R^2	0.789	0.769	0.807	0.811	0.776	0.765
Sample	January 6, 2011–October 4, 2018	January 6, 2011–October 4, 2018	January 6, 2011–October 4, 2018	January 6, 2011–October 4, 2018	October 25, 2012–October 4, 2018	October 11, 2012–October 4, 2018
n	403	403	403	403	309	311

Note: Gulf denotes U.S. Gulf of Mexico at New Orleans and NYH denotes New York Harbor. Standard errors are Newey-West with 20 lags.

frequency differences among the predicted values, presumably due to transient local supply or demand conditions. At medium and low frequencies, however, the predicted values are essentially the same for all three spreads.

Finally, we computed the spread using the energy discount arising because biodiesel has an energy content that is 92.7% that of ULSD. Because this energy discount is identified in our model, we can use our model to see whether the RIN market actually incorporates the biodiesel energy discount when pricing RINs.[13] Specifically, recall that the spread is defined in energy-adjusted terms, that is, $S_t = P_t^{Biodiesel} - 0.927P_t^{ULSD}$. Instead, write this as $S_t = P_t^{Biodiesel} - \lambda P_t^{ULSD}$, where λ is a factor to be estimated. Rearranging the terms in the linear model of equation (6) yields

$$
\begin{aligned}
(7)\qquad & E\left[1.5e^{r(T-t)}P_t^{D4}|S_t,B_t\right] \\
&= S_t - [pB_t - (1-q)(1-B_t)] \\
&= P_t^{Biodiesel} - \lambda P_t^{ULSD} - [pB_t - (1-q)(1-B_t)].
\end{aligned}
$$

The coefficient λ thus can be estimated from a regression of $1.5e^{r(T-t)}P_t^{D4} - P_t^{Biodiesel}$ on P_T^{ULSD}, B_t, and $1 - B_t$ (with no constant). The resulting estimated values for the Gulf and NYH are remarkably similar to the actual value of the energy discount of 0.927, respectively being 0.920 (SE = 0.054) and 0.927 (0.045). For Chicago, the estimated discount of 0.809 (0.047) is less similar to the theoretical value. This departure appears to be due to downward errors-in-variables bias driven by infrequent spikes in the Chicago ULSD price. If the energy discount is estimated for Chicago from 2016 onward, a period with few spikes, the estimated discount is 0.887 (0.097), close to the theoretical value of 0.927. These results suggest that, in general, the market correctly incorporates the reduced energy value of biodiesel when pricing RINs, providing further evidence in support of the market RIN price reflecting fundamentals.

Real-Time Probabilities that the Tax Credit Will Be in Effect

The Markov model for the tax credit assumes that market predictions of whether the credit

[13] We thank an anonymous referee for this suggestion.

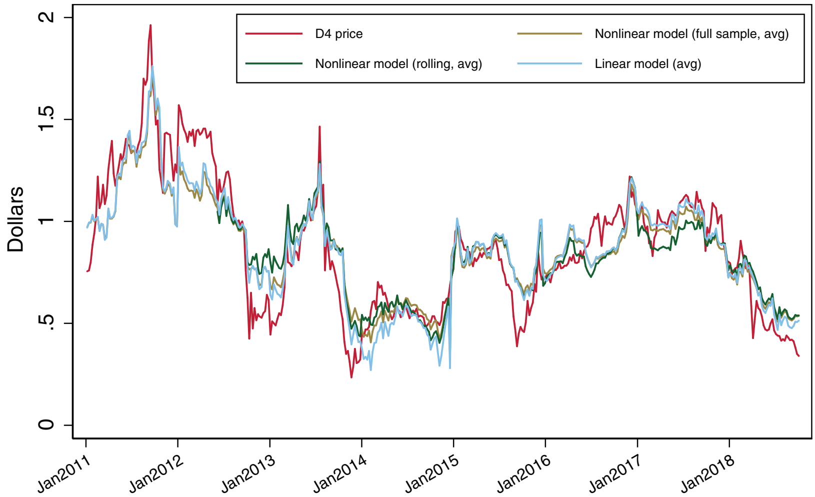

Figure 5. Weekly D4 biodiesel RIN price and predicted value based on linear and nonlinear models, both full-sample and the real-time rolling nonlinear model, January 6, 2011–October 4, 2018

