Do firms play Markov strategies?*

Nicolas de Roos†
Sydney University

April 23, 2010

PRELIMINARY DRAFT
Please do not cite or circulate without permission

Repeated interaction between agents is ubiquitous, yet the road to a compelling theory of play under repeated interaction remains a perilous one. The supergame approach is bedevilled by issues of tractability and a continuum of equilibria. In most applications, a mechanism for limiting the set of admissable equilibria is a necessity. Restriction to Markov strategies presents a beguiling prospect. By conditioning only on payoff-relevant information, Markov strategies are simple and transparent. Reactions are also more immediately tied to individual self interest than the sometimes tortuous logic behind threats and punishments in the supergames literature.[1] From a more practical perspective, restriction to Markov strategies can do wonders for tractability.

We examine an unusual field setting that positions us well to test whether Markov strategies are played. Our setting is the market for retail petroleum in Perth, the capital city of Western Australia (WA). This market is a participant in an unusual policy intervention which permits us to observe the exact timing of price changes by firms. The market is also a prime exhibit of regular asymmetric price cycles resembling Edgeworth cycles. We exploit the model of Maskin and Tirole (1988) (hereafter MT) for our purpose because Markov strategies play a pivotal role in the model's predictions, and because the MT model has been adopted by the literature as the leading explanator of petrol price cycles.

Studies addressing similar questions have often resorted to a laboratory experimental setting or to field experiments in somewhat unconventional environments. For example, sporting contests and game shows are often exploited. By directly examining firm behaviour in an important industry, there is no need for us to ponder the portability of our results to a market environment.

Our field study has two main implications. First, firms do not play Markov strategies in this environment. This conclusion rests essentially on a simple observation: Edgeworth cycles are observed in a setting in which firms choose to act simultaneously. We expand on this argument below. Second, the predictions of the MT model are rejected for this setting, casting some doubt on the presumption by the literature that MT provides the theoretic undergirdings for empirically observed Edgeworth cycles.

*This project has benefitted from financial support from University of Sydney Bridging Support Grants in 2008 and 2009. We are grateful to Murali Agastya, Hajime Katayama, and Abhijit Sengupta for valuable comments and encouragement.

†Corresponding author: Faculty of Economics and Business, Merewether Building H04, University of Sydney, NSW 2006, Australia, e-mail: nicolas.deroos@sydney.edu.au, phone: +61 (2) 9351 7079, fax: +61 (2) 9351 4341.

[1] See Maskin and Tirole (1988) for additional discussion of the attractiveness of Markov strategies.

The rest of the paper is organised as follows. Below, we set the scene by sketching the model of MT, upon which our discussion is based. In section 2, we discuss the market setting. In section 3, we test the predictions of MT and examine whether Markov strategies (in the sense of MT) survive our empirical observations. We then turn to the literature on Edgeworth cycles in retail petroleum markets in section 4 where we argue that MT may not provide the foundations for empirically observed price cycles, contrary to the presumptions of the literature.

1 The Maskin and Tirole (1988) model

MT consider the following environment. Two firms compete repeatedly for a homogeneous product by choosing prices from a discrete price grid. Price commitment is introduced by permitting each firm to adjust price every second period, with this privelege alternating between firms. Thus, in period t, firm i chooses price p^i_t, while firm j remains committed to p^j_{t-1}. Players discount the future at fixed rate δ. Finally, MT restrict attention to Markovian strategies in which firms condition only on payoff-relevant states. In this setting, if i is called on to choose price in period t, the only payoff-relevant information is p^j_{t-1}.

