

The robustness of industrial commodity oligopoly pricing strategies

EPRG Working Paper 1522

Cambridge Working Paper in Economics 1540

David M. Newbery and Thomas Greve

Abstract

Industrial commodity markets are typically oligopolies in which firms set prices but need to make sunk and durable investment decisions, requiring them to make predictions of future prices. Mark-up pricing models are attractive both for setting prices and predicting future prices for investment analysis. Simple algorithms can find Nash equilibria, but these equilibria are not necessarily robust. This paper examines fixed and proportional mark-up models and demonstrates that they are robust to single firm Nash Cournot deviations but not against more sophisticated deviations in the deterministic case. Cournot equilibria are not robust under demand uncertainty, where proportional mark-up models emerge as the most robust when marginal costs are increasing.

Keywords market modelling, mark-up equilibria, robustness, oligopoly

JEL Classification C63, C73, D43, L10, L13, L94

Contact
Publication
Financial Support

dmgn @ cam.ac.uk
November 2015

The robustness of industrial commodity oligopoly pricing strategies*

David M. Newbery[†] and Thomas Greve[‡]
Imperial College London and EPRG Cambridge

December 10, 2015

Abstract

Industrial commodity markets are typically oligopolies in which firms set prices but need to make sunk and durable investment decisions, requiring them to make predictions of future prices. Mark-up pricing models are attractive both for setting prices and predicting future prices for investment analysis. Simple algorithms can find Nash equilibria, but these equilibria are not necessarily robust. This paper examines fixed and proportional mark-up models and demonstrates that they are robust to single firm Nash Cournot deviations but not against more sophisticated deviations in the deterministic case. Cournot equilibria are not robust under demand uncertainty, where proportional mark-up models emerge as the most robust when marginal costs are increasing.

Keywords: market modelling, mark-up equilibria, robustness, oligopoly

JEL Classification: C63, C73, D43, L10, L13, L94

1 Introduction

Industrial commodity markets such as those for electricity, gas, chemicals, aluminium and steel¹ are typically oligopolies, despite the homogeneity of their products that would seem to favour

*This paper is a generalisation of the earlier paper “The strategic robustness of mark-up equilibria,” EPRG1318. We are indebted to a number of earlier referees’ comments as well as to Robert Ritz and Marta Rocha.

[†]Corresponding author: Faculty of Economics, Sidgwick Ave, Cambridge, CB3 9DE, UK, ph: +44 1223 335248; email: dmgn@cam.ac.uk.

[‡]EPRG, Cambridge; contact address: Faculty of Economics, Sidgwick Avenue, Cambridge UK CB3 9DE; email: tg336@cam.ac.uk.

¹Commodities that can be easily stored require an intertemporal analysis with additional complications (see e.g. Baldursson, 1999). We therefore restrict attention to markets for goods like electricity that cannot easily be stored.

intense price competition. A key feature of these markets is that variable costs are typically only a fraction of total costs, and high fixed, or more accurately, sunk costs create a barrier to entry that supports the oligopoly (Grossman, 1990). In addition, different plants often have different costs (through differences in the date of investment combined with technical progress, different input and/or transport costs, etc.). Over time supply and demand conditions change, but firms are imperfectly informed about their rivals' costs and profits. At any moment, firms in such industries are likely to be well-informed about their own variable costs, but aware that their fixed costs are often accounting artifices, or more accurately as their investments are sunk, their current value is only what they are worth in delivering future profits.

Tirole's classic text on Industrial Organization notes the difficulty in theorizing about firms with market power that repeatedly interact in markets, and notes (Tirole, 2000, p240) that "the proliferation of theories is mirrored by an equally rich array of behavioral patterns actually observed under oligopoly." After discussing the delicacy of game theoretic models with reputation, he turns to the Darwinian or evolutionary approach, in which agents may not consciously maximize but suboptimal rules will be rejected or lead to exit. "This approach ... looks at strategies, or (better terminology) rules, that are 'robust' in the sense that they do relatively well against a variety of other rules. ... and that "only actors who use robust rules will stay around. ... What is mean by 'robust rule' is *a priori* contingent on the set and the probability distribution of mutations, and on the set of strategies the robust rule is allowed to be compared to."

Somewhere between the super-rational game players and a population of naive agents following differing rules that is winnowed out by natural selection lie models of learning in which agents modify their rules in the light of their observations and experience. The simplest of these that retain a degree of plausibility are mark-up rules, in which firms mark-up prices on either variable or average costs. They have the obvious attraction that they focus on the key strategic parameter, the price, and keep open the prospect that firms can move to the optimal exploitation of their market power. They are particularly appealing for industrial commodities with homogenous products where price, rather than quality or reputation (needed if quality is hard to perceive *ex ante*), is the focus of competition. Such firms are more likely to choose pricing strategies than choosing quantities to sell, although decisions on investment and capacity are necessarily quantity decisions. More complex versions of such rules include supply function models (Klemperer and Meyer, 1989) which have been applied extensively to electricity wholesale markets (e.g. Green and Newbery, 1992). They have attractions in that they can reconcile capacity constraints, contracting, entry and repeated price setting (Newbery, 1998), but they run into two serious objections: they suffer from a multiplicity of solutions, and they are extremely hard to solve. As such they do not lend themselves to plausible adaptive learning rules, but their

theoretical underpinnings nevertheless provide a lens through which to examine other equilibria, as discussed below.

This dichotomy between setting prices in the short run but choosing quantities in the longer run has been analyzed by Kreps and Sheinkman (1983), and at least under certainty,² allows the short-run decision to be considered as equivalent to a quantity choice, leading to a Cournot equilibrium as a plausible oligopoly equilibrium concept. However, variations in demand without instantaneous changes in capacity still force firms to choose prices rather than quantities, except in rather special circumstances. Long-term contracts for metals and fuels may specify the quantity but link the contract price to some index (e.g. gas is often linked to a lagged oil price), and some spot markets operate as auctions in which some firms offer quantities but other firms offer price-quantity pairs that set the clearing price for all (e.g. spot electricity auction platforms such as EUPHEMIA).

For most industrial commodities, firms set prices and are aware that setting prices at short-run marginal cost not only foregoes evident market power, but risks bankruptcy if variable cost differences are modest. Firms therefore need to adopt pricing rules which can accommodate these market features, and will therefore need to add a mark-up on variable costs (or possibly average costs, where these are known) to reflect their market power and earn a return on their sunk investment costs. Also, what firms observe, and whether they can imitate strategies or only choose between strategies and learn of their outcome, clearly matters. Tirole (2000) cited above makes the point that we should look for robust rules or strategies. Hence, it is clearly timely to examine what form of mark-up to choose and test the robustness of these models to alternative pricing strategies.

This paper examines two particular forms of mark-up pricing – fixed and proportional – that seem well-suited to modeling market equilibria in homogenous industrial oligopolies, and asks whether these models are robust against more sophisticated strategies. If so, they pass the first test of plausibility, but if not, then marketers and modelers need to be aware of their fragility and perhaps consider alternative strategies. The paper follows the tradition of looking at the interaction of a small number of agents, each of which potentially has market power, but who are not initially well-informed about the decisions their competitors or customers are making. They therefore start out of equilibrium and must learn from their observations. Even if there is no uncertainty about demand and technology (and hence costs), each agent is uncertain about the choices its competitors will make, and therefore what is its own best strategy.

The literature on oligopoly pricing is huge. However, the literature is, as far as we know,

²Hviid (1991) noted that under uncertainty this result is threatened by the non-existence of pure-strategy equilibria, a conjecture confirmed by Grant and Quiggen (1996).

silent on the analysis of the robustness of oligopoly models. The papers of Herweg and Mierendorff (2013) and Atalay (2014) highlight the importance of stability in pricing. Atalay (2014) studies the consequences of substantial variation in input prices in an industry of homogeneous products and the survival of firms. Herweg and Mierendorff (2013) go a step further and consider how to price in a market with demand uncertainty. They solve for the profit-maximizing tariff to suggest a flat rate. This paper also considers demand uncertainty with increasing variable costs and finds that proportional mark-ups are both more profitable and more robust.

