Innovate or imitate?
Behavioural Technological Change*

Cars Hommes[a,b†] and Paolo Zeppini[a,c ‡]

[a] *CeNDEF, Amsterdam School of Economics, University of Amsterdam*

[b] *Tinbergen Institute*

[c] *School of Innovation Sciences, Eindhoven University of Technology*

JEL classification: C62, C73, D21, O33.

Key words: discrete choice, innovation patterns, learning curves, switching behaviour.

Abstract

We propose a behavioural model of technological change with evolutionary switching between boundedly rational costly innovators and free imitators, and study the endogenous interplay of innovation decisions, market price dynamics and technological progress. Innovation and imitation are strategic substitutes and exhibit negative feedback. Endogenous technological change is the cumulative outcome of innovation decisions. There are three scenarios: market breakdown, Schumpeterian rents and learning curves. The latter is characterised by an increasing fraction of innovators when demand is elastic, while inelastic demand allows technological progress with shrinking innovation effort. Model simulations are compared to empirical data of two industrial sectors.

*We are grateful to Jeroen van den Bergh and Koen Frenken for many discussions. We appreciate comments and suggestions from Giovanni Dosi and Luigi Marengo. Previous versions of this paper were presented at the DIMETIC School in Strasbourg 2009, the Tinbergen Institute Amsterdam, the Amsterdam School of Economics, the Innovation Studies group of the University of Utrecht, the 15^{th} and the 17^{th} editions of the Conference for Computational Economics and Finance in Sydney and in San Francisco, the KITeS centre of Bocconi University in Milano, the NAKE Research Day 2010 in Utrecht, the 11^{th} Workshop on Optimal Control, Dynamics Games and Nonlinear Dynamics in Amsterdam, the WEHIA 2010 conference in Alessandria, the 25^{th} EEA congress in Glasgow, the 65^{th} ESEM congress in Oslo and the 27^{th} EEA congress in Malaga.

†email: C.H.Hommes@uva.nl

‡Email: P.Zeppini@tue.nl (corresponding author)

1 Introduction

In this article we investigate the dynamics of innovation and imitation as two market strategies that affect total factor productivity in a perfectly competitive market, using a discrete choice mechanism. Our main focus is the interplay between market and behaviours, and its effects on innovation intensity and technological progress.

The approach proposed here is different and complementary to the Endogenous Technological Change literature (Romer, 1990; Grossman and Helpman, 1991). In particular, it is related to the Schumpeterian Growth type of models (Aghion and Howitt, 1992, 1998). Instead of a production function with expanding products variety, or input goods variety, we use a market dynamics with Walrasian equilibrium of (homogeneous) demand and (heterogeneous) supply, with only one homogenous good but with differentiated production technology. This way of modelling the price effect on technological change distinguishes our model also from the recent theory on Directed Technical Change (Acemoglu, 2002, 2007).

The empirical evidence shows a substantial unexplained inter-firm and intra-sectoral variability of innovation proxies as $R\&D$ expenditure, innovative output, patenting activity, etc. (Dosi, 1988). This indicates that firms' heterogeneity regarding innovation behaviour may be important in modelling technological change. Technology is a non-rival partially excludable good (Romer, 1990), which makes direct imitation possible. In some cases intellectual property rights pose a limit to imitation. Benoit (1985) addresses non-patentable innovations and studies the interplay of innovators and imitators in the strategic setting of a duopoly. With our model we adopt an adaptive behavioural approach, as, for instance, in Arthur (1989), and consider a population of firms where innovation and imitation are two alternative strategies.

Our model of technological change is based on switching behaviour of costly innovators and cheap imitators. The interplay of innovation and imitation plays an important role in the dynamics of industry evolution, in particular affecting the incentives for costs reduction effort (Ceccagnoli, 2005). Imitation in a broad sense is the exploitation of external knowledge sources. This can involve public knowledge such as published research but also spillovers and leakages from private knowledge (Spence, 1984). Considering the taxonomy of Malerba (1992), innovation and imitation refer to *learning by searching* and *learning from spillovers*. In the latter there are all different kinds of information flows, from knowledge leakages to pure copying activity. Modelling innovation and imitation as two different strategies relies on Schumpeter's hypothesis of *routinisation* of innovation (Schumpeter, 1942), and more generally on Simon's view about bounded rationality of agents (Simon, 1957).

Behavioural heterogeneity and switching behaviour are empirically relevant in other applications, as testified by survey data (Branch, 2004), market data (de Jong et al., 2009) and laboratory experiments Hommes (2011). And although the literature on heterogenous agents models is now quite vast, little has been tried in this direction to model technological change. We intend to cover this gap.

We model behavioural diversity and switching behaviour using the discrete choice framework of Brock and Hommes (1997). Our model addresses interacting firms that make a choice about whether or not to invest in innovation in order to be more productive. The idea of imitation as a cheap heuristic opposed to a costly sophisticated innovative strategy is similar in spirit to Grossman and Stiglitz (1976)'s model of informed and uninformed agents in a competitive asset market. In our model this idea can be expressed by saying that it may be more efficient for some firms to exploit other firms than to invest in innovation themselves. Because of these different elements, our model of innovation combines the approach of neoclassical economics with the evolutionary-economic approach of dynamic heterogeneous populations.

The literature on innovations diffusion has addressed the role of imitation. Two seminal papers here are Mansfield (1961) and Bass (1969). Those models mainly look at the demand side, and the focus are the timing of adoption and diffusion rates. Our evolutionary selection based on production cost reduction shares some elements with Iwai (1984), where firms are described by a distribution of production costs. In our model behavioural heterogeneity leads to a negative feedback that makes it profitable to switch strategy in an environment where a strategy becomes dominant. This feature may be interpreted as a minority game, and finds a parallel in models with strategic complements and substitutes (Bulow et al., 1985). Conlisk (1980) has a negative feedback with costly optimisers and cheap imitators. An important difference of our model is the endogenous interplay of market and firms' choices, without exogenous stochastic process. Another endogenous model of interacting sophisticated and naive agents is Sethi and Franke (1995). However, this model and Conlisk's model are globally stable: if not for exogenous random shocks, the economy would converge to an equilibrium where all agents use the cheap strategy. In our model there may be a stable equilibrium with coexistence of strategies, or even cyclical or chaotic dynamics without any exogenous shocks.

We have two aims in this paper. The first is to study how the behavioural decision mechanism and market dynamics interact, and what are the factors that make one strategy, innovation or imitation, prevail. The second is to address the mutual effects of behaviours and technological change, to see how different innovation patterns endogenously depend on market factors and behavioural regimes. In a first basic version of the model we focus on equilibrium stability and on the main factors driving market and strategy dynamics. In a more elaborated version of the model we focus on technological change and innovation behavioural regimes.

The model with endogenous technological change presents three scenarios: *market breakdown*, where depreciation of technology is too strong compared to knowledge cumulation. This is the story of shrinking sectors. *Balanced technological change*, where technological growth is just enough to offset depreciation. The price decreases but sets to a positive limit and the technological frontier is limited. The third scenario is *technological progress*, with a price falling to zero and a technological frontier that grows unboundedly.

The scenario with technological progress represents the main and final focus of this paper. Here we show two main results. First, the key-role of demand elasticity in explaining innovation patterns. Second, the ability of the model to reproduce learning curves. The role of demand in technological change has been widely overlooked in the literature. Our model shows that when the demand is elastic, technological progress leads to an ever increasing fraction of innovators. With inelastic demand, technological progress is characterised by less and less innovators, instead. These two different outcomes are much in line with the patterns of innovation of Schumpeterian tradition: the Schumpeterian Mark I pattern, that is referred to as *widening*, is characterised by a decreasing concentration of patenting firms, and is obtained with elastic demand. The Schumpeterian Mark II pattern, referred to as *deepening*, is the opposite, and in our model it realises with an inelastic demand. This explanation of innovation patters complements the technological regimes explanation of Breschi et al. (2000), and advocates the potential of a behavioural approach to endogenous technological change.

The second result of the scenario with technological progress is a behavioural micro-foundation of learning curves. These are a stylised fact of technological change (Hartley, 1965; Lieberman, 1984; Argote and Epple, 1990). The empirical literature on learning curves is vast (Berndt, 1991), but on the other hand models that include this factor in their analysis are only a few, and usually devoted to study the implications of learning curves for pricing, market equilibrium and social welfare (Spence, 1981; Cabral and Riordan, 1994; Petrakis et al., 1997). A common feature of these models is that learning curves are taken as exogenous. McCabe (1996) makes learning curves endogenous in a learning models based on a principal-agent approach. Our model constitutes an alternative endogenous explanation of learning curves that is based on the interplay between agents decision making and market dynamics. Simulated time series of price and production are compared to empirical data from two different industrial sectors, the US tire industry and a global index of solar power technology.

The chapter is organised as follows. Section 2 introduces the general framework and presents a basic model, with a stability analysis of market dynamics. Section 3 describes the full model with behavioural technological change, presenting the different innovation patterns and the setting that reproduces learning curves. Section 4 concludes.

