WORKING PAPERS

Simulating a Homogeneous Product Merger: A Case Study on Model Fit and Performance

Daniel Greenfield
Nicholas Kreisle
Mark Williams

WORKING PAPER NO. 327

October 2015

BUREAU OF ECONOMICS
FEDERAL TRADE COMMISSION
WASHINGTON, DC 20580

Simulating a Homogeneous Product Merger:

A Case Study on Model Fit and Performance

Daniel Greenfield
Nicholas Kreisle
Mark Williams[1]

October 28, 2015

Abstract

This paper studies Tesoro's 2013 acquisition of British Petroleum's Los Angeles refinery. We present a merger simulation model tailored to the gasoline market, which includes Cournot firms and a price-taking fringe. This hybrid model generates margins that are more plausible than those generated by the standard Cournot model. We also test the predictive accuracy of the models relative to empirical estimates of the acquisition's price effect. We estimate the effect of the acquisition using both difference-in-differences estimation and the synthetic control method. Both methods suggest the acquisition had little if any effect on Los Angeles gasoline prices. We can reject the price effect predicted by the standard Cournot model, but not that of the hybrid model.

[1] The views expressed in this article are those of the authors and do not necessarily reflect those of the Federal Trade Commission.

1. Introduction

This paper studies Tesoro's 2013 acquisition of British Petroleum's (BP's) Los Angeles refinery, a transaction the United States Federal Trade Commission investigated for nine months before concluding the acquisition was not likely to lessen competition substantially.[2] The paper has two primary goals. First, we simulate the price effect of the acquisition using a variant of the Cournot model, which includes a competitive fringe of price-taking firms.[3] This hybrid model generates margins that are more plausible than those generated by the standard Cournot model. Second, we empirically estimate the price effect of the acquisition. We can reject the price effect predicted by a standard Cournot model but not that of the hybrid model.

The Cournot model is a standard framework for analyzing issues of market power in homogeneous goods industries.[4] However, margins in the California gasoline market are likely smaller than market shares and demand elasticity estimates imply in Cournot equilibrium. Specifically, the Cournot model implies marginal costs below the cost of crude oil. We do not observe refinery marginal costs, yet the cost of crude oil is a plausible lower bound, given that it is the primary input for gasoline production. That the standard Cournot model implies marginal cost below the cost of crude oil suggests the model does not fit the California gasoline market.

Introducing a competitive fringe of price-taking firms breaks the strict relationship between overall market shares and margins found in the standard Cournot model. We calibrate the hybrid model restricting equilibrium marginal costs to the range of plausible values, between

---

[2] See FTC press release, *FTC Closes Investigation into Tesoro's Acquisition of BP Refinery*, May 17, 2013. Available at https://www.ftc.gov/news-events/press-releases/2013/05/ftc-closes-investigation-tesoros-acquisition-bp-refinery.

[3] Borenstein and Bushnell (1999) use a similar model to study market power in a wholesale electricity market. Also see Borenstein, Bushnell, and Knittel (1999) and Borenstein, Bushnell, and Stoft (2000).

[4] Surveys on economic theories relevant to antitrust emphasize the importance of Cournot models for homogeneous goods industries. See, for example, Werden and Froeb (2008) and Kaplow and Shapiro (2007).

the cost of crude oil and the price of gasoline. After calibrating both versions of the model, we simulate the effect of the acquisition. The hybrid model predicts a price increase between zero and one percent, depending on assumed pre-merger margins, while the standard Cournot model predicts a price increase of six percent.

We use two methodologies to empirically estimate the causal effect of the transaction and test the models' predictions. First, we estimate a difference-in-differences regression model using data from Los Angeles and 18 control cities. Second, we apply the synthetic control methodology introduced by Abadie and Gardeazabal (2003) and Abadie, Diamond and Hainmueller (2010). Using two methodologies allows us to test the robustness of our conclusions to alternative identifying assumptions and alternative forms of statistical inference.

The difference-in-difference regressions apply equal weight to each city. They assume Los Angeles would have experienced the same conditional mean price change as the 18 control cities but for the acquisition. Alternatively, the synthetic control method constructs a weighted average control city that best approximates pre-acquisition prices and predictors of price in Los Angeles. The synthetic control city's prices then serve as the post-acquisition counterfactual prices for Los Angeles.

Following Abadie, Diamond and Hainmueller (2010), we perform a placebo exercise to evaluate the significance of our synthetic control estimates. We test for a placebo effect in each of the control cities, iteratively re-implementing the synthetic control method for each control city. This generates an empirical distribution of effects. If the effect estimated for Los Angeles is an outlier among the placebo effects, we can be confident that there is evidence of an acquisition effect. The empirical distribution also reveals the probability of finding an effect comparable to that found in Los Angeles if we were to randomly select the treatment (acquisition) city from our data. This form of exact inference contrasts with traditional hypothesis testing, where standard errors reflect sampling uncertainty, under the assumption that the data are a sample from a larger population.

Neither empirical method provides robust evidence of a price change after the acquisition. All of our difference-in-differences estimates are statistically indistinguishable from zero. We can reject the 6% price increase predicted by the Cournot model, but not the 1% price increase predicted by the hybrid model. With the synthetic control method, the placebo exercise indicates that there is a better than 50% chance of finding as large an effect as that found in Los Angeles if we were to randomly select the treatment city from our sample.

Our results contribute to the merger retrospective literature, a growing body of studies that estimate the price effects of mergers. Recent examples on mergers in the petroleum industry include Hosken, Silvia and Taylor (2011), Silvia and Taylor (2013), and Kreisle (2015).[5] Hosken, Olson, and Smith (2012) study horizontal mergers in the U.S. grocery industry and appear to be the first authors to use synthetic control methods for a merger retrospective.[6]

Our study also contributes to the industrial organization literature by evaluating a simulation model tailored to a specific market. Carlton (2010) and Werden (2015) note the importance of evaluating specific analytic tools used during the merger review process. Other papers that evaluate the accuracy of merger simulation models relative to empirical estimates of merger price effects include Peters (2006), Weinberg (2011), and Weinberg and Hosken (2013).[7]

---

[5] Earlier examples include Taylor and Hosken (2007) and Simpson and Taylor (2008).

[6] Also, see Hosken, Olson, and Smith (2015), who use the synthetic control method to estimate price changes following entry and exit by grocery retailers. See Kwoka (2014) for a recent survey of the merger retrospective literature.

[7] Budzinski and Ruhmer (2010) survey the broader merger simulation literature.

The next section of the paper outlines the standard Cournot model and introduces our hybrid variant. Section 3 describes the California gasoline market and refinery production technology. Section 4 applies both models to industry data and demonstrates the superior *ex ante* fit of the hybrid model. Section 5 simulates the Tesoro/BP transaction under our baseline assumptions and section 6 explores how the simulations change with alternative parameterizations. Section 7 presents the empirical analysis, which we use to assess the *ex post* accuracy of the simulation models. Finally, Section 8 concludes.

2. The Models

Consider a homogeneous goods market with inverse demand $P(Q)$ and N Cournot firms simultaneously choosing output. We can write firm i's first order condition, $P(Q) + q_i P'(Q) - C_i{}'(q_i) = 0$, in Lerner index form:

$$\frac{P - C_i{}'}{P} = \frac{s_i}{\varepsilon}, \tag{1}$$

where $Q = \sum_i q_i$ is total output, C_i' is the marginal cost of firm i at equilibrium, s_i is firm i's share and ε is the absolute value of the elasticity of market demand. In equilibrium, margins are proportional to market shares and inversely proportional to ε.

Less elastic market demand implies larger margins and larger differences in marginal costs between firms with different market shares. Firms with larger market shares operate at lower equilibrium values of marginal costs because marginal revenue is decreasing in q_i. This equilibrium relationship is true regardless of the functional form of demand or costs.

