Edgeworth Cycles and Focal Prices:

Computational Dynamic Markov Equilibria

Michael D. Noel

9500 Gilman Dr. 0508, UC San Diego

La Jolla, CA 92093-0508

mdnoel@ucsd.edu

September 28, 2004

Abstract

Motivated by the discovery of apparent Edgeworth Cycles in many retail gasoline markets, this paper extends the Maskin & Tirole [1988] theory of Edgeworth Cycles to a wide range of more complicated and realistic settings. Taking a computational approach to search for Markov Perfect Equilibria, I examine models involving duopoly and triopoly, differentiation, capacity constraints, and different sharing rules, discount factors and initial beliefs about price leading behavior. I find Edgeworth Cycles in equilibrium in many scenarios outside the homogenous-good Bertrand mold. Cycle characteristics and average markups depend on the scenario.

JEL Classification: D43, L11, L13

1 Introduction

Several authors have recently documented an interesting price cycling phenomenon in many retail gasoline markets in Canada. The price cycles are high-frequency, tall relative to marginal costs and sharply asymmetric. These authors (Noel[2003a], Noel[2003b], Eckert[2003]) argue the cycles both appear and behave very similar to the Edgeworth Cycles of Maskin & Tirole[1988].

To see the visual similarity between the empirical and theoretical cycles, compare Figure 1, taken from Noel[2003b], with Figure 2 taken from Maskin & Tirole[1988]. Figure 1 shows retail prices for a representative major branded station and a representative independent along with the wholesale price for the city of Toronto during 2001. The interval between consecutive data points is just 12 hours. The graph clearly shows the asymmetry in price movements. When a station increases price, it does so by 13% on average and as much as 25%. Two consecutive increases are effectively non-existent. Yet other stations respond to a competitor increase almost immediately. In contrast, when prices decrease, they do so by less than 2% per period and the overall decrease is spread out over many periods. The pricing pattern is similar to that of Figure 2, which shows the path of the market price in a theoretical Edgeworth Cycle.[1] I refer the reader to the abovementioned papers for further evidence that the cycles behave consistent with the theory of Edgeworth Cycles.

In this article, I turn my attention to the underlying theory. The theoretical market structure under which Edgeworth Cycles are derived is highly simplified

[1] For a specific numerical example. Marginal costs in the example are zero.

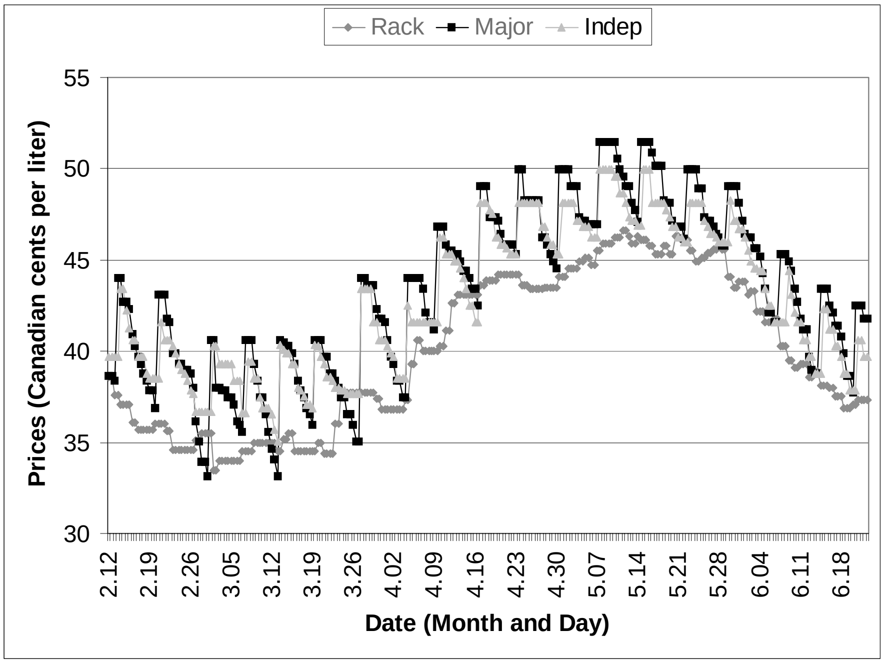

Figure 1: Retail Prices (Major, Independent) and Rack Price

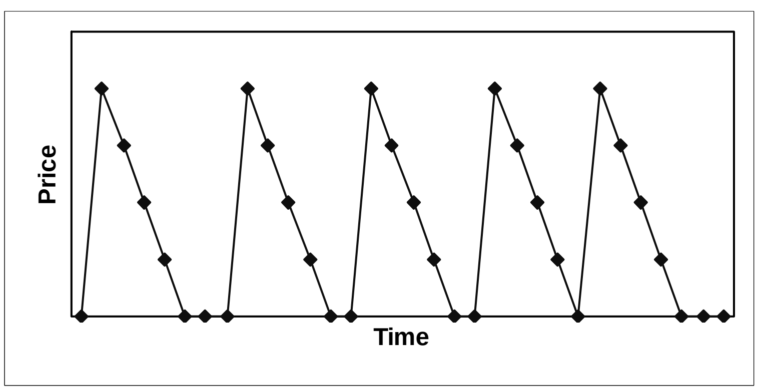

Figure 2: Edgeworth Cycle

– typically it assumes a dynamic, symmetric Bertrand duopoly game with perfectly homogeneous goods. Yet the apparent discovery of Edgeworth Cycles in retail gasoline markets underscores the point that Edgeworth Cycles can exist in real world environments that do not fit the simple Bertrand mold. Hence, it is important to develop the theory more fully to learn under which competitive environments we can observe Edgeworth Cycles. Conversely, we will be able to make inferences about the underlying environment in markets where these cycles are found.

To answer these questions, I extend the original theory of Edgeworth Cycles along many new dimensions. I consider models involving such real world complications as horizontal differentiation, capacity constraints, and markets with more than two firms. I also consider different assumptions about discount factors, market elasticities, and beliefs about leader-follower behavior. The main goal of this article is to see under which models (and model-specific parameter assumptions) Edgeworth Cycles can still exist in equilibrium, and how the characteristics of the cycles differ with the environment. In order to investigate so many different scenarios, I use a computational dynamic programming algorithm to search for equilibria.

The results show that Edgeworth Cycles are an equilibrium under a wide range of scenarios, not just the simple Bertrand assumption. They can easily be found in modestly differentiated goods markets. They can also easily be found in the presence of capacity constraints as long as those constraints are not very tight. Moreover, I show that cycles are not confined to duopoly settings but exist in triopoly as well. In the tripoly case, interesting coordination challenges arise. Delayed starts and even false starts in resetting cycles become a part of the equilibrium cycle process. The emergence of a price leader can help in this regard.

The shape of the Edgeworth Cycles also varies in interesting ways with the specific scenario. Variation in the aggressiveness of firms directly impacts the amplitude, period, and asymmetry of the cycles. In many scenarios, price leaders and followers emerge and impact the cycle in different ways. Therefore the cycle shape contains important information about the underlying competitive environment. Also, average markups vary dramatically across scenarios.

Taken together, the results show that Edgeworth Cycles are robust to many more scenarios than previously assumed. Real world markets in which they have been documented indeed do not fit the simple Bertrand mold. Edgeworth-like cycles are regularly observed in many Canadian retail gasoline markets and have also been observed in several west coast U.S. retail gasoline markets in the 1960s and 1970s.[2]

As mentioned, previous theoretical work has maintained the straight Bertrand assumption for analytical tractability. In the original game-theoretic model of Maskin & Tirole[1988], the authors assume a symmetric Bertrand duopoly, in which firms alternate in setting prices. They show that there are two sets of

[2] For the U.S., see Allvine & Peterson(1974) and Castanias & Johnson(1993). For Canada, see Noel[2003a], Noel[2003b], Eckert[2002] and Eckert[2003]. Prior research on oligopoly pricing in Canadian retail gasoline markets (for example, Slade[1987,1992] or Godby et al.[1997]) find asymmetric price movements but do not report Edgeworth Cycles as a possible source. For Edgeworth Cycles in an experimental setting, see Kruse et al. (1994).

possible Markov Perfect equilibria under these conditions – focal price equilibria and "Edgeworth Cycle" equilibria. The former is characterized by constant prices in equilibrium, the latter by an interesting and asymmetric cycle. In an Edgeworth Cycle, firms repeatedly undercut one another to steal the market, until price reaches marginal cost. At that point, a war of attrition ensues with each firm mixing between raising price and remaining at marginal cost. Eckert[2003] extends this to the case when two firms share the market unequally at equal prices and finds Edgeworth Cycles still exist in equilibrium.

The concept of Edgeworth Cycles dates back to Edgeworth[1925] who considers two identically capacity constrained firms. The author postulates that after undercutting brings firms close to their capacity constraints, one could raise price and profitably serve the residual demand. Maskin & Tirole[1988] show that capacity constraints are not a prerequisite, and one of the findings in the current article is that, contrary to Edgeworth, stronger capacity constraints make Edgeworth Cycles less likely, not more.

The paper is organized as follows. Extensions of a homogeneous-good Bertrand duopoly in section 2. The differentiated goods model appears in section 3, and the Bertrand triopoly model appears in section 4. Section 5 concludes.

2 Bertrand Duopoly

The basic model I consider is the following extension of Maskin & Tirole[1988]. Two infinitely-lived profit-maximizing firms compete in a homogeneous Bertrand pricing game by setting prices in an alternating fashion – one firm sets its price in even periods and the other in odd periods – and once set, the price for that firm is fixed for two periods. Therefore, if firm 1 adjusts its price in period t, $p_t^1 = p_{t+1}^1$ and $p_t^2 = p_{t-1}^2$. Prices are chosen from a discrete price grid. Marginal cost, c_t, is also allowed to vary over time, and is chosen by nature from a discrete cost grid. Each firm earns current period profits of

$$\pi_t^i(p_t^1, p_t^2, c_t) = D^i(p_t^1, p_t^2) * (p_t^i - c_t) \tag{1}$$

where

$$D^i(p_t^1, p_t^2) = \begin{cases} D(p_t^i) & \text{if } p_t^i < p_t^j \\ \theta^i D(p_t^i) & \text{if } p_t^i = p_t^j \quad \text{for } i \neq j \\ 0 & p_t^i > p_t^j \end{cases} \tag{2}$$

The parameter $\theta^i \in [0,1]$, $\sum_i \theta^i = 1$ specifies the fraction of market demand firm i receives at equal prices.

The strategies of each firm are allowed to depend only on the payoff-relevant state in each period. Therefore, a firm's strategy depends only on the price set by the other firm in the previous period, and current marginal cost which it learns prior to setting its price. The equilibrium concept is that of a Markov Perfect Equilibrium (MPE), and the equilibrium strategies are given by R^1, R^2, where $(p_t^1)^* = R^1(p_{t-1}^2, c_t)$, $(p_t^2)^* = R^2(p_{t-1}^1, c_t)$ and $p_{t-1}^j$is the price chosen by firm j is period $t-1$ which remains in effect in period t.[3]

Let $V^1(p_{t-1}^2)$ be the firm 1's value function when firm 2 adjusted its price to p_{t-1}^2 in the previous period, firm 1 adjusts its price in the current period, and the current marginal cost c_t is not yet known. Let $W^1(p_{s-1}^1)$ be firm 1's value function when it has set price p_{s-1}^1 in the previous period, firm 2 is about to adjust its price, and the current cost is not yet known. V^1 and W^1 can be written as

$$V^1(p_{t-1}^2) = \mathop{\mathrm{E}}_{c}\left(\max_{p_t}\left[\pi_t^1(p_t, p_{t-1}^2, c_t) + \delta_1 W^1(p_t)\right]\right) \tag{3}$$

$$W^1(p_{s-1}^1) = \mathop{\mathrm{E}}_{c}\left(\mathop{\mathrm{E}}_{p_s}\left[\pi_s^1(p_{s-1}^1, p_s, c_s) + \delta_1 V^1(p_s)\right]\right) \tag{4}$$

and similar equations are found for V^2 and W^2. The firm-specific discount factor is δ_i. The inside expectation in W^1 is taken with respect to the distribution of R^2 and both the outside expectation in W^1 and the expectation in V^1 is taken with respect to the distribution of c.[4] To choose the best response price, given the current rival price p_{t-1}^2 and current cost c_t, firm 1 maximizes $\pi_t^1(p_t, p_{t-1}^2, c_t) + \delta_1 W^1(p_t)$ (ie. V^1 without the expectation.) Firm 2 acts in a similar way.

