

(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation

Zhongmin Wang

Northeastern University

This paper studies oligopoly firms' dynamic pricing strategies in a gasoline market before and after the introduction of a unique law that constrains firms to set price simultaneously and only once per day. The observed gasoline pricing behavior, both before and under the law, is well captured by the Edgeworth price cycle equilibrium in the Maskin and Tirole dynamic oligopoly model. My results highlight the importance of price commitment in tacit collusion. I also find evidence that the price leadership outcome under the law is better predicted by mixed strategies play than by alternative hypotheses.

I. Introduction

This paper examines the oligopoly problem by studying firms' dynamic pricing strategies in a retail gasoline market before and after the introduction of a unique law that constrains firms to set price *simultaneously* and *only once per day*. The observed gasoline pricing behavior, both before

I thank the editor and two anonymous referees for their comments, which significantly improved this paper. I am especially grateful to Jim Dana for many helpful and stimulating discussions. I thank Kamran Dadkhah, Juan Dubra, John Kwoka, Christian Rojas, and seminar participants at Clark University for their comments on an earlier version of this paper. I also thank Simon Goss, Sean Isakower, Shu Li, and Jingjing Yang for their assistance with collecting data; Allan Price at Informed Sources; staff members at the Western Australia Department of Consumer and Employment Protection; the owners of three gasoline retailers for providing data; several industry participants for explaining the Perth gasoline market and providing evidence on the special forms of vertical restraints documented in this paper; and the Faculty of Business and Economics at Monash University for providing financial support. Any errors are mine only.

[*Journal of Political Economy*, 2009, vol. 117, no. 6]
© 2009 by The University of Chicago. All rights reserved. 0022-3808/2009/11706-0005\$10.00

and under the law, is well characterized by the Edgeworth price cycle equilibrium in the Maskin and Tirole (1988) model that features short-run price commitment. My results thus highlight empirically the importance of price commitment in tacit collusion. I also find evidence that the price leadership outcome under the law is better predicted by mixed strategies play than by alternative hypotheses even though firms have the incentive not to deliberately randomize.

Central to the oligopoly problem is the question of how price is formed in oligopoly markets. Given a market setting, strategy describes behavior. The question then becomes what strategies oligopoly firms use to form price. Game-theoretic oligopoly theories have mainly considered two classes of dynamic strategies. The Maskin-Tirole approach presumes that once a price is set, it cannot be changed in the short run, because of price rigidity or commitment, so that another firm can subsequently react to that price. Thus, a firm reacts to a past price because that price affects its payoff. In contrast, in the canonical repeated game or supergame model, a Bertrand game is independently repeated so that past prices have no tangible effect on the current market: nothing prevents firms from changing price every period. Hence, a firm elects to condition its current price on past prices only because it decides to do so, not because past prices directly affect its payoff.

There are many game-theoretic models of tacit collusion in the literature (see Tirole [1988], Shapiro [1989], and Vives [2001] for references), but there is relatively little empirical evidence on oligopoly firms' actual dynamic pricing strategies in real markets.¹ A major reason for this imbalance between theory and evidence is that much of the discussion on tacit collusion is based on the supergame approach, but the strategies used in supergame models are typically not intended as descriptions of what firms actually do in reality. For example, the folk theorem is a key result of the supergame approach, yet according to Mailath and Samuelson (2006, 73), it "says nothing about behavior. The strategy profiles used in proving folk theorems are chosen for analytical ease, not for any putative content, and make no claims to descriptions of what players are likely to do." This paper presents evidence that the strategies featured in the Maskin-Tirole model effectively describe firms' dynamic pricing behavior in a real market.

This paper's window into firms' pricing strategies is a unique law that regulates the *timing* and *frequency* of price setting. The law, a remarkable market experiment, took effect in Western Australia in January 2001. It

¹ There is a literature on how cartels work (see Genesove and Mullin [2001] for an excellent example). My focus is on tacit collusion. There is also anecdotal evidence on facilitating practices such as price leadership. I analyze price leadership as part of firms' dynamic pricing strategy.

requires every gasoline station in the Perth metropolitan area to (1) notify the government of its next day's retail prices by 2:00 p.m. each day so that these prices can be published on an Internet Web site, and (2) post the published prices on its price board at the start of the next day for a duration of at least 24 hours. This law thus forces firms to set gasoline prices simultaneously (without knowing rivals' prices) and at most once every 24 hours.

In this paper, I document firms' pricing strategies before and under the law using a rich and high-frequency data set that tracks the price changes of nearly every gasoline site in the Perth area and the daily changes in marginal costs of supplying wholesale and retail gasoline. Figure 1 shows the hourly (from 6:00 a.m. to 6:00 p.m.) brand average retail gasoline prices of three firms in the Perth area for a period of 39 days before the law. Figure 2 shows the daily brand average retail prices of three firms and a wholesale gasoline price series for a period of 57 days under the law. Figure 3 shows a theoretical example of Maskin and Tirole's Edgeworth price cycle equilibrium. In all three figures, firms hike price sequentially and then decrease price gradually, and at the bottom of each cycle, a war of attrition problem arises. Price increases by all would benefit all, but none would like to be the first to hike price.

My empirical evidence supports Maskin and Tirole's theory that price commitment is important to tacit collusion. Before the law, a lead price hike would stick for a couple of hours; it is simply too costly to change price every hour or minute, even in the gasoline market. Thus, rival firms can observe and react to the lead price hike within hours before the law. On the other hand, a firm must be committed to its price increase for at least a day under the law. This large increase in the length of price commitment implies a much higher cost for price leadership: a leader has to lose market share for an entire day before any rival firm can respond. Consistent with this change, I find that a single large firm was nearly always the first to hike price before the law, but price leadership under the law is distributed among the three largest firms.

Mixed strategy is part of the dynamic strategies used by Maskin and Tirole to derive the Edgeworth price cycle equilibrium. Specifically, they presume that firms play a mixed strategy to decide price leadership² at the bottom of each price cycle. In this paper, I take this assumption seriously.

Mixed strategy is a fundamental concept in game theory, widely used in both zero-sum and nonzero-sum games. Because naturally occurring

² Price leadership is "one of the most important institutions facilitating tacitly collusive pricing behavior" (Scherer and Ross 1990, 346). The Maskin-Tirole model generates price leadership as part of the price cycle equilibrium.

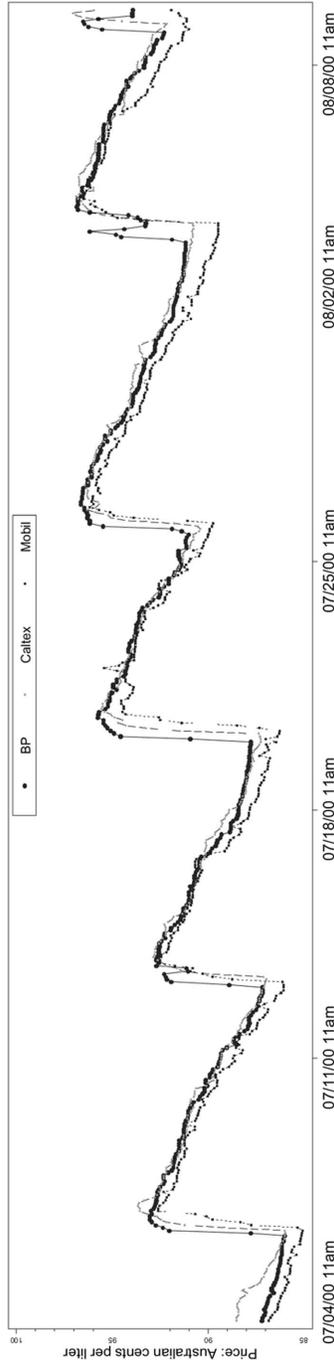


FIG. 1.—Hourly brand average gasoline prices over six cycles before the law

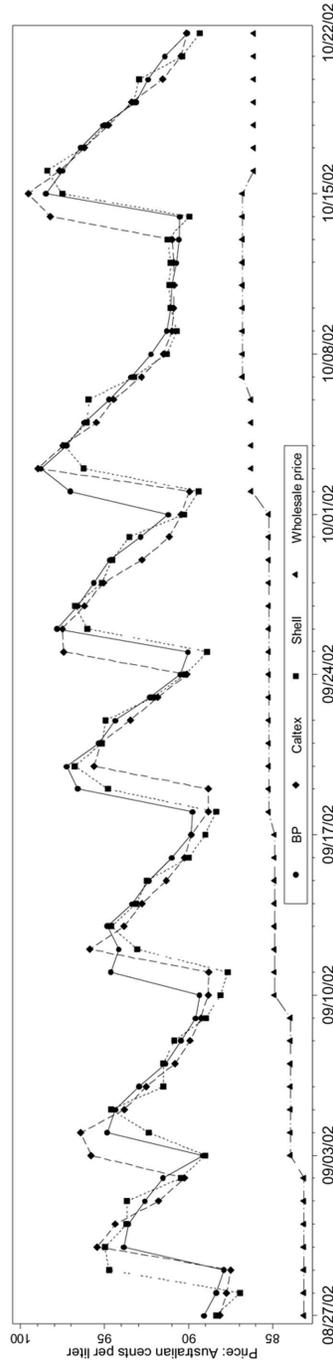


FIG. 2.—Daily brand average gasoline prices over seven cycles under the law

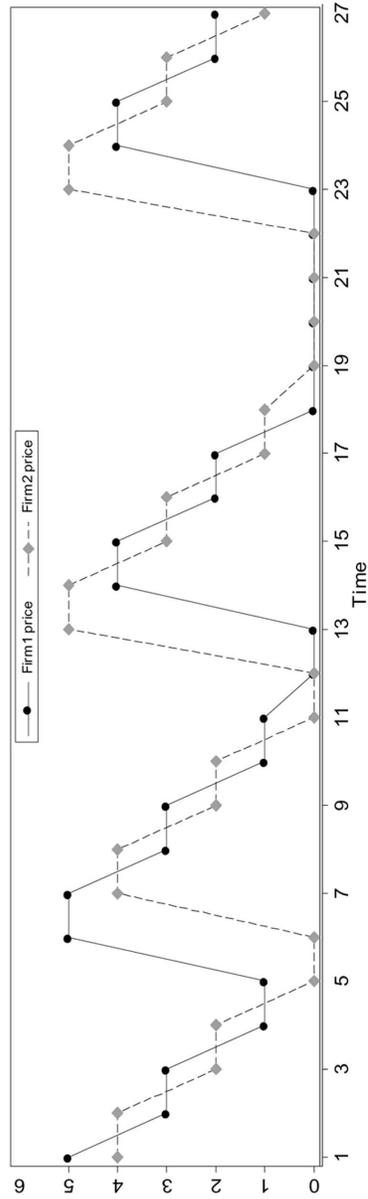


FIG. 3.—Maskin and Tirole (1988) Edgeworth price cycle

strategic situations are often too complicated to be suitable settings for empirical testing, nonexperimental tests of the mixed strategy concept have been limited to the zero-sum games of tennis serves and soccer penalty kicks (Walker and Wooders 2001; Chiappori, Levitt, and Groseclose 2002; Palacios-Huerta 2003; Hsu, Huang, and Tang 2007).³ The use of mixed strategy in zero-sum games is relatively intuitive since players in such games have the incentive to deliberately randomize to remain unpredictable (von Neumann and Morgenstern 1944, 146). However, the use of mixed strategies in nonzero-sum games is rather counterintuitive since players in such games often have the incentive to avoid randomization (Schelling 1960, 175). This raises the question of whether there are any empirically tractable strategic situations in which players do not have the incentive to randomize, but their actions are well characterized by equilibrium mixed strategies.

This is the first paper in the literature to study such a strategic situation: the war of attrition game at the bottom of the gasoline price cycles. Firms in this game have the incentive to avoid randomization since pure strategies would end a war of attrition immediately without incurring any cost of delay. **In this paper, I can test the mixed strategy hypothesis using techniques similar to those used to test mixed strategies in sports games.** First, the actions in this game are discrete (either relent or fight), and the outcomes are the identities of the price leaders. Second, the same game is repeated many times, providing rich variations in the observed outcome.

