



Almost-Maximization as a Behavioral Theory of the Firm: Static, Dynamic and Evolutionary Perspectives

Huw Dixon¹

Published online: 26 September 2019
© The Author(s) 2019

Abstract

In this paper we consider the effect of epsilon maximization on firm behavior. In particular we focus on the dynamic behavior of firms with the use of the example of price-setting: We show how almost-rational firms can be more volatile in their behavior. However, if a lexicographic preference for simplicity is made, then we can explain nominal price rigidity as a result of epsilon optimization. The behavior of the firm—which is consistent with its long-term survival—is examined. We argue that epsilon-optimization is consistent with survival in any context in which something is optimized: such as sales revenue.

Keywords Bounded rationality · Dynamic choice · Inertia · Price rigidity · Volatility

JEL Classification B52 · C61 · D21

The idea of bounded rationality became popular in the economics of the firm with the publication in Cyert and Marsh (1963) of the book *A Behavioral Theory of the Firm* by Richard M. Cyert and James G. March.¹ Whilst the dominant model of rationality has remained constrained-optimization 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 that asks how we should understand near rationality in our modelling of firms.

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

¹ In the wider social sciences, the modelling of bounded rationality in Simon (1947) had already begun to have an effect earlier. Simon (1982) represents a more mature summation of Simon's thinking.

✉ Huw Dixon
dixonh@cardiff.ac.uk

¹ Cardiff Business School, Colum Drive, Cardiff CF14 0RZ, UK

1. What are the objectives of firms (for example profits, sales, managerial utility)?²
2. As organizations, do firms seek to maximize in the pursuit of their goals, or are they boundedly rational in some sense (such as satisficing)? and
3. How does our conception of the *behavior* of the firm fit with the explanation for the *existence* or *survival* of firms—are firms generally efficient institutions in terms of minimizing transactions costs and responding to changing market conditions?

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 that interact with each other or as incumbents that compete with potential entrants) and at higher levels in capital markets where firms' overall performance is characterised (for example) by share prices and related profit metrics. Poorly managed firms cannot survive: They are driven out of business by more able competitors, or perhaps they 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 Alchian (1950, p. 213): “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 Hick’s (1935, p. 8) thought that “the best monopoly profit is a quiet life”, there are limits to under-performance that make even the most sedentary manager subject to scrutiny. Barnard (1938) argued that one of the main reasons that firms do not last forever was that over time they cease to be effective.

Whilst external forces operate on firms, most firms face sufficient slack to allow some degree of “managerial discretion”. This is particularly true for owner-managed firms: So long as it remains solvent, an owner-managed firm can continue to operate. Indeed, some owners might even subsidise a loss-making firm from other revenues for reasons beyond financial considerations.

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 behavioral theories in economics (Seth and Thomas 1994, p. 175):

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 behavior, 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.

² This of course depends on the key individuals and groups within the firm—for example, senior management or owners—and the interactions of these individuals within the firm.

We see the challenge for a behavioral theory of the firm in terms of providing equilibrium decision-making processes, which are summarised as heuristic decision-rules. The decision rules should arise from within the firm in the pursuit of its goals and enable the firm to survive in its competitive environment.

In Dixon (1992) I argued for the concept of near-optimization, which is sometimes referred to epsilon-optimization.³ Epsilon-optimization means that the firm chooses an action that yields a payoff within epsilon of its optimum. The epsilon is assumed to be small but strictly positive, with strict optimization being the limiting case where epsilon is zero. 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”, and thus enables managerial discretion. This general approach leaves open the first issue of behavioral theory: What is almost-optimised: managerial utility, sales, or profits?

This paper will concentrate on these general issues with a focus on two particular angles: First, the interpretation of bounded rationality in dynamic Settings; second, how rivalry in an oligopolistic setting and the need to survive interact with the goals of the firm.

We apply the theory of epsilon optimization to a dynamic model 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 behavioral theory of the firm: In practice, we often observe firms’ setting the same price for prolonged periods of time.⁴ Akerlof and Yellen (1985a, b) put forward the argument that near-rational firms would display nominal price rigidity rather than vary their price to track the optimum perfectly. An alternative put forward by Mankiw (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 menu 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 obtain inertia from bounded rationality one needs to add a preference for simplicity of strategy, preferring repetition to variation.

The second issue that we examine is that of *what* is nearly maximized.

The long-run survival of firms depends on earning enough profits (normal risk-adjusted profits). Only certain forms of behavior—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;

³ This is often written as ϵ optimization.