The two papers closest to ours are Grant and Quiggin (1994; 1996). Grant and Quiggin (1994) examine the case in which each firm chooses its profit-maximizing mark-up over average cost instead of choosing the optimal (Cournot) output level. This has the attraction that firms normally choose prices, not outputs, and relates back to the earlier Hall and Hitch (1939) findings about firm behaviour. Firms are still assumed to set the mark-up where marginal cost is equal to marginal residual revenue, although that residual revenue is governed by how other firms set prices. However, once imperfect information is accepted, and that agents need to experiment to discover the relevant features of their market environment, and specifically the residual demand they face given final demand and the supplies of their competitors, the way is open to modeling how this learning will take place. That will depend on what is known, what is unchanging (e.g. final demand) and what may evolve (e.g. the actions of competitors). Several papers touch on this way of adapting price to rivals' actions. Bagwell (1995) looks at one-shot Stackelberg games and finds that the first-mover advantage is eliminated if there is even a slight amount of noise in observing the leader's choice. Vega-Redondo (1997), Schipper (2009) and Duersch, Oechssler, and Schipper (2012) look at dynamic stochastic Cournot strategies with varying learning and adjustment strategies. Appendix B surveys various approaches to learning and their application to oligopoly models, particularly those extensively used in the electricity industry. However, as far as we know, no paper offers a theoretical exploration into the robustness of different mark-up equilibria to alternative strategies such as Cournot and Stackelberg.

We show that the two mark-up strategies considered are more competitive than Nash-Cournot behaviour, with the Nash choice of the optimal proportional mark-up on marginal costs yielding lower prices and profits than the Cournot oligopoly but higher prices and profits than the optimum fixed mark-up on marginal costs. In deterministic cases, while these mark-up equilibria are robust against Nash deviations by single firms choosing quantities (or any other actions) instead of mark-ups, they are not robust to more sophisticated single-firm Stackelberg deviations in which the deviant maintains her output and the remaining players adapt to that and find the corresponding mark-up equilibrium output levels. The deviant player makes higher profits following this Cournot Stackelberg strategy. If demand is stochastic, then a fixed quantity

response (deterministic Cournot) is strictly inferior to either mark-up equilibria, and in the case of linear marginal costs and linear demand, the proportional mark-up equilibrium as a supply function equilibrium is robust against any deviation, while a fixed mark-up is vulnerable to a proportional mark-up deviant.

This paper makes a number of contributions to the literature on oligopoly pricing. It first sets the scene by ranking in order of profitability three common market equilibrium models: the standard Cournot model (which can be rationalized in a world of certainty as a capacity constrained short-run equilibrium), and two price-setting models in which firms set prices as mark-ups on their marginal costs. This allows us to test their robustness against various kinds of deviations by single and multiple firms, first on the assumption of certainty, and then under uncertainty about demand levels. If, as in most industrial commodity industries, marginal costs are increasing, we find that proportional mark-ups emerge as the most attractive pricing model, which we relate to the literature on supply functions. These findings are relevant to the considerable literature on simulation modeling of such industries, widely used in investment analysis, policy reform (e.g. Green and Newbery, 1992, commenting on restructuring state-owned monopoly generation companies) and anti-trust investigations.

When making investment decisions, such firms need to be able to model the future strategic price-setting behaviour of their rivals in order to decide whether it is profitable to invest. Firms devote considerable resources to modeling their industry, or hire specialist consultancy firms to provide such models. To take just one example, the UK Government commissioned the LCP Consulting to construct the Dynamic Dispatch Model to study the impact of various electricity market reforms on investment incentives and hence future prices.³ Models are also used by anti-trust authorities examining cases of suspected market abuse. In 2005, after energy prices moved sharply upwards, the European Commission launched a *Sector Inquiry* into competition in gas and electricity markets, pursuant to Article 17 of Regulation 1/2003 EC. The final report, published as DG COMP (2007), identified serious shortcomings in the electricity and gas markets, and was based on extensive modelling of these markets. Given the number of possible interactions (between firms and over time), useful models need to strike a balance between simplicity in specifying the rules and the complexity of strategic interactions. Increasingly, agent-based model simulations are constructed to mimic the interaction of agents in quite complex market settings (e.g. Bunn and Oliveira, 2001). Mark-up models emerge as very attractive candidates and have seen widespread analysis and use in such models. Weidlich and Veit (2008) give an excellent

³See <http://www.lcp.uk.com/news-publications/case-studies/2013/decc-dynamic-dispatch-model-envision/> and a review of the model at https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/65711/5427-ddm-peer-review.pdf

survey of wholesale electricity market models and compare different learning strategies and their results.

The paper is organized as follows. In Section 2, we introduce our mark-up models. In Section 3, we present and test the robustness of Cournot deviations, while in Section 4 we test our models against Stackelberg deviations.

2 Mark-up models

The next section sets out the model in which demand and cost functions are assumed non-stochastic. The need to model behaviour as learning is motivated by the imperfect knowledge agents have about the shape of the residual demand schedule they face, which depends on the choices of their competitors — all they observe are the *ex post* price realizations. Different mark-up formulations lead to different equilibria that can be ranked in terms of profit, motivating the reason why a sophisticated agent may wish to choose a different strategy from her rivals. The two mark-up models are examined for their robustness against the first type of Nash-Cournot deviation, demonstrating their robustness against single-firm deviations but not to multi-firm deviations. Section 4 then constructs a counter-example of a beneficial deviation by a single firm in the duopoly case in which a sophisticated leader plays a Stackelberg strategy – in this case she can profitably commit to an inelastic output level or proportional mark-up that differs from the mark-up chosen by the more naive follower.

3 The market model

Consider a market of n identical firms, $i = 1, \dots, n$, producing a non-storable homogenous output, q_i , (e.g. electricity) selling at price p , each with cost function $C(q_i)$, marginal cost, MC, $C'(q_i) > 0$, $C''(q_i) \geq 0$, facing an aggregate demand schedule $Q(p)$, $Q'(p) < 0$, $Q''(p) \geq 0$, $Q = \sum q_i$. The residual demand schedule facing any single firm i , $R_i(Q_{-i}(\cdot))$, $Q_{-i} = \sum_{j \neq i} q_j$, will depend on the supplies offered by the $n - 1$ other firms at the market clearing price. Working with price as the dependent variable, firm i 's residual demand will be $R_i(p) = Q(p) - (n - 1)\bar{q}(p)$, where $\bar{q}(p)$ is the equilibrium supply offered by each of the $n - 1$ remaining firms. Firm i 's profit will be $\pi_i(p) = pR_i(p) - C(R_i(p))$, and the first order condition (f.o.c.) for profit maximization will be

$$\frac{d\pi_i}{dp} = R_i(p) + (p - C')\frac{dR_i}{dp} = 0, \quad p - C' = \frac{R_i}{-R'_i}.$$

This gives the Lerner Index, one of the standard measures of market power, as

$$\frac{p - \text{MC}}{p} = -\frac{R_i}{p} \frac{dp}{dR_i}.$$

3.1 The robustness of competitive and Cournot equilibrium

The benchmark perfectly competitive (Walrasian) equilibrium has $p(n\bar{q}) = C'(\bar{q})$. As this paper is concerned to examine the robustness of various equilibrium concepts, the logical place to start is with an exploration of the competitive equilibrium, as Vega-Redondo (1997) argues for the Walrasian (competitive) Equilibrium as the only equilibrium if firms adopt Evolutionarily Stable Strategies (ESS). In the n -firm case, would it be advantageous for one firm to choose a quantity different from that implied by offering to supply at its marginal cost, MC? The answer is clearly yes, as the supply of the remaining $n - 1$ firms can be subtracted from total demand to give a downward-sloping (and hence well-behaved) residual demand facing the firm, who will then choose output to maximize profit (i.e. where $MC = MR$). However, the Cournot deviant raises the price for all other price-taking firms, who, as they outproduce the deviant, will make higher profits than the deviant, raising the question whether firms would defect from a Cournot oligopoly to become price takers. This is the situation examined by Duersch et al. (2012) and others studying ESS, who argued that copying the actions of the most successful firm (in this case the price-takers) would lead to a competitive outcome (see Appendix B for a more detailed discussion of learning and ESS models).

One can turn the question round and ask whether, starting from the n -firm Cournot oligopoly, it would pay $r \geq 1$ firms to deviate to become price-taking competitive firms and setting output where MC equals price. Clearly as the deviant receives the same price as the Cournot firms, he makes more profit than the other Cournot firms and might seem to find it more attractive to act competitively than stay with the oligopolists. However, Appendix A shows (in a linear-quadratic model) that the profit the deviant makes as a competitor facing $r - 1$ other competitive firms and $n - r$ Cournot firms is less than the profit he would make as one of $n - r + 1$ Cournot firms facing $r - 1$ competitive firms, so there is no advantage in deviating.