The firms did not play mixed strategies before the law when a single large firm appeared to serve as the market leader. However, I find empirical evidence for mixing behavior under the law. First, the leadership outcome of the individual wars of attrition under the law, once conditional on the outcome of the previous war, is random. Second, the stochastic regularities of the leadership outcomes (leadership types and their frequencies) are captured reasonably well by the mixed strategy presumed by Maskin and Tirole. The observed mixing behavior differs from that presumed by them in one important aspect: The outcomes of the wars of attrition are serially correlated, suggesting that the firms may have attempted to coordinate over the wars of attrition even under the law. The finding of mixing behavior under but not before the law is consistent with Harsanyi's (1973) Bayesian justification for mixed strategy: A player's mixed strategy represents other players' uncertainty of

³ The results from experimental tests are rather mixed (see, e.g., Walker and Wooders [2001] for a review). A literature estimates game-theoretical models that may involve mixed strategies (e.g., Hendricks and Porter 1988).

that player's pure choice.⁴ Under the law, firms are uncertain about rivals' actions when deciding whether to hike price, but that uncertainty does not exist before the law.

A regular gasoline price cycle is not unique to the Perth market. It appeared in many U.S. cities in the 1960s (Castanias and Johnson 1993), and it is currently occurring in many U.S. cities in the Midwest (Lewis 2009) and in Canada (e.g., Eckert and West 2004; Noel 2007). Without observing a market experiment or being free from data constraints, these studies do not highlight the importance of short-run price commitment or test the mixed strategy hypothesis.

The rest of the paper proceeds as follows. Section II discusses Maskin and Tirole's model of dynamic oligopoly pricing. Section III describes the market setting and the data set. Section IV documents that key regularities of firms' pricing behavior are characterized by Maskin and Tirole's Edgeworth price cycle equilibrium. Section V tests the mixed strategy hypothesis under the law. Section VI evaluates the welfare impact of the law, and Section VII presents conclusions.

II. Maskin and Tirole's Model of Dynamic Oligopoly Pricing

The idea of reaction based on commitment traces back at least to the classic Stackelberg model. Ever since Schelling (1960), commitment has been recognized as a central feature of much strategic behavior. Maskin and Tirole (1988) formalize the idea of reaction based on commitment in a fully dynamic setting. To capture the existence of short-run price commitment, they assume that a price, once set, lasts for two periods. One justification for this assumption is that price is rigid in the short run because of exogenous price adjustment costs. Price rigidity thus can serve as a commitment device. To capture the idea of reaction, Maskin and Tirole presume that two firms set price alternately when competing repeatedly over the price of a homogeneous good. Firms follow Markov reaction strategies in that a firm responds only to the price set by the rival firm in the previous period, which is the only payoff-relevant variable. Either the Edgeworth price cycle equilibrium or the kinked demand curve equilibrium arises as a Markov perfect Nash equilibrium in this model. Since equilibrium profit in either case is well above the one-shot Bertrand profit, tacit collusion is not only possible but necessary in this model.

⁴ Harsanyi (1973) shows that almost any mixed strategy equilibrium can be viewed as a pure strategy Bayesian equilibrium in a nearby game in which the payoffs to each player are subject to small private random variations. Aumann (1987) takes Harsanyi's idea further and directly interprets a player's mixed strategy as an expression of other players' uncertainty of that player's choice of pure strategy. See Reny and Robson (2004) for a recent discussion of the interpretation of mixed strategy.

Recall figure 3 for an example of the Edgeworth price cycle equilibrium, which has two rare features. First, equilibrium prices change over time even though the underlying demand and cost remain constant, and second, equilibrium prices are directly indicative of the underlying pricing strategies. If price is at marginal cost, firms are in the *war of attrition phase*: Both firms want to hike price, but neither wants to be the first to do so because the leader loses market share whereas the follower gets a free ride. For the cycle equilibrium to arise, however, the public good of price leadership must be provided. Maskin and Tirole presume that firms play a mixed strategy to decide price leadership. This technical assumption is elaborated immediately below. Once a firm relents by hiking price, the other firm reacts with a slightly smaller increase in the following period. These two price increases constitute the *rising phase* of a price cycle. In the subsequent *falling phase*, the two firms undercut each other gradually until price reaches marginal cost.

Maskin and Tirole presume that firms decide price leadership by playing the standard stationary mixed strategies in wars of attrition: firm i always relents with probability p_i in period t conditional on no firm having relented before then. When a rival's price is at the competitive level, a firm's best response is to attach probability p_i to relent (i.e., increase price) and $1 - p_i$ to fight (i.e., keep price at marginal cost). In any period of an attrition war, a firm is indifferent between relenting and fighting, with the value of an action given by a Bellman equation. Mixed strategy is a convenient technical device here. By playing a mixed strategy, both firms have a chance to be the leader for any cycle so that the burden of price leadership is shared over time.

In the prelaw Perth market, however, firms were unlikely to play mixed strategies. Players in wars of attrition have the incentive not to randomize: mixed strategies lead to costs of delay whereas pure strategy equilibria end a war of attrition game immediately. In addition, the disincentive to be a leader before the law is quite small. The gasoline firms constantly monitor each other's price changes, so once a firm relents, rival firms can quickly follow by hiking their price as well. This reasoning suggests that a single firm may be willing to serve as the price leader.

The law does not change the basic intuition behind the price cycle equilibrium, which, in the words of Tirole (1988, 256), is "that if firms were stuck in the competitive price region, with the prospects of small profits, a firm could raise its price dramatically and lure its rival to charge a high price for at least some time (the rival would not hurry back to nearly competitive prices)." Indeed, most postlaw changes in the observed gasoline price cycles are intuitive to the point of being self-explanatory. For this reason, I discuss only briefly the impact of the law on the war of attrition game and on welfare.

The timing law changes the war of attrition game at the bottom of

each price cycle in significant ways. As mentioned earlier, the cost of providing price leadership is much greater now; the leader must lose market share for at least 24 hours before any rival firm can respond. This implies that price leadership needs to be allocated among the firms. Indeed, postlaw price leadership in Perth is distributed among the three largest firms. Under the law, firms still have the incentive to coordinate on pure strategy equilibria. To do that, however, the three largest firms must be certain about each other's pure action at the beginning of each war, which is implausible under the law. Firms cannot credibly signal their intent at the bottom of the price cycles under the law. The timing of the wars of attrition game is thus forced to be discrete and simultaneous. Consequently, on a day of a war of attrition, a firm may decide to relent or fight, but rival firms are uncertain about that firm's pure action because firm-specific private information always exists. Therefore, the presumed mixed strategies may describe the leadership pattern under the law.

In the short run, the welfare impact of the law is clear. It lowers average retail gasoline price because it disrupted the gasoline firms' pricing coordination. Section IV.B.1 documents that regular gasoline price cycles disappeared after the law took effect (see App. fig. A1). However, after the firms succeeded in coordinating on the price cycle equilibrium under the law, it is not clear whether the average retail price is higher or lower than it was before the law.

III. Market Setting and Data Set

The Perth gasoline market resembles Maskin and Tirole's (1988) model environment. Only a few firms are in the market during the sample period. Price is the primary strategic variable, and retail gasoline price is publicly observable. Gasoline is a relatively homogeneous product,⁵ and the demand for gasoline is stable on a daily basis.⁶ The only significant discrepancy is that the cost of gasoline varies significantly over time as oil price fluctuates. A major feature of the market setting is that it allows one to construct cost measures that closely track the daily changes in marginal costs of supplying wholesale and retail gasoline.

During the sample period July 1, 2000, through October 31, 2003, the major gasoline firms in the Perth market include four oil firms (BP,

⁵ Wang (2009) collects daily station-specific gasoline sales data in the Perth market and uses the timing of the regular price cycles as instruments to estimate station-level gasoline demand. His elasticity estimates confirm that drivers in the Perth market are highly price sensitive to station-level gasoline price differentials: the estimated own price elasticity is as large as -18.8 .

⁶ In any case, changes in gasoline demand have little effect on the price cycle dynamics. Gasoline firms face the same demand shocks, yet it is often the case that some firms increase price on a day but others do not.

Caltex, Shell, and Mobil) and two independent firms (Gull and Peak). BP operates the only refinery in Western Australia. Before July 2002, Caltex, Shell, and Mobil all obtained fuel from BP through reciprocal refinery exchange programs. Since July 2002, Caltex and Shell have been buying fuel from BP, and Mobil has been importing fuel from Singapore. During the sample period, Gull bought fuel from BP and Peak bought fuel from Mobil.

Caltex and Shell buy fuel from BP through contracts that are typically renewed every 6 months. It is important to note that BP's gasoline pricing is constrained by potential import from Singapore because the BP refinery is "small in scale and less efficient than refineries in the Asia-Pacific region, particularly the large modern refineries in Singapore" (Australian Competition and Consumer Commission 2007, 100). For this reason, Caltex's and Shell's purchase contracts specify that the price that they pay BP is determined by a formula tracking the potential cost of importing from Singapore. This feature of the market allows one to construct a cost measure that closely tracks the movements in Caltex's and Shell's marginal costs of supplying wholesale gasoline.

By law, oil firms in Australia can own or operate only a small number of retail sites.⁷ It is widely acknowledged, however, that the law is ineffective in preventing oil firms from controlling retail gasoline price.⁸ BP, Shell, and Mobil used what are called multisite franchise agreements to control their retail gasoline price.⁹ For example, Shell had a single franchisee that essentially operated all the Shell branded retail sites in the Perth market, and a report by the WA government concluded unequivocally that BP directly controlled the retail prices of its multisite franchisees (Western Australia Select Committee on Pricing of Petroleum Products 2000, 38).¹⁰ Prohibited from using multisite franchise agreements, Caltex used a different mechanism, called the conditional price support, to effectively control the retail price of its many franchisees. Appendix B documents this interesting form of vertical restraint.

I describe below the retail price and cost indicators used in this paper. The price data used to evaluate the welfare impact of the timing law are described in Section VI. The unit of all price and cost data throughout the paper is Australian cents per liter.

⁷ The law is called the Petroleum Retail Marketing Sites Act 1980 (the Sites Act). According to the Sites Act returns for March 2002, BP owned six sites, Caltex owned 22 sites, and Mobil and Shell owned no sites in the Perth market.

⁸ The Sites Act was repealed in 2007, perhaps an admission of its ineffectiveness.

⁹ The multisite franchise agreements are known to be consignment agreements that allow the oil firms to retain the ownership of fuel until it is sold at retail, thus giving the oil firms the right to set retail prices.

¹⁰ The multisite franchise agreement "exhibited characteristics more consistent with commission agency than franchise operations."

A. Retail Price

To study the price cycles under the law,¹¹ I use a panel data set that records the daily (regular unleaded gasoline) price of all retail sites in the Perth market from the start of the law on January 3, 2001, through October 31, 2003.¹² During this period, BP, Caltex, Shell, and Mobil operated or controlled an average of 67, 88, 46, and 23 sites per day, respectively. Gull and Peak operated an average of 38 and 18 sites per day. These six brands accounted for about 85 percent of the retail sites in the Perth metropolitan area. Appendix figure C1 shows the daily market average price from the start of the law through October 31, 2003.

To study the price cycles before the law, I use a panel data set that covers the retail price of nearly every gasoline site in the Perth area for the period July 1 to December 20, 2000.¹³ For three brands (BP, Caltex, and Mobil), the price data are hourly—between 6:00 a.m. and 6:00 p.m. each day, 7 days a week. The hourly prices, critical to the identification of price leaders, were electronically sourced from purchase transactions with gasoline credit cards.¹⁴ For Shell and the independent brands, retail prices were collected (via drive-by) twice a day—between 5:00 a.m. and 9:00 a.m., and noon and 3:00 p.m., Monday through Friday. The number of gasoline sites covered in the data set before the law is slightly smaller than that under the law. For example, on July 1, 2000, the prelaw data set covers 286 stations: 89 Caltex sites, 73 BP, 35 Shell, 26 Mobil, 30 Gull, 7 Peak, and 26 sites of several small independent brands.