⁴ The famous example is the 5-cent bottle of Coke that lasted over 60 years. However, the price-quote data that are used to construct CPI data in many countries have revealed a long fat tail of prices that change very infrequently. See Klenow and Malin (2010) for a survey of the evidence.

and so on. Short-run profit maximization as a form of behavior may be inconsistent with long-run survival, and other forms of behavior (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 boundedly rational firms is about more than just looking at decision-making processes either in the head of an individual or in the heads of a 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 its interactions with the wider economy through capital markets. We will see that the “decision rules” that emerge result from the interaction of all three levels: the individual, the collective, and the wider economy.

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

1 Almost-Optimality

Perfect rationality⁵ has been the approach that has been 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. The agent has 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 a technological constraint. 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 as 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:

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

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^* ; both can be seen as depending on S . In the case of the consumer, the budget set is determined by prices and income: Indirect utility identifies the maximum utility given prices and income. The Marshallian demand is the optimal choice of consumption given prices and income. The supply and input-demand functions of competitive firms are likewise determined.

⁵ This is often called the “neoclassical” model when applied to the firm, or *hyper-rationality* in general.

Now, let us see under what external circumstances this sort of behavior might arise in a firm. Perfect competition is the market structure that would lead to this sort of firm behavior. 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 this textbook scenario, 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 in the long run (unless, perhaps, in the special case where it is universal across all firms).

1.1 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 they 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, which were first introduced by 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 , the chosen action x' is within a distance κ of x^* ⁶:

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

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 (e.g., Bronze instead of Gold). However, a small error in action that gives rise to a big loss can occur only 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,⁷ then we can ensure that small mistakes in the implementation of the strategy will give rise to small losses in payoff. Put simply, if U is continuously differentiable, then for any $\varepsilon > 0$, there exists $\kappa > 0$ such that: if $\kappa > M(x', x^*)$, then $U(x^*) - U(x') < \varepsilon$

We can go further and think of this as defining the distance κ as an (increasing) function of ε : $\kappa = \kappa(\varepsilon)$; $d\kappa/d\varepsilon > 0$: If we want the loss in terms of payoff to be smaller than a certain level, then we need the strategy to be sufficiently close to the optimal strategy.⁸

⁶ For example, if x is an n -dimensional vector, a possible metric could 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.

⁷ This is usually written as $U \in C^2$.

⁸ In an online appendix, we have a brief formal explanation of these ideas at <http://huwdixon.org/online%20appendix.pdf>.

Can we go in the other direction and claim 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 be only one local maximum, which is the global maximum. If there is weak concavity or even some convexity, then there can be local maxima that are close to the global maximum in terms of payoff but a long way away in terms of action. Think of two hills that are separated by a valley; hill A is slightly taller than hill B. We can be almost as high as the summit of A by being close to the summit of A or by 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.

2 The Implications of Almost-Optimization in Dynamic Models

If agents are almost-maximizing, what are the implications for how firms and markets behave over time? This has been the focus of research in the past: notably by Akerlof and Yellen (1985a, b) and Conlisk (Conlisk 1996).

There are two contrasting views:

- *Bounded rationality increases volatility.* Bounded rationality produces an additional source of noise, which can be thought of as an “erratic archer”.
- *Bounded rationality increases inertia.* Rather than responding to all shocks or changes, agents will just respond to those that take it out of its “comfort zone” or “Band of Inertia” as in Akerlof and Yellen.

The notion that bounded rationality might increase volatility arises from the following simple line of reasoning. Suppose that the optimal level of x at time t is driven by some random shock⁹ ε_t : $x_t^* = \bar{x} + \varepsilon_t$ for example. With full rationality, the choice would be the optimal value: the volatility of x_t would simply be the volatility of the “shock” ε_t . With bounded rationality, there is another random “shock” or fault in execution in the form of “trembles” v_t :

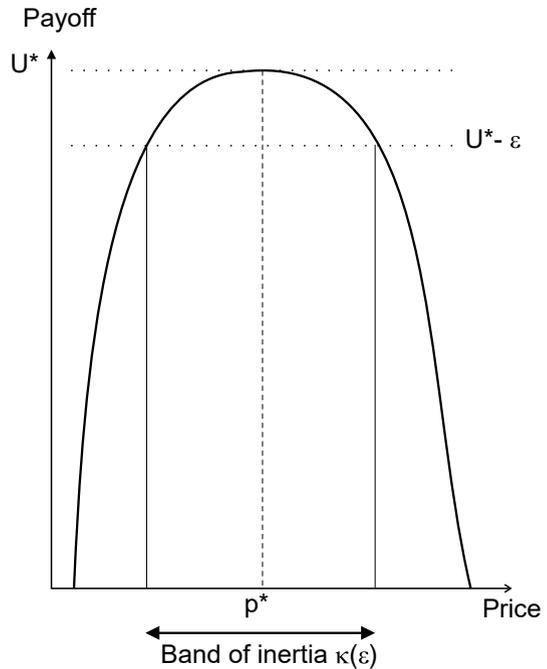
$$x_t = x_t^* + v_t = \bar{x} + \varepsilon_t + v_t$$

