

# Almost-maximization as a behavioral theory of the firm: from the static, dynamic and evolutionary perspectives.

Huw Dixon<sup>1</sup>

## Abstract

In this paper we consider the approach of epsilon maximization to how the firm behaves. In particular we focus on the dynamic behaviour of firms using the example of price-setting. We show how almost rational firms can be more volatile in their behaviour. However, if a lexicographic preference for simplicity is made, then we can explain nominal rigidity as a result of epsilon optimisation. The behaviour of the firm which is consistent with its long-term survival is examined and we argue that epsilon optimisation is consistent with survival in any context in which something is optimised (such as sales revenue).

Key words: bounded rationality, dynamic choice, inertia volatility.

JEL codes: B52, C61, D21

---

<sup>1</sup> Cardiff Business School, [dixonh@cardiff.ac.uk](mailto:dixonh@cardiff.ac.uk).

The idea of bounded rationality became popular in the economics of the firm with the publication in 1963 of the book *A Behavioral Theory of the Firm* by Richard M. Cyert and James G. March<sup>2</sup>. Whilst the dominant model of rationality has remained constrained-optimisation in most fields of economics, there has been a gradual spread of alternatives. Today behavioral economics is an important component of economics across a wide range of fields. In this paper, I seek to adopt a primarily methodological approach asking how we should understand near rationality in our modelling of firms.

Firms represent a range of challenges for economics. These can be simplified into perhaps three main categories:

Firstly, what are the objectives of firms (for example profits, sales, managerial utility);

Secondly, as organizations, do firms seek to maximize in the pursuit of their goals, or are they boundedly rational in some sense (such as satisficing);

Thirdly, how does our conception of the *behavior* of the firm fit with the explanation for the *existence or survival* of firms - are firms efficient institutions in terms of transactions costs and markets in general.

---

<sup>2</sup> In the wider social sciences, the modelling of bounded rationality in Herbert Simon's *Administrative behaviour* published in 1947 had already begun to have an effect earlier. His 1982 book *Models of Rationality* represents a more mature summation of Simon's thinking.

A related way of thinking is in terms of the level of analysis: the key individuals within the firm (for example, senior management or owners), the interactions of these individuals within the firm and lastly the interaction of the firm with the wider economy.

These issues are all linked in practice. Firms are not isolated institutions: they compete both in the particular markets in which they operate (for example as oligopolists interacting with each other or incumbents competing with potential entrants) or at higher levels in capital markets where firms overall performance is characterised for example by share prices and related profit metrics. Really stupid firms cannot survive: they are driven out of business by more able competitors, or perhaps are taken over by a private equity firm for a remake and remodel. We can think of these external constraints in terms of Darwinian evolutionary forces, as in Armen Alchian (1950): “This is the criterion by which the economic system selects survivors: those who realize positive profits are the survivors; those who suffer losses disappear”. Thus, although John Hick’s (1935) thought that “the best monopoly profit is a quiet life”, there are limits to under-performance which make even the most sedentary manager subject to scrutiny. One of the main reasons Chester Barnard (1938) argued that firms do not last forever was that over time they cease to be effective.

Whilst external forces operate on firms, there is some degree of slack for most firms that creates some degree of “managerial discretion”. This is particularly true for owner-managed firms: so long it remains solvent, an owner managed firm can keep going. Indeed, some owners might even subsidise a loss-making firm from other revenues for reasons beyond economics.

In relation to the theory of the firm, the influential survey of Seth and Thomas (1994) laid out different approaches to modelling firms across the disciplines of economics and management. It is worth quoting at length their thoughts on fully rational and behavioural theories in economics:

*We note that all versions of rationality discussed above have in common that economic agents are viewed as purposeful and intelligent, and assumed to follow reasonable and logical procedures in making decisions. Some version of rationality underlies all economic explanation, to allow prediction of the relevant outcomes of the decision-making process: if economic agents are permitted arbitrary behaviour, the outcomes of their actions are necessarily indeterminate. Rather, the essential difference lies in the types of decision-rules they are assumed to use: maximizing, satisficing, or habit<sup>3</sup>.*

---

<sup>3</sup> Seth and Thomas (1994), page 175.

We see the challenge is to provide decision making processes (summarised as heuristic decision-rules) that arise from within the firm in the pursuit of its goals and which enable the firm to survive in its competitive environment.

In Dixon (1992) I argued for the concept of near-optimisation (epsilon-optimisation). The argument is that in order for a decision rule to survive in the long-run, it must be almost optimal. Too far away, and it will fail the “survival test”. Close enough, and it falls within the range of the “satisfactory, good enough”, thus enabling managerial discretion. Now, this general approach leaves open the first issue of behavioural theory, namely what is almost optimised? Managerial utility, sales, or profits

The plan of this paper is first to set out the pure theory of epsilon optimisation and focus on its link to inertia or the tendency for agents to choose almost optimal strategies that involve holding actions constant over time. This is a very important aspect for the behavioural theory of the firm: in practice, we often observe firms setting the same price for prolonged periods of time<sup>4</sup>. George Akerlof and Janet Yellen put forward the argument in 1985 that near-rational firms would display nominal rigidity and rather than vary their price to track the optimum perfectly. An alternative put forward by

---

<sup>4</sup> The famous example is the 50 cent bottle of coke that lasted over 60 years. However, the price-quote data used to construct CPI data in many countries has revealed a long fat tail of prices that change very infrequently. See Klenow and Malin (2010) for a survey of the evidence.

Gregory Mankiw also in 1985 was that menu-costs would lead to perfectly rational firms setting the same price for two or more periods. To what extent is a boundedly rational firm the same as a perfectly rational firm with a decision cost? I argue that bounded rationality can lead to very different outcomes to the case of perfect rationality with menu costs in a dynamic setting. Bounded rationality does not in itself lead to inertia: there are many almost rational strategies in dynamic settings, only a subset of which involve choosing the same action over time. To get inertia from bounded rationality you need to add a preference for simplicity of strategy, preferring repetition to variation.

The second part of the paper will look at the issue of what is nearly maximized. Long-run survival of firms depends on them earning enough profits (normal risk adjusted profits). Only certain forms of behaviour, in terms of “decision rules” will be consistent with this survival criterion. Markets and firm environments differ, in terms of the number of active incumbents, the extent of product differentiation (or ability to create the differentiation through branding), the ease of entry and exit and so on. Profit maximization as a form of behaviour may be inconsistent with long-run survival, and other forms of behaviour (decision rules) may emerge: from sales maximisation to

cooperation (joint profit maximisation). These can then be the objectives to be almost optimised.

The main lesson of this paper is that studying bounded rationality of the firm is about more than looking just looking at decision-making processes either in the head of an individual or group of individuals. The processes within the firm are also determined by the wider environment of the firm: the market(s) it operates in, its interactions with competitors and the wider economy through capital markets. We can see the “decision rules” that emerge as resulting from the interaction of all three levels: the individual, the collective and the wider economy.

The plan of the paper is as follows. Section 1 outlines the basic theory of approximate optimisation and its special case of perfect optimisation. Section 2 looks at almost optimisation in the context of dynamic decision making, with particular focus on whether bounded rationality generates inertia in behaviour. Section 3 goes on to look at the survival of the firm in the long-run and how this constrains the behaviour of firms. Section 4 concludes.

## 1.. Almost optimality.

Perfect rationality<sup>5</sup> is the approach adopted by most economics for over a century. A stereotypical perfectly rational agent has the ability to calculate the answer to any well-defined problem. It will have a well-defined objective function that enables the best solution to be identified from a given range of possibilities. An objective function should be able to rank all possible outcomes in a way that is transitive (or at least acyclic). Choice is subject to some constraint: for example, a budget constraint or production function. Uncertainty can be introduced if there is a known probability distribution and the objective function satisfies the Von Neumann-Morgenstern properties. In its simplest form, we can think of agent utility being defined over a vector  $x$ ,  $U : \mathfrak{R}^n \rightarrow \mathfrak{R}$ , where  $x$  is chosen from some compact convex set  $S$ . The agent then solves:

$$\text{Max } U(x) \text{ subject to } x \in S \subset \mathfrak{R}^n$$

The set  $S$  may itself be determined by some parameters (prices and income in the case of the budget set). The solution to the problem is a choice of action  $x^*$  and resultant utility  $U^*$  which can be both seen as depending on  $S$ , In the

---

<sup>5</sup> This is often called the “neoclassical” model when applied to the firm.

case of the consumer, the budget set is determined by prices and income: the indirect utility gives the maximum utility given the prices and income. The Marshallian demand is the optimal choice of consumption given prices and income.

