

The Price of Biodiesel RINs and Economic Fundamentals

October 28, 2018

Scott H. Irwin

Department of Agricultural and Consumer Economics, University of Illinois – Urbana

Kristen McCormack

Harvard Kennedy School, Harvard University

James H. Stock*

Department of Economics and Harvard Kennedy School, Harvard University
and the National Bureau of Economic Research

*We thank Sandra Dunphy, Cynthia Lin Lawell, Eben Lazarus, and Aaron Smith for helpful comments and/or guidance.

The Price of Biodiesel RINs and Economic Fundamentals

Abstract

The D4 RIN is the tradable compliance certificate for the biomass-based diesel 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 prices 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 US fuel supply system.

Keywords: biodiesel, D4, fundamental, option, RIN, option, tax credit

JEL Categories: G13, Q4

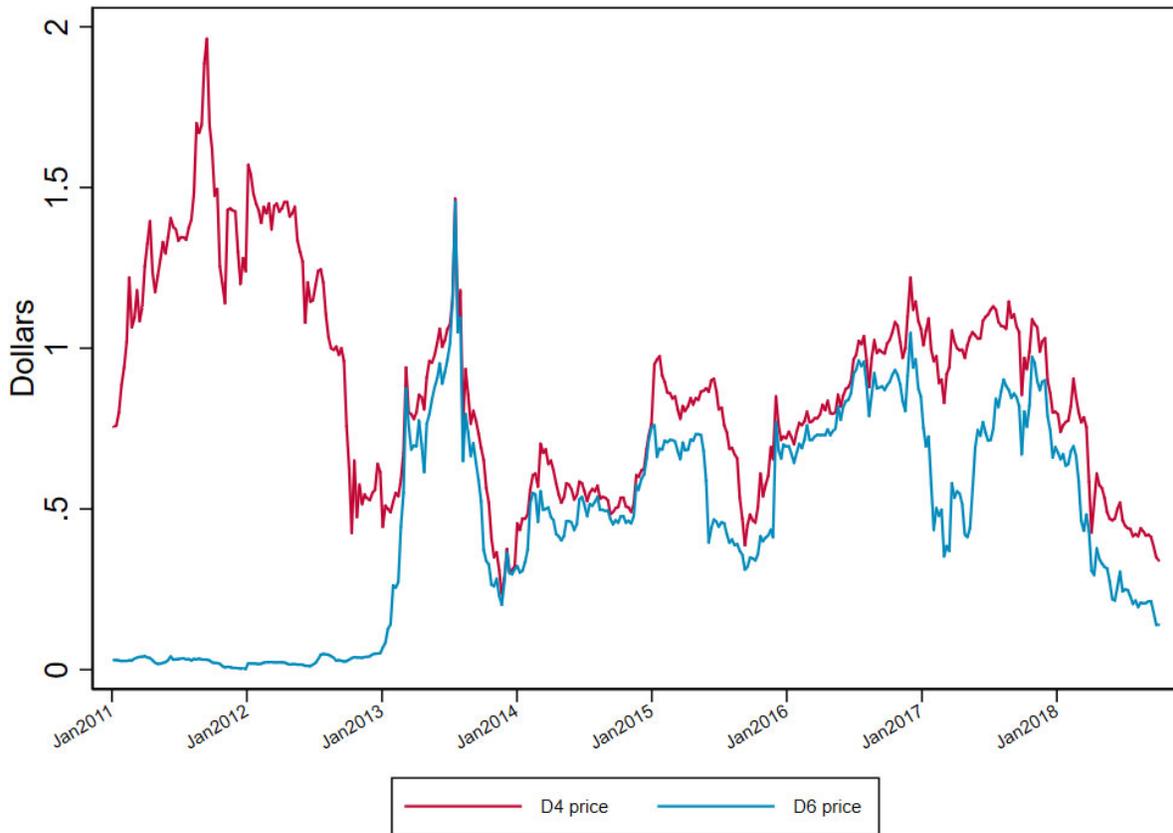
1. Introduction and Summary of Results

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 by the US Environmental Protection Agency (EPA). Refiners and importers of gasoline and diesel fuel (“obligated parties”) demonstrate compliance with the RFS using the RIN system. A RIN (Renewable Identification Number) is a unique electronic certificate that is created (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 EPA to demonstrate compliance.

The total market value of RINs retired in 2017 was \$14 billion (Irwin and Stock 2018). 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 is mainly comprised 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. This volatility has 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.

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 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. As can be seen in Figure 1, this cap was binding much of the time from February 2013 to June 2018. 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 RIN pricing, whereas this is not possible for D6 RINs, as explained in Section 2. 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.

