J Evol Econ (2001) 11: 433–456

Journal of
Evolutionary
Economics

Firm routines, customer switching and market selection under duopoly*

Martin Currie and Stan Metcalfe

School of Economic Studies and ESRC Centre for Research on Innovation and Competition, Victoria University of Manchester, Manchester M13 9PL, UK

Abstract. This paper explores the dynamics of market selection for an industry in which firms employ relatively simple pricing, production and investment routines and in which consumers switch between rival firms in response to price differentials but do not all do so instantaneously. The key issue is whether market processes result in the elimination of less efficient firms by their more efficient rivals. That is to say, do such processes unfailingly increase the efficiency with which available economic resources are used? In the context of duopoly, we show that the survival of the more efficient firm is not guaranteed and that, more generally, the outcome depends upon the speeds with which firms adjust prices and capacities and with which customers switch between rival firms.

Key words: Non-linear dynamics – Firm adjustment routines – Market selection

JEL classification: C61, D21, D42, D43

1 Introduction

This paper explores the dynamics of market selection for an industry in which firms employ behavioural routines and in which consumers switch between rival firms in response to price differentials but do not all do so instantaneously. In so doing, this paper confronts a central issue in the evolutionary analysis of market processes; namely, do such processes inexorably result in the elimination of less efficient firms by their more efficient rivals? That is to say, do such processes unfailingly increase the efficiency with which available economic resources are used? The framework is evolutionary in the very precise sense that it is concerned with the dynamic processes that flow from the existence of micro-diversity. It is

* The authors are very grateful for valuable and constructive suggestions received from Pasquale Commendatore, Ingrid Kubin, Ian Steedman and an anonymous referee.
Correspondence to: M. Currie

also evolutionary in its insistence that firms are necessarily constrained to employ relatively simple behavioural routines.

Our firms, which use machines to produce a homogeneous product, compete directly through price setting. Each firm's pricing routine involves increasing its price in response to buoyant demand (manifested in low and falling stock levels) and reducing price in response to depressed demand (manifested in high and rising stock levels). A firm's production routine simply entails fully using its current capacity. Its investment routine involves expanding (contracting) its productive capacity if its going rate of return on machine ownership exceeds (is less than) its target rate of return. Consumers switch between firms in response to price differentials but do so with some degree of inertia. The speeds with which firms adjust their prices and their stocks of machines and the speed with which customers switch in response to price differentials are naturally crucial determinants of the dynamical behaviour of the industry.

Our model differs fundamentally from the recent dynamic duopoly model of Puu (1998). He demonstrates the possibility of chaotic behaviour under duopoly but for an industry in which – *à la* Cournot and Stackelberg – profit-maximising firms set quantities and each receives the market-clearing price. In contrast, our firms themselves set prices. Furthermore, in contrast both to the model of Puu and to the recent game-theoretic equilibrium models of oligopoly based on the concept of Markov perfect equilibria, our analysis assumes that the firms employ relatively simple behavioural routines that reflect the fact that they do not have as much information about their rivals' circumstances and motivation as they do about their own. Specifically, rather than acting on conjectures about the concurrent decisions of rivals, our firms act on *internally available information*, notably on knowledge of the levels and movements in their own product stocks. Our insistence on modelling in terms of algorithmic, routinised behaviours follows the tradition in economic theorising that dates from the famous work of Cyert and March (1963). A central presumption of this approach is its reliance on information arising *within* the firm itself as its actions work themselves out in the competitive process. Our model also differs sharply from Bertrand-like models that assume that every consumer always patronises the firm setting the lowest price. Recognising that such models cannot capture the sort of price competition that characterises many imperfectly competitive industries, some economists have explored the implications of the reluctance of consumers to change their suppliers. Rosenthal (1982) assumes that consumers stay with their existing supplier unless the latter raises price, in which case all its previous customers look for the lowest price *à la* Bertrand. More recently, Chen and Rosenthal (1996) assume that in any period a constant number of consumers switches from the high-price firm to the low-price firm. Instead, in a hypothesis similar to that invoked *inter alia* by Phelps and Winter (1970) and by Metcalfe (1998), we assume that a firm's share in the total number of customers changes in response to the difference between its own price and a weighted average price for the industry.

The model is explained in Section 2. Even though the firms' routines are relatively simple, their interaction with each other and with the process of customer switching gives rise to a highly complex non-linear dynamical system. Given this complexity, a natural way of proceeding is, first, to focus on an individual firm assuming that it monopolises the market. Accordingly, in Section 3, we characterise a stationary equilibrium for a monopolist and then, using simulations, explore whether the routines would converge on that equilibrium. A powerful tool for investigating dynamical behaviour is a bifurcation diagram – that is, a diagram that shows the qualitative long-term behaviour of one of the system's variables as a multi-valued function of one of the system's parameters. We use this technique to explore the comparative dynamic impact of changes in the monopolist's price adjustment speed on the nature of the time path of its price. We also identify those combinations of price and capital stock adjustment speeds that would result in convergence on equilibrium. In Section 4, we focus on a duopoly in which micro diversity is manifested solely in terms of differences in the firms' costs of operating machines. The central question is whether the market process results in the selection of the more efficient firm. Notwithstanding that the firms produce a homogenous commodity, that consumers have identical linear demand curves, that there are no stochastic forces, that there is no form of increasing returns and that firms are identical except for a (non-trivial) cost difference, the dynamical process of firm adjustment and customer switching may result in the 'selection' of the high-cost firm. Indeed, there may be perpetual coexistence of the firms, necessarily involving periodic or chaotic industry behaviour. The analysis suggests certain *ceteris paribus* tendencies. Specifically, the low-cost firm is more likely to survive (a) for adjustment speeds that would lead to convergence under monopoly; (b) for lower customer switching speeds; (c) for greater cost differentials; and (d) for higher initial market shares for the low-cost firm. In Section 5, we explore the robustness of these propositions by examining briefly alternative specifications of the model. This includes considering the implications of alternative firm routines. However, since something has to be taken as given in any model, we do not consider 'meta-routines' that would amount to specifying rules for altering rules in response to changing circumstances. In sum, our approach is consonant with the claim of Winter (1971) that '. . . the proposition "firms establish decision rules and apply them routinely over extended periods" is sufficiently significant, obvious, and well documented to deserve a prominent place in theoretical characterisations of firm behaviour' (p. 239).

2 Model

2.1 Period t

The industry comprises F firms that use the services of machines to manufacture a homogeneous product.[1] The timings of the firms' decisions and activities are

[1] Although we are primarily interested in market selection for the case of duopoly, the model is applicable to the case where $F > 2$. In Section 5.1, we consider very briefly simulations for the case where there are initially three firms in the industry.

synchronised. At date t, which constitutes the transition between period $(t-1)$ and period t, firm f owns both a stock of machines, k_t^f, which determines its productive capacity over period t, and a stock of the product, s_t^f, carried over from period $(t-1)$. At date t, the firm takes three decisions. First, it sets a selling price, p_t^f, to which it commits itself for the duration of the t^{th} period. Second, it decides its production flow, q_t^f, for the t^{th} period. Assuming that fully using one machine for the duration of one period yields one unit of the product, the firm's capacity constraint is $q_t^f \leq k_t^f$. Finally, at date t, the firm orders new machines, I_t^f, to be delivered, with an institutionally given lag, at date $(t+1)$. In Section 2.2, we specify the firms' pricing, production and investment routines.