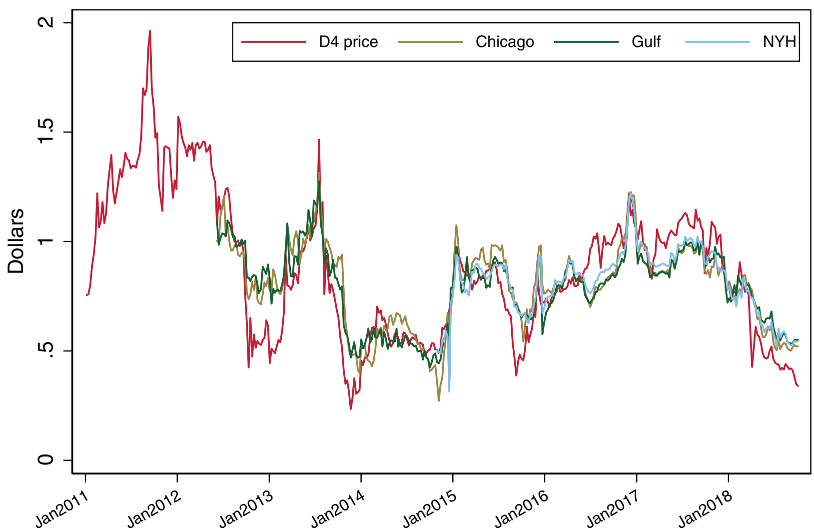

Figure 6. Weekly D4 biodiesel RIN price and predicted value based on nonlinear rolling models for Chicago, Gulf, and NYH spreads, January 6, 2011–October 4, 2018

will be in effect next year depend only on whether it is currently in effect. In practice, over this period the trade news regularly tracked developments in the U.S. Congress that affect the probability that the tax credit would be extended or reinstated, and indeed it is rational for market participants to assess that probability using all available information, not just the current status of the credit. This observation has two implications. First, a real-time market probability of $B_T = 1$ can be estimated directly from pricing discrepancies. Second, that real-time probability can then be used to construct an alternative, perhaps improved, real-time D4 price.

We investigate these two implications by modifying the model to allow for probabilities of $B_T = 1$ to evolve continuously over time.

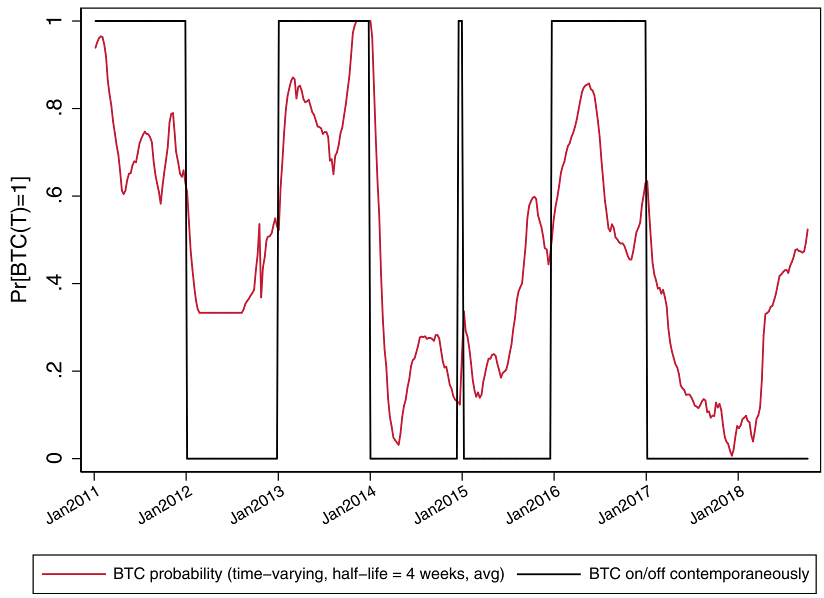

Figure 7. Estimated real-time market probabilities $\hat{p}^{EWMA}_{T|t-1}$ that the biodiesel tax credit will be in place in the next year