Under these assumptions, we can summarise the dynamic problem faced by firm i when contemplating their period t choice of price as follows:

$$V^i\left(p^j_{t-1}\right) = \max_p \left[\pi^i\left(p, p^j_{t-1}\right) + \delta\, W^i(p)\right] \tag{1}$$

$$W^i\left(p^i_t\right) = E_{p^j_{t+1}}\left[\pi^i\left(p^i_t, p^j_{t+1}\right) + \delta\, V^i\left(p^j_{t+1}\right)\right], \tag{2}$$

where $\pi^i(p^i, p^j)$ represents firm i's profits at the price vector $\{p^i, p^j\}$. That is, $\pi^i(p^i, p^j) = 1\{p^i < p^j\}\pi(p^i) + 1\{p^i = p^j\}\frac{\pi(p^i)}{2}$, where $\pi(p)$ is industry profit at price $p = \min\{p^i, p^j\}$, and $1\{x\} = 1 \Leftrightarrow x$ is true. $V^i(p)$ is i's valuation of the expected discounted stream of profits when i chooses price while j is committed to price p. $W^i(p)$ is i's valuation when it is committed to price p and it is j's turn to move.

Informally, in an Edgeworth cycle equilibrium, firms undercut when their rival is committed to a high price, and relent by raising price with some probability when their rival is committed to a price equal to marginal cost. This results in a price path exhibiting an undercutting phase in which prices steadily decline and a relenting phase in which they rise in a single move. Transition from the undercutting to the relenting phase is punctuated by a war of attrition. Because products are homogeneous, the incentive to undercut is compelling when one's rival is committed to a high price. The incentive to relent is more tenuous, and stems from the anticipation of future opportunities to undercut. At the bottom of the cycle, firms would like to transition to the top of the cycle. However, each firm would rather their rival raised price first and thereby relinquish the first opportunity to undercut at the top of the cycle. This results in the war of attrition.

Relative to repeated Bertrand competition, MT adopt three primary modelling innovations. Firms choose prices from a discrete grid rather than a continuum; firms alternate in price setting; and firms condition only on payoff-relevant information (in this instance their rival's price). Superficially, the first two elements appear to play a critical role in the viability of an Edgeworth cycle equilibrium, while the third merely aids equilibrium selection. However, Proposition 9 of MT suggests that the alternating moves timing restriction emerges endogenously. The intuition is as follows. In the context of simultaneous moves play, if firms can only condition strategies on payoff-relevant information, then we are left with approximately the Bertrand outcome. With this weak opponent, sequential moves play emerges the victor. Importantly, it is the restriction to Markov strategies that plays a critical role.

We should also note that the predictions of the MT model are robust to a number of generalisations. Eckert (2003) finds cycle equilibria for a wide range of firm sizes. Noel (2008) finds computationally that Edgeworth cycle equilibria are robust to product differentiation, i.i.d. cost shocks, capacity constraints, and three firms. Similarly, Robertson (2008) uses computational methods to identify Edgeworth cycle equilibria in the presence of a mixture of searching and non-searching consumers.

We close this section by outlining some empirical predictions of the MT model.

Prediction 1. Firms alternate in setting prices.

As discussed above, for markets with homogeneous products, this follows directly from Proposition 9 of MT. Based on the same logic, we conjecture that this result is robust to mild product differentiation.

Prediction 2. If firms move simultaneously, then the equilibrium is in fixed prices. Moreover, if products are homogeneous, price is approximately equal to marginal cost.

This follows directly from the restriction to Markov strategies. There is no payoff-relevant information if play is simultaneous, and we are left with the one shot Bertrand outcome. Price will deviate from marginal cost only due to the coarseness of the price grid.[2]

Prediction 3. If firms alternate in setting prices, equilibrium will take the form of either focal pricing or Edgeworth cycles.[3]

For homogeneous product environments, this is shown analytically by MT. As discussed above, this survives a number of generalisations.

Prediction 4. If Edgeworth cycles are observed, firms will undercut by a single grid point in the undercutting phase, and the peak of the cycle will be above the monopoly price.

Because Edgeworth cycles are predicted only in an alternating moves environment, in the event that firms move simultaneously, we view a test of this prediction as a test of the importance of the timing prediction. This prediction is based on the strategies used by MT (see equation (23)) to construct the existence proof of an Edgeworth cycle equilibrium. The first part of our claim stems directly from these strategies. The strategies involve undercutting by a single grid point at the top of the cycle, with the final undercut being down to marginal cost.

