

Testing Optimal Punishment Mechanisms under Price Regulation: the Case of the Retail Market for Gasoline

by

Robert Gagné*, HEC Montréal and CIRPÉE
Simon van Norden, HEC Montréal, CIRANO and CIREQ
Bruno Versaevel, EM Lyon and GATE

October 2006
Revised July 2007

ABSTRACT

Following a severe price war in the retail market for gasoline in 1996, the Quebec provincial government introduced a price floor. We analyse the effects of this new regulation on the pricing behaviour of a sample of gas stations in Montreal over the 1994-2002 period. We use a Markov Switching Model with two latent states to simultaneously identify the periods of low/high margins and estimate the parameters characterizing each state. We find that the net effect of the price floor on average margins is near zero, as the impact on retail prices in low margin periods is offset by the rise in the average duration of these periods. The increase of the duration of low margin periods lends supports to the presumption of collusive pricing behavior.

Keywords : price regulation, Markov switching model, gasoline prices.

JEL codes: L13, L81, C32.

1. Introduction

The behaviour of retail gasoline prices has long been and still is the object of fierce public debate. Given its importance for the consumers and the apparently “suspect” behaviour of the integrated oil companies (the “majors”), many jurisdictions have regulated aspects of gasoline retailing. For example, refiners are forbidden to operate retail outlets in some U.S. states. Different types of price regulations are also enforced in several U.S. states and Canadian provinces. In this paper, we provide evidence of the effects of price floor regulation on the behaviour of retail gasoline prices in Montreal, the largest market in the province of Quebec.

The May 4, 1992 edition of Bloomberg’s Oil Buyers Guide reported the conclusions of market studies indicating that “major Canadian oil companies were going to use price wars, new credit terms, and the strategic closure of service stations and refineries to squeeze independent gasoline retailers out of the market in central Canada”. A few months later, following a severe price war in the summer of 1996, the Quebec provincial government responded to the lobbying of independent gasoline retailers by establishing a price floor in December 1996. The regulation was motivated by the claim that price wars formed as a discipline device for the implementation of anticompetitive strategies.

The observed deep price cuts were commented quite extensively. For example, the local association of independent gasoline retailers reported that “[t]he summer 1996 episode was very harmful for suppliers in Quebec. The price war, triggered by an integrated major, resulted in retail prices that were observed *well below wholesale prices*. It was so severe as to force several independent retailers either to close down temporarily or to exit the market” (added emphasis, pp. 7-8).¹ The Liberal Committee on Gasoline Pricing in Canada explained that “(...) prices stayed at levels *below anyone’s cost of crude and taxes* with the independents continuing to be charged prices well in excess of the retail pump price for months” (added emphasis, p. 24).²

By limiting the severity of price wars, the floor was thus seen as a mean to reduce the ability of firms to punish retailers deviating from a high price strategy. When the latter retailers were small independent suppliers, typically characterized by less favourable financial conditions than majors, the introduction of a floor was rationalized as a form of protection that helps to maintain competition in the market. Since its introduction, the price floor has been computed weekly and regionally as the sum of the wholesale (rack) price, transportation costs and taxes. The floor level computed in week t is the one which applies in the market during week $t+1$. It is the only type of economic regulation specific to Quebec's retail market for gasoline.

Our objective in this paper is to determine whether the price floor actually reached the goal of the regulator, i.e. more competitive prices. The data we use consists of gasoline retail prices collected weekly in a sample of self-service gas stations on the Montreal market from 1994 to 2001. Although prices partly display cycling features, the specific feature of the data is the occurrence of several deep price reductions. They lead to relatively large *negative* margins, in early 1995 and mid-1996.³ After the introduction of the price floor regulation, although negative margins are ruled out, severe price cuts can still be observed. These cuts are irregular, seem unpredictable, and may take place over several weeks.

To investigate the impact of the price floor regulation on the formation of prices, we use a Markov Switching Regression framework with two latent states to simultaneously identify the periods of low/high margins and estimate the parameters characterizing each state. We allow regulation to influence both the state-conditional prices and the expected duration of each state. The switching regression is then estimated on weekly data for retail gasoline price margins in Montreal from 1994 to 2001. In summary, we find that the introduction of a price floor significantly impacts retail prices in low margin periods, and also increases the persistence of those periods. The net effect on average margins is near zero, as the impact on retail prices in low margin periods is offset by the rise in their average duration.

Section 2, below, describes the data and exploits some insights from theoretical contributions on the possible impact of a price floor on the formation of retail prices. The empirical model and data are described in Section 3. The results are discussed in Section 4, and Section 5 presents concluding remarks.

2. The Data, Theory, and Related Literature

A recent literature explains features of retail gasoline prices by using data for Canadian cities. It refers to alternating-move price-setting duopoly models of the kind introduced by Maskin and Tirole (1988), in which a class of equilibria can be constructed that consists of price cycles. In these cycles, starting from a high price, firms repeatedly undercut one another in the downward portion of the cycle. When the price reaches the competitive value, each firm participates in a so-called “war of attrition” of variable length, by waiting for the rival to raise its price back to the high level, so that the cycle is repeated. In this theoretical framework, a price floor has no impact on the formation of prices when it is lower than or equal to the marginal cost.