B. Cost Indicators

Since the marginal cost of supplying gasoline varies considerably even in the short run, it is important to study the potential impact of cost changes on pricing dynamics. To do so, I use two types of cost measures that capture the short-run variations in suppliers' marginal costs of sup-

¹¹ The timing law is called the 24-hour rule. The Internet Web site established by the law mentions that "motorists' frustration at intra-day price fluctuations and the significant difference between city and country fuel prices" were the political reasons that led to the 24-hour rule. There was a loophole to the law prior to August 24, 2001. During that period, a station must *nominate* its next-day retail price but is not required to move to the nominated price. Since the major firms in the market did not take advantage of this loophole, it does not affect the analysis in this paper.

¹² The data were downloaded from the Internet Web site (<http://www.fuelwatch.wa.gov.au>) established by the law.

¹³ The data were collected by Informed Sources, a market research firm in Australia.

¹⁴ Many drivers in Australia purchase gasoline using gasoline credit cards (e.g., Caltex Star Card). Each time a gasoline credit card is used to purchase gasoline at a retail site, the retail price at the pump is sent electronically to Informed Sources. The price data provided to the author are the latest prices for each station at each hour between 6:00 a.m. and 6:00 p.m.

plying gasoline: the confidential wholesale transaction prices paid by three retailers and a cost measure that closely tracks the movements of the actual gasoline price that Caltex and Shell pay BP. Other cost components, such as labor, inventory, and storage, are presumably fixed in the short run.

The three retailers are a BP franchisee that owns several sites and sets retail price independently, a small independent retailer,¹⁵ and a major independent retailer.¹⁶ The wholesale price shown in figure 2 is the wholesale price paid by the small independent retailer. I am able to estimate the gasoline price that Caltex and Shell pay BP because it is determined by a pricing formula that tracks the potential cost of importing gasoline from Singapore. The pricing formula, called the import parity pricing (IPP) formula, is the following:

IPP-based import cost

= a benchmark gasoline wholesale price in Singapore
+ shipping cost + quality premium + wharfage
+ insurance and loss + tax.

Detailed explanations of this formula can be found in Western Australia Department of Consumer and Employment Protection (2007, 19–22) or Australian Competition and Consumer Commission (2007, chap. 7). The Platts quote for the gasoline specification of Mogas 95 is the Singapore benchmark wholesale gasoline price, which drives the vast majority of the variations in the import cost. I have access to the daily IPP-based import cost for the 10-month period January 1 through October 31, 2003.¹⁷ I am not able to estimate BP's marginal cost of refining

¹⁵ This retailer initially bought fuel from Shell but later changed its fuel supplier to some other firm.

¹⁶ The major independent retailer's wholesale buying price is available for the period January 1, 2001, through June 30, 2002. The wholesale buying prices paid by the BP franchisee and the small independent retailer, respectively, are available for the period January 1, 2001, through October 31, 2003. The BP franchisee's wholesale buying price is missing for three months.

¹⁷ The import cost data were obtained from the Western Australia Department of Consumer and Employment Protection. The department did not provide the import cost series directly. Instead, the department provided the wholesale margin series, which is defined as the difference between the average terminal gate price and the import cost. By a regulation similar to the 24-hour rule, oil firms in Western Australia must report to the government by 2:00 p.m. each day their next day's terminal gate price (which is the maximum wholesale price an oil firm can charge any retailer). Since the terminal gate prices are published on the Internet, the import cost series can be calculated.

regular unleaded gasoline,¹⁸ but this does not affect the analyses significantly.¹⁹

Figure 4 shows Caltex's and Shell's brand average retail price, two retailers' wholesale buying price, and the import cost for the period July 20 through October 20, 2003. The import cost explains over 99 percent of the variations in the level of the BP franchisee's wholesale price, and the first difference in import cost explains 29 percent of the first difference in the BP franchisee's wholesale price. The regular cycle in retail price is clearly absent from the cost series. Indeed, the first difference in import cost explains 0.00 percent of the first difference in the daily retail price of either Caltex or Shell. To emphasize this point, figure 5 plots the corresponding daily margins for Caltex and Shell, defined as the difference between brand average retail price and the import cost. The margin series exhibit a cycle that is essentially identical to the retail price cycle except that the wholesale price trend is removed from the margin series. Moreover, import cost decreased monotonically, with a single exception, over the 27-day period September 4 through September 30, 2003, yet retail price continued to cycle during this period. Hence, cost changes cannot explain the existence of the price cycles. Nonetheless, the logic of the Maskin-Tirole model implies that cost changes may affect the price cycles in subtle ways, which I document in Section IV.C.

IV. Pricing Dynamics

Figures 1–5 suggest that the basic patterns of the gasoline price cycles, both before and under the law, are well captured by the Edgeworth price cycle equilibrium. The gasoline firms hike prices sequentially, decrease prices gradually, and confront the issue of which firm will be the first to increase price at the bottom of each cycle. This section uses the rich data set to document the gasoline firms' pricing strategies. In particular, I highlight the critical role of short-run price commitments in generating the regular price cycles. Subsection A documents the pricing behavior before the law. Subsection B describes how the firms adapted

¹⁸ It is difficult to estimate accurately a refiner's daily marginal cost of refining regular unleaded gasoline. First, the actual crude oil price paid by BP and other components of BP's production costs are not publicly available. Second, BP's refinery in the Perth area refines several types of petroleum products simultaneously, so it is difficult to determine the cost of a particular refined product.

¹⁹ The import cost should also track the movement in the actual price Gull pays to BP and Mobil's actual cost of importing. Gull has import facilities, so it pays a price that tracks the potential cost of importing as well. The Australian Competition and Consumer Commission (2007, 208) suggests that, as an independent gasoline chain, Gull's potential import cost is slightly higher than that of Caltex and Shell. The report also suggests that Mobil's actual cost of importing is not substantially different from the IPP-based import cost.

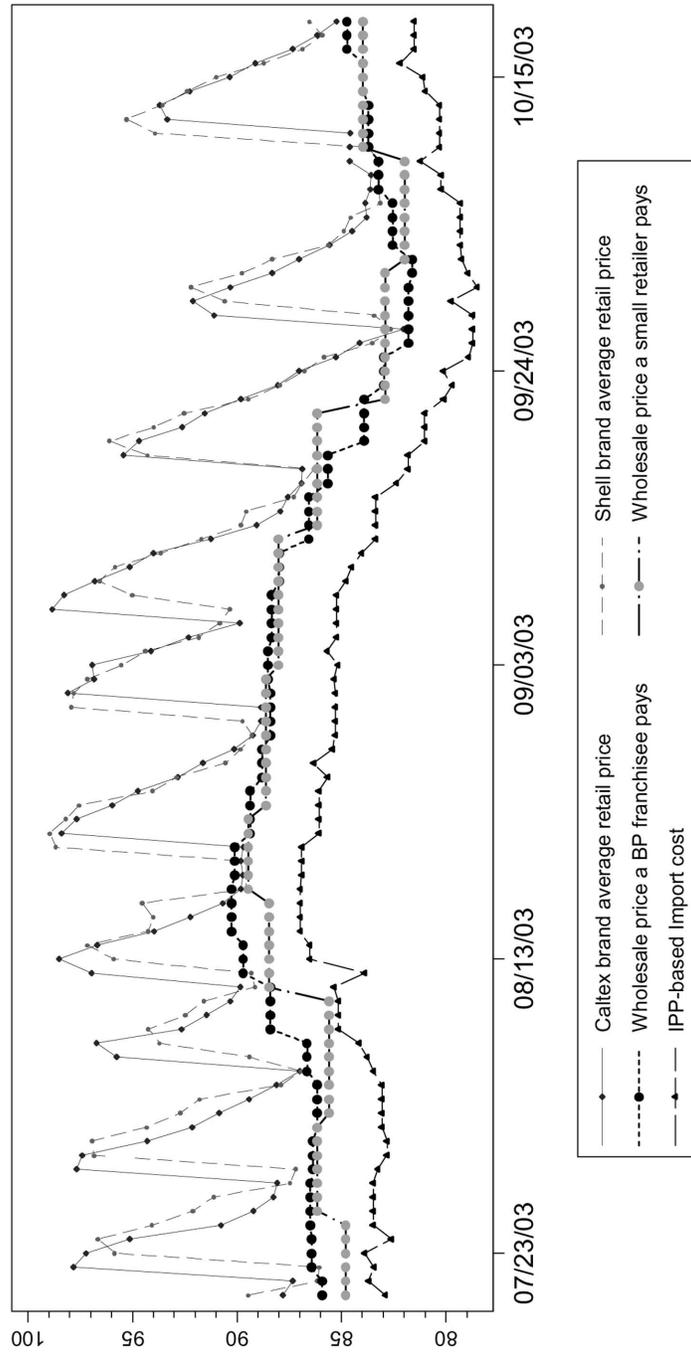


FIG. 4.—Daily brand average retail price, wholesale price, and import cost, July 20 to October 20, 2003

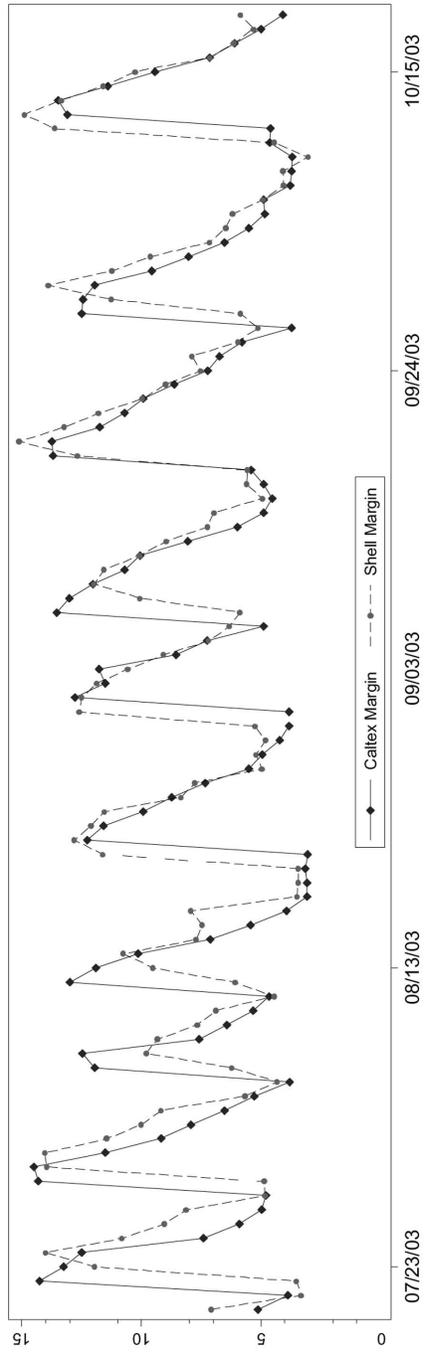


FIG. 5.—Daily margins for Caltex and Shell, July 20 to October 20, 2003

to the law and compares the pricing behavior before and under the law. Subsection C studies the impact of cost changes on the price cycles.

A. *Pricing Behavior before the Law*

1. Intra-brand Synchronization and Uniformity in Price Hikes

Figures 1, 2, and 4 indicate that the decision makers in the Perth market are the small number of oil and independent gasoline firms, not the hundreds of retail gasoline sites. The price hikes exhibit a conspicuous sequential pattern by brand, a pattern that can arise only if there is strong intra-brand synchronization. Indeed, the gasoline firms synchronize and homogenize intra-brand price hikes to facilitate inter-brand pricing coordination. To see intra-brand synchronization and uniformity directly, consider the rising phase of the first price cycle shown in figure 1, which took place on July 13, 2000. The average BP price increased from 87.2 cents per liter at 11:00 a.m. to 88.9 cents at noon and then to 91.9 cents at 1:00 p.m. The average increase from 87.2 to 88.9 cents occurred because the prices at 13 BP sites were increased from 86.5 to 92.9 cents between 11:00 a.m. and noon whereas the other 60 BP sites' prices were not changed. The average increase to 91.9 cents occurred because the prices at another 35 BP sites were increased to 92.9 cents (mostly from 86.5 cents) between noon and 1:00 p.m. Thus, between 11:00 a.m. and 1:00 p.m., 48 of 73 BP sites hiked price to exactly 92.9 cents. Since no sites from other brands increased prices by 1:00 p.m. that day, we observe strong (albeit not perfect) intra-brand synchronization and uniformity in price hikes.²⁰

Intra-brand synchronization and uniformity in price increases are important to inter-brand pricing coordination. By hiking the price of a large number of retail sites to a uniform level, a price leader makes a credible commitment, thereby inducing rival firms to feel the effect of strategic complementarity and establishing a single target at which its rivals will aim. Through intra-brand synchronization and uniformity, rival firms can react quickly to the first price hike, and the price leader can easily verify if its price increase has been matched in a timely manner.