Figure 1. Weekly D4 and D6 RIN prices, January 6, 2011 – October 4, 2018



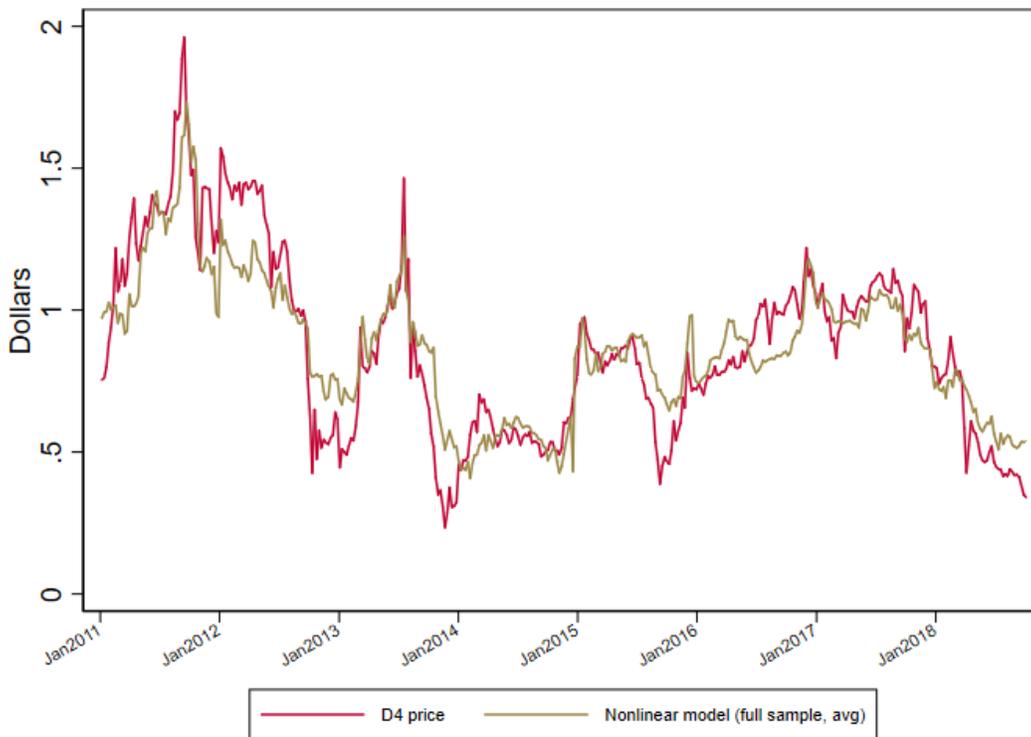
Notes: Weekly data (Thursday) are from OPIS. Prices are of RINs generated in the current calendar year (current-year vintage RINs).

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. Because it can be retired any time during this window, the D4 RIN is in effect an American call option. As explained in Section 2, economic theory indicates that the price of the underlying asset depends on (i) the price spread between biodiesel and its petroleum substitute, ultra low-sulfur diesel (ULSD), and on (ii) whether or not the 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—which 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 the additional assumption that the spread process is Gaussian. The second

expression does not require normality but assumes that the probability of a negative fundamental is negligible.

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 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. 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).

Figure 2. Weekly D4 RIN price and its predicted value based on the nonlinear economic fundamentals model, averaged over three markets (Chicago, Gulf, New York Harbor)



Notes: The predicted price from the nonlinear option pricing model is the average of three predicted prices based on three biodiesel-ULSD spot spreads: Chicago, Gulf, and (starting Oct. 25, 2012), New York Harbor.

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) use an option pricing framework to develop a joint model for pricing multiple RINs including the nesting cap. Relative to their paper, we use a more immediate measure of fundamentals which we would expect should improve fit (biodiesel prices and ULSD, whereas they use soybean oil and crude oil prices), we incorporate the biodiesel tax credit, and, by focusing solely on the D4 RIN, we are able to obtain a closed-form solution for the option price. There is also a growing literature on the pass-through of RIN prices through the fuel supply chain (see Knittel, Meiselman and Stock (2017) and Lade and Bushnell (forthcoming) for references); 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.

2. D4 RIN Pricing Model

At current blend ratios, pure petroleum diesel (ULSD) and a blend of biodiesel and ULSD are effectively perfect substitutes, after adjusting for biodiesel having 92.7% the energy content of ULSD. Because biodiesel is more expensive than ULSD, it would not enter the market were it not for subsidies. The two national-level subsidies are through the RFS, in the form of the D4 RIN, and the 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, then with the tax credit in place. Because the biorefiner produces both the wet 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 biodiesel generates 1.5 D4 RINs, absent the tax credit the price based on contemporaneous economic fundamentals at date t is $P_t^* = \max(P_t^{Biodiesel} - 0.927P_t^{ULSD}) / 1.5, 0$, 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.¹ The fundamental price is truncated at zero because the, if the BBD price is less than the energy-adjusted ULSD price, the BBD mandate will not be binding so the D4 RIN price be zero.