Eckert (2002) formulates an econometric version of the model of price cycles, and estimates it with daily data that describe prices in the city of Windsor (Ontario) from 1989 to 1994. He finds results consistent with the theoretical predictions that price reductions result when firms undercut rivals, and that the size of price restorations depends on the wholesale price. Noel (2006b) develops an empirical framework to separate out different pricing phenomena, and applies it to the analysis of dynamic pricing behaviour in nineteen Canadian retail gasoline markets from 1989 to 1999. He finds that the most frequent pattern is one in which prices cycle rapidly, beginning with a large increase of short duration, before declining gradually over a longer time period. However, the evolution of prices in Montreal – in the period that precedes the 1996 events mentioned in the introduction – appears less easily classifiable in a given category of pricing phenomena than the patterns observed in other Canadian marketplaces.⁴

In the present paper, the weekly data we use have been provided by M.J. Ervin Inc., a Calgary-based firm which conducts a weekly survey on gasoline prices in all major Canadian markets. In each market surveyed, retail prices are collected by gasoline grade using a sample of self-service gas stations. Whenever possible, the same stations are surveyed each week. In the Montreal market, the survey covers approximately 20 stations. For our analysis, we use the average retail price for unleaded regular gasoline computed from all stations in the Montreal survey. The data cover the 1994-2001 period (416 weekly observations). Our analysis is limited to the price of unleaded regular gasoline as retail prices for all other grades follow closely unleaded regular gasoline prices. Wholesale prices are “rack” prices (excluding taxes) posted everyday at wholesale distribution outlets.

[Figure 1 here]

The retail price series, which displays highly frequent and irregular fluctuations of limited magnitude, seems to share only partly the main features of the pricing patterns analyzed by Eckert (2002, 2003) and Noel (2006a,b). Indeed theory predicts that the ascending phase of a cycle starts when the retail price reaches the marginal cost from above. In Noel (2006a,b), the dataset of retail gasoline prices collected in a sample of service stations in Toronto, Ontario, reveals clearly an asymmetric price cycle, with price falling gradually down to a level that is extremely close to the rack price, before a sudden and large increase occurs that starts the cycle anew. According to theory again, the dynamic of cycling retail prices may occur in complete independence of changes in the wholesale price. In Eckert (2003), a key feature of the data collected from a sample of stations in Windsor, Ontario, is that the asymmetric changes in retail prices are not present in the wholesale price. To compare, consider Figure 1, which shows the evolution of the weekly retail prices we examine, and the rack price. In our data the changes in retail prices are not characterized by a regular series of asymmetric upward and downward movement. There is no period in which wholesale prices are constant, while retail prices change. Most movements in the retail price series seem to amplify the fluctuations that occur in the wholesale price. Moreover, ascending movements, which do

not seem to be systematically more sudden than descending movements, do occur when the retail price is relatively high above the rack price.⁵

The key distinctive feature of our data is that sudden price cuts of large magnitude occur in addition to more limited and frequent price fluctuations. These deep price cuts are irregular, seem unpredictable and, surprisingly, lead to *large negative margins* on several occasions. This can be observed in Figure 2, which shows the evolution of retail margins over time. Since a margin is computed as a retail price in a given week, minus the rack price observed the same week, a negative margin means that firms charge below marginal costs. At first glance, periods in which margins are either negative or very close to zero do not last more than a few weeks. However, given the weekly nature of the data, each point of the graph may represent a constant price for up to seven days, implying that episodes may take place for over tens of days. To see that, consider the circumstances that motivated the introduction of a price floor, in mid-1996. The margin fell from 8.72 ¢/litre in week 22 down to a minimum of -5.875 ¢/litre in week 26. It remained below 1 ¢/litre from week 25 to week 27. Another example can be observed in late 1996, shortly before the introduction of the price-floor, about week 37 (the margin is -1.1 ¢/litre). Again the margin is strictly below 1 ¢/litre in three successive observations, from 1996-36 to 1996-38. After 1996, in practice the *Régie de l'Énergie du Québec*, Quebec's regulatory board in charge of implementing the price floor regulation, considers that stations are not covering their fixed operating costs if the difference between the retail price and the floor price (computed as the sum of the wholesale price observed during the previous week, corresponding taxes, and transportation costs) is below 3 ¢/litre for a month or more.⁶ There are five episodes with three successive points that describe a margin below the threshold, and three episodes with four observations in a row below the threshold (in early 1997 and mid-2000). Of course, large negative margins are no longer observed with a price floor (only small negative margins are still observed under regulation in 2000-17 and 2000-22 because the price floor is computed from the wholesale price observed the week before).