2. Price Leadership and Followership

There are 21 regular price cycles for the prelaw sample period. BP was the first to hike price in 18 of the 21 cycles, and Shell was the leader for the other three cycles. In those three cases, BP started to increase price within an hour. Caltex, the largest firm in the market, was never

²⁰ Price decreases typically do not exhibit strong, if any, intra-brand synchronization or uniformity.

the first to increase price before the timing law. Since Caltex served as a leader most often under the law, I do not consider this prelaw price leadership pattern to be consistent with mixed strategy play. This leadership pattern probably reflects BP's position as the market leader since it owns the only refinery in the Perth area. This pattern also reflects the fact that BP's major rivals typically followed BP's initial price hike very quickly and that BP temporarily retracted its price hike if a major rival did not quickly follow.

The first price hike for each of the 21 cycles always occurred between 11:00 a.m. and 2:00 p.m. on Tuesday (five cycles), Wednesday (eight cycles), or Thursday (eight cycles). The length of a cycle, defined as the period between two lead price hikes, was 6 days (eight cycles), 7 days (seven cycles), 8 days (three cycles), or 9 days (three cycles). Caltex typically followed within 2 or 3 hours and Shell within 3 or 4 hours. The amount of time it took Mobil, Gull, and Peak to follow suit tends to be less precise, but for the vast majority of the 21 cycles, Mobil followed by 6:00 a.m. of the second day, and Gull and Peak followed by 9:00 a.m. of the second day.²¹ The oil firms tend to increase price to match the price leader whereas the independent firms tend to slightly undercut the price leader (typically by 0.2 cent). For example, on July 13, 2000, after BP increased most of its sites' prices to 92.9 cents per liter, most Caltex, Shell, and Mobil sites matched this price, but most Gull and Peak sites increased prices to only 92.7 cents.

3. Price Retraction

BP retracted its price increase temporarily if either Caltex or Shell did not follow quickly. Price retractions, visible in figure 2, may be partial or full. For instance, between 4:00 p.m. and 5:00 p.m. on July 13, 2000, 15 of the 56 BP sites that had hiked prices retracted those hikes and returned to their previous price levels. BP's price retractions over six other cycles are much more pronounced as many more BP sites retracted price hikes, and in two of these six cases, all sites returned to their prehike levels.

Price retraction is important to our understanding of the firms' pricing strategy. Price retractions are part of the Edgeworth price cycle equilibrium if the Maskin-Tirole model has three or more firms (Noel

²¹ For about half of the cycles, Mobil followed the leader within a few hours, but for the other cycles, the data are such that Mobil started to hike price at 6:00 a.m. the next day. Because the data between 6:00 p.m. and 6:00 a.m. are not available, it is possible that Mobil may have started to hike price before 6:00 a.m. In the data, Gull and Peak always started to hike price on the second day, mostly between 7:00 a.m. and 9:00 a.m. Because the data for Gull and Peak were collected only twice a day, it is possible that these two brands may have sometimes started to hike price earlier than what the data indicate.

2008). In such a model, after one firm relents by increasing its price, the two remaining firms may not follow immediately because they still have the incentive to be the last to increase price. Hence, the leader may temporarily retract its price increase. Price retraction highlights that the lead price hike represents a *short-run* price commitment. Temporary price retraction also suggests that BP's price commitment is partly endogenous. BP's lead price hike is a commitment partly because it involves a loss of market share that cannot be recovered, not purely because BP cannot change its price quickly.

B. Pricing Behavior under the Law

1. Short-Run Price Commitment during the Adjustment Period

Short-run price commitments emerged quickly after the law took effect, adding to the evidence that short-run price commitment is central to the gasoline firms' pricing behavior. Appendix figure A1 indicates that regular gasoline price cycles disappeared after the law took effect on January 3, 2001, and were not reestablished until early May 2001. However, Shell started to initiate large price increases as early as February 12, and Caltex followed suit in early March. Furthermore, it is important to note that both Shell and Caltex kept their price hikes on the second day, a display of price commitment beyond a single day. Shell retracted its price hike on the third day if either BP or Caltex did not fully follow, and Caltex did the same if either BP or Shell did not fully follow. Thus, it appears that short-run price commitments and subsequent price retractions were used to communicate intent and to help coordinate the new price cycles. BP did not initiate a price hike until April 18, and when it did, both Caltex and Shell followed on the second day, and a new regular price cycle started to emerge.

2. Pricing Dynamics after the Adjustment Period

There are 102 regular gasoline price cycles from May 10, 2001, through October 21, 2003. Intra-brand synchronization in price hikes is stronger in these cycles because of the need to increase price quickly. Before the law, a firm could hike some of its sites' prices one hour and others the following hours; a similar staggering of increases would take days to implement under the law. Once the new price cycle is established, lead price hikes are rarely retracted. This is expected since temporary retraction that lasts a few hours cannot exist under the law. There are three cases of price retraction over the 102 postlaw price cycles.²² In

²² All three price retractions (on June 16, 2001, November 3, 2001, and April 25, 2003) were by BP and were in full. The price retraction on November 3, 2001, by BP differs from the other two in that BP alone hiked its price just the day before.

two cases, BP, as the price leader, fully retracted its price hike on the third day after either Caltex or Shell failed to follow on the second day. Note that even in these two cases, BP kept its price hike on the second day. Over the rest of the price cycles, Caltex or Shell, if not a leader, always followed the leader on the second day. However, BP appears to receive special treatment: it did not follow on the second day in five cases, but neither Caltex nor Shell retracted prices in response. Mobil mostly followed on the second day, and the independent firms largely followed on the third day.

3. Drastic Changes in Price Leadership Pattern

As in figure 2, the price leaders for the postlaw cycles can also be clearly identified.²³ Figure 6 displays the leaders of the 102 price cycles between May 10, 2001, and October 21, 2003. Caltex, never before a leader, now initiates price hikes for 52 of the 102 cycles. BP, almost always a leader before, now initiates price hikes for only 49 of the 102 cycles. There are seven mutually exclusive and exhaustive leadership types: (1) BP leads alone (27 cycles), (2) Caltex leads alone (37 cycles), (3) Shell leads alone (15 cycles), (4) BP and Caltex lead together (eight cycles), (5) BP and Shell lead together (eight cycles), (6) Caltex and Shell lead together (one cycle), and (7) BP, Caltex, and Shell all lead simultaneously (six cycles). These observations are consistent with the hypothesis that price leadership under the law needs to be allocated among the firms.

4. Cycle Length Becomes More Unpredictable

Under the law, the length of a price cycle becomes much more unpredictable. Define the day on which the initiating price hike(s) took place as the *start day* of a price cycle and the day immediately before as the *last day* of the previous cycle. Table 1 shows the weekday frequency distribution of the cycle start day under the timing law. Over a third of the postlaw cycles start on a weekday other than the three days on which the prelaw cycles always started.

Define the *length* of a cycle as the number of days from the start day through the last day of the cycle. Table 2 shows the postlaw cycle length

²³ For 94 out of the 102 price cycles, the price leaders are as clear-cut as those for the cycles shown in fig. 2. The leaders for these 94 cycles are always BP, Caltex, or Shell. None of the other firms in the market led any of these 94 cycles. For the other eight cycles, one or more independent firms had positive average price changes on the day when one or more of the three largest firms hiked price. Because the independent firms' price increases are much smaller in size, they are not considered as price leaders. There is a full price cycle between October 22, 2003, and the end of the sample period. This last cycle is ignored in the analysis because Mobil co-led this cycle.

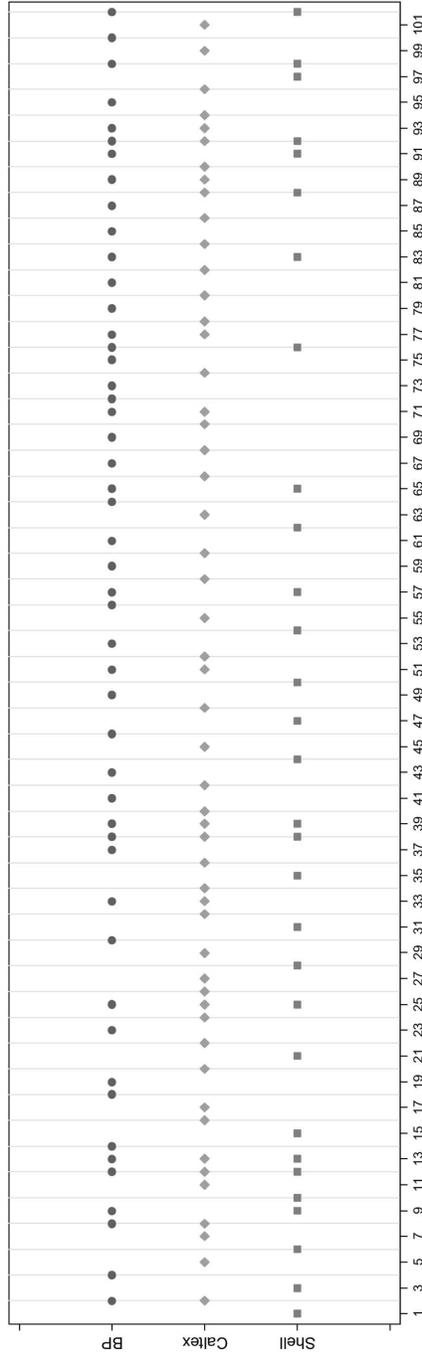


FIG. 6.—Price leadership pattern in the 102 wars of attrition under the law

TABLE 1
DAY OF WEEK FREQUENCY DISTRIBUTION OF CYCLE START DAY UNDER THE LAW

	DAY OF THE WEEK							1-7
	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday	
Start day frequency	7	25	20	19	5	1	25	102

distribution by leadership type. While the length of all prelaw cycles is between 6 and 9 days, the length distribution of the postlaw cycles is much more dispersed; about a third of the cycles have a length equal to or longer than 10 days. The length of postlaw cycles, even after breaking them down by leadership type, is still more dispersed than the length of all prelaw cycles. The average cycle length is not statistically different across three leadership types (BP leads alone, Shell leads alone, and multiple firms lead), and the average cycle length for the Caltex leading alone type is slightly bigger because the three longest cycles all happen to have been led by Caltex alone. While the null that the prelaw cycle length time series is generated by a white-noise process of uncorrelated random variables with a constant mean and a constant variance is soundly rejected by the Barlett periodogram-based test or the Box and Pierce Q test for white noise (the p -values are 0.0006 and 0.0000, respectively), the same null for cycle length under the law cannot be rejected by the same two tests (the p -values are 0.16 and 0.84, respectively).

C. *Impact of Cost Changes*

Section III.B presented evidence that cost changes cannot explain the existence of the regular price cycles because the cycles continue to exist even when cost is monotonically decreasing and gasoline margins exhibit a cycle that is essentially identical to the retail price cycle. Sections IV.A and IV.B provide additional evidence. If cost changes are driving the price cycles, why should the regular price cycles disappear in the first four months of the timing law? Furthermore, why should the price leadership pattern differ dramatically before and under the law? These facts are all consistent with the Edgeworth price cycle equilibrium. The firms face a war of attrition problem when price is in the competitive region, and this war of attrition game, which determines the timing of price hikes, is altered by the timing law. This explains the drastic change in price leadership patterns.