Following the setting of prices at date t, each consumer decides whether to patronise the same firm as in the previous period or to switch to some other firm. In Section 2.3, we explain the process whereby consumers shift between firms in response to price differentials. The stationary total number of consumers in the market being n, the number of customers who patronise firm f during period t is $n_t^f = \omega_t^f n$, where ω_t^f denotes the firm's 'market share'. Each consumer, who does not carry stocks, has the same stationary demand curve, $a - p_t$ for $p_t \leq a$. The (aggregate) demand curve facing firm f in period t is then:

$$d_t^f = \begin{cases} n_t^f \left(a - p_t^f\right) & \text{if } p_t^f \leq a \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where d_t^f denotes the demand for the firm's product over the period. Since the firm has a maximum of $\left(s_t^f + q_t^f\right)$ available for sale during period t, the quantity that it sells is:

$$x_t^f = \min\left\{d_t^f;\ s_t^f + q_t^f\right\} \tag{2}$$

where demand, output and sales are accumulated flows over the period. Where the firm cannot meet demand, it may ration consumers on a first-come-first-served basis or on a *pro rata* basis. Firm f will begin the next production period at date $(t+1)$ with a product stock given by:

$$s_{t+1}^f = s_t^f + q_t^f - x_t^f \tag{3}$$

and with a stock of machines given by:

$$k_{t+1}^f = (1-\delta)\, k_t^f + I_t^f \tag{4}$$

where $0 < \delta < 1$ is the rate at which machines depreciate per period.

A key feature of the model is that there can be efficiency differences between firms. These take the form of differences in their costs of operating machines. The cost incurred by firm f in operating one machine per period (including wages, the cost of raw materials, etc.) is denoted by ν^f. In contrast, all the firms face the same acquisition cost for machines. The stationary cost of a machine, denoted by m, is incurred at the time of delivery and includes any installation cost. If necessary, a firm can finance the acquisition of machines by borrowing at a given stationary market rate of interest, where i is the rate applicable to the period. Whereas firm f's short-period (marginal and average) cost is ν^f, its *full* long-period cost is $\nu^f + (\delta + i)m$, where $(\delta + i)m$ constitutes the 'owner cost of capital' for the period, δm being the machine replacement cost and im being the interest cost associated with machine ownership *per se*.[2]

2.2 Firms' pricing, production and investment routines

At date t, firm f decides its price, production and investment in that sequence. Consider firm f's pricing routine. The price set by the firm impacts on the demand it faces during the ensuing t^{th} period by influencing both the number of its customers and the demand per customer. However, since the number of its customers depends on the prices being set *simultaneously* by rival firms, the firm cannot know at date t the demand curve that it faces. Furthermore, given that it does not have information on the current product stocks or current productive capacities of rival firms, the firm does not indulge in speculations about their concurrent pricing decisions. Instead, the firm's pricing routine is based entirely on *internally available* information, namely, on its directly observable product stocks. Specifically, the firm takes account *both* of the discrepancy between its current stock level and its desired or target stock level *and* of the direction of change in its stock level over the previous period. Low and falling stocks – signalling buoyant demand – would invite a price rise. High and rising stocks – signifying depressed demand – would invite a price cut. However, the decision to change price is not taken lightly. Even if stocks differ from the desired level, this does *not* trigger a price change provided that stocks moved in the right direction over the previous period.

We further suppose that the Pricing Department is instructed not to set a price below a certain minimum. At one extreme, the constraint on the Pricing Department might be that it simply cover operating costs, that is, $p_t^f \geq \nu^f$. At the other extreme, it might be charged with covering the full cost, that is, $p_t^f \geq \nu^f + (\delta + i)m$. Initially, we assume a procedural rule, between these extremes, whereby the Pricing Department is expected to cover machine replacement costs as well as operating costs, that is, $p_t^f \geq \nu^f + \delta m$. It seems plausible that, notwithstanding the fact that the use of its current machines imposes no opportunity cost, a firm with some degree of market power would employ a pricing routine that would take account of its machine replacement cost as well as its operating cost. At the same time, we would not expect a firm to refuse to contemplate a price that did not cover its full cost, including the cost of machine ownership. However, these are not matters on which we would wish to be dogmatic and, in Section 5.3, we contrast briefly the implications of imposing either the weaker lower-bound, $p_t^f \geq \nu^f$, or the stronger lower-bound, $p_t^f \geq \nu^f + (\delta + i)m$.

We formalise firm f's price adjustment process as follows. Its notion at date t of what would constitute an appropriate product stock level, denoted by $\hat{s}_t^f$,

[2] See Winston (1982, pp. 54–56) for a careful delineation of different 'prices of capital'.

depends on its sales in the previous period:

$$\hat{s}_t^f = \gamma^f x_{t-1}^f \tag{5}$$

where $\gamma^f > 0$. Expressing its actual and desired stocks relative to its current productive capacity, the firm perceives its stocks to be too high if s_t^f / k_t^f exceeds $\gamma^f \left(x_{t-1}^f / k_t^f \right)$. Thus γ^f may be interpreted as the desired ratio of stocks to productive capacity in the case where the firm's current capacity equals its sales over the previous period. The firm's price adjustment routine is then:

$$\begin{aligned} & p_t^f = p_{t-1}^f && \text{for } \left(s_t^f - \hat{s}_t^f \right) \left(s_t^f - s_{t-1}^f \right) < 0 \\ & \frac{p_t^f - p_{t-1}^f}{p_{t-1}^f} = \theta^f \left(\frac{s_t^f - \hat{s}_t^f}{k_t^f} \right) && \text{otherwise} \\ & \text{subject to } p_t^f \geq \nu^f + \delta m \end{aligned} \tag{6}$$

where $\theta^f > 0$ is firm f's price adjustment 'speed'. Thus if stocks are too high ($s_t^f > \hat{s}_t^f$), the firm would not alter price provided that stocks had fallen over the previous period ($s_t^f < s_{t-1}^f$) but it would reduce price if they had risen ($s_t^f > s_{t-1}^f$). Conversely if stocks are too low, the firm would not alter price provided that stocks had risen over the previous period but it would increase price if they had fallen.

Given that the minimum price that the firm would contemplate is strictly above its machine operating cost and that machines depreciate irrespective of use, there is no reason for leaving machines idle. That is, as long as firm f remains in the industry, its production routine is simply:

$$q_t^f = k_t^f \tag{7}$$

Consider finally firm f's investment routine. The firm's investment decisions are driven by a *target* rate of return on the ownership of a machine, denoted by τ^f. On the basis of the newly-set price, the firm's *going* rate of return per machine is:

$$\rho_t^f = \frac{p_t^f - \nu^f - \delta m}{m} \tag{8}$$

If this going rate exceeds (is less than) the target rate, the firm expands (contracts) its stock of machines. Specifically:

$$\frac{k_{t+1}^f - k_t^f}{k_t^f} = \kappa^f \left(\rho_t^f - \tau^f \right) \quad \text{subject to} \quad \frac{k_{t+1}^f - k_t^f}{k_t^f} \geq -\delta \tag{9}$$

where $\kappa^f > 0$ denotes the firm's capital stock adjustment speed and where the lower bound, equivalent to $k_{t+1}^f \geq (1 - \delta) k_t^f$, reflects its inability to sell second-hand machines. The corresponding gross investment in machines is:

$$I_t^f = \delta k_t^f + \kappa^f \left(\rho_t^f - \tau^f \right) k_t^f \quad \text{subject to} \quad I_t^f \geq 0 \tag{10}$$

Replacement investment, δk_t^f, is thus adjusted (upwards or downwards) in the light of the difference between the going and target rates of return.

Since the investment routine is so central to the dynamical behaviour of the firm (and industry), it is worth making explicit two *equivalent* specifications. For the first, define the firm's Marshallian long-period *supply price*, p_s^f, as:

$$p_s^f = \nu^f + \left(\delta + \tau^f\right) m \tag{11}$$

that is, as the product price that not only covers the machine operation and replacement costs but also yields the target return on machine ownership. Substituting (8) into (9) and using (11) gives:

$$\frac{k_{t+1}^f - k_t^f}{k_t^f} = \frac{\kappa^f}{m} \left(p_t^f - p_s^f\right) \tag{12}$$

subject to $k_{t+1}^f \geq (1 - \delta) k_t^f$. That is, the firm increases (decreases) its stock of machines if its current product price exceeds (is less than) its long-period supply price.