[Figure 2 here]

The deep and sudden price cuts we observe are not explained by Edgeworth cycles theory. Interestingly, they are consistent with models of tacit collusion, in which price wars alternate with longer collusive episodes. The specifications of these models, in their Bertrand version and with complete information, reflect the main properties of gasoline markets: sales are repeated, retailers use price as a strategic variable, prices are observable by all, rival sales can be monitored, price adjustment can be small and occur at low cost. In particular, Lambson (1987, 1994) characterizes optimal punishments for a large class of infinitely-repeated games, with price-setting sellers of a homogenous good and a constant marginal (and average) cost c .⁷ Häckner (1996) builds on this by demonstrating that a symmetric two-phase “stick-and-carrot” structure is an optimal price path in a specific supergame duopoly model with differentiated products.⁸ This means that, in a collusive period, all firms sell at a price that yields the highest sustainable level of profits.⁹ When firms detect a deviation from the collusive strategy, they all switch to a lower price in a one-period severe punishment phase (the “stick”), before returning to the collusive price to earn super-normal profits again (the “carrot”).¹⁰ The higher the deviation payoffs, the lower the optimal punishment price must be for collusion to remain sustainable. This price may therefore need to fall below marginal cost c , in which case the price floor constraint $p \geq c$ would be binding.¹¹ When this occurs, one-period punishments are not sufficiently large to outweigh the gain from undercutting the collusive price. However, Häckner shows that, by constraining the punishment price to remain above the marginal cost, it is possible that a *prolonged* price war can sustain the same collusive price as in the unrestricted case.¹² The empirical implication is that a price floor, when it prevents below marginal cost pricing, should lengthen the duration of price wars, to support an unchanged level of high margins in collusive periods. This should be reflected in the data by longer low margin periods.

In the models of collusion we mentioned, punishment prices may be charged below marginal costs, preventing price wars to take place in equilibrium. Other well known contributions by Porter (1983a) and Green and Porter (1984) have established that unobserved random shocks on demand may induce price wars in equilibrium. Slade

(1989, 1992) has extended the analysis to model specifications which conform closely to the characteristics of a retail gasoline market.^{13,14}

To summarize, the dynamic of our retail price series seems to combine in a complex way some aspects of cycling behaviour with the below marginal cost pricing behaviour explained by models of collusion. In theory, a price floor that equates marginal cost – and thus is not above it – will not impact the formation of prices in Edgeworth cycles, but may lengthen the duration of punishment strategies in tacitly collusive schemes.¹⁵ In practice, although the separation into war of attrition/price and upper part of a cycle/collusion is uncertain, one may attempt to establish whether the regulation has actually modified or not the formation of prices. In the empirical work that follows, we examine the impact of the price floor on the magnitude of margins, and also on the duration of high margin and low margin periods. Only a lengthening of the latter periods may come in support to a presumption of collusive behaviour.

3. Econometric Implementation

Variables and Descriptive Statistics

As already mentioned, retail prices (P_t) are computed as the average of the posted prices collected weekly on approximately 20 gasoline stations. Wholesale prices (W_t) are “rack” prices, excluding taxes, and posted everyday at wholesale distribution outlets. Weekly wholesale prices are obtained by averaging the daily prices. No transportation costs are considered given the proximity of the retail market to the different wholesale distribution outlets.

In our empirical analysis, we use retail margins (M_t) rather than prices in order to eliminate effects coming from the wholesale market and thereby concentrate on retail market effects. Retail margins are computed as retail prices in week t minus wholesale prices observed during the same week. Recall that the price floor for week t is computed as the sum of the wholesale price observed during week $t-1$, corresponding taxes and

transportation costs. Finally, our empirical analysis includes a regulation dummy (R_t). It is equal to 0 until the price floor regulation was introduced during the last week of December 1996, and equal to 1 thereafter.

Table 1 presents some descriptive statistics. We see that average margins are roughly the same before and after regulation but significantly less dispersed after regulation (the standard deviation is 2.61 before regulation and 2.00 after). Given that peaks are approximately of the same magnitude before and after regulation (see Figure 2) and that average margins are about the same size, small margins (i.e. deep price cuts) should be more prevalent after regulation, in accordance with theoretical predictions. This higher prevalence of small margins is not obvious in Figure 2, which calls for a more elaborate data analysis.

Toward this aim, in Table 2 we present the number of weeks (in %) where the margins are below and above average margins, before and after regulation. Before regulation, margins were below average 43% of the time while after regulation this number increased at 48%. Again, given that average margins were nearly the same before and after regulation, this result is also an indication (albeit a weak one) that periods of small margins are more prevalent since regulation. This provides some modest albeit encouraging indirect support for our theoretical prediction. We therefore now turn to a more formal statistical analysis of retail margins using a Markov Switching regression approach.