For the second part of our claim, some intuition is provided by the following informal proof by contradiction. Adopting the notation of MT, let the peak of the price cycle be $\bar{p}+k$, where k is the size of a grid point, and let $R(.)$ denote strategies in an Edgeworth cycle equilibrium. Suppose that $\bar{p}+k < p^m$, the monopoly price. Consider a state $\hat{p} > p^m + k$. Equilibrium play (see equation (23) of MT) calls for $R(\hat{p}) = \bar{p}$, yielding payoff $\pi(\bar{p},\hat{p}) + \delta^2 V(\bar{p}-k)$. Playing p^m instead yields $\pi(p^m,\hat{p}) + \delta^2 V(\bar{p})$. Because $V(.)$ is nondecreasing (MT Lemma A), this is a profitable deviation.

Prediction 5. If focal prices are observed, profits will be bounded substantially above the competitive level.

This prediction follows from Proposition 3 of MT which suggests that industry profits will be bounded above $\frac{4}{7}$ of industry profits. As above, this prediction rests on an alternating moves environment. In the context of simultaneous play, violation of this prediction highlights the importance of the timing prediction.

[2] See footnote 13 of MT for details.

[3] As we shall see, firms do not alternate in setting prices, so we are unable to test this prediction. We include it for completeness.

2 The Perth market environment

There are three main stages to production in the petrol market: refining, wholesale distribution, and retailing.[4] In Australia, four *oil major* are integrated into all three stages: BP, Caltex, Mobil and Shell. Supermarket chains play a growing role in the retail market, with the remainder accounted for by independent firms. BP operates the only refinery in WA. In Perth, wholesale petrol is available at terminal gates operated by each of the oil majors and the independent firm, Gull.

Beginning on 3 January 2001, a unique policy environment called Fuelwatch was created in the state of Western Australia (WA). Under the scheme, retailers are required to notify Fuelwatch of their intended price for the next day by 2pm. This price must then be offered at the pump for 24 hours, starting 6am the next day. At 2:30pm the price is posted on a publicly available website. Most of the retail petroleum outlets in WA are covered by the scheme. Consequently, consumers can obtain price information before hopping into their cars by clicking on a website.

From our perspective, Fuelwatch has two immediate implications. First, firms are restricted to play in a discrete time environment, with each day corresponding to a choice node. Second, we can observe the entire history of prices chosen by a retailer rather than merely a sample. Consequently, we are uniquely positioned to study firm behaviour. We observe every choice made by firms and the information they possess at the time of choosing. We have access to pricing data from the Fuelwatch program for the period January 3 2001 to January 24, 2007. This dataset is described in de Roos and Katayama (2010).

If we abstract from locational differentiation, retail petrol is very close to our ideal of a homogeneous product. The physical product is, to most consumer perceptions, identical; and the service experience is similarly bland across locations. In most retail petrol markets, price information is highly visible through display on large street-side billboards. In Perth, price information is also readily accessible online through the Fuelwatch program.

We believe the restriction of price choices to a discrete grid suits the retail petrol market. Much of advertising is through street-side billboards making it both impractical for firms to display prices to a large number of decimal places and unrealistic for consumers to be able to obtain and retain this information while driving. In practice, stations report prices in increments of at least 0.1 cents per litre.

3 Do firms play Markov strategies?

The above discussion suggests that the retail petrol market in Perth satisfies most of the assumptions of the MT model. Firms compete by choosing prices, plausibly from a discrete price grid. There is very little product differentiation; and the predictions of MT appear robust to mild differentiation in any case. We also have no cause to question the presumption of discounting. The only assumption upon which we are unable to form a judgement is the restriction to Markov strategies. The Fuelwatch policy context enables us to probe this restriction.