Now consider the same homogeneous goods market, yet with two sets of firms: F price-taking firms and S strategic firms ($F + S = N$). Price-taking firms choose output where price equals marginal cost. We assume strategic firms choose output in a Cournot fashion under the belief that price-takers expand production until their marginal costs equal the market price.

The mechanics of the hybrid model are analogous to those of a dominant firm model. The residual demand facing the strategic firms is the difference between market demand and the supply from price-takers, $Q^r(P) = Q(P) - \sum_{i \in F} q_i\,(P)$. The following identity defines the price elasticity of residual demand:

$$\varepsilon^S s^S = \varepsilon + s^F \eta^F \tag{2}$$

where ε is the absolute value of the market demand elasticity, ε^S is the absolute value of the residual demand elasticity facing the combined strategic firms, s^S is the combined market share of the strategic firms, s^F is the combined share of price-taking firms and η^F is the elasticity of price-taking firms' supply.[8]

Just as in the standard Cournot model, each strategic firm's first order condition implies an equilibrium relationship between market shares and margins:

$$\frac{P - C_i{}'}{P} = \frac{s_i^S}{\varepsilon^S} \tag{3}$$

where s_i^S is the share of firm i within the group of all strategic firms. This implies that each strategic firm's price cost margin remains proportional to its share among the strategic firms. Because marginal cost curves are upward sloping, price takers can earn positive variable or total profits, but zero incremental profits.

[8] This identity holds when the price-taking firms have increasing marginal cost curves. Constant marginal costs would imply perfectly elastic supply and residual demand.

Assignment of firms to the strategic or price-taking groups is a challenge that dates at least as far back as Stigler's (1947) investigation of oligopoly pricing.[9] However, this assignment is in principle no more challenging than deciding when the textbook dominant firm model may apply. We initially calibrate the hybrid model assuming all firms with a market share greater than 10% are strategic and explore how the simulation results change as we alter the strategic group.[10]

3. The California Gasoline Market

Gasoline sold in California must meet emissions standards set by the California Air Resources Board (CARB). Thus, refiners intending to produce gasoline for the California market produce California Reformulated Blendstock for Oxygenate Blending (CARBOB). The additional processing required to meet California specifications tends to make CARBOB more costly to produce than blendstocks used in other parts of the country. Distribution terminals typically combine nine parts CARBOB and one part ethanol to make finished California gasoline.[11]

CARBOB is a Homogeneous product with a distinct wholesale market given that the specification is unique to California. According to the California Energy Commission (CEC), in 2012 nine firms produced CARBOB in fourteen California refineries and sales of finished CARB gasoline totaled 14 billion gallons.[12] We calculate market shares based on crude oil

---

[9] While Stigler admits his criteria for inclusion in the oligopolistic group are "an unsatisfactory substitute" he notes that they "focus attention on the relevant variables," namely, the combined share of the oligopolistic (here, "strategic") group and the share of the largest excluded firm relative to the market leader.

[10] Tesoro and BP are the second and third largest firms, and are always strategic firms in our analysis. A merger between two price-taking firms (not becoming strategic as a result of a merger) would have no impact on the market price.

[11] The market for ethanol is unconcentrated as discussed in Federal Trade Commission (2014), which summarizes ethanol capacity and concentration since 1998.

[12] See http://energyalmanac.ca.gov/petroleum/refineries.html

distillation capacities, which, according to the CEC, are roughly proportional to gasoline production shares. Table 1 lists the refineries and these crude distillation capacities.

TABLE 1—Crude Distillation Capacity of CARBOB-producing California Refineries, 2012

Refinery	Barrels per Day	Market Share
Chevron, U.S.A., Inc.		27.3%
Richmond Refinery	276,000	
El Segundo Refinery	245,271	
Tesoro Refining & Marketing Company		14.2%
Golden Eagle Martinez/Avon Refinery	166,000	
Wilmington Refinery	103,800	
BP West Coast Producs LLC		12.6%
Carson Refinery	240,000	
ConocoPhillips		11.4%
Wilmington Refinery	139,000	
Rodeo San Francisco Refinery	78,400	
Valero		11.0%
Benicia Refinery	132,000	
Wilmington Refinery	78,000	
Shell Oil Products US		8.2%
Martinez Refinery	156,400	
ExxonMobil Refining & Supply Company		7.8%
Torrance Refinery	149,500	
ALON USA		6.1%
Bakersfield Refinery	66,000	
Paramount Refinery	50,000	
Kern Oil & Refining Company		1.4%
Bakersfield Refinery	26,000	

Source: CEC

To calibrate the models to the pre-acquisition data, we use a market price of $127 per barrel, the average Los Angeles spot price of CARBOB in 2012.[13] We assume a crude oil cost

[13] Source: http://www.eia.gov/dnav/pet/pet_pri_spt_s1_d.htm.

of $105 per barrel, the average West Coast crude acquisition cost in 2012.[14] Lastly, consistent with prior literature we assume a demand elasticity of -0.3.[15]

Refineries produce multiple products, including gasoline of different grades, diesel fuel, jet fuel, heating oil, and other heavier products.[16] We focus on gasoline because it is generally the highest-valued product, and because it represents the plurality if not majority of production at California refineries.[17]

A typical refinery has thousands of production constraints, including the capacity of individual processing units, available tankage, and blending requirements of inputs and outputs. For a given stock of capital, the marginal cost of producing gasoline tends to increase as output expands and additional production constraints bind. Moreover, for a given amount of crude oil and other feedstocks running through the refinery, marginal cost increases as the refinery can only increase gasoline production at the expense of reducing the production of other products.

Oil refiners use linear programming (LP) models to choose a profit-maximizing vector of inputs and outputs, subject to the model's constraints and expected market prices (e.g. the prices of crude oils, intermediates, and finished products).[18] That refinery models maximize profits subject to fixed market prices suggests price-taking behavior, at least if refineries strictly follow the production plans generated by these LP models (using unbiased estimates of future expected pricing). What is more, investigations of refinery pricing have found evidence that refineries

[14] Source: http://www.eia.gov/dnav/pet/pet_pri_rac2_dcu_nus_m.htm.
[15] We believe -0.3 is a reasonable estimate of the price elasticity of demand over a yearly period, based on research including Lin and Prince (2013), Levin, Lewis, and Wolak (2015) and Coglianese, Davis, Kilian, and Stock (2015).
[16] As this paper uses crude distillation capacity as its measure of market share of gasoline production, if we were instead to model the acquisition's effect on diesel prices (also using distillation capacity as the relevant measure of diesel production), we would get the same theoretical results, subject to a different choice for the demand elasticity for diesel fuel.
[17] See http://energyalmanac.ca.gov/petroleum/refineries.html.
[18] See Federal Trade Commission (2006), Leffler (2008), and Gary, Handwerk and Kaiser (2010).

generally take prices as given. For example, a recent FTC investigation of gasoline pricing found that refiners typically "assume their short-run operational decisions do not affect market prices."[19]

By contrast, Borenstein, Bushnell and Lewis (2004) argue that despite evidence consistent with price-taking behavior, it would be surprising if producers did not consider the effect of their own production levels on California gasoline market prices. Our hybrid model allows for both price-taking and strategic behavior to exist simultaneously. This additional degree of freedom allows the hybrid model to better match plausible industry margins, given market shares and demand elasticity estimates.

4. Fitting the Models to the California Gasoline Market

In the standard Cournot model, a firm's margin is a function of its market share and the price elasticity of demand. Hence, we can calculate equilibrium marginal costs with information on market shares, price, and the price elasticity of demand. Again, we use crude distillation capacity shares as a proxy for gasoline output shares. Table 2 shows the quantity, implied margins, and implied marginal costs for each firm in California assuming Cournot equilibrium. The standard Cournot assumption implies equilibrium marginal costs that are lower than the cost of crude oil for all firms except Kern Oil.