[Tables 1 and 2 here]

Econometric Model

To assess the impact of the price floor on margins, we need to allow for the structural relationships determining margins to vary depending on whether the industry is in a period of low retail prices, possibly under the wholesale price (the low margin state), or not (the high margin state). We further need to investigate how the duration of these

different states may be influenced by changes in the regulatory environment. Estimation and inference is complicated by the fact that, while regulatory changes are directly observed, the state must be inferred indirectly.¹⁶

Of course, one could simply construct a binary variable to indicate which observations appear to correspond to a low pricing period as predicted by theory (a war of attrition in the case of cycles, or a price war in the case of tacit collusion), then use traditional methods (e.g. Ordinary Least Squares) separately on the two distinct subsets of observations (for example, see Borenstein (1991, 1996)). While intuitive, this approach has serious problems. First, since the separation into war of attrition/price and upper part of a cycle/collusion is somewhat uncertain, some observations will be misclassified. That means this approach will produce biased and inconsistent estimates of the underlying relationships. Second, standard errors for the resulting regressions will ignore the contribution of uncertainty about the sample separation, making inference unreliable as well. The approach we adopt avoids both of these problems.

Porter (1983) and Lee and Porter (1984) address the same kind of problem using regime switching techniques, which estimate the structural parameters of each pricing regime together with probability that each observation may have been produced by a price war. We use an extension of their approach, inspired by Hamilton (1993)'s Markov Switching Models with time-varying transition probabilities.¹⁷

As our baseline model, we estimate the following system of equations by maximum likelihood

$$M_t = \alpha_i + \rho_i M_{t-1} + \beta_i R_t + \gamma_i W_t + \varepsilon_t \quad (1)$$

$$\Pr(S_t = i | S_{t-1} = i) = \Lambda(\varphi_i + \theta_i R_t + \omega_i W_t) \quad (2)$$

where M_t , W_t and R_t are defined as earlier, $i = 1$ for the low margin state and 0 otherwise, S_t is the market state (1 for the low margin state and 0 otherwise) at time t , $\Lambda(\cdot)$ is the logit cumulative distribution function, ε_t is an i.i.d. mean-zero normally-distributed

error term with a standard deviation of σ_ε , and $\{\alpha_0, \alpha_1, \rho_0, \rho_1, \beta_0, \beta_1, \gamma_0, \gamma_1, \varphi_0, \varphi_1, \theta_0, \theta_1, \omega_0, \omega_1, \sigma_\varepsilon\}$ is the vector of unknown parameters to be estimated by maximum likelihood. We performed the usual diagnostics tests suggested by Hamilton (1996) for the fit of such models.¹⁸

Since the two states in this model follow a first-order Markov chain, we can calculate the half-life of a regime (the length of time over which the probability of remaining in the same regime has fallen to 50%) as $\ln 0.5 / \ln a$, where a is the probability given by equation (2). Similarly, we can calculate the expected duration of the regime (in periods) as $1/(1 - a)$. Furthermore, if $\{a_0, a_1\}$ are the probabilities of remaining in regimes 0 and 1 for one more period, then on average the market will spend fraction $(1 - a_1)/(2 - a_0 - a_1)$ of the time in the “low” regime 0 and the remainder in the “high” regime 1.

The sign and significance of β_i determine whether the introduction of a price floor has an impact on margins either during low margin periods ($i = 1$) or high margin periods ($i = 0$). Due to the dynamic nature of the model, the long-run impact of the regulation on margins in each regime will be $\beta_i / (1 - \rho_i)$. The parameter θ_i tells us whether the price floor raises the probability of being in the same regime the following period. To determine the average margin before and after the introduction of a price floor, we therefore need to take account of the floor’s effect on average margins in each state (e.g. if there is a price war, it can be less intense) as well as on the average fraction of the time the market will spend in that regime (e.g. the price wars may last longer.) This average will be given by

$$\frac{\alpha_0 + \gamma_0 \bar{W}}{1 - \rho_0} \cdot \frac{1 - \Lambda(\varphi_1 + \omega_1 \bar{W})}{2 - \Lambda(\varphi_0 + \omega_0 \bar{W}) - \Lambda(\varphi_1 + \omega_1 \bar{W})} + \frac{\alpha_1 + \gamma_1 \bar{W}}{1 - \rho_1} \cdot \frac{1 - \Lambda(\varphi_0 + \omega_0 \bar{W})}{2 - \Lambda(\varphi_0 + \omega_0 \bar{W}) - \Lambda(\varphi_1 + \omega_1 \bar{W})} \quad (3)$$

before the introduction of the price floor and by

$$\frac{\alpha_0 + \beta_0 + \gamma_0 \bar{W}}{1 - \rho_0} \cdot \frac{1 - \Lambda(\varphi_1 + \theta_1 + \omega_1 \bar{W})}{2 - \Lambda(\varphi_0 + \theta_0 + \omega_0 \bar{W}) - \Lambda(\varphi_1 + \theta_1 + \omega_1 \bar{W})} +$$

$$\frac{\alpha_1 + \beta_1 + \gamma_1 \bar{W}}{1 - \rho_1} \cdot \frac{1 - \Lambda(\varphi_1 + \theta_1 + \omega_1 \bar{W})}{2 - \Lambda(\varphi_0 + \theta_0 + \omega_0 \bar{W}) - \Lambda(\varphi_1 + \theta_1 + \omega_1 \bar{W})} \quad (4)$$

after, where \bar{W} is the average wholesale price.¹⁹

4. Empirical Results

The results of the estimation of equations (1) and (2) by maximum likelihood are presented in Table 3. The first two columns present estimates using margin and regime selection specifications without W ; the last two columns present estimates with W included in both margin and regime selection specifications. We test the fit of the above models using the score-based tests proposed by Hamilton (1996). Results are shown in Table 4. These tests have power against unmodeled serial correlation and heteroscedasticity in equation (1). They also test for omitted higher-order Markov dependence in equation (2). Since a higher-order Markov chain may be rewritten as a first-order chain with a larger number of states, the latter test also has power against omitted states. We test for each effect in each regime and also present a test for the joint null hypothesis of none of these forms of misspecification in any of the model's equations.