For the second equivalent specification, recall that the firm can borrow or lend at a given stationary market rate of interest, i. The firm's target rate of return, τ^f, embodies both the rate of interest and a target *pure* rate of return, $\tilde{\tau}^f$, that is, $\tau^f = i + \tilde{\tau}^f$. The going *pure* rate of return per machine is:

$$\tilde{\rho}_t^f = \frac{p_t^f - \nu^f - (\delta + i)\, m}{m} = \rho_t^f - i \tag{13}$$

Since $\rho_t^f - \tau^f = \tilde{\rho}_t^f - \tilde{\tau}^f$, routine (9) is equivalent to:

$$\frac{k_{t+1}^f - k_t^f}{k_t^f} = \kappa^f \left(\tilde{\rho}_t^f - \tilde{\tau}^f\right) \tag{14}$$

subject to $k_{t+1}^f \geq (1 - \delta) k_t^f$. That is, firm f expands (contracts) its stock of machines if its going *pure* rate of return exceeds (is less than) its target *pure* rate of return.

2.3 Customer switching and demand

Following the setting of prices at date t, each consumer decides whether to patronise the same firm as in the previous period or to switch to some other firm. Having decided which firm to patronise, a consumer does not switch to another firm within the period. Our customer switching hypothesis, for $\omega_{t-1}^f > 0$, is:

$$\frac{\omega_t^f - \omega_{t-1}^f}{\omega_{t-1}^f} = \sigma \, \frac{\sum_j \omega_{t-1}^j p_t^j - p_t^f}{\sum_j \omega_{t-1}^j p_t^j} \tag{15}$$

where $\sigma \geq 0$ is the customer switching speed. It is easily confirmed that, for (15), $\sum_{f=1}^{F} \omega_{t-1}^{f} = 1$ implies $\sum_{f=1}^{F} \omega_{t}^{f} = 1$. The speed, σ, reflects the degree of inertia in the market. *Ceteris paribus* the greater consumer loyalty and / or the greater the impediments to the flow of information, the lower would be σ. For the limiting case of $\sigma = 0$, each firm would be a monopolist in its own segment of the market, price differentials being immaterial. For $\sigma > 0$, the market share of a firm changes in response to any differential between its own price and the weighted average price, $\sum_{j} \omega_{t-1}^{j} p_{t}^{j}$, where the latter gives a greater weight to a firm's price the larger its market share in the previous period.[3] A simple rationale for this formulation is that the rapidity with which information is disseminated through contacts between customers of different firms depends on the firms' market shares. If a firm loses all its customers, it cannot re-capture them, that is, if $\omega_{t'}^{f} = 0$ at some date t', then $\omega_{t}^{f} = 0$ for $t > t'$. We would not deny the possibility that, in reality, a firm that has lost all its customers might re-enter a market. However, we regard entry and re-entry as idiosyncratic acts that (similar to innovation) are not amenable to representation by simple decision criteria.

3 Dynamics of the individual firm

As a heuristic device to disentangle the dynamic operation of the pricing, production and investment routines from the process of customer switching, we first explore the behaviour of an individual firm on the assumption that there is no customer switching ($\sigma = 0$), so that the firm in question monopolises (a segment of) the market.[4] Assume then that the firm has a (stationary) number of customers, n. Once we specify initial conditions $\{p_0;\ x_0;\ s_0; s_1;\ k_1\}$ encapsulating the relevant pre-history of the firm, its future behaviour follows deterministically.[5] In Section 3.1, we characterise a stationary equilibrium for the monopolist and in Section 3.2, using simulations, we identify those combinations of the price and capital stock adjustment speeds that would result in convergence on the stationary equilibrium.

3.1 Long-period equilibrium

Figure 1 shows a stationary equilibrium for the firm. The price in a thoroughgoing stationary state, denoted by $\overline{p}$, is determined by the investment routine. Thus, attainment of the target rate of return requires:

[3] The model could be modified to allow customer switching to depend also on the extent of unsatisfied demands in the previous period. In this case, the ways in which firms ration customers would be relevant.

[4] We are not proposing here a model of pure monopoly *per se*. In justifying the pricing routine, our argument as to why the individual firm cannot know the demand curve that it faces was couched in terms of a lack of information about the concurrent decisions of rival firms.

[5] Since we are concerned with an individual firm in this Section, we dispense with the firm superscript.

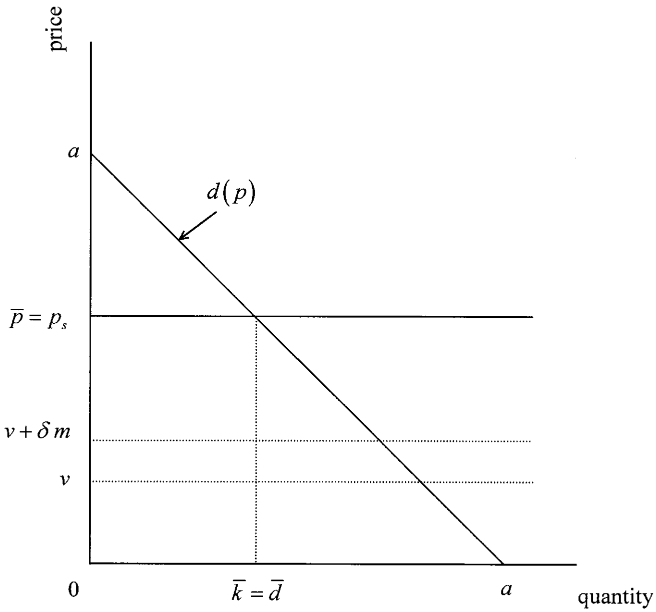

Fig. 1

$$\bar{p} = v + (\delta + \tau)\ m = p_s \tag{16}$$

that is, the stationary price equals the firm's long-period supply price, p_S. At this price, the stream of quasi-rents from the future use of a new machine, discounted back to the delivery date using the firm's (target / achieved) rate of return, is given by:

$$(\bar{p} - v)\left(\frac{1}{1+\tau} + \frac{1-\delta}{(1+\tau)^2} + \frac{(1-\delta)^2}{(1+\tau)^3} + \ldots\right) = \frac{\bar{p} - v}{1-\delta}\sum_{j=1}^{\infty}\left(\frac{1-\delta}{1+\tau}\right)^j = \frac{\bar{p} - v}{\delta + \tau} \tag{17}$$

In the stationary state, this present value equals the cost of the machine at the delivery date. Demand at $\bar{p}$ determines the stationary capacity, production and sales per period:

$$\bar{k} = \bar{q} = \bar{x} = \bar{d} = n\,(a - \bar{p}) \tag{18}$$

where, by assumption, $a > \bar{p}$. The firm's replacement investment is $\bar{I} = \delta\bar{k}$ and its product stock is $\bar{s} = \gamma\bar{x}$. Henceforth we refer to the values $\{\bar{p};\, \bar{k};\, \bar{q};\, \bar{x};\, \bar{d};\, \bar{I};\, \bar{s}\}$ as the 'fixed point' for the dynamical system for the monopoly.