2 Costly innovators versus cheap imitators

2.1 The basic model

Consider an industry with N firms producing the same good in a perfectly competitive market. *Innovation* means to reduce the production cost, while *imitation* means to adopt the currently available technology. Firms are either *innovators*, with fraction n_t, or *imitators*, with fraction $1 - n_t$. Choosing the strategy (innovation or imitation) sets the production technology and the cost structure or total factor productivity (TFP) of a firm. The quantity $S^h(p_t)$ supplied in period t by a firm choosing strategy h is a function of price and depends on the cost structure of strategy h. In each period the market clears in a Walrasian equilibrium:

$$D(p_t) = n_t S_t^{INN}(p_t) + (1 - n_t) S_t^{IM}(p_t), \tag{1}$$

where $h = INN$ stands for innovation, and $h = IM$ for imitation. Eq. (1) results from the aggregation of demand over consumers and supply over firms, and then dividing by the total number of firms N.[1] The supply is a convex combination of innovators' and imitators' production, with n_t and $1 - n_t$ the fractions of innovators and imitators, respectively. Profits of an individual firm of type h in period t are $\pi_t^h = p_t q_t^h - c^h(q_t^h)$, with $q_t^h \equiv S_t^h(p_t)$. We choose a quadratic cost function as in Jovanovic and MacDonald (1994): the cost of producing quantity q for a firm adopting strategy h is $c^h(q) = \frac{q^2}{2s^h} + C^h$, where C^h represents the fixed costs of the strategy. This choice keeps the model as simple as possible: maximisation of profits with respect to quantity q gives a linear supply:

$$S_t^{INN}(p_t) = s_t^{INN} p_t, \qquad S_t^{IM}(p_t) = s_t^{IM} p_t. \tag{2}$$

The parameters s_t^{INN} and s_t^{IM} are proportional to TFP, and consequently depend on the production technology of the firm.[2] An innovator invests $C^{INN} = C > 0$ and increases TFP, expressed by $s^{INN} > s^{IM}$, cutting down the production cost $c(q)$ (see Jovanovic and MacDonald (1994)). Cost reduction is larger for larger values of output: $\Delta c = -\frac{q^2}{2s^2}\Delta s$. This means that larger firms profit more from innovation. Imitation is free ($C^{IM} = 0$) and it amounts to using the state-of-the-art technology, a sort of publicly available technological frontier.[3] This setting is similar to Iwai (1984), the difference being that here we have

[1] Aggregation of supply gives $S_t = \sum_{i=1}^{N_t^{INN}} S_{i,t}^{INN} + \sum_{j=1}^{N_t^{IM}} S_{j,t}^{IM}$. Subgroups of innovators (imitators) are homogeneous i.e. $S_{i,t}^{INN} = S_t^{INN}$ ($S_{j,t}^{IM} = S_t^{IM}$) for all i (j). Hence $S_t = N_t^{INN} S_t^{INN} + N_t^{IM} S_t^{IM}$. Dividing by the number of firms N one gets the right-hand side of (1).

[2] If we think in terms of a production function like $q = A\phi(K, L)$, where ϕ is a function of capital and labour, the parameter s is positively related to the production technology factor A.

[3] In principle imitators have the advantage of not replicating an unsuccessful innovation. Here we assume that innovation is always successful. One can also interpret the model in a slightly different way,

two types of firms instead of a continuous distribution. If we focus on TFP, our model resembles the model of competition driven by $R\&D$ in Spence (1984), provided that time is discrete and firms are homogeneous but for their choice about innovation, as in Llerena and Oltra (2002). The competitive advantage of innovators over imitators is expressed by specifying the production cost structure. Assume TFP's of innovators and imitators do not depend on time, and $R\&D$ expenditure enhances the TFP of innovators by an exponential factor (Nelson and Winter, 1982; Dosi et al., 2005): $s_t^{INN} = se^{bC}$ and $s_t^{IM} = s$, where $b > 0$ represents the benefits of the innovation investment. It follows that marginal production costs are $c'(q) = \frac{q}{s}$ for imitators and $c'(q) = \frac{q}{se^{bC}}$ for innovators. Average costs are $\gamma^{INN} \equiv \frac{c^{INN}(q)}{q} = \frac{p}{2} + \frac{C}{s^{INN}p}$ and $\gamma^{IM} = \frac{p}{2}$, with $\gamma^{INN} \geq \gamma^{IM}$ and $\gamma^{INN} = \gamma^{IM}$ in the limit of infinite price. This is an indication that innovators benefit from a high price, although their aggregate effect is exactly in the opposite direction, i.e. more innovators lower the price.

Firms switch between innovation and imitation based on the evaluation of profits. For a quadratic cost function, profits of innovation and imitation are:

$$\begin{aligned} \pi_t^{INN} &= \frac{1}{2}s_t^{INN}p_t^2 - C = \frac{1}{2}se^{bC}p_t^2 - C, \\ \pi_t^{IM} &= \frac{1}{2}s_t^{IM}p_t^2 = \frac{1}{2}sp_t^2. \end{aligned} \tag{3}$$

In particular $\Delta\pi \equiv \pi^{INN} - \pi^{IM} = 0$ for $p = \overline{p} \equiv \sqrt{2C/s(e^{bC} - 1)}$. We model agents' decision using the discrete choice framework of Brock and Hommes (1997) (BH henceforth), with an endogenous evolutionary selection between costly innovation and cheap imitation. This framework is based on the concept of random utility (see Hommes (2006) for an extensive survey and discussion). The fraction of innovators at time t is given as:

$$n_t = \frac{e^{\beta\pi_{t-1}^{INN}}}{e^{\beta\pi_{t-1}^{INN}} + e^{\beta\pi_{t-1}^{IM}}}. \tag{4}$$

If we use the difference of profits $\Delta\pi_t \equiv \pi_t^{INN} - \pi_t^{IM} = \frac{1}{2}s(e^{bC} - 1)p_t^2 - C$, we obtain the following function $n_t = \hat{g}(p_{t-1})$:

$$n_t = \frac{1}{1 + e^{-\beta\left[\frac{1}{2}s(e^{bC}-1)p_{t-1}^2 - C\right]}} \equiv \hat{g}(p_{t-1}). \tag{5}$$

A higher price creates incentives to innovate, because of a larger $\Delta\pi$. The intensity of choice β is inversely proportional to the variance of the utility noise, and measures the ability of firms to choose the best strategy. In the limit $\beta = 0$ agents are myopic, and split

thinking that innovation is an uncertain event, and that innovators improve their productivity with a given (exogenous) probability. Say that S^{INN} is the expected value of productivity from this innovation process. With a large number of identical innovating agents, everything goes as if all innovating agents are given the improved productivity S^{INN}

equally among the different types ($n = 1/2$). On the other hand, $\beta = \infty$ represents the rational limit, where all agents choose the optimal strategy.

In this basic specification of the model we ignore technological advance and focus on the interplay between strategy switching behaviour and market dynamics. We assume that innovation is like buying a shortcut which results in lower production costs in one period. A similar assumption is in Aghion et al. (2005), where profits depend only on the gap between leading and laggard firms, and not on the absolute level of technology. Section 3 relaxes this hypothesis, and considers technological progress.

Consider a hyperbolic demand $D(p_t) = \frac{a}{p_t^d}$, with price elasticity equal to $-d$ ($d > 0$).[4] Solving the market equilibrium equation (1) with $s_t^{INN} = se^{bC}$ and $s_t^{IM} = s$ we get

$$p_t = \left\{ \frac{a}{s\left[(e^{bC} - 1)n_t + 1\right]} \right\}^{\frac{1}{1+d}} \equiv \hat{f}(n_t), \qquad (6)$$

where fractions n_t and $1 - n_t$ depend on last period price according to (4). The function $\hat{f}(n)$ is decreasing because $e^{bC} > 1$: *an increase in the density of innovators drives down the price.* When everybody innovates the price reaches its minimum value $p^*_{INN} = \left(\frac{a}{se^{bC}}\right)^{\frac{1}{1+d}}$. On the other hand, the maximum value $p^*_{IM} = \left(\frac{a}{s}\right)^{\frac{1}{1+d}}$ is obtained with only imitators,[5] as illustrated in Fig. 1. The more innovators, the steeper is the aggregate supply curve and

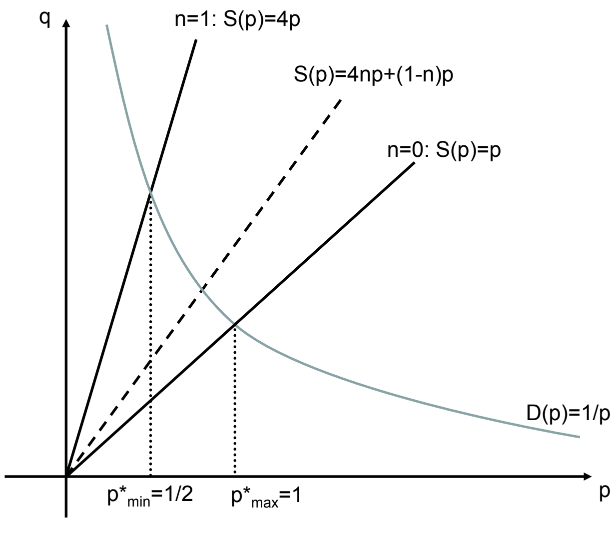

Figure 1: Demand and supply curves with $D(p) = 1/p$, $s^{INN} = 4$ and $s^{IM} = 1$.

the lower is the price. The intuition behind this mechanism is that innovation is defined as cost reduction, so that a positive mass of innovators lower the average production cost of the industry, which translates into a lower market price.