Let us now return to our earlier predictions. The prevalence of stations exhibiting an Edgeworth cycle is immediately apparent from casual inspection of the data. de Roos and Katayama (2010) (hereafter DK) argue that stations exhibiting an oil major brand tend to exhibit strongly asymmetric price cycles of a similar character to those observed recently in other retail petrol markets. Additional discussion is also presented in Wang (2009). Casual inspection of the data also reveals that firms adjust prices simultaneously. This is evident in the graphical depiction of data in DK and in Wang (2009), and confirmed by the empirical work of DK. Wang (2009) also confirms that oil majors often simultaneously raise price at the

[4] Additional details of the WA petrol market can be found in Wang (2009) and de Roos and Katayama (2010).

base of the cycle, a critical strategic juncture of the cycle. We also note that price adjustments are not simply a response to changes in cost conditions; it is the retail markup that adjusts each day.

This observation is clearly at odds with Prediction 1. Prediction 1 rests on the comparative unattractiveness of equilibria in simultaneous move strategies, which stems directly from the restriction to Markov strategies. We note in passing that the discreteness of the price grid is not pivotal here. Lau (2001) shows that alternate timing emerges if we permit firms to choose from a continuum of prices. Similarly, Markov strategies play a critical supporting role in Prediction 2, which is ruled out by the prevalence of cycling stations. If we relax the Markov restriction, and firms can condition on the history of prices, it is straightforward to identify equilibria exhibiting price cycles. Taken together, we see violation of Predictions 1 and 2 as clear evidence against the assumption of Markov strategies.

To test Prediction 4, we draw on the observations of DK relating to the characteristics of the cycle. Consistent with the results of Noel (2007), DK find that firms tend to undercut more substantially at the top of the cycle. By contrast, MT predict undercutting at a constant rate, with a potentially larger final undercut to marginal cost. By itself, we do not view this as strong evidence against MT.

Our next observation is that the peak of the Edgeworth cycle appears to be well below the monopoly price. DK suggest that the trough of the cycle occurs at approximately the wholesale list price called the "terminal gate price". This terminal gate price plausibly approximates the opportunity cost of petrol. Wholesale transactions may occur below the terminal gate price, but are unlikely to occur much below this price. As Wang (2009) notes, BP is constrained in wholesale pricing by the prospect of imports. Consequently, terminal gate prices are closely related to an "import parity pricing" formula based on Singapore prices, shipping costs and incidentals. [Check with Fuelwatch people about this.] DK show that the amplitude of the cycle averages around 10 cents for oil majors, or between approximately 7 and 11 per cent of the retail price. Hence, the markup at the cycle peak appears to be below 15 per cent. Given that, in the short term, petrol is a necessity for most commuters, this seems implausibly low for a monopoly markup, and we interpret this as a violation of Prediction 4.

We view this last result as evidence that the timing prediction has some import. Consider for a moment a candidate simultaneous moves equilibrium in which price cycles are exhibited. Notice that a cycle peak that is above the monopoly price will present enormous strain on the candidate equilibrium. A deviator is tempted not just by the prospect of stealing market share from its rivals, but also doing so at a price at which industry profits are higher than the status quo.

To examine Prediction 5, we examine the stations classified by DK as non-cycling. For these stations, the average retail margin appears to be of the order of 5-10 per cent [CHECK]. We view this as a violation of 5. Again, violation of the timing prediction is important. Notice, however, that for these stations, violation of Prediction 2 is not acute.

4 Implications for retail petrol price cycles

At least since Castanias and Johnson (1993), price cycles in retail petroleum markets have been linked to the price commitment model of MT. This association is natural. The Edgeworth cycle equilibrium of MT yields a distinctive price path with highly asymmetric cycles of substantial amplitude. Further, these same characteristics are observed in many retail petrol markets. In this section, we make two main points. First, in the policy context of Western Australia, MT does not provide an explanation of the petrol price cycle. Second, we may wish to reevaluate whether petrol price cycles in other environments can be accounted for by MT.