Note: The red line is the probability $\hat{p}^{EWMA}_{T|t-1}$ that the biodiesel tax credit will be in effect on December 31 of the next calendar year, based on information through week $t-1$, as estimated using an exponentially weighted moving average filter of $p_{T|t}$ in equation (8). The binary line indicates whether the tax credit is currently in effect at date t.

Specifically, use equation (3) to rewrite the second expression in equation (4) as $P_t^{D4} = e^{-r(T-t)}\{f_{0,t} - (f_{0,t} - f_{1,t})\Pr[B_T = 1|S_t, B_t]\}/1.5$. Let $p_{T|t} = \Pr[B_T = 1|\Omega_t]$ denote the probability at time t that the tax credit is in effect at date T, given all information Ω_t available to the market at time t. Under the assumption that the spread has the Gaussian distribution $N[s_t, \sigma^2(T-t)]$ conditonal on Ω_t, then an extension of equation (4) to this larger information set implies that $P_t^{D4} = e^{-r(T-t)}\{f_{0,t} - (f_{0,t} - f_{1,t}) p_{T|t}\}/1.5$. Rearranging this expression yields

$$(8) \quad p_{T|t} = \frac{1.5e^{r(T-t)}P_t^{D4} - f_{0,t}}{f_{1,t} - f_{0,t}}.$$

The probability $p_{T|t}$ uses the current value of the D4 price, so it is not available in real time. A real-time estimate of $p_{T|t}$ can be constructed either by using covariates observed in real time or by using a time-series filter. Because such covariates (e.g., real-time rumors of congressional proclivities toward biodiesel) are not readily available, we adopt the latter approach and estimate $p_{T|t}$ using an exponentially weighted moving average time-series filter with a weighting half-life of four weeks, computed by averaging data on $p_{T|t}$ through week t.[14] Because $f_{0,t}$ and $f_{1,t}$ are available in real time and because the lag of the resulting estimated probability, $\hat{p}^{EWMA}_{T|t-1}$, is a function only of data through time t-1, these filtered probabilities can be used to produce the real-time D4 price:

$$(9) \quad P_t^{D4,EWMA} = e^{-r(T-t)}\left\{f_{0,t} - (f_{0,t} - f_{1,t})\hat{p}^{EWMA}_{T|t-1}\right\}/1.5.$$

The resulting real-time probabilities $\hat{p}^{EWMA}_{T|t-1}$ and real-time implied price $P_t^{D4,EWMA}$ are plotted in figures 7 and 8, respectively. Three features bear emphasis. First, consistent with the Markov model estimates, the filtered probability is greater if the biodiesel tax credit is currently in effect than if it is not. Second, unlike the Markov model estimate, there is substantial month-to-month variation in the real-time estimate of the

[14] Specifically, let $T(t)$ denote December 31 of the calendar year following week t. Then, ignoring initial conditions for convenience, the filtered probability is $\hat{p}^{EWMA}_{T(t)|t} = (1-\alpha)^{-1}\sum_{i=0}^{\infty}\alpha^i p_{T(t-i)|t-i}$, where α is the weighting parameter. We estimate the probability by applying the filter separately to the Chicago, Gulf, and NYH data, then averaging the probabilities. To avoid probabilities outside the unit interval, $\hat{p}^{EWMA}_{T(t)|t}$ was truncated at 0 and 1, although this truncation was only binding in late 2013. The lag of this series, that is, $\hat{p}^{EWMA}_{T(t-1)|t-1}$, uses data only from week t-1 and earlier, so it is available in real time in week t.