Now, let us see under what external circumstances this sort of behaviour might arise in a firm. Perfect competition is the market structure that would lead to this sort of firm behaviour. With free zero cost entry, (supernormal) profits will be zero in the long-run. Firms that do not minimize costs and maximize profits will not survive against maximizers. In the textbook scenario of perfect competition, profit maximization is required for survival (except perhaps in the case of an owner manager who is willing to subsidise the firm out of other income). Since profit maximization is required for survival, there is no discretion for managers and bounded rationality is not possible.

### 1.2 Almost correct choices: trembling hands.

We can think of almost correct choices in terms of two metrics: the closeness of the *payoff* relative to the optimal payoff, or alternatively the closeness of the *action* to the optimal action. Under certain assumptions, these two metrics are equivalent, but need not be.

Consider first the case of an action. The idea here is analogous to the “trembling hand” of equilibrium refinements in Game Theory first introduced by Reinhardt Selten (1975). The agent tries to implement the optimal action  $x^*$ , but by a mistake in execution chooses some other action. We might want to say that the mistake in action is “close” to the optimal action. That means that with an appropriate metric  $M^6$ , the chosen action  $x'$  is within a distance  $\kappa$  of  $x^*$ :

$$\kappa > M(x', x^*)$$

The fact that the error might be “small” when measured in the metric  $M$  need not imply that the loss in *utility* is small. For example, in an Olympic final, a small error can give rise to a huge difference in utility (Bronze instead of Gold). However, a small error giving rise to a big loss can only occur if there is a *discontinuity* in the objective function. If we make the standard assumption that the payoff function is continuously differentiable for at least the first two derivatives<sup>7</sup>, then we can ensure that small mistakes in implementing the strategy will give rise to small losses in payoff. Put simply, with if  $U$  is continuously differentiable, then for any  $\varepsilon > 0$ , there exists  $\kappa > 0$  such that: if

---

<sup>6</sup> For example, if  $X$  is an  $n$  dimensional vector, a possible metric would be:  $M(x', x^*) = \sum_{i=1}^n \|x'_i - x_i^*\|$ .

The distance between two points is the sum of the absolute differences across each of the  $n$  dimensions.

<sup>7</sup> usually written as  $U \in C^2$ .

$\kappa > M(x', x^*)$  then  $U(x^*) - U(x') < \varepsilon$ . We can go further and think of this as defining the distance  $\kappa$  as an (increasing) function of  $\varepsilon$ , i.e.  $\kappa = \kappa(\varepsilon)$ . That is, if we want our loss in terms of payoff to be smaller than a certain level, then we need our strategy to be sufficiently close to the optimal strategy.

Can we go in the other direction, and say that if the action is close in terms of payoff, then the action must be close to optimal? If the payoff is strictly concave then the answer is yes, since there can only be one local maximum which is the global maximum. If you have weak concavity or even some convexity, then there can be local maxima which are close to the global maximum in terms of payoff but a long way away in terms of action. Think of two hills separated by a valley, with hill A slightly taller than hill B. you can be almost as high as the summit of A by being close to the summit of A, or being across the valley near the top of summit B. The valley is the convexity. If we assume strict concavity, then there is in effect only one hill and no local maxima other than the global maximum.

To make this idea concrete, let us consider the simple example of a monopolist. The monopolist faces an inverse demand curve of the form:  $P = P(x)$  where price depends on output, with . The cost of producing output is given by the convex cost function  $C(x)$ . Payoff (in this case profits) is then:

$$U(x) = x.P(x) - C(x)$$

Assuming that Revenue is strictly concave in output and costs are weakly convex, the utility function is strictly concave and will have a unique maximum when  $x$  is chosen from compact convex set.<sup>8</sup> If an interior solution exists, it will satisfy the usual “marginal revenue equals marginal” cost condition.

$$\frac{dU(x^*)}{dx} = 0$$

$$P(x^*) \left[ 1 + \frac{x^*}{P(x^*)} \frac{dP(x^*)}{dx} \right] = \frac{dC(x^*)}{dx}$$

We can make a Taylor expansion around the optimum output to give us the profit from a mistaken output  $x'$ :

$$U(x') = U(x^*) + \frac{dU(x^*)}{dx} (x' - x^*) + \frac{d^2U(x^*)}{dx^2} \frac{(x' - x^*)^2}{2}$$

Since  $x^*$  is optimal, the second term drops out, leaving us with:

$$U(x^*) - U(x') = -\frac{d^2U(x^*)}{dx^2} \frac{(x' - x^*)^2}{2}$$

The left hand gives us the distance in payoff space: the right hand gives us the distance in action (output) space. We can see if we want to ensure that the

---

<sup>8</sup> For example, if there is a saturation output  $\bar{x} > 0$ , such that

$P(x) = 0$  for  $x \geq \bar{x}$ , then  $S = [0, \bar{x}]$  is compact and convex.

loss in payoff is less than  $\varepsilon > 0$ , then we need to ensure that output is close to the optimum:

$$U(x^*) - U(x') < \varepsilon \Rightarrow -\frac{d^2U(x^*)}{dx^2} \frac{(x' - x^*)^2}{2} < \varepsilon$$

Which holds only if:

$$|x' - x^*| < \left(\sqrt[3]{2\varepsilon}\right) \cdot \left(-\frac{d^2U(x^*)}{dx^2}\right)^{-1} = \kappa(\varepsilon)$$

We can think of the right hand side of the inequality as representing the  $\kappa$  mentioned before. In the appendix, we give an explicit form for a firm with linear demand and costs.

In our examples, we have measured distance in absolute terms.

However, much the same concepts apply if we express the distance in terms of payoff and output in terms of the *proportional* distance. In general terms we have the “proportional” version of the previous Taylor expansion:

$$\frac{U(x^*) - U(x')}{U(x^*)} = -\left[\frac{d^2U(x^*)}{dx^2} \frac{x^{*2}}{U(x^*)}\right] \frac{(x' - x^*)^2}{2x^{*2}}$$

This formulation allows us to relate the percentage deviation of profits to the percentage deviation of output squared from the optimal.

## 2. The implications of almost optimisation in dynamic models.

We can think of agents as almost maximizing. What are the implications of how firms and markets behave over time? This has been the focus of research in the past, notably by George Akerlof and Janet Yellen (1985a and 1985b) and John Conlisk (1996). There are two contrasting views.

*Bounded rationality increases inertia.* Rather than responding to all shocks or changes, agents will just respond to those that take it out its comfort zone or “Band of Inertia”.

*Bounded rationality increases volatility.* It is an additional source of noise.

The notion that bounded rationality might increase volatility arises from the following simple line of reasoning. Suppose that the optimal level of  $x$  is driven by some shock which we can assume is random:  $x_t^* = \bar{x} + \varepsilon_t$  for example. With full rationality, the choice would be optimal value: the volatility of  $x_t$  would simply be the volatility of the “shock”  $\varepsilon_t$ . With bounded rationality, there is another “shock” in the form of “trembles”  $\nu_t$  :

$$x_t = x_t^* + \nu_t = \bar{x} + \varepsilon_t + \nu_t$$

Assuming that the two error terms are uncorrelated, the variance of  $x_t$  will be equal to the sum of the variances of the real shock and the trembles.

The story behind the bounded rationality increasing volatility rests on the idea that the decision maker is trying to hit the target each period, but most of the time will miss it due to faulty execution. We can think of faulty execution either as being that the agent knows the optimal value but cannot quite hit it, or that it simply misperceives the optimal value (for example, it cannot calculate the optimal value accurately in real time) or indeed a mixture. An inexperienced archer knows where the bulls-eye is, but due to lack of skill the arrows are spread around the bulls-eye. A card player might not be able to calculate the odds accurately and so not know the optimal play is at each stage in the card game.