The decision mechanism (5) and the market mechanism (6) express a negative relationship between price p and innovation n. These two opposing forces feed the dynamic

[4] In an earlier version of the model we have considered a linear demand (Zeppini, 2011).

[5] We can think of this limit as a situation with only one innovator: If $N \gg 1$ we have $n \simeq 0$.

equilibrium (1). There are conditions for a stable equilibrium, where fractions and price remain unchanged through time. The system under study is one-dimensional, and the equilibrium can be found either using the price p_t or the innovators fraction n_t as state variable. By substituting Eq. (5) into (6) we obtain a flow map for the price:

$$p_t = \left(\frac{a}{s}\right)^{\frac{1}{1+d}} \left\{ \frac{1 + e^{\beta[\frac{1}{2}s(e^{bC}-1)p_{t-1}^2 - C]}}{1 + e^{\beta[\frac{1}{2}s(e^{bC}-1)p_{t-1}^2 - C] + bC}} \right\}^{\frac{1}{1+d}} \equiv f(p_{t-1}). \tag{7}$$

If instead we substitute (6) into (5), we obtain a map for the fraction of innovators:

$$n_t = \frac{1}{1 + e^{-\beta\left\{\frac{1}{2}s^{\frac{d-1}{d+1}}(e^{bC}-1)\left[\frac{a}{n_{t-1}(e^{bC}-1)+1}\right]^{\frac{2}{1+d}} - C\right\}}} \equiv g(n_{t-1}). \tag{8}$$

In Eq. (8) the factor s does not play any role when the demand is unit elastic ($d = 1$). This fact is important when we introduce endogenous technological progress (Section 3).

2.2 Steady states and stability

An equilibrium is expressed by a fixed point of function f (or g), that is a value of the price p^* such that $p^* = f(p^*)$ (or $n^* = g(n^*)$).

Proposition 2.1 *There is a unique steady steady state p^* (or n^*).*

This is because the map f (or g) is monotonically decreasing (Appendix A). The stability of p^* depends on the parameters setting:

Proposition 2.2 *p^* (or n^*) is stable in the limit $\omega \to 0$ for $\omega = a, b, C, s, \beta$.*

The proof is given in Appendix A.

The intensity of choice β expresses the extent to which agents make the optimal decision between innovation and imitation (Eq. 4). In the limit $\beta = \infty$ the price map is a step function. Consider the price $\overline{p}$ where imitators and innovators have the same profit, $\overline{p} = \sqrt{2C/[s(e^{bC} - 1)]}$. Whenever $p > \overline{p}$, it holds $\Delta\pi > 0$, and $\beta = \infty$ in Eq. (5) gives $n_t = 1$. This means that $f(p) = p^*_{INN}$ $\forall p$, by Eq. (6). On the contrary, for $p < \overline{p}$, $\beta = \infty$ gives $f(p) = p^*_{IM}$ $\forall p$, with $n_t = 0$. Consequently, the price map (7) is a decreasing step, with a discontinuity at $p = \overline{p}$.

For finite values of β different situations may occur. The first derivative of the map (8) at the equilibrium n^* is

$$g'(n^*) = -n^*[1 - n^*] \frac{\beta s^{\frac{d-1}{d+1}} a^2 (e^{bC} - 1)^2}{\left[(e^{bC} - 1)n^* + 1\right]^{\frac{d+3}{d+1}}}. \tag{9}$$

The stability condition $-1 < g'(n^*) < 0$ is satisfied in particular when there is a sufficiently large prevalence of innovators ($n^* \simeq 1$) or imitators ($n^* \simeq 0$).

The qualitative change from stable equilibrium to period 2 cycle is a *period-doubling bifurcation*. This may occur by changing any of the parameters a, d, b, C, s or β. Although an analytic computation of bifurcation values is not feasible, Prop. 2.2 summarises how the stability of the steady state depends on parameters. Changes in the demand parameters a and d only affect the price range defined by p^*_{INN} and p^*_{IM}, and leave $\overline{p}$ unaffected. An increase of a (positive demand shock) moves the demand curve outwards (Fig. 1), and enlarges the gap $p^*_{IM} - p^*_{INN}$. This change is destabilizing (Zeppini, 2011). An increase of d (price elasticity of demand) reduces the gap $p^*_{IM} - p^*_{INN}$, and tend to be stabilizing instead. The supply parameters s, b and C affect both $\overline{p}$ and the range $[p^*_{IM}, p^*_{INN}]$. Their effect on the equilibrium is not obvious, then.

If the map is steep enough in the fixed point and $|g'(p^*)| > 1$, the market does not attain a stable equilibrium. Since g is decreasing and bounded, when the equilibrium is unstable a (stable) 2-cycle occurs. In Fig. 2 we report two examples of time series of the innovators fraction n_t (upper panels). On the left we have a case where the market converges to a

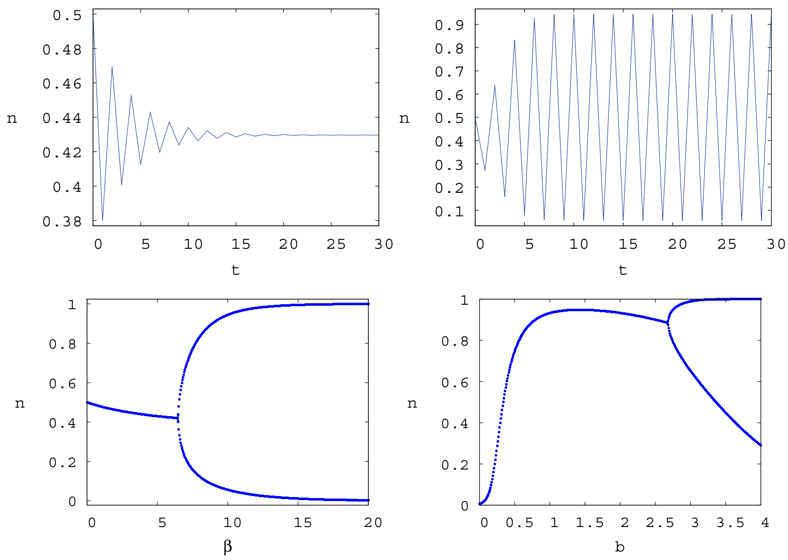

Figure 2: Upper panels: examples of time series of innovators fraction n. Left: Stable equilibrium ($\beta = 5$). Right: 2-cycle ($\beta = 10$). Here $a = d = 2$, $b = C = s = 1$. Lower panels: long run values of n. Left: bifurcation diagram of the intensity of choice β ($b = 1$, $a = 2$, $d = 2$). Right: bifurcation diagram of the innovation benefit b ($\beta = 5$, $a = 4$, $d = 0.5$). Here $s = C = 1$.

stable equilibrium $n^* \simeq 0.43$. On the right we have a stable 2-cycle, obtained increasing the intensity of choice from $\beta = 5$ to $\beta = 10$.

The intuition for cyclical dynamics is as follows. Innovation drives down the price, and at some point the profits from innovation become too low (even negative, due to the fixed costs C), so that imitation becomes preferable. Agents start switching to imitative behaviour then, and the price goes up. An increasing price boosts innovators' profits more than imitators', because of larger TFP. When innovators profits become largest, agents switch back to innovation again, and the story repeats. This cyclical behaviour reflects a "minority game" dynamics, in that innovation and imitation show strategic substitutability (Bulow et al., 1985): a strategy adopted by the minority is more appealing. Stated differently: innovation works better in a market dominated by imitators, while imitation is more profitable in an environment dominated by innovators. Hence, there is a negative feedback from strategy adoption. Such a negative feedback mechanism resembles the dynamic counterpart of the inverted-U relationship between competition and innovation studied in Aghion et al. (2005): a fall of the price means stronger competition and it is associated with a surge in innovation, but at the same time it creates incentives for imitation, and innovation slows down.