3.2 The robustness of mark-up equilibria

In a symmetric n -firm oligopoly the possible Nash Equilibria (NE) will depend on the strategy set, and whether each firm chooses its strategic variable simultaneously. The three cases considered are the Nash-Cournot equilibrium, in which firms choose quantities, the fixed mark-up equilibrium in which firms add a fixed mark-up m_i to MC: $C'(q_i) + m_i$, and the proportional mark-up to marginal cost, in which firms offer supplies in proportion to MC, $\theta C'(q_i)$, $\theta > 1$. Proposition 1 ranks these equilibria.

Proposition 1 *In the simultaneous move symmetric n -firm oligopoly, the Nash Equilibrium price is highest in the Nash-Cournot equilibrium, higher than in the proportional mark-up equi-*

librium, which in turn is higher than the fixed mark-up equilibrium, which is higher than the competitive price.

Proof. (i) If the other firms choose quantities (Cournot), $\bar{q}(p) = \bar{q}^c$, and firm i 's residual demand will be $R_i^c(p) = Q(p) - (n-1)\bar{q}^c$, so $R_i^c = Q'(p)$, and the f.o.c. for profit maximization will define the equilibrium output \bar{q}^c :

$$p(n\bar{q}^c) - C'(\bar{q}^c) = \frac{\bar{q}^c}{-Q'(p)}. \quad (1)$$

(ii) If firms choose a fixed mark-up on marginal cost, MC , their offer schedules will be $P_j(q_j) = C'(q_j) + m_j$, the market clearing condition will be $p(\sum q_j) = P_j(q_j)$, all j , and $\bar{q}^m(p) = C'^{-1}(p - \bar{m})$, so $R_i^m = Q' - (n-1)dC'^{-1}/dp = Q' - (n-1)/C''(\bar{q}^m)$ so the f.o.c. for profit maximization will be

$$p - C' = \frac{q_i}{-Q' + (n-1)/C''(\bar{q}^m)}. \quad (2)$$

As $Q' < 0$ and $C'' \geq 0$, the RHS of (2) is less than that of (1), but replacing \bar{q}^c by \bar{q}^m . But as demand slopes down and the MC slopes (weakly) up, the only way that the Nash-Cournot mark-up could be higher would be if $\bar{q}^c < \bar{q}^m$ and $p^c > p^m$. It follows that the fixed mark-up price- MC margin will be lower than under Nash-Cournot competition.

(iii) If firms choose a proportional mark-up, then their offer schedules will be $P_j(q_j) = \theta C'(q_j)$, $p(\sum q_j) = P_j(q_j)$, all j , $\theta > 1$, and hence in equilibrium $\bar{q}^p(p) = C'^{-1}(p^\theta/\theta)$ and so

$$p - C' = \frac{q_i}{-Q' + (n-1)/(\theta C''(\bar{q}^p))}. \quad (3)$$

As $\theta > 1$, the RHS of (3) is larger than (2) but smaller than (1), each evaluated at \bar{q}^p , so reasoning as in (ii), the proportional price- MC margin will be in between the mark-up margin and the Nash-Cournot margin, demonstrating that the proportional MC strategy is less competitive than the fixed mark-up strategy, but more competitive than the Nash-Cournot strategy, and the equilibrium prices will lie between these two equilibria. Finally, in each case the price- MC margin is strictly positive and hence prices will be higher than the competitive equilibrium price. ■

Figure 1 shows the result of the deviant choosing output when the remaining firms choose to offer at a fixed mark-up on marginal cost for the simple case of linear demand ($Q = A - p$) and quadratic costs ($C(q_i) = \frac{1}{2}cnq_i^2$) as set out in Appendix A. The deviant's optimal response is to choose q_d to maximize profit, $p q_d - C(q_d)$, given the residual demand generated by the fixed mark-up behaviour, which in this linear case can be written $p = (\alpha - q_d)/\beta$ for some α and β . The question to address is whether choosing the Nash mark-up given in Appendix A equation (13) is robust against a player optimizing against this strategy. Proposition 2 shows, as the figure clearly demonstrates, that the answer is yes.




Figure 1 $A = 140$, $c = 1$, $n = 2$, $q^m = 30$.

Proposition 2 *In a deterministic Nash game, if players assume that the strategy space is the choice of the fixed or proportional mark-up on its marginal cost, then that player will maximize her profits regardless of whether another player chooses the same strategy choice as other players (fixed or proportional mark-up on marginal cost) or an optimal quantity to supply, and hence the two Nash mark-up equilibria are robust against Cournot deviations.*

Proof. *Each firm faces the same residual demand schedule and will choose the same optimal output whether they choose the optimal fixed or proportional mark-up, or the optimal quantity (or any other choice variable such as price) that corresponds to MC set equal to the marginal residual demand revenue. QED. ■*

Thus the Nash equilibrium behaviour in mark-ups or slopes is robust against a Nash-Cournot deviant who can choose from a broader set of strategies that also includes quantities, at least in a deterministic setting. Figure 2 illustrates this for the case in which all but the deviant firm choose their supply as a proportional mark-up on their linear MC.

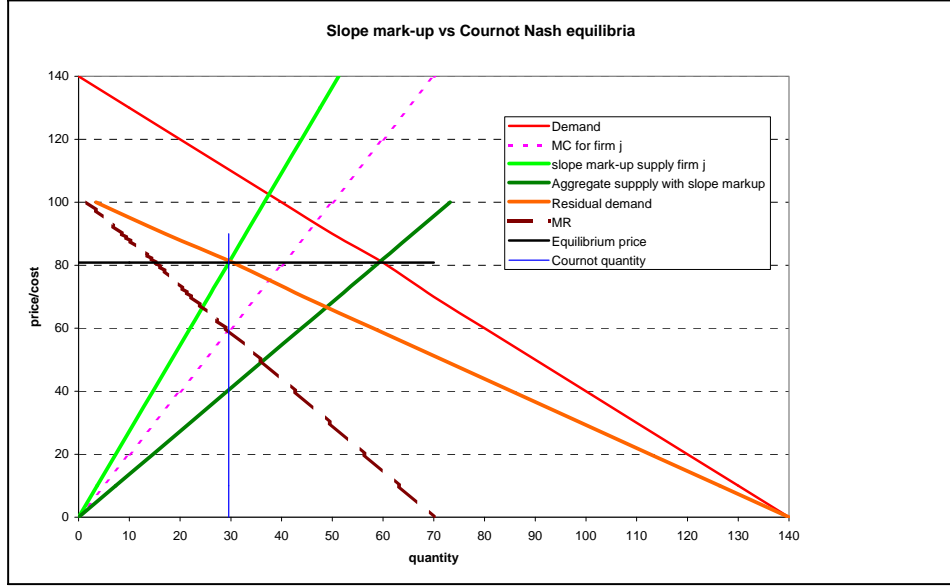


Figure 2 $A = 140$, $c = 1$, $n = 2$, $q^s = 29.59$.

On reflection, this should not be surprising, as once the other firms have chosen their optimal mark-up (or slope) given the residual demand *they assume faces them*, the deviant faces the same residual demand and hence chooses the same mark-up, which is where the residual marginal revenue meets MC, shown in Figures 1 and 2. One might reasonably argue that the resulting outcome is no longer a true Nash equilibrium, in that while the deviant firm correctly predicts what the other firms will do, these mark-up firms are not correctly predicting what strategy the deviant is following and hence not correctly predicting the residual demand they face, although they are predicting the mark-up she will actually choose. We will explore this further in section 4. It is also important to note that although the equilibrium is robust to deviations, the resulting price depends on the strategy space, in this case mark-ups or supply slopes rather than quantities (and is lower as a result).

3.3 Multi-firm deviations

Suppose that r firms decide to play a Cournot strategy, knowing that $n - r$ firms will continue with their fixed mark-up strategies but the other deviants will choose the same output as the first deviant (but independently, each taking the other's output as given). For the linear-quadratic example set out in Appendix A, the market clearing condition is $p = A - q - (r - 1)q_o - \frac{n-r}{nc}(p - m)$, where m is the fixed mark-up, q_o is the output of each of the other Cournot deviants, taken as given, and q is the output choice to be made. Then

$$p\gamma = A - q - (r - 1)q_o + \frac{(n - r)m}{nc}, \quad \frac{dp}{dq} = -\frac{1}{\gamma}, \quad \gamma = 1 + (1 - \frac{r}{n})/c.$$

As before the f.o.c. for the deviant is given by $p = MC - Rdp/dR$ or

$$p\gamma = ncq\gamma + q = A - q - (r - 1)q_o + \frac{m(n - r)}{nc}.$$

Set $q_o = q$ and substitute for m to give

$$\begin{aligned} (r + 1 + nc\gamma)q_o &= A \frac{n^2(1 + c)^2 - r}{n^2(1 + c)^2 - n}. \\ \gamma p &= A - rq_o + \frac{A(n - r)}{n^2(1 + c)^2 - n}. \end{aligned}$$

As a numerical example, let $n = 5$, $r = 2$, $c = 1$, $A = 100$, so that the symmetric mark-up equilibrium has $q_i = 9A/95 = 9.47$, $p = 52.6$. The deviants' output will be $98A/1045 = 9.378$, which is smaller, so the equilibrium will be different. The price will be $p = 52.8$ and so the profit of a deviating firm rather than conforming to the original strategy will be 274.8 rather than 274.2, or 0.22% higher. Thus there is a (small) incentive for a subset of more than one (very) sophisticated firms to deviate from a Nash mark-up equilibrium. This is reminiscent of Delgado and Moreno (2004), who show that only the Cournot outcome can be sustained by a coalition-proof supply function equilibrium, in configurations such as this, with linear demand and at least three players. However, the meaning of the Cournot outcome with supply functions is rather different, as it is the least competitive supply function, which is not a constant level of output.