The notion that bounded rationality leads to *inertia* rests on the idea that the agent is deciding whether to *change* its action from a previous value. The action will remain the same unless the agent sees a large enough *advantage* in changing it. More specifically, the agent might decide to change its action only if the gains exceed a certain threshold. This is the notion that underlies the menu-cost approach to nominal rigidity in macroeconomics. Under the standard interpretation, agents are fully rational but take into account the lump-sum costs of changing price. In a dynamic-stochastic

continuous time set up, agents choose when to change price and by how much in response to evolving shocks altering the optimal price. This sort of model cannot be solved analytically except for some special cases (usually the optimal price is Brownian motion or Wiener process, the continuous time equivalent of a random walk)<sup>9</sup>.

However, in their seminal 1985 paper, Goerge Akerlof and Janet Yellen put it in terms of bounded rationality:

“Near-rational behavior is nonmaximizing behavior in which the gains from maximizing rather than nonmaximizing are small in a well-defined sense. It is argued that in a wide class of models—those models in which objective functions are differentiable with respect to agents' own wages or prices—the cost of inertial money wage and price behavior as opposed to maximizing behavior, is small when a long-run equilibrium with full maximization has been perturbed by a shock. If wages and prices were initially at an optimum, the loss from failure to adjust them will be smaller, by an order of magnitude, than the shock.”<sup>10</sup>

In a static framework the argument is very much as in the Figure 1 “Hill” diagram, that is understandable in terms of a simple one dimensional metric in

---

<sup>9</sup> To see a recent example of this approach, see Alvarez et al (2014). The basic technique dates back to Sheshinski and Weiss (1977).

<sup>10</sup> Akerlof and Yellen 1985b, pages 825-4.

terms of choice variable and payoff. On the vertical axis we have the payoff, on the horizontal the action (price).

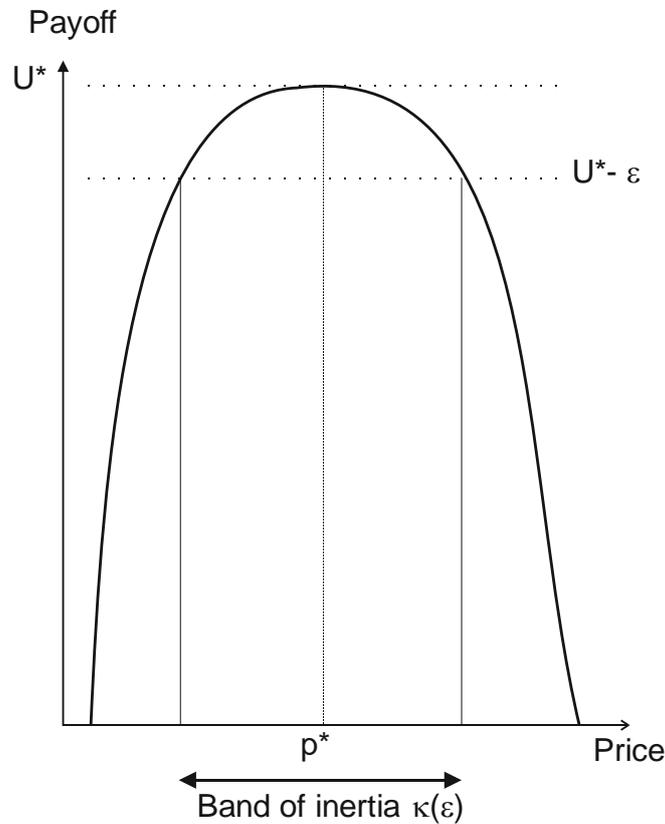


Figure 1: Epsilon optimization and the Band of Inertia

The payoff function is some concave function of price. The optimal price  $p^*$  gets you to the top of the payoff hill, yielding payoff  $U^*$ . However, any price in the “band of inertia” (size  $\kappa(\varepsilon)$ ) will get you to within  $\varepsilon$  of the optimum.

## 2.1 Dynamic models with no uncertainty.

In a purely static setting, the idea of menu costs and bounded rationality look pretty similar: you can think of the  $\varepsilon$  as a lump sum cost with perfect

optimization or as an imperfect optimization. In a dynamic setting, matters are more complicated. First we need to define bounded rationality across time: given an appropriate choice of metric to measure the distance between different strategies across time, bounded rationality would mean a choice of strategy that yields a payoff close to the optimum in terms of appropriately discounted payoff<sup>11</sup>. We can then compare this to the perfectly rational outcome with or without menu costs.

First we will consider a simple two period dynamic problem with no discounting and no uncertainty. The payoff depends on a shift variable  $e$  (cost or demand) and there is a menu cost  $\gamma \geq 0$  to pay if the choice of  $x$  differs between the two periods. The fully rational firm chooses its choice  $(x_1, x_2)$  to solve the maximization problem:

$$\max_{x_1, x_2} U(x_1, e_1) + U(x_2, e_2) - \gamma(x_1 - x_2)$$

Where  $\gamma$  is zero if  $x_1 = x_2$ , or is a positive constant otherwise. We can first define the optimal flexible action for each period (the case of  $\gamma = 0$ ), simply derived from the first order conditions given the shift variable, and the corresponding payoffs:

---

<sup>11</sup> Of course, bounded rationality might take the form of inappropriate discounting. For example, there exists plenty of evidence that humans have hyperbolic discounting and/or myopic planning horizons. This is a topic we do not have space to explore in this paper.

$$x^* = x(e)$$

$$U^* = U(e)$$

If we ignore menu costs, we can define our metric in terms of a loss function, giving the lost payoff for any choice of  $(x_1, x_2)$  and the optimum:

$$L(x_1, x_2) = U(x_1, e_1) + U(x_2, e_2) - U^*(e_1) + U^*(e_2)$$

In Figure 2 we depict the strategy space for the choice of  $(x_1, x_2)$ , with the optimal flexible value being  $x^*$ . We depict the set of pairs with a loss less than or equal to  $\kappa$  as the shaded circle. The 45-degree line is then combinations where the action is unchanged  $x_1 = x_2$ .

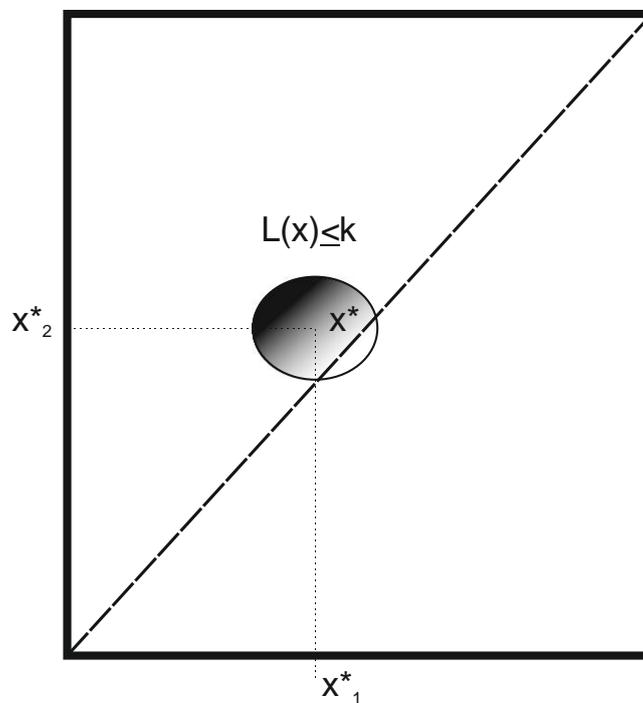


Figure 2: Two period choice.

As depicted, the 45% line intersects the shaded circle, so that there are constant price strategies that involve a small enough loss. For all of the other elements inside the circle, the action is different. Of course, there may be no intersection of the set  $L(x) \leq \kappa$  with the 45% line: if the values of  $e$  are sufficiently different, then there will be no “fix price” with loss less than or equal to  $\kappa$ . However, let us focus on cases where the 45% intersects the shaded area.

From figure 2, we can see immediately that there if there is a menu-cost  $\gamma = \kappa$ , then the optimal choice of actions across the two periods will be from the subset of price pairs on the 45% line, which will earn strictly higher payoffs than all of the other pairs in the circle (including the optimal flexible solution). Choosing from the 45% line means that no menu cost is incurred: in the rest of the shaded area the menu cost has to be paid since prices differ in each period.