A *bifurcation diagram* shows the qualitative changes in the dynamics due to changes in parameters.[6] The lower-left panel of Fig. 2 reports an example of bifurcation diagram of the intensity of choice. For $\beta \simeq 7$ the steady state loses stability and a stable 2-cycle occurs. As β gets larger, the cycle approaches values $\{0, 1\}$, meaning that almost all agents switch between innovation and imitation. The lower-right panel of Fig. 2 reports a bifurcation diagram of innovation return b, with a period doubling at $b \simeq 2.7$. This diagram shows a trade-off in b: larger innovation benefits do not necessarily mean more innovation. The effect of b on n is positive for small values of b, but negative for large values. This is due to a double effect of innovation on innovators' profits: a positive direct effect comes from the exponential factor that makes profitability larger, $s^{INN} = se^{bC}$. A negative indirect effect is from the price: the price reduction of innovation has a stronger effect on innovators themselves, because of their larger productivity, which also means a higher price elasticity of supply. If the price effect is prevailing, innovators become less frequent as b gets larger.[7]

2.3 Asynchronous updating of strategies and chaos

So far we have assumed that in each period all agents evaluate the payoff from innovation and imitation, and switch to the optimal strategy with a probability that depends on the intensity of choice β. This may not be realistic. Firms show a good degree of persistence (Dosi, 1988), and the empirical evidence of persistence in firms' propensity to innovate or not-innovate holds across countries and industrial sectors (Cefis and Orsenigo, 2001). It is therefore useful to introduce a hypothesis of inertia, as in evolutionary game theory learning models (Kandori et al., 1993). In discrete choice models this is implemented through *asynchronous updating* (Diks and van der Weide, 2005; Hommes et al., 2005): in

[6] A bifurcation diagram is obtained as a collection of long run values of the state variable for a set of different initial conditions and a specified range of the parameter under study. Here the transient time is 100 periods. The numerical implementation has been done with *E&F Chaos* (Diks et al., 2008).

[7] Beside period doubling bifurcations, also period halving is possible, where increasing one parameter moves the market from 2-cycles to stable equilibrium. The joint effect of any two parameters can be analysed with a phase diagram. A detailed analysis of different bifurcation scenarios is in Zeppini (2011).

each period the fraction $1-\alpha$ ($\alpha \in [0,1]$) of agents update strategy, while the rest stick to the previous strategy. Consequently, the fraction of innovators is as follows:

$$\begin{aligned} n_t &= \alpha n_{t-1} + (1-\alpha)\frac{e^{\beta\pi_{t-1}^{INN}}}{e^{\beta\pi_{t-1}^{INN}} + e^{\beta\pi_{t-1}^{IM}}} \\ &= \alpha n_{t-1} + (1-\alpha) g(n_{t-1}) \equiv \hat{g}(n_{t-1}), \end{aligned} \tag{10}$$

where the function g is the map (8) of the basic model with synchronous updating (that we obtain with $\alpha = 0$). This system is still one-dimensional. The map $\hat{g}$ in (10) is a convex combination of an increasing function, n_{t-1}, and a decreasing function, $g(n_{t-1})$, and therefore can be non-monotonic depending on the value of α (Fig. 3, upper-left panel). In

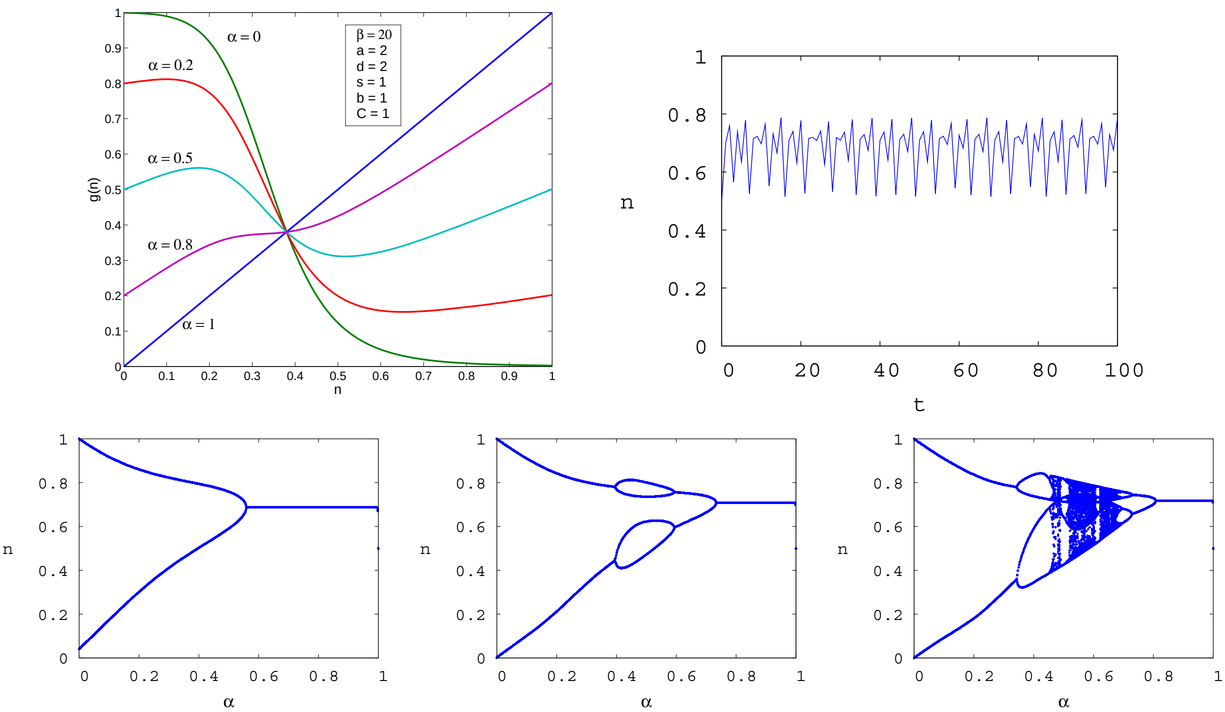

Figure 3: The model with asynchronous updating. Upper-left: Examples of the map $\hat{g}$ for different values of the weight α. Upper-right: time series n_t ($\alpha = 0.6$, $\beta = 30$, $a = 4$, $d = C = 2$, $b = 1$, $s = 0.5$). Lower panels: bifurcation diagrams of the updating weight α. Left: $\beta = 10$. Centre: $\beta = 20$. Right: $\beta = 30$ (with $a = 4$, $d = 2$, $b = 1$, $s = 0.5$, $C = 2$).

particular, $\hat{g}$ is decreasing for $\alpha = 0$, it becomes non-monotonic for intermediate values of α and it is increasing for α close to 1. The non-monotonicity of the map $\hat{g}$ leads to complicated dynamics when the steady state is unstable (Fig. 3, upper-right panel). Indeed, chaotic dynamics can arise, as illustrated in the bifurcation diagrams of Fig. 3 (lower panels). When β is relatively small (lower-left panel), either a 2-cycle or a stable equilibrium are possible. Increasing β, cycles of period 4 appear for mid values of α (lower-middle panel). A larger β further destabilises the market introducing irregular dynamics for $\alpha > 0.5$ (lower-right panel). These examples indicate that in general, when most agents stick to their strategy (large α), the industry converges to a stable equilibrium. When only a small fraction of agents update strategy (low α) instead, the market converges to a period 2-cycle. Intermediate values of the updating fraction α may present a period doubling bifurcation route to irregular chaotic dynamics. Nevertheless, the variability of n decreases with a larger α. This means that asynchronous updating is quantitatively stabilizing, but qualitatively destabilizing: it dampens the amplitude of the orbit oscillations, but at the same time chaos may occur. This global dynamics is similar to the cobweb model with adaptive expectations of Hommes (1994), with the asynchronous updating fraction α playing the role of the adaptive expectations weight factor. Prop. 2.3 shows the occurrence of chaos with asynchronous strategy updating.

Proposition 2.3 *Let $\hat{g}$ be the map (10). If β and C are sufficiently large, there exist values α_1, α_2 and α_3 with $0 < \alpha_1 < \alpha_2 < \alpha_3 < 1$ such that the following holds true:*

- *(A1) $\hat{g}$ has a stable period 2 orbit for $\alpha \in [0, \alpha_1)$,*
- *(A2) the map $\hat{g}$ is chaotic in some interval $[\alpha_2 - \epsilon, \alpha_2 + \epsilon]$,*
- *(A3) $\hat{g}$ has a stable equilibrium for $\alpha \in (\alpha_3, 1]$.*

A proof is given in Appendix B. Asynchronous updating increases persistence of strategies. The time series in the upper-right panel of Fig. 3 is an example where oscillations of innovators fraction are strongly and irregularly dampened in several periods.

3 Technological change

In the previous sections we have studied the dynamics of the interplay between innovation and imitation assuming that strategy switching and price dynamics do not interfere with the underlying technological progress. In this section we study the mutual effects of *technological progress* and strategy switching, proposing a behavioural model of technological change. The closest reference to this model is the "Schumpeterian" version of endogenous growth theory (Aghion and Howitt, 1992, 1998). There are two main differences in our model: first, we have behavioural heterogeneity of firms, leading to a differentiated production cost, in place of quality ladder of technology vintages. Second, we rely on the market dynamics of supply and demand, and not on the concept of production function and factors prices.