[Tables 3 and 4 here]

The joint misspecification test in Table 4 finds some evidence of misspecification without W . However, this evidence appears to be entirely confined to higher-order Markov dependence in the high regime; the test statistics are more than three times larger than those for any other single test and are the only ones significant at the 5% level. There is no significant evidence of misspecification in the unrestricted model. Although we have provided no theoretical justification for the inclusion of wholesale prices in the regime

switching model, we include these results to check whether our conclusions are robust to minor changes in specification, and whether the limited misspecification detected in our restricted model seriously affects the reliability of our statistical inferences. We also find that the parameters associated with the wholesale price are negative and significant during high regimes. A reduction in margins in the face of increased costs during high margin periods is consistent with monopolistic pricing behaviour.

The Effects of Regulation on Price Formation

As far as the effects of regulation are concerned, results are similar for both specifications shown in Table 3. Significant effects of regulation are found in the margin equation as well as in the state transition probability equation, but only in the low regime. The introduction of a price floor significantly increases margins during low margin periods and significantly increased the persistence of these periods. While models of Edgeworth cycles, as evoked in Section 2, are silent on the duration of “wars of attrition”, this finding is consistent with the theoretical prediction of models of collusion, in which the length of the punishments increases when the price floor constraint is binding.

As already mentioned, in both specifications (with and without W) and in both sets of equations (margin and regime selection) considered, the parameters associated with the regulation dummy (R_t) are positive and significant only during low regimes. This result is consistent with the theoretical predictions where a price floor regulation has no effect on prices (or margins), both in the upper part of a cycle and during a collusive period, since a price close to the monopoly level is assumed to be charged in the two cases. Moreover, our results show that during low regimes, the price floor regulation increases both margins and probabilities of continuing to charge a low retail price. Table 5 reports estimated regime-dependent conditional probabilities (using equation 2), expected durations (using $1/(1 - q)$, where q is the regime dependent conditional probability) and margins (using equations 3 and 4), all computed from the estimates presented in Table 3.

[Table 5 here]

During high margin regimes, the transition probabilities are similar before and after regulation and their estimates are quite robust to the inclusion/exclusion of the wholesale price. However, during price war regimes, transition probabilities are significantly higher after regulation again, regardless of the specification considered. For example, in the unrestricted model the transition probabilities increased from 0.34 to 0.90 after regulation during price wars while the same probabilities stayed at around 0.90 during collusive periods.

Regulation therefore significantly increased the expected duration of low margin periods, while reducing sharply the expected duration of high margin periods (although not significantly.) Before regulation, the expected duration of low margin periods was a little less than two weeks, which rises to ten weeks after regulation.

Results on estimated margins are fully consistent with the theory. On one hand, in high regimes, estimated margins are about the same magnitudes with and without a price floor regulation. On the other hand, regulation increases significantly the margins during low margin periods: from approximately -1.5 cents before the price floor regulation to 3.5 cents with a price floor. However, and unsurprisingly, we verify that regulation did not raise low margin levels to high margin levels. Furthermore, since the price floor raises margins in the low margin state, but increases the time spent in this state, its net effect on the average margin (i.e. its unconditional expectation) in this market is potentially ambiguous. To answer this question, Table 6 presents estimated unconditional probabilities of the high regime (this is 1 minus the corresponding probability for the low regime) as well as the unconditional expected margin. From those figures, it appears that the increase in margins during high regimes has been almost exactly offset by the increase in the average duration of these periods, resulting in no significant change in the average margins in the industry. In other words, the price floor regulation had little or no effect on average margins (and therefore prices) even if margins are now higher during low margin periods, simply because those periods last longer.

[Table 6 here]

5. Concluding Remarks

The empirical results obtained with a Markov Switching Model using data on the Montreal retail market for gasoline show that the introduction of a price floor regulation limits the decrease of retail prices in low margin regimes, but raises their expected duration.

Our results lead to three important conclusions. First, given the coherence between the estimated results and the theoretical predictions offered by models of tacit collusion, the empirical analysis lends support to the presumption that some collusive behaviour took place on the retail market for gasoline in Montreal in the period 1994-2001. This does not imply that other pricing dynamics, including cycles, never occurred in the same period. However, models of Edgeworth cycles do not predict that a price floor – when it does not exceed the marginal cost – will increase the persistence of low margin periods. Second, the government regulation succeeded in reducing the severity of price-wars; we see far less pricing below marginal cost after the introduction of regulation than before. Third, the net effect of the regulation on the competitiveness of the industry is ambiguous. There is no demonstrable change in the average price margin after regulation. This would appear to have benefited gasoline consumers only if one assumes that margins would have risen had regulation not been introduced.