3.4 Robustness under uncertainty

The results derived above assume no uncertainty in the level of demand nor of cost, and the results are not robust to uncertainty. We know from Weitzman (1974) and Klemperer and Meyer (1986) that the choice of quantity or price may be a matter of indifference under perfect certainty, but not with uncertainty. Consider the simple case of additive risk:

$$\tilde{D}(Q, \varepsilon) = \tilde{p} = D(Q) + \tilde{\varepsilon}, \quad E\tilde{\varepsilon} = 0, \quad (4)$$

where $\tilde{\varepsilon}$ is a mean zero random variable (indicated with a tilde) and $D(Q)$ is deterministic.

Proposition 3 *Under additive risk, a Cournot deviant in a simultaneous play Nash game will earn lower profits than players choosing either a constant or proportional mark-up on their marginal costs.*

Proof. *Under additive risk as in (4) the optimal mark-up is the same as under uncertainty, as the partial derivatives of price and quantity with respect to the fixed mark-up, m_i , or proportional mark-up parameter, θ_i , are deterministic. Expected profit is $E p(Q, \varepsilon) q_i(\phi_i, \varepsilon) - C(q_i(\phi_i, \varepsilon))$, where $\phi_i = m_i$ or θ_i . But $\partial q_i / \partial \phi_i$ and $\partial p / \partial \phi_i$ are both deterministic, and the f.o.c. are*

$$C'(q_i) \frac{\partial q_i}{\partial \phi_i} - E q_i \frac{\partial p}{\partial \phi_i} = \frac{\partial q_i}{\partial \phi_i} E(p) = D(Q) \frac{\partial q_i}{\partial \phi_i},$$

which is the same deterministic equation for Eq_i as for q_i in the deterministic case. Additivity also implies the same deterministic f.o.c. for the deviant's optimal choice of quantity, $q_d = Eq_i$. The difference in profit between the deviant choosing a fixed quantity and her mark-up rivals is

$$\pi_d = \pi_i = Eq_i \cdot Ep - Epq_i < 0,$$

as p and q_i are positively correlated. ■

Whether a proportional mark-up equilibrium is robust against a fixed mark-up deviant, and a fixed mark-up equilibrium is robust against a proportional mark-up deviant is less clear, except in the linear marginal cost case.

Proposition 4 *In the case of stochastic demand and linear marginal costs, the proportional mark-up equilibrium is robust against a fixed mark-up deviant but the fixed mark-up equilibrium is vulnerable to a proportional mark-up deviant.*

Proof. We know from Klemperer and Meyer (1989) that there is a unique linear supply function equilibrium in the case of linear marginal costs which has constant slope, so the proportional mark-up equilibrium is also the supply function equilibrium and hence optimal against any deviation. ■

4 Robustness to Stackelberg deviations

Although a single deviant was unable to improve on her profits by choosing quantities rather than mark-ups in the deterministic case, knowing that the remaining firms were acting on the (mistaken) assumption that all firms were choosing their mark-up facing the same residual demand schedule, there remains a question whether this is a consistently formulated equilibrium for a sophisticated deviant. If all that firms observe is the consequences of their choices in the market price, then they are correctly choosing the optimal choice of mark-up (or output). If they are basing their choice of mark-up on assumptions about the shape of the residual demand they face, then the assumed residual demand will be incorrect in the face of a Cournot deviant. One way round this inconsistency is to suppose that the deviant firm's strategy choice is known to the remaining firms, who nevertheless continue to choose their mark-up (and similarly the deviant knows that the other firms will behave that way). In a learning context, this would require the leader to stick to her output strategy, while the followers learned that they could then improve their profits by adapting to the new environment. The resulting equilibrium is most simply modelled as the outcome of a Stackelberg game in which the deviant is the leader who can commit in this case to her output, and to which the followers respond. As the aim is to demonstrate that deterministic mark-up equilibria are not robust to this more sophisticated deviation, this

section considers the simpler duopoly case ($n = 2$) for the linear quadratic example of Appendix A. The first step reproduces the classic Cournot Stackelberg oligopoly. The leader can commit to her output level, q_l , before the other (the follower) makes his choice, q_f , so that the leader can optimize against the follower's reaction function (10) given in Appendix A. The resulting equilibrium has

$$q_l = \frac{A(1+2c)}{2(1+4c+2c^2)} > q_f = \frac{A(1+6c+4c^2)}{4(1+c)(1+4c+2c^2)}. \quad (5)$$

These expressions are somewhat opaque, but simplify in the constant returns case in which $c = 0$ to the familiar solution $q_l = A/2$, $q_f = A/4 = p$, and the profit of the leader is $3A^2/16$, larger than the follower's profit, who receives only one-third as much or $A^2/16$, and also higher than under the symmetric Nash Cournot equilibrium of $A^2/9$. If $c = 1$, then $q_l = 12A/56 > q_o = A/5 > q_f = 11A/56$ and the leader's profit is 0.45% higher than in the symmetric duopoly.

4.1 Stackelberg quantity deviations from the mark-up equilibrium

The question we now address is whether choosing the optimal Nash mark-up is robust against a more sophisticated player who in a sequential setting can stick to her optimal output while the other player continues to mark up on marginal cost but learns the optimal mark-up (effectively the correct position of his residual demand schedule). The previous duopoly example demonstrates that mark-up behaviour it is not robust (in a deterministic setting at least) against a Cournot (quantity-fixing) deviation and that is also the case when the follower is choosing his mark-up. Thus if the leader commits to a quantity, q_l , the follower chooses a mark-up m given the residual demand schedule $q_f(q_l) = A - p - q_l$. The follower's supply schedule is given by Appendix A equation (12) with $n = 2$, so market clearing yields

$$p = A - q_l - \frac{p - m}{2c}, \text{ or } p = \frac{m + 2(A - q_l)c}{1 + 2c}. \quad (6)$$

The follower chooses m to maximize profit (see Appendix A) resulting in a reaction function

$$q_f = \frac{p - m}{2c} = \frac{p}{1 + 2c}.$$

The leader's optimal response (and the follower's output) are exactly as in the Stackelberg Cournot equilibrium (5) in which both agents choose quantities, as shown in Figure 3.

Stackelberg quantity vs mark-up equilibrium




Figure 3 $A = 140$, $c = 1$, $n = 2$, $q_l = 16.2$, $q_f = 31.0$

Proposition 5 *If one player assumes that the strategy space is the choice of the mark-up on its marginal cost, then with a quadratic cost function, linear demand and two identical players, the other player will find it profitable to commit to choosing a (different) optimal quantity to supply.*

Proof. *From equation (5) $q_l > q_f$. Since the leader could have chosen the same output as the follower but chose not to, she must be making higher profits. ■*

4.2 Stackelberg quantity deviations from the proportional mark-up equilibrium

If the leader offers a fixed quantity q_l and the follower offers the supply schedule $p = \theta C'(q_f)$, then $q_f = sp/2$, replacing $\theta = 1/s$. The market clearing price is

$$p = \frac{2(A - q_l)}{2 + s}, \quad \frac{ds}{dp} = -\frac{(2 + s)}{p}.$$

The follower's problem is to maximize $\pi_f = p^2(2s - cs^2)/4$ for which the f.o.c. is

$$\begin{aligned} p(2s - cs^2) &= p(1 - cs)(2 + s), \\ s &= \frac{2}{1 + 2c}. \end{aligned} \tag{7}$$

It may seem surprising that this does not depend on the leader's choice, but the follower's actual mark-up will be lower the larger the output of the leader and hence the lower the price, which gives the leader the advantage. The leader chooses q_l to maximize profit, given that

$p = 2(A - q_l)/(2 + s)$. The profit maximizing solution is again the Cournot Stackelberg equilibrium (5).