The biodiesel blenders' tax credit provides a tax credit of \$1 for each gallon of BBD that is blended with ULSD. Because ULSD and BBD are perfect substitutes (energy-adjusted), under perfect competition the blenders' tax credit will accrue to the biorefiner (and thus to the feed stock producer). Thus, in our base model, the fundamental value of the D4 RIN at date t is

$P_t^* = \max(P_t^{Biodiesel} - 0.927P_t^{ULSD} - B_t) / 1.5, 0$, 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.² Thus, it is an American option with no dividend and expiration date of December 31 in the year after it was generated. As a result, the D4 RIN can be priced as a European option, using its fundamental value on the expiration date as the price of the underlying asset. For a risk-neutral firm³, this gives the pricing formula in Lade, Lin Lawell, and Smith (2018, equation (2)),

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

where $S_t = P_t^{Biodiesel} - 0.927P_t^{ULSD}$, T is December 31 of the year after it was generated, E_t denotes the expectation conditional on information at time t , and the factor of 1/1.5 adjusts for the fact

¹ The fundamental price here is expressed for biodiesel, which 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 conventional 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.

² For example, a 2017 obligation can be met using a RIN generated in 2016 or one generated in 2017. Thus, the final date for generating a RIN to meet a 2017 obligation is Dec. 31, 2017.

³ Risk neutrality is not needed to obtain (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.

that one gallon of biodiesel generates 1.5 RINs. The maximum in (1) imposes the condition that the price of the D4 RIN cannot be negative.

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

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

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

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, it is in place for a 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. 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 (1) – (3).

We provide two closed-form solutions for the D4 RIN price. The first further assumes that the spread innovations are Gaussian, so that the conditional distribution of S_T given $S_t = s_t$ and $B_t = b_t$ is $N[s_t, \sigma^2(T-t)]$, 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). Under these assumptions, a calculation yields,

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

where

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

where $\phi(\cdot)$ is the normal density and $\Phi(\cdot)$ is the cumulative normal distribution.⁴

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

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

Equation (6) also obtains as a limiting approximation to (4) and (5) for small σ .

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 a higher octane value than petroleum gasoline so at blends less than 10% it is used as an octane booster. At blends greater than 10%, it faces the so-called D10 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 on commodity exchanges, the ethanol demand price depends on the blend ratio and is not observed. In addition, the nesting structure of the RFS allows D4 RINs to be used to meet the conventional mandate, further complicating the analysis. For additional discussion, see Lade, Lin Lawell, and Smith (2018).

3. Data and Empirical Results

We use weekly 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,

⁴ 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 (4) and (5).

and the New York Harbor. The D4 price data and the Chicago and Gulf fuel price data span September 3, 2009 – Oct. 4, 2018. The New York Harbor data span Oct. 19, 2012 – Oct. 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.⁵ For the interest rate we use the 6-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 . 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.⁶

Columns (4) and (5) in Table 1 examine the possibility that the level of the spread depends on whether or not the biodiesel tax credit is in place. The evidence suggests that (a) the Chicago spread averages \$0.74 higher if the \$1 tax credit is in effect (\$0.77 for the Gulf and \$0.67 for the NYH spread, which are over a shorter sample), and (b) the residual from regression

⁵ The biodiesel blenders’ tax credit was in place contemporaneously for the full calendar years of 2007-2009, 2011, 2013, and 2016. For calendar years 2010, 2012, 2014-15, and 2017, the tax credit was restored retroactively. For 2018, the tax credit has expired and, as of this writing, it is not known whether it will be restored retroactively. Thus, for 2010, 2012, 2014-15, and 2017, 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 reinstated 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.

⁶ 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.