3.1 The model

The fundamental assumption of this extension of the model is that innovation cumulates: in each period the achievements of innovators contribute to a technological frontier. The frontier consists of all past innovations, and has the connotation of a learning curve. Imitators have access to the technological frontier, while innovators expand it, obtaining a better production technology due to their innovation investment. We introduce a cumulation rate γ for innovations, and also a depreciation rate δ. Based on this we define the technological frontier:

$$s(t) = se^{\sum_{i=1}^{t-1}[\gamma n_i - \delta]}. \tag{11}$$

This technological frontier grows over time exponentially by a time-varying factor $\gamma n_i - \delta$, where n_i is the fraction of innovators in period i. Imitators exploit the frontier technology, while innovators build on it, getting a competitive advantage. Consequently the two productivity levels are as follows:

$$s_t^{INN} = s(t)e^{bC}, \qquad s_t^{IM} = s(t). \tag{12}$$

Formally nothing changes with respect to the basic model: innovators increase TFP by the factor e^{bC}, after investing C in innovation. This advantage lasts one period, because it becomes publicly available afterwards. The difference with the basic model is that *innovation now exhibits endogenous growth* and cumulates at a rate γ, resulting in an advancing technological frontier. An agent can innovate today and imitate tomorrow, without loosing the benefits from its previous innovation, although everybody else can use it as well. It has to be noted that a dynamic technological frontier $s(t)$ makes the technological gap $\Delta s(t) = s(t)(e^{bC} - 1)$ change over time. In particular, technological progress causes Δs to enlarge.

The rate γ measures two effects, namely cumulativeness of knowledge and spillovers of technological innovations. The implicit assumptions here are that innovation always cumulates and spills over at the same rates, in line with the assumption of our model that innovation benefits b and costs C are the same in every period.

Let's consider synchronous updating ($\alpha = 0$) for the moment. The introduction of a technological frontier in the basic model of Section 2 amounts to substitute parameter s with $s(t)$ in the distribution of agents' fractions (5) and in the market equilibrium equation (6), which become, respectively,

$$n_t = \frac{1}{1 + e^{-\beta\left[\frac{1}{2}s(t)(e^{bC}-1)p_{t-1}^2 - C\right]}}, \tag{13}$$

$$p_t = \left\{\frac{a}{s(t)\left[(e^{bC} - 1)n_t + 1\right]}\right\}^{\frac{1}{1+d}}. \tag{14}$$

By substituting (14) into (13) we obtain a new map of the market system:

$$n_t = G(n_{t-1}; s(t)) \equiv \frac{1}{1 + e^{-\beta\left\{\frac{1}{2}s(t)^{\frac{d-1}{d+1}}(e^{bC}-1)\left[\frac{a}{n_{t-1}(e^{bC}-1)+1}\right]^{\frac{2}{1+d}} - C\right\}}}. \tag{15}$$

Similarly, we obtain a map for the price $F(p_t; s(t))$ by substituting (13) into (14). The technological frontier $s(t)$ works as a "slowly changing parameter" that spans the technology dimension of the model. Fig. 4 illustrates how the map G evolves due to changes in $s(t)$. The effect of technological change strongly depends on the elasticity of demand:

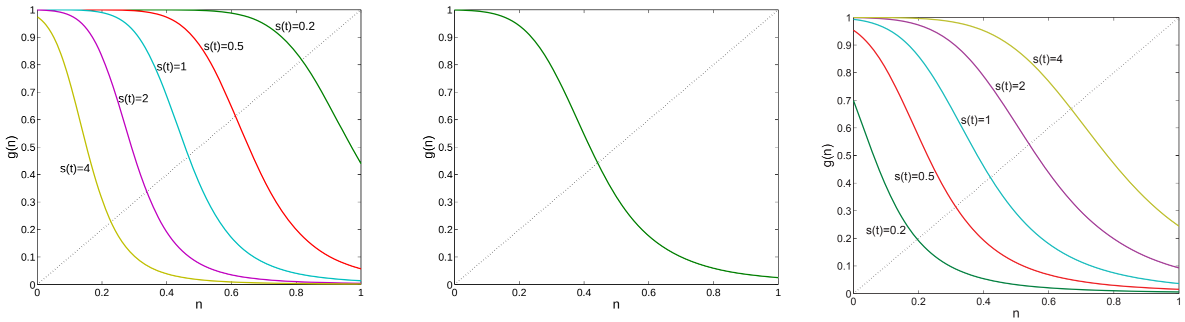

Figure 4: Graph of the map of innovators fraction with technological change $G(x; s(t))$. Left: inelastic demand (here $d = 0.5$). Centre: unit elastic demand ($d = 1$). Right: elastic demand (here $d = 1.5$). Other parameters are $\beta = 10$, $b = C = 1$, and $a = 2$.

Proposition 3.1 *Consider the market of innovators and imitators with technological change, represented by Eq. (15), and assume technological progress ($s'(t) > 0$):*

1. *for inelastic demand ($d < 1$), technological progress goes with less innovators n^*,*

2. *for unit elastic demand ($d = 1$) technological progress does not affect the market,*

3. *for elastic demand ($d > 1$) technological progress goes with more innovators.*

A proof is in Appendix C. When the demand is elastic, technological progress leads to an ever increasing fraction of innovators. With inelastic demand, technological progress is characterised by less and less innovators, instead. The intuition is based on the differentiated price elasticity of supply, which is larger for innovators. A price reduction hurts innovators more than imitators (see Section 2.2), but at the same time innovation increases the quantity exchanged in equilibrium, which rewards innovators more than imitators. When the demand is elastic, the second effect overcomes the first, because the marginal increase in exchanged quantity from a price reduction is relatively larger. The opposite is true with inelastic demand, while the two effects offset each other when the demand is unit elastic.[8]

These two different conditions substantially match the patterns of innovation of Schumpeterian tradition. the Schumpeterian Mark I pattern, *widening*, which is characterised by an increasing concentration of patenting firms, is obtained with elastic demand. The Schumpeterian Mark II pattern, *deepening*, in our model realises with an inelastic demand.

[8]The statement of Proposition 3.1 is absolute in all cases of stable equilibrium n^*, while it holds on average (over time) whenever n^* is unstable and the dynamics of the system is cyclical.

This explanation of innovation patters adds to the explanation based on technological regimes that is proposed inBreschi et al. (2000).

The time evolution of the technological frontier $s(t)$ requires some analysis. Let us write $s(t)$ as follows:

$$s(t) = se^{-\delta(t-1)}e^{\gamma \sum_{i=1}^{t-1} n_i}. \tag{16}$$

The rate of change of $s(t)$ is bounded. In the long run the lower bound is $-\delta$, which is attained when innovators disappear ($n_t \to 0$). The upper bound is $\gamma - \delta$, at which all agents become innovators ($n_t \to 1$).

Depending on the value of lower and upper bounds we have a number of different scenarios, summarised by the following proposition:

Proposition 3.2 *The long run dynamics of the market with technological change (15) presents six different scenarios:*

1. *for $\gamma < \delta$: $s(t) \to 0$, $p_t \to \infty$ and $q_t \equiv D(p_t) \to 0$ (**market breakdown**).*
2. *for $\gamma = \delta$: $s(t) = se^{-\gamma \sum_1^{t-1}(1-n_i)}$ and we have two subcases:*
   (a) *if $\sum_1^{\infty}(1-n_i) \to \infty$, then $s(t) \to 0$, $p_t \to \infty$, $q_t \to 0$ (**market breakdown**).*
   (b) *if $\sum_1^{\infty}(1-n_i) \to \Sigma < \infty$, then $s(t) \to se^{-\gamma\Sigma}$ and $p \to p^* > 0$ stable or unstable (**balanced technological change**).*
3. *for $\gamma > \delta$ we have three subcases:*
   (a) *if $\gamma \sum_{i=1}^{t-1} n_i < \delta t$, then $s(t) \to 0$, $p_t \to \infty$, $q_t \to 0$ (**market breakdown**).*
   (b) *if $\gamma \sum_{i=1}^{t-1} n_i \sim \delta t$, then $\exists \Sigma \in (0, \infty)$ with $s(t) \to se^{-\gamma\Sigma}$ and $p \to p^* > 0$ stable or unstable (**balanced technological change**).*
   (c) *if $\gamma \sum_{i=1}^{t-1} n_i > \delta t$, then $s(t) \to \infty$, $p_t \to 0$, $q_t \to \infty$ (**technological progress**).*

For cases $2b$ and $3b$, the following applies:

Corollary 3.1 *Balanced technological change occurs $\Leftrightarrow n^* = \frac{\delta}{\gamma}$.*

Proofs are in Appendix D. Case 1 is trivial, because technical progress can never counter depreciation. In case 2 all boils down to the convergence of the series n_t. If innovators do not take the whole market in the long run, but some imitators are always present, then we have a net depreciation of the frontier, $s(t) \to 0$. Otherwise, if innovators conquer the market fast enough, then $s(t)$ converges to a positive value, and so does the price. Case (3) is the most realistic, but also the most uncertain, because three scenarios are possible. If the process of knowledge accumulation is not strong enough to compensate technological depreciation, a market breakdown occurs (case $3a$). This is the case if γ is only slightly larger than δ. If instead knowledge accumulation goes at a rate similar to δt (case $3b$), we are in a situation similar to case ($2b$), where depreciation and technological progress offset each other. In case ($3c$) technological accumulation is stronger than depreciation, and price and marginal cost $c'(q) = p/s$ fall down to zero. This case occurs when $\gamma \gg \delta$, for instance, and on average there are enough innovators in the history of the market. Notice that scenario $3a$ can realise with a divergent series $\sum_{i=1}^{t-1} n_i$ if δ is too large. On the other hand, scenario $3c$ can occur even with a steadily diminishing fraction of innovators $n_t \to 0$, if it is slow enough. What matters is the relative value of accumulated innovation compared to the linear depreciation δt.