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 implication that bounded rationality increases 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 the exact location of the bulls-eye, but due to lack of skill the arrows are spread around the

⁹ For example, we could assume that the shock is normally distributed with mean zero.

Fig. 1 Epsilon optimization and the Band of Inertia



bulls-eye. A card player might not be able to calculate the odds accurately and so not know the optimal play 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 that alter 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, which is the continuous-time equivalent of a random walk).¹⁰

However, in their seminal paper, Akerlof and Yellen (1985b, pp. 823–824) put the argument 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

¹⁰ To see a recent example of this approach, see Alvarez and Lippi (2014). The basic technique dates back to Sheshinski and Weiss (1977).

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.

In a static framework the argument is very much as in the Fig. 1 "Hill" diagram, which is understandable in terms of a simple one-dimensional metric where the payoff on the vertical axis is a concave function of price on the horizontal axis

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 ε of the optimum.

2.1 Dynamic Models with no Uncertainty

In a purely static setting, models with menu costs and bounded rationality look very similar: One can think of ε as a lump sum cost with perfect optimization or originating from imperfect optimization. In a dynamic setting, matters are more complicated, and we can distinguish between a menu cost and approximate optimization. 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 that is close to the optimum in terms of an appropriately discounted payoff.¹¹ We then compare this bounded-rationality outcome to the alternative perfectly-rational outcomes with or without menu costs.

First we will consider a simple two period ($t=1, 2$) 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 (in this example, the menu cost can be thought of as a *switching* cost). The fully-rational firm chooses x_1 and x_2 to solve the maximization problem:

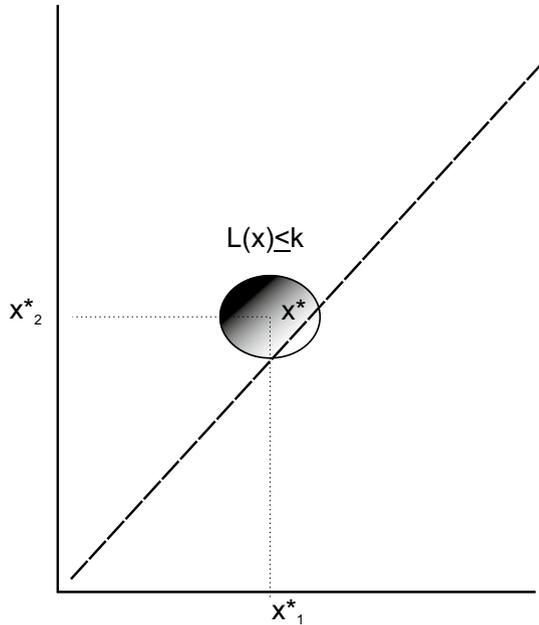
$$\max_{x_1, x_2} U(x_1, e_1) + U(x_2, e_2) - \gamma I(x_1, x_2)$$

where $I(x_1, x_2)$ is an index function which is zero if $x_1 = x_2$, or one otherwise. We can first define the optimal flexible action for each period (the case of $\gamma = 0$ for all x_1, x_2). This is simply derived from the first order conditions given the shift variable, and the corresponding optimal action and payoff x^* and U^* as a function of e :

$$\begin{aligned} x^* &= x(e) \\ U^* &= U(e) \end{aligned}$$

¹¹ Of course, bounded rationality might take the form of inappropriate discounting. For example, there exists considerable evidence that humans have hyperbolic discounting and/or myopic planning horizons (see review by Frederick et al. 2002). This is a topic we do not have space to explore in this paper.

Fig. 2 Two period choice



13 Of course, bounded rationality might take the form of inappropriate discounting. For example, there exists considerable evidence that humans have hyperbolic discounting and/or myopic planning horizons (see review by Frederick et al. 2002). This is a topic we do not have space to explore in this paper.