Figure 1 Retail Prices and Wholesale Prices, Montreal 1994-2001




Figure 2 Retail Margins, Montreal 1994-2001




Table 1 Descriptive Statistics

Variable	Mean	Minimum	Maximum	Standard Deviation
Price : P_t (¢/litre)	28.7738	15.4000	49.7000	7.1932
Wholesale price: W_t (¢/litre)	24.3781	13.7250	42.8250	6.8470
Margin : M_t (¢/litre)	4.3957	-5.8750	8.8000	2.2452
Margin before regulation	4.3504	-5.8750	8.7200	2.6105
Margin after regulation	4.4229	-1.2000	8.8000	1.9993
Regulation : R_t (0,1)	0.6274	0	1	0.4841

Table 2 Percentage of time with margins below or above mean margin

	Before regulation (156 weeks)	After regulation (260 weeks)	Total (416 weeks)
Margin < Mean Margin	43%	48%	46%
Margin > Mean Margin	57%	52%	54%

Table 3 Parameter Estimates (Dependent variable is M_t)

Variable	AR(1) without W		AR(1) with W	
	Estimated value	Standard error	Estimated value	Standard error
Margin Equation				
High Margin Regime				
<i>Constant</i>	3.7034*	0.5926	4.0761*	0.3621
M_{t-1}	0.2546*	0.0947	0.2088*	0.0598
R_t	0.1679	0.5060	0.7352	0.4501
W_t	-	-	0.0067	0.0324
Low Margin Regime				
<i>Constant</i>	-1.3040	0.9118	-1.1864*	0.5928
M_{t-1}	0.1242	0.1107	0.2591*	0.0739
R_t	4.4870*	1.1564	3.7059*	0.6684
W_t	-	-	-0.0684*	0.0214
Regime Selection Equation				
High Margin Regime				
<i>Constant</i>	1.6296*	0.3169	1.1344*	0.2617
R_t	-0.3362	0.6831	-0.1476	0.4610
W_t	-	-	-0.1078*	0.0343
Low Margin Regime				
<i>Constant</i>	-0.0950	0.3983	-0.4228	0.3535
R_t	1.3906*	0.7181	1.7070*	0.4707
W_t	-	-	-0.1451*	0.0511
Error Variance				
σ_ε	1.8386*	0.1233	1.6395*	0.0817

* Denotes statistical significance at the 5% confidence level

Table 4 Diagnostic Tests

Test for ²⁰	Without W		With W	
	Statistic	p-value	Statistic	p-value
Serial Correlation – High Margin Regime	0.13	0.718	3.45	0.063
Serial Correlation – Low Margin Regime	0.20	0.655	0.00	0.976
ARCH	5.76	0.016	2.59	0.107
Higher-order Markov Dependence – High Margin Regime	18.17	0.000	3.56	0.059
Higher-order Markov Dependence – Low Margin Regime	3.21	0.073	0.04	0.842
Joint Test	46.75	0.000	10.50	0.062

Table 5 Regime Dependent Statistics

Regime	P($S_t=i S_{t-1}=i$)		E(duration)		Estimated Margins	
	without W	with W	without W	with W	without W	with W
High Margins ($R_t=0$)	0.94841	0.87170	19.382	7.794	4.9683	5.1520
High Margins ($R_t=1$)	0.90206	0.83814	10.210	6.178	5.1936	6.0813
Low Margins ($R_t=0$)	0.46217	0.33624	1.859	1.507	-1.4889	-1.6013
Low Margins ($R_t=1$)	0.90246	0.90047	10.252	10.048	3.6344	3.4007

Table 6 Unconditional Probabilities and Margins

State	P($S_t=$ high margins)		Estimated Margins		Sample Margins
	without W	with W	without W	with W	
$R_t=0$	0.91247	0.83801	4.4031	4.0581	4.3504
$R_t=1$	0.49899	0.38077	4.4124	4.4214	4.4229