The scenarios with balanced technological change ($2b$ and $3b$) can present either stable equilibrium or 2-cycles in the long run, depending on the stability of the limit value p^*. The other scenarios are less obvious. The price converges either to 0 or to ∞, but the long run value of n_t depends on two unbounded quantities, $s(t)$ and p_t (Eq. 13), which are one diverging and one converging to 0. In all cases of stable equilibrium we can simplify Prop. 3.2 in the following way:

Proposition 3.3 *Assume that the model converges to a stable equilibrium, with $n_i \to n^*$. Consider the quantity $\nu^* \equiv \gamma n^* - \delta$. Three cases are possible:*

*(i) $\nu^* < 0$, then $s(t) \sim s e^{\nu^*(t-1)} \to 0$, $p_t \to \infty$ and $q_t \to 0$ (**market breakdown**).*

*(ii) $\nu^* = 0$, then $s(t) \to s e^{-\Sigma}$, $p_t \to p^* > 0$ (**balanced technological change**).*

*(iii) $\nu^* > 0$, then $s(t) \sim s e^{\nu^*(t-1)} \to \infty$, $p_t \to 0$ (**technological progress**).*

Case (i) can occur in all three cases of Prop. 3.2. In particular it coincides with cases (1), ($2a$) and ($3a$). Case (ii) implies an equilibrium value of the innovators fraction $n^* = \frac{\delta}{\gamma} \leq 1$, and may occur in cases (2) and (3) of Prop. 3.2. Case (ii) falls in (but does not coincide with) cases ($2b$) and ($3b$) of Prop. 3.2. Finally, case (iii) implies $\gamma > \delta$ and implies case ($3c$) of Prop. 3.2.

Market breakdown concerns shrinking industrial sectors, where the accumulation of knowledge does not keep the pace of depreciation. An example are the artisan productions that enriched aristocratic residences in the past centuries. *Balanced technological change* has multiple interpretations. It describes industries where real technological progress is limited. This can be the case of consolidated industrial sectors, which have already experienced a technological progress phase, and where currently innovation is like "re-novation". Alternatively, this scenario reproduces the so-called "Schumpeterian rents", where a rent is earned by the innovator in the period following innovation, before imitation occur, and further innovation is just enough to compensate for depreciation. Notice how in this scenario the higher the depreciation rate δ relative to accumulation γ, the more innovators are in the market. In particular, all agents can be innovator when $\delta = \gamma$. This is an ill adapted situation where a high number of innovators does not translate into real progress, and fails to drive the price down to zero. *Technological progress* extinguishes entrepreneurial rents with a falling price, that follows after the unlimited reduction of production costs. This is the case of *learning curves*, that we address with attention in the final part of this section.

Technological progress can be sustained with a small fraction of innovators, when the demand is inelastic (Prop. 3.1). In general, high cumulativeness and strong spillovers (large γ) reduce the comparative advantage of innovators (Eq, 11 and Eq. 12). When the demand is inelastic this translates into more concentrated industries, because selection is tougher (Dosi, 1988). When the demand is elastic, the opposite is true, and technological progress characterises a market that converges to a complete dominance of innovators. These considerations are relevant to the question whether more competition is good or bad for innovation (Aghion et al., 2005). If one measures competition by the number of innovating firms (the total number of firms is fixed and large, by assumption), and innovation by price reduction, than the answer depends on the elasticity of the demand. Our model allows to capture this mechanism thanks to the interplay between technological dynamics and market dynamics with supply and demand.

A further message of our model is the following. The quantities n_t and $s(t)$ represent $R\&D$ intensity and innovation in an industry, respectively (Nelson, 1988). Our model describes exactly their relationship, by mean of an endogenous interplay between decisions n_t and technological change $s(t)$. Such a behavioural model of technological change allows to see how decisions translates into technological change and similarly how technological change affects agents' decisions. One important message from the model is that not necessarily many innovators make a competitive market together with sustained technological progress, as the scenario of balanced technological change shows. Often, a concentrated industry with few innovators does better in terms of competition intended as a falling price, which translates into higher consumer surplus. The relationship between innovation n and technological progress $s(t)$ is dictated by the elasticity of the demand, as explained by Prop. 3.1.

The model is simulated in different conditions that illustrate the scenarios described above (Fig. 5).[9] In the first scenario (top panels), the market presents an oscillatory phase before converging to the breakdown where $s(t) = 0$. This is the effect of the slowly varying frontier factor, which takes the model to a periodic orbit first, and then back again to a stable steady state condition. In the example with balanced technological change (middle panels), the fraction of innovators converge to $\frac{\delta}{\gamma} = 0.1$, while the price converge to a value near 0.6. Finally, the example of technological progress (bottom panels) presents a steadily

[9]Zeppini (2011) contains many simulation examples for the model with a linear demand curve.

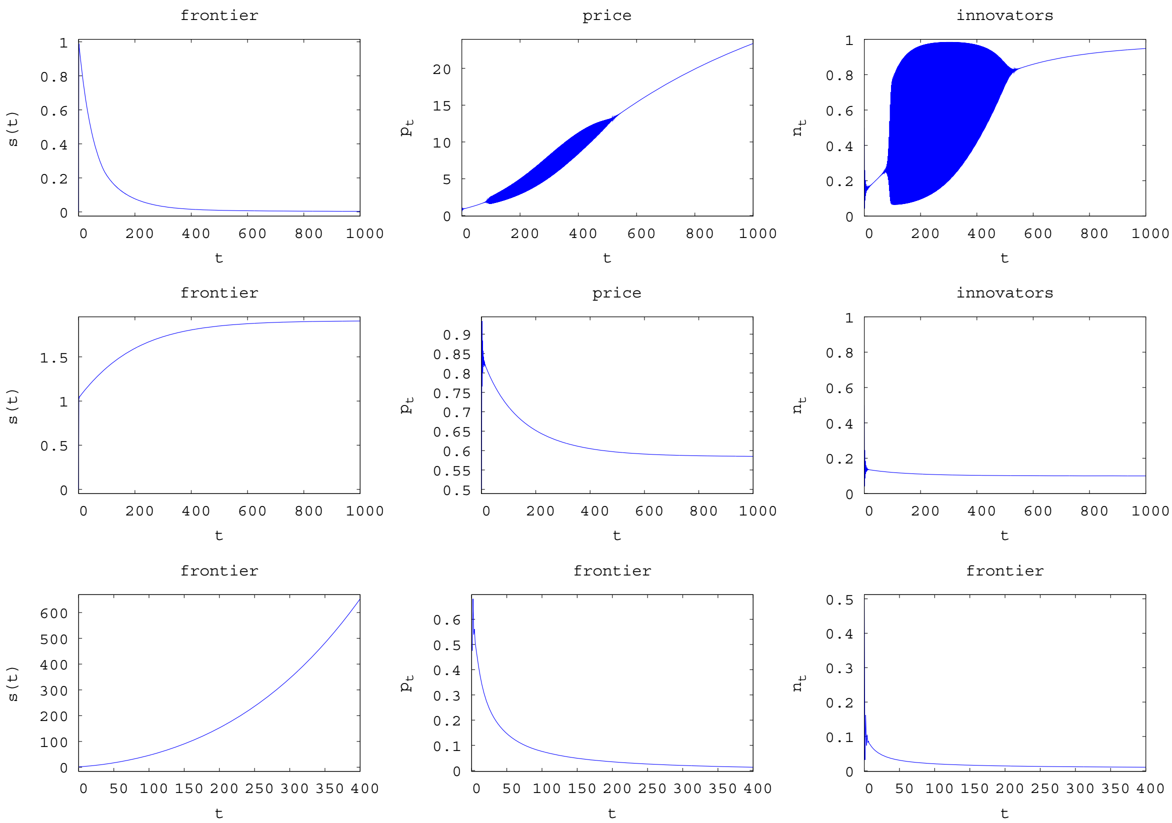

Figure 5: Model with technological change (inelastic demand). Top panels: example of market breakdown, with $\gamma = \delta = 0.02$. Middle panels: example of balanced technological change, with $\gamma = 0.1$ and $\delta = 0.01$. Bottom panels: example of technological progress, with $\gamma = 1$ and $\delta = 0.005$ (notice a different time scale). Here $\beta = 5$, $a = 1$, $d = 0.5$, $b = C = s = 1$.

declining price with an ever diminishing fraction of innovators, in accordance with Prop. 3.1.

The examples of Fig. 5 make use of an inelastic demand curve. In this setting, the scenario of balanced technological change turns out to be quite robust, and arises for a vast range of parameters settings. This is by no means the case with an elastic demand curve. In this case the model is much more sensitive to the parameters γ and δ, and presents sudden regime shifts from a market breakdown to a technological progress scenario for very small changes of γ and δ, making it very hard to find the right setting for a balanced technological change scenario. Fig. 6 reports an example for elastic demand. This is a technological progress scenario (exponentially growing frontier $s(t)$ and falling price p_t) that is characterised by an ever increasing fraction of innovators, as Prop. 3.1 indicates: with elastic demand technological progress requires many innovators.