If we ignore menu costs ($\gamma = 0$), we can define our metric in terms of a loss function, giving the lost payoff for any choice of (x_1, x_2) relative to the flexible optimum:

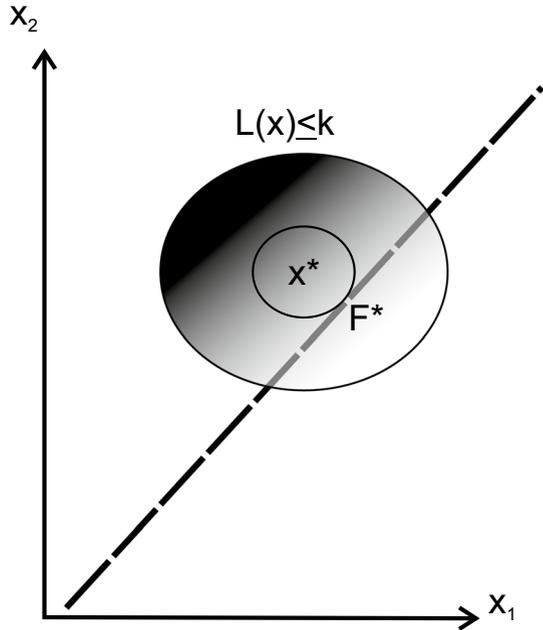
$$L(x_1, x_2) = U(x_1, e_1) + U(x_2, e_2) - [U^*(e_1) + U^*(e_2)]$$

In Fig. 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 κ as the shaded circle. The 45-degree line identifies all combinations where the action is unchanged, $x_1 = x_2$. If we are thinking of the action as setting a price, then the 45-degree line represents fixed prices across the two periods. Whilst we will interpret x as “price” in what follows, in fact it can be any “action” that determines payoffs.

As depicted, the 45° line intersects the shaded circle, so that there are constant or fixed prices that involve a small enough loss. For all of the other elements inside the circle, the action is different (prices change). 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 fixed price with loss less than or equal to κ . However, let us focus on cases where the 45° line intersects the shaded area.

From Fig. 2, we can see immediately that if there is a menu-cost $\gamma = \kappa$, then the optimal choice of prices across the two periods will be the subset of price pairs on the 45° line. This subset will earn strictly higher payoffs than all of the

Fig. 3 The optimal choice with menu costs



other pairs in the circle (including the optimal flexible solution). By choosing from the 45° line, 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 with the same x in both periods:

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

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

The optimum fixed-price action is where the 45° line is tangent to the isopayoff circle at point F^* , which gives the maximum payoff from 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. We can now turn to the set of almost rational choices. The whole set $L(x) \leq \kappa$ represents the set of almost rational choices. Within this set, there are two possibilities:

1. The near-rational firm 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 firm may choose combinations in the set $L(x) \leq \kappa$, but prefer to choose simple strategies in which actions are the same in each period, i.e. lying on the 45° line.

In order to get outcome 2, we need not only near-rationality, but also a lexicographic preference for simple strategies that involve constant actions over time. We can claim 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” as being a small cost of changing prices over time. The presence of this cost causes the firm to prefer strategies along the 45-degree line that are a tiny bit better and will be preferred and chosen.

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 on which side of the 45-degree line it is)—even when the optimal flexible action shows a particular pattern.

In Fig. 3, the optimal flexible action involves *increasing* x , since $x_2^* > x_1^*$ for points above the 45-degree line; but there are many pairs in the shaded region below the 45-degree line that go the opposite way ($x_2 < x_1$). It is only with option 2 combined with a preference for simple strategies that we will see a fixed price across the two periods ($x_2 = x_1$). The “band of inertia” is then equivalent to the intersection of the 45-degree line and the shaded circle: The menu-cost optimizer in contrast will pick the single point F^* on that line. In terms of the hill in Fig. 1, the band of inertia is the 45-degree line, and point F^* is value of x that gets us to the top of the hill.

2.2 Dynamic Models with Uncertainty

We will now see how uncertainty alters our analysis and concentrate on the case of price-setting (although the argument applies to any variable(s) of choice). Most dynamic models with menu costs allow for uncertainty. It is common to assume that the optimal price follows a random walk¹²:

$$p_{t+1}^* = p_t^* + e_t \tag{1.1}$$

where e_t is a white noise error.¹³ In each period t , 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 of e_t in both periods (effectively there was perfect foresight).