From within the set of prices which are unchanged, there will be an optimal price which maximizes the profits subject to the constraint that the price is fixed. In effect, this solves the optimization:

$$\text{Max}_x \quad U(x, e_1) + U(x, e_2)$$

We can depict this in Figure 3, zooming in on the set  $L(x) \leq \kappa = \gamma$ .

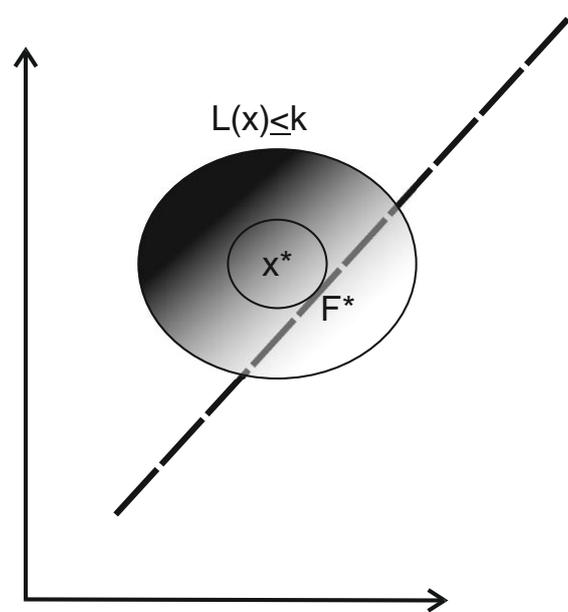


Figure 3: The optimal choice with menu costs

The optimum fix-price action is where the 45-degree line is tangential to the iso-payoff circle at point  $F^*$ . This gives the maximum payoff from amongst the set of fixed actions. This is the choice of the perfectly rational firm with a menu cost: it will choose the unique optimum  $F^*$ .

We can see how the optimizing firm behaves with or without menu-costs.

In this case, the set  $L(x) \leq \kappa$  represents the set of almost rational choices.

Within this set, there are two possibilities:

1. The near rational firms may choose any near rational combination in the set  $L(x) \leq \kappa$ , most of which involve choosing different actions in each period.

2. The near rational firms may choose combination in the set  $L(x) \leq \kappa$ , but *prefer* to choose simple strategies in which actions are the same in each period.

In order to get outcome 2, we need not only near-rationality, but also a *lexicographic preference* for simple strategies involving constant actions over time. We can say that from a set of strategies that are good enough, the agent will prefer a strategy that involves the same action in each period. This can be seen as a preference for simplicity. We can think of an infinitesimal “*menu cost*”, of a tiny magnitude for changing price. The near-rational strategies that are good enough that involve a change are equal in payoff: the strategies along the 45% line that are in the set are a tiny little bit better and will be chosen in preference.

If we compare the bounded rationality outcomes 1 and 2 to the fully rational menu-cost outcome, we can see that in a dynamic setting perfect rationality and bounded rationality are quite different. There is no inherent reason for the boundedly rational firm to choose the same action in each period. The action might increase or decrease across time (depending which side of the 45% line it is), even when the optimal flexible action shows a particular pattern. In Figure 3, the optimal flexible action involves increasing  $x$  (since  $x^*$  lies above the 45% line), but there are all of the pairs in the shaded

region below the 45% line which go the opposite way. It is only with option 2 and a preference for simple strategies that we will see a fixed price across the two periods. The “band of inertia” is then equivalent to the intersection of the 45% line and the shaded circle: the menu-cost optimiser is the single point  $F^*$  on that line.

## 2.2 Dynamic models with uncertainty.

Most dynamic menu-cost models are models with uncertainty. It is common to assume that the optimal price follows a random walk<sup>12</sup>:

$$x_{t+1}^* = x_t^* + e_t \quad (1.1)$$

where  $e_t$  is a white noise error<sup>13</sup>. Each period, a new shock is realised and the firm updates its plans. This contrasts with the previous analysis where we assumed that the firms knew the value the value of  $e_t$  in both periods (effectively there was perfect foresight).

From (1.1) the expected value of all future errors is zero and that the optimal price is expected to be the same in all future periods:

---

<sup>12</sup> This is the discrete time version of the continuous time model used (see for example Alvarez et al 2014 or Shesinski and Weiss 1978).

<sup>13</sup> Technically,  $e_t$  has an expected value of zero and is not serially correlated.

$$E_t [x_{t+i}^*] = E_t [x_t^*]$$

That is, in terms of our simple two period diagram, the *optimal* actions lie on the 45% line  $x_1^* = x_2^*$ , the optimal strategy involves *planning* to set the same price in both periods 1 and 2. Thus the optimal fix-price strategy and the optimal flex-price strategy coincide – in terms of figure 3,  $F^*$  and  $x^*$  coincide. This is depicted in Figure 4: the 45% line now lies in the middle of the shaded zone, with the optimal strategy lying on the 45% line.

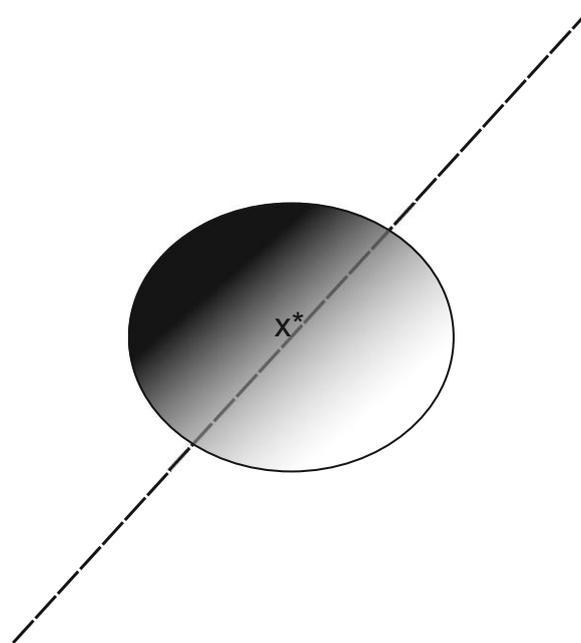


Figure 4: two period choice when the optimal price follows a random walk.

Now, let us for now assume that the firm chooses the optimal strategy in period 1. When period 2 arrives, the second period shock  $e_2$  is realized and in

$(x_1, x_2)$  space, this is a vertical shift equal to the shock  $e_2$ . Since the period 1 payoff is already a given, in period 2 the agent has the simple choice whether to leave the price at its planned value, or to change it to the optimal price given the realised shock. The problem then is exactly as we had in the one period case: if the loss in payoff is sufficiently small, then the nearly rational firm will not change the price.

$$U(x^*(e_2)) - U(x^*(0)) < \kappa$$

As in the one period case, the difference between a menu cost and bounded rationality is not clear – they look the same.

If we extend the number of periods to some finite  $T$ , we can work backwards. At any moment  $t=1\dots T$ , the firm will expect the current optimal price to extend for all of the remaining periods. In the case of a fully rational firm without menu costs, the firm will plan to set its future prices equal to the current optimum,  $x_{t+i}^P = E_t x_{t+i}^* = x_t^*$ , where  $x_{t+i}^P$  is the planned price in period  $t+i$ . Of course, the actual path of prices will follow a random walk, as the optimal price varies with the realisation of each shock.

With a menu-cost, the perfectly rational firm will consider at any time after the first period  $t=2\dots T$  whether to leave its price where it is (the price the previous period  $t-1$ ) or switch to the current optimal price. Without discounting, this is a very simple problem. For menu cost  $\gamma$ , the firm will

compare the profits it will earn over the remaining  $T-t+1$  periods from sticking with the old price (now it knows the current shock  $e_t$ ), or incurring the one-off menu cost and switching to the current optimal price<sup>14</sup>. Since the expected value of future shocks is zero, this takes the simple form:

$$U(x_t^*, e_t) - U(x_{t-1}, e_t) < \frac{\gamma}{T-t+1}$$

For each period there is a band of inertia around the optimum, which becomes larger over time. You are comparing the one-off menu cost to the stream of losses over  $T-t+1$  periods: you will choose not to change your current price only if the current loss times the number of remaining periods is less than the menu-cost. At time  $t=2\dots T$ , you know the current shock  $e_t$ , which you expect to remain in place for the remaining  $T-t+1$  periods. You can stay put, incur no menu cost and then the per-period loss is the difference between the optimal (with no menu costs) and what you earn at your current price. You are then considering this per period loss times the remaining number of periods to the one-off menu cost. The earlier you are, the more likely you are to change: in period 2, the current loss is multiplied by  $T-2$ , in the last period there is just the current loss to consider.