The original Maskin & Tirole[1988] model can be recovered from this setup simply by setting $\theta^i = \frac{1}{2}$, $\delta_1 = \delta_2$, and $c_t = c$ for all t. The Eckert[2003] extension can be recovered by setting $\delta_i = \delta$ for all i, and $c_t = c$ for all t but allowing θ^i to differ from $\frac{1}{2}$ so that firms share the market unequally at equal

[3] See Maskin & Tirole (2001) for general properties of MPEs (not simply with alternating moves games).

[4] This formulation implicitly assumes there is no persistence in c, a point to which I will return.

prices. Since marginal costs are constant and known in these models, each R^i does not depend on c and the expectations in V^i and W^i over c vanish.

From this starting point, I extend the basic model to a wide range of scenarios. Because of the number of variants, I employ a computational dynamic programming algorithm to solve for the value functions V^i and W^i and the best response functions R^i. There is a challenge in doing this because the MPE are generally mixed strategy equilibria and an algorithm based on pure strategies will not converge. One could solve for the probability (or probabilities) analytically if one knew exactly at which opponent price (or prices) the mixing would occur, but this is not known ex ante.

To circumvent this problem, I adapt the model in such a way to ensure it always has a pure strategy equilibrium. I do this by allowing marginal costs to fluctuate within a specified band period by period. To see how this works, imagine the equilibrium in a constant marginal cost world would call for a firm to set p_H with probability $\frac{1}{2}$ and p_L with probability $\frac{1}{2}$. I replicate this by having the firm draw a new c_t from the marginal cost distribution in the current period. If that draw is above the median, the firm in equilibrium will choose p_H and if below, it will choose p_L. Effectively, the current realization of marginal cost does the required "mixing". The resulting pure strategy equilibrium involves a best response for firm i, $R^i(p^j_{t-1}, c_t)$ that depends on both p^j_{t-1} and c_t.

Given that the mixing probabilities are unknown ex ante, I choose a very fine grid of possible marginal costs (to replicate the mixing probability to within 0.00025) and a uniform distribution to cover the band evenly. While the purpose of including fluctuating marginal costs is to solve the "mixing problem", the resulting cost process is necessarily an abstraction. For example, current period marginal cost is assumed independent of previous draws, the distribution is assumed uniform and finally, marginal cost can change each period although any given price changes every two. These simplifications avoid significant dimensionality problems in the computation, as explained in the footnote.[5] As my purpose is to study Edgeworth Cycles under a variety of market structures, these assumptions are not important to the results. The cost band is narrow and the particular process within that band is of little importance.

The system is converged when fixed point vectors V^i and W^i are found. At each iteration k, I update all four value functions as follows. To update V^1 and W^2, I calculate the best response function of firm 1, $(p^1)^* = R^1(p^2, c)$, to every possible prior period price p^2 and cost realization c. I then calculate the net present value of profits for each firm (given each p^2, c, $R^1(p^2, c)$, and $V^2_{k-1}(p^1)$ and $W^1_{k-1}(p^1)$ from the previous iteration). Next, for each p^2, I calculate the expectation over c of each firm's profits. This expectation is the new $V^1_k(p^2)$ and $W^2_k(p^2)$. I update V^2 and W^1 similarly.[6]

[5]In the current setup, V^i and W^i depend only on previous price and so are $\#_p$ vectors, where $\#_p$ is the number of points on the price grid. Adding persistence in marginal costs – perhaps as a simply random walk process – means that V^i and W^i must each contain $\#_p * \#_c$ elements (and must converge elementwise), where $\#_c$ is the number of points on the cost grid. In examples that follow, $\#_p$ is 20 but $\#_c * \#_p$ is 40,000. Also, with any distribution other than uniform, either a greater $\#_c$ would be necessary to achieve the same 0.00025 standard error, or the cost grid would have to be spaced unevenly with more points near the peak of the distribution.

[6]Pakes & McGuire (1994,2001), Ericson & Pakes (1995), Pakes(2000) and others suggest techniques for reducing computational burden such as using an polynomial approximation for the value function and making efficient use of symmetry. Because of the discrete and jumpy nature of the best response functions I describe below (and resulting "waviness" of the value function), I choose to use the precise, but slow, algorithm in the text. I use symmetry between firms to reduce the burden where possible, but since most specifications are asymmetric in

In the Maskin & Tirole model, multiple Edgeworth Cycle equilibria and multiple focal price equilibria exist. Therefore I test a wide range of starting values in each scenario in case there are multiple equilibria. Interesting, I routinely find the same equilibrium each time regardless of the starting values attempted, except as noted. In particular, I have not replicated both Edgeworth Cycle and Focal Price equilibria in the same scenario. This may be due to the fluctuating marginal cost assumption or because of the instability of the one of the equilibria. For consistency, all reported results use starting values for V^i and W^i that would be the outcome in a single period static game (ie. if discount factors were zero) except as noted.

2.1 Undercutting, Matching,and Relenting

The baseline case is the simple two-firm Bertrand model with linear inverse demand curve given by $D(p) = a - bp$. To fix an example for discussion, let $a = 20$ and $b = 1$. Prices are chosen from a discrete price grid $p^i = \{x\},\ x = 0..20$. Marginal cost in each new period is drawn randomly from a discrete uniform cost grid in the range $c_t = \{x/2000\},\ x = 0..2000$. Also, set the equal-price sharing rule $\theta_i = \frac{1}{2}$ and the discount factor $\delta_i = 0.95$ for each firm.

The key difference between this specification and previous works is that I allow for fluctuating marginal costs. In this case, I still easily find Edgeworth Cycles in equilibrium. In the figure appendix, I report the equilibrium best response functions $R^i(p^j_{t-1}, c_t)$ in the top panel of Figure 3 and the equilibrium

nature, this is of limited value.

price paths over 40 periods in the bottom panel.[7]

With minimal practice, one can glean all the pertinent information in the best response figure in a single glance. For this first example, I discuss the mechanics of firm 2's best response function, depicted with circles, in more detail. First, any circle placed immediately below the 45° line represents an undercut of one notch on the price grid by firm 2. For example, firm 2's response to any $p_1 \in [9, 14]$ is to undercut by a single notch. A circle placed further below the 45° line represents a larger and more aggressive undercut. The figure shows that firm 2 aggressively undercuts any $p_1 \in [15, 20]$ down to 14 and undercuts any $p_1 \in [5, 7]$ down to 4 (and as explained below, it sometimes undercuts $p_1 = 8$ to 4 as well.)

Matching firm 1's price is represented by a circle directly on the 45° line, as can occur in response to $p_1 = 1$ or 2. When firm 2 responds by raising its price back to the top of the cycle, which I call "relenting", we observe circles far above the line, as can happen in response to $p_1 \in [0, 4]$. Had firm 2 wanted to respond by raising its price only slightly instead, which I call "stepping up" (and not observed in this example), we would see a circle only slightly above the line.

Often, firm 2 will respond to firm 1 by mixing between two or more actions (more accurately, its action depends on realized marginal cost). This is seen as multiple circles on the same vertical line. In this example, it happens at $p_1 = 8$ (where firm 2 mixes between undercutting to 7 and undercutting to

[7] Because the price grid is discrete, the displayed best response functions are also discrete. I have, however, connected the dots in the price path figures to facilitate presentation.

4), $p_1 = 3$ or 4 (mixing between undercutting and relenting), $p_1 = 2$ (mixing between undercutting, matching, and relenting), and $p_1 = 1$ (mixing between matching and relenting).[8]

Reading firm 1's best response function is similar: undercuts are represented by triangles to the left of the 45° line, relenting and stepping up are to the right, and multiple triangles on the same horizontal line represent mixing.

Firms exhibit several interesting behaviors along the cycle. First, since the monopoly price in the static Bertrand game is 10 to 11, the first undercut from the top of the cycle is to a price substantially greater than the static monopoly price. This allows firms to operate more often through the most profitable section of the demand curve.

Second, undercutting proceeds in an orderly fashion, one notch at a time, through the most profitable prices. Then firms become more aggressive when price reaches 8. If marginal cost is high enough, a firm will aggressively undercut to 4 (and if not, it undercuts to 7 and the opponent with certainty undercuts to 4 the next period). A price of 4 opens up the possibility that the other firm will relent (setting price back to 15) the very next period.[9]

Third, once prices become low, firms play an increasingly passive strategy. Each firm would prefer not to be first to raise price since it would mean losing its turn of being the low price supplier to the market. Should price reach 1 or

[8] For readability, I do not report the exact mixing probabilities (or scale the size of the symbols) in the graphs when multiple responses are possible. It is always the case that undercutting occurs at lower marginal costs than matching, and matching occurs at lower marginal costs than relenting or stepping up.

[9] Although excluded from the diagram to reduce clutter, the probability of relenting increases as the price falls from 4 to 1, and is certain from a price of 0. (Since 0 is not a response to any price, though, it does not occur on the equilibrium path.)

2, firms often use a matching strategy, accepting a split of the market at very unprofitable prices and then passing the turn back to its opponent. Unlike the quick and certain length of the relenting phase, the length of the undercutting phase is greater and more uncertain.

Finally, relative a static Bertrand game, the average market price consumers face and firms receive in the simulations is relatively high. At 8.5, it is close to the average static monopoly price of 10.5 but off the highest static competitive price of 2.

2.2 A Price Leader and the Step-Up Strategy

In the previous example, the firms were symmetric and equally likely to relent first. But one can easily imagine situations in which otherwise symmetric firms fall into a pattern in which the same firm becomes a consistent price leader in raising prices. The leader relents first each time because it knows the follower never will; the follower never relents because it knows the leader always will. Computationally, I find this asymmetric equilibrium after making a small, random perturbation in the starting values in V^i for one of the firms.

In Figure 4, firm 1 is the consistent price leader and firm 2 is the follower. (Note the absence of a horizontal line of circles in the upper left.)

The behaviors of the firms now differ in interesting ways that impact the shape of the cycle. Relative to the leader and to itself in the previous case, follower behavior is especially aggressive at high and moderate prices and especially passive at low prices. From the top of the cycle, its first undercut is to a price very close to the static monopoly price. Since firm 1 will surely be first to relent, firm 2 has less incentive to delay the bottom of the cycle into the future and instead takes a large profit currently. Firm 2 may then undercut by as much as five notches, causing firm 1 to relent again sooner.[10] In contrast, firm 1 is less aggressive than in the previous case, preferring to undercut more slowly and relent less often.

Once prices get low, firm 2 suddenly becomes passive. In response to low prices, it may either match price or instead simply "step up" its price just above that of firm 1 (as shown by circles just above the 45° line). Though not observed on the equilibrium path, this step up strategy is important as it guarantees the leader will relent.

The resulting cycle is more rapid, smaller in amplitude, and less asymmetric relative to the previous case.