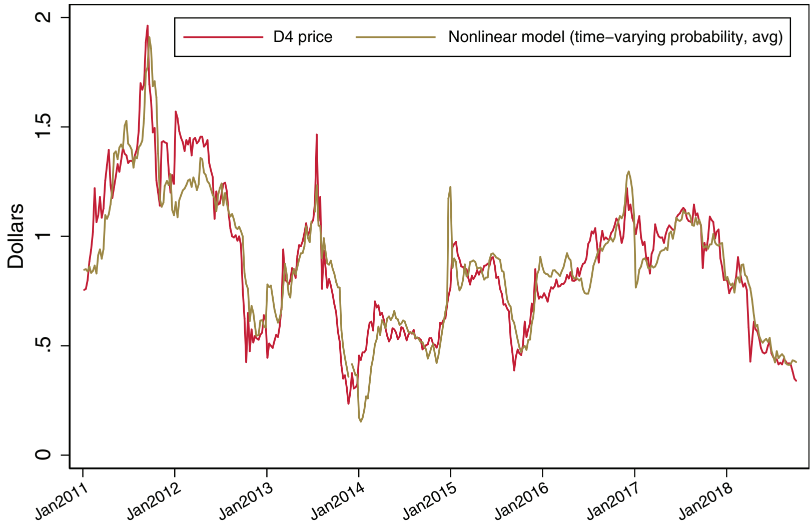

Figure 8. D4 RIN prices predicted using the real-time feasible probability $\hat{p}^{EWMA}_{T|t-1}$

Note: The predicted price is $P_t^{D4,EWMA}$, computed using the real-time feasible probability $\hat{p}^{EWMA}_{T|t-1}$ (shown in figure 7) according to equation (9).

probability. Third, as seen in figure 8, using these probabilities essentially eliminates the multi-week departures of predicted from actual prices.

One limitation of this approach is that the time-series filter does not adjust for a discontinuous jump in the probability when t changes calendar years, which one would expect to happen because then T also changes calendar years. Indeed, this is likely to be the reason for transitory pricing errors evident in figure 8 at the start of 2014, 2015, 2016, and 2017. Even with this caveat in mind, these estimates suggest that market participants track real-time information to update their beliefs about whether the tax credit will be in effect on the RIN expiration date, consistent with market rationality and fundamentals-based pricing. Indeed, this evidence points to changing beliefs about the tax credit as another fundamental source of RIN price variability.

Extension to the Spread Depending on the Tax Credit

The analysis so far assumes that the status of the biodiesel tax credit does not affect the spread. However, the estimates in the final column of table 1 provide some weak evidence that the level of the spread depends on whether the tax credit is in effect. In this section, we provide two possible explanations for this dependence and then extend the earlier model to allow for the level of the spread to depend on whether the tax credit is in place.

One explanation for this dependence is a "race" by diesel blenders to take advantage of the \$1 per gallon blenders' tax credit that expired at the end of 2011, 2013, and 2016 (Irwin 2017). If blenders perceive that there is a substantial probability that the expiring credit will not be renewed, then, in the face of a binding and continuing BBD mandate, it is rational for blenders to take advantage of the tax credit while it is still in place and thus to purchase biodiesel at a discount in the current year in excess of this year's BBD mandate. Because excess D4 RINs detached in this way can be used to meet next year's mandate, and because any blending limit on BBD is not binding during this period (no so-called BBD "blend wall"), this increase in blenders' demand will bid up the price of biodiesel in the current year. If blenders were confident that the tax credit would not be renewed, blenders would bid up the price by as much as \$1 over what would otherwise prevail; if they were uncertain, they would still have the incentive to bid up the price by \$1 times the probability that it would not be renewed.