---

<sup>14</sup> The following is a slightly simplified argument, which is only completely accurate if the payoff function is quadratic so that certainty equivalence holds.

Hence there is a clear prediction that with a finite time horizon, the band of inertia shrinks as you get closer to the final period  $T$ . Existing models of menu costs do not have this feature, because they assume an essentially stationary problem with an infinite time horizon. It might be argued that a finite time horizon is arbitrary. However, in the case of price-setting, there is clear evidence of time-dependence in pricing – there are regular cycles of price-setting that we can think of as opportunities to change price for free. Menu costs are incurred when you change price out of the regular cycle.  $T$  would then be the length of the regular cycle. The probability of changing price would be highest near the start of the cycle, because the costs of getting it wrong would cumulate for longer. Near the end of the cycle you know you can change your price soon, so there is less to worry about<sup>15</sup>.

Turning next to the epsilon optimising near rational firm without menu costs, it will here make a big difference if we define the payoff metric in terms of absolute or relative differences in payoffs. The reason that the band of inertia grew with menu costs was that the absolute size of the menu costs remained the same, so it got larger relative to the absolute size of the loss as  $t$  increased (and hence  $T-t$  decreased). If the epsilon is a proportion of optimal profits, then it varies in absolute size with profits. Let us take the case of an

---

<sup>15</sup> In fact the phenomenon of a decreasing hazard rate is a robust empirical finding that this can explain.

absolute epsilon. At time  $t$ , we then have an almost optimal sequence of planned prices  $(x_t, x_{t+1} \dots x_T)$  such that

$$(T - t + 1)U(x_t^*, e_t) - \sum_{i=0}^{T-t} U(x_{t+i}, e_t) \leq \varepsilon$$

As in the case with certainty, the almost rational price plans may bounce around. We need to add a lexicographic preference from simple strategies that involve prices being fixed, in which case we restrict our attention to plans with  $x_{t+i} = x_t$ , so that:

$$(T - t + 1)(U(x_t^*, e_t) - U(x_t, e_t)) \leq \varepsilon$$

This is almost the same as the menu cost equation above, except that here the current price  $x_t$  is chosen rather than the inherited  $x_{t-1}$ . However, invoking again the lexicographic preference for simplicity our firm will prefer to set  $x_t = x_{t-1}$  rather than any other almost rational price. In period 1, when the menu cost optimiser gets to freely choose its opening price it will set  $x_1 = x_1^*$  whereas the almost optimiser will have a band of choices around  $x_1^*$  satisfying:

$$(U(x_1^*, e_1) - U(x_1, e_1)) \leq \frac{\varepsilon}{T}$$

If we compare the almost rational firm it will have a shrinking band of inertia similar in type to the rational firm with menu costs. However, the

behaviour is less predictable, since when it does change price, the near optimiser does not choose the unique optimal price, but a near-optimal price from the “band of inertia”, the range of acceptable epsilon optimal prices.

If the near rational firm has a proportional epsilon, then the band is fairly stable. For any  $t$  the band is defined by:

$$\left( \frac{U(x_t^*, e_t) - U(x_t, e_t)}{U(x_t^*, e_t)} \right) \leq 1 - \varepsilon \quad (1.2)$$

Hence a firm that gets to within 90% of the optimal ( $\varepsilon = 0.1$ ) will not necessarily have a decreasing hazard rate.

### 2.3 Inertia versus volatility.

If we contrast the Akerlof-Yellen inertia story with the inexperienced archer, the key difference is that in the inertia story the choice of action is like a *state variable*. It does not change from its current value without the explicit action or decision of the firm. The archer, however, has to shoot a different arrow each time and will almost never hit the same spot as the previous shot (with the exception of Robin Hood). There is no “simplicity” in hitting exactly the same spot as was hit the previous time: it is just as difficult as hitting the bulls-eye (or even more so, since the bulls-eye is easier to see). To be more precise, if it takes no effort to do the same as before, then we are in an

Akerlof-Yellen world of inertia. Where it takes just as much effort to do exactly the same as before as for any other specific action, we are in a world of the inexperienced archer.

If we look at the strategic choices of a firm, are we in the world of archery or the Akerlof and Yellen inertia? If the variable in question is something clearly measurable *and* under precise control, then the firm can in principle easily choose to do exactly the same as it has done before.

We can think of price and output for example. The farmer is like an archer. Both the price and output can be measured, but they are not under the farmer's control. The output of the farmer is stochastic and depends on weather and other factors and the price might well be determined by the market at the time of sale. To get the same price (or output) would be almost impossible to achieve.

Of course, at a more fundamental level, you can say that the farmer does control something that can be measured and controlled (e.g. how many hours are worked, how many seedlings planted, animals reared). At this level, the farmer is also in an Akerlof-Yellen world. However, when we look at the farm from the perspective of price and output, neither is under the farmer's control. So, whether we classify a particular enterprise as belonging to the Archer/farmer set or Akerlof-Yellen inertia set will depend on how we describe

and specify the activities in terms of our economic model. It is not an absolute classification that is invariant to the purpose of economic modelling.

For example, if we think of explaining how much activity a farmer puts into producing wheat and how much into pigs, we might be in an Akerlof-Yellen world. The farmer controls its inputs into these activities, and its choice is in terms of full or almost optimisation given the known distributions of supply and demand side variables beyond the farmer's control. However, if we are looking at price and output, the farmer controls neither and we are in the world of the archer.

A restaurant is more predictable than the farm. It can certainly choose its menu and prices given a demand curve. However, even the restaurant does not have full control of output. There might be a bus strike meaning some staff cannot get to the restaurant; staff might fall ill; the credit card connection might fail meaning only cash can be used. From an economic modelling perspective, we might well decide that the "stochastic" element in the restaurant's production function might not be an important part of the story and so can be ignored. Typically, as economists when we model firms we make a decision over how much uncertainty we build into our modelling and our choice will depend on the context in terms of what we are trying model. For example, it is not usual to model the probability of a nuclear war that will

wipe out the human race. Clearly the probability of such a nuclear war varies over time. It was zero prior to the invention of the Bomb; since then it has varied, and has clearly been higher at times such as the Cuban Missile crisis than others. Usually, economists focus on just one or two main sources of uncertainty in a particular model: a technology or cost shock on the one hand and a demand shock on the other. If the firms are choosing conditional on the current shock (in our case the  $e$ ), then they are an Akerlof-Yellen world. If they choose prior to its realization, they are more like the archer.

To give a very simple example, suppose we have a monopolist setting the price with the given demand and no costs.

$$P = \max[A - x, 0]$$

Now suppose the intercept term  $A$  is a random variable which can take two values, high  $A_H$  and low  $A_L$  each with a probability 0.5. The profit is equal to revenue. The optimal flexible price  $P^*$  is equal to

$$P^* = \frac{A_j}{2}; j = H \text{ or } L$$

If the fully rational firm sets price knowing the value of  $A$ , it will set the price each period equal to the value of  $A/2$ , with output being the same value (since the demand curve has a slope of unity optimal output is also  $A/2$ ).

If we look at the almost rational firm, there will be a band of inertia around each optimal flex price: for the high intercept the set  $H = [P_H^* - s, P_H^* + s]$  and for the low intercept  $L = [P_L^* - s, P_L^* + s]$ . Since the payoff is quadratic, the distance  $s$  will be the same for both sets. Now, if the gap between  $A_H$  and  $A_L$  is large enough, then  $H \cap L = \emptyset$  - the almost-maximizer will move between the sets  $H$  and  $L$ . This is depicted in Figure 5. In this case we can think of the almost maximizing firm as the archer: it is trying to hit the optimal price as it switches between the two values. Whereas the optimiser divides its time between the two optimal prices, the almost maximizer will divide time between the two ranges of prices  $H$  and  $L$ .