As expected in a price game, there is a profit advantage to the follower. Firm 2 enjoys profits 50% higher than the leader, and 24% higher than in the symmetric case. The leader's profits fall by 17%. Consumers are also better off, with an average price of 7.7, or 10% lower than in the symmetric case.[11]

2.3 Elasticities, Discount Factors, and Sharing Rules

To see how the aggregate elasticity of demand affects cycles, I examine a series of demand curves by changing the parameters a and b to pivot around the point

[10] Undercutting to 5 yields a $\frac{2}{3}$ probability that firm 1 will relent in the next period. Undercutting to 3 guarantees it.

[11] For each symmetric equilibria, these two additional equilibria – one for each firm as the price leader – are easily found. I return to the practice of symmetric starting values hereafter.

(10,10). Cycles are easily generated for all values of a and b, consistent with the notion that the primary gain to undercutting is from stealing the existing consumer base, rather than creating new consumers. But the shapes differ. With relatively more elastic demand curves, firms are *less* aggressive in undercutting, more often proceeding by one notch at a time. This is because they can serve a relatively larger market at low prices and can thus be more patient during the war of attrition. The best response function and cycles for the more elastic case $a = 15,\ b = 0.5$ are shown in Figure 5. With relatively less elastic curves, aggressive multi-notch undercutting near the bottom is more common and the cycles progress more quickly.[12]

I also find Edgeworth Cycles in equilibrium for all tested pairs of discount factors $\delta_i \in (0, 1)$.[13] Consider first the case of a common discount factor. When the discount factor is relatively lower, the first undercut from the top of the cycle will come closer to the static monopoly price. Then, since firms are less interested in hastening the next relent (and the greater profits that comes with it), undercutting proceeds more slowly, by one notch at a time even through lower prices. The downward slope of the cycle loses its concavity and becomes more linear (until matching at very low prices adds some convexity). Extended matching at low prices is more likely.[14]

[12] Rather than pivot around (10,10), I also simulated parallel shifts in demand and pivots around the quantity intercept (of 20) with descriptively similar results – less aggressive the more elastic. Pivoting around the price intercept yields a set of identical best response functions since the elasticity at a given price is unchanged.

[13] $\delta_i = \{0.01, 0.05x,\ 0.99\}, .x = 1..19$. Even when $\delta_i = 0$, Edgeworth Cycles is still one possible type of equilibrium, since when prices are at marginal cost, a firm is indifferent across all $p \geq c$. Any $p > c$ will create a cycle.

[14] It is interesting to note that while economists describe such firms (with low δ) as less patient, some non-economists may view these firms as being *more* patient since they routinely

For example, in the case of $\delta_i = 0.5$ (not shown), the first undercut is to 11 or 12 and undercutting is always by one notch at a time down to 2. Matching or undercutting still occurs with 61% probability from a price of 2 and matching still occurs with 22% probability from 1. The cycles are thus shifted downward in price, and longer in duration. Consumers are better off: the average price is 5.8 in this scenario, 44% below the static monopoly price and 31% below the baseline $\delta_i = 0.95$ case.[15]

If firms have different discount factors, the equilibrium is a mix between the asymmetric equilibria of the previous section and low discount factor equilibria discussed above. First, the most patient firm (with the highest δ) emerges as the consistent price leader and relents first every time. For example, assume this is firm 1 with $\delta_1 = 0.95 > \delta_2$ (cases not shown). Because firm 1 puts greater weight on future profits, it has the greater incentive to move prices back to the top of the cycle. From there, firm 2 undercuts to a price close to the static monopoly price – and the lower the discount factor, the closer – until at a discount factor of about $\delta_2 = 0.5$, it undercuts directly to it. And while firm 1 undercuts by one notch at a time except near the very bottom, firm 2 plays an aggressive stategy at moderate prices, often undercutting by several notches at once. The higher firm 2's discount factor, the more it does of this, the lower its δ_2, the less it does (choosing instead to undercut by one notch). Firm 2 turns passive at low prices, either matching or standing ready to play the step up strategy at very low prices to avoid relenting.

wait out a war of attrition longer than high-δ firms.

[15] When $\delta_i = 0.99$, average price rises to 9.2, 11% higher than the $\delta_i = 0.95$ case.

Conditional on δ_1, changing δ_2 has little effect on average prices, which is about 7.7 for $\delta_1 = 0.95$. Although the cycle top is lower with lower δ_2, firm 2 does not aggressively undercut from moderately high prices as much, and these effects tend to cancel out. Lower δ_1 (with $\delta_1 > \delta_2$) leads to lower average prices as before.

Next I briefly consider an unequal sharing rule at equal prices, θ^i. Eckert[2003] also considers unequal sharing rules and solves for an analytical equilibrium in the context of constant marginal costs. An unequal sharing rule would arise if some consumers have a slight preference for one firm at equal prices, or perhaps if consumers randomly choose equally priced stores but the stores are owned in different numbers by the firms. I find that Edgeworth Cycles exist in equilibrium with unequal sharing rules and fluctuating marginal costs.

Consider the extreme case of $\theta^1 = 1$ (and $\delta_1 = \delta_2 = 0.95$), shown in Figure 6. In this case, firm 1 captures the entire market at equal prices. Because it can earn more profit at low prices than its opponent (simply by matching), firm 1 has less incentive to relent first. Therefore firm 2 emerges as the consistent price leader, relenting to a price of 12 each time.

The undercutting phase is especially interesting. Firm 1 chooses to match relatively high prices rather than undercut them, since it can steal the market at this higher price anyway. Firm 2 continues to undercut by one notch every time (unless it relents). When price falls to 7, firm 1 then aggressively undercuts to 4 to hasten firm 2's next relent. Average price along the cycle is similar to the baseline case at 8.3. While the top of the cycle is lower, high prices are matched and in effect longer.

With $\theta^1 = 1$ but lower discount factors, firm 1 matches through increasingly more moderate prices. When $\delta_i = 0.5$, for example (not shown), firm 2 undercuts all prices by one notch and firm 1 matches all prices down to 2 (or until firm 2 relents). The cycle is visually pleasing with a gentle downward slope that is linear through all prices, until matching adds convexity at the bottom. Because firm 1 matches everywhere, the cycle period is longer. Average market price falls to 6.8.

For intermediate values of $\theta^1 \in (0.5, 1)$, firm 2 continues to be the consistent price leader and continues to undercut by one notch at a time when not relenting. The follower firm 1 continues to aggressively undercut moderate prices to hurry firm 2's next relent. But in response to high prices, firm 1's behavior depends on θ^1. With θ^1 closer to 0.5, firm 1 more often undercuts high prices by one notch, and with θ^1 closer to 1, it more often matches.

In summary, Edgeworth Cycles are robust to a wide range of elasticities, discount factors, sharing rules, and price leading behavior. The shape and speed of the cycles depends upon the scenario.

2.4 Capacity Constraints

Edgeworth Cycles result from a firm's short term incentive to steal market share by undercutting price. However, this incentive is diminished if the firm is capacity constrained and cannot serve the entire market at the new lower price. Moreover, there is residual demand left over when an opponent's constraint binds so a firm may prefer a relatively high price for this reason. Many industries, including gasoline markets, have capacity constraints on at least some firms. Can Edgeworth Cycles still exist when one or more firms operate under capacity constraints? The answer is yes, provided the constraints are not too tight.

Demand is recast as

$$D^i(p_t^1, p_t^2) = \begin{cases} \min\{K^i, D(p_t^i)\} & \text{if } p_t^i < p_t^j \\ \min\{K^i, D(p_t^i) - \min[K^j, \theta^j D(p_t^j)]\} & \text{if } p_t^i = p_t^j \\ \min\{K^i, \max[0, D(p_t^i) - \min[K^j, D(p_t^j)]]\} & p_t^i > p_t^j \end{cases} \quad \text{for } i \neq j \tag{5}$$

where K^i is the maximum output, or capacity, of firm i. Capacities are exogeneously given.[16] Marginal cost is assumed to be c_t for all units up to K^i and infinity thereafter. I return to the situation where $\theta^i = \frac{1}{2}$ and $\delta^i = 0.95$, and allow K^i to take on integer values in $[0, 20]$.

I first consider the case of equal capacities, $K^1 = K^2$. For large capacities, $K^i \geq 10$, I easily find Edgeworth Cycles in equilibrium. However, for relatively small capacities, $K^i < 10$, the incentive to undercut is so diminished that only Focal Price equilibria appear instead.

With capacities large enough to support Edgeworth Cycles, the shape of the cycles itself depends on the tightness of the constraints. With $K^i \geq 15$, the best

[16] Some models of endogeneous capacity choice in other contexts include Kreps & Scheinkman (1983), Saloner(1987), and Kovenock & Roy (1998), and Reynolds & Wilson (2000).

response function diagram is identical to the basecase of Figure 3 and the cycles very similar.[17] As K^i falls toward 10, the undercutting phase becomes more linear and longer and the Edgeworth Cycles thus become less rapid and more asymmetric. The reason is that when a firm's capacity-constrained opponent undercuts the firm in the following period, it leaves residual demand for the firm. Having made a smaller undercut in the previous period allows the firm to serve that residual demand at a higher price. As a result, firms are more reluctant to undercut by multiple notches at once, and the undercutting phase is more linear. Because constrained firms make fewer profits at low prices, firms are more likely to relent a bit sooner as well. Relents are still to a price of 15.

In Figure 7, I show the case of $K^i = 10$. The sales-weighted average price in this case is 8.27, only 3% below that the baseline case. The smoother undercutting through the bottom half of the cycle tends to lower the average, but the fact that firms no longer set very low prices on the equilibrium path and the fact that the higher priced firm still makes sales tends to raise it.

Using a wide range of starting values, I find only Focal Price equilibria below $K^i = 10$.[18] The best response functions and equilibrium price path are shown in Figure 8 for the case of $K^i = 9$. In this case, the focal price is 7 and firms carry 28% excess capacity. The excess capacity allows firms to stand ready to undercut a price of 6 down to 5 (with positive probability) to punish defection

[17] The mixing probabilities differ.

[18] While I cannot claim each equilibrium to be unique, the fact that all attempted sets of starting values reach the same one suggests it that the derived equilibria may be most stable. See Maskin & Tirole (1988) for a discussion of multiple equilibria with constant marginal costs.

from the original focal price. From 5, firms play a war of attrition, mixing between matching or resetting the focal price.

As K^i decreases below 9, there are two opposing effects on the focal price. First, it becomes more costly to punish defections from a given focal price (by undercutting further), since a firm has less capacity. This suggests a lower focal price with lower K^i (reducing the incentive to defect in the first place) and firms producing closer to capacity.[19] However, because capacities are smaller, the market-clearing price that would occur with full capacity production rises. The former effect dominates with high K^i and the latter dominates at lower K.

For example, when $K^i = 8$, the focal price falls to 6 and excess capacity falls to 13%. Thereafter, focal prices begin to rise again, but at a relatively slower rate than if the first effect were not present. At $K^i = 7$, the focal price is 7 although excess capacity is now only 7%. Once $K^i \leq 5$, excess capacity falls to zero and only the second effect remains. The focal price is then the same as the full-capacity market clearing price ($p^i = 10$ at $K^i = 5$, $p^i = 12$ at $K^i = 4$, etc.).

For all cases ($K^2 = j,\ j = 0..19$) in which only one firm (firm 2) is constrained, I find Edgeworth Cycles in equilibrium (except $K^2 = 0$). The strong asymmetric shape, however, becomes less and less pronounced as the constraint on firm 2 tightens.