3.2 Empirical learning curves

Learning curves are usually proposed in two versions, namely a relationship between marginal cost and output quantity (Argote and Epple, 1990), or a relationship between

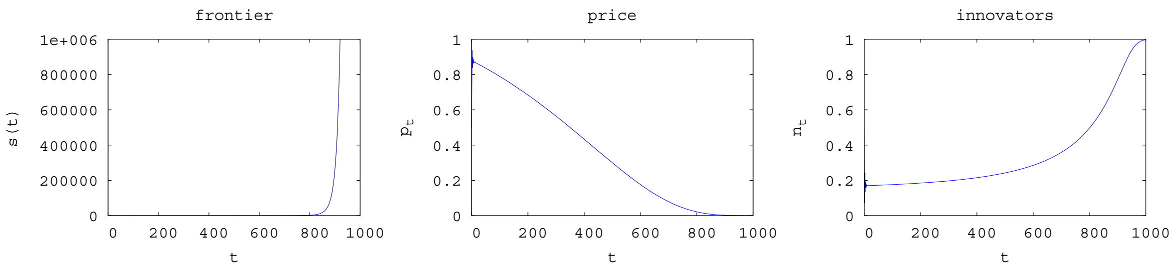

Figure 6: Model with technological change and elastic demand, technological progress scenario. Notice the increasing fraction of innovators and compare to the lower-right panel of Fig. 5. Here $\beta = 5$, $a = 1$, $d = 1.2$, $b = C = s = 1$, $\gamma = 0.1$, $\delta = 0.015$.

marginal cost and time like Moore's law (Koh and Magee, 2006). The latter is the version that we consider in this article, since the price reflects marginal costs.

Let introduce asynchronous strategy updating in the model with technological change. The resulting model reproduces the time pattern of learning curves with an irregular market variability. Both features are obtained through endogenous mechanisms based on agents' decision making and market dynamics. This full model is then used to match the empirical evidence from two examples of industrial sectors: the tyre industry and the solar module technology.

The cumulative process of technological change (11) works in the same way as before. In particular, the frontier $s(t)$ slowly changes the law of motion and possibly takes it through regions of different qualitative dynamics. Under asynchronous updating the dynamics is enriched with irregular chaotic orbits (see Section 2.3). It may be that a chaotic orbit is the long run outcome of a the model with technological change. It is exactly this condition that we will implement in order to reproduce the empirical time pattern of prices in an industrial sector. The variability of market dynamics is obtained endogenously from switching behaviour, without any exogenous noise factor.

It goes without saying that different industrial sectors require different settings of the model. It is not the purpose of this article to perform a model calibration. Nevertheless we can consider two examples of industrial sectors, namely the tire industry and solar technology, and show how the model can qualitatively reproduce the empirical evidence of market time series. The upper part of Fig. 7 refers to the tire industry. On the upper-left panel we have the empirical time series of the price index and exchanged quantity for the automobile tire industry in the US (Jovanovic and MacDonald, 1994). On the upper-right panel of Fig. 7 there is a simulation of the model for the same two time series, price p_t and quantity $q_t = D(p_t)$. The qualitative match of this example is obtained with an inelastic demand curve in a setting of balanced technological change. Both price and quantity match qualitatively empirical data. The fast oscillations of simulated time series can be averaged away by just sampling selected periods. Notice that firms can not scale up production in

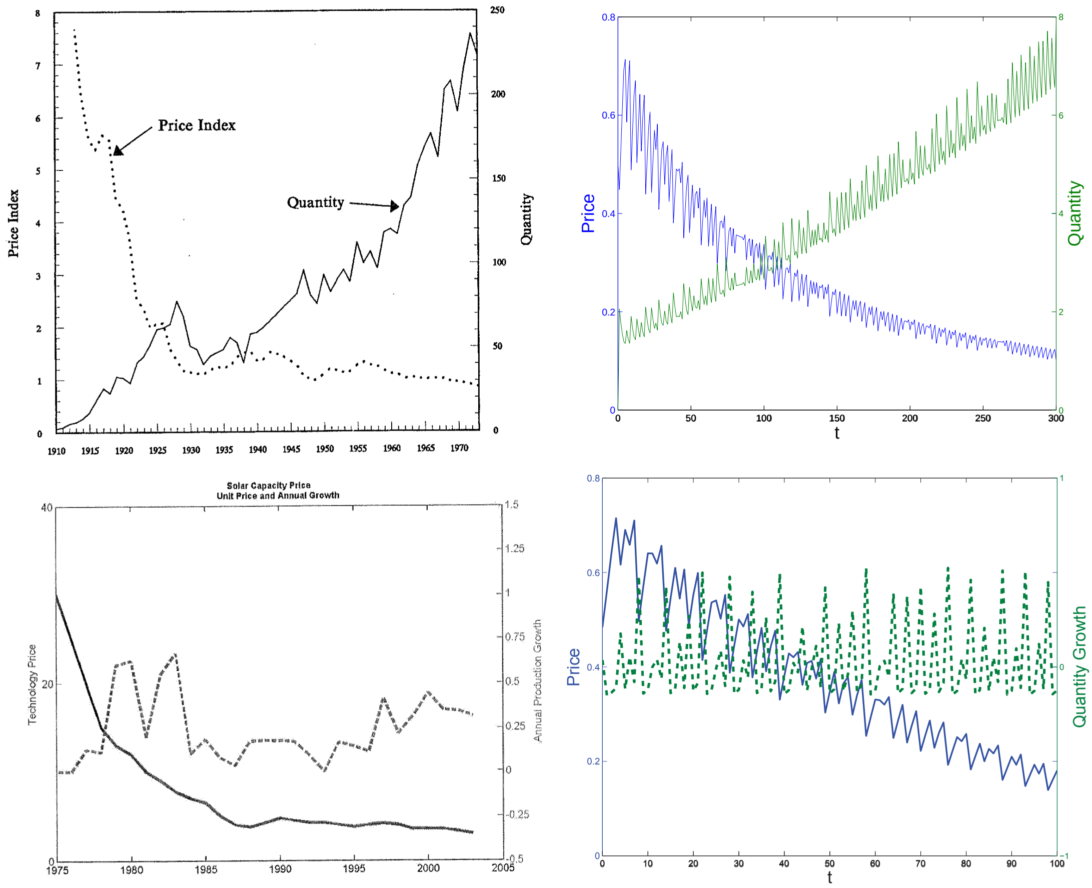

Figure 7: Upper-left: empirical time series of the price index in the US automobile tire industry (Jovanovic and MacDonald, 1994). Upper-right: simulated time series of market price and quantity. Model setting: $\beta = 15$, $a = 1$, $d = 0.9$, $b = s = 1$, $C = 2$, $\alpha = 0.6$, $\gamma = 0.2$, $\delta = 0.01$. Lower-left: empirical time series of a price index and production growth for solar module technology (Alberth, 2008). Lower-right: simulated time series of price and quantity growth rate. Model setting: $\beta = 15$, $a = 1$, $d = 1.1$, $b = s = 1$, $C = 2$, $\alpha = 0.6$, $\gamma = 0.2$, $\delta = 0.01$.

the model, so that an increased quantity is obtained only with higher productivity. While production scaling could be obtained by adjusting installed capacity, the actual model can better be compared to data on quantity per unit of production. Nevertheless, economies of scale are often less important than learning in reducing market price (Lieberman, 1984).

The lower part of Fig. 7 addresses solar technology. The lower-left panel contains the empirical time series of a price index for solar modules (referred to as "solar capacity unit price"), together with annual production growth. In the lower-right panel of Fig. 7 we report simulated time series for price and quantity growth rate ($\frac{q_t - q_{t-1}}{q_{t-1}}$). The match is again good. Notice that for this second example we have changed the model settings only in the price elasticity parameter, from $d = 0.9$ to $d = 1.1$.

The settings used to match examples of empirical time series should be compared by also looking at the technological frontier $s(t)$ and the fraction of innovators n_t. In the upper panels of Fig. 8 are simulations from the setting used for the US tire industry. In accordance with Prop. 3.1, an increasing technological frontier is accompanied here by a decreasing number of innovators, due to the inelastic demand. The fraction of innovators converges to $n^* = \delta/\gamma = 5\%$ (Fig. 8, upper-right panel), as Corollary 3.1 requires. Only few players are able to pursue innovation. Technological progress is limited (Fig. 8, upper-left panel), as it happens in a consolidated sector. This is reflected in the time series of the price, which converges to 0.05 (Fig. 8, upper-central panel). The lower panels of Fig.