Figure 4 in Section III.B suggests that the level of retail price varies with the level of import cost. Import cost affects the floor against which the retail price cycles bounce. The floor, in theory, is the marginal cost

TABLE 2
FREQUENCY DISTRIBUTION OF CYCLE LENGTH UNDER THE LAW

	DAY																
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	1-16
Cycles BP leads alone			1		1	4	5	4	1	4	6	1					27
Cycles Caltex leads alone						1	7	10	3	7	2	3	1	1	1	1	37
Cycles Shell leads alone					1	1	1	2	7	1			2				15
Cycles multiple firms lead					1	1	10	3	3	4	1						23
All cycles			1		3	7	23	19	14	16	9	4	3	1	1	1	102

of supplying gasoline. Since variations in import cost track the changes in marginal cost of supplying gasoline, we can study how the variations in import cost affect the retail price cycle dynamics.

The basic logic of Maskin and Tirole's (1988) model implies that cost changes can have subtle impacts on retail price. The core feature of the Edgeworth price cycle equilibrium is that price increases quickly and decreases gradually, and a war of attrition problem is embedded at the bottom of the price cycles. In their model, the falling phase of a cycle ends when the gasoline margin is zero or when a war of attrition starts. This implies that gasoline price and marginal cost at the end of a price cycle should be close. Thus, if marginal cost has decreased since the start of a price cycle, it would require a greater reduction of, and take a longer time for, retail price to be close to marginal cost again. This implies that cost changes between the start and end of a cycle affect the amplitude and length of this cycle.

The amplitude of a price cycle can be measured by the height of either the rising phase or the falling phase. If marginal cost is constant over time, these two measures yield the same result; if marginal cost changes over time, they yield different results. Define the amplitude of the rising phase as the market average price at the top of the cycle minus the price at the end of the previous cycle, and define the amplitude of the falling phase as the price at the top of the cycle minus the price at the end of this cycle. Thus, the difference between the two cycle amplitude measures—the height of the rising phase minus the height of the falling phase—is the same as the market price at the end of the current cycle minus the market price at the end of the previous cycle. The logic of the price cycle equilibrium thus implies the following testable implication: the difference between the two cycle amplitude measures varies positively with cost changes (the cost at the end of a cycle minus the cost at the end of the previous cycle).

Figure 7 confirms this implication. The sample period in figure 7 is January 1 through October 31, 2003 (for which the import cost series is available). If the difference between the two measures of cycle amplitude is regressed on a constant term and cost change, the coefficient on cost change is 0.83, with a standard error of 0.10.

Cost change also affects the length of the price cycles (defined as the number of days between two lead price hikes) since it takes longer for the retail price to fall close to the marginal cost again if the marginal cost itself has decreased since the start of the cycle. This logic suggests another testable implication: the length of a cycle tends to be longer if the cost at the end of the cycle is lower than the cost at the start of the cycle.

Figure 8 shows the relationship between cycle length and cost change over the postlaw period May 15, 2001, through October 21, 2003. The

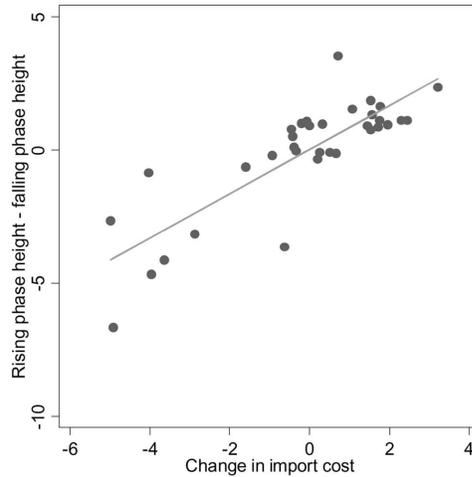


FIG. 7.—The impact of cost changes on cycle amplitude, January 1 to October 31, 2003. The x axis is the import cost at the end of a cycle minus the import cost at the end of the previous cycle. The y axis is the height of the rising phase of a cycle minus the height of the falling phase of the same cycle. This is the same as the market average price at the end of a cycle minus the market price at the end of the previous cycle.

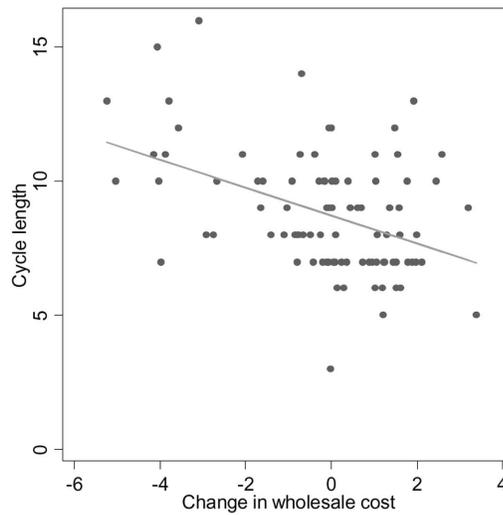


FIG. 8.—The impact of cost changes on cycle length, May 15, 2001, to October 21, 2003. The x axis is the wholesale price paid by the BP retailer at the end of a cycle minus that at the beginning of the cycle. The length of a cycle, the y axis, is the number of days from the start through the end of a cycle.

cost change is the wholesale price paid by the BP retailer at the end of a cycle minus the wholesale price paid at the beginning of the cycle. If cycle length is regressed on a constant and cost change, the cost change coefficient is -0.52 , with a standard error of 0.12 . If cost change is measured by the wholesale price paid by the small independent retailer, the cost change coefficient is -0.32 , with a standard error of 0.11 . If the sample period is restricted to January 1 through October 21, 2003, for which the import cost series is available, the coefficient for (import) cost change is negative (-0.25), though not statistically significant (p -value is 0.16).

V. Mixed Strategies in Price Leadership

The results in the previous section suggest that the Maskin-Tirole price cycle equilibrium captures the fact that a war of attrition game is at the bottom of each of the gasoline price cycles. In this section, I present evidence that the observed price leadership outcomes under the law are better characterized by mixed strategy plays than by two alternative hypotheses: (1) the firms play a pure strategy in each war of attrition or (2) they simply alternate as price leader over the wars.

In Maskin and Tirole's (1988) model, firms always move sequentially and play the same mixed strategies in each war of attrition. This implies that there is always a single leader and the leadership outcome does not exhibit any serial correlation. The postwar price leadership patterns shown in figure 6 clearly reject these implications. Two or three firms may relent simultaneously, and a firm tends to be less likely to lead again if it led in the previous war. This is not surprising given that (1) the timing of the war of attrition under the law, as argued in Section II, is simultaneous instead of sequential, and (2) the gasoline firms have the incentive to coordinate over the wars of attrition. My point is that once these two factors are taken into consideration, the mixed strategy presumed by Maskin and Tirole can generate the postwar price leadership patterns.

I hypothesize that the timing of the attrition wars is simultaneous and that the firms play the mixed strategy presumed in Maskin and Tirole's model (firm i always relents with probability p_i on date t conditional on no firm having relented before then) within each war, but the probability with which each firm relents is affected by the outcome of the previous war.

This hypothesis implies that the leadership outcome of each war, a random realization of mixed strategy play, must be random once conditional on the outcome of the previous war. In addition, the presumed mixed strategies impose strong restrictions on the leadership types and their frequencies. I consider these restrictions in subsection A and the

alternative hypotheses in subsection B. Here I present evidence that the outcome of this war, once conditional on the outcome of the previous war, is an independent draw from a random process.

First note from figure 6 that the leadership outcome of this war, conditional on the outcome of the previous war, is far from deterministic. For example, the leadership outcome of the 27 wars immediately preceded by a war in which BP led alone is that (1) BP leads alone twice, (2) Caltex leads alone 13 times, (3) Shell leads alone seven times, and (4) multiple firms lead simultaneously five times. To test conditional serial independence, I separate the wars of attrition in the sample into four subsamples or four types according to the leadership outcome of the previous war: (1) BP alone led (27 wars), (2) Caltex alone led (37 wars), (3) Shell alone led (15 wars), and (4) multiple firms led together (22 wars). I use the nonparametric run test (see, e.g., Gibbons and Chakraborti 2003), which has been used in the recent literature that tests mixed strategies in sports games. The run test of serial independence is based on the number of runs in a sequence.²⁴ A small (large) number of runs indicates positive (negative) serial correlation.

Let binary variable l_i^k equal one if event i is true for the k th war in a sample and zero otherwise. Event i indicates each of the seven leadership types or whether a firm (BP, Caltex, or Shell) is a leader. For example, $l_{BP}^k = 1$ if BP is a leader of the k th war of a sample, and $l_{BP\text{alone}}^k = 1$ if BP alone is the leader of the k th war of a sample. Consider the subsample of wars of attrition that are preceded by a war in which BP led alone. Then, the Bernoulli sequence $l_i = \{l_i^1, l_i^2, \dots, l_i^{27}\}$ denotes if event i is true in each of the 27 wars in this subsample. Other Bernoulli leadership sequences can be similarly defined for this subsample and other samples.

Let r_i be the number of runs in the Bernoulli sequence $l_i = \{l_i^1, l_i^2, \dots, l_i^K\}$, where K is the number of wars in a sample, and let $L_i = \sum_{k=1}^K l_i^k$ be the number of successes in this sequence. The expected number of runs in the sequence under the null of serial independence is $\mu_r = 2L_i(K - L_i)/K + 1$, and the variance is $\sigma_r^2 = 2L_i(K - L_i)[2L_i(K - L_i) - K]/[K^2(K - 1)]$. The test statistic, $z = (r - \mu_r)/\sigma_r$, is approximately distributed as the standard normal distribution.

When the run test is applied to the sequences $l_i = \{l_i^1, l_i^2, \dots, l_i^{102}\}$ of the entire sample, the null of serial independence is rejected at normal significance levels for the sequences of l_{BP} , l_{Caltex} , $l_{BP\text{alone}}$, and $l_{Caltex\text{alone}}$, which is consistent with the observation that BP or Caltex is less likely to lead again if it led in the previous war. Table 3 reports the test results for the four subsamples. The null of serial independence cannot be rejected at the 5 percent level for any of the 18 leadership sequences

²⁴ A run is a sequence of identical symbols. For example, the sequence (1, 1, 0, 1) has three runs.

TABLE 3
 RUN TESTS OF SERIAL INDEPENDENCE WITHIN WARS OF ATTRITION OF THE SAME TYPE

	Number of 1's	Number of Runs	z-Statistic	p-Value	Number of 1's	Number of Runs	z-Statistic	p-Value
BP Leads Alone Previously: 27 Wars				Caltex Leads Alone Previously: 37 Wars				
BP is a leader	6	13	1.54	.12	30	12	-.19	.85
Caltex is a leader	15	13	-.53	.60	7	9	-1.86	.06
Shell is a leader	12	13	-.53	.60	10	19	1.45	.15
BP leads alone	2	5	.48	.63	22	17	-.64	.52
Caltex leads alone	13	15	.20	.84	2	5	.40	.69
Shell leads alone	7	9	-1.23	.22	5	10	.26	.80
BP and Caltex lead	0	1	3	7	.59	.56
BP and Shell lead	3	7	.71	.48	3	6	-.62	.54
Caltex and Shell lead	1	3	.28	.78	0	1
BP, Caltex, and Shell lead	1	3	.28	.78	2	5	.40	.69
Shell Leads Alone Previously: 15 Wars				Multiple Firms Lead Previously: 22 Wars				
BP is a leader	4	5	-1.31	.19	9	13	.62	.54
Caltex is a leader	13	4	-.60	.55	17	6	-1.73	.08
Shell is a leader	1	2	-2.55	.01	6	6	-2.08	.04
BP leads alone	1	3	.39	.69	2	5	.54	.59
Caltex leads alone	11	5	-1.31	.19	11	10	-.87	.38
Shell leads alone	0	1	2	4	-.94	.35
BP and Caltex lead	2	4	.60	.55	3	7	.81	.42
BP and Shell lead	1	2	-2.55	.01	1	3	.32	.75
Caltex and Shell lead	0	1	0	1
BP, Caltex, and Shell lead	0	1	3	7	.81	.42

NOTE.—The reported number of runs is the observed ones.

with at least one success if BP or Caltex led alone in the previous war. For the other two war types, the null cannot be rejected at the 5 percent level for 13 out of the 16 leadership sequences with at least one success, and none of the sequences are rejected at the 1 percent level. These results suggest that the leadership outcome of wars of attrition of the same type is serially independent.