References

- Abreu, Dilip J. (1986), "Extremal Equilibria of Oligopolistic Supergames", *Journal of Economic Theory* 39(1), 191-225.
- Abreu, Dilip J., Pearce, D. and E. Stacchetti (1986), "Optimal Cartel Equilibria with Imperfect Monitoring", *Journal of Economic Theory* 39(1), 251-269.
- Bloomberg (1992), *Bloomberg Oil Buyers' Guide*, May 4, 21(1066).
- Borenstein, Severin (1991), "Selling Costs and Switching Costs: Explaining Retail Gasoline Margins", *RAND Journal of Economics* 22(3), 354-369.
- Borenstein, Severin and Andrea Shepard (1996), "Dynamic Pricing in Retail Gasoline Markets", *RAND Journal of Economics* 27(3), 429-451.
- Eckert, Andrew (2002), "Retail Price Cycles and Response Asymmetry", *Canadian Journal of Economics* 35(1), 52-77.
- Eckert, Andrew (2003), "Retail Price Cycles and the presence of Small Firms", *International Journal of Industrial Organization* 21(2), 151-170.
- Filardo, Andrew J. (1994), "Business-Cycle Phases and Their Transitional Dynamics", *Journal of Business and Economic Statistics* 12(3), 299-308.
- Filardo, Andrew J. and Stephen Gordon (1998), "Business Cycle Durations", *Journal of Econometrics* 85(1), 99-123.
- Green, Edward J. and Robert H. Porter (1984), "Noncooperative Collusion under Imperfect Price Information", *Econometrica* 52(1), 87-100.
- Haltiwanger, John and Joseph E. Harrington (1991), "The Impact of Cyclical Demand Movements on Collusive Behavior", *RAND Journal of Economics* 22(1), 89-106.
- Häckner, Jonas (1996), "Optimal Symmetric Punishments in a Bertrand Differentiated Products Duopoly", *International Journal of Industrial Organization* 14(5), 611-630.
- Hamilton, James D. (1993), "Estimation, Inference and Forecasting of Time Series Subject to Changes in Regime". In G. S. Maddala, C. R. Rao, and H. D. Vinod, eds., *Handbook of Statistics*, Vol. 11, North-Holland.
- Hamilton, James D. (1994), *Time Series Analysis*, Princeton University Press, 799 p.
- Houde, Jean-François (2005), "Entry and Exit in a Price Regulated Industry: Gasoline Retailing in Quebec", *Queen's University*, mimeo.

Kim, Chang-Jin and Charles R. Nelson (2001), *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, MIT Press, 1999, 297 p.

Lambson, Val E. (1987), "Optimal Penal Codes in Price-setting Supergames with Capacity Constraints", *Review of Economic Studies* 54(3), 385-397.

Lambson, Val E. (1994), "Some results on Optimal Penal Codes in Asymmetric Bertrand Supergames", *Journal of Economic Theory* 62(2), 444-468.

Lee, Lung-Fei and Robert H. Porter (1984), "Switching Regression Models With Imperfect Sample Separation Information – With An Application on Cartel Stability", *Econometrica* 52(2), 391-418.

Maskin, Eric and Jean Tirole (1988), "A Theory of Dynamic Oligopoly II: price Competition, Kinked Demand Curves and Edgeworth Cycles", *Econometrica* 56(3), 571-599.

Noel, Michael D. (2006a), "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market", *Journal of Industrial Economics* 55(1), 69-92.

Noel, Michael D. (2006b), "Edgeworth Price Cycles, Cost-based Pricing, and Sticky Pricing in Retail Gasoline Markets", *Review of Economics and Statistics* 89(2), 324-334.

Porter, Robert H. (1983a), "Optimal Cartel Trigger-Price Strategies", *Journal of Economic Theory* 29(2), 313-338.

Porter, Robert H. (1983b), "A Study of Cartel Stability: the Joint Executive Committee, 1880-1886", *The Bell Journal of Economics* 14(2), 301-314.

Rotemberg, Julio J. and Garth Saloner (1986), "A Supergame-Theoretic Model of price Wars during Booms", *American Economic Review* 76(3), 390-407.

Seegerstrom, Paul S. (1988), "Demons and Repentance", *Journal of Economic Theory* 45(1), 32-52.

Slade, Margaret E. (1987), "Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion", *Journal of Industrial Economics* 35(4), 499-516.

Slade, Margaret E. (1989), "Price Wars in Price-setting Supergames", *Economica* 56(223), 295-310.

Slade, Margaret E. (1992), "Vancouver's Gasoline-Price Wars: An Empirical Exercise in Uncovering Supergame Strategies", *Review of Economic Studies* 59(2), 257-276.

van Norden, Simon and Robert Vigfusson (1996), "Regime-Switching Models, A Guide to the Bank of Canada Gauss Procedures", Working Paper 96-3, Bank of Canada.

* Corresponding author. Address : Institut d'économie appliquée, HEC Montréal, 3000 chemin de la Côte-Sainte-Catherine, Montréal, Québec, Canada H3T 2A7. Email : robert.gagne@hec.ca.

¹ Translated from the *Mémoire de l'Association Québécoise des Indépendants du Pétrole*, June 1998.

² Quoted from the *Report of the Liberal Committee on Gasoline Pricing in Canada*, Ottawa, 1998.

³ In this paper, we follow Slade (1992, footnote 4, p. 260) by using the term “price war” to denote a punishment strategy of some kind in which prices are below their collusive level.

⁴ In a second price pattern, retail prices remain “fixed” for months at a time. In a third “normal” pattern, retail prices follow closely wholesale prices. In Noel (2006b), retail gasoline prices in Montreal are reported to cycle 66.5% of the time, to be fixed 17.9% of the time, and to display a normal behavior in the remaining 15.6%. These figures (in Table 6, p. 35) appear to be intermediate with respect to the cases of Vancouver (43.7%; 30.0%; 27.2%) and Toronto (83.9%; 14.6%; 1.4%).