Proposition 6 *If one player assumes that the strategy space is the choice of the slope (or proportional mark-up on marginal cost) of its offer, then with a quadratic cost function, linear demand and two identical players, the other player will find it profitable to commit to choosing a (different) optimal quantity to supply.*

Proof. *From equation (5), $q_l > q_f$. Since the leader could have chosen the same output as the follower but chose not to, she must be making higher profits. ■*

4.3 Stackelberg proportional deviations from the mark-up equilibrium

If the leader chooses a proportional supply schedule given by $q_l = sp/2$ and the follower chooses his mark-up m , the market clearing price is given by

$$p = A - \frac{sp}{2} - \frac{p - m}{2c}, \text{ or } p = \frac{2Ac + m}{2c + 1 + sc}. \quad (8)$$

Figure 4 illustrates and Appendix A demonstrates that this can be a profitable deviation, although not as profitable as choosing output rather than slope when confronting mark-up followers.




Figure 4 $A = 140$, $c = 1$, $n = 2$, $q_l = 29.0$, $q_f = 29.7$

5 Conclusion

Mark-up pricing simulation models are attractive in attempting to model outcomes in complex markets where some agents can act strategically, and there has been considerable interest in whether adaptive learning will lead to Nash equilibria (Weidlich and Veit, 2008), as these would seem natural equilibrium concepts. However, as with all attempts to model strategic behaviour, the resulting equilibrium is sensitive to the action space from which agents choose. Standard oligopoly models consider actions to be either quantities (supplies to the market), as in the Cournot formulation, or prices offered to the market (the Bertrand assumption). In the presence of uncertain or varying demand, supply function models, developed by Klemperer and Meyer (1989) and applied to electricity markets by Green and Newbery (1992), are attractive intermediate formulations, and their linear solutions⁴ have been influential in motivating the kind of agent-based models considered here.

All these specifications assume a unitary or owner-managed firm pursuing maximum profit but under managerial capitalism the ultimate owners need to motivate managers. Ritz (2008) argues that rewarding managers for increasing their market share is consistent with the evidence and can be a useful in pursuing more collusive strategies. Ritz concludes that though competing for market shares seems more aggressive, it is indeed more “robust” to strategic manipulations. Following the same line, Vickers (1985) and Fershtman and Judd (1987) study the strategic distortion of preferences. These models compare different proximate objective functions with the same choice variable (and with the same goal of ultimately maximizing profits). In contrast, our paper has the same ultimate objective function – profit maximization – but compares different choice variables.

While the choice of action space in optimizing models is normally guided by the market structure and the actions that agents have, the choice of action space in simulation and particularly agent-based models is normally guided by tractability, where a choice of a single parameter (such as the fixed or proportional mark-up over marginal cost) considerably simplifies the problem. This paper has shown that the two mark-up strategies considered are more competitive than Nash-Cournot behaviour, with the Nash choice of the optimal proportional mark-up on marginal costs of the offers yielding lower prices than the Cournot oligopoly prices but higher prices than the optimum fixed mark-up on marginal costs. In deterministic cases, while these mark-up equilibria are robust against Nash deviations by single firms choosing quantities instead of mark-ups (so they are in that sense Nash equilibria), they are not robust to either group deviations or to

⁴Supply function models typically have a continuum of solutions, one of which may be linear, providing there are no relevant capacity constraints. Where capacity constraints are important, there may be unique but non-linear solutions.

more sophisticated single firm Stackelberg deviations in which the deviant maintains her output and the remaining players adapt to that and find the corresponding mark-up equilibrium output levels. The deviant player makes higher profits following this Cournot Stackelberg strategy (or the proportional mark-up against a fixed mark-up strategy), casting doubt on the robustness of simultaneous move Nash mark-up equilibria. On the other hand, if demand is stochastic, then a fixed quantity response (deterministic Cournot) is strictly inferior to either mark-up equilibria, and in the case of linear marginal costs, the proportional or proportional mark-up equilibrium as a supply function equilibrium is robust against any deviation, while a fixed mark-up is vulnerable to a proportional mark-up deviant.

The problem becomes more complex if multiple deviants collude, in which case Delgado and Moreno (2004) show that only the least competitive supply function equilibrium is coalition-proof. But that is to address a different set of questions, and raises the issue of whether entry is free or not, which considerably alters the set of sustainable equilibria, as Newbery (1998) shows.

The implication of our findings for agent-based modeling, which is where mark-up modeling is most widely used, is that it would be sensible to test any resulting numerically found convergence results for robustness against different choice variables by one or more (large) agents, who would need to maintain their choice for sufficiently many iterations to induce responses by the remaining agents. A finding that the resulting equilibria were close to the original solution would attest to its robustness, but if not the proposed solution would remain suspect. One moderately robust conclusion is that proportional or proportional mark-ups on marginal cost are likely to be superior to and more robust than alternative modeling assumptions.

References

- Atalay, E., (2014). "Materials Pricing and Productivity." *Journal of the European Economic Association*, 12, 575–611.
- Bagwell, K., (1995). "Commitment and Observability in Games." *Games and Economic Behavior*, 8, 271-280.
- Baldursson, F.M., (1999). "Modelling the price of industrial commodities." *Economic Modelling*. 16. 331-353.
- Bunn, D.W. & Oliveira, F.S., (2001). "Agent-based simulation: an application to the new electricity trading arrangements of England and Wales." *IEEE Transactions on Evolutionary Computation*, Special issue: Agent Based Computational Economics 5 5, 493-503.
- Bunn, D.W., Oliveira, F.S., (2003). "Evaluating individual market power in electricity markets via agent-based simulation." *Annals of Operations Research*, 121 1-4, 57–77.
- Delgado, J. & Moreno, D., (2004). "Coalition-proof supply function equilibria in oligopoly." *J. Econ. Theory*, 114. 231-254.

- DG COMP, (2007). *DG Competition report on energy sector inquiry*. SEC20061724, 10 January
- Duersch, P., J. Oechssler, and B.C. Schipper., (2012). “Unbeatable imitation.” *Games and Economic Behavior*, 761, 88-96.
- Fershtman, C. and K. Judd (1987). “Equilibrium incentives in oligopoly.” *American Economic Review*, 77, 927-940.
- Grant, S. and Quiggin, J. (1994). “Nash equilibrium with markup-pricing oligopolists.” *Economics Letters*, 45, 245-51
- Grant, S. and Quiggin, J. (1996). “Capital Precommitment and Competition in Supply Schedules.” *The Journal of Industrial Economics*, 44(4), 427-441.
- Green, R. J. & Newbery, D.M., (1992). “Competition in the British Electricity Spot Market.” *Journal of Political Economy*, 100 (5), 929-53.
- Grossman, S.J., (1990), “Nash Equilibrium and the Industrial Organization of Markets with Large Fixed Costs”, *Econometrica*, 49, 1149-1172.
- Hall, R. and Hitch, C., (1939). “Price Theory and Business Behaviour.” *Oxford Economic Papers*. 2, 12-45
- Herweg, F. and Mierendorff, K., (2013). “Uncertain Demand, Consumer Loss Aversion, and Flat-Rate Tariffs.” *Journal of the European Economic Association* ,11, 399–432.
- Hviid, M., (1991). “Capacity Constrained Duopolies, Uncertain Demand and Nonexistence of Pure Strategy Equilibria”, *European Journal of Political Economy*, 7, 183-190.
- Klemperer, P. D., & Meyer, M.A., (1986). “Price competition vs. quantity competition: the role of uncertainty.” *Rand Journal of Economics*, 17(4), 618-638.
- Klemperer, P. D. and Meyer, M.A., (1989). “Supply Function Equilibria in Oligopoly under Uncertainty.” *Econometrica*, 57 (6), 1243-1277.
- Kreps, D.M. & J. A. Scheinkman, (1983). “Quantity precommitment and Bertrand competition yield Cournot outcomes.” *The Bell Journal of Economics*, 14(2), 326-337.
- Littman, M., (1994). “Markov games as a framework for multiagent reinforcement learning”, in *Proceedings of the Eleventh International Conference on Machine Learning*. pp.157-163, San Francisco, CA: Morgan Kaufmann.
- Newbery, D.M., (1998). “Competition, Contracts and Entry in the Electricity Spot Market.” *RAND Journal of Economics*, 29(4), 726-49.
- Ritz, R. A., (2008). “Strategic incentives for market share.” *International Journal of Industrial Organization*, 26, 586-597.
- Schipper, B.C., (2009). “Imitators and optimizers in Cournot oligopoly.” *Journal of Economic Dynamics and Control*, 33, 1981-1990.