A. *The Multinomial Distribution Test*

In this subsection, I derive and test the implications of the mixed strategy hypothesis on the leadership types and their frequencies. Note that only three firms (A , B , and C) potentially play the presumed mixed strategies since the three largest firms lead all the price cycles under the law. Qualitatively, the presumed mixed strategies imply that one of the following seven mutually exclusive and exhaustive leadership types should arise in each war: (1) firm A leads alone, (2) firm B leads alone, (3) firm C leads alone, (4) firms A and B lead together, (5) firms A and C lead together, (6) firms B and C lead together, and (7) firms A , B , and C all lead together. On a particular day of a war of attrition, it may be the case that none of the three firms relent. If such an event happens, the firms would continue to play the same mixed strategies until at least one firm relents. These seven price leadership types are precisely what we observe in Section IV.B.3.

Quantitatively, the presumed mixed strategies imply specific multinomial distributions over the frequencies of the seven outcomes. Each war can be viewed as a random multinomial experiment with seven outcomes. Let firm i play relent with probability p_i , $i = A, B, C$. Then, the mixed strategy hypothesis implies that the multinomial distribution over the seven leadership types must be specified by three probability parameters (p_A , p_B , and p_C). For example, the probability with which firm A relents alone is

$$\begin{aligned} & p_A(1 - p_B)(1 - p_C) + p_A(1 - p_B)(1 - p_C) \times [(1 - p_A)(1 - p_B)(1 - p_C)] \\ & + \dots + p_A(1 - p_B)(1 - p_C) \times [(1 - p_A)(1 - p_B)(1 - p_C)]^t + \dots \\ & = \frac{p_A(1 - p_B)(1 - p_C)}{1 - (1 - p_A)(1 - p_B)(1 - p_C)}, \end{aligned} \quad (1)$$

where $[(1 - p_A)(1 - p_B)(1 - p_C)]^t$ is the probability that none of the three firms has relented up to period t of a war. Similarly, the probabilities with which the other six leadership types arise are

$$\frac{(1 - p_A)p_B(1 - p_C)}{1 - (1 - p_A)(1 - p_B)(1 - p_C)}, \quad (2)$$

$$\frac{(1 - p_A)(1 - p_B)p_C}{1 - (1 - p_A)(1 - p_B)(1 - p_C)}, \quad (3)$$

$$\frac{p_A p_B (1 - p_C)}{1 - (1 - p_A)(1 - p_B)(1 - p_C)}, \quad (4)$$

$$\frac{p_A(1 - p_B)p_C}{1 - (1 - p_A)(1 - p_B)(1 - p_C)}, \quad (5)$$

$$\frac{(1 - p_A)p_B p_C}{1 - (1 - p_A)(1 - p_B)(1 - p_C)}, \quad (6)$$

and

$$\frac{p_A p_B p_C}{1 - (1 - p_A)(1 - p_B)(1 - p_C)}. \quad (7)$$

While a typical multinomial distribution with seven outcomes has six free parameters, the multinomial distribution predicted by the presumed mixed strategy equilibrium must be specified by three parameters p_A , p_B , and p_C only. That is, the mixed strategy equilibrium imposes quite strong restrictions on the multinomial distribution. The Pearson χ^2 goodness-of-fit statistic can then be used to test if the observed leadership frequency distribution is the same as the distribution predicted by the presumed mixed strategy equilibrium. This test must be applied separately to the various types of wars of attrition to satisfy the assumption of serial independence.

To implement this test, I first estimate the three parameters by the maximum likelihood method for each of the four types of wars of attrition. Suppose that in a sample of wars of attrition of the same type, the observed frequencies of the seven leadership types are n_i , $i = 1, \dots, 7$. Let $n = \sum_{i=1}^7 n_i$. Then, the likelihood function for the multinomial distribution is

$$l(p_A, p_B, p_C; n, n_i) = \left(\frac{n!}{n_1! n_2! \cdots n_7!} \right) \left[\frac{p_A(1 - p_B)(1 - p_C)}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} \right]^{n_1} \cdots \left[\frac{p_A p_B p_C}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} \right]^{n_7}.$$

The three parameters (p_A , p_B , and p_C) can be obtained by solving the three first-order conditions:

$$\frac{n_1 + n_4 + n_5 + n_7}{p_A} - \frac{n_2 + n_3 + n_6}{1 - p_A} - \frac{n(1 - p_B)(1 - p_C)}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} = 0, \quad (8)$$

$$\frac{n_2 + n_4 + n_6 + n_7}{p_B} - \frac{n_1 + n_3 + n_5}{1 - p_B} - \frac{n(1 - p_A)(1 - p_C)}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} = 0, \quad (9)$$

$$\frac{n_3 + n_5 + n_6 + n_7}{p_C} - \frac{n_1 + n_2 + n_4}{1 - p_C} - \frac{n(1 - p_A)(1 - p_B)}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} = 0. \quad (10)$$

By substituting the estimated parameters (\hat{p}_A , \hat{p}_B , and \hat{p}_C) into equations (1)–(7), we can obtain the predicted frequencies of the seven price leadership types (\hat{n}_i , $i = 1, \dots, 7$). The Pearson χ^2 goodness-of-fit statistic is $\sum_{i=1}^7 (n_i - \hat{n}_i)^2 / \hat{n}_i$. We observe seven alternative outcomes but estimate only three parameters; thus the test has three degrees of freedom.

This multinomial distribution test is similar in spirit to the method used in the recent literature to test the indifference property of mixed strategies in sports games. Walker and Wooders (2001), for example, estimate the winning probabilities of two discrete actions (e.g., serving the tennis ball to the right or left), which are taken as the parameters of two binomial distributions, and then use the χ^2 statistic (with one degree of freedom) to test whether the two binomial distributions are the same. Unfortunately, it is not feasible to test the indifference property in a market setting because the value of an action (relent or fight) to a firm is not observable.²⁵

To better understand the test, it is useful to first consider the following proposition. It can be easily shown that the maximum likelihood esti-

²⁵ The value of an action on a particular day of a war of attrition is characterized by a dynamic equation. To test the indifference property, we have to estimate the value of an action. To do so, we have to observe at least (1) the date on which a war of attrition starts and (2) the firms' *daily* economic profit over an extended period of time. Unfortunately, we observe neither. The advantage of this paper's empirical strategy is that it relies only on the price leadership outcome, which is observable.

mates of the three parameters are identical to the solutions to the following three equations:

$$\frac{p_A}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} = \frac{n_1 + n_4 + n_5 + n_7}{n}, \quad (11)$$

$$\frac{p_B}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} = \frac{n_2 + n_4 + n_6 + n_7}{n}, \quad (12)$$

$$\frac{p_C}{1 - (1 - p_A)(1 - p_B)(1 - p_C)} = \frac{n_3 + n_5 + n_6 + n_7}{n}. \quad (13)$$

The left-hand side of equation (11) is the probability predicted by the presumed mixed strategy equilibrium that firm *A* is a price leader (alone or with other firms), and the right-hand side is the observed probability that firm *A* is a price leader. Equations (12) and (13) can be similarly interpreted. The first useful observation is that the ratio of the probabilities with which two firms play relent is the same as the ratio of the frequencies with which they are a leader. A firm is less likely than another firm to be a leader in a war because it relents with a smaller probability. Because of this property, these three equations boil down to one nonlinear equation and two linear ones. Second, note that the information used to estimate the three parameters is *four frequencies*: the total number of wars and the frequencies with which firms *A*, *B*, and *C* are each a leader. The leadership multinomial distribution is characterized by *seven frequencies*. This is why the χ^2 goodness-of-fit test has three degrees of freedom.

It is also useful to recognize that this test is a joint test of the presumed equilibrium mixed strategies and other implicit assumptions. The timing regulation and the special properties of the war of attrition game embedded at the bottom of the price cycle equilibrium make it possible to test the mixed strategy hypothesis in a market setting. However, oligopoly behavior in a market is quite different from athletes' play inside an arena: markets, and consequently firms' pricing decisions, are subject to significant and stochastic shocks. In the gasoline market studied here, the significant changes in cost may affect the firms' assessment of their expected profits from relenting or not, thus affecting the firms' relenting decision over time. In implementing the distribution test, I have implicitly made the strong assumption that market shocks do not alter firms' mixing behavior at all.

Table 4 reports by war type the estimated probabilities with which BP, Caltex, and Shell each play relent. Table 5 reports by war type the expected and observed frequencies of each leadership type and the

TABLE 4
ESTIMATED PROBABILITIES WITH WHICH A FIRM PLAYS RELENT
BY WAR TYPE

	PROBABILITY PLAYS RELENT		
	BP	Caltex	Shell
BP alone leads previously	.1118	.2796	.2237
Caltex alone leads previously	.5543	.1293	.1848
Shell alone leads previously	.1800	.5850	.0450
Multiple firms lead previously	.3264	.6166	.2176

χ^2 distribution test results. The null hypothesis for two types of wars (Caltex led alone or multiple firms led in the previous war) cannot be rejected at the conventional significance levels. The null hypothesis for the other two types of wars is rejected at the 5 percent level. The distribution test is stringent in the sense that slight changes in the observed frequencies could change the χ^2 statistic substantially. For example, the war type for which Shell led alone in the previous war has a χ^2 statistic of 12.70. If the observed frequency of BP and Shell leading together were reduced by one and the frequencies of BP or Shell each leading alone were increased by one, the predicted distribution would not change since the right-hand sides of equations (11)–(13) do not change, but the χ^2 statistic would be reduced to 2.1.

B. *Alternative Hypotheses*

My results suggest that the hypothesized mixed strategy plays capture reasonably well both the randomness in the leadership outcome of the individual wars and the stochastic regularities of the leadership outcome over the wars. This subsection considers two alternative hypotheses.

One alternative hypothesis is that the three largest gasoline firms simply alternate in being a price leader. It is true that a firm is less likely to lead again in this war if it led alone in the previous war. However, we also find that the outcome of this war, conditional on the outcome of the previous war, is stochastic, and this is inconsistent with simple alternating behavior.

Another alternative hypothesis is that all the observed leadership outcomes result from pure strategy equilibria. It is not clear how this hypothesis explains both the randomness in the outcome of the individual wars and the stochastic regularities of the leadership outcome over the wars. In addition, pure strategy equilibria in wars of attrition require at least one firm to relent immediately. This implies that the duration of

TABLE 5
FREQUENCY DISTRIBUTION OF LEADERSHIP TYPES BY WAR TYPE

IN THE CURRENT WAR	IN THE PREVIOUS WAR							
	BP Leads Alone		Caltex Leads Alone		Shell Leads Alone		Multiple Firms Lead	
	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed
BP leads alone	3.36	2	21.29	22	1.59	1	2.70	2
Caltex leads alone	10.34	13	2.54	2	10.18	11	8.96	11
Shell leads alone	7.68	7	3.88	5	.34	0	1.55	2
BP and Caltex lead	1.30	0	3.16	3	2.23	2	4.34	3
BP and Shell lead	.97	3	4.83	3	.07	1	.75	1
Caltex and Shell lead	2.98	1	.58	0	.48	0	2.49	0
BP, Caltex, and Shell lead	.38	1	.72	2	.11	0	1.21	3
Total frequency	27	27	37	37	15	15	22	22
χ^2 statistic	9.22		4.03		12.70		6.43	

NOTE.—The critical values of the χ^2 distribution with three degrees of freedom are 6.25 and 7.82 at the 10 percent and 5 percent significance levels, respectively.