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 3 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 BBD tax credit is large economically as well as statistically (the R^2 of regression (4) for Chicago is 0.45); however, because the spread is integrated of order one and there are only a few times that the tax credit turns off and on contemporaneously, the coefficient on the tax credit is estimated imprecisely.⁷ Because this dependence of the spread on B_t is perfectly colinear 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 it was regularly reinstated retroactively after it expired. Thus, with the benefit of the full data set, it looks like the probability of the tax credit being on in the future was always 1 regardless of whether it was currently in effect. In real time, however, there was always uncertainty as to whether Congress would in fact enact the

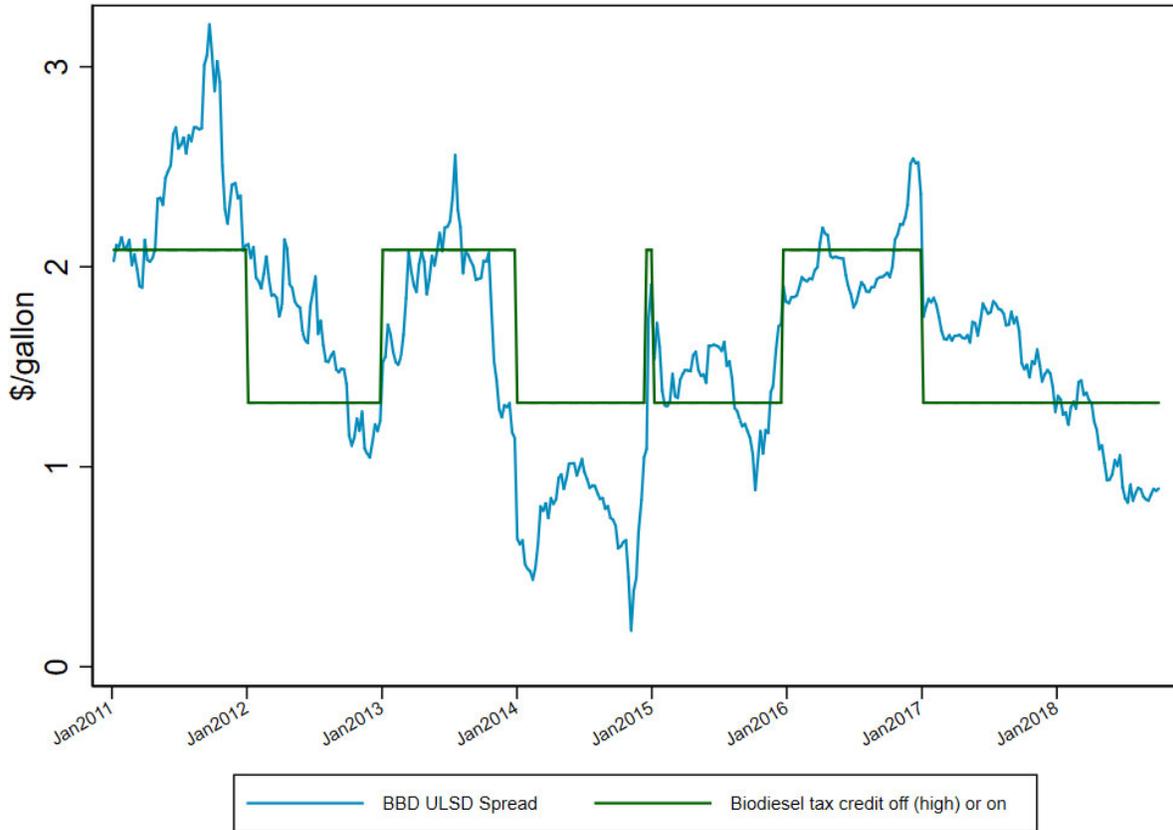
⁷ The regression in column (4) is $S_t = \alpha + \beta_B B_t + u_t$, where (under the assumptions of the text) u_t follows a random walk with $\text{var}(\Delta u_t) = \sigma_{\Delta u}^2$. The persistence of the tax credit and the random walk assumption for the error term leads to a nonstandard sampling distribution for $\hat{\beta}_B$, the OLS estimator of β_B . One can 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 $[\cdot]$ 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, $T^{-1/2}(\hat{\beta}_B - \beta) \Rightarrow \sigma_{\Delta u} \int_0^1 \delta(\tau) W^\mu(\tau) d\tau / \left(\int_0^1 \delta^2(\tau) d\tau \right)^{1/2}$, 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 $\text{var}(\hat{\beta}_B) = T \sigma_{\Delta u}^2 \int_0^1 \int_0^1 \delta(\tau) \delta(r) \min(\tau, r) d\tau dr / \left(\int_0^1 \delta^2(\tau) d\tau \right)^2$ (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 $\sigma_{\Delta u}^2$ and by numerical evaluation of the double integral.