Tirole, J., (2000). *The Theory of Industrial Organization*, Cambridge MA. MIT Press

Vega-Redondo, F., (1997). “The evolution of Walrasian behavior.” *Econometrica*, 65, 375-384.

Vickers, J., (1987). “Delegation and the theory of the firm.” *Economic Journal*, 95, 138-147.

Watkins, C., (1989). “Learning from Delayed Rewards,” Ph.D. dissertation, Cambridge University, Cambridge, England.

Weidlich, A. and D. Veit, (2008). “A critical survey of agent-based wholesale electricity market models”, *Energy Economics*, 30, 1728-59.

Weitzman, M., (1974). “Prices vs quantities.” *Rev. Econ. Stud.* 477-491.

Appendix A Derivations of results

The linear demand quadratic cost case

If the cost function is quadratic: $C(q_i) = aq_i + \frac{1}{2}cnq_i^2$, then without loss of generality set $a = 0$ by measuring prices relative to the level a . The marginal cost, MC, is cnq_i , which gives an aggregate competitive linear supply schedule cQ , independent of n . If demand is also linear, $Q(p) = A - bp$, then the inverse slope parameter can be set at $b = 1$ by a suitable choice of price units, so that $p = A - Q$. The perfectly competitive solution is $p_c = MC = cQ$, with

$$p_c = \frac{Ac}{(1+c)}, \quad q_c = \frac{A}{n(1+c)}. \quad (9)$$

The Nash-Cournot solution is found from the reaction function derived from the first order condition (f.o.c.):

$$q_i = (A - Q_{-i})/(2 + cn), \quad (10)$$

so the symmetric oligopoly solution is

$$p_o = \frac{A(\frac{1}{n} + c)}{(1+c + \frac{1}{n})} > p_c, \quad q_o = \frac{A}{n(1+c) + 1} < q_c. \quad (11)$$

Deviations from competitive equilibrium

The supply function of each competitive firm j is given by $ncq_j = p$, so residual demand is $q = A - p - (n-1)q_j = A - p(\frac{nc+n-1}{nc})$. The deviant firm d then maximizes profit, for which the f.o.c. gives

$$q_d = \frac{A}{1+n+nc}, \quad p_d = \frac{Ac(1+c)}{(1+c)^2 - \frac{1}{n^2}} > p_c = \frac{Ac}{1+c}.$$

Note that the optimal output is the same as the Nash Cournot output give in (11). Profit for the deviant is

$$\pi = \frac{A^2c(n(1+c)^2 + 2(1+c) + 1/n)}{2(1+n+nc)^2((1+c)^2 - 1/n^2)} > \pi_c = \frac{A^2c}{2n(1+c)}.$$

This last inequality can be demonstrated by subtracting competitive profit from deviant profit, which, after simplification, has the sign of

$$c(1 + c + 1/n)^2 \geq 0.$$

Thus, provided the competitive equilibrium has a positive price (i.e. $c > 0$), it is more profitable to set quantity (or offer whatever the market demands at the price p_d above) than to act competitively. This is equivalent to observing that the competitive equilibrium here is not a Nash equilibrium with finitely many firms.

Fixed mark-up models

If each firm offers its supply at a fixed mark-up above marginal cost, MC, its offer price $p = MC + m_i$, (so $p - MC = m_i$). As $MC = ncq_i = p - m_i$, the supply schedule in price space is

$$q_i = \frac{(p - m_i)}{nc}, \quad (i = 1, \dots, i, \dots, n). \quad (12)$$

The market clearing price, MCP, $p(m_i, \mathbf{m}_{-i})$ solves, after substituting q_i from (12),

$$\begin{aligned} p &= A - \sum q_i = \frac{(Anc + \sum m_i)}{n(1 + c)}, \text{ so } \frac{\partial p}{\partial m_i} = \frac{1}{n(1 + c)}, \\ ncq_i &= p - m_i, \text{ so } \frac{\partial q_i}{\partial m_i} = \frac{1}{nc} \left(\frac{\partial p}{\partial m_i} - 1 \right). \end{aligned}$$

The f.o.c. from maximizing profit w.r.t. m_i gives

$$\begin{aligned} \frac{\partial \pi_i}{\partial m_i} &= (p - MC) \frac{\partial q_i}{\partial m_i} + q_i \frac{\partial p}{\partial m_i} = \frac{m_i}{nc} \left(\frac{\partial p}{\partial m_i} - 1 \right) + q_i \frac{\partial p}{\partial m_i}, \\ &= \frac{-m_i(1 - \frac{1}{n} + c)/c + q_i}{n(1 + c)} = 0, \text{ so } q_i = \frac{m_i}{nc} (n - 1 + nc), \\ ncq_i &= p - m_i = m_i(n - 1 + nc), \end{aligned}$$

$$\begin{aligned} nm_i(1 + c) &= p = \frac{(Anc + m_i + \sum_{j \neq i} m_j)}{n(1 + c)}, \\ m_i(n^2(1 + c)^2 - 1) &= Anc + \sum_{j \neq i} m_j. \end{aligned}$$

The symmetric equilibrium has $m_i = m$:

$$\begin{aligned} m(n^2(1 + c)^2 - 1) &= Anc + (n - 1)m, \\ m &= \frac{Ac}{n(1 + c)^2 - 1}, \end{aligned} \quad (13)$$

Stackelberg proportional deviations from the proportional mark-up equilibrium

It is convenient to work in terms of the inverse of the proportional markup, $s = 1/\theta$. From (8), $\partial p/\partial m = 1/(1+2c+sc)$ and the follower maximizes profit $pq_f - cq_f^2$ by his choice of m as before, noting that $\partial q_f/\partial m = (\partial p/\partial m - 1)/2c$ and $p - MC = m$, giving f.o.c

$$\begin{aligned}\frac{\partial \pi_f}{\partial m} &= \frac{m}{2c} \left(\frac{\partial p}{\partial m} - 1 \right) + q_f \frac{\partial p}{\partial m} = 0, \\ m &= p \frac{\partial p}{\partial m} = \frac{p}{1+2c+sc}, \text{ but} \\ p &= \frac{2Ac+m}{2c+1+sc},\end{aligned}$$

so

$$p = \frac{2Ac(1+2c+sc)}{c(2+s)(2+2c+sc)}.$$

The leader's profit function is proportional to $(2s - cs^2)p^2$ for which the f.o.c. can be written

$$\begin{aligned}\frac{d \log p}{ds} &= \frac{cs-1}{s(2-cs)}, \text{ but} \\ \frac{d \log p}{ds} &= \frac{c}{1+2c+sc} - \frac{1}{2+s} - \frac{c}{2+2c+sc},\end{aligned}$$

from which the optimal value of s can be determined. For example if $c = 1$, $A = 100$, the symmetric proportional mark-up equilibrium has $s = \sqrt{3} - 1 = 0.732$ and $q = 21.1$, but in this case the solution is $s = 0.7142$, which is a slightly steeper slope. Instead of the symmetric price being $A/\sqrt{3} = 57.7$, the new price is 58.1 and the leader's output is 20.7, profit 774. Had the leader accepted the fixed mark-up equilibrium, the price would have been 57.1, the output 21.4 and profit 765, which is less. If the leader had chosen output rather than the proportional mark-up the equilibrium price would have been 58.9, quantity 21.4 and profit 800, higher than choosing the proportional mark-up.

Appendix B Learning, adjustment, convergence and stability

In an oligopoly, agents compete against others who also possess market power, and the resulting equilibrium, if one exists, will depend on the actions that agents take (setting prices, offering quantities, choosing quality, investing, advertising, etc.), and the extent to which they take account of strategic interactions, which will depend on what they know, what they can observe, and what rules of behaviour they adopt. The wide range of oligopoly equilibria depends on the rules followed, the strategy sets, and information available to players. This appendix briefly surveys some of these options and their implications for the choice for modelling markets.

The standard economic model assumes that firms have full information about costs and demand and rationally decide on their profit-maximizing choice (of output or price in the simplest case of homogenous products). Empirical industry studies in the 1940's were concerned with the apparent mismatch between the theory of profit maximization and the evidence that managers had little if any concept of marginal cost and revenue. Instead they followed rules of thumb in setting prices at a mark-up over average cost (Hall and Hitch, 1939). As Grant and Quiggin (1994) observed, this led to a methodological debate in which Friedman (1953) argued that managers could be following rules of thumb but still be pushed towards profit-maximizing behaviour, as those who chose non profit-maximizing mark-ups would earn lower profits and lose market share and/or exit. Grant and Quiggin further argued that this argument does not immediately apply to oligopolistic markets in which managers need to anticipate how their rivals will respond to their actions. If other firms were choosing mark-ups on costs, then this would affect the choice of a profit maximizing firm.