TABLE 6
EXPECTED DURATION OF WARS OF ATTRITION BY TYPE

	IN THE PREVIOUS WAR			
	BP Leads Alone	Caltex Leads Alone	Shell Leads Alone	Multiple Firms Lead
Mean duration (days)	1.99	1.46	1.48	1.25
Probability that duration is 1 day	.50	.68	.68	.80
Probability that duration is 2 days or shorter	.75	.90	.89	.96
Probability that duration is 3 days or shorter	.88	.97	.97	.99

each of the 102 wars should be exactly 1 day, which appears implausible. It is inconsistent with the observation that a small number of wars appear to last longer than 2 or 3 days (see the last war of attrition in fig. 2 and App. fig. C1). The presumed mixed strategies, on the other hand, have plausible predictions about the duration of the wars of attrition. The probability that a war of attrition lasts exactly T days has a geometric distribution,

$$[(1 - p_A)(1 - p_B)(1 - p_C)]^{T-1} [1 - (1 - p_A)(1 - p_B)(1 - p_C)].$$

Therefore, the expected mean duration is

$$\frac{1}{1 - (1 - p_A)(1 - p_B)(1 - p_C)},$$

and the probability the war lasts T days or shorter is

$$1 - [(1 - p_A)(1 - p_B)(1 - p_C)]^T.$$

A war of attrition starts when the retail price is near or at the competitive level, but the start date of the wars of attrition is not directly observable. For this reason, I cannot formally test these war duration predictions. However, I do know that a small number of wars appear to last longer than 3 days, and the duration of a war should be smaller than the duration of a price cycle (because a war of attrition is only part of a price cycle). These observations allow for checking whether the presumed equilibrium mixed strategies are reasonable from a perspective other than the leadership outcome.

Table 6 reports the predicted mean duration of each type of war and the probabilities that each war type lasts 1, 2, or 3 days or shorter. The predicted mean duration ranges from 1.25 days to 1.99 days, indeed much smaller than the average length of the price cycles. The predicted probabilities also suggest that relenting should occur immediately in about 50–80 percent of the wars of attrition, and wars of attrition that last longer than 3 days should occur, but very infrequently.

VI. Welfare Impact of the Timing Law

To evaluate the welfare effect of the timing law, I compare the Perth market average retail gasoline price with three control variables before and after the law took effect. The first two controls are the weekly market average retail gasoline prices in Sydney and Melbourne, the two largest cities in Australia, where a regular gasoline price cycle exists but the timing law does not. The third control is an import parity indicator constructed by the Australian Competition and Consumer Commission according to the IPP principle discussed in the data section.²⁶ All three controls are available for the period July 1, 2000, through December 31, 2001.²⁷ Figure 9 shows the market average prices in Perth, Melbourne, and Sydney, as well as the import parity indicator, all in a weekly frequency.²⁸

The most significant feature to note in figure 9 is that the Perth market average retail price is significantly below the three control variables in the first 3 weeks of the law, which was not the case before the law.²⁹ It is also visible that in the first 4 months of the law, the Perth market average prices were consistently below or near the minimum of the Melbourne and Sydney market average prices, which has not been the case since May 2001, when the regular price cycles reappeared. These observations are consistent with the fact that the timing law disrupted gasoline firms' pricing coordination in the short run.

A simple before and after comparison indicates that in the first 4 months of the law, the treatment effect of the timing law is -1.8 cents per liter, relative to the average of the three control variables. However, if we restrict the treatment period to May through December 2001, the same before and after comparison indicates that the treatment effect of the law is 0.09 cent per liter. I also estimated a standard fixed-effect difference-in-difference model and found similar results. In particular,

²⁶ This import parity indicator differs from the previously used import cost series in two components: it is not adjusted for quality premium, and it includes marketing and distribution costs as well as a return on assets. For details of the pricing formula, see Australian Competition and Consumer Commission (2002, 27).

²⁷ All three controls were published in Australian Competition and Consumer Commission (2002).

²⁸ The Melbourne and Sydney retail prices were published weekly, and the import parity indicator was published daily. I convert the daily import parity indicator and the daily and hourly Perth retail prices into weekly sequences according to the week definition used to report the Melbourne and Sydney retail price series. The fuel standard in Western Australia was higher than that in Sydney and Melbourne by 0.35 cent per liter in 2000 and by 0.85 cent per liter in 2001. All three controls shown in fig. 9 are adjusted by this change in fuel quality premium. The three controls are also available for the 9-month period January 1 through September 30, 2002. I do not consider this period because I do not know the precise quality premium for this period.

²⁹ For the three control variables, the last data point before the law was for the 7 days starting on December 16, 2000, and for the Perth retail price, the last data point was for the 5 days starting on December 16, 2005.

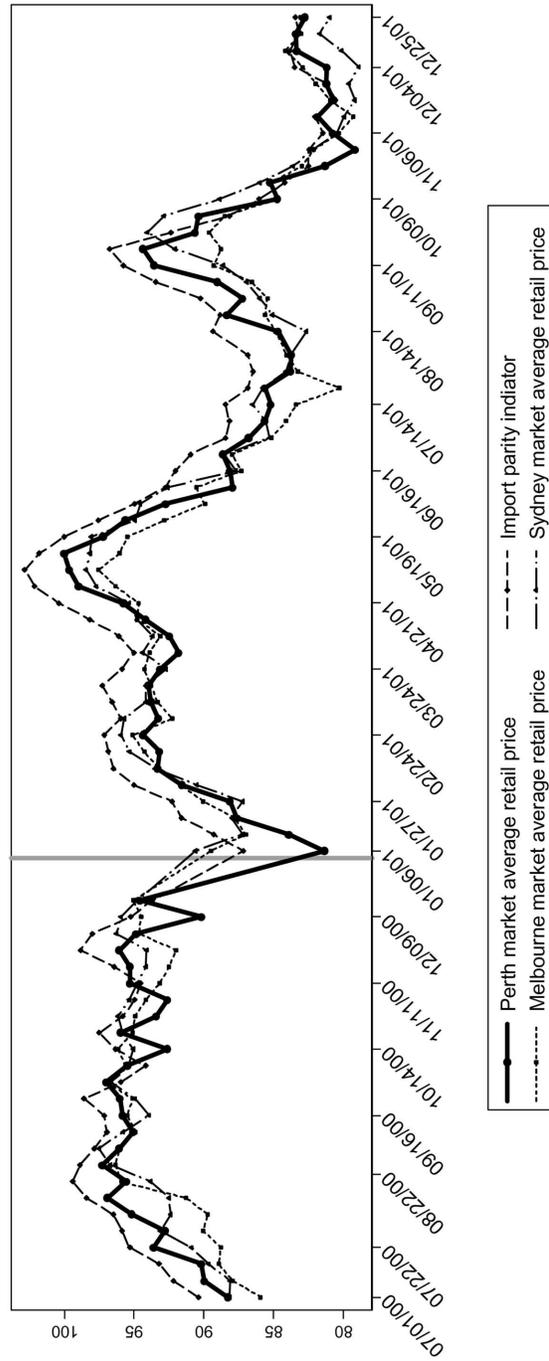


Fig. 9.—Weekly Perth, Melbourne, and Sydney market average prices and an import parity indicator

the treatment effect is statistically significant for the first 4 months of the law but statistically insignificant since then. These results are very robust. They are not sensitive to the cutoff week when retail price cycles are thought to reappear under the law, nor are they sensitive to whether the import parity indicator is the only control or the Sydney and Melbourne prices are the only controls. Thus, the evidence suggests that the timing law lowered the Perth market average retail price in the first 4 months of the law but did not have a statistically significant effect once the gasoline firms succeeded in reestablishing the price cycle equilibrium.

VII. Conclusion

This paper has studied oligopoly firms' pricing strategies in a gasoline market before and under a law that regulates firms' timing and frequency of price changes. I find that the gasoline pricing behavior, both before and under the law, is well captured by the pricing strategies used in Maskin and Tirole's (1988) model to generate the Edgeworth price cycle equilibrium. My results thus highlight the importance of short-run price commitment in tacit collusion. Since price rigidity exists in many industries, the results suggest that the idea of reaction based on short-run price commitment may be applicable in many other settings as well. Regular price cycles have been observed only in gasoline markets, but Maskin and Tirole's model also predicts the kinked demand curve equilibrium.

This is the first paper in the literature that tests the concept of mixed strategy in a market setting. I find that the price leadership outcomes observed in the Perth market under the law, though exhibiting first-order serial correlation, are better captured by mixed strategies than by alternative hypotheses. The firms have the incentive not to randomize, but under the timing law they are naturally uncertain about rivals' action at the bottom of each price cycle. Mixed strategies capture players' uncertainty of each other's pure actions. My results also suggest that mixed strategy can be a useful technical device even if agents do not actually play mixed strategies. The gasoline firms did not play mixed strategies before the law, but the Edgeworth price cycle equilibrium captures key regularities of the gasoline price cycles before the law.

My results suggest that tacit collusion is inevitable if the market condition is close to the Bertrand environment of the Maskin-Tirole model: if price is close to marginal cost, a firm can always increase price dramatically and induce rivals to charge a high price for at least a short period of time. This reinforces the challenge that tacit collusion poses for antitrust policy. My evidence shows unambiguously that the gasoline firms in the Perth market engaged in "conscious parallel pricing" or

tacit pricing coordination. My results suggest that some forms of vertical control are needed to generate the price cycles. Without vertical control, intrabrand synchronization and uniformity in price increases cannot exist. Why do the oil firms in the Perth market get involved in retail price coordination? Is the theory of Telser (1960)—upstream firms coordinate on retail price rather than on wholesale price because the latter is subject to secret price cuts—relevant in the Perth market? Should there be any policy response to the tacit pricing coordination documented in this paper? These questions are easy to ask but difficult to answer. The law and economics literature is far from a consensus on the appropriate antitrust policy with regard to tacit collusion.³⁰ Some advocate that tacit collusion should be subject to antitrust penalty (e.g., Posner 2001), whereas others argue that it is best left out of antitrust enforcement (e.g., Turner 1962).

It is also interesting to note that the ex ante timing structure of the postlaw Perth gasoline market matches exactly the timing structure of the standard supergame model: a simultaneous-move Bertrand pricing game is discretely repeated once a day for an indefinite number of days. How can the gasoline firms continue to increase price sequentially ex post and coordinate on the price cycle equilibrium under the law? Maskin and Tirole (1988, sec. 9) show that the price cycle equilibrium is robust to endogenous timing under the standard assumption that short-run price commitment exists exogenously. However, no exogenous factors appear to exist to prevent the gasoline firms from changing price every 24 hours. Where then does short-run price commitment under the law, the basis for sequential price increases, originate? One possibility is that commitment arises endogenously. Only recently did economists start to model commitment as an endogenous decision (e.g., Lau 2001; Henkel 2002; Caruana and Einav 2008). Further theoretical work on endogenous commitment is clearly needed.

³⁰ Even the meaning of horizontal agreement in antitrust is subject to debate (e.g., Kaplow 2009).

Appendix A

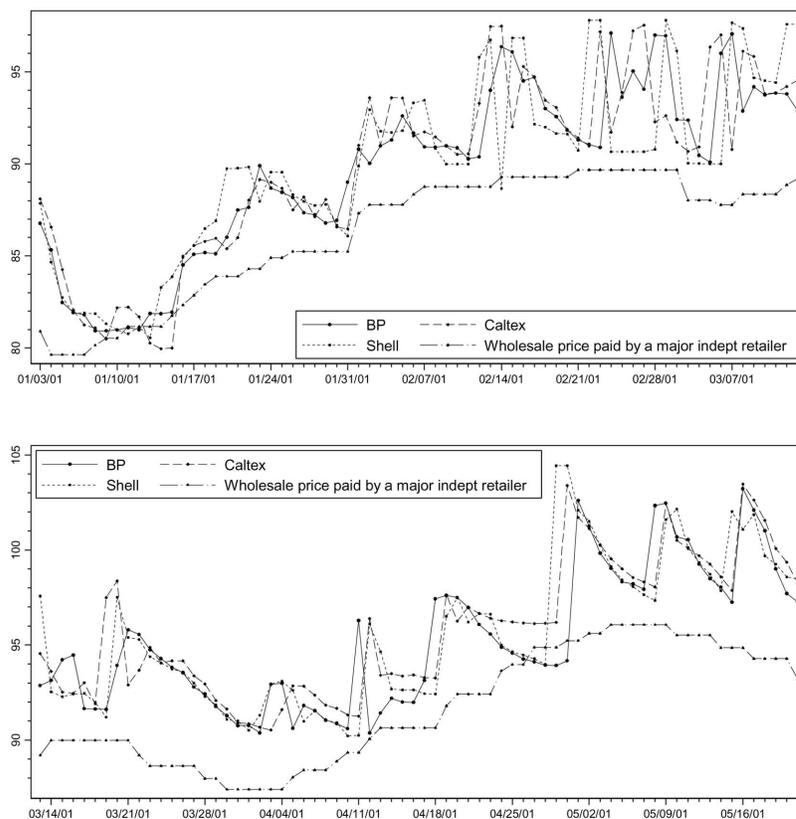


FIG. A1.—Daily brand average retail price of three firms and the wholesale price during and shortly after the adjustment period: January 3 to May 21, 2001. During the adjustment period, the significant initiating price hikes took effect on February 12 (Shell), February 22 (Shell), March 5 (Caltex), March 19 (Caltex), April 10 (BP), and April 18 (BP). In all cases but the BP price hike on April 10, the initiating firm kept its price increase the following day.