⁵ Eckert (2002) explains that, when the current price lies above marginal cost, and an unanticipated increase in the wholesale price occurs, “it is reasonable to expect that firms will continue to undercut each other until the new lower bound is reached” (p. 64, added emphasis).

⁶ For a detailed description of the law on petroleum products administrated by the *Régie de l'Énergie du Québec*, see Houde (2005).

⁷ In Lambson (1994), firm-specific capacities are introduced to make firms dissimilar, with no particular reference to retail gasoline markets. Eckert (2003) and Noel (2006b) mention the assumption that gasoline retailers are capacity constrained, but do not formalize it.

⁸ In Häckner (1996), the chosen formulation of the differentiation assumption fits well with the fact that, in retail gasoline markets, product varieties are not only differentiated by brand or service, but also (most importantly) by location. This is emphasized in an empirical analysis by Borenstein and Shepard (1996, p. 430), who argue that retailers compete with local rivals only, and thereby justify the use of oligopolistic settings in urban areas where stations are many.

⁹ The most severe punishments are defined as “optimal” in the sense that they lead firms to obtain the highest level of sustainable collusive profits. This follows Abreu (1986).

¹⁰ The optimal symmetric punishment must satisfy two conditions. First, the one-period gain from a deviation from the collusive price must be less than (or equal to) what is lost when the punishment price is charged in the next period. Second, the one-period gain from such a deviation must be less than (or equal to) what is lost by prolonging the punishment by one more period. When these two conditions hold with equality, and for a given discount factor δ , they constitute a system of two equations which is solved by the pair (p^*, p) . In the absence of restriction on δ , all firms would charge the joint profit-maximizing price as a collusive price, that is $p^* = p^m$. Häckner (1996) demonstrates that there exists a threshold value δ' such that $\delta < \delta'$ implies $p^* < p^m$.

¹¹ In Häckner (1996), the constant marginal (and average) cost is normalized to zero (i.e. $c = 0$) for simplicity, and the case of positive marginal costs is treated as a generalization (see Häckner (1996), footnote 13, p. 623).

¹² Note that the relatively large number of geographically close outlets that are available to retail gasoline customers in an urban area contributes to a high level of product substitutability among consumers. In this theoretical framework, such substitutability implies larger market share gains (and related profits) for those who deviate from collusive pricing. Large incentives for deviations from collusive pricing in turn can call for a severe (below marginal cost) punishment price. This may motivate a competition authority to establish a price floor.

¹³ The main idea is that, although all prices are observable and rival sales can be monitored, when a shock has occurred firms do not know the true demand conditions unless they change their price and thereby precipitate a price war, which forms as a learning process for parameter identification. Slade (1989) models the process of adjustment to new stationary-equilibrium prices. Slade (1992) builds on this using price-war data to estimate different specifications of the inter-temporal reaction function. In Slade (1989, 1992), continuous strategies are distinguished from discontinuous strategies, such as reversion to the one-shot Nash equilibrium in Green and Porter (1984), in that large (small) deviations lead to large (small) punishments. This formulation is justified by Slade (1987), who produces an empirical test of tacit collusion in the Vancouver retail gasoline market. It is found that continuous strategies provide a better

model of the price-war dynamics than discontinuous price-reversions to the one-shot Nash equilibrium. In Slade (1992), again with data on retail gasoline activity in a region of Vancouver, reaction functions that are piecewise-linear in the previous-period prices of rivals are given empirical support. Pricing below marginal cost is not discussed.

¹⁴ Other models produce price wars at equilibrium with more severe punishments than Nash reversion. For example, Abreu, Pearce, and Stachetti (1986) modify the model of Green and Porter (1984) by replacing the continuum of possible output levels with a discrete production set for each firm. Another example is Segerstrom (1988), who introduces a small probability that a player will deviate, and derives conditions under which the deviating firm will find it optimal to “repent” by reducing its own output level for a certain number of subsequent periods.

¹⁵ In this study, we do not attempt to integrate the many other models of price formation in gasoline markets (e.g., Rotemberg and Saloner (1986), Haltiwanger and Harrington (1991)). Their relevance to the present analysis of the impact of a price floor regulation is somewhat limited, as they do not allow for the possibility of pricing below marginal costs (not even out of equilibrium), nor rationalize the occurrence of price wars (in equilibrium).

¹⁶ Slade (1992) addresses this issue using Kalman filtering methods for unobserved component models. However, those are not quite appropriate here since they are designed for unobserved variables that are continuous, not dichotomous. See Hamilton (1994) or Kim and Nelson (2001) for a discussion of the relationship between Kalman filtering approach and the regime switching techniques we use.

¹⁷ See also Filardo (1994) and Filardo and Gordon (1998).

¹⁸ All the technical details regarding the estimation procedure, likelihood function, derivatives and computer code are presented in van Norden and Vigfusson (1996).

¹⁹ This average should correspond closely to the sample average of margins before or after regulation.

²⁰ Each individual test statistic has an asymptotic $\chi^2(1)$ distribution under the null hypothesis of no misspecification, while the joint test is asymptotically distributed as a $\chi^2(5)$ under the null.