In all cases, the unconstrained firm 1 emerges as the consistent price leader

[19] The result is similar to Benoit & Krishna (1987) and Davidson & Denekere (1990) who show that excess capacities can support higher collusive prices in pricing supergames. Also see Brock & Scheinkman (1985), Rotemberg & Saloner (1986), Haltiwanger & Harrington (1991) and Bagwell & Staiger (1996) for discussions of maximum sustainable cartel prices in various contexts.

in relenting. The cost to relenting first is relatively lower for firm 1 since it will still serve residual demand in the current period at its new higher prices. Given that firm 1 will relent first, firm 2 plays a familiar strategy: undercutting less aggressively near the top, more aggressively near the bottom, and then matching or stepping up at very low prices. For $K^2 \geq 15$, the best response function looks very similar to that in Figure 4.

As constraints tighten on firm 2 from $K^2 = 15$ down to 6, the cycle becomes more rapid, smaller in amplitude, and less asymmetric. Since firm 1 is willing to relent sooner (since more residual demand is available), firm 2 is willing to undercut more aggressively through lower prices to hasten it. Firm 1 relents to a price closer to the monopoly price each time. In Figure 9, I report the case of $K^2 = 7$ showing the rapid, low amplitude cycle.

There is one-time discrete upward shift in cycle duration when K^2 falls from 6 to 5. The capacity constraint is now tight enough that firm 2 elects to match very high prices rather than undercut, since it can still sell to capacity when matching. Duration increases from 4 periods on average to 6. However, firm 2 is still increasingly aggressive in undercutting moderate and low prices, and cycle duration begins to decrease again as K^2 falls toward zero.

It may at first appear that the cycle is no longer present by the time $K^2 = 2$. Firm 1 always sets price to 9 or 10 depending on cost, and firm 2 always sets a price of 7. But these are not different "focal" prices – in fact, it is a hyper cycle where firm 2 undercuts on every turn by just enough to force firm 1 to relent on every turn. While firm 2 could have sold to capacity by just matching price, firm 1 would have then responded with an undercut and leave zero sales for firm 2 in the next period. Hence, it undercuts to force the relent. (Off the equilibrium path, firm 2 matches most prices as one may have expected.) It is similar for $K^2 = 1$. Once $K^2 \cong 0$, we observe firm 1 simply setting the static monopoly price each period.

Average sales-weighted prices edge up as well with lower K^2. At $K^2 = 7$, it is 7.3 but then rises to 9.8 when $K^2 = 3$ and ultimately 10.5 (the static monopoly price) when $K^2 = 0$.

In summary, when firms have identical capacity constraints, I find Edgeworth Cycles in equilibrium unless the constraints are too tight. I find focal prices that first fall and then rise as constraints tighten even more. When only one firm is capacity constrained, I again find Edgeworth Cycles at all constraint levels except very close to zero. The cycles become more rapid, smaller in amplitude and less asymmetric with a tighter constraint. High-price matching can also occur with relatively tight constraints.

3 Differentiated Product Model

Most industries are characterized by at least some product differentiation. With differentiation, the incentive to undercut is diminished since some consumers will remain loyal to the opponent's product. Can Edgeworth Cycles still exist when the goods are differentiated? The answer is yes, provided the differentiation is not too great. I illustrate with an alternating moves Hotelling model of horizontal differentiation.

Consumers tastes are uniformly distributed over a univariate product space that has support [0,1]. Firm 1's product is located at point 0 and firm 2's product is at point 1. Each consumer h has unit demand and receives utility in time t of

$$u_t^{ih} = \begin{cases} v^i - p_t^i - \tau z_t^{ih} & \text{if purchase from firm } i \\ 0 & \text{if do not purchase at all} \end{cases} \tag{6}$$

where v^i is the intrinstic value of the product of firm i, p_t^i is the price charged by firm i, and z^{ih} is the distance in product space between consumer h's most preferred product (which may not be produced) and product of firm i. Let τ is the disutility per unit of distance between the preferred and the purchased product.[20] Since no firm will price above the intrinsic value of its product, let the top of the grid $sup(p^i) = v^i$.

If the prices are low enough that all consumers make a purchase, which occurs when $(v^1 + v^2) - (p_t^1 - p_t^2) \geq \tau$, the market share of firm i is given by

$$s_t^i(p_t^1, p_t^2) = \begin{cases} \frac{1}{2} + \frac{(v^i - v^j)}{2\tau} - \frac{(p_t^i - p_t^j)}{2\tau} & \text{if } \left|(v^i - p_t^i) - (v^j - p_t^j)\right| \leq \tau \\ 1 & \text{if } (v^i - p_t^i) - (v^j - p_t^j) > \tau \qquad i \neq j \\ 0 & \text{if } (v^j - p_t^j) - (v^i - p_t^i) > \tau \end{cases} \tag{7}$$

If prices are high enough that not all consumers are served at those prices, then

[20] For a consumer whose preferred product is at $y^h \in [0,1]$, $z^{1h} = y^h$ and $z^{2h} = 1 - y^h$.

firm i's share is

$$s^i_t(p^i_t) = \frac{v^i - p^i_t}{\tau} \tag{8}$$

Letting H be the total number of consumers in each period, current period profits to firm i is

$$\pi^i_t(p^1_t, p^2_t, c_t) = H * s^i_t(p^1_t, p^2_t) * (p^i_t - c_t) \tag{9}$$

which is substituted into the equations for V^i and W^i given above.

By construction in this model, aggregate elasticity is zero while all consumers are served, and $-\frac{2}{t}$ when not all served.

3.1 Cycles in Differentiated Markets

I find Edgeworth Cycles in equilibrium for mildly differentiated goods markets, but not for more highly differentiated goods markets. To illustrate, assume $v^1 = v^2 = 10$, so the price grid is $p^i = \{x\},\ x = 0..10$, the marginal cost grid as before, $\delta_i = 0.95$ and $\theta^i = \frac{1}{2}$. Now consider variation in the degree of differentation, τ.[21]

First, because an undercut of one notch steals the entire market for all $\tau \leq 1$, all such cases are identical and the Edgeworth Cycles found replicate the homogeneous goods situation (with a zero aggregate elasticity demand curve). Firms undercut by one notch at a time through all prices, except from 6 and 4 where it is by two, and they may match at prices of 2 and 1. Firms relent back

[21] I examine the set $\tau = \{1 + 0.05x,\ 2 + x\}$, $x = 0..10$.

to 10 from prices as high as 4. (Case not shown.)

When $\tau > 1$, the market is meaningfully differentiated in that not all consumers will switch in response to a minimum undercut. In Figure 10, I describe the case of $\tau = 1.1$. Edgeworth Cycles can indeed exist in mildly differentiated goods markets. Given this τ, an undercut of one notch steals an additional 45%, of the market, leaving the other firm with 5%.[22] Average sales weighted price is 5.2, well below the static monopoly price of $p^m(v,\tau) = v - \frac{\tau}{2} = 9.45$, but above the average static competitive price of 2.1.

However, with a bit more differentiation ($\tau \geq 1.25$), Edgeworth Cycles are replaced by Focal Price equilibria. The case of $\tau = 1.25$, when a one-notch undercut steals an additional 40% of the market, is shown in Figure 11. The focal price is 7. Off the equilibrium path, firms stand ready to undercut a price of 6 down to 3 (with positive probability) to punish defection from the original focal price. From 3, firms mix between matching or resetting the focal price.

As τ increases higher above 1.25 (but still below 6 when the no purchase option will become a factor), there are two opposing effects on the equilibrium focal price. First, it becomes more costly to punish defections from a given focal price since a greater undercut would be needed to have an impact. This works to reduce the sustainable focal price. However, consumers are willing to pay a higher price as products become more differentiated. The former effect dominates with lower τ and the latter dominates at higher τ. For example, when $\tau \geq 2$, firms no longer credibly threaten to further undercut an undercut from

[22] This best response function diagram is identical to the previous $\tau \leq 1$ case. Small differences in the mixing probabilities still exist.

the focal price...only to match it with positive probability. Hence, the focal price falls to 5 when $\tau = 2$ but then climbs back up to 7 as τ rises up to 6.[23]

When τ exceeds 6, firms compete against the no purchase option rather than each other and prices must fall gradually to provide zero utility to the consumer at $\frac{1}{2}$. Once $\tau \geq 9$, firms no longer serve the middle consumers in equilibrium and each sets its own monopoly price of 5 or 6 (depending on cost) thereafter.[24]

Therefore, Edgeworth Cycles are robust to a small amount of product differentiation. However, we should not expect to find Edgeworth Cycles in markets that sell highly differentiated products. Focal prices are more likely.

4 Bertrand Triopoly

Earlier studies have focused on the duopoly Bertrand model for its analytical tractability. But can Edgeworth Cycles still exist when there are more than two firms? In this section, I show that Edgeworth Cycles are still an equilibrium outcome in a Bertrand triopoly. However, in contrast to the two firm case, coordination problems in relenting can occur. There are "delayed starts" in which following firms do not follow immediately and even "false starts" in which the relenting firm returns to the bottom when others do not follow soon enough.

In the three firm game, each firm can adjust its price every third period and its price is fixed for the following two. Firm 1 adjusts its price in period t, firm 2

[23] At $\tau = 2$, a one-notch undercut steals an additional 25% of the market. At $\tau = 6$, it steals only 8%.

[24] Recall that when all consumers are served, the monopoly price is $p^m = v - \frac{t}{2}$ and falls with τ. When not all consumers are served, which requires $\tau > v - c$, the static monopoly price is $p^m = \frac{(v+c)}{2}$ and is independent of τ.

in $t+1$, and firm 3 in $t+2$ before returning to firm 1 again. The value functions for firm 1 are:

$$V^1(p^2_{t-2}, p^3_{t-1}) = \underset{c}{\mathrm{E}} \left(\max_{p_t} \left[\pi^1_t(p_t, p^2_{t-2}, p^3_{t-1}, c_t) + \delta_1 W^1(p^3_{t-1}, p_t) \right] \right) \tag{10}$$

$$W^1(p^3_{s-2}, p^1_{s-1}) = \underset{c}{\mathrm{E}} \left(\underset{p_s}{\mathrm{E}} \left[\pi^1_s(p^1_{s-1}, p_s, p^3_{s-2}, c_s) + \delta_1 U^1(p^1_{s-1}, p_s) \right] \right) \tag{11}$$

$$U^1(p^1_{r-2}, p^2_{r-1}) = \underset{c}{\mathrm{E}} \left(\underset{p_r}{\mathrm{E}} \left[\pi^1_r(p^1_{r-2}, p^2_{r-1}, p_r, c_r) + \delta_1 V^1(p^2_{r-1}, p_r) \right] \right) \tag{12}$$

The value function $V^1(p^2_{t-2}, p^3_{t-1})$ is the expected future profits of firm 1 at a time t when it is firm 1's turn to adjust its price, given that firm 2 set price p^2_{t-2} two periods before ($p^2_{t-2} = p^2_{t-1} = p^2_t$), firm 3 set price p^3_{t-1} in the previous period ($p^3_{t-1} = p^3_t = p^3_{t+1}$), and c_t is not yet known. Similarly, the value function $W^1(p^3_{s-2}, p^1_{s-1})$ is firm 1's expected future profits at a time s when it is firm 2's turn to adjust its price and $U^1(p^1_{r-2}, p^2_{r-1})$ is its expected future profits at time r when it is firm 3's turn to adjust price. V^2, W^2, U^2, V^3, W^3, and U^3 are similarly defined.