If the firm's target rate of return were equal to the market rate of interest (i.e. if its target *pure* rate of return were zero), this stationary state would correspond to a 'competitive equilibrium'. However, we would expect that the firm's target rate of return would reflect *inter alia* its perception of its own product market

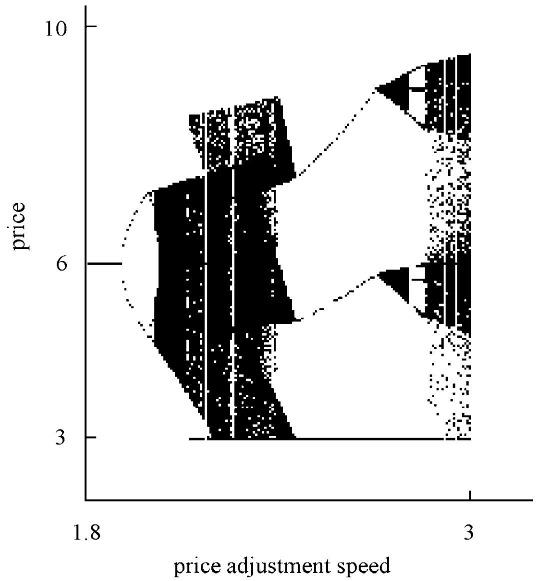

Fig. 2

power and, as such, would exceed the market rate of interest. How would our stationary equilibrium compare with an equilibrium for an *omniscient* monopolist? Maximising the discounted present value of the firm would require that marginal revenue equal $(\nu + (\delta + i)\ m)$, where the latter is the full marginal cost.[6] Depending on the parameters, the corresponding rate of return of the omniscient monopolist could be above, equal to or below the target rate of our routine-based firm. Accordingly, the former's price could be above, equal to or below $\bar{p}$.

3.2 Monopoly dynamics

But would the routines converge on a stationary equilibrium? To identify the possible long-term dynamics, we employ simulations based on $n = 5,000$, $a = 10$, $m = 10$, $\delta = 0.1$, $\nu = 2$, $\gamma = 0.2$ and $\tau = 0.3$. These parameter values imply a cost, $\nu + \delta m = 3$; a fixed-point price, $\bar{p} = 6$; and a fixed-point capacity, $\bar{k} = 20,000$. In order to separate the pricing and investment routines, consider initially the operation of the pricing routine on the assumption that capacity and production are stationary at the fixed-point level. Figure 2 is a bifurcation diagram showing price as a multi-valued function of the price adjustment speed, θ. Assuming

[6] Maximising the rate of return per machine would, of course, not be a meaningful objective.

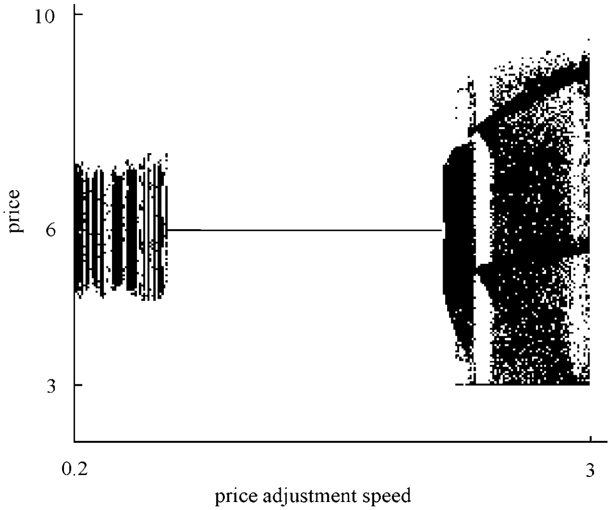

mu = 1
tau = 0.3
gamma = 0.2
delta = 0.1
v = 2

m = 10
T = 3000 plot = 1000
p-min = v + m x delta
pinit = 0.99 p*

Fig. 3

initial conditions $p_0 = 0.99\bar{p}$, $x_0 = n(a - p_0)$ and $s_1 = s_0 = \gamma x_0$, a sequence of prices is generated by iterating the system for 3,000 periods for values of θ between 1.8 and 3. In order to identify the *long-term* behaviour of price, the first 1,000 periods are discarded. The bifurcation diagram is generated by plotting the ensuing 2,000 prices as a function of θ. At 'slow' price adjustment speeds, the firm's price converges on $\bar{p} = 6$. However, as θ increases through 1.91, the fixed point becomes unstable and the system is attracted to a period-two cycle. At $\theta \cong 1.99$, the period-two cycle bifurcates into a period-four cycle. For θ above 2.13, the constraint that the firm would not set a price below 3 impacts on its behaviour. A period-three cycle, which can be seen in the diagram, extends from $\theta \cong 2.46$ to $\theta \cong 2.71$. Figure 2 confirms that, even with a stationary production level, the operation of the price adjustment routine can by itself generate complex long-term dynamics.

We now re-introduce the investment routine. In the full system, the pricing routine (6) plays more than simply a passive role of striving for market-clearance; the going price itself feeds back on investment. Thus, whilst the investment routine is the real driving-force in that it determines the fixed-point price, the interaction of the routines determines whether or not the system converges on that price. Figure 3 is a bifurcation diagram for price with respect to the price adjustment speed, assuming a capital stock adjustment speed of $\kappa = 1$ and an initial capital stock of $k_1 = n(a - p_0)$. The full system exhibits complex dynamics

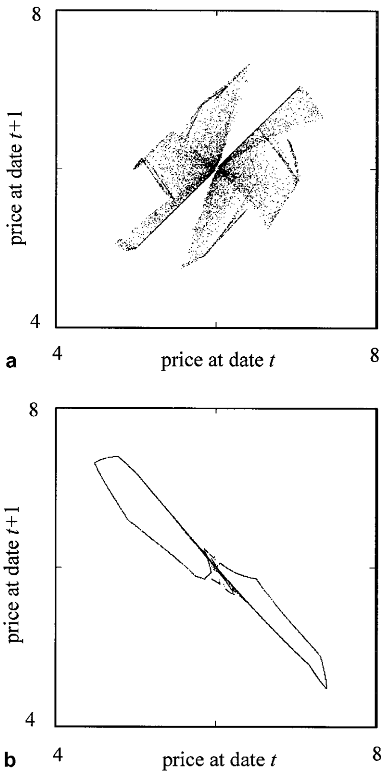

Fig. 4a,b

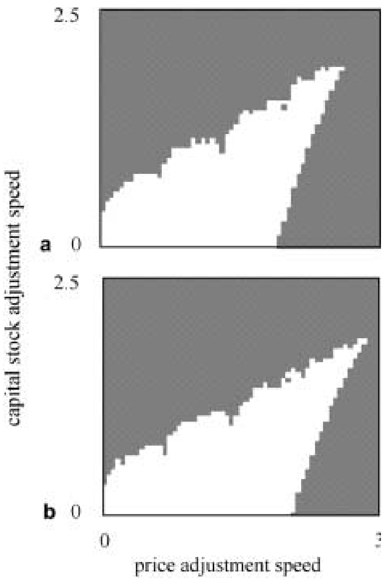

Fig. 5. a $\nu = 2$. **b** $\nu = 1.8$

for price adjustment speeds below 0.79 and above 2.22 but converges on the fixed point for intermediate speeds. Figure 4a and b shows the price attractors[7] at $\theta = 0.75$ and $\theta = 2.25$, respectively. The stark contrast between the attractors exhibits clearly the possible sensitivity of the price dynamics to relatively small changes in the operation of the routines.

Figure 5 provides a point of reference for our subsequent analysis. Figure 5a is a ‘convergence bitmap’ for $\nu = 2$. Specifically, the price adjustment speed θ ranges (in increments of 0.05) from 0.1 to 3 and the capital stock adjustment speed κ ranges (in increments of 0.05) from 0.1 to 2.5. For each combination of speeds, the system is iterated for 1,000 periods. A white square indicates that, by the 1,000$^{\text{th}}$ period, a stationary price of 6 has been achieved (within + or -0.0001), whereas a grey square designates cyclical or chaotic price behaviour. That is, the white region in Figure 5a shows those combinations of the price and capital stock adjustment speeds that result in convergence on the fixed-point price. The diagram confirms the complex nature of the dynamic interaction between the routines. For example, for a given price adjustment speed, the system may

[7] For our attractors, we discard the first 1,000 iterations and, to obtain sufficient detail, plot the subsequent 10,000 iterations.

exhibit complex dynamics for both ‘low’ and ‘high’ capital stock adjustment speeds but converge on the fixed point for intermediate speeds.