Grant and Quiggin (1994) examined the case in which each firm chose its profit-maximizing mark-up over average cost instead of choosing the optimal output level. This relates firms' actions more closely to what they observe, but they are still assumed to set the mark-up where marginal cost is equal to marginal residual revenue (although their opening remarks suggest that firms actually adjust their mark-ups in response to market conditions and so grope towards this profit-maximizing position, in the spirit of Friedman's defence).

Once imperfect information is accepted, and that agents need to experiment to discover the relevant features of their market environment, and specifically the residual demand they face given final demand and the suppliers of their competitors, the way is open to modelling how this learning will take place. That will depend on what is known, what is unchanging (e.g. final demand) and what may evolve (e.g. the actions of competitors). One approach has been to explore the consequences of imperfectly observing rivals' actions but nevertheless having otherwise full information about the consequences of any set of actions by the players. Bagwell (1995) explores one-shot Stackelberg games to disentangle the implications of the two standard assumptions - that one agent can move first, and that all other agents perfectly observe her choice. He claimed that the first-mover advantage is eliminated if there is even a slight amount of noise in observing the leader's choice. van Damme and Hurkens (1997) criticized this conclusion by noting that it depended on restricting choices to pure strategy equilibria, and that Bagwell's game always had mixed strategy equilibria close to the Stackelberg equilibrium when the noise was small. Their paper contains an extensive discussion of equilibrium selection, but remains firmly within the standard game-theoretic approach.

An alternative approach is to suppose that the information available to firms is limited to

their own costs and market outcomes, which they can observe in a sequence of choices by all the firms. Again there are several ways this can be modelled, depending on what firms can observe. A number of papers have assumed that firms can observe the actions and subsequent profits earned by their rivals, so that they can imitate their behaviour. Vega-Redondo (1997) explores the classic Cournot setting of a market of quantity-setting firms producing a homogenous output in which firms experiment with a different level of output with a ε -(small) probability, and successful firms win out over less successful firms, as Friedman (1953) argued. He concludes that the final resting place of this stochastic dynamic process was the unique Walrasian (i.e. perfectly competitive) outcome. This approach was extended by Schipper (2009), who allowed firms either to imitate the output choices of the most profitable firms, or to optimize against the other firms, but with all firms making small mistakes. This time the long-run state is one in which imitators are better off than optimizers. A subsequent paper by Duersch, Oechssler, and Schipper (2012) extends this idea, which was prompted by observing experimental subjects playing a Cournot duopoly against a computer programmed with a variety of learning algorithms. The computer could easily be beaten, except when it followed the rule of copying the action of the most successful player in the previous round.

Schipper and his colleagues assume that all firms can identify both the actions and the resulting profits of their rivals, while Vega-Redondo (1997) adopted a more Darwinian approach in that exit or death is more likely with lower profits. But in a world of strategic rivalry, firms may go to considerable lengths to conceal both profits resulting from specific actions and the actions themselves, and with imperfect competition they may be able to survive without necessarily maximizing profits.

The idea of looking for equilibria as the outcome of Darwinian selection leads to the concept of an Evolutionarily Stable Strategy (ESS), and this appears to be an attractive equilibrium concept in a world of imperfect information in which agents experiment, and either prosper or suffer in the resulting competition. Much of the subsequent literature has been driven to characterize the resulting equilibrium: will interaction lead to competitive outcomes (WE) or will firms learn to sustain imperfectly competitive outcomes, such as the Cournot Equilibrium (CE), or the even more profitable tacit collusion equilibrium (TCE)?

Thus Vallée and Yildozoglu (2009) note that an action that harms the performer less than it harms its rivals differentially advantages the performer. In some cases that may lead the performer to have a higher survival probability even though it makes itself worse off. Hamilton (1970) calls this a spite effect (familiar from trade literature as a “beggar-my-neighbour” policy). But if firms are reasonably capital intensive and have high sunk costs (as generators are), they are unlikely to follow a beggar-my-neighbour policy as the cost in lost variable profits may be high

compared to the gain of bankrupting competitors, particularly if there is free entry. Although in general not all Nash Equilibria (NE) are ESS, in the absence of the spite effect, Vallée and Yildozoglu (2009, p6) prove that a strict NE is an ESS in a symmetric finite population oligopoly game.

However, if survival depends on *relative* rather than absolute profitability - as in imitation-based *social learning* - the result is convergence to the WE. Specifically, Vallée and Yildozoglu (2009) show that strategies that imply a deviation from the WE will be eliminated while the strategies that imply a deviation from the CE in the direction of the WE will diffuse in the population. Social learning leads to WE because increasing output below the WE leads to a relatively greater profit for the deviant than the original population, so it is advantaged. Note that this requires less successful firms to exit unless they increase their output. Not surprisingly, if firms take the price as given they are likely eventually to behave competitively.

Vriend (2000) draws the distinction between *social learning* where survival depends on relative performance, and *individual learning* in which firms cannot observe and therefore imitate the actions of others, but instead experiment and learn from their own performance. If firms assume that they can change their output holding constant the price enjoyed last period (i.e. they ignore the effect their output will have on price next period) the WE is the ESS and the CE is not an ESS. But if firms ignore the (misleading) price last period and continue to try out strategies (output levels) for long enough to learn how they compare, they can converge to the CE. Models of social learning therefore seem inappropriate for modelling oligopolistic markets.

Waltman and Kaymark (2008, p3277) note two approaches to individual learning: belief-based learning and reinforcement learning. “Examples of belief-based learning models are Cournot adjustment and fictitious play (Fudenberg and Levine, 1998). These two models assume that an agent has the ability both to observe its opponents’ action choices and to calculate best responses. In a Cournot oligopoly game, the models predict that firm behavior can converge only to the Nash equilibrium. Compared with belief-based learning models, reinforcement learning models make few assumptions about both the information available to an agent and the cognitive abilities of an agent. . . . An agent is only assumed to have knowledge of the strategies that it can play and, after playing a strategy, of the payoff that it has obtained from that strategy.” Given that, they show (under some rather delicate conditions) that reinforcement learning can result in a TCE, the least competitive alternative to a WE.

What firms observe, and whether they can imitate strategies or only choose between strategies and learn of their outcome, clearly matters. Tirole (2000, p240) notes that “the proliferation of theories is mirrored by an equally rich array of behavioral patterns actually observed under oligopoly.” Tirole (p261) also notes that the evolutionary approach “looks at strategies, or (better

terminology) rules, that are “robust” in the sense that they do relatively well against a variety of other rules. . . . and that “only actors who use robust rules will stay around. . . . What is meant by “robust rule” is *a priori* contingent on the set and the probability distribution of mutations, and on the set of strategies the robust rule is allowed to be compared to.”

All this argues for individual learning in mark-up models as appropriate for modelling electricity markets. There are several papers on the theory of mark-up pricing in learning models, mostly based on mark-ups on average costs (AC). One that would seem to parallel our interests closely is the paper by Pasche (2002), which looks at two forms of mark-up: either as an absolute difference between price and AC, or as a proportion of the AC. He defines a number of equilibrium concepts for Stackelberg mark-up deviations from Cournot or Bertrand equilibria for the differentiated duopoly case.

Al-Najjar et al. (2008) look at mark-up choices when firms follow “naïve” adaptive learning to adjust prices and periodically try different “costing methodologies” to see whether they can improve on their price searching behaviour, as part of a research programme to reconcile cost misallocation and “irrational” pricing that appear to characterize actual firm behaviour and even the management textbook advice. This follows from observations that companies “use full costs rather than variable costs” (Maher, Stickney and Weil, 2004) but the same authors also note that “cost-based pricing is far less prevalent in Japanese process-type industries (for example, chemicals, oil and steel)” to which we would add electricity.