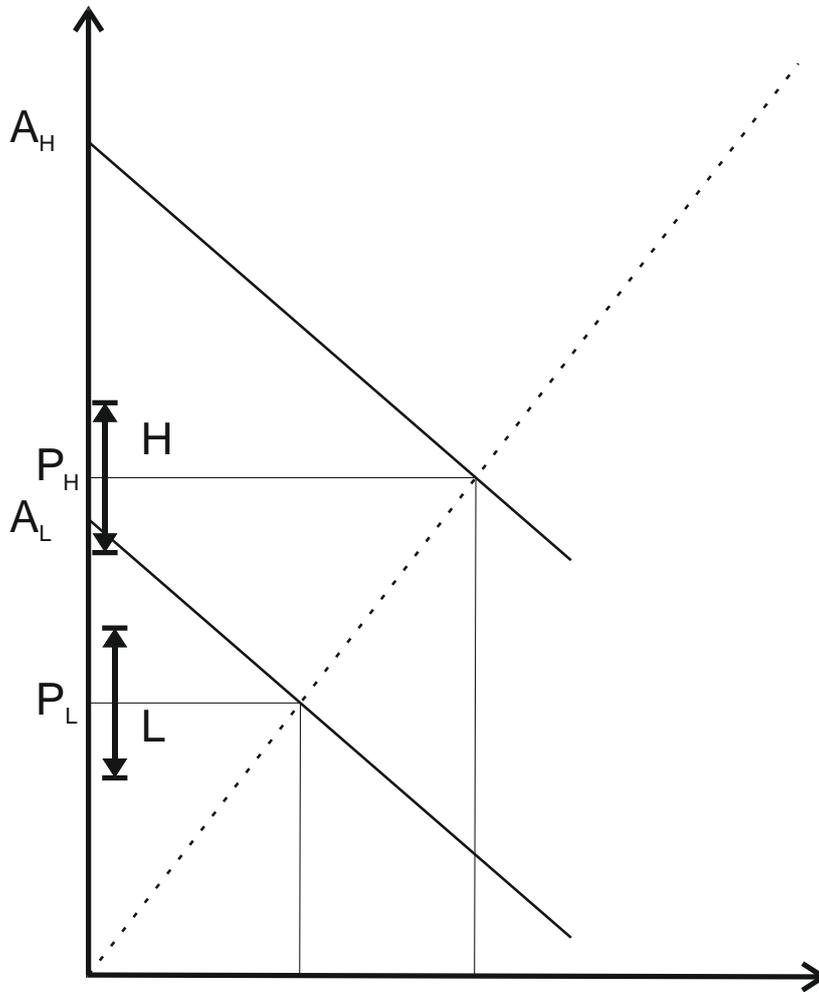


Fig 5:  $H$  and  $L$  no common elements.

However, suppose that the high and the low intercepts are close together so that  $H \cap L \neq \emptyset$ . In this case the almost optimal firm might adopt an entirely different type of behaviour: setting the same price  $P \in H \cap L$  for both realisations of intercepts  $A$ . This is depicted in Figure 6, with the intersection of  $H$  and  $L$  being represented by the thick dotted line. If we add a lexicographic preference for simplicity in terms of price stability, prices in the intersection will be the preferred choice for the almost maximizer.

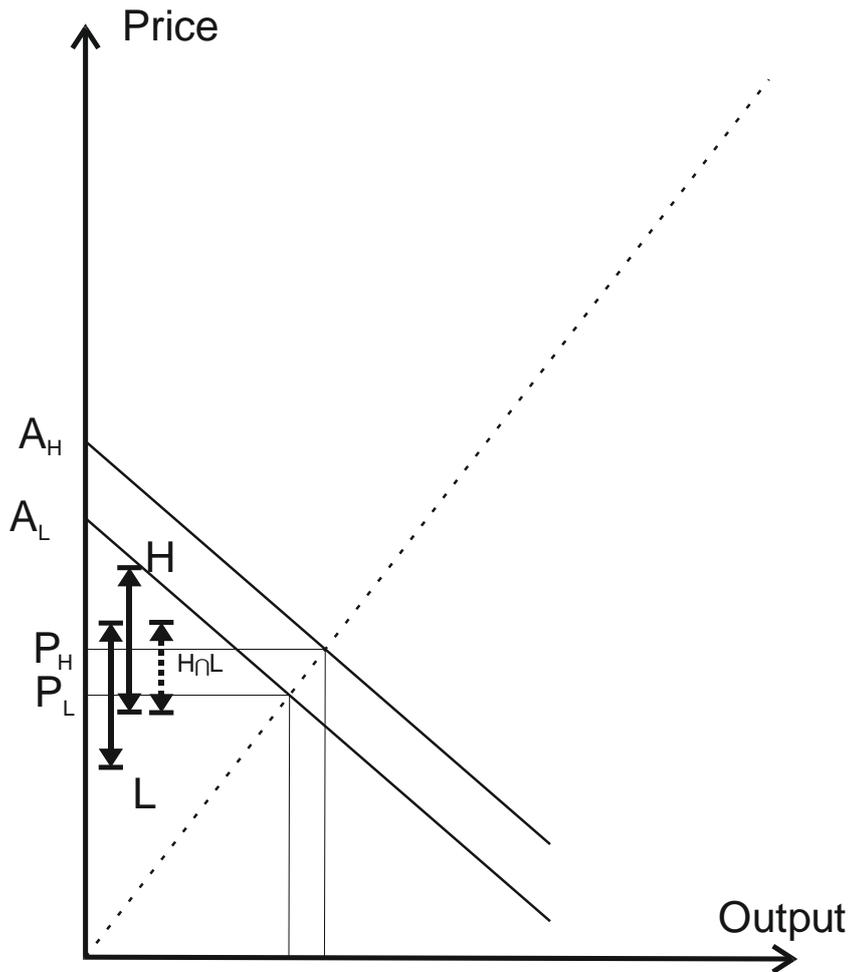


Fig 6: H and L common elements.

Now, suppose that the optimiser has to choose its price before it knows what value  $A$  will take. It will of course choose the price that maximize the expected profit:

$$E\Pi = 0.5P(A_L - P) + 0.5P(A_H - P)$$

The solution is to set the “average price”, since the payoff is quadratic and certainty equivalence applies.

$$P^* = \frac{0.5(A_L + A_H)}{2}$$

If we repeatedly observe the ex-ante fully rational price setter, since the problem is unchanged over time, we will see the same price set each period.

If we look at the almost rational firm in exactly the same position, we can see that there are a range of prices that can almost maximize profits: a distance plus or minus  $s$  from the optimal. The almost rational firm might choose to bounce around this set over time, or if it has a preference for simplicity choose to stick to just one price from within a range (which of course includes the unique optimum). As we saw before, there will be almost rational sequences of prices that bounce around the fully rational optimum and also fixed prices that remain the same but close to the fully rational optimum.

We have considered a range of dynamic scenarios. Whilst fully rational pricing usually indicates a unique path of prices (or unique conditional on errors), almost rationality generates a range of possibilities. There are many paths of prices that are approximately rational, and at each time there will be a range of prices that can occur. In this case we have “extra volatility” potentially generated by almost rationality. However, assuming that there is a desire for simple strategies, it is possible for prices to remain fixed with almost rationality where they are not fixed for fully rational firms.

### 3 Almost optimising what?

Managers of firms might have very different objectives to shareholders. However, in an oligopolistic environment maximizing profits is not necessarily the best way to maximise profits. The relation between the objectives of managers and profits is a potentially complex one, but we will consider it in a simple setting. What I want to argue is that we should see the profit motive as a long-run force, whilst the short-run objectives may differ. This is hard to model in an explicit real-time model. In terms of the structure we outlined in the introduction, we are moving from looking at the decision making process within the firm to how this is influenced by the external markets and economy as a whole. How the need to survive is a determining factor in the decision making of the firm in terms of its “decision rules”.