In Section 4, we explore the dynamics of a duopoly where one firm has an operating cost of 2 and the other an operating cost of 1.8. Accordingly Figure 5b is a convergence bitmap for $v = 1.8$ (and a corresponding fixed-point price of 5.8). Limitations of space preclude illustrating the comparative dynamic effects of changes in other parameters. However, we should note that – given constant costs, a linear consumer demand curve and the specified firm routines – a change in the (given) number of customers has no impact on the qualitative nature of the long-term dynamical behaviour of a monopolist.

4 Duopoly

In analysing duopoly, our primary concern is whether the dynamical process involving the interaction between the firms’ adjustment routines and the consumers’ response to price differentials ensures the selection of the more efficient firm. The interesting question is whether this occurs if the firms differ *only* in respect of machine operating costs. Accordingly, we assume that the duopolists have a common target rate of return, τ; a common stock coefficient, γ; and common price and capital stock adjustment speeds, θ and κ.

We assume that initially the duopolists incur the same operating cost and are in equilibrium. At date 1, there is a (permanent) fall in the operating cost of firm 2. Firm 2’s going rate of return exceeds its target rate of return and it expands its stock of machines. The resulting build-up of product stocks prompts firm 2 to reduce price. Losing some customers to firm 2, firm 1 reacts to increased stocks by cutting its price and by not replenishing all its stock of machines. The dynamic process is underway. There are three possible long-term outcomes of this process. First, the high-cost firm 1 may be driven from the industry, that is, at some date t', $\omega^1_{t'} = 0$ implying $\omega^1_t = 0$ for $t > t'$, with the low-cost firm thereafter monopolising the market. Second, the low-cost firm 2 may be driven from the market. The final possibility is perpetual co-existence of the firms.

4.1 Firm survival

Unless otherwise stated, we assume throughout that $n = 10,000$, $a = 10$, $m = 10$, $\delta = 0.1$, $\gamma = 0.2$ and $\tau = 0.3$. Prior to date 1, the firms incur the same operating cost of 2 and are both in stationary equilibrium, charging a price of 6 and meeting the demands of their customers at that price. At date 1, the operating cost of firm 2 falls. To examine the impact of the cost reduction, we identify the state of the industry after the elapse of T periods. We use a ‘survival bitmap’ to show, for specified parameter combinations, whether the market process has selected one of the firms by the T^{th} period. For each parameter combination, the system is iterated for T periods. A white square indicates selection of the efficient firm; that is, only the low-cost firm 2 has survived to the T^{th} period,

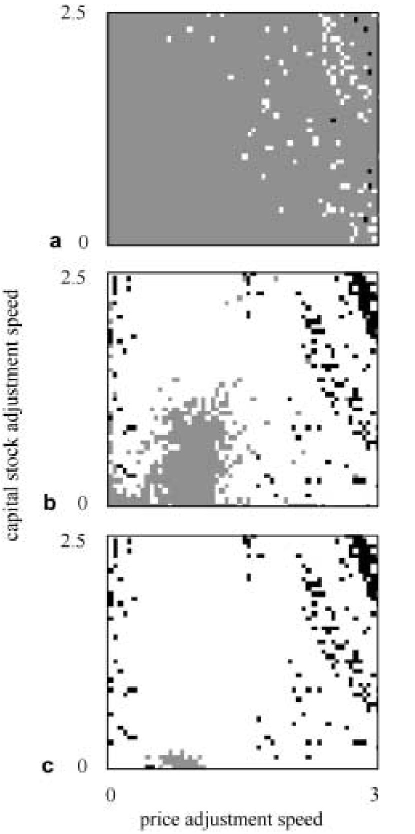

Fig. 6. **a** $T = 20$. **b** $T = 100$. **c** $T = 1,000$

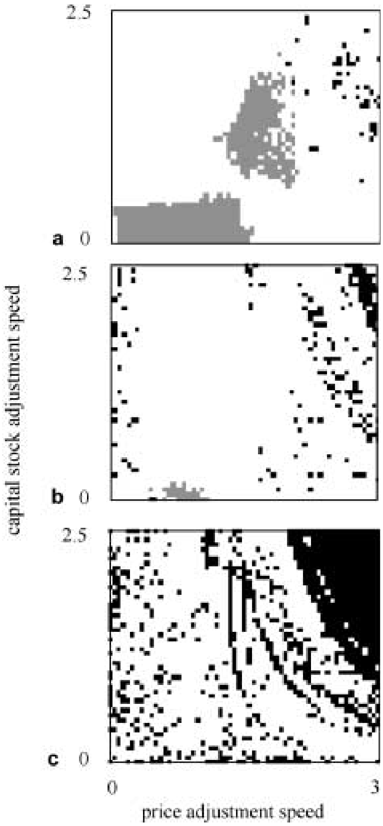

Fig. 7. **a** $\sigma = 1$. **b** $\sigma = 2$. **c** $\sigma = 3$

firm 1 having being driven from the industry by the loss of all its customers. In contrast, a black square signifies inefficient selection: only the high-cost firm 1 has survived. A grey square indicates that both firms are still operating in the T^{th} period. Figure 6a–c are survival bitmaps based on horizons $T = 20$, $T = 100$ and $T = 1,000$ respectively, and they serve to indicate the sorts of speeds at which the market process operates. Figure 6 assumes that the firms initially share the market equally; that the reduction in firm 2's operating cost is 0.2; and that the customer switching speed is $\sigma = 2$. In each diagram, θ ranges (in increments of 0.05) from 0.1 to 3 and κ ranges (in increments of 0.05) from 0.1 to 2.5. In the transition from (a) $T = 20$ to (b) $T = 100$ and to (c) $T = 1,000$, a white square must remain white; a black square must remain black; but a grey square may turn white, turn black or remain grey. For most combinations of θ and κ, both firms are still in operation after 20 periods. However, for most speed combinations, the market process has selected one of the firms by the 100^{th} period.[8] Henceforth,

[8] It is interesting to note that, for $\sigma = 2$, if both firms were to set and maintain prices equal to their supply prices from the outset, the resulting price differential of 0.2 would, from initially equal

in order to focus on 'long-term' outcomes, our survival bitmaps are based on a horizon of $T = 1,000$.[9]

The importance of the customer switching speed is demonstrated by the survival bitmaps in Figure 7a–c, which are based on $\sigma = 1$, $\sigma = 2$ and $\sigma = 3$, respectively, and otherwise on the same parameters as Figure 6. A negative but nevertheless important inference from Figures 6 and 7 is that – even though we have taken the simplest case of a homogenous product, consumers with identical linear demand curves, no form of increasing returns, no stochastic forces and firms that are identical except for a (non-trivial) cost difference – there is no guarantee that the low-cost firm will survive the dynamical process of firm adjustment and customer switching.

Figure 7a–c should be compared with the white regions in Figure 5a and b, that is, with the sets of adjustment speeds that would (in the absence of any customer switching) result in the firms converging on their respective fixed points. For convergent adjustment speeds, the instances of 'wrong' selection are rare for $\sigma = 1$ and $\sigma = 2$. However, for $\sigma = 3$, inefficient selection is more frequent. More generally, Figure 7 suggests that some degree of market inertia may be desirable from an efficiency perspective: as σ increases, the instances of inefficient selection increase. Thus the customer switching speed can be too rapid, at least in that the elimination of the low-cost firm may be more likely.

The survival prospects of the low-cost firm depend crucially on the *pre-history* of the industry as reflected in the initial market shares. Figure 8a–c are survival bitmaps, for a customer switching speed of $\sigma = 1$, based on initial market shares for the low-cost firm of $\omega_0^2 = 1/2$, $\omega_0^2 = 1/3$ and $\omega_0^2 = 1/5$, respectively.[10] As the initial share of the low-cost firm falls, the instances of inefficient selection increase significantly; in particular, the set of adjustment speeds that lead to inefficient selection encroaches on the set of speeds that would imply convergence in the absence of customer switching.[11]

Figure 9 is a survival bitmap, based on $\theta = 1$, $\kappa = 1$ and $\sigma = 2$, showing the state of the industry after the elapse of 1,000 periods contingent on firm 2's initial market share, ω_0^2, and on the size of its cost reduction, $\left(\nu^1 - \nu^2\right)$. Specifically, firm 2's initial market share is increased (in steps of 0.02) from 0.1 to 0.9 and its cost reduction is increased (in steps of 0.01) from 0.01 to 0.4. The lower the initial share of the low-cost firm 2, the greater its new cost advantage has to be for there to be a reliable prospect of it surviving the dynamic process initiated by the cost reduction.