As shown in figure 9, there is in fact a clear pattern of biodiesel production increasing in anticipation of the expiration of the biodiesel tax credit. Here, we proxy biodiesel production by monthly D4 RIN generation (annualized)

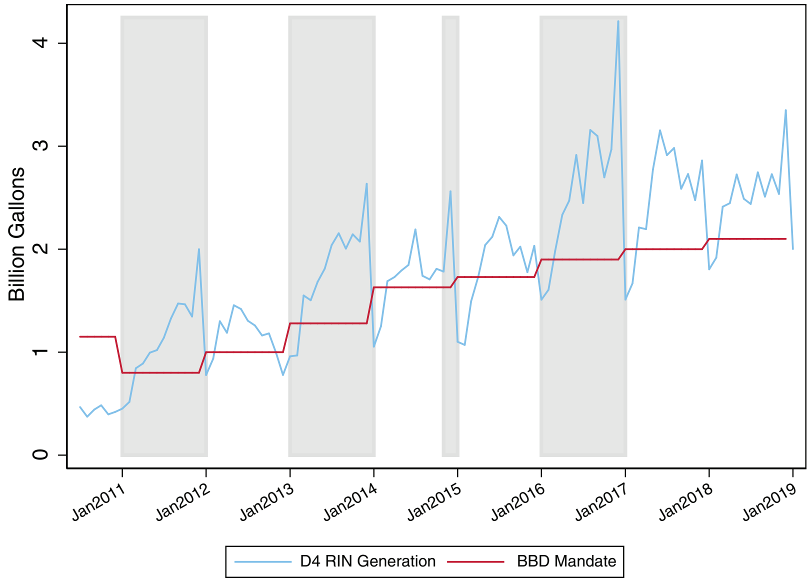

Figure 9. Monthly D4 biodiesel RIN generation (annualized) and biomass-based diesel renewable volume obligation, July 2010–October 2018 D4 RIN generation is collected from the EPA website: https://www.epa.gov/fuels-registration-reporting-and-compliance-help/rins-generated-transactions. Shaded areas represent periods when the biodiesel tax credit was in place but expired at the end of the calendar year.

and compare it to the annual BBD RVO from July 2010 through October 2018. There is a general pattern of biodiesel production exceeding the BBD RVO because BBD was the marginal gallon for most of this period for filling the advanced and conventional RVOs in addition to the BBD RVO. In the years when the biodiesel tax credit was in place for the entire calendar year and expired at the end of the year (2011, 2013, and 2016), there is a clear pattern of rising D4 RIN generation throughout the year and a dramatic upward spike in December just before the credit expires. It is especially interesting to observe that this pattern even holds at the end of 2014 when the tax credit was reinstated for just the last three weeks of the calendar year, which produced another dramatic upward spike in biodiesel production in December 2014.

Figure 10 provides evidence on the impact of increasing biodiesel production on prices and profits in years when the tax credit expired. Specifically, the figure plots the biodiesel price versus a simple break-even relationship between the biodiesel price and the price of the marginal feedstock during this period, soybean oil. This simple model posits that the break-even price for a representative Iowa biodiesel producer is $0.6+7.55P_t^{SoyOil}$, where 7.55 is the number of pounds of soybean oil assumed to produce a gallon of biodiesel, P_t^{SoyOil} is the Iowa price of soybean oil, and the intercept captures the non-oil variable costs of the plant, estimated to be \$0.60 per gallon. This simple break-even price tracks the biodiesel price very closely outside of the spikes in 2011, 2013, and 2016, the three years in which the tax credit was in place but was slated to expire. Note that the spike in biodiesel prices relative to costs builds within each of the three years, consistent with increasing pressure by blenders to take advantage of a tax credit that might not be reinstated. If the reinstatement of the tax credit at the beginning of 2011, 2013, and 2016 drove biodiesel prices upward, we should observe a large spike in biodiesel prices relative to costs early in the calendar year, but we do not.[15]

A second possible explanation for the dependence of the spread on the tax credit is that the EPA takes the presence of the biodiesel tax credit into account in its annual rulemakings that set the renewable volume obligations and percentage standards, and indeed it has at times

[15] Biodiesel prices also exceed the simple break-even price starting in the second half of 2017. We attribute this to the imposition of countervailing import duties on biodiesel imports from Argentina and Indonesia, the two largest importers to the United States.

acknowledged doing so.[16] If the EPA treats the tax credit as effectively shifting out available supply when setting the percentage standard, then some or all of the tax credit accrues to biorefiners and to feedstock producers, such as soybean farmers.