Table 1. Autoregressive models of the spread S_t

	(1)	(2)	(3)	(4)	(5)
Dependent variable	S_t	ΔS_t	ΔS_t	S_t	\hat{u}_t
Regressor	S_{t-1}, \dots, S_{t-6}	$\Delta S_{t-1}, \dots, \Delta S_{t-5}$	$\Delta S_{t-1}, \dots, \Delta S_{t-5}$ $B_{t-1} \times \Delta S_{t-1}, \dots, B_{t-5} \times \Delta S_{t-5}, B_t$	B_t	$\hat{u}_{t-1}, \dots, \hat{u}_{t-6}$
A. Chicago Jan. 6, 2011- Oct 4, 2018 ($n = 405$)					
Intercept	0.033 (0.021)	-0.003 (0.005)	-0.013 (0.007)	1.623 ^b	-0.001 (0.006)
Coefficient on B_t	--	--	0.018 (0.011)	0.764 (0.570) ^b	
ADF test	-2.16*	--	--	--	--
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)	--	
B. Gulf Jan. 6, 2011- Oct 4, 2018 ($n = 405$)					
Intercept	0.030 (0.020)	-0.003 (0.006)	-0.015 (0.007)	1.650 ^b	-0.002 (0.006)
Coefficient on B_t	--	--	0.020 (0.011)	0.796 (0.570) ^b	
ADF test	-1.97*	--	--	--	--
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 (0.145)	--	0.74 (0.596)
F -test, all interactions (p -value)	--	--	1.03 (0.400)	--	
C. New York Harbor Oct. 19, 2012- Oct 4, 2018 ($n = 307$)					
Intercept	0.044 (0.028)	-0.001 (0.006)	-0.012 (0.007)	1.477 ^b	-0.001 (0.006)
Coefficient on B_t	--	--	0.026 (0.014)	0.691 (0.422) ^b	
ADF test	-2.09*	--	--	--	*
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 (0.277)	--	0.70 (0.623)
F -test, all interactions (p -value)	--	--	1.27 (0.277)	--	

Notes: Standard errors are in parentheses below coefficients, p -values are in parentheses below F -statistics; $n = 376$. All regressions include an intercept. \hat{u}_t in column (5) are the residuals from regression (4) of S_t on B_t . ^a indicates that a value was imposed. ^b indicates that standard errors for coefficient on B_t in the levels regression (4) are computed using Gaussian functional limit described in text; standard error for the intercept is not substantively relevant and not computed. ADF test rejects the null of a unit root at the **1% *5% significance level.

Figure 3. Weekly Chicago BBD-ULSD spread and its predicted value based on whether the biodiesel tax credit is in effect contemporaneously



Notes: Weekly data (Thursday) are from OPIS. Prices are of RINs generated in the current calendar year (current-year vintage RINs).

credit in the next year or restore it retroactively (even though it always did take one of these actions).

D4 pricing: nonlinear and linear models. For the Chicago and Gulf spreads, the linear models were estimated over Jan. 6, 2011 – Oct. 4, 2018. For the NYH spread, the linear model was estimated over the full span of the available data, Oct. 19, 2012 – Oct. 4, 2018. Constructing the terms $f_{0,t}$ and $f_{1,t}$ in the nonlinear model requires an additional parameter, the variance of ΔS_t . We estimated this variance using a rolling 52-week retrospective window. For the Chicago and Gulf spreads, we have more than a year of pre-sample data available to estimate the initial variance, so the model estimation sample is January 6, 2011 – Oct. 4, 2018 and all observations use the 52-week retrospective rolling variance. For the NYH spread, the first observation for our

NYH data is Oct. 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 52 weeks, and thereafter a 52-week rolling estimator. This allows us to estimate the NYH nonlinear model over the span Oct. 25, 2012 – May 31, 2018. The remaining two free parameters in the nonlinear model, q and p , are estimated by OLS estimation of (4).

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 Jan. 6, 2011 to Oct. 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 approximately a 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 a standard error of the estimates for the Gulf.

Table 2. Results for D4 pricing model

	(1)	(2)	(3)	(4)	(5)	(6)
	Chicago	Chicago	Gulf	Gulf	NYH	NYH
	nonlinear	linear	nonlinear	linear	nonlinear	linear
$1-q$	0.186 (0.064)	0.098 (0.048)	0.183 (0.065)	0.111 (0.045)	0.253 (0.046)	0.146 (0.038)
p	0.664 (0.065)	0.614 (0.054)	0.697 (0.058)	0.658 (0.048)	0.751 (0.060)	0.696 (0.043)
R^2	0.789	0.769	0.807	0.811	0.776	0.763
<i>Sample</i>	1/6/11 – 10/4/18	1/6/11 – 10/4/18	1/6/11 – 10/4/18	1/6/11 – 10/4/18	10/25/12 – 10/4/18	10/11/12 – 10/4/18
n	403	403	403	403	309	311