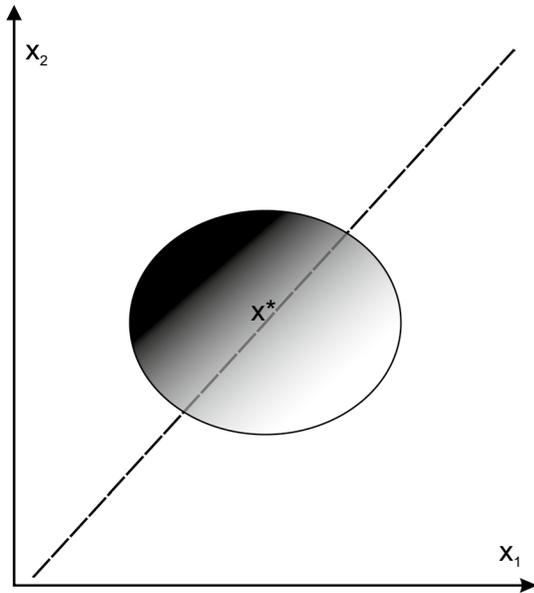
From (1.1) the expected value of all future errors is zero and the optimal price is expected to be the same in all future periods $i \geq 0$:

$$E_t [p_{t+i}^*] = E_t [p_t^*]$$

¹² This is the discrete time version of the continuous-time model used (see, for example, Alvarez and Lippi 2014 or Sheshinski and Weiss, 1978). Note that in the previous section, e was a parameter that influenced the optimal price. In a dynamic setting, we allow e to vary over time and possibly be stochastic.

¹³ Technically, e_t has an expected value of zero and is not serially correlated.

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



That is, in terms of our simple two period diagram, the *optimal* actions lie on the 45° line where the optimal strategy involves *planning* to set the same price in both periods 1 and 2 ($p_1^* = p_2^*$). Thus the optimal fix-price strategy and the optimal flex-price strategy coincide—in terms of Fig. 3; F^* and x^* coincide. This is depicted in Fig. 4: the 45° line now lies in the middle of the shaded zone and the optimal strategy is on the 45° line.

Let us now assume that the firm chooses the optimal strategy in period 1. When period 2 arrives, the second-period shock e_2 is realized, which is represented by a vertical shift in the optimal price equal to e_2 . Since the period 1 payoff is already a given, in period 2 the agent has the simple choice of whether to leave the price at its planned value or to change it to the optimal price given the realized 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.

$$U(p * (e_2)) - U(p * (0)) < \kappa$$

As in the one-period case, the difference between a menu cost and bounded rationality is unobservable, because 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 $p_{t+i}^p = E_t p_{t+i}^* = p_t^*$, where p_{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 realization of each shock.

With a menu-cost, the perfectly-rational firm will consider at any time after the first period, $t=2...T$, whether to leave its price where it is (at the price from the previous period $t-1$) or switch to the current optimal price. Without discounting, this is a simple problem: For menu cost γ , 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.¹⁴ Since the expected value of future shocks is zero, this takes the simple form:

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

For each period there is a band of inertia around the optimum, which becomes larger over time. Comparing the one-off menu cost with the stream of losses over $T-t+1$ periods, the firm will hold constant its current price only if the current loss times the number of remaining periods is less than the menu cost.

At time $t=2...T$, the current shock e_t is known and is expected to remain in place for the remaining $T-t+1$ periods. The firm can stay put and incur no menu cost, and then the per-period loss is the difference between the optimal (with no menu costs) and what is earned at the current price. This per-period loss times the remaining number of periods is then compared to the one-off menu cost. *The earlier that this is considered, the more likely the firm is 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.

Hence there is a clear prediction: With a finite time horizon, the band of inertia gets larger as the firm gets 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 the firm changes 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 accumulate for a longer period of time. Near the end of the cycle, the firm will be changing the price soon for free, so there is less to worry about.

We now turn to the case of an almost-rational ϵ -optimizing firm. At time t , we then have an almost optimal sequence of planned prices $(p_t, p_{t+1} \dots p_T)$ such that:

$$(T - t + 1)U(p_t^*, e_t) - \sum_{i=0}^{T-t} U(p_{t+i}, e_t) \leq \epsilon$$

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

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 that are held fixed, which restricts our attention to plans with $p_{t+i} = p_t$, so that:

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

This is almost the same as the menu cost equation above, except that here the current price p_t is chosen rather than the inherited p_{t-1} . However, invoking again the lexicographic preference for simplicity our firm will prefer to set $p_t = p_{t-1}$ rather than any other almost-rational price. In period 1, when the menu cost optimizer gets to freely choose its opening price it will set the unique optimal price $p_1 = p_1^*$; whereas the almost optimizer will have a band of choices around p_1^* satisfying the almost optimality condition:

$$(U(p_1^*, e_1) - U(p_1, e_1)) \leq \frac{\varepsilon}{T}$$

The almost-rational firm will have an expanding band of inertia as time moves towards T that is similar in type to the rational firm with menu costs. However, the behavior of the near-rational firm is less predictable. When it does change price, the near-optimizer does not choose the unique optimal price, but instead chooses a near-optimal price from the range of acceptable epsilon-optimal prices.