[19] See Federal Trade Commission. (2006). The report does note that "although refiners state that they generally adhere to the LP model, some refiners occasionally deviate from the model if their internal analyses and judgment indicate that such a deviation is necessary to maximize refinery profitability." A refinery planner may run the LP model with the quantity of gasoline constrained based on feedback from marketing and trading personnel on the price sensitivity of demand. If a refinery planner introduces such a "demand" or "marketing" constraint into the model, the LP model output will show a positive incremental margin on the constrained product. Positive incremental margins can also arise due to capacity constraints on refining units, pipelines, and storage facilities.

TABLE 2—Implied Margins in the Standard Cournot Model

	Quantity*	C'	$(P - C')/P$
Chevron	27.34	11.28	0.91
Tesoro	14.15	67.31	0.47
BP	12.59	73.95	0.42
ConocoPhillips	11.40	78.99	0.38
Valero	11.02	80.64	0.37
Shell	8.20	92.58	0.27
Exxon	7.84	94.12	0.26
ALON	6.08	101.58	0.20
Kern Oil & Refining	1.36	121.63	0.05
Total	**100.00**		

*Total market output normalized to 100. We use crude distillation capacity shares as a proxy for gasoline output shares.

Surely, models must abstract from some industry details, but such a glaring inconsistency with basic industry facts suggests that the Cournot model is not appropriate for the California gasoline market.[20] It is important to emphasize the poor fit of the Cournot model does not depend on the functional form of demand or costs. The same relationship between equilibrium marginal costs and market shares exists whether marginal costs are constant or increase rapidly beyond some level of output.

Fitting the standard Cournot model to pre-merger data required solving for the unknown values of equilibrium marginal costs. The hybrid model, however, has an additional unknown, the elasticity of residual demand (defined by the elasticity of supply from the price-taking firms). The extra unknown necessitates additional information or assumptions.

[20] Werden, Froeb, and Scheffman (2003) and Werden and Froeb (2008) discuss the importance of model fit in policy settings. In two non-merger antitrust cases where economic experts used a Cournot model, judges excluded expert testimony due to poor model fit with industry facts. In *Heary Brothers*, for example, the judge noted that profit margins were similar for firms with vastly different market shares, a fact inconsistent with Cournot equilibrium. See, *Concord Boat Corp. v. Brunswick Corp.*, 207 F.3d 1039 (8th Cir. 2000) and *Heary Bros. Lightning Prot. Co. v. LIGHTNING PROT. INSTITUTE*, 287 F. Supp. 2d 1038 (D. Ariz. 2003). Also, see Coate and Fischer (2000).

For this case study, we take advantage of the fact that the price of crude oil is a reasonable lower bound for the marginal cost of producing gasoline. In our baseline calibration, we assume that the largest strategic firm (Chevron) has equilibrium marginal costs equal to the cost of crude oil, the lower bound of the plausible range for marginal costs. In section 6, we relax this assumption and show how the predicted merger effect varies with assumed marginal costs, a relationship illustrated in Figure 1.

Table 3 shows the baseline calibration of the hybrid model, assuming all firms with a market share below 10% are price-takers and the largest firm's marginal cost is equal to the cost of crude oil. Here, the strategic firms are BP, Chevron, ConocoPhillips, Tesoro, and Valero.[21] By design, every firm's marginal cost falls between the cost of crude oil and the market price of gasoline. The implied residual demand elasticity is 2.0 and the implied supply elasticity from the price-taking firms is 5.2.[22]

TABLE 3—Implied Margins in the Hybrid Model

	Quantity*	C'	$(P - C')/P$
Chevron	27.34	104.69	0.18
Tesoro	14.15	115.66	0.09
BP	12.59	116.96	0.08
ConocoPhillips	11.40	117.94	0.07
Valero	11.02	118.27	0.07
Shell	8.20	127.43	0.00
Exxon	7.84	127.43	0.00
ALON	6.08	127.43	0.00
Kern Oil & Refining	1.36	127.43	0.00
Total	**100.00**		

*Total market output normalized to 100. We use crude distillation capacity shares as a proxy for gasoline output shares.

[21] Table 6 examines the effect of changing the composition of the strategic group.

[22] Firms can change production levels of all products, substitute CARBOB production for other grades of gasoline, substitute gasoline production for other types of refined products (e.g. changing the gas/diesel split), or import CARBOB from non-California sources.

5. Simulating Tesoro's Acquisition of BP Carson

Tesoro acquired BP's Southern California refining, marketing and logistics business on June 1, 2013. The assets included BP's Carson refinery, marine terminals, land terminals, and pipelines, all of which were adjacent to Tesoro's Wilmington refinery. Tesoro announced plans to gradually integrate the two refineries over a number of years and undertake significant investments in pipeline interconnections and refinery capacity improvements in 2015.[23] Our simulations examine the potential for a price increase due to the Tesoro/BP combination but do not consider possible efficiencies related to refinery integration.

To simulate the effects of the acquisition, we assume both market demand and firm marginal costs are linear. Inverse demand is $P = a - bQ$, where total market output is $Q = \sum_{i=1}^{N} q_i$. We consider a cost function that has been widely used for merger analysis.[24] Firm i's total cost function is $q_i^2/2k_i$, with marginal costs, $MC_i = q_i/k_i$.[25] The reciprocal of the slope term, k_i, can be viewed as representative of each firm's capital stock. A merger between firms 1 and 2 creates a new firm with $k_{12} = k_1 + k_2$. The resulting marginal cost function is thus the horizontal summation of the pre-merger firms' marginal cost curves.

With these demand and cost functions, each Cournot firm maximizes profits with respect to their own quantity, yielding the following first order conditions: $a - bq_i - bQ = q_i/k_i$. Solving the n first order conditions results in equilibrium quantities:

[23] See, Thomson StreetEvents (2012).

[24] See, for example. Perry and Porter (1985), Farrell and Shapiro (1990), McAfee and Williams (1992), and Werden and Froeb (2008).

[25] To be sure, the quadratic function is stylized approach to modeling costs. Assuming constant marginal costs, an even more stylized approach, would suggest the merged firm simply transfers all production to the lower marginal cost firm.

$$q_i = \frac{a}{b}\left(\frac{\frac{bk_i}{bk_i + 1}}{1 + \sum_{i=1}^{n} \frac{bk_i}{bk_i + 1}}\right).^{26} \tag{4}$$

It is straightforward to solve for the parameters of the demand function based on the market price, quantity, and assumed elasticity of demand. Simulating a merger occurs in two steps. First, we recover the k_i using the first order conditions. Second, we calculate the post-merger equilibrium q_i given the post-merger values of k_i.

TABLE 4—Simulation Results from the Standard Cournot Model

	Quantity*	C'	$(P - C')/P$
Chevron	29.19	12.04	0.91
Tesoro	**19.78**	**51.99**	**0.62**
ConocoPhillips	12.17	84.31	0.38
Valero	11.76	86.07	0.37
Shell	8.76	98.82	0.27
Exxon	8.37	100.46	0.26
ALON	6.50	108.43	0.20
Kern Oil & Refining	1.46	129.83	0.05
Total	**97.98**		

*Pre-acquisition total market output normalized to 100.

Table 4 shows the results of the Cournot simulation. Tesoro's predicted post-merger output is 26% less than the combined output of BP and Tesoro before the acquisition. Equilibrium marginal costs are lower as the merged entity restricts output along its new marginal cost function. Conversely, all of the non-merging firms increase output in response, with their marginal costs increasing as they expand output. Overall, industry output decreases by 2%. Based on the demand parameters, the Cournot model predicts the merger will increase prices by 6%.

[26] McAfee and Williams (1992) use the same notation. They derive this equilibrium and study the welfare implications of horizontal mergers.