The per period profit function is

$$\pi^i_t(p^1_t, p^2_t, p^3_t, c_t) = \begin{cases} D(p^i_t)(p^i_t - c_t) & \text{if } p^i_t < p^j_t \text{ and } p^i_t < p^k_t \\ \frac{\theta^i}{\theta^i + \theta^j} D(p^i_t)(p^i_t - c_t) & \text{if } p^i_t = p^j_t,\ p^i_t < p^k_t \\ \theta^i D(p^i_t)(p^i_t - c_t) & \text{if } p^i_t = p^j_t = p^k_t \\ 0 & \text{if } p^i_t > p^j_t \text{ or } p^i_t > p^j_t \end{cases} \quad \text{for } i \neq j \neq k \tag{13}$$

Therefore the lowest priced firm serves the market, or if two or more firms have the lowest price, they split the market according to the sharing rule $\theta^i,\ i = 1..3$. Again, to fix an example for discussion, let $D(p) = 20 - p$ and $\theta^i = \frac{1}{3}$. Because of the additional computational demands of the three firms model, I allow for 200 points on the cost grid, $c_t = \{x/200\},\ x = 1..200$.[25]

If we believe that the interval of time between firm i's moves should not change regardless of how many firms there are, then firms should care about its profits three periods hence in the three firm model as it would two periods hence in the two firm model. To adjust for this, I use a discount factor of $\delta_1 = \delta_2 = \delta_3 = 0.967$ in the base case.[26] If instead we believe that the time interval between consecutive price changes by different firms should not change, $\delta_i = 0.95$ would again be used. Results are very similar between the two.

4.1 Cycles in Triopoly

I find Edgeworth Cycles continue to be an equilibrium in triopoly. An example of the market price path is given in the top panel of Figure 12. Best response functions are not shown.

The process underlying the three-firm cycle is the same as the two firm case. If the minimum price of the other two firms are greater than 14, the third firm undercuts to 14 and captures the market. The undercut is to a price substantially greater than the monopoly price of 10 to 11.

From there, the active firm undercuts the lowest priced firm by one notch on

[25] Mixing probabilities can be replicated to within 0.0025.

[26] $0.967^3 \cong 0.95^2$

the grid in response to prices from 14 all the way down to 7. Hence, undercutting is orderly and firms take turns serving the market through the most profitable range of prices. Once the minimum price of the other two reaches 6, a firm may undercut by one notch as usual if cost is high, but would instead aggressively undercut to 3 if cost is low. The aggressive play pushes the market through the low prices faster and also pressures opponents into relenting earlier. If the minimum price of the other firms reaches 4 or lower, the active firm – if it does not relent – responds by undercutting by one notch or by passively matching the lowest price. Matching can occur down to a price of 1 on the equilibrium path, as observed in the last trough of the diagram. When a firm attempts to lead prices back to the top of the cycle, it sets its price to 16.

4.2 Delayed and False Starts

I say "attempts to lead prices back to the top" because in contrast to the two firm case, immediate following by the other two firms is no longer guaranteed. There can be "delayed starts" in cycle resetting, and in some instances "false starts."

A delayed start occurs when a firm must wait more than one turn (three periods) for others to follow it to the top of the cycle. The top panel of Figure 12 shows an example of a delayed start around the third peak. It is easily identified by an extended flat line at the top of the cycle. After a high cost draw and facing a minimum opponent price of 3 (previously set by firm 2), firm 3 is the first to relent to the top. But with a low cost draw in the next period, firm 1 finds it more profitable to undercut to a price of 2 rather than follow firm 3. The result is that firm 3 sits at the top of the cycle and makes no sales for two consecutive turns (six periods in all) instead of the usual one. Longer delays can also occur.

A false start occurs when a firm abandons its effort to reset prices higher altogether and returns immediately to the bottom with the other firms. Two examples of false starts are shown in the bottom panel of Figure 12. They take on the appearance of double peaks along the price path – the first and third main peaks show false starts. (The reader will note the second peak is another example of a delayed start.)

Consider the first false start in the figure. In this case, firm 2 relents first after facing a low minimum opponent price of 3 (set by firm 1) and suffering a very high cost draw. Unfortunately for firm 2, firm 3 and then firm 1 receive favorably low cost draws in the following two periods and – rather than follow – continue to undercut each other. Had its next cost draw been high, firm 2 would have remained at the top for another turn, as occurs with a delayed start. In this example, however, its receives a cost draw low enough that it is more profitable to abandon its position at the top of the cycle and match firm 1's price at the bottom. This action delays the resetting of the cycle but greatly increases the probability (up from zero) that the others will relent first.

In simulations, false starts occurred in 6% of all attempted relents (ie. all peaks) and delayed starts occurred in an additional 13%.

The coordination problems of delayed starts and false starts make it more challenging and costly to be first to reset the cycle. As a result, firms hesitate in relenting and market prices tend to be fall closer toward the band of marginal costs.[27] This is easily seen in a comparison of Figures 3 and 12. The average market price on the equilibrium path is now 7.2, 16% lower than in the two firm case.

When $\delta^i = 0.63$, average market price falls further to 5.51.[28] Firms undercut by only one notch even through lower prices and extended matching at prices of 1 and 2 are commonplace. False starts and delayed starts are also more common, and the delays are often two or more turns instead of one.

In summary, Edgeworth Cycles can exist in equilibrium even in a triopoly setting. They progress similar to the two firm case, except that coordination problems in cycle resetting can occur. The emergence of a consistent price leader – or more specifically a price leading order – is useful for reducing these coordination problems.

5 Conclusion

Prior theoretical work on Edgeworth Cycles has largely been restricted to the symmetric homogeneous-goods Bertrand duopoly case for its analytic tractability. In this article, I employ a computational approach in order to search for Edgeworth Cycles under a wide assortment of competitive models involving such real world complications as differentiated goods, capacity constraints, and

[27] With lower discount factors (for example, $\delta_i = 0.5$), the troughs become deeper and coordination problems become more prevalent.

[28] Comparable to the $\delta^i = 0.5$ two firm case, as $(0.63)^3 = (0.5)^2$.

a triopoly market structure.

In a framework that allows for fluctuating marginal costs, I show that Edgeworth Cycles can exist in many scenarios beyond the simple Bertrand mold. They can exist in a differentiated goods market provided the differentiation is not too great. They can exist in capacity constrained markets as long as the constraints are not too tight. If differentiation is sufficiently great or capacity constraints sufficiently tight, focal prices appear instead. Edgeworth Cycles can also exist in triopoly situations, although firms deal with coordination problems – delayed starts and false starts – that do not occur in the two firm model. I also find the existence of Edgeworth Cycles is robust to assumptions about discount factors, elasticities, sharing rules, and price leading behavior. The shape of the cycle is impacted by the aggressiveness of the firms and varies across scenarios.

Therefore, the discovery of Edgeworth Cycles in a particular market and its shape carries important information about the competitive environment in that market.

This article was motivated by the discovery of apparent Edgeworth Cycles in some Canadian retail gasoline markets. This article shows that consumers may still consider gasoline as a differentiated product, but the differentiation must be relatively small. Capacity constraints on major firms can exist but cannot be too tight (the symmetric constraint case). Small independents can be very tightly constrained but still induce Edgeworth Cycles (the asymmetric constraint case) – the constraint just cannot be very close to zero.

The absence of coordination problems in these retail gasoline markets suggests the emergence of a consistent price leader and relenting order. Noel[2003b] argues this is true in the case of Toronto. In Figure 1 above, we observe the major firm depicted always relents before the independent and there are no false starts. We also observe that the independent exhibits follower attributes, often aggressively undercutting through moderate prices and then passively matching (or making smaller undercuts) at low prices.

This article is an important step to understanding the range of environments condusive to Edgeworth Cycle activity. Edgeworth Cycles are indeed a real and important economic phenomenon in need of more study. This is especially true as new technologies create increasingly real-time markets (in electricity, long distance telephone, internet shopping, etc.) where relatively homogenous products, frequent purchases, and low switching costs may work to generate Edgeworth Cycles in these areas.

References

[1] Allvine, F. & J. Patterson (1974). "Highway Robbery, An Analysis of the Gasoline Crisis", Bloomington: Indiana University Press, pp. 192-197, p.243.

[2] Bagwell, K. and R.W. Staiger (1997). "Collusion over the Business Cycle", RAND Journal of Economics, 28.

[3] Benoit, J.P. and V. Krishna (1987). "Dynamic Duopoly: Prices and Quantities", Review of Economic Studies, 54.

[4] Brock, W. and J. Scheinkman (1985). "Price Setting Supergames with Capacity Constraints", Review of Economic Studies, 52.

[5] Castanias, R. and H. Johnson (1993). "Gas Wars: Retail Gasoline Price Fluctuations", Review of Economics and Statistics, 75.

[6] Davidson, C. and R.J. Deneckere (1990). "Excess Capacity and Collusion", International Economic Review, 31.

[7] Eckert, A. (2002). "Retail Price Cycles and Response Asymmetry", Canadian Journal of Economics, 35.

[8] Eckert, A (2003). "Retail Price Cycles and the Presence of Small Firms", International Journal of Industrial Organization, 21.

[9] Edgeworth, F.Y. (1925). "The Pure Theory of Monopoly" in Papers Relating to Political Economy, Volume 1", McMillan.

[10] Ericson, R. and A. Pakes (1995). "Markov-Perfect Industry Dynamics: A Framework for Empirical Work", Review of Economic Studies, 62-1.

[11] Godby, R., A.M. Lintner, T. Stengos, and B. Wandschneider (2000). "Testing for Asymmetric Pricing in the Canadian Retail Gasoline Market", Energy Economics, 22.

[12] Haltiwanger, J. and J.E. Harrington (1991). "The Impact of Cyclical Demand Movements on Collusive Behavior", RAND Journal of Economics, 22.

[13] Kovenock, D. and S. Roy (1998). "Dynamic Capacity Choice in a Bertrand-Edgeworth Framework", Journal of Mathematical Economics, 29.

[14] Kreps, D. and J. Scheinkman (1983). "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes", Bell Journal of Economics, 14.

[15] Kruse, J.B., S. Rassenti, S. S. Reynolds, V. L. Smith (1994). "Bertrand-Edgeworth Competition in Experimental Markets", Econometrica, 62.

[16] Maskin E., and J. Tirole (1988). "A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves and Edgeworth Cycles", Econometrica, 56-3.

[17] Maskin, E. and J. Tirole (2001). "Markov Perfect Equilibrium", Journal of Economic Theory, 100-2.

[18] Noel, M. (2003a). "Edgeworth Price Cycles, Cost-based Pricing and Sticky Pricing in Retail Gasoline Markets", UCSD Working Paper.

[19] Noel, M. (2003b). "Edgeworth Price Cycles: Firm Interaction in the Toronto Retail Gasoline Market", UCSD Working Paper.

[20] Pakes, A. (2000) "A Framework for Applied Dynamic Analysis in I.O", NBER Working Paper 8024.

[21] Pakes, A. and P. McGuire (1994). "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model", RAND Journal of Economics, 25-4.

[22] Pakes, A. and P. McGuire (2001). "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality", Econometrica, 69.

[23] Reynolds, S. and B. Wilson (2000). "Bertrand-Edgeworth Competition, Demand Uncertainty, and Asymmetric Outcomes", Journal of Economic Theory, 92.

[24] Rotemberg, J. and G. Saloner (1986). "A Supergame-Theoretic Model of Price Wars during Booms", American Economic Review, 76.

[25] Slade, M. (1987). "Interfirm Rivalry in a Repeated Game: An Empirical Test of Tacit Collusion", Journal of Industrial Economics, 35.

[26] Slade, M. (1992). "Vancouver's Gasoline Price Wars: An Empirical Exercise in Uncovering Supergame Strategies", Review of Economic Studies, 59.