Agent-based models and learning

Agent-based models define a number of agents (e.g. firms, traders, consumers) and specify their rules of behaviour – the actions they can take and the information available and on the basis of which they learn and update their behaviour. The literature on models of learning (see also Camerer, 2003) argues for the suitability of reinforcement learning for modeling industrial commodity and particularly electricity markets. Reinforcement learning models have been studied in economics (e.g. Roth and Erev, 1995) and in the artificial intelligence literature (e.g. Sutton and Barto, 1998). Q-learning is a reinforcement learning model originally developed in the field of artificial intelligence (Watkins, 1989, further developed by Littman, 1994 and Hu and Wellman, 2006) and appears to be the preferred model for electricity markets. An agent using a Q-learning algorithm keeps in memory a Q-value function of the weighted average of the payoffs obtained by playing a certain action in the past. The agent then plays with high probability the action that gives the highest payoff and with a small probability a randomly chosen different action (to test that any optimum found is not just a local maximum), observes the payoff it obtains and then

updates its Q-value (Krause et al., 2006; Waltman and Kaymak, 2008).⁵ Supply function models, developed by Klemperer and Meyer (1989) and applied to electricity markets by Green and Newbery (1992), are attractive intermediate formulations. In particular, their linear solutions have been influential in motivating the kind of agent-based models often chosen. In the simple quadratic cost, linear demand model there are two simple types of deviation from competitive bidding - marking-up the offer schedule by a constant amount, or proportionally marking up the slope of the offer schedule, as in Hobbs et al. (2000). If marginal costs are linear, then there is a supply function equilibrium which is a proportional mark-up on the marginal cost, and this solution has been widely used in simple analytical models (e.g. Green, 1999).⁶

To summarize, a number of markets, for example electricity markets, can be explored through agent-based modelling. In such models, agents learn by testing out deviations from past strategy choices to see if they can increase profits, and continue to experiment until further improvements can no longer be found. The model must specify exactly what can and cannot be observed, and in many cases this is restricted to information about the firm's profit and his past choices (which can be computed from past outputs, the cost function and the associated market prices). As such the information is even more imperfect than normally assumed in formal game-theoretic studies, but may avoid some of the evident sensitivities of such studies to the precise form of the information assumed.

Agent-based modelling of wholesale electricity markets

There are a wealth of examples of using such models to explore learning in complex markets. Wholesale electricity markets are one example. Learning is seen where generators (and in some cases other agents such as retailers or suppliers) might choose a variety of actions in their pursuit of profit or advantage. In some, agents can sign contracts ahead of time as well as deciding at what price to offer electricity into the spot market. Except for very stylized (e.g. Cournot) models, it has been analytically very difficult to solve for the optimal combination of forward contracts and spot sales in the presence of market power. Newbery (1998) was able to do this for a supply function equilibrium model by assuming that existing generators would choose contracts to deter entry that might otherwise lead to excess capacity. Unless average prices can be externally specified (in this case by free entry), there are too many possible combinations of

⁵See Littman (1994) and Hu and Wellman (2003). Weidlich and Veit (2008) critically compare these with other reinforcement learning models.

⁶With finite support to the distribution of demand and no capacity limits as here, there will be a continuum of supply function equilibria, of which one is the linear supply function, which is the optimal response of any firm provided all other firms have chosen that supply function.

contracts and spot offers to be tractable in a supply function setting (arguably the most natural choice for modelling electricity wholesale markets).

Bunn and Oliveira (2001) adopt an agent-based model to explore the possible implications of moving from the original Electricity Pool for England and Wales to the New Electricity Trading Arrangements (NETA) to see how this might impact on both generators and suppliers, given that generators may have market power and can choose to contract ahead but also be exposed to the Balancing Mechanism. Given that their model was developed before NETA went live some of their predictions were remarkably prescient (such as the predicted low volume of trading in the Balancing Mechanism), and illustrate the strength of this approach for studying proposed market design changes in a realistic setting that includes contracting.

Bunn and Oliveira (2003) extend this model to explicitly consider market power in a model of the English wholesale market in which the regulator, Ofgem, had imposed a Market Abuse Licence Condition that was appealed to the Competition Commission. Again, the development of their model was motivated by a practical policy question. Veit et al (2006) also study the case in which an oligopoly of generators sign contracts ahead of time and then compete in a transmission constrained spot market. In this case agent-based modelling is chosen to handle the complexity of the decision process (complicated by the transmission constraints), and confirms the prediction of simpler analytical models (e.g. Newbery, 1998) that forward contracting to sell leads to more competitive spot market behaviour and hence lower prices.³ These and other agent-based models of wholesale electricity market are discussed in the excellent survey by Weidlich and Veit (2008), who compare different learning strategies and their results.

A key question facing such modelers is whether the resulting equilibrium is indeed a Nash equilibrium (where that is unique) in the space of actions allowed in the formulation of the game, and indeed what happens where there are multiple Nash equilibria. This is the question that Krause et al. (2005) address in the context of a simplified power market, and answer affirmatively for unique Nash equilibria.

The obvious problem with agent-based modelling where agents are assumed to learn about the profit consequences and adapt their strategies to increase profits is that the action space over which they make choices may be too restrictive and may allow other more sophisticated agents to exploit this limited strategy choice by choosing from a wider range of actions. In that sense the models may be dismissed as too simplistic to model the behaviour of more sophisticated firms (who certainly hire the brightest and best to examine their strategic choices). A good defence of adaptive learning would be to show that the equilibrium of the form of learning were robust against more sophisticated players choosing from a wider set of actions.

Additional References for Appendix B

- Al-Najjar, N., Baliga, S. and Besanko, D., (2008), “Market forces meet behavioral biases: cost misallocation and irrational pricing.” *RAND Journal of Economics*, 39, 214-237.
- Bunn, D.W. & Oliveira, F.S., (2001). “Agent-based simulation: an application to the new electricity trading arrangements of England and Wales.” *IEEE Transactions on Evolutionary Computation*, Special issue: Agent Based Computational Economics 5 5, 493-503
- Bunn, D.W. & Oliveira, F.S., (2003). “Evaluating individual market power in electricity markets via agent-based simulation.” *Annals of Operations Research*, 121 1-4, 57–77.
- Camerer, C. F., (2003), *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, Princeton, New Jersey.
- Friedman, M., (1953). *Essays in Positive Economics*, University of Chicago Press, Chicago.
- Fudenberg, D. & D. Levine, (1998), “Learning in games.” *European Economic Review*, 42(3–5), 631–639.
- Green R., (1999). “The electricity contract market in England and Wales.” *The Journal of Industrial Economics*, 47, 107-124.
- Hamilton, W.D. , (1970). “Selfish and Spiteful Behavior in an Evolutionary Model.” *Nature*, 228, 1218-1220.
- Hobbs, B. C. Metzler, and Pang, J.-S., (2000). “Strategic gaming analysis for electric power systems: an MPEC approach.” *Power Systems, IEEE Transactions*, 15, 638-645.
- Hu, J. and M. Wellman, (2003). “Nash Q-learning for general-sum stochastic games”, *Journal of Machine Learning Research*, vol. 4, pp. 1039–1069.
- Krause, T., G. Andersson, D. Ernst, E. Vdovina-Beck, R. Cherkaoui, and A. Germond, (2005). “A comparison of nash equilibria analysis and agent-based modelling for power markets.” in Power Systems Computation Conference (PSCC), Liege.
- Littman, M. (1994). “Markov games as a framework for multiagent reinforcement learning.” in *Proceedings of the Eleventh International Conference on Machine Learning*. 157-163, San Francisco, CA: Morgan Kaufmann.
- Maher, M., Stickney C. and Weil P.L., (2004). “Market forces meet behavioral biases: cost misallocation and irrational pricing.” *Managerial Accounting*, 8th ed. Mason, OH: South-Western, 2004.
- Pasche M. , (2002). “Heterogeneous Behavioral Rules in the Oligopolistic Case.” University of Jens WP 2002/01.
- Roth, A. E. and Erev, I. , (1995). “Learning in extensive form games: experimental data and simple dynamic models in the intermediate term.” *Games and Economic Behavior*, 8, 164-212.
- Sutton, R. S. and Barto, A. G., (1998). *Reinforcement learning: An introduction*. MIT

Press, Cambridge, MA.

Vallée, T., and Yildozoglu M., (2009). “Convergence in the Finite Cournot Oligopoly with Social and Individual Learning.” GREQAM 2009-04.

Van Damme and Hurkens, (1997). “Games with Imperfectly Observable Commitment.” *Games and Economic Behavior*, 21, 282-308.

Veit, D.J., Weidlich, A., Yao, J., Oren, S.S., (2006). “Simulating the dynamics in two- settlement electricity markets via an agent-based approach.” *International Journal of Management Science and Engineering Management*, 1, 83—97.

Vriend, N.J., (2000). “An illustration of the essential difference between individual and social learning, and its consequences for computational analyses.” *Journal of Economic Dynamics & Control*, 24,1-19.

Waltman, L. and Kaymak, U., (2008). “Q-learning agents in a Cournot oligopoly.” *Journal of Economic Dynamics and Control*, 32, 3275-3293.