To simulate the merger with the hybrid model, we assume a linear supply function from the price-taking firms, with parameters defined by the pre-merger equilibrium price, quantity, and the same supply elasticity for each price-taking firm. From this, one can algebraically derive the residual demand function facing the strategic firms. The simulation proceeds as under the standard Cournot model above, with the exception that we replace the overall market demand function with the residual demand function, and replace the overall market shares with the share of each strategic firm within the strategic group. Table 5 presents the results.

TABLE 5—Simulation Results from the Hybrid Model

	Quantity*	C'	$(P - C')/P$
Chevron	27.60	105.67	0.18
Tesoro	**24.83**	**107.97**	**0.16**
ConocoPhillips	11.51	119.05	0.07
Valero	11.12	119.38	0.07
Shell	8.59	128.63	0.00
Exxon	8.21	128.63	0.00
ALON	6.37	128.63	0.00
Kern Oil & Refining	1.43	128.63	0.00
Total	**99.66**		

*Pre-acquisition total market output normalized to 100.

Tesoro reduces output less than predicted with the standard Cournot merger simulation: 7% rather than 26%. This difference reflects the more aggressive output expansion of the price-taking firms. In the hybrid simulation, strategic firms outside the merger increase output by 1%, while the non-strategic firms increase output by 2%. Overall, output falls by 0.3% and the hybrid model predicts that price increases by approximately 1%.

6. Alternative Parameterizations

In Table 6, we derive simulated price effects as we change the composition of the strategic group. The first column assumes the strategic group includes only Chevron, BP, and

Tesoro. Moving to the right, the strategic group includes additional firms, added in descending market share order. The middle column, which includes five strategic firms, is the baseline case presented in section 5. Notably, the order-of-magnitude of the simulated price effect does not change with the composition of the strategic group.

TABLE 6—Price Effect with Different Strategic Groups

Strategic Firms	Chevron, Tesoro, BP	+ Cononco	+Valero	+Shell	+Exxon
Price Effect (% change)	1.04%	0.99%	0.94%	0.91%	0.88%
Supply Elasticity of Price-Takers	2.68	3.57	5.25	8.06	16.54

One might expect that the price effect would increase as additional firms join the strategic group and the hybrid model ostensibly looks more like a standard Cournot model. However, note that the margin of the largest strategic firm remains constant as we move to the right in the table.[27] In order to maintain the level of margins, the elasticity of supply from the price-taking firms must increase, as defined by equations (2) and (3). Intuitively, as the share of the competitive fringe shrinks, the remaining price-taking firms must possess greater supply elasticity to maintain the assumed margin of the largest strategic firm.[28]

In our baseline results, we calibrate the residual demand elasticity by assuming the largest firm's equilibrium marginal cost equals the cost of crude oil (the bottom of the plausible range of values for marginal costs), with smaller strategic firms' having higher marginal costs, as implied by equation (3). Crude oil is the primary input for gasoline and a given volume of crude oil will produce an approximately equal volume of refined products. Thus, the cost of crude oil is a

[27] Residual demand elasticity declines as you move right in the table, ranging from 2.8 to 1.7. The decline reflects the fact that the largest firm's share of the strategic group decreases while holding its margin constant.
[28] A researcher with access to firm-level production data could validate the implied supply elasticity of the price-taking fringe, and perhaps pin down the most realistic grouping. Qualitative evidence about firms' business strategies could also guide assignment to the strategic and price-taking groups.

proxy for the minimum marginal cost of producing gasoline. Relevant marginal costs may also include additional purchases of materials, labor, and capital, or the opportunity cost of forgone output of other products. Hence, equilibrium marginal cost likely falls somewhere between the cost of crude oil and the price of gasoline.

Figure 1 illustrates the relationship between the largest firm's pre-acquisition equilibrium marginal cost and the predicted price effect from the Tesoro/BP acquisition. Smaller assumed margins translate to more elastic residual demand, greater supply elasticity from price-taking firms, and smaller predicted price effects. The dashed vertical line, indicating where marginal cost equals the cost of crude oil, intersects the curve at our baseline predicted price effect, 0.94%. The price effect declines to zero as assumed marginal cost increases from a neighborhood of crude oil costs to that of the market price ($127). Smaller margins imply greater supply elasticity from the price-taking fringe, which limits the market power of strategic firms. To summarize, the predicted acquisition effect from the hybrid model depends critically on assumed margins. However, nowhere in this plausible range of marginal costs does the hybrid model's predicted merger effect approach the 6% level implied by the standard Cournot model.

FIGURE 1—Relationship between Assumed Margins and the Predicted Acquisition Effect

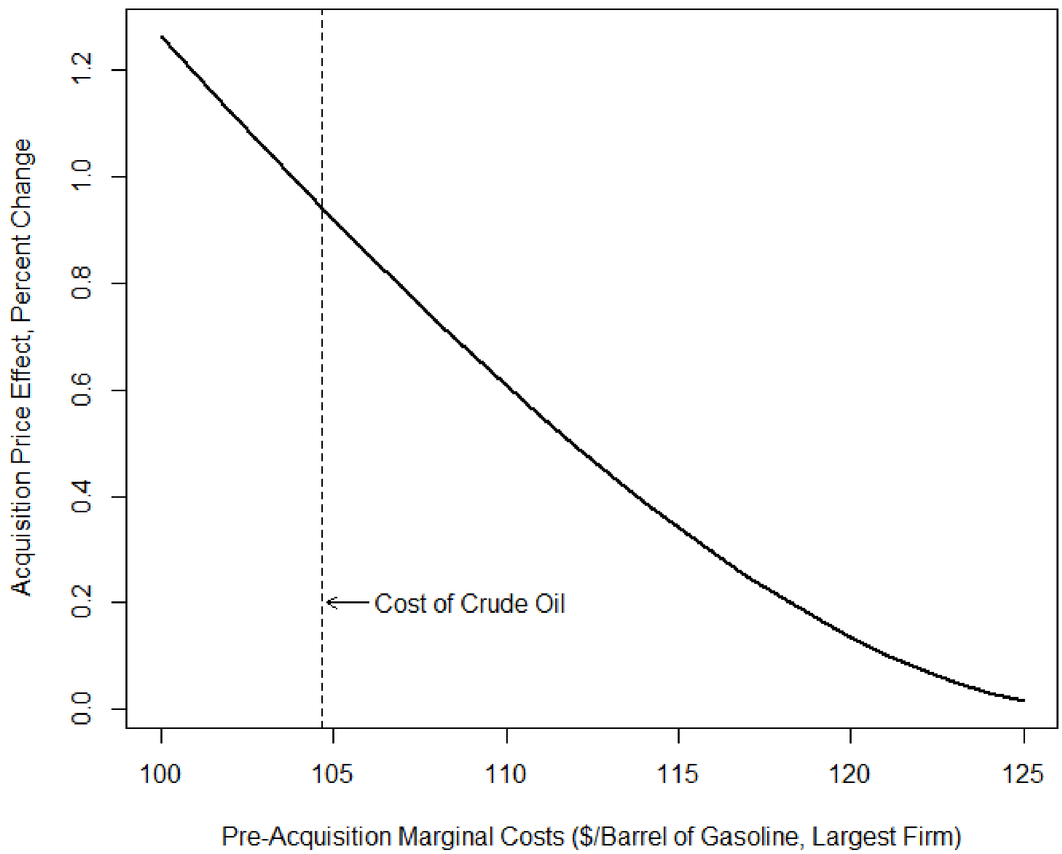

The choice of market demand elasticity does not affect the size of the merger's price effect in the hybrid model, given our strategy of identifying the model's parameters by restricting the largest firm's marginal cost to a plausible range. As long as the group of strategic firms is unchanged, equation (3) sets the relationship between the residual demand elasticity, market shares, and margins. In other words, assumed margins imply the residual demand elasticity facing the strategic firms. Changing the overall market demand elasticity will change the supply elasticity of the price-taking firms (and therefore their quantity response), but will not change the predicted price effect of the acquisition.