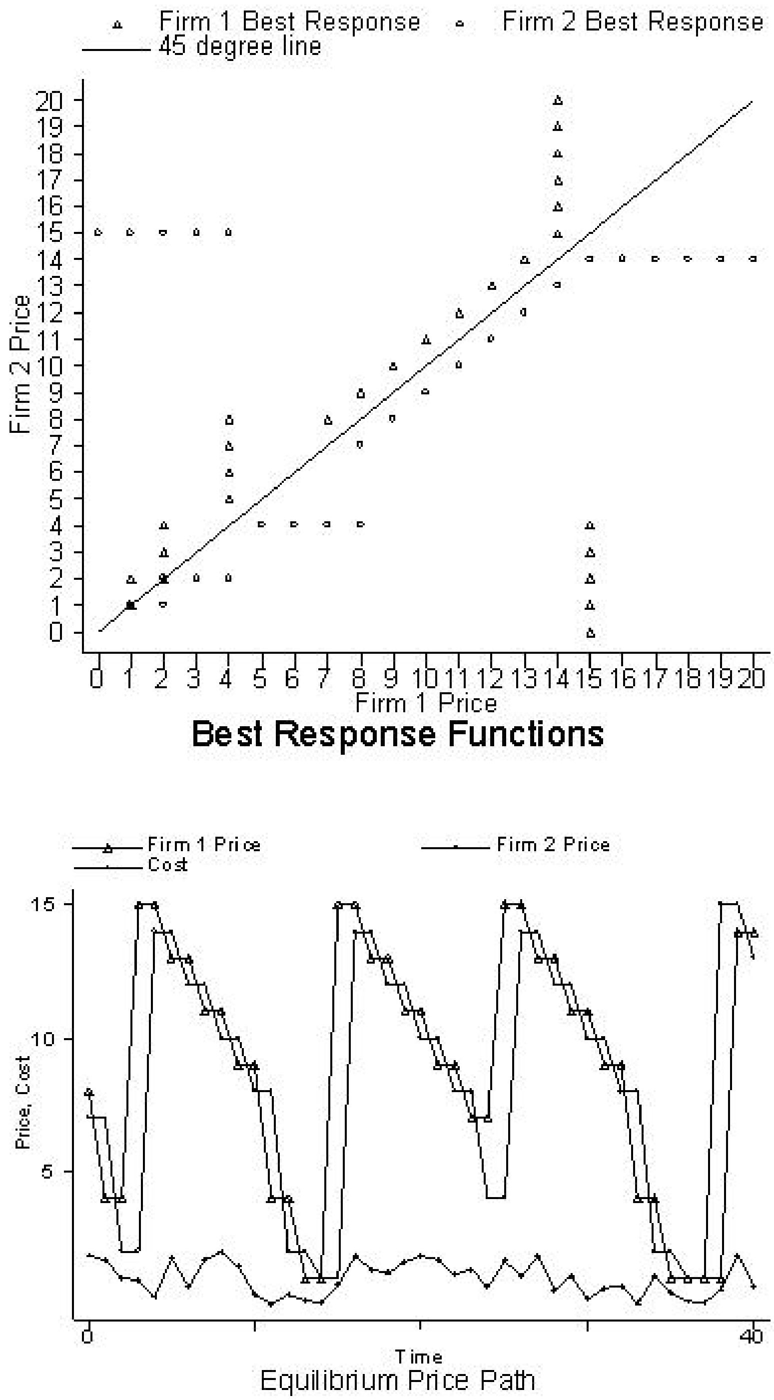

Figure 3: Symmetric Bertrand Duopoly

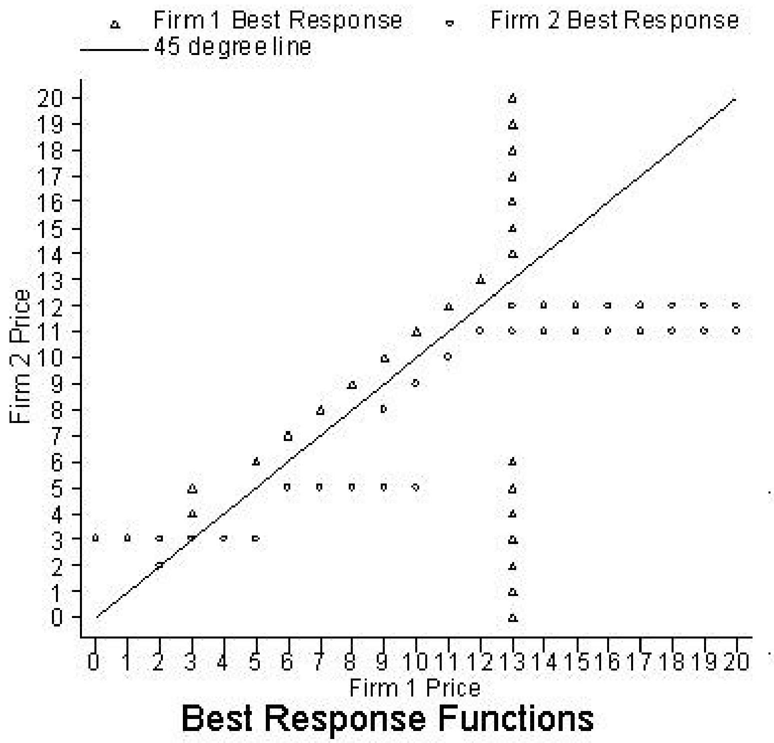

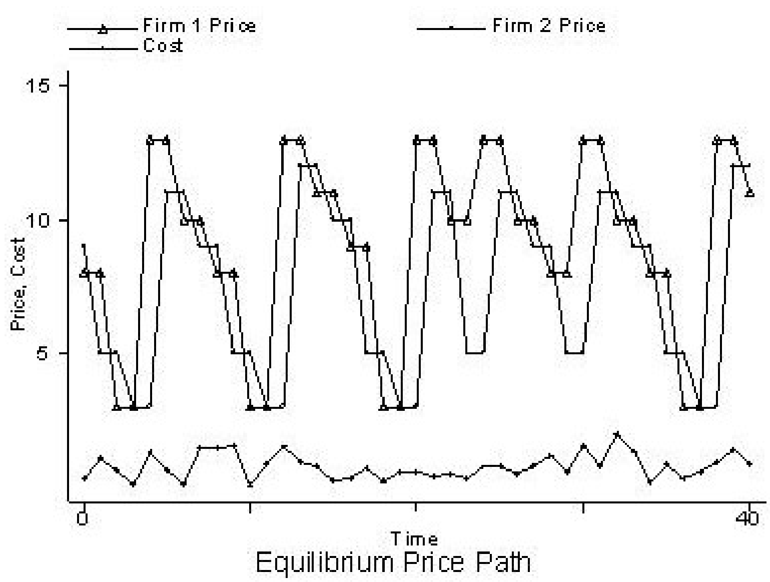

Figure 4: Consistent Price Leader

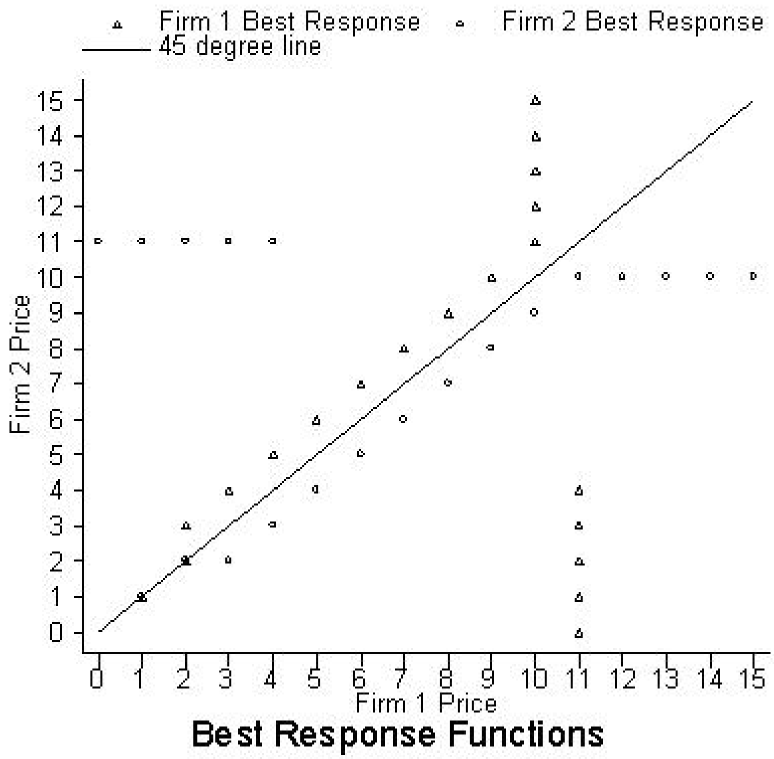

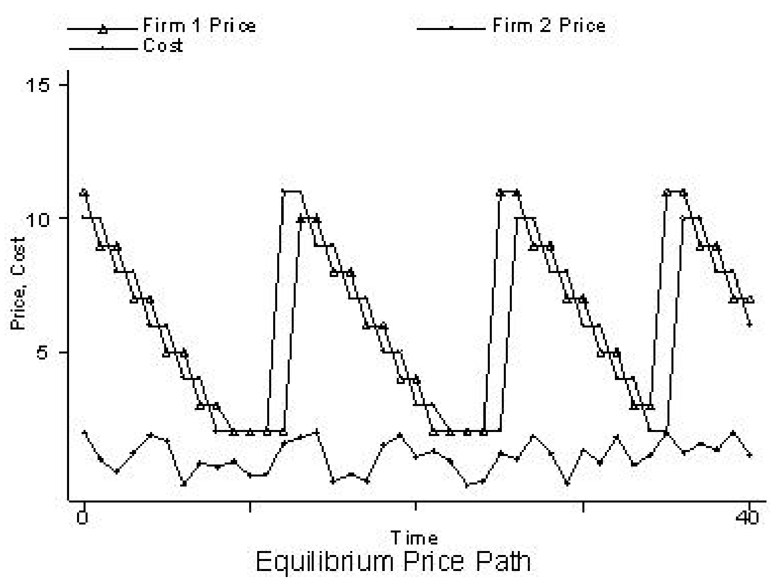

Figure 5: Elastic Demand

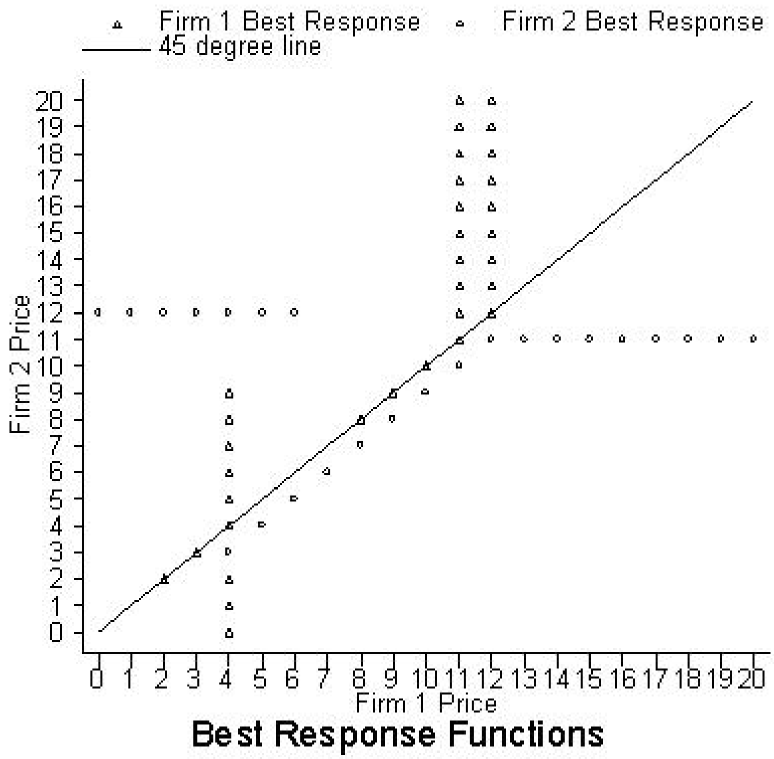