There is substantial empirical evidence that the firms are less likely to change price as time goes on (at least initially).¹⁵ We find that a possible explanation is almost-rationality with a fixed time horizon: As we approach the terminal period, there is less to gain by changing price.

However, almost-rationality on its own will not guarantee this: We need the additional assumption that there is a preference for simple strategies: in this context: keeping the price unchanged. To some extent there is a degree of “observational equivalence” between the menu cost interpretation with perfect rationality and the almost-rational. However, the almost-rational firm has a wider range of prices to change to when it does make the change. This result could be used to observationally or experimentally differentiate between the two types of behavior.

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 similar to 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 twice (with the exception of Robin Hood). In this case, there is no “simplicity” in hitting the exact same spot as was hit previously.

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

¹⁵ In fact, this can explain the robust empirical finding of a decreasing hazard rate.

the same as before as for any other specific action, we are in a world of considerable volatility due to the inexperienced archer.

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

Consider the price and output decisions of a farmer, 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, as it depends on weather and other factors. Price might well be determined by the market at the time of sale. To get the same price (or output) at two different points in time would be almost impossible to achieve.

Of course, at a more fundamental level, one can say that there is something that the farmer can measure and control: e.g., how many hours are worked, how many seedlings planted, how many animals reared, etc. 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.

If we think of explaining how much activity a farmer puts into producing wheat and how much into raising pigs, we might be in an Akerlof-Yellen world. The farmer controls the inputs into these activities, and the choice is in terms of full- or almost-optimization—given the known distributions of supply and demand side variables that are 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.

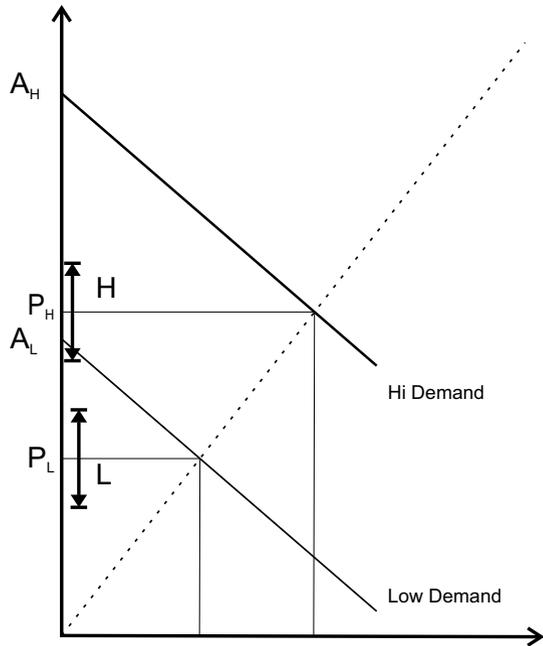
A restaurant is more predictable than the farm. It can certainly choose its menu and prices, given a demand curve. However, even the restaurant does not have full control of output. There might be a bus strike, which would mean that some staff cannot get to the restaurant; staff might fall ill; the credit card connection so that 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 model, and this decision will depend on the context in terms of what we are trying to 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 war varies over time. It was zero

prior to the invention of the atomic bomb; since then it has varied, and has clearly been higher at times such as the Cuban Missile crisis.

Usually, economists focus on just one or two main sources of uncertainty in a particular model, such as 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 in an Akerlof-Yellen world. If they choose prior to its realization, they are more like the archer.

Fig. 5 Hi and Low demand with no common elements



To provide a very simple modelling example: Suppose that we have a monopolist that sets the price with the given demand and no costs.