Changing the choice of demand elasticity will change the predicted price effect in the standard Cournot model. Implied margins would also change. Indeed, it is possible to obtain implied marginal costs for the largest firm close to the cost of crude oil if one is willing to assume a demand elasticity of -1.3. In that case, the predicted price effect would be 1.2%. Hence, in some sense, the practical difference between the hybrid and standard Cournot models rests on the assumed elasticity of demand. However, we believe a market demand elasticity of -1.3 is implausible, and other researchers share this view. Coglianese, Davis, Kilian, and Stock (2015) describe a price elasticity of -0.46 as "the upper end of the range of elasticity values that seem economically plausible and indeed is higher than many economists would be comfortable with."

7. Estimating Merger Price Effects

The standard Cournot and the hybrid model predict respective price effects of 6% and 1% using our baseline calibrations. To assess the accuracy of these predictions we estimate the actual effect of the acquisition on Los Angeles wholesale gasoline prices. The standard approach in the merger retrospective literature is to use a difference-in-differences model, comparing the evolution of prices in markets affected by the transaction with the evolution of prices in markets not affected. Lacking an ideal control for Los Angeles prices, we supplement the difference-in-differences approach with the synthetic control method.

We start by comparing the average pre/post-acquisition price difference in Los Angeles with the same difference in all 18 cities for which we have data. This difference-in-differences regression analysis allows us to conduct traditional statistical inference and test the null hypotheses from the simulation models. In this approach, we assume the 18 controls cities are a sample drawn from a larger population, and that the conditional mean price change in that population is the same change that would have occurred in Los Angeles but for the acquisition.[29]

The primary drawback of this approach is that Los Angeles prices might have followed a different path than the average price in the control cities if the acquisition had not occurred. Unobserved variables that differentially affect Los Angeles and one or more control cities will lead to correlation between the error term and treatment dummy. We could select—based on our best judgment—a particular control city that we believe experienced the same time-varying shocks. Yet this approach may raise concerns about the objectivity of our choice of control city and the degree to which its prices are a credible proxy for Los Angeles's counterfactual prices.

The synthetic control methodology, which selects a control group using a data-driven procedure, reduces the ability of the researcher to influence results and makes the selection process more objective. After finding the optimal synthetic control city and estimating the effect of the transaction on Los Angeles prices, we conduct a placebo experiment that applies the synthetic control method to each of the control cities. We can then assess whether the estimated effect in Los Angeles is large relative to the distribution of placebo effects.

Section 7.1 describes our data set. Section 7.2 presents the difference-in-differences regression analysis. Section 7.3 describes the synthetic control methodology. Section 7.4 presents the synthetic control results.

[29] Assuming prices in Los Angeles are precisely measured and the price effect in Los Angeles is the relevant statistic, the standard errors reflect sampling uncertainty arising because we do not observe every control city in the population. Abadie, Athey, Imbens, and Wooldridge (2014) discuss how traditional standard errors can also reflect the fact that, even if the researcher observes outcomes for all subjects in the population, there are additional potential outcomes for each subject with different levels of treatment.

7.1. Data and Sample

We construct a panel dataset with 19 cities and 48 monthly observations from January 2011 to December of 2014. The variable of interest is the wholesale price of gasoline. We construct monthly averages using the daily average rack price of branded and unbranded gasoline published by Oil Price Information Service (OPIS). Gasoline retailers and distributors can purchase gasoline at distribution terminals referred to as racks.[30]

We include data on Metropolitan Statistical Area annual per capita personal income, from the U.S. Bureau of Economic Analysis, to capture demand. That these data are annual is not problematic as both the synthetic control method and the regression analysis use pre-acquisition and post-acquisition averages.[31] To control for costs, we use Refiner Acquisition Cost of Crude Oil compiled by the U.S. Energy Information Administration (EIA). EIA calculates these data for five Petroleum Administration for Defense Districts (PADDs) and these data are unavailable on a finer geographic basis.

Table 7 shows the average crude oil acquisition cost, per-capita income, pre-acquisition price, and post-acquisition price for each city in our sample. The final column shows the percent change in average price between the pre- and post-acquisition periods. This percent change is approximately equal for Los Angeles and the average of all the control cities. By definition, the point estimates in the first specification of our regression analysis, which includes no controls for demand or cost, reproduce the 8.1% decrease in prices for the average control city and a difference-in-differences estimate that rounds to 0.0%.

[30] Seven of the potential control cities have reformulated gasoline. In those cities, we use the average wholesale price of reformulated gasoline. We use the average wholesale price of conventional gasoline in the remaining cities where reformulated gasoline is not sold.

[31] The 2014 personal income data are not yet available from the BEA. We set the 2014 annual per capita personal income equal to the 2013 value in each metro area. Note that the synthetic control results, and the regressions specifications presented in columns 1 and 3 of table 8, do not depend on 2014 income data.

Table 7—Control Variables and Pre- and Post-acquisition Prices, Averages by City

	Crude Cost ($/barrel)	Per-Capita Income	Pre-Acquisition Price (cpg)	Post-Acquisition Price (cpg)	Percent Price Change
Los Angeles	**102.2**	**$47,373**	**302.4**	**278.8**	**-8.1%**
Atlanta	106.7	$40,795	287.0	268.2	-6.8%
Boston	106.7	$60,824	296.9	275.1	-7.6%
Chicago	90.8	$48,217	297.1	273.0	-8.4%
Cleveland	90.8	$44,849	294.6	275.1	-6.8%
Dallas	101.4	$46,207	291.0	264.2	-9.7%
Denver	84.8	$51,055	287.6	270.2	-6.2%
Detroit	101.4	$42,272	293.2	269.4	-8.5%
Fairfax	106.7	$61,367	292.0	270.3	-7.7%
Houston	101.4	$50,893	288.8	261.1	-10.1%
Louisville	90.8	$40,863	299.8	279.6	-7.0%
Miami	106.7	$44,743	288.8	266.3	-8.1%
Minneapolis	90.8	$50,502	289.2	261.3	-10.2%
New Orleans	101.4	$44,174	283.2	257.3	-9.6%
Newark	106.7	$58,575	294.5	271.5	-8.1%
Phoenix	102.2	$38,292	303.0	278.4	-8.5%
Salt Lake City	84.8	$40,684	288.2	275.3	-4.6%
Seattle	102.2	$54,080	296.3	270.0	-9.3%
St. Louis	90.8	$45,056	295.1	270.0	-8.9%
Average Control City	**98.2**	**$47,969**	**292.6**	**269.8**	**-8.1%**

7.2. Difference-in-Differences Regressions

We estimate the regression models using time-differenced data that capture the change in each variable from the pre-merger period to the post-merger period. The regression equation takes the form:

$$\ln\left(\frac{P_i^{post}}{P_i^{pre}}\right) = \alpha + \beta ACQUISITION_i + \gamma \ln\left(\frac{INCOME_i^{post}}{INCOME_i^{pre}}\right) + \delta \ln\left(\frac{COST_i^{post}}{COST_i^{pre}}\right) + \varepsilon_i,$$

where P_i^{post} and P_i^{pre} respectively are the post- and pre-acquisition average prices for city i. The variables $INCOME_i$ and $COST_i$ represent per-capita income and crude oil acquisition costs. Finally, $ACQUISITION_i$ is an indicator variable for Los Angeles. The parameter β is the coefficient of interest: the log-point change in prices attributed to the merger.

The goal is to capture the price effect from Tesoro's acquisition. We need a window of time long enough for firms to adjust output in response to the post-merger allocation of capital, yet short enough to avoid contaminating effects from unrelated changes in the industry. The possibility that firms change their behavior before the actual transaction further complicates matters. For example, Kim and Singal (1993) find that merging airlines increase prices after the firms reach an agreement but before they actually consummate their deal.