These two explanations—blenders bidding up the price of biodiesel in advance of the tax credit expiration and the EPA taking the tax credit status into account—are not mutually exclusive, and there is evidence that, in fact, both channels were operating. We therefore extend our model to allow for the biodiesel tax credit to have an effect on the BBD price and thus on the spread.

Specifically, consistent with equations (4) and (5) in table 1, write the spread as

$$(10) \quad S_t = \mu + \beta B_t + u_t, \text{where } u_t = u_{t-1} + \varepsilon_t,$$

where ε_t is serially uncorrelated. Because a fraction, β, of the tax credit accrues to biorefiners in the form of higher biodiesel prices, only a fraction, $1 - \beta$, remains to offset the price of the D4 RIN. Accordingly, the D4 RIN fundamental price is given by

$$(11) \quad P_t^{D4} = e^{-r(T-t)} E_t \max[S_T - (1-\beta)B_T, 0]/1.5,$$

where S_t follows from equation (10) and B_t continues to follow the Markov process in equation (3).

Although the economics of the pricing formula in equation (11) are quite different from our base model, it turns out that the pricing formula and predicted prices are identical in the linear model and are nearly identical in the nonlinear model. We show this in the linear case, in which the probability of a negative fundamental is assumed to be zero. Then

$$\begin{aligned}(12) \quad P_t^{D4} &= e^{-r(T-t)}[E_t S_T - (1-\beta)E_t B_T]/1.5 \\ &= e^{-r(T-t)}\{S_t - (1-2\beta)(1-q)(1-B_t) - [(1-\beta)p + \beta]B_t\}/1.5,\end{aligned}$$

where the second line of equation (12) follows by substituting equation (3), $E_t S_T = \mu + \beta E_t B_T + u_t$, and equation (10) into the first line of equation (12) and simplifying.

The key observation is that the terms in brackets in the second line of equation (12) are the same as in the baseline linear equation (6), except that the coefficients have a different interpretation. Because the terms are the same, the predicted prices are the same in the linear model. In the nonlinear model, the predicted price depends on the value of β; however, the fact that the nonlinear and linear models produce very similar predicted prices indicates that in practice this dependence is very weak, so the alternative nonlinear model based on equation (11) will differ negligibly from the base nonlinear model.

When the estimate of β from table 1 is used, along with the expressions for the coefficients in equation (12), one obtains different estimates of q and p than in the base model. For example, estimated over the full sample using the Chicago data, the resulting estimates are $\hat{q} = 1.19$ and $\hat{p} = 0.28$. The estimated value of q exceeds 1, which is not sensible; in any event, both these estimates suggest substantially lower market assessments of whether the tax credit is in effect in the coming year.

We stress that although β is identified from equation (10), in practice it is very imprecisely estimated: formally, because it compares regime means of random walks, and informally, because the tax credit only shifted a few times in our sample, so there are few "experiments" with which to estimate β. Indeed, a 95% confidence interval for β includes both 0 and 1 for each of the three spreads, respectively, corresponding to the cases that none and all of the tax credit accrues to the biodiesel producer. In the (nonrejected) case that all of the tax credit accrues to the

[16] For example, in the 2013 rulemaking, the EPA discusses public comments on whether the tax credit should be taken into account in its rulemakings:

> Recently, the tax credit for biodiesel was reinstated after having expired at the end of 2011. This tax credit, applicable retroactively to 2012 and through the end of 2013, may provide additional incentive to produce and consume biodiesel volumes in excess of the 1.28 bill gal requirement. While one party commented that the biodiesel tax credit should not be a relevant factor, the existence of a tax credit affects the likelihood that biodiesel volumes in excess of 1.28 bill gal will be produced. Therefore, it is a relevant consideration in determining whether there are likely to be sufficient volumes of advanced biofuel available to meet the statutory volume requirement of 2.75 bill gal. (78 FR 49813, Aug. 15, 2013)