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

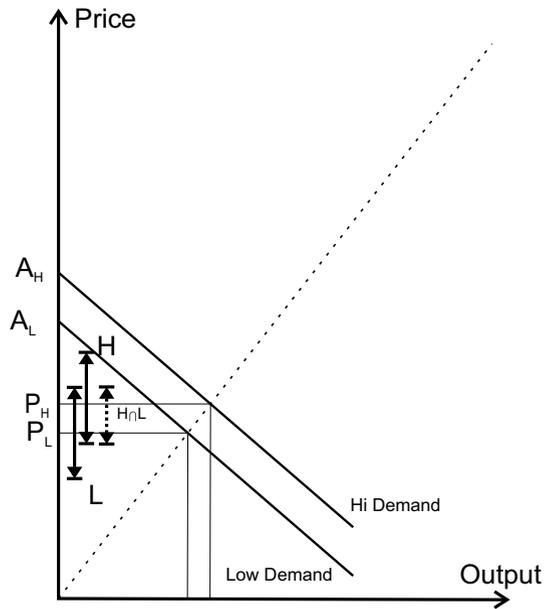
Now suppose the intercept term A is a random variable which can take two values, high A_H and low A_L , each with a probability 0.5. The firm's profit is equal to revenue. The optimal flexible price P^* is equal to

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

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

For the almost-rational firm, there will be a band of inertia around each optimal flex price. The maximum distance between this price and the optimum is s : 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, 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 Fig. 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. 6 H and L with 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 behavior: setting the same price $P \in H \cap L$ for both realisations of intercepts A . This is depicted in Fig. 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.

Now, suppose that the optimizer 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 that is 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, it might choose to stick to just one price from within a range (which of course includes the unique optimum). As we saw before, there are almost-rational sequences of prices that bounce

around the fully rational optimum; and there are also fixed prices that remain the same but close to the fully rational optimum.

We have considered a range of alternative 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” that is potentially generated by almost-rationality. However, on the assumption that there is a desire for simple strategies, it is possible for prices to remain fixed for almost-rational firms but not for fully rational firms.

3 Almost-Optimising What?

Managers of firms might have very different objectives from shareholders. However, in an oligopolistic environment the *attempt* to maximize profits is not necessarily the best way to *achieve* maximum 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 the profit motive is a long-run force, whilst 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 examining the decision-making process within the firm to how this is influenced by external market forces and the economy as a whole. We must consider how the need to survive influences the “decision rules” of the firm.

As a starting point, assume that the managerial objective is sales maximization. Since Vickers (1985), it has been known that in a Cournot oligopolistic setting a sales-maximizing firm can earn more than can a profit-maximizing firm. 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 the firm to choose a larger output than does a profit-maximising firm. This can lead to the reaction function of the sales-maximiser to shift out and move the Nash equilibrium towards the Stackelberg point, which increases the profits that are earned by the sales-maximizer over the profits that are earned by the profit-maximizer.

Thus we can see that it is possible for the non-profit-maximizer actually to earn higher profits than does 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 managers’ objectives and the profitability of firms, as is 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 behavior—such as sales maximization—can increase profit.

Whereas, if their choice variables are strategic complements—as in a standard Bertrand framework—the opposite will be true.¹⁶

Another example is the evolution of cooperation with satisficing firms, which is explored in Dixon (2000). We can now introduce the capital market as an explicit force and require that firms earn at least normal profits, which is defined to be average firm profit across the whole economy.

Consider a simple prisoner's dilemma type situation, where both firms earn 2 when they cooperate; if they both compete (defect), they earn 1. If one defects whilst the other cooperates, it earns 3 and the other earns 0.

Now, suppose that we think of firms that are playing this game 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. 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 strategy.¹⁷

In this simple example of two possible actions (cooperate, defect), a firm with below-average returns will switch to the other strategy with a high probability and stick with its current strategy with a low probability.

In Dixon (2000) I show that in this set up, collusion will come to predominate in each industry (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 are 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 a firm is in an industry where both firms are cooperating, the firm's 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 that is doing badly will have to change its strategy from cooperate to defect. Hence, industries 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. This will cause both firms to 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 to cooperation has relied on firms' remaining fixed together in the same industry and competing over time. However, as Bendor et al. (2001) showed, even with random matching—firms' being randomly paired each period—the result of evolving cooperation can still come about (although it is not so inevitable).

¹⁶ For a full discussion, see Dixon (2001, ch. 6).

¹⁷ Technically, it is assumed that the firm randomly switches and all strategies have a non-zero probability of being chosen.

The important point here is that cooperation is a dominated strategy, and firms have pressure coming from capital markets that forces firms to choose the cooperative strategy. The prisoner's dilemma is a simple two-choice example. However, even in more complicated settings there will be a tendency toward 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—even being “almost” optimal—if we interpret optimal as profit- or payoff-maximizing. Is there any way of thinking of this sort of behavior as near-rational or rational in a different sense?