The resulting challenge of determining the appropriate window of analysis has been widely acknowledged in the merger retrospective literature.[32] We proceed by specifying three different pre/post periods. The first window of analysis utilizes our full dataset, spanning January 2011 through December 2014, with June 2013 as the first post-acquisition period. The second window uses a 12-month period immediately preceding the acquisition and a 12-month period immediately following the acquisition. The third window has a 12-month period that begins 18 months before the acquisition and a 12-month period that begins 6 months after the acquisition. Using a 6-month transition window on either side of the acquisition should minimize the influence of any anticipatory or delayed acquisition effects.

Averaging the data across the various pre- and post-acquisition periods obviates the need to account for serial correlation. Bertrand, Duflo and Mullainathan (2004) demonstrate that serial correlation can lead to significant biases in OLS standard errors and examine the efficacy of several solutions. They find that collapsing the time series is the most reliable solution when the number of cross sectional subjects is relatively small.[33] Other merger retrospectives that rely

[32] See Ashenfelter, Hosken, and Weinberg (2009) and Greenfield (2014) for overviews.
[33] Donald and Lang (2007) propose an analogous two-step regression procedure.

on time-aggregated data include Tenn (2011), Tenn and Yun (2011), Thompson (2011), and Haas Wilson and Garmon (2011).[34]

TABLE 8—Difference-in-Differences Regression Models

Variables	1	2	3	4
Full Sample (Jan '11-May '13 pre period and Jun '13-Dec '14 post period)				
ACQUISITION	0.000	-0.002	0.001	0.003
	(0.015)	(0.016)	(0.015)	(0.017)
Δ ln(*INCOME*)		0.115		-0.2
		(0.475)		(0.539)
Δ ln(*COST*)			0.106	0.125
			(0.089)	(0.106)
Constant	-0.081***	-0.0841***	-0.0755***	-0.0693***
	(0.003)	(0.013)	(0.006)	(0.018)
R-squared	0.000	0.004	0.081	0.089
12-month pre and post periods, no transition window				
ACQUISITION	-0.002	-0.003	-0.002	-0.004
	(0.011)	(0.012)	(0.012)	(0.012)
Δ ln(*INCOME*)		0.455		0.479
		(0.775)		(0.845)
Δ ln(*COST*)			0.007	-0.006
			(0.064)	(0.069)
Constant	-0.037***	-0.039***	-0.037***	-0.039***
	(0.003)	(0.005)	(0.003)	(0.005)
R-squared	0.002	0.023	0.003	0.024
12-month pre and post periods, 12-month transition window				
ACQUISITION	-0.008	-0.01	-0.007	-0.008
	(0.014)	(0.015)	(0.014)	(0.015)
Δ ln(*INCOME*)		0.397		0.167
		(0.545)		(0.590)
Δ ln(*COST*)			0.099	0.089
			(0.079)	(0.088)
Constant	-0.079***	-0.084***	-0.073***	-0.076***
	(0.003)	(0.008)	(0.006)	(0.012)
R-squared	0.018	0.049	0.106	0.111
N	19	19	19	19

*** p<0.01, ** p<0.05, * p<0.1

[34] Our time-differenced model is equivalent to a two-period city fixed effects model.

Table 8 presents the regression results. There are four different specifications, with different combinations of the control variables. The same qualitative conclusion arises from every specification. The change in Los Angeles prices is statistically insignificant. We cannot reject the 1% increase predicted by the hybrid model, yet can reject the 6% increase predicted by Cournot model, in all specifications.

7.3. Synthetic Control Method

The synthetic control method constructs a weighted average of the 18 available control cities. It selects weights so that the synthetic control city best approximates pre-acquisition characteristics of the Los Angeles market. The synthetic control city's post-acquisition prices then form the estimate of Los Angeles's post-acquisition prices but for the merger.

Let $j = 1, 2, \ldots, J + 1$ index cities and $t = 1, 2, \ldots, T$ index time. There are J potential control cities (Los Angeles is $j = 1$) and T_o pre–acquisition periods. A synthetic control city is defined by a ($J \times 1$) vector of weights $\mathbf{W} = \left(w_2, \ldots, w_{J+1}\right)'$, which are non-negative and sum to one. Each element, w_j, represents the weight assigned to potential control city j. $\mathbf{X_1}$ is a $(k \times 1)$ vector of pre-merger characteristics in Los Angeles. In our study, $\mathbf{X_1}$ is a 3 x 1 vector including the pre-acquisition average of per-capita income, the pre-acquisition average of refinery crude acquisition costs, and the pre-acquisition average of wholesale prices. The $(k \times J)$ matrix $\mathbf{X_0}$ contains the same variables as $\mathbf{X_1}$ for the potential control cities.

The optimal set of weights, $\mathbf{W}^*$, is the vector that minimizes the distance

$$\|\mathbf{X_1} - \mathbf{X_0}\mathbf{W}\|_{\mathbf{V}} = \sqrt{(\mathbf{X_1} - \mathbf{X_0}\mathbf{W})'\mathbf{V}(\mathbf{X_1} - \mathbf{X_0}\mathbf{W})},$$

subject to $w_1 + \cdots + w_J = 1$ and $w_J \geq 0$. Hence, $\mathbf{W}^*$ defines the synthetic control city that most closely resembles Los Angeles as measured by the $\mathbf{V}$-weighted distance between the vectors $\mathbf{X_1}$

and $\mathbf{X_0W}$. $\mathbf{V}$ is a diagonal matrix with elements that apply weights to the variables in $\mathbf{X_1}$ and $\mathbf{X_0}$, determining their importance as relevant characteristics.

We choose $\mathbf{V}^*$ by minimizing the mean square prediction error (MSPE) of the synthetic control city relative to Los Angeles during the pre-acquisition period. That is, we choose $\mathbf{V}^*$ such that it minimizes

$$\left(\mathbf{Y_1} - \mathbf{Y_0}\mathbf{W}^*(\mathbf{V})\right)'\left(\mathbf{Y_1} - \mathbf{Y_0}\mathbf{W}^*(\mathbf{V})\right),$$

where $\mathbf{Y_1}$ is a $T_0 \times 1$ vector of pre-acquisition prices in Los Angeles and $\mathbf{Y_0}$ is a $(T_o \times J)$ matrix of the pre-acquisition prices in the control cities.[35]

The synthetic control estimator for the acquisition's effect is the difference between the Los Angeles price and the synthetic control city price in a given post-acquisition period: $P_{1t} - \sum_{j=2}^{J+1} w_j^* P_{jt}$. Abadie, Diamond and Hainmueller (2010) derive the synthetic control estimator and discuss its theoretical properties using a model that generalizes the standard difference-in-differences model. They demonstrate how it accounts for time-varying confounders. The intuition is that the synthetic control group can fit $\mathbf{X_1}$ only if it also fits the unobserved predictors of price. They are able to put bounds on the associated bias and demonstrate that the synthetic control estimator does a better job of accounting for transitory shocks if the number of pre-acquisition periods is large relative to the scale of unobserved confounders.

[35] Finding $\mathbf{W}^*$ and $\mathbf{V}^*$ is a nested optimization problem. We use a software package developed in the R programing language to implement the synthetic control method and solve the nested optimization problem. The software runs the optimization twice. It uses two sets of starting values (equal **V** weights and a regression-based method of selecting initial **V** weights) and returns the solution with the lower loss. Abadie, Diamond and Hainmueller (2011) explain the functionality of the package in detail. Our results are qualitatively unchanged if we assign equal weight to each variable, solving for $\mathbf{W}^*$ with a predefined **V.**

7.4. Synthetic Control Results

The optimal synthetic control city for Los Angeles applies weights of 0.73 to Phoenix, 0.27 to Boston, and less than 0.001 to all other potential control cities. The components of $\mathbf{V}^*$ assign weights of 0.84, 0.15, and 0.01 to the pre-acquisition price level, per-capita income, and crude costs, respectively. Boston and Phoenix have relatively high gasoline prices and crude costs, similar to Los Angeles. While Boston's income is higher than Los Angeles, Phoenix's is lower, resulting in a weighted average that approximates Los Angeles.