Our first point follows directly from our discussion in the previous section. Let us briefly summarise.

Firms adjust prices simultaneously not sequentially, violating the Markov assumption of MT. Further, Edgeworth cycles are observed despite simultaneity of play, again violating the Markov assumption. One could argue that MT still delivers a reasonable approximation by predicting the strong asymmetry observed in retail petrol markets. We have two objections. First, the qualitative predictions of MT are imperfect. The most striking violation is the simultaneity of price changes. In addition, the amplitude of the cycle predicted by MT appears too large with a peak above the monopoly price, and a trough at marginal cost. Second, and more fundamentally, we believe the MT model approximates the asymmetric cycle of retail petrol markets for the wrong reasons. A model of price cycles must at a minimum deliver two components: an undercutting phase in which firms have a strong incentive to lower price; and a relenting phase in which firms have an incentive to raise price. The MT model delivers these features, but otherwise does not match empirical observation in Perth. In this matter, we disagree with Wang (2009) who finds support for MT in the same context.

Our second point must be more speculative. The restriction to discrete time play is unique to the WA petrol market. In all other environments, play is in continuous time. This makes it challenging to infer whether play is simultaneous or sequential. A simultaneous moves game has two well known interpretations: players either move literaterally simultaneously, or they move before observing the action of their rival. In the context of continuous time play, the latter is the only viable interpretation. The interpretation of a sequential moves game is clear: firms move after observing their rival's choice. The distinction then boils down to the following question: do firms observe rival prices before implementing their own price changes?

To answer this question in a particular environment, we must understand the mechanics of price changing. A change in price requires two main elements: i) the recognition that a price change is called for and the calculation of the new optimal price; and ii) the implementation of that price. Most retail petrol markets are highly concentrated, with each major brand associated with a large number of stations. If prices are centrally controlled as they are for the major brands in WA (see Wang (2009) for details) and locally implemented, there is plausibly a substantial recognition and implementation lag. This suggests to us that we should seriously consider whether a simultaneous play environment is also appropriate for other retail petrol markets. Concomitantly, we should also revisit the presumption that MT provides the foundations for empirically observed petrol price cycles in other markets.

5 Conclusion

We harness an unusual policy context to examine whether firms play Markov strategies. We find that the MT model is clearly rejected by observations in the WA retail petrol market. Further, central to our conclusions is the restriction to Markov strategies imposed by MT. In this environment, we must conclude that firms do not play Markov strategies. The inapplicability of the MT model in WA also forces us to question whether this model provides the theoretic underpinnings for Edgeworth cycles in retail petrol markets as has been widely accepted in the literature.

References

Rick Castanias and Herb Johnson. Gas wars: Retail gasoline price fluctuations. *The Review of Economics and Statistics*, 75:171–174, 1993.

Nicolas de Roos and Hajime Katayama. Retail petrol price cycles in Western Australia. Working paper, 2010.

Andrew Eckert. Retail price cycles and the presence of small firms. *International Journal of Industrial Organization*, 21:151–170, 2003.

Sau-Him Paul Lau. Aggregate pattern of time-dependent adjustment rules, II: Strategic complementarity and endogenous nonsynchronization. *Journal of Economic Theory*, 98:199–231, 2001.

Eric Maskin and Jean Tirole. A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica*, 56:571–599, 1988.

Michael Noel. Edgeworth price cycles: Evidence from the Toronto retail gasoline market. *Journal of Industrial Economics*, 55:69–92, 2007.

Michael Noel. Edgeworth price cycles and focal prices: Computational dynamic Markov equilibria. *Journal of Economics & Management Strategy*, 17:345–377, 2008.

David Robertson. Price cycles in the retail petroleum market and the role of searching consumers. University of Sydney Honours Thesis, 2008.

Zhongmin Wang. (Mixed) strategy in oligopoly pricing: Evidence from gasoline price cycles before and under a timing regulation. *Journal of Political Economy*, 117:987–1030, 2009.