The idea that cooperative behavior in the prisoner's dilemma can be thought of as almost-rational goes back at least to Radner (1980). His paper analyses behavior in a repeated game with common knowledge, which is a different setting from those considered previously. 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 final stage of the game. We can think of the simple decision rules such as tit-for-tat (or indeed tit for two tats) as being approximately rational—at least in the early stages.

We can conclude this section by saying that when firms are concerned with survival in an environment where they are interacting with other firms, they may end up using decision rules that are different from profit-maximization. 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 the epsilon-maximization of this objective.

However, will evolutionary forces of the market drive out epsilon-maximization or drive epsilon to zero? I would argue no. The epsilon is envisaged as a small number or proportion. As we argued in the introduction, a firm with a large epsilon (an error-prone firm) would indeed fail. However, I would argue that the market itself has some grit or “epsilon” that will allow long-run (risk-adjusted) profits to vary among firms. 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 in miniature.

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-optimization or near-rationality is easy to understand. The agent reasons or reaches its decision by some process that gets close to the maximum—as in the Fig. 1 “Hill diagram”. In this paper, I have sought to examine what epsilon-optimization means in a dynamic setting where a firm is deciding what to do over time and under changing conditions. Near-rationality can

give rise to erratic behavior, as is described by 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 its action constant. This is similar to the menu-cost model of the rational-optimiser that is subject to lump-sum menu costs; but as I showed it differs since there is no “menu cost”. In addition, whilst the menu-cost optimum might be unique, the near-rational outcomes are never unique.

When we turn 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 that are consistent with long-run survival depend on the strategic environment in which the firm finds itself and may have no obvious interpretation. The behavior of firms may be consistent and rule-following; but the rule cannot be understood as resulting from maximizing the firm’s objectives except in the context of the interaction of the firm with the wider market and economy as a whole.

In that context the behavior of the firm can be understood only by explaining how the decision rule arises—not as the solution (exact or approximate) to an optimization problem.

Acknowledgements I would like to thank Victor Tremblay and Mo Xian and Larry White for the insightful comments and advice. Faults remain my own.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Akerlof, G., & Yellen, J. (1985a). Can small deviations from rationality make significant differences to economic equilibria? *American Economic Review*, 75, 708–720.
- Akerlof, G., & 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., & Lippi, F. (2014). Price setting with menu costs for multi-product firms. *Econometrica*, 82, 89–135.
- Barnard, C. (1938). *The functions of the executive*. Cambridge: Harvard University Press.
- Bendor, J., Mookherjee, D., & Ray, D. (2001). Reinforcement learning in repeated interaction games. *The B.E. Journal of Theoretical Economics, De Gruyter*, 1(1), 1–44.
- Conlisk, J. (1996). Bounded rationality and market fluctuations. *Journal of Economic Behavior and Organisation*, 29, 233–250.
- Cyert, R., & Marsh, J. (1963). *A behavioral theory of the firm*. Hoboken: Wiley-Blackwell.
- Dixon, H. (1992). Some thoughts on artificial intelligence and economic theory. In S. J. Moss & J. Rae (Eds.), *Artificial intelligence economic analysis* (pp. 131–154). Cheltenham: Edward Elgar.

- Dixon, H. (2000). Keeping up with the Joneses: competition and the evolution of collusion. *Journal of Economic Behavior and Organisation*, 43, 223–238.
- Dixon, H. (2001). *Surfing economics: Essays for the inquiring economist*. London: Palgrave.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351–401.
- Hicks, J. (1935). Annual survey of economic theory: The theory of monopoly. *Econometrica*, 3, 1–20.
- Klenow, P., & Malin, B. (2010). Microeconomic evidence on price-setting. In B. M. Friedman & M. Woodford (Eds.), *Handbook of monetary economics. Chapter 5* (Vol. 3, pp. 231–284). Dublin: Elsevier.
- Mankiw, N. G. (1985). Small menu costs and large business cycles: A macroeconomic model of monopoly. *The Quarterly Journal of Economics*, 100(2), 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.
- Sheshinski, E., & Weiss, Y. (1977). Inflation and costs of price adjustment. *Review of Economic Studies*, 44, 287–303.
- Simon, H. (1947). *Administrative behavior, A study of decision-making processes in administrative organization*. Mumbai: Free Press.
- Simon, H. (1982). *Models of bounded rationality* (Vol. 1 and 2). Cambridge: MIT Press.
- Vickers, J. (1985). Delegation and the theory of the firm. *Economic Journal*, 95, 138–147.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.