However, we can start by assuming that the managerial objectives are sales revenue. We use this as a simple example. Since John Vickers (1985), it has been known that in an oligopolistic setting that a sales-maximizing firm can earn more than a profit-maximizing. If we are modelling a duopoly, the individual firm’s reaction function will depend on the objective (sales, profits etc.). In a Cournot setting, sales maximization will lead to the firm wanting to choose larger outputs than profit maximisation. This can lead to the reaction function of the sales maximiser shifting out and moving the market Nash

equilibrium towards the Stackelberg point, increasing the profits earned by the sales-maximizer over the profits earned by the profit-maximizer.

To take the simple linear duopoly model of the previous section, suppose we have the industry inverse-demand function, with price depending on the total output produced by two firms  $i=1,2$ :

$$P = \max[1 - x_1 - x_2, 0]$$

With both firms having constant unit cost  $0 \leq c < 1$ . The reaction-function (best-response function) of firm 1 gives its optimal response given the output of the other firm. The sales and profit maximizing reaction functions are given by:

$$x_1^S = \frac{1 - x_2}{2}; \quad x_1^\Pi = \frac{1 - c - x_2}{2}$$

If we consider the Augustain Cournot (1838) case where there are no costs,  $c = 0$  then since profits are sales, the two outcomes coincide and  $x_1^S = x_1^\Pi$ .

However, as unit cost  $c$  increases, the two diverge: the sales maximising output does not vary, whilst the profit maximizing output declines. If we look at the Nash equilibrium, where firm 1 is a sales maximiser and firm 2 a profit maximiser, the relative profits will depend on the value of unit costs. If unit costs are zero, as in Cournot's original example, the two will be equal (since profit and sales maximization coincide). As unit cost increases, at first the sales

maximizer starts to earn more. This “peaks” and the advantage declines until eventually the sales maximizer earns less than the profit maximizer (and eventually starts to earn negative profits).

Thus we can see that it is perfectly possible for the non-profit maximizer to actually earn higher profits than the profit maximizer. Hence the capital market or stock-market requirement for survival will be satisfied by the sales maximizer, and indeed an evolutionary process could drive out profit maximizers in this market environment. There has been much research on the subject of delegation and the relationship between manager’s objectives and the profitability of firms, as surveyed for example in Sengul et al (2012). There are many possible outcomes, and crucially the nature of the game played by the duopolists matters. If the two firms are playing a game where their choice of actions are strategic substitutes (as in Cournot oligopoly) more aggressive behaviour such as sales maximization can prosper, whereas if they are playing strategic complements as in a standard Bertrand framework, the opposite will be true<sup>16</sup>.

However, our focus of interest is to see how almost optimisation will fit into this story. Sticking to the example of sales maximization we can see that

---

<sup>16</sup> For a full discussion see Huw Dixon (2001), chapter 6 Oligopoly theory made simple.  
<http://huwdixon.org/surfing-economics/surfing-economics-chapter-six.html>

almost sales optimisation will lead to a best-response *correspondence* i.e. a range of optimal responses to any particular output. Let us stick to the mixed-motive oligopoly where we have one sales near-optimiser and one profit maximiser. In figure 7 we show the reduced form profits of the sales maximiser and the profit maximiser as unit cost varies from 0 to 0.5 in units of 0.1 (for unit cost above 0.5 profits become *negative* for both types of firm).

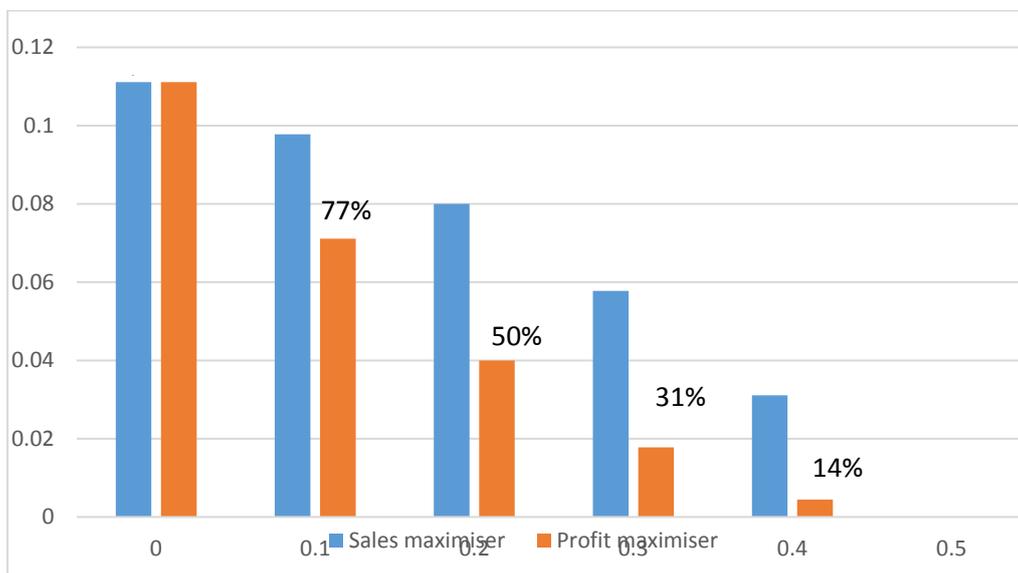


Fig 7 Mixed motive duopoly.

We can see that as unit cost increase, the sales maximiser does better except when unit cost is zero or 0.5, when they are equal. For values of unit cost in between 0 and 0.5 sales maximisation earns more at each value of unit cost, with profit maximisation earning a smaller percentage of the sale maximiser.

Hence, in a simple Cournot environment, we can see that *near maximisation* by a firm orientated towards sales (i.e. an almost sales

maximiser) will do better in terms of profits than the corresponding strictly rational profit maximiser<sup>17</sup>. Indeed, the near sales maximize will earn more than a near profit maximizer (so long as the epsilon is small enough).

In terms of the big picture, we can see that decision rules can survive and enable the firm to thrive even when they are not profit oriented or perfectly rational.

Another example is the evolution of cooperation with satisficing firms I explored in Dixon (2000). We can now introduce the capital market as an explicit force and require that firms earn at least normal profits, where we can interpret normal profits as the average across the whole economy. Think in terms of a simple prisoner dilemma type situation, where if both firms cooperate they get 2; if they both compete (defect), they get 1. If one defects whilst the other cooperates, it gets 3 and the other 0. Now, suppose that we think of firms playing this across the whole economy in pairs, locked into their local industry. The capital market is present in the sense that the firms are required (at least in the long-run) to earn normal profits, which are interpreted as average profits across the whole economy. As in satisficing models, I assume that if the firm earns at least average profits it continues to pursue its current strategy. If it is not earning normal profits, it will switch to some other

---

<sup>17</sup> I am assuming that the nearness or epsilon is small of course.

strategy<sup>18</sup>. In this simple example of two strategies (cooperate, defect), so if you are below average you could switch to the other strategy with a high probability and stick with the current strategy with a low probability.

In Dixon (2000) I show that in this set up, collusion will come to predominate (Theorem 1). The reasoning behind this result is simple. In the prisoner's dilemma example, there are three possible states: both firms cooperate and earn 2; both firms defect and earn 1; and two mixed states where one firm defects and the other cooperates and on average the two firms earn 1.5. The average payoff in the economy will then be the weighted average of these three payoffs, the weights being the proportion of industries in each state. Average profits in the economy will thus range from 2 to 1. Cooperation becomes an absorbing state: if you are in an industry where you are both cooperating, your payoff will be above average if some proportion of industries are in the defect or the mixed state.

Now, if we look at the industries in the mixed state, one firm is doing very well (3 will necessarily be above average) and one will be doing very badly and earning below average profits. The firm doing badly will have to change its strategy: for example changing from cooperate to defect. Hence, industries

---

<sup>18</sup> Technically, it is assumed that the firm randomly switches and all strategies have a non-zero probability of being chosen.

that are in a mixed state will transition to the defect state. Given that there is a measurable proportion of industries in the cooperation state, then the average payoff must exceed 1. Hence, for industries in the defect state both firms will be earning below average profits. Hence, both will switch strategy and with some non-zero probability will move to the cooperative state. Thus, in the long-run, all industries become cooperative.