Appendix B

The Conditional Price Support Mechanism

This appendix documents the conditional price support mechanism used by Caltex to control the retail price of the many Caltex franchisees during the sample period of this study. Caltex was prohibited from using multisite franchise agreements by a constraint imposed by the Australian federal government when Caltex merged with another petroleum firm in 1995.

To see the mechanism, suppose that the competitors of a Caltex franchisee are charging the retail price of 95 cents per liter, but the Caltex franchisee has the contractual obligation to pay Caltex a wholesale price of \$1.00 per liter! To

survive, this franchisee has to receive a wholesale price reduction from Caltex. Caltex provides such a price reduction or price support (e.g., 8 cents per liter) under the condition that the franchisee agrees to a maximum resale price (e.g., 95 cents). Through this mechanism, Caltex forces this franchisee to set a price between 92 cents, the effective wholesale price, and 95 cents per liter. Most franchisees in such a situation would set their retail prices at 95 cents, thus earning a retail margin of 3 cents per liter, though some aggressive sites may set a price a bit lower than 95 cents. By changing the size of the price support and the maximum retail price allowed, Caltex is able to effectively control the price movement of its franchisees. A retail price increase occurs when Caltex decides to provide no price support and suggests a retail price to be set by Caltex franchisees. Most franchises set their price at the suggested level, thus leading to the synchronization and uniformity of Caltex retail price increases. During the cycle falling phase, the size of price supports and the maximum retail prices allowed may vary from site to site.

A Caltex branded retail site receives two electronic messages from Caltex each day that informs the retail site of the contract price, the size of the price support, and the maximum retail price allowed for the following day. Below are two messages received by a Caltex franchisee on May 7, 2004, regarding the price of regular unleaded gasoline.

91.25 08/05 6:00 TEMP REDUCTION OFF CONTRACT BUY PRICE
IN RESPONSE TO COMPETITOR PRICES CONDITIONAL ON MAX
RESALE PRICE NOT EXCEEDING 93.95 UNLEAD

EFF 08/05/04 0600 HRS DEALER METER MAP PRICE FOR UN-
LEADED = 91.25 CPL, ADJUSTMENT = 12.05 CPL.

To protect the identity of the retail site, the prices shown in these two messages are slightly altered. The retail site was instructed that the effective wholesale price on May 8, 2004, starting from 6:00 a.m. was 91.25 cents per liter (CPL), and the retailer had to set a retail price lower than 93.95 cents in order to receive a price support of 12.05 cents. Otherwise, the retailer had to pay the "contract buy price" of 103.30 (= 91.25 + 12.05) cents. In order to receive the price support, the retailer had to set a price between 91.25 and 93.95 cents.

Appendix C

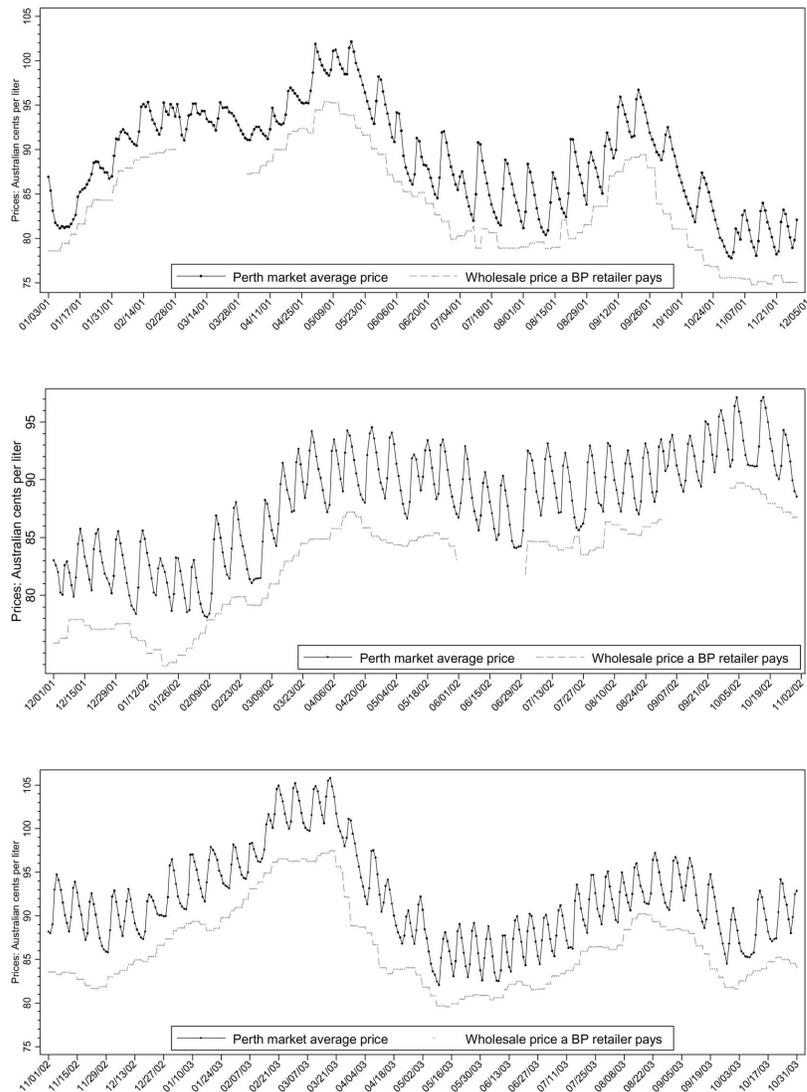


FIG. C1.—Daily market average retail gasoline price in Perth, January 3, 2001, to October 31, 2003.

References

Aumann, Robert J. 1987. "Correlated Equilibrium as an Expression of Bayesian Rationality." *Econometrica* 55 (January): 1–18.

- Australian Competition and Consumer Commission. 2002. "Terminal Gate Pricing Arrangements in Australia and Other Fuel Pricing Arrangements in Western Australia." <http://www.accc.gov.au/content/index.phtml/itemId/321849>.
- . 2007. "Petrol Prices and Australian Consumers—Report of the ACCC Inquiry into the Price of Unleaded Petrol." <http://www.accc.gov.au/content/index.phtml/itemId/806216>.
- Caruana, Guillermo, and Liran Einav. 2008. "A Theory of Endogenous Commitment." *Rev. Econ. Studies* 75 (January): 99–116.
- Castanias, Rick, and Herb Johnson. 1993. "Gas Wars: Retail Gasoline Price Fluctuations." *Rev. Econ. and Statis.* 75 (February): 171–74.
- Chiappori, Pierre-André, Steven Levitt, and Timothy Groseclose. 2002. "Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer." *A.E.R.* 92 (September): 1138–51.
- Eckert, Andrew, and Douglass West. 2004. "Retail Gasoline Price Cycles across Spatially Dispersed Gasoline Stations." *J. Law and Econ.* 47 (April): 245–71.
- Genesove, David, and Wallace Mullin. 2001. "Rules, Communication, and Collusion: Narrative Evidence from the Sugar Institute Case." *A.E.R.* 91 (June): 379–98.
- Gibbons, Jean, and Subhabrata Chakraborti. 2003. *Nonparametric Statistical Inference*. New York: Dekker.
- Harsanyi, John C. 1973. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points." *Internat. J. Game Theory* 2 (December): 1–23.
- Hendricks, Kenneth, and Robert H. Porter. 1988. "An Empirical Study of an Auction with Asymmetric Information." *A.E.R.* 78 (December): 865–83.
- Henkel, Joachim. 2002. "The 1.5th Mover Advantage." *Rand. J. Econ.* 33 (Spring): 156–70.
- Hsu, Shih-Hsun, Chen-Ying Huang, and Cheng-Tao Tang. 2007. "Minimax Play at Wimbledon: Comment." *A.E.R.* 97 (March): 517–23.
- Kaplow, Louis. 2009. "On the Meaning of Horizontal Agreements in Antitrust." Manuscript, Law School, Harvard Univ.
- Lau, Sau-Him P. 2001. "Aggregate Pattern of Time-Dependent Adjustment Rules, II: Strategic Complementarity and Endogenous Nonsynchronization." *J. Econ. Theory* 98 (June): 199–231.
- Lewis, Matthew. 2009. "Temporary Wholesale Gasoline Price Spikes Have Long-Lasting Retail Effects: The Aftermath of Hurricane Rita." *J. Law and Econ.* 52 (August): 581–605.
- Mailath, George J., and Larry Samuelson. 2006. *Repeated Games and Reputations: Long-Run Relationships*. Oxford: Oxford Univ. Press.
- Maskin, Eric, and Jean Tirole. 1988. "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves and Edgeworth Cycles." *Econometrica* 56 (May): 571–99.
- Noel, Michael. 2007. "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market." *J. Indus. Econ.* 55 (April): 69–92.
- . 2008. "Edgeworth Price Cycles and Focal Prices: Computational Dynamic Markov Equilibria." *J. Econ. and Management Strategy* 17 (May): 345–77.
- Palacios-Huerta, Ignacio. 2003. "Professionals Play Minimax." *Rev. Econ. Studies* 70 (April): 395–415.
- Posner, Richard A. 2001. *Antitrust Law*. 2nd ed. Chicago: Univ. Chicago Press.
- Reny, Phillip J., and Arthur J. Robson. 2004. "Reinterpreting Mixed Strategy

- Equilibria: A Unification of the Classical and Bayesian Views." *Games and Econ. Behavior* 48 (August): 355–84.
- Schelling, Thomas. 1960. *The Strategy of Conflict*. Cambridge, MA: Harvard Univ. Press
- Scherer, F. M., and David Ross. 1990. *Industrial Market Structure and Economic Performance*. 3rd ed. Boston: Houghton Mifflin.
- Shapiro, Carl. 1989. "Theories of Oligopoly Behavior." In *Handbook of Industrial Organization*, vol. 1, edited by R. Schmalensee and R. Willig. Amsterdam: North-Holland.
- Telser, Lester. 1960. "Why Should Manufacturers Want Fair Trade?" *J. Law and Econ.* 3 (October): 86–105.
- Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge, MA: MIT Press.
- Turner, Donald F. 1962. "The Definition of Agreement under the Sherman Act: Conscious Parallelism and Refusals to Deal." *Harvard Law Rev.* 75 (February): 655–706.
- Vives, Xavier. 2001. *Oligopoly Pricing: Old Ideas and New Tools*. Cambridge, MA: MIT Press.
- von Neumann, John, and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton Univ. Press.
- Walker, Mark, and John Wooders. 2001. "Minimax Play at Wimbledon." *A.E.R.* 91 (December): 1521–38.
- Wang, Zhongmin. 2009. "Station Level Gasoline Demand in an Australian Market with Regular Price Cycles." *Australian J. Agricultural and Resources Econ.* 53 (October): 467–83.
- Western Australia Department of Consumer and Employment Protection. 2007. "Submission to the ACCC Inquiry into the Price of Unleaded Petrol." <http://www.accc.gov.au/content/index.phtml/itemId/790926>, 15th submission.
- Western Australia Select Committee on Pricing of Petroleum Product. 2000. "Getting a Fair Deal for Western Australia Motorists." <http://www.parliament.wa.gov.au>.