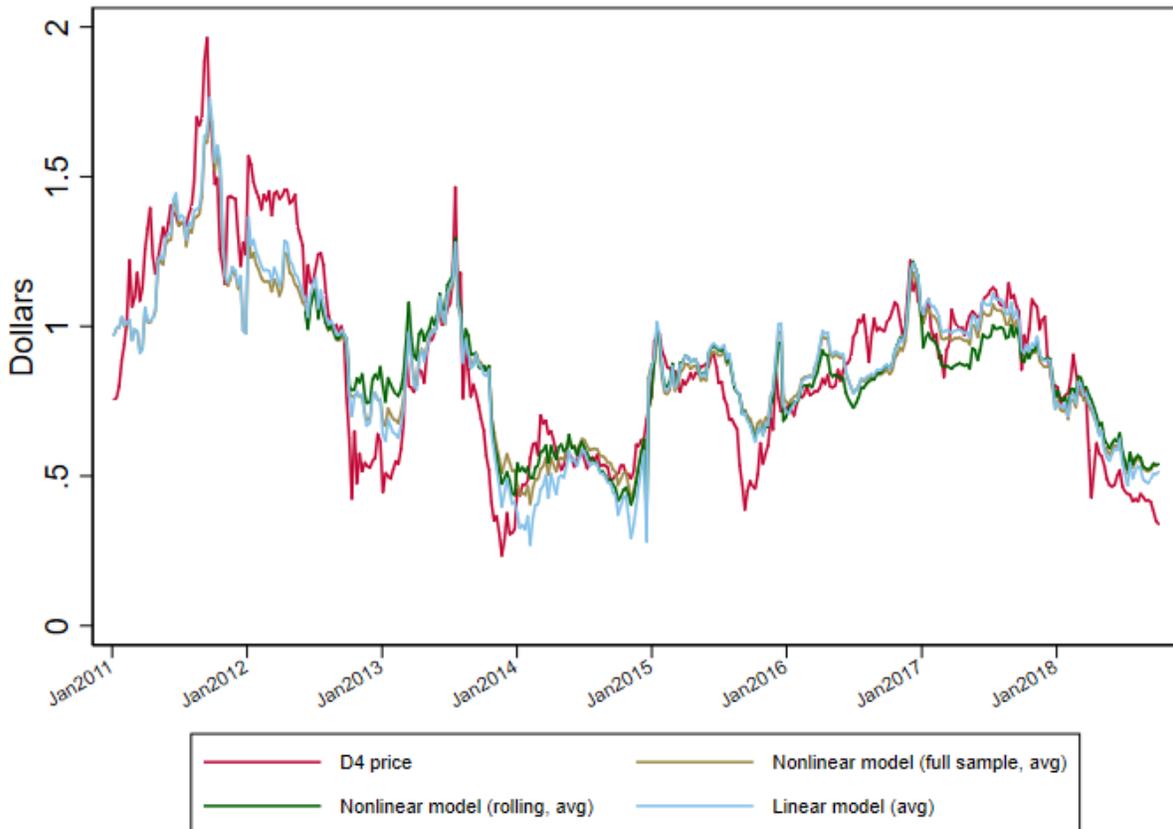
Notes: Standard errors are Newey-West with 20 lags.

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 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, estimated over $t-105, \dots, t-1$; 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 52-week retrospective window ending in $t-1$). 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 4 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.

Figure 4. Weekly D4 RIN price and predicted value based on linear and nonlinear models, both full-sample, and the real-time rolling nonlinear model.