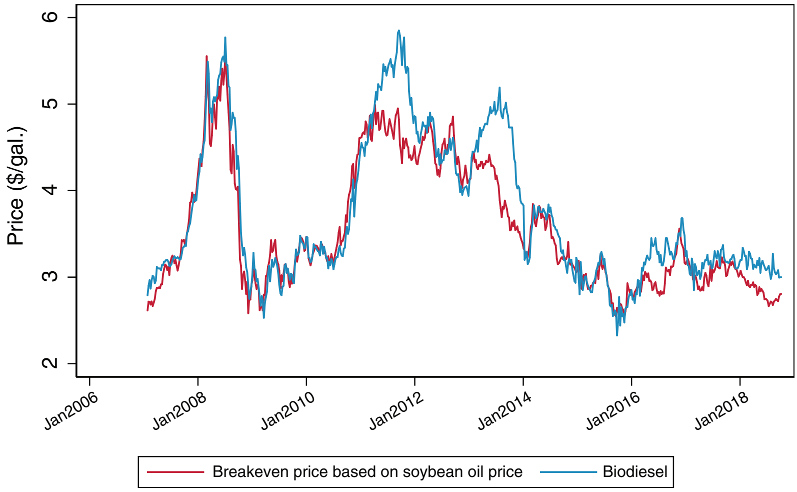

Figure 10. Weekly (Friday) biodiesel price and simple break-even price at a representative Iowa plant, January 26, 2007–October 4, 2018

biodiesel producer, the Markov probability p is not identified (p drops out if $\beta = 1$ in equation (12)). Thus, p is weakly identified in this application.

Conclusions

There has been considerable discussion in recent years about the need for regulation of the renewable identification number (RIN) market in order to deter manipulation. For example, President Trump issued a statement in October 2018 directing the Environmental Protection Agency (EPA) to initiate a rulemaking to address RIN price manipulation claims and increase transparency in the RIN market. The most important conclusion from our work is that movements in D4 biodiesel RIN prices at frequencies of a month or longer are well explained by two economic fundamentals: (*a*) the spread between the biodiesel and ULSD prices and (*b*) whether the $1 per gallon biodiesel tax credit is in effect. To explain RIN price volatility, one does not need to resort to market irrationality or market manipulation; rather, one need look no further than the supply and demand for biodiesel, the setting of statutory volumes in the renewable fuel standard (RFS), and the history of the U.S. Congress intermittently extending, or not, the biodiesel tax credit. Hence, new regulations that restrict trading in the RIN market could end up reducing liquidity, increasing price volatility, and reducing the efficiency of price discovery, all without addressing the real reasons for RIN price volatility.

We also analyze three economic channels whereby the biodiesel tax credit affects RIN prices: (*a*) an expectational channel in which the tax credit does not affect the spread, but affects the D4 price by reducing the subsidy that the D4 RIN would otherwise provide; (*b*) an expectational channel in which the imminent expiration of the tax credit induces buying biodiesel before the deadline and thus increases the spread; and (*c*) a regulatory channel in which the EPA sets the biomass-based diesel (BBD) mandate based on whether the tax credit is likely to be in effect. All three channels provide predicted D4 prices that are identical in the linear model and are nearly so in the nonlinear model; nevertheless, the parameters have different interpretations under the first channel alone than if the second two are operational. Unfortunately, the relevant parameters differentiating these models are weakly identified because of the persistence of the spread and the infrequency with which the tax credit regime changes. We provided evidence, both econometric and institutional, that all three of these channels are in operation; however, sorting out their relative contributions is left to further research.