FIGURE 2—Price Levels in Los Angeles and the Control Cities

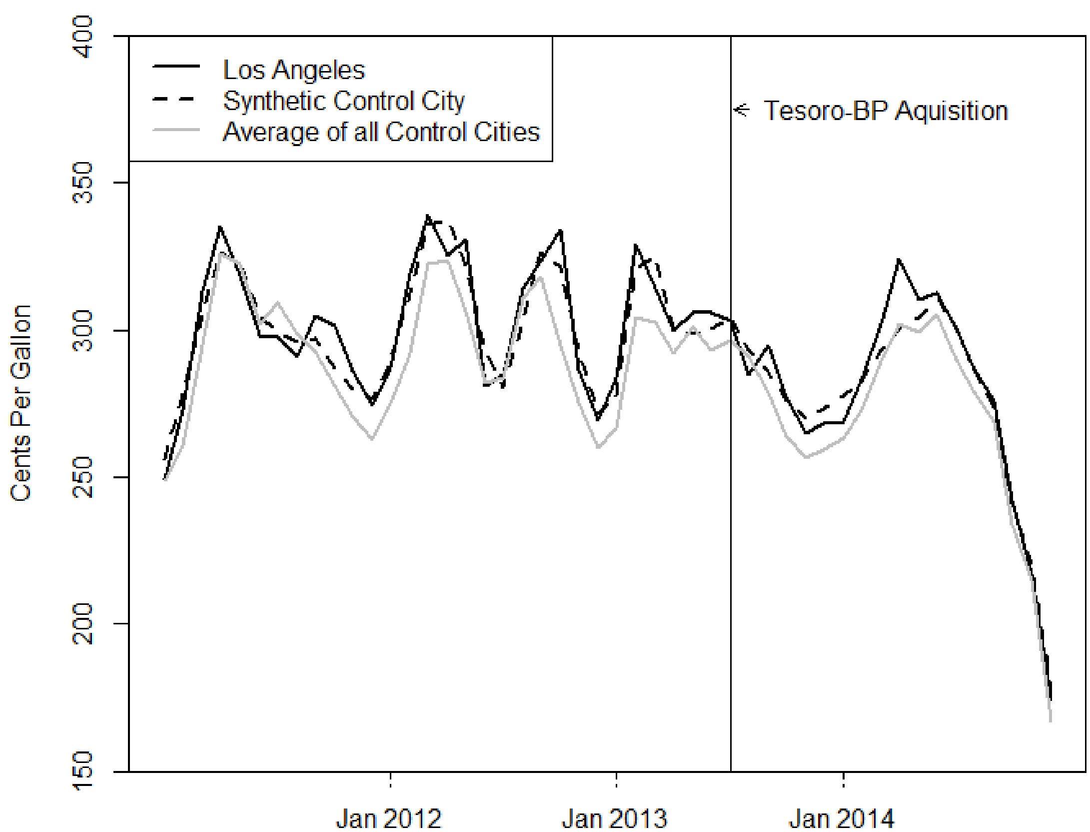

Figure 2 plots the level of prices in Los Angeles over time, as well as those of the synthetic control group and an equally weighted average across all 18 control cities. The synthetic control city's prices clearly track Los Angeles better than the average control city does.

In addition, there is no obvious change in the slope or level of Los Angeles prices relative to the control groups.

The average pre-acquisition price gap is 1.12 cents per gallon and the pre-acquisition MSPE (the mean of the squared price gap) is 57.68. Post-acquisition, the average price gap is 1.32 cents per gallon and the MSPE is 54.40. The slight decline in prediction error suggests that the synthetic control group's ability to replicate Los Angeles did not worsen. Hence, there is no evidence that the acquisition affected prices.

An increase in MSPE from the pre- to post-acquisition period is an important condition for reaching a conclusion that the acquisition had a causal effect on prices. The pre-acquisition MSPE quantifies the synthetic control city's ability to replicate the Los Angeles gasoline prices. If the MSPE does not increase following the acquisition, any post-acquisition gap may well reflect the inherent inability of the synthetic control group to replicate Los Angeles prices precisely, rather than a causal effect form the acquisition.

We further assess the significance of our estimates by performing a series of placebo tests. We use the synthetic control method to test for a price change during the sample period in each of the potential control cities. We iteratively apply the synthetic control method to every city in the data. For each iteration, we reassign the acquisition city to one of the 18 control cities and shift Los Angeles to the pool of potential cities for the synthetic control. Each iteration is a placebo test where we test for a merger effect in a city where no merger occurred. This generates an empirical distribution of estimated price effects. We can be confident there is evidence of an acquisition effect if the estimated effect in Los Angeles is an outlier among the distribution of placebo estimates.

Table 9 reports the mean gap between Los Angeles and the synthetic control city prices and the MSPE for the pre- and several post-acquisition periods. It also summarizes the distribution of these statistics from the placebo exercise. The post/pre MSPE ratio facilitates comparison between the Los Angeles estimate and the placebo estimates by creating a strictly positive measure that is normalized based on a particular synthetic control city's ability to predict prices before the acquisition.

Table 9—Mean Price Gap and MSPE, Los Angeles and the Placebo Tests

	Pre-Acquisition Period		**Three Alternative Post-Acquisition Periods**					
	Jan 2011 - May 2013		Jun 2013 - Dec 2014 Full Sample		June 2013 - May 2014 First 12 Months		Dec 2013 - Nov 2014 12 Months, 6-month transition	
	Mean Price Gap	MSPE	Mean Price Gap	Post/Pre MSPE	Mean Price Gap	Post/Pre MSPE	Mean Price Gap	Post/Pre MSPE
Los Angeles	1.12	57.68	1.32	0.94	2.22	1.44	2.21	1.16
Distribution from Placebo Exercise (includes Los Angeles estimate)								
Mean	-0.01	46.82	-1.32	2.47	-1.10	2.12	-1.31	2.78
5th Percentile	-1.13	5.11	-9.30	0.52	-8.10	0.44	-11.10	0.56
10th Percentile	-0.75	8.48	-5.45	0.66	-6.76	0.57	-7.22	0.84
25th Percentile	-0.04	17.05	-3.60	0.93	-2.87	0.95	-3.81	1.08
Median	0.00	31.58	-2.00	1.67	-0.96	1.55	-1.77	1.68
75th Percentile	0.03	74.07	2.16	2.74	2.15	2.62	1.77	3.28
90th Percentile	1.17	95.16	4.33	5.53	2.92	3.81	6.05	6.69
95th Percentile	1.41	111.66	5.19	6.91	3.14	5.31	7.54	7.65

Again, the average gap between the synthetic control city and Los Angeles during the post-acquisition period is 1.32 cents per gallon, approximately 0.4%. That falls within the interquartile range of the empirical distribution from the placebo exercise. Moreover, the post/pre MSPE ratio of 0.94 falls below the median of the distribution of the post/pre MSPE ratio. Together these facts suggest there is a greater than 50% chance of finding an effect of equal or greater magnitude if we were to relabel the acquistion city in our dataset at random. Table 9 reports similar statistics for two alternative post-acquisition periods; the qualitative results are unchanged.

Similar to the regresion analysis, there is no evidence of a price change. That is not to say that the placebo exercise provides a weak test. If we had found a 3% change in Los Angeles prices (9 cpg), the distribution from the placebo exercise would allow us to infer that there is a 5.2% (1/19 cities) chance of finding a comparable effect if we were to randomly select the treatment city from our dataset.

8. Conclusions

We model the California gasoline market using a variant of the Cournot model, which includes a competitive fringe of price-taking firms. This hybrid model performs better than the standard Cournot model for two reasons. First, the hybrid model implies more plausible margins than does the standard Cournot model, given widely accepted estimates of market demand elasticity. Second, the hybrid model more accurately predicts the price effect of the Tesoro/BP acquisition, relative to empirical estimates of the acquisition's price effect.