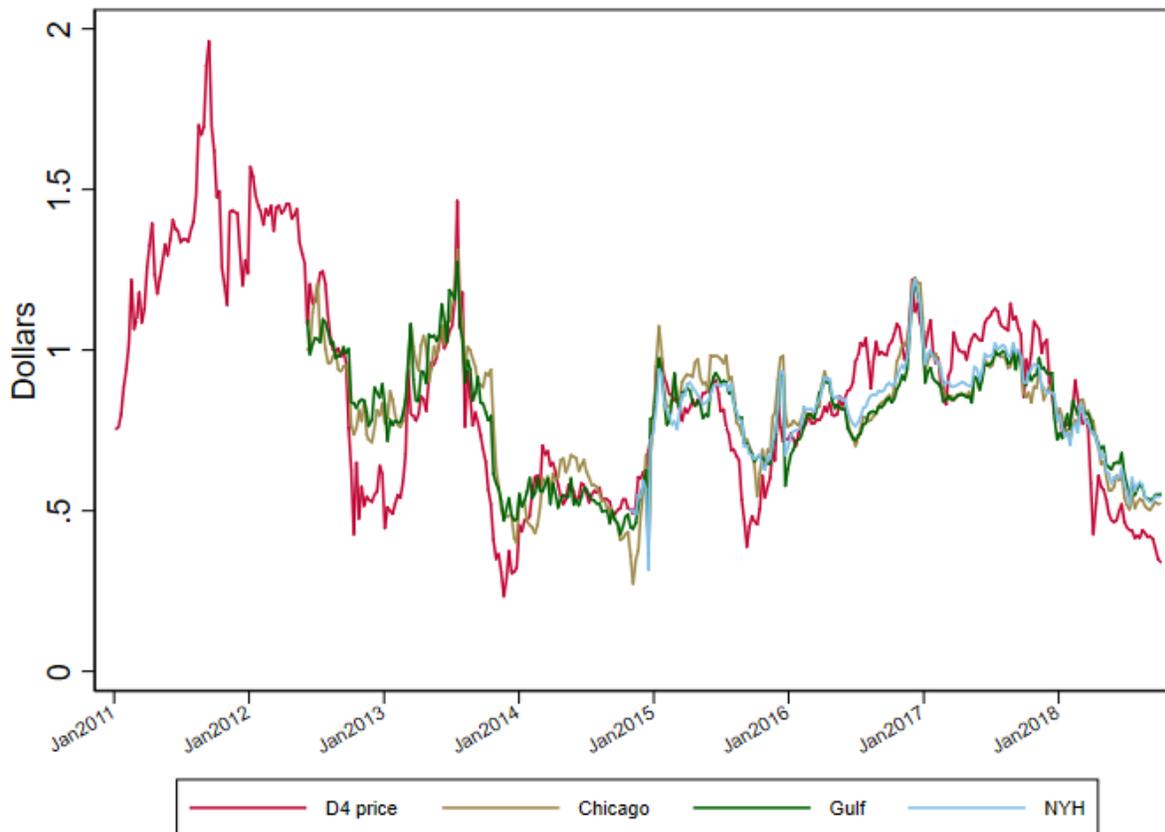
First, at the monthly frequency, the models generally track each other closely.

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, where 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 appear to be inappropriate descriptions of the actual market probabilities of reinstatement of the tax credit during this period.

Finally, Figure 5 plots the real-time predicted values from the nonlinear rolling models for the three spreads separately. Evidently there are high 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.

Figure 5. Weekly D4 RIN price and predicted value based on nonlinear rolling models for Chicago, Gulf, and NYH spreads.



4. Extension to the Spread Depending on the Tax Credit

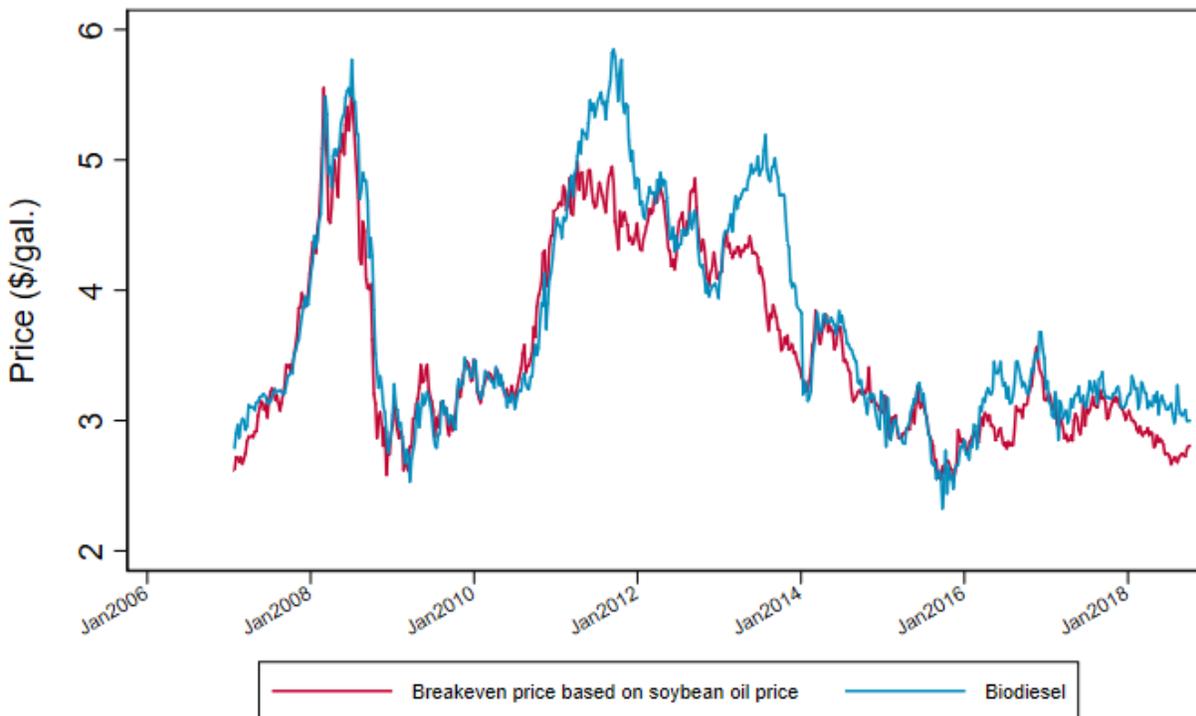
The model laid out in Section 2 assumes that the status of the biodiesel tax credit does not affect the spread. However, the estimates in the final column of Table 1 provide some weak evidence that the level of the spread depends on whether the tax credit is in effect. In this section, we provide two possible explanations for this dependence and then extend the model in Section 2 to allow for the level of the spread to depend on whether or not the tax credit is in place.