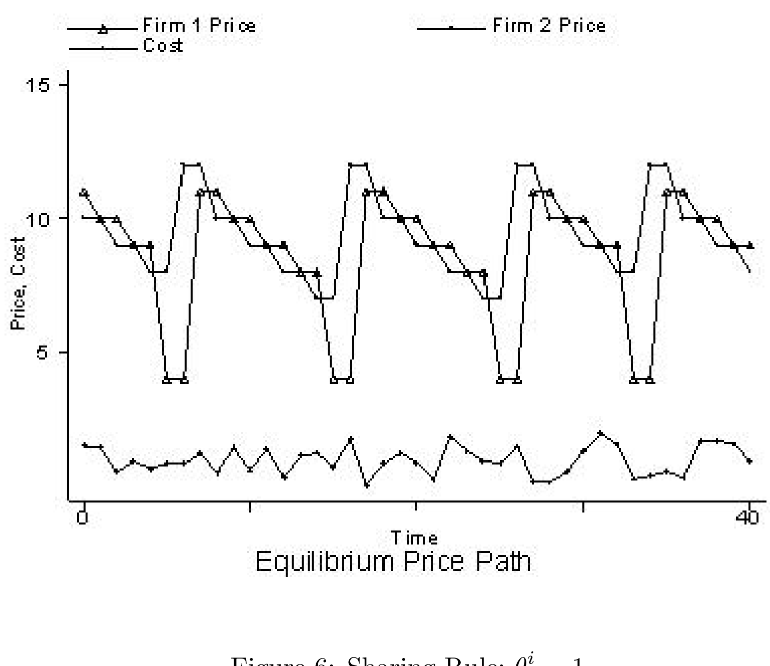

Figure 6: Sharing Rule: $\theta^i = 1$

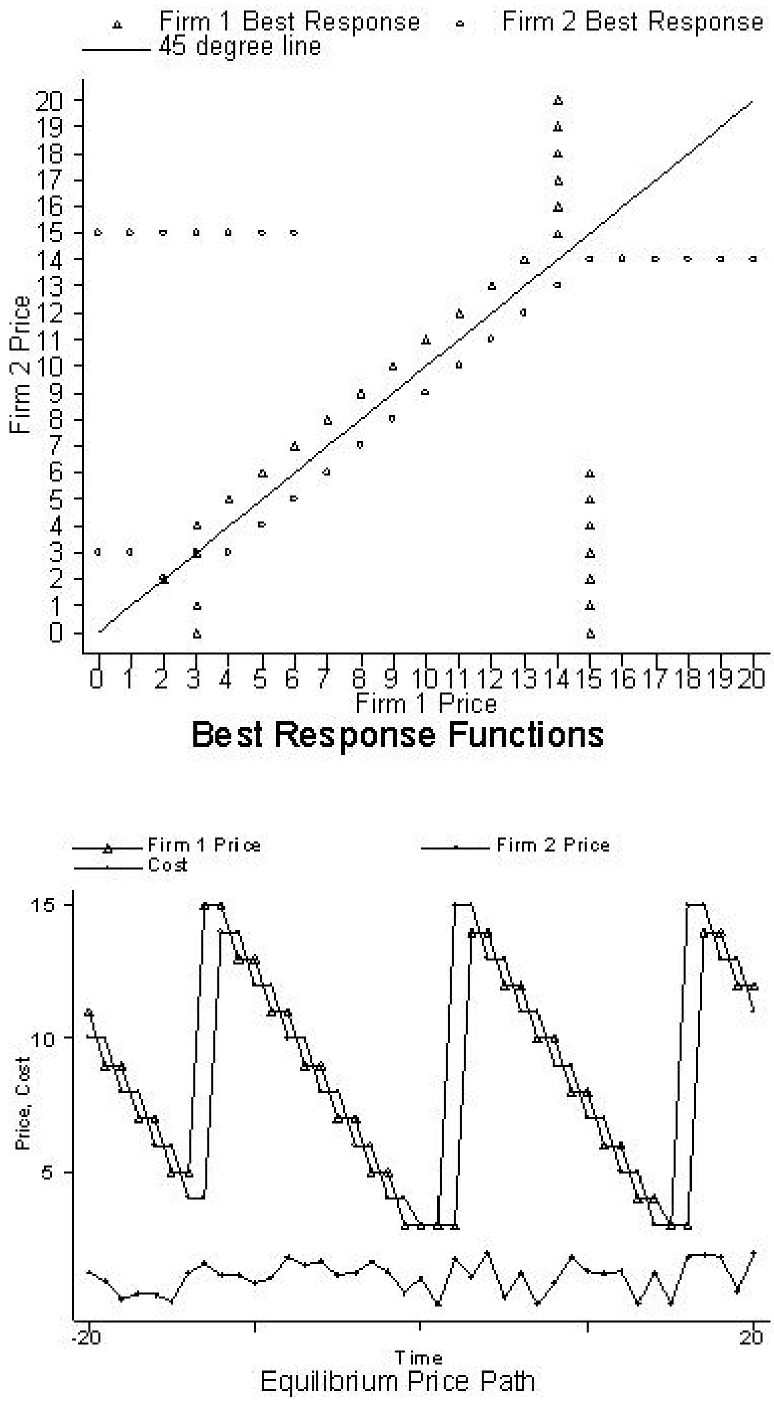

Figure 7: Capacity Constraints: $K^i = 10$

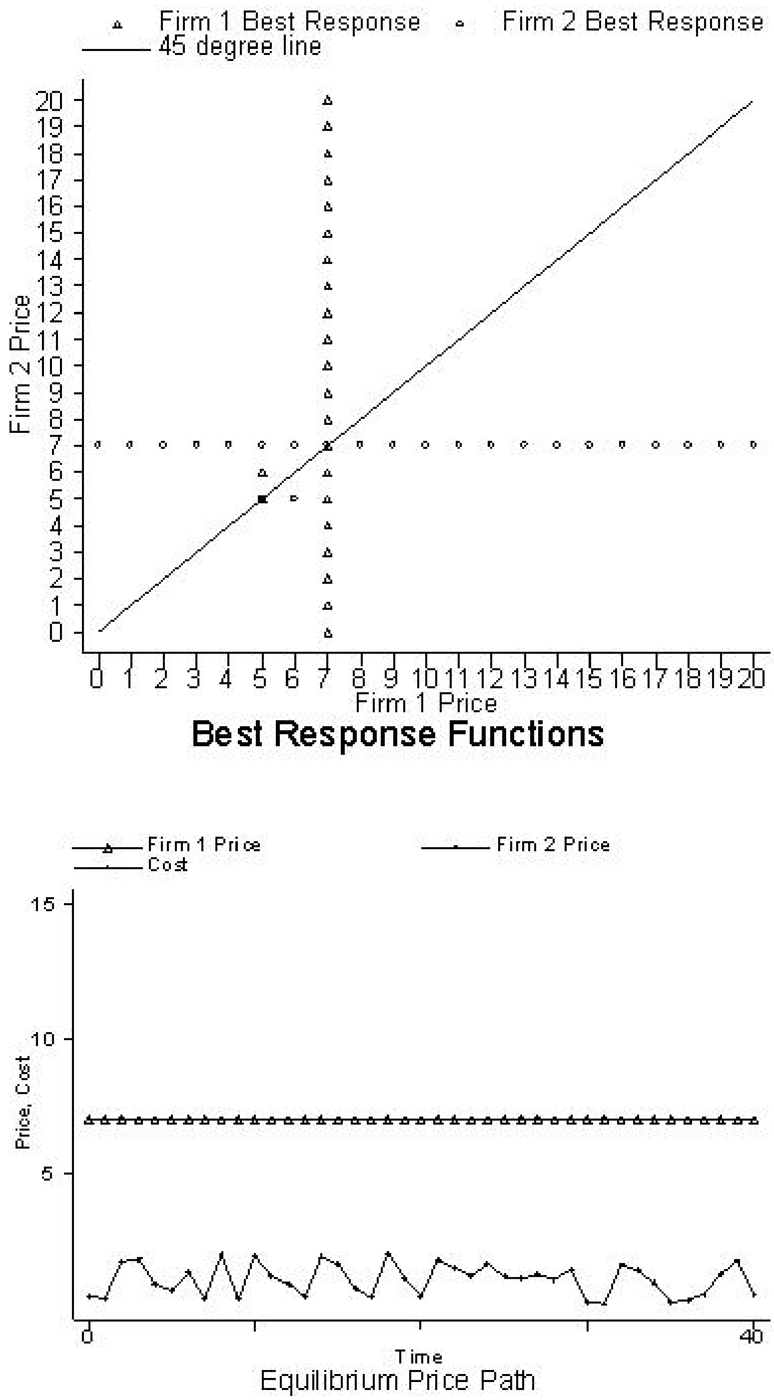

Figure 8: Capacity Constraints: $K^i = 9$

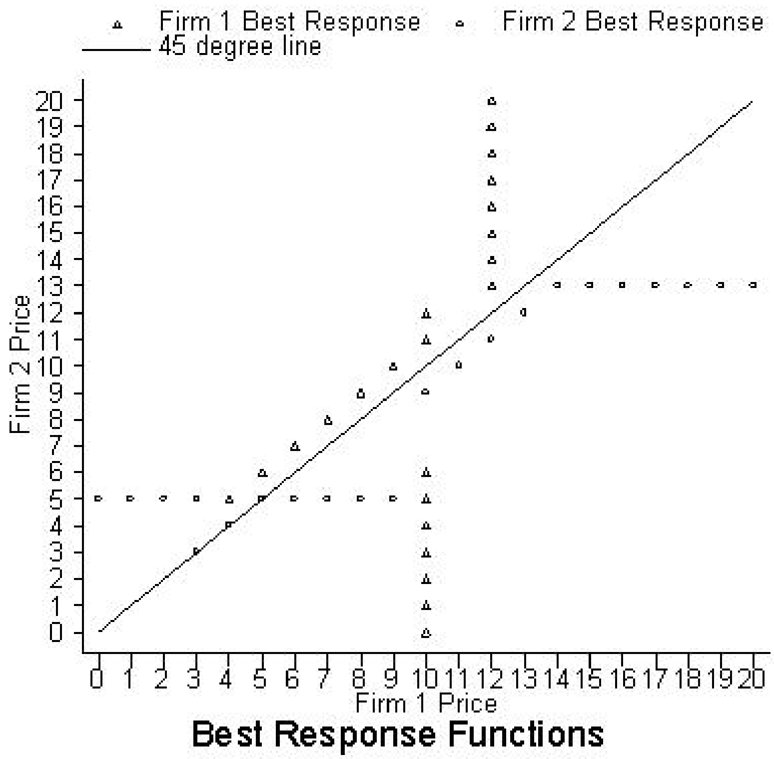

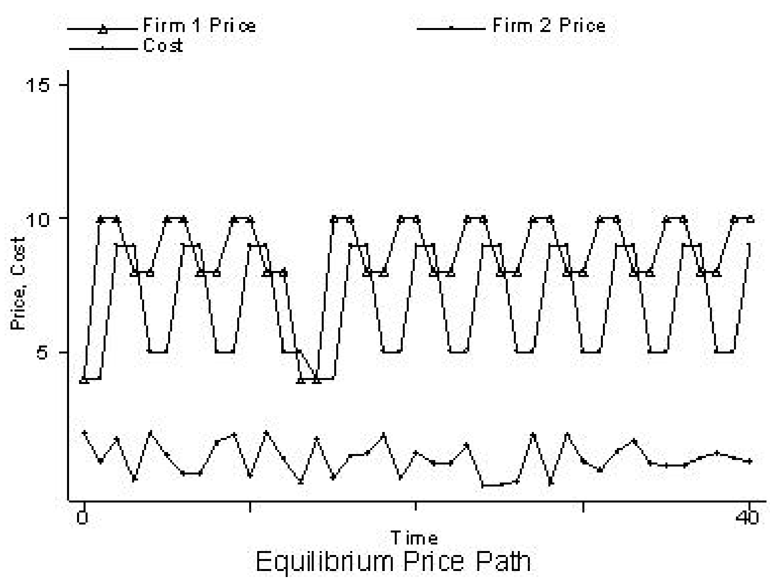