---

market shares, lead to the elimination of the high-cost firm after 129 periods. That is, for most of the speed combinations, selection is more rapid when the firms employ our behavioural routines.

[9] Given that $n = 10,000$, it makes little difference whether one imposes on the simulations that the number of customers patronising any firm be an integer. In our simulations, the market shares are treated as continuous variables subject to no firm having less than one customer.

[10] For ease of comparison, Figure 8a repeats Figure 7a.

[11] Figure 8 confirms that re-entry would be very difficult for a firm that has been driven from the market.

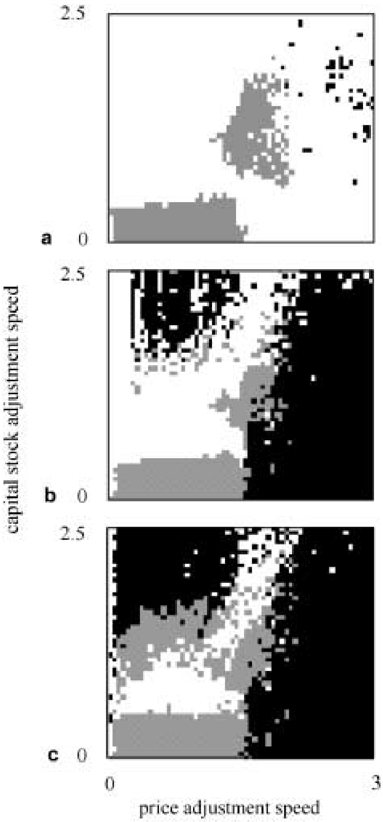

Fig. 8. a $\omega^2 = 1/2$. **b** $\omega^2 = 1/3$. **c** $\omega^2 = 1/5$

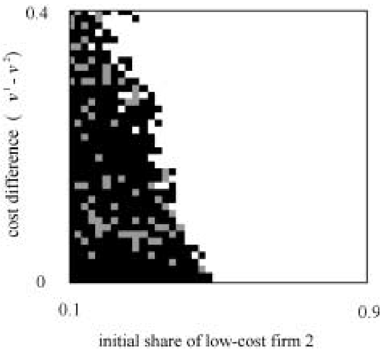

Fig. 9

4.2 Long-term co-existence

Hitherto we have used the term 'equilibrium' to denote a state-of-rest corresponding to a stationary fixed point for the dynamical system. Given a difference in operating cost, there are only two possible fixed points: one in which the low-cost firm monopolises the market and one in which the high-cost firm does so. That is, there cannot be a *stationary* equilibrium in which both firms co-exist, since, for there to be no consumer switching, the firms would have to set the same product price, whereas for each firm to be in equilibrium would require that its price equal its long-period supply price (and, following the cost change, their supply prices differ). But perpetual co-existence – necessarily involving cyclical or chaotic industry behaviour – is possible. For example, for $\omega_0^1 = \omega_0^2 = 1/2$, $\nu^1 = 2$, $\nu^2 = 1.8$, $\theta = 1.6$, $\kappa = 1.5$, $\delta = 0.05$ and $\sigma = 0.5$, both firms are still in operation after 40,000 time periods – by anyone's standards a 'long' horizon.[12] Figure 10a and b presents the time paths over the last 200 periods

[12] The depreciation rate and the customer switching speed are both lower than those assumed in Figures 6 to 9.

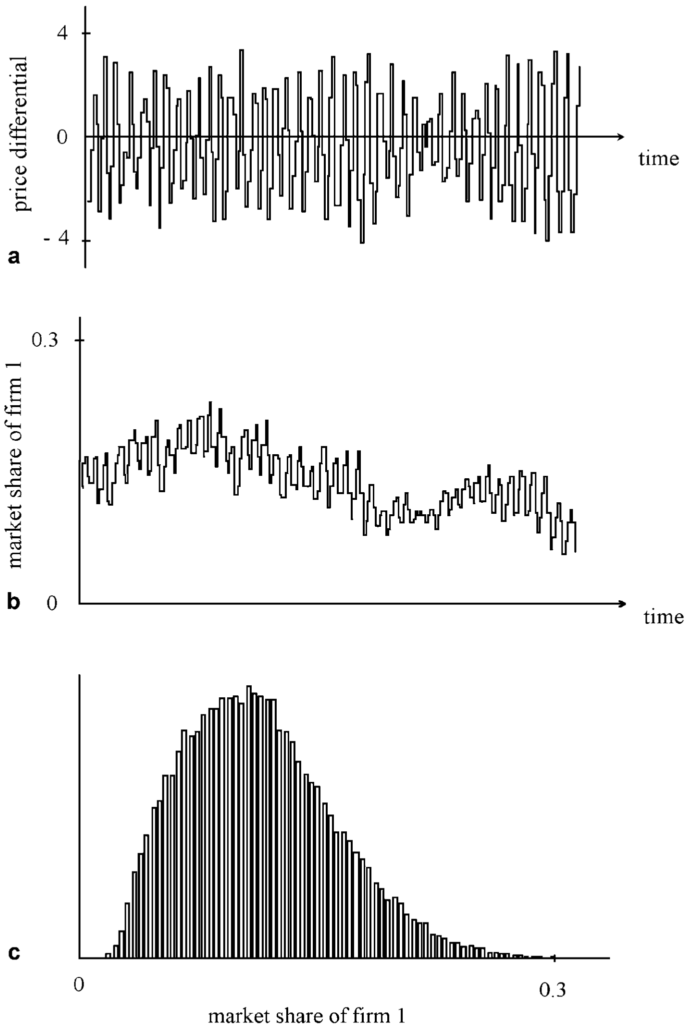

Fig. 10a–c

of that horizon of (a) the price differential $\left(p_t^1 - p_t^2\right)$ and (b) the market share of the high-cost firm 1. Figure 10c shows the relative frequencies of firm 1's market share over the entire horizon. Such perpetual co-existence constitutes a form of industry equilibrium, albeit implying a broader equilibrium notion than the conventional one (and a different equilibrium notion from that of statistical regularity, as invoked in models with recurrent random shocks). Continued co-existence, notwithstanding both a non-trivial cost difference and the absence of stochastic forces, derives here from non-linearities in the system.[13]

[13] Rothschild (1973) provides a useful survey of early attempts to develop models in which 'equilibrium' can involve perpetual price variability.

5 Robustness of the results

Our analysis of the dynamic behaviour of duopoly has suggested certain *ceteris paribus* tendencies. Specifically, the low-cost firm is more likely to survive (a) for adjustment speeds that would lead to convergence under monopoly; (b) for lower customer switching speeds; (c) for greater cost differentials; and (d) for higher initial market shares for the low-cost firm. We now consider, albeit briefly, whether modifications to the model would alter these conclusions.

5.1 Three firms

First consider very briefly the case where there are initially *three* firms in equilibrium incurring the same operating cost of 2 and sharing the market equally. At date 1, the cost of firm 2 falls to 1.8 and the cost of firm 3 rises to 2.2, initiating a dynamic process. Figure 11a–c are survival bitmaps based on customer switching speeds of $\sigma = 1$, $\sigma = 2$ and $\sigma = 3$, respectively. A white square indicates efficient selection: after 1,000 periods only the low-cost firm remains in the industry. A grey square indicates that, whilst the low-cost firm is still in operation, either or both of the other firms survive. A black square indicates that the low-cost firm has been driven from the industry. These survival bitmaps, which exhibit a similar pattern to those in Figure 7a–c, confirm the *ceteris paribus* tendencies for instances of inefficient selection to be more common for a higher customer switching speed and for firm adjustment speeds that would imply volatile prices in the absence of customer switching.