One plausible 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 RFS biodiesel 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 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 "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.

The impact of the expiring biodiesel tax credit on biodiesel prices is readily seen with the aid of Figure 6, which plots the biodiesel price versus a simple breakeven relationship between the biodiesel price and the price of the marginal feedstock during this period, soybean oil. This simple model posits that the breakeven 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 (OPIS), and the intercept captures the non-oil variable costs of the plant, estimated to be \$0.60 per gallon. This simple breakeven 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.

Figure 6. Weekly (Friday) Biodiesel Price and Simple Breakeven Price at a Representative Iowa Plant, 01/26/2007 – 10/04/2018



Source: AMS/USDA.

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 sets the renewable volume obligations and percentage standards, and indeed the EPA has at times explicitly taken the tax credit into account in its rulemakings.⁸ If the EPA treats the tax

⁸ For example, in the 2013 rulemaking, 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)

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 feed stock producers, such as soybean farmers.

These two explanations – blenders bidding up the price of biodiesel in advance of the tax credit expiration and 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 regressions (4) and (5) in Table 1, write the spread as,

$$S_t = \mu + \beta B_t + u_t, \text{ where } u_t = u_{t-1} + \varepsilon_t, \quad (7)$$

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

$$P_t^{D4} = e^{-r(T-t)} E_t \max[S_T - (1-\beta)B_T, 0] / 1.5, \quad (8)$$

where S_t follows (7) and B_t continues to follow the Markov process (3).

Although the economics of the pricing formula (8) 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,

$$\begin{aligned} P_t^{D4} &= e^{-r(T-t)} [E_t S_T - (1-\beta)E_t B_T] / 1.5 \\ &= e^{-r(T-t)} \{S_t - (1-2\beta)(1-q)(1-B_t) - [(1-\beta)p + \beta]B_t\} / 1.5, \end{aligned} \quad (9)$$

where the second line of (9) follows by substituting (3), $E_t S_T = \mu + \beta E_t B_T + u_t$ and (7) into the first line of (9) and simplifying.

The key observation is that the terms in brackets in the second line of (9) are the same as in the baseline linear model (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 (8) 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 (9), 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 one 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 while β is identified from (7), in practice it is very imprecisely estimated: formally, because it compares regime means of random walks, 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 (non-rejected) case that all of the tax credit accrues to the biodiesel producer, the Markov probability p is not identified (p drops out if $\beta = 1$ in (9)). Thus, p is weakly identified in this application.

5. Discussion and Conclusions

The most important conclusion from this work, shown in Figures 2, 4, and 5, is that movements in the D4 RIN price at frequencies of a month or longer are well explained by two economic fundamentals: the spread between the biodiesel and ULSD prices and whether the 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 RFS, and the history of Congress intermittently extending, or not, the biodiesel tax credit.

We have laid out three economic channels whereby the tax credit affects RIN prices: 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; an expectational channel in

which the imminent expiration of the tax credit induces buying BBD before the deadline and thus increases the spread; and a regulatory channel in which the EPA sets the 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, however 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.

References

- Irwin, S. (2017). “The Profitability of Biodiesel Production in 2016.” *farmdoc daily* (7):38. Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, March 1, 2017.
- Irwin, S., and D. Good (2017). “How to Think About Biodiesel RINs Prices under Different Policies.” *farmdoc daily* (7):154. Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, August 23, 2017.
- Irwin, S.H. and J.H. Stock (2018). *The Renewable Fuel Standard: The Case for Reform*. manuscript, in process.
- Lade, G.E., and J. Bushnell. (forthcoming). “Fuel Subsidy Pass-Through and Market Structure: Evidence from the Renewable Fuel Standard.” *Journal of the Association of Environmental and Resource Economists*.
- Lade, G.E., C-Y C. 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), 707-731.
- Knittel, C.R., B.S. Meiselman, and J.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), 1081-1119.