Figure 9: Capacity Constraint: $K^2 = 7$

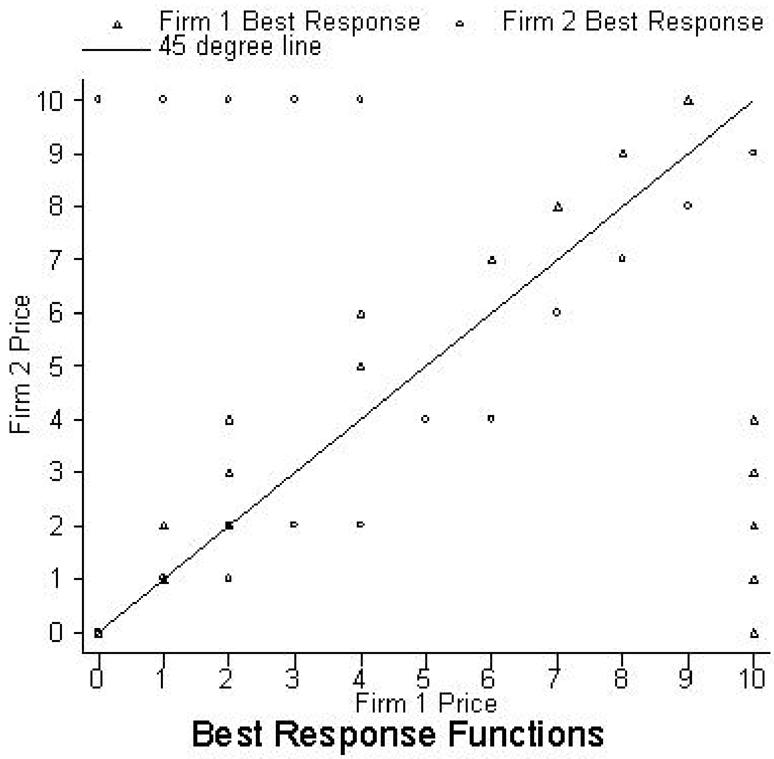

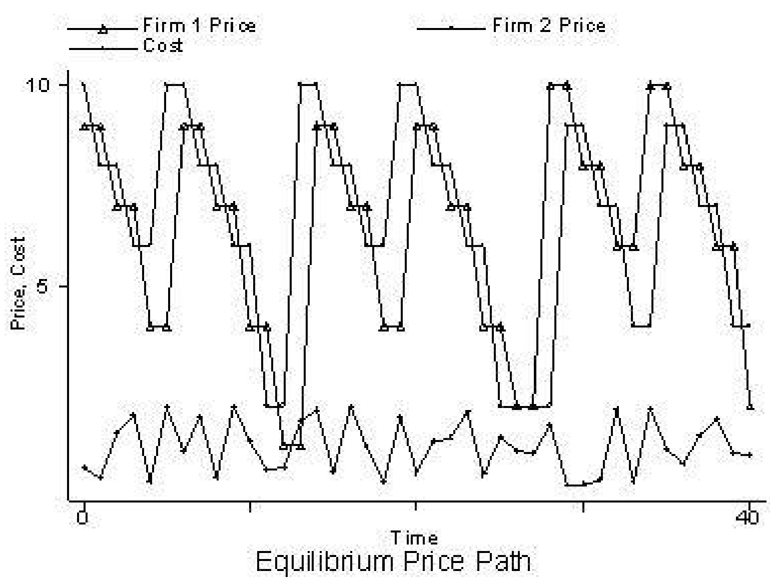

Figure 10: Differentiated Goods: $\tau = 1.1$

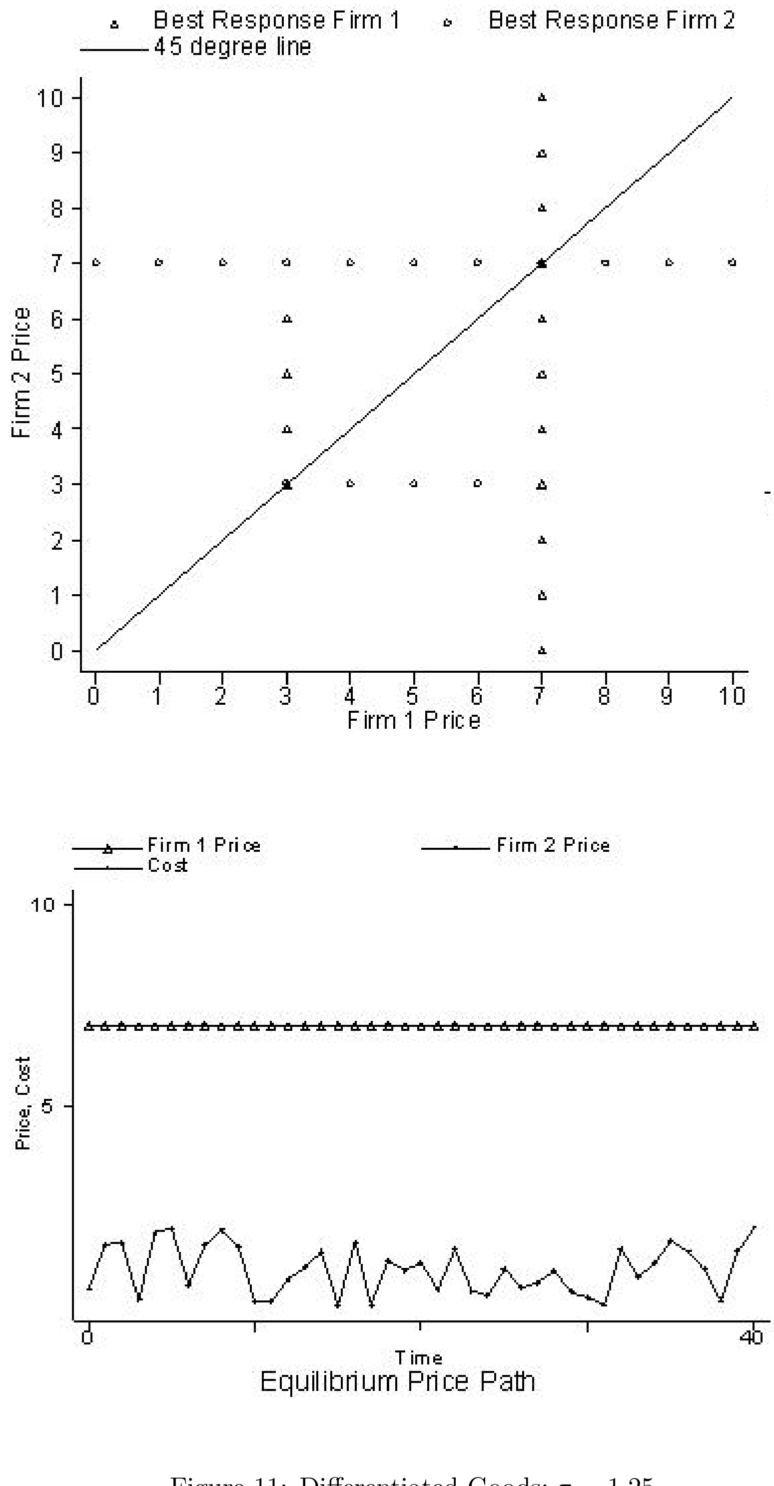

Figure 11: Differentiated Goods: $\tau = 1.25$

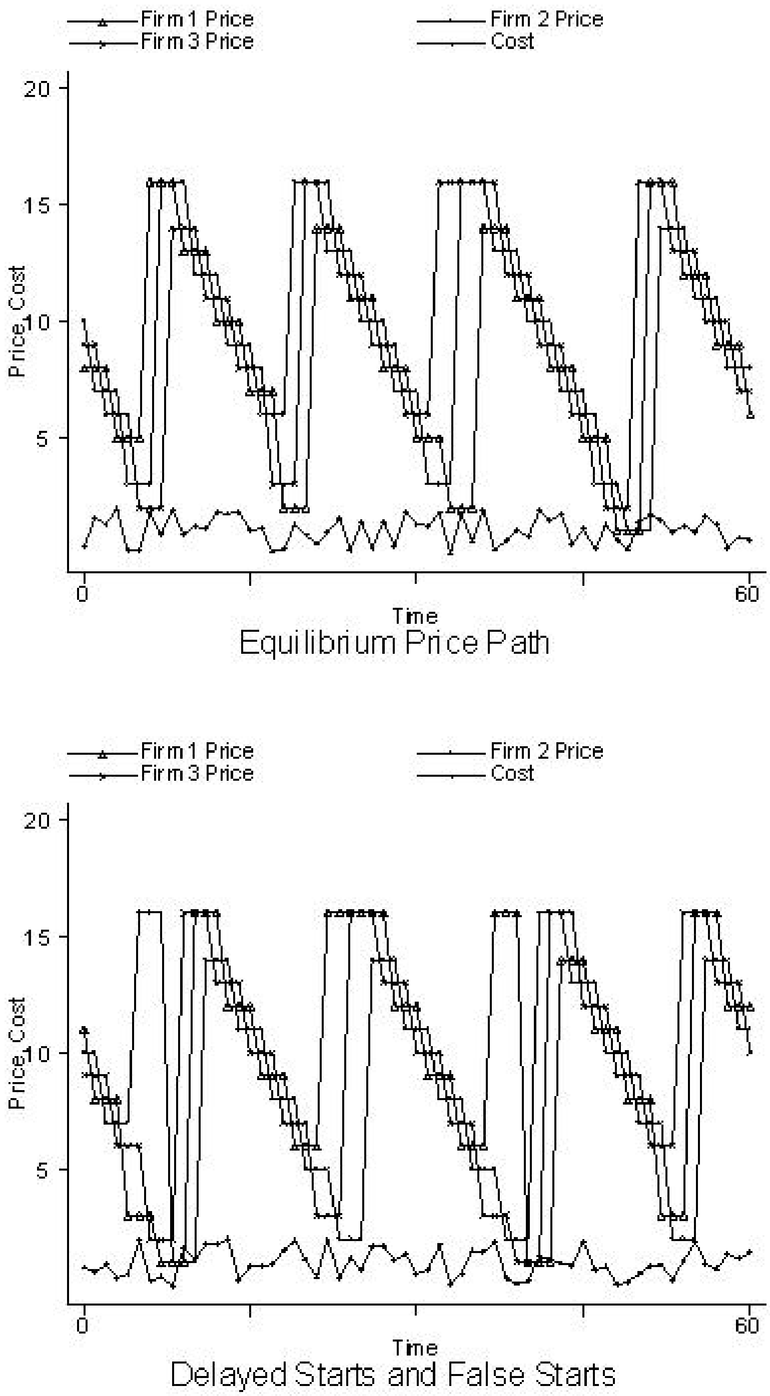

Figure 12: Symmetric Bertrand Triopoly