5.2 Stochastic customer switching

Returning to the case of duopoly, consider now the introduction of a stochastic element into the response of consumers to price differentials: at each date t, the customer switching speed is randomly and independently drawn from a uniform distribution between $\underline{\sigma}$ and $\overline{\sigma}$. Figure 12 assumes that σ_t is uniformly distributed between 1 and 3; otherwise it assumes the same parameters as Figure 7 in Section 4.1. Comparing Figure 12a to Figure 7b, for which σ is fixed at 2, the introduction of a stochastic element means that inefficient selection is more common at 'convergent' adjustment speeds. Comparing Figure 12b with Figure 9, at a given initial share for the low-cost firm, the cost reduction needed for there to be a reliable prospect of the low-cost surviving is marginally greater with a stochastic element to customer switching.

5.3 Minimum price

Recall that the pricing routine (6) embodies a rule-of-thumb that price at least cover machine operation and replacement costs. Consider possible alternative

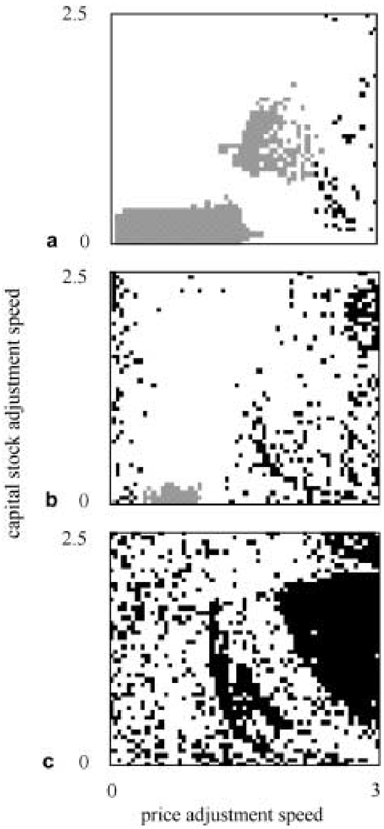

Fig. 11 a σ = 1. **b** σ = 2. **c** σ = 3

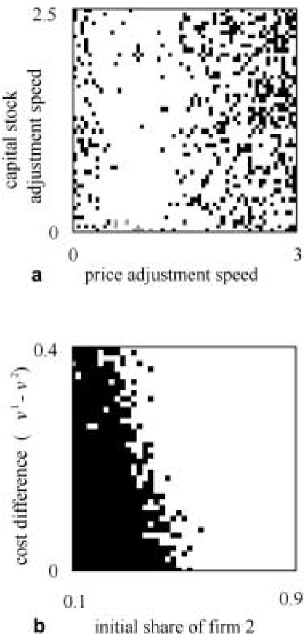

Fig. 12a,b

procedural rules. At one extreme, the constraint on the Pricing Department might be that it simply cover its operating costs, that is, $p_t^f \geq \nu^f$. At the other extreme, for a firm with a positive pure target rate of return, it might be charged with covering the full cost, that is, $p_t^f \geq \nu^f + (\delta + i)\, m$.

Altering the minimum price rule can impact significantly on the interaction between duopolists and thereby on the efficiency of the selection process. The possible effect on firm survival is shown in Figure 13. As in Section 4.1, we assume $n = 10{,}000$, $a = 10$, $m = 10$, $\delta = 0.1$, $\gamma = 0.2$, $\nu^1 = 2$, $\nu^2 = 1.8$, $\omega_0^1 = \omega_0^2 = 1/2$ and $\tau = 0.3$. In addition, we assume a market rate of interest of $i = 0.15$(so that the target pure rate of return is 0.15). Figure 13a–c are survival bitmaps for $\sigma = 1$, $\sigma = 2$ and $\sigma = 3$, respectively, for $p_t^f \geq \nu^f$; Figure 13d–f are the corresponding survival bitmaps for $p_t^f \geq \nu^f + (\delta + i)\, m$. The simulations confirm that, for a given procedural rule, increasing the customer switching speed increases the instances of inefficient selection.[14] More significantly but not unexpectedly, for a given customer switching speed, increasing the minimum price

[14] Recall that Figure 7 shows the corresponding survival bitmaps for $p_t^f \geq \nu^f + \delta m$.

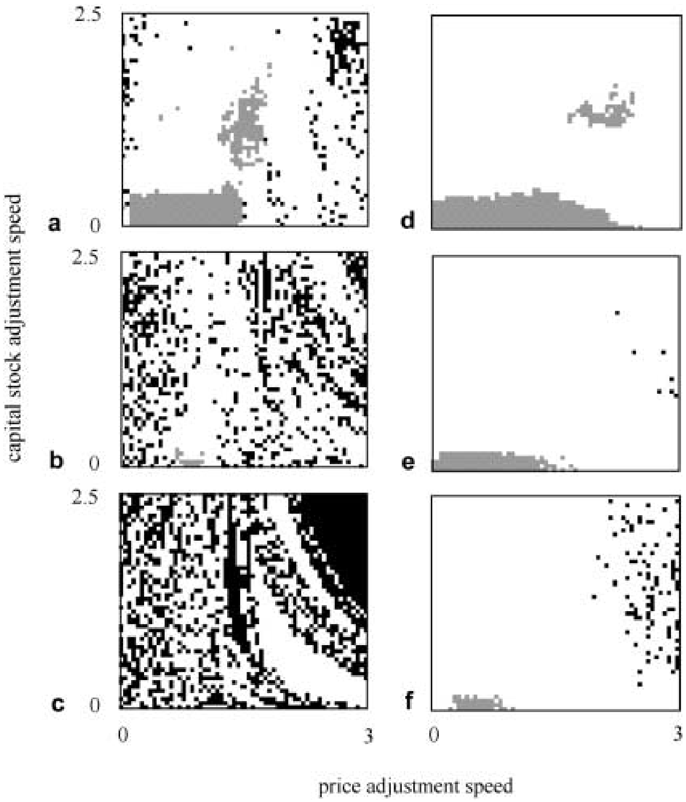

Fig. 13a–f

reduces the potential range of price variation for an individual firm; it reduces the potential magnitude of price differentials; and it thereby reduces the instances of inefficient selection.

5.4 Pricing and perceived excess demand

For behavioural models, it is particularly important to consider whether the conclusions are sensitive to the specifications of the participants' routines. A stern test of the robustness of our conclusions is provided by postulating a very different pricing routine, one involving the firm's perception of the excess demand that it faces. Suppose then that firm f raises (lowers) price at date t if the demand that it faced over the previous period exceeds (is less than) the quantity that it will have available for sale over the ensuing period. Specifically:

$$\frac{p_t^f - p_{t-1}^f}{p_{t-1}^f} = \varphi^f \left(\frac{d_{t-1}^f - s_t^f - k_t^f}{s_t^f + k_t^f} \right) \tag{19}$$

where $\varphi^f > 0$ is the firm's price adjustment speed and where the constraint on price is $p_t^f \geq \nu^f + \delta m$. In contrast to (6), this routine presupposes that, even