We estimate the effect of the acquisition using two methodologies. Difference-in-differences yields an estimate that is not statistically different from zero, nor is it statistically different from the 1% price increase predicted by our hybrid model. We can, however, reject the 6% price increase predicted by the standard Cournot model. Inference with the synthetic control method takes a different approach but yields a similar conclusion. Our placebo exercise indicates that if we relabel the acquisition city in our dataset at random, the probability of obtaining a result of the magnitude that we obtained for Los Angeles is greater than 0.5.

.

9) References

Abadie, A., Athey, S., Imbens, G. W., & Wooldridge, J. M. (2014). Finite Population Causal Standard Errors, *National Bureau of Economic Research,* working paper.

Abadie, A., & Gardeazabal, J. (2003). The economic costs of conflict: A case study of the Basque Country. *American economic review*, *93*(1).

Abadie, A., Diamond, A., & Hainmueller, J. (2011). Synth: An r package for synthetic control methods in comparative case studies. *Journal of Statistical Software*, *42*(13).

Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, *105*(490).

Ashenfelter, O. C., Hosken, D., & Weinberg, M. (2009). *Generating evidence to guide merger enforcement*. National Bureau of Economic Research.

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates?. *The Quarterly Journal of Economics*, *119*(1).

Borenstein, S., Bushnell, J., & Lewis, M. (2004). Market power in California's gasoline market. *Center for the Study of Energy Markets*.

Borenstein, S., & Bushnell, J. (1999). An empirical analysis of the potential for market power in California's electricity industry. *The Journal of Industrial Economics*, *47*(3).

Borenstein, S., Bushnell, J., & Knittel, C. R. (1999). Market power in electricity markets: beyond concentration measures. *The Energy Journal*, *20*(4).

Borenstein, S., Bushnell, J., & Stoft, S. (2000). The Competitive Effects of Transmission Capacity in a Deregulated Electricity market. *Rand Journal of Economics*, *31*(2).

Budzinski, O., & Ruhmer, I. (2010). Merger simulation in competition policy: A survey. *Journal of Competition Law and Economics*,*6*(2).

Carlton, D. W. (2010). Revising the horizontal merger guidelines. *Journal of competition Law and Economics*, *6*(3).

Coate, M. B., & Fischer, J. H. (2000). Can Post-Chicago Economics Survive Daubert? *Akron L. Rev.*, *34*.

Coglianese, J., Davis, L. W., Kilian, L., & Stock, J. H. (2015). Anticipation, tax avoidance, and the price elasticity of gasoline demand. *National Bureau of Economic Research.*

Donald, S. G., & Lang, K. (2007). Inference with difference-in-differences and other panel data. *The review of Economics and Statistics*, *89*(2).

Farrell, J., & Shapiro, C. (1990). Horizontal mergers: an equilibrium analysis. *The American Economic Review*, *80*(1).

Federal Trade Commission (2014). 2014 Report to Congress on Ethanol Market Concentration.

Federal Trade Commission (2006). Investigation of Gasoline Price Manipulation and Post-Katrina Gasoline Price Increases.

Gary, J. H., Handwerk, G. E., & Kaiser, M. J. (2010). *Petroleum refining: technology and economics*. CRC press.

Greenfield, D. (2014). Guide to Merger Analysis using Difference-in-Differences, Chapter 4 in *Mergers, Merger Control, and Remedies in the United States: A Retrospective Analysis*, John Kwoka, MIT Press.

Haas Wilson, D., & Garmon, C. (2011). Hospital mergers and competitive effects: Two retrospective analyses. *International Journal of the Economics of Business, 18*(1).

Hosken, D. S., Olson, L., & Smith, L. (2012). Do Retail Mergers Affect Competition? Evidence from Grocery Retailing. *Federal Trade Commission, Bureau of Economics,* working paper.

Hosken, D. S., Olson, L. M., & Smith, L. K. (2015). Can Entry or Exit Event Studies Inform Horizontal Merger Analysis? Evidence from Grocery Retailing. *Economic Inquiry*.

Hosken, D., Silvia, L., & Taylor, C. (2011). Does concentration matter? Measurement of petroleum merger price effects. *The American Economic Review*, *101*(3).

Kaplow, L., & Shapiro, C. (2007). Antitrust, *Handbook of law and economics* (Vol. 27) Polinsky, A. M., & Shavell, S. (Eds.), Elsevier.

Kwoka, J. (2014). *Mergers, Merger Control, and Remedies: A Retrospective Analysis of US Policy*. MIT Press.

Kim, E. H., & Singal, V. (1993). Mergers and market power: Evidence from the airline industry. *The American Economic Review*, 549-569.

Kreisle, N. (2015). Merger Policy at the Margin: Western Refining's Acquisition of Giant Industries. *Review of Industrial Organization*, *47*(1).

Leffler, W. L. (2008). *Petroleum refining in nontechnical language*. PennWell Books.

Levin, L., Lewis, M. S., & Wolak, F. A. (2015). *High frequency evidence on the demand for gasoline*. working paper.

Lin, C. Y. C., & Prince, L. (2013). Gasoline price volatility and the elasticity of demand for gasoline. *Energy Economics, 38*.

McAfee, R. P., & Williams, M. A. (1992). Horizontal mergers and antitrust policy. *The Journal of Industrial Economics*, *40*(2).

Perry, M. K., & Porter, R. H. (1985). Oligopoly and the incentive for horizontal merger. *The American Economic Review, 75*(1).

Peters, C. (2006). Evaluating the Performance of Merger Simulation: Evidence from the US Airline Industry. *Journal of Law and Economics, 49*(2).

Silvia, L., & Taylor, C. T. (2013). Petroleum mergers and competition in the northeast United States. *International Journal of the Economics of Business*, *20*(1).

Simpson, J., & Taylor, C. (2008). Do gasoline mergers affect consumer prices? The Marathon Ashland petroleum and Ultramar Diamond Shamrock transaction. *Journal of Law and Economics*, *51*(1).

Stigler, G. J. (1947). The kinky oligopoly demand curve and rigid prices. *The Journal of Political Economy, 55*(5).

Taylor, C. T., & Hosken, D. S. (2007). The Economic Effects of the Marathon-Ashland Joint Venture: The Importance of Industry Supply Shocks And Vertical Market Structure. *The Journal of Industrial Economics, 55*(3).

Tenn, S. (2011). The price effects of hospital mergers: a case study of the Sutter–Summit transaction. *International Journal of the Economics of Business*, *18*(1).

Tenn, S., & Yun, J. M. (2011). The success of divestitures in merger enforcement: Evidence from the J&J–Pfizer transaction. *International journal of industrial organization*, *29*(2).

Thompson, A. (2011). The effect of hospital mergers on inpatient prices: a case study of the New Hanover–Cape Fear transaction. *International Journal of the Economics of Business*, *18*(1).

Thomson StreetEvents. (2012). Tesoro Corporation to Purchase BP's Fully Integrated Southern California Refining and Marketing Business - Conference Call Transcript Aug 13, 2012. New York, NY.

Weinberg, M. C. (2011). More evidence on the performance of merger simulations. *The American Economic Review, 101*(3).

Weinberg, M. C., & Hosken, D. (2013). Evidence on the accuracy of merger simulations. *Review of Economics and Statistics*, *95*(5).

Werden, G. J. (2015). Inconvenient Truths on Merger Retrospective Studies. Available at SSRN: http://ssrn.com/abstract=2545343

Werden, G. J., & Froeb, L. M. (2008). Unilateral competitive effects of horizontal mergers. *Handbook of antitrust economics*, *43*.

Werden, G. J., Froeb, L. M., & Scheffman, D. T. (2003). Daubert Discipline for Merger Simulation, A. *Antitrust*, *18*.