The evolution of cooperation has relied on firms remaining fixed together in the same industry competing over time. However, as Jonathan Bendor, Dilip Mookherjee, and Debraj Ray (2001) showed, even with random matching (i.e. firms being randomly paired each period) the result of evolving cooperation can still come about (although is not so inevitable).

The important point here is that cooperation is a *dominated* strategy: it is not that choice of the rational agent. However, the external agency of the capital market acts to force firms to choose the cooperative strategy. The prisoner's dilemma is a simple two choice example. However, even in more complicated examples there will be a tendency to collusion (for example, Cournot duopoly with a large but finite number of output choices).

The cooperative output or strategy might be a long way from the optimal, not even being "almost" optimal, if we interpret optimal as profit or

payoff maximizing. Is there any way of thinking of this sort of behaviour as near rational?

The idea that cooperative behaviour in the prisoners dilemma can be thought of as almost rational goes back at least to Roy Radner (1980). His paper analyses behaviour in a repeated game setting with common knowledge and so is not directly applicable to the evolutionary models we are considering. He showed that with a finitely repeated game, cooperation could be a subgame perfect epsilon-equilibria at the beginning of the game up to a point somewhere near the end.

However, we can think of the simple decision rules as tit-for-tat (or indeed tit for two tats) as being approximately rational. Keeping things simple, suppose that both players are following the grim trigger strategy: cooperate until the other firm defects, after which you defect until the end of the game. Over  $T$  periods, with the payoffs of the prisoner's dilemma you will expect at period  $t$  to earn  $(T-t) \cdot 2$  if both firms follow grim trigger strategy. You can defect in the last period and increase your payoff to 3 for the last period, a gain of 1. If we express the epsilon as a proportion of remaining profits, then without discounting, you will plan to defect in the last period if:

$$\frac{(T-t) \cdot 2}{(T-t) \cdot 2 + 1} \geq 1 - \epsilon$$

The numerator is the payoff from cooperating until the end: the denominator is the payoff from defecting in the last period and gaining an extra unit (3 as opposed to 2). If the numerator is within proportion  $\varepsilon$  of the denominator, then at time  $t$  it is almost rational to plan to cooperate until the end. As in Radner's rather more detailed modelling, we can see that as  $t$  gets closer to  $T$ , cooperation until the end will cease to be  $\varepsilon$  optimal for  $t$  close to  $T$ .

We can conclude this section by saying that when we look at the need for firms to survive in an environment where they are interacting with other firms they may end up with decision rules that are different to the profit-maximizing rule. From the perspective of near-rationality, if we can characterise the optimal decision rule as optimising something (sales, joint-profits etc.), then we can also think of epsilon maximization of this objective. However, will evolutionary forces of the market drive out the epsilon maximize or drive the epsilon to zero? I would argue not. The epsilon is envisaged as a small number or proportion. As we argued in the introduction a firm with a large epsilon (a stupid firm) would indeed not survive. However, I would argue that the market itself have some grit or "epsilon" that will not drive all firms to have exactly the same (risk adjusted) profits. Small differences in profits will pass by largely unnoticed and not lead to any response. That is the essence of the managerial discretion idea and Hick's quiet life writ small.

There are, I believe, counteracting tendencies in the interaction of firms within their own markets and the economy as a whole. As competitors within an industry, firms want to outdo their rivals and this tends to make them more competitive (as in the sales maximiser example). However, the capital markets want industries that are less competitive. The balancing of these local and global forces is a subject worthy of future research.

#### 4 Conclusion.

In a simple static framework, epsilon or near rationality is easy to understand. The agent reasons or reaches its decision by some process which gets it close to the maximum - as in the Figure 1 "Hill diagram". In this paper, I have sought to examine what epsilon optimisation means in a dynamic setting where a firm is deciding what to do over time under changing conditions. Near rationality can give rise to erratic behaviour I have described with the Archer metaphor: the choices can be more variable than the optimal. However, we can also see the emergence of inertia or keeping to the same strategy over time, particularly if we assume that the firm has a lexicographic preference for simplicity and prefers the almost rational strategies that involve holding the strategy constant. This is similar to the menu-cost model of the rational optimiser subject to lump-sum menu costs, but as I showed it differs since

there is no “menu cost” and whilst the menu-cost optimum might be unique, the near rational outcomes are never unique.

Turning to the survival of the firm over time, near rationality is only nearly as good as its fully rational counterpart. If we can say that long-run survival depends on maximizing some objective, then we can also say that long-run survival may be consistent with near maximization of the same objective. However, the decision rules consistent with long-run survival depend on the strategic environment the firm finds itself in and may have no obvious interpretation. The behaviour of firms may be consistent and rule following, but the rule cannot be understood as resulting from maximizing the firms objectives except in the context of the interaction of the firm with the wider market and economy as a whole. In that context the behaviour of the firm can only be understood by explaining how the decision rule arises, not as the solution (exact or approximate) to an optimisation problem.

## Bibliography.

- Akerlof G and Yellen J (1985a). Can small deviations from rationality make significant differences to economic equilibria?, *American Economic Review*, 75, 708-720.
- Akerlof G and Yellen J (1985b). A Near-Rational Model of the Business Cycle, With Wage and Price Inertia, *Quarterly Journal of Economics*, 100, 823-838
- Alchian A. (1950), *Uncertainty, evolution and economic theory*. *Journal of Political Economy* 58, 211–222.
- Alvarez F and Lippi F (2014), "Price setting with menu costs for multi-product firms", *Econometrica*, 82, 89-135.
- Barnard C (1938), *The functions of the executive*, Harvard University press.
- Conlisk J (1998), *Bounded rationality and market fluctuations*, *Journal of Economic Behaviour and Organisation*, 29, 233-250.
- Cournot A (1838), *Researches on the Mathematical Principles of the Theory of Wealth*, Macmillan, London.
- Cyert R and Marsh J (1962), *A behavioural theory of the firm*, Wiley-Blackwell.
- Hicks J. (1935), *Annual survey of economic theory: the theory of monopoly*. *Econometrica* 3, 1–20.
- Klenow P and Malin B (2010), "Microeconomic evidence on price-setting", Chapter 5 (Pages 231-284), *Hanbook of Monetary Economics*, Volume 3, Edited by Benjamin M. Friedman and Michael Woodford. Elsevier.
- Mankiw, N. G. (1985). 'Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly', *The Quarterly Journal of Economics*, vol. 100(2), pp. 529-537.
- Radner R (1980), *Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives*, *Journal of Economic Theory*, 22, 136-154.
- Selten, R (1975), *Reexamination of the perfectness concept for equilibrium points in extensive games*, *International journal of game theory* 4(1), 25-55
- Sengul M, Gimeno J, Dial J. (2012) *Strategic Delegation: A Review, Theoretical Integration, and Research Agenda*, *Journal of Management*, 38(1), 375-414.
- Seth A, Thomas H (1994), *Theories of the firm: implications for strategy research*, *Journal of management studies* 31, 139-301.
- Simon H (1947), *Administrative behaviour, A Study of Decision-Making Processes in Administrative Organization*, Free Press.
- Simon H (1982), *Models of Bounded Rationality*, Vols. 1 and 2. MIT Press.

## Appendix: epsilon optimisation.

In order to make the examples in section 1.2 even more concrete, we can think of the explicit functional forms for cost and demand:

$$P = \max[1 - x, 0]$$

$$C = c.x$$

Where marginal cost is  $c < 1$  so that the marginal cost at zero output is less than marginal revenue. In this case, we have:

$$U(x) = x.(1 - x) - c.x$$

$$\frac{dU}{dx} = 1 - (2 - c).x$$

$$\frac{d^2U}{dx} = -(2 - c)$$

Since  $c < 1$ , payoff is strictly concave, so that there exists a unique interior

maximum:  $x^* = \frac{1}{2 - c}$

Since the payoff is quadratic, the second order Taylor expansion is exact, so that

$$U(x^*) - U(x') = \frac{2 - c}{2}(x' - x^*)^2$$

Hence we have

$$U(x^*) - U(x') < \varepsilon \Leftrightarrow |x^* - x'| < \frac{\sqrt{2\varepsilon}}{2-c} = \kappa(\varepsilon)$$