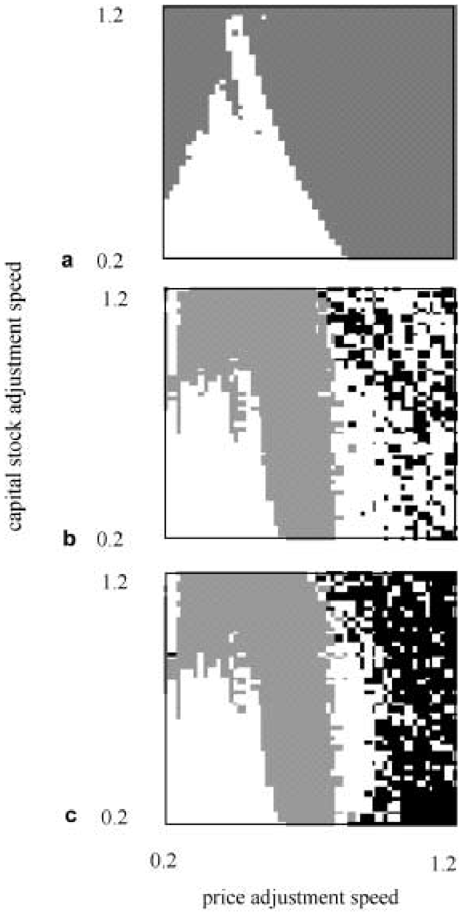

Fig. 14. a $\nu = 2$. **b** $\omega_0^2 = 1/2$. **c** $\omega_0^2 = 1/3$

if the firm was unable to meet all the demand in the previous period, it knows what it could have sold. Assume that the production routine involves the full use of available capacity and that the investment routine, expressed here in terms of the difference between the firm's current price and its supply price as defined by (11), is:

$$\frac{k_{t+1}^f - k_t^f}{k_t^f} = \mu^f \left(p_t^f - p_s^f \right) \tag{20}$$

subject to $k_{t+1}^f \geq (1 - \delta) k_t^f$, where $\mu^f > 0$ is the capital stock adjustment speed. The fixed-point price and capacity for a monopolist using (19) and (20) are the same as those depicted in Figure 1.

Figure 14a is a convergence bitmap for a monopolist with an operating cost of 2. Figure 14b is a survival bitmap for duopoly. The firms initially face the same operating cost of 2 and share the market equally; at date 1, the operating cost of firm 2 falls to 1.8. Given an assumed customer switching speed of $\sigma = 2$, the instances of inefficient selection over the course of 1,000 periods are confined to speed combinations that would not result in convergence to the fixed point under monopoly. Figure 14c involves the same customer switching speed, but assumes that the initial share of the low-cost firm 2 is 1/3. A comparison of the difference between Figure 14b and c with the difference between Figure 8a and 8b suggests rather less sensitivity to the initial shares with pricing routine (19) than with pricing routine (6).

5.5 Fixed-price routine

A common-place observation is that firms with market power often set prices on the basis of a fixed mark-up. Suppose then that firm f's (degenerate) pricing routine is that, as long as it remains in the industry, it sets and maintains a price equal to its long-period supply price:

$$p_t^f = \nu^f + \left(\delta + \tau^f\right) m \tag{21}$$

In this case, the investment routine can only play a passive role of adjusting capacity to demand at that price. For example, firm f might adjust its capital stock in the light of perceived excess demand, $\left(d_{t-1}^f - s_t^f - k_t^f\right)$.

Consider the scenario where two firms initially incur the same operating cost of 2 and share the market equally and suppose that firm 2 experiences a cost reduction. If *both* firms were to use the fixed-price routine (21), the high-cost firm would inexorably be driven from the market by the process of customer switching. But suppose that one firm uses the fixed-price routine (21), whereas the other employs routines (6), (7) and (9). Would the elimination of the high-cost firm still be inevitable? If the high-cost firm were to maintain a fixed price and the low-cost firm were to use routines (6), (7) and (9), the process of customer switching would result in efficient selection, that is, the high-cost firm 1 would be driven out of the market. In contrast, if the low-cost firm were the one to use the fixed-price routine, the outcome would depend on firm 1's price and capital

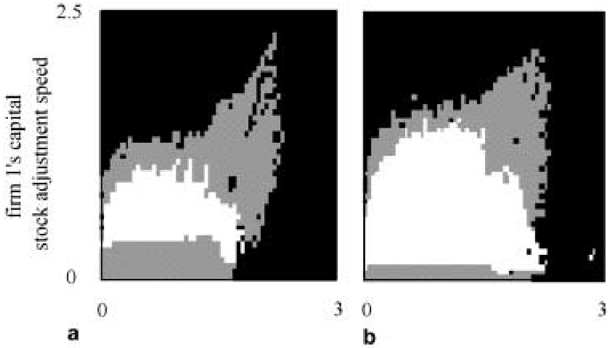

Fig. 15a,b. Firm 1's price adjustment speed. **a** $\nu^1 - \nu^2 = 0.2$. **b** $\nu^1 - \nu^2 = 0.4$

stock adjustment speeds. Figure 15a and b are survival bitmaps for combinations of firm 1's adjustment speeds for a cost reduction to firm 2 of (a) 0.2 and (b) 0.4, respectively, for a customer switching speed of $\sigma = 1$. For adjustment speeds that (under monopoly) would imply convergence for firm 1, the high-cost firm 1 would typically be driven from the market by the process of customer switching. However, for adjustment speeds that (under monopoly) would imply *volatility* in the price of firm 1, it is typically the low-cost firm 2 that is driven from the market by the process of customer switching. Thus, when confronted by a more efficient rival setting a fixed price, a firm's likelihood of survival may be enhanced by the use of a pricing routine, such as (6), which implies variability in the firm's own price.

This brief examination of the use by one firm of a fixed-price routine versus the use by the other of a routine in which price responds to stock movements merely hints at the possible insights from a systematic analysis of competition between different routines. But the latter is beyond the scope of the present paper.

6 Some concluding comments

In this paper, we have explored the dynamical behaviour of an industry on the basis of two central premises. The first is that, in reality, firms with some degree of market power frequently employ relatively simple behavioural routines: lacking information about the concurrent decisions of rivals and about how consumers would react to price differentials, our firms employ algorithmic rules based on internally available information. The second premise is that customers switch between firms in response to price differentials but do not all do so instantaneously. The frequent assumption, as in the simple Bertrand model, that all consumers always patronise the lowest price firm is at variance with reality and it cannot provide an adequate basis for analysing processes of market selection and firm survival under oligopolistic conditions.

The speeds at which firms adjust to changing circumstances and at which customers respond to price differentials matter for the efficiency of the process of market selection. There is no guarantee that the low-cost duopolist will survive, even for the simplest case where rival firms differ only in their operating costs. It should not be inferred, however, that the conclusion from our analysis is that ‘anything might happen’. On the contrary, the simulations suggest certain *ceteris paribus* tendencies, ones which appear to be robust with respect to the specifications of the firms' routines. First, the prospects of efficient selection are greater for price and capital stock adjustment speeds that would imply convergence for the individual firm in the absence of customer switching. Second, survival of the low-cost firm may be jeopardised by ‘rapid’ customer switching, that is, a little ‘grit’ in the system may have beneficial consequences. These conclusions are certainly consonant with economic intuition (and with common-sense). And yet many studies purport to establish convergence to market equilibria without any explicit regard to the speeds at which participants react to changes in their circumstances.

References

Chen Y, Rosenthal RW (1996) Dynamic duopoly with slowly changing customer loyalties. International Journal of Industrial Organization 14: 269–296

Cyert RM, March JG, (1963) A behavioural theory of the firm. Prentice-Hall, New Jersey

Metcalfe JS (1998) Evolutionary economics and creative destruction. Routledge, London

Phelps ES, Winter S (1970) Optimal price policy under atomistic competition. In: Phelps ES (ed) Microeconomic foundations of employment and inflation theory. Norton, New York

Puu T (1998) The chaotic duopolists revisited. Journal of Economic Behavior and Organization 33: 385–394

Rosenthal RW (1982) A dynamic model of duopoly with customer loyalties. Journal of Economic Theory 27: 69–76

Rothschild M (1973) Models of market organization with imperfect information: a survey. Journal of Political Economy 81: 1283–1308

Winston GC (1982) The timing of economic activities. Cambridge University Press, Cambridge

Winter SG (1964) Economic "natural selection" and the theory of the firm. Yale Economic Essays 4: 225–272

Winter SG (1971) Satisficing, selection and the innovating remnant. Quarterly Journal of Economics 2: 237–261