The Markov model for the biodiesel tax credit makes the very simple assumption that market predictions of whether the credit will be in effect next year depend only on whether it is currently in effect. In practice, the trade news over this period regularly tracked developments in the U.S. Congress that affect the probability that the tax credit would be extended or reinstated, and indeed we find evidence that market participants update their beliefs about the future of the tax credit in real time. The additional RIN market volatility arising from these evolving beliefs is not an indication of market failure; rather, it reflects the rational response of market participants to on-again, off-again congressional treatment of the biodiesel tax credit. Taken together, we interpret the good fit of our predicted prices, especially of the real-time pricing model, as a strong indication that the RIN market is driven by fundamentals and that any actions to reduce price volatility need to look at the intrinsic structure of the RFS, not at market failures.

Supplementary Material

Supplementary material are available at *American Journal of Agricultural Economics* online.

References

Blewitt, Laura, and Zachary Mider. 2016. Icahn Calls on EPA to Fix "Mother of All Short Squeezes." *Bloomberg*, August 15. Available at: https://www.bloomberg.com/news/articles/2016-08-15/carl-icahn-calls-on-epa-to-fix-mother-of-all-short-squeezes. Accessed December 23, 2019.

Bracmort, Kelsi. 2019. Renewable Fuel Standard (RFS): An Overview. Congressional Research Service Report for Congress R43325. Available at: https://fas.org/sgp/crs/misc/R43325.pdf. Accessed December 23, 2019.

Cronshaw, Mark B, and Jamie Brown Kruse. 1996. Regulated Firms in Pollution Permit Markets with Banking. *Journal of Regulatory Economics* 9(2): 179–86.

Irwin, Scott. 2014. Understanding the Behavior of Biodiesel RINs Prices. *farmdoc daily*, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, 4(196), October 10. Available at: https://farmdocdaily.illinois.edu/2014/10/understanding-behavior-of-biodiesel-rins-prices.html.

———. 2017. The Profitability of Biodiesel Production in 2016. *farmdoc daily*, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, 7(38), March 1. Available at: https://farmdocdaily.illinois.edu/2017/03/the-profitability-of-biodiesel-production-in-2016.html.

———. 2019. Revisiting the Value of Ethanol in E10 Gasoline Blends. *farmdoc daily*, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, 9(60), April 4. Available at: https://farmdocdaily.illinois.edu/2019/04/revisiting-the-value-of-ethanol-in-e10-gasoline-blends.html. Accessed December 23, 2019.

Irwin, Scott, and Darrel Good. 2013. RINs Gone Wild? *farmdoc daily*, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, 3(138), July 19. Available at: https://farmdocdaily.illinois.edu/2013/07/rins-gone-wild.html.

———. 2017. How to Think about Biodiesel RINs Prices under Different Policies. *farmdoc daily*, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, 7 (154), August 23. Available at: https://farmdocdaily.illinois.edu/2017/08/how-to-think-about-biodiesel-rins-prices.html. Accessed December 23, 2019.

Kling, Catherine, and Jonathan Rubin. 1997. Bankable Permits for the Control of Environmental Pollution. *Journal of Public Economics* 64(1): 101–15.

Knittel, Christopher R, Ben S Meiselman, and James H Stock. 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): 1081–119.

Lade, Gabriel E, and James Bushnell. 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): 563–92.

Lade, Gabriel E, Cynthia Lin Lawell, and Aaron Smith. 2018. Policy Shocks and Market-based Regulations: Evidence from the Renewable Fuel Standard. *American Journal of Agricultural Economics* 100(3): 707–31.

Litzenberger, Robert H, and Nir Rabinowitz. 1995. Backwardation in Oil Futures Markets: Theory and Empirical

Evidence. *Journal of Finance* 50(5): 1517–45.

Schnepf, Randy, and Brent D. Yacobucci. 2013. Renewable Fuel Standard (RFS): Overview and Issues. Washington, DC: U.S. Congressional Research Service, Report for Congress R40155, March. Available at: https://fas.org/sgp/crs/misc/R40155.pdf. Accessed December 23, 2019.

Voegele, Erin. 2013. Members of Congress Ask for RIN Market Investigation. *Ethanol Producer Magazine*, October 23. Available at: http://ethanolproducer.com/articles/10387/members-of-congress-ask-for-rin-market-investigation. Accessed December 23, 2019.