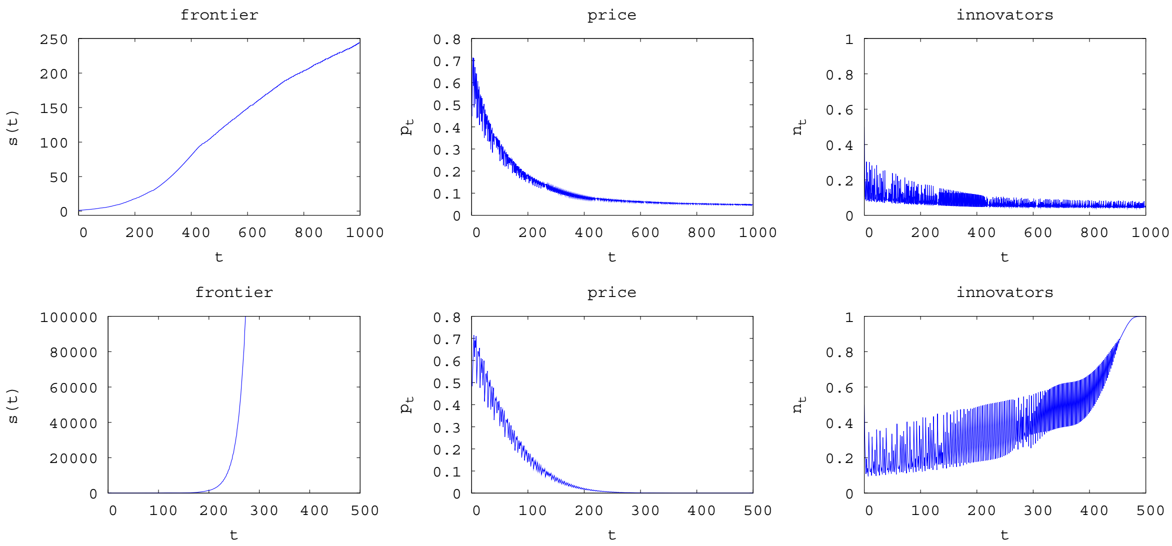

Figure 8: Upper panels: simulated time series in the setting used for the tire industry in Fig. 7. Lower panels: simulated time series in the setting used for the Solar technology in Fig. 7

8 report simulations with the setting used for the Solar modules technology. As expected, the elastic demand leads to an increasing concentration of innovators, in accordance with a Schumpeter Mark II *widening* pattern. It would be interesting to know data on the innovation behaviour of market participants in industrial sectors based on solar technology, in order to check this prediction of the model. In any case, an elastic demand makes sense, since renewable technologies are still not necessities, and their market penetration depends to a large extent on the price.

4 Conclusion

The model proposed in this article describes the effects of behavioural heterogeneity on technological change, with an endogenous interplay between adaptive heterogeneous firms, which either innovate or imitate, and a technological frontier that builds on firms' innovation decisions.

The core mechanism of the model is an evolutionary selection of agents' choices that affects endogenously the production technology. Similarly to a minority game, one strategy (innovation or imitation) is more profitable when the opponent strategy is dominant. Innovators drive down the market price because of cost reduction, but on the other hand they profit more from a high price. These two opposite incentives may end up offsetting each other in a stable equilibrium where both strategies coexist in some proportion. Alternatively, the model exhibits cyclical dynamics. Such a negative feedback mechanism is the dynamic counterpart of an inverted-U relationship between competition and innovation: a fall of price means stronger competition and it is associated with a surge in innovation, but at the same time it creates incentives for imitation.

The basic version of the model is extended first with asynchronous updating, second with technological change. The first is a more realistic assumption, where only a fraction of agents switch strategy in a given period. With asynchronous updating the dynamics of agents' choices and market price may turn chaotic. Although qualitatively destabilizing, asynchronous updating is quantitatively stabilizing, because it reduces the amplitude of market oscillations and increases persistence of strategies.

Technological change is introduced with a technological frontier that builds on agents' innovation decisions. Repeated choices between innovation and imitation shape dynamically the technological environment, and technological change feeds back into agents choices. This behavioural model of endogenous technological change presents three alternative scenarios: market breakdown, balanced technological change and technological progress. The first scenario describes abandoned industrial sectors. The second and third are more relevant to actual economic systems. Balanced technological change describes consolidated industrial sectors, where progress is relatively slow and price reduction is limited. Alternatively, this scenario fits so-called Schumpeterian rents, where innovating firms slow down technological investments in order to profit from innovation before it gets imitated. A technological progress scenario is characterised by an unbounded technological frontier, with market price falling to zero in the limit. This scenario fits better young and competitive sectors as hi-tech industries.

The model's scenario of technological progress is the more complex and rich one. Here the price elasticity of demand is a key factor. An elastic demand leads to a widening pattern of technological progress (Schumpeter Mark I) with increasing fraction of innovators. An inelastic demand does the opposite, leading to a deepening pattern (Schumpeter Mark II), where innovation is more concentrated. This result is relevant for understanding the relationship between competition and innovation. First, our model gives a behavioural explanation of the mechanism linking $R\&D$ intensity (fraction of innovators) and innovation outcome (the productivity technological frontier). Second, an elastic demand creates conditions where more competition is good for technological progress, while the opposite is true with inelastic demand.

The stylised fact of learning curves can be reproduced by the model, together with a market variability generated endogenously by agents' decisions, which can explain, at least in part, observed market variability. Models simulations are compared to two examples of industrial sectors. the US tire industry and global solar technology. A good match of both sets of time series is obtained by adjusting the price elasticity of demand. This evidence shows how the model's interplay between market conditions and agents' behaviours is a powerful mechanism for reproducing different patterns of technological change.

Appendix A Steady states and stability: proofs

Let us use the fraction n as state variable, and consider the map g of Eq. (8):

$$g(n) = \frac{1}{1 + e^{-\beta\left\{\frac{1}{2}s^{\frac{d-1}{d+1}}(e^{bC}-1)\left[\frac{a}{n(e^{bC}-1)+1}\right]^{\frac{2}{1+d}} - C\right\}}}. \tag{17}$$

This map is such that $0 < g(n) < 1$ for $n \in [0, 1]$. The first derivative is as follows:

$$g'(n) = -g(n)[1 - g(n)] \frac{\beta s^{\frac{d-1}{d+1}} a^2 (e^{bC} - 1)^2}{\left[(e^{bC} - 1)n + 1\right]^{\frac{d+3}{d+1}}}. \tag{18}$$

Since all parameters are positive, it holds $g'(n) < 0$. One and only one fixed point $n^* = g(n^*)$ exists, then. The same applies for the map f of Eq. (7). This proves Prop. 2.1.

The equilibrium corresponding to the fixed point n^* is stable whenever at least one of the parameters a, b, C, s, β is small enough, because $\lim_{\omega\to 0} g'(n) = 0$ for $\omega = a, b, C, s, \beta$. In particular, $\lim_{\omega\to 0} g'(n^*) = 0$. This proves Prop. 2.2.

Appendix B Conditions for chaotic dynamics

The model with asynchronous updating is specified by the map $\hat{g}$ of Eq. (10):

$$\hat{g}(n) = \alpha x + (1 - \alpha) g(n), \tag{19}$$

where g is the map (17) of the basic model with synchronous updating. Consider property $(A3)$ of Prop. 2.3, first. The stability condition of the steady state n^* is $-1 < \alpha + (1 - \alpha) g'(n^*)$. Since $\hat{g}'$ is bounded for finite values of β, a, d, s, b, C, there will always be a value of α close to 1 which makes the stability condition $|\hat{g}'| < 1$ hold true. Regarding property $(A1)$, the lower α, the closer the map $\hat{g}$ is to the map g of the basic model with synchronous updating. This means that in all situations where g has a stable 2-cycle, $\hat{g}$ has the same type of dynamics whenever α is close enough to 0. Finally, to prove $(A2)$ we follow Hommes (1994) p. 370. The map $\hat{g}$ of Eq. (19) is in the same class of functions of Eq. (12) in Hommes (1994), because it is obtained as a convex combination of a linear map (the diagonal) and a decreasing S-shaped map. Such functions have two critical points, c_1 and c_2, such that $\hat{g}'$ is decreasing in $[c_1, c_2]$ whenever one (or more) among β, a, d, s, b, C is sufficiently large, and it is increasing outside this interval with $0 < \hat{g}' < 1$. For intermediate values of α the map $\hat{g}$ has a 3-cycle (see Hommes (1994)) and chaotic behaviour then follows by applying the Li-Yorke "Period 3 implies chaos" theorem (Li and Yorke, 1975). Shifting the graph of such a map leads to bifurcations from a stable 2-cycle to chaos, and back to stable steady state (se Fig. 3).

Appendix C Technical change and demand elasticity

Consider a technological frontier $s(t)$ given by Eq. (11), and assume technological progress ($s'(t) > 0$). If we substitute s with $s(t)$ in Eq. (17) we can evaluate the effect of technological progress by differentiating the equilibrium value $n^* = g(n^*)$ with respect to s. Whenever $d < 1$, $\frac{\partial n^*}{\partial s} < 0$, while $d > 1$ gives $\frac{\partial n^*}{\partial s} > 0$. In the special case $d = 1$ we have $\frac{\partial n^*}{\partial s} = 0$.

Appendix D Proof of Proposition 3.2 and Corollary 3.1

Let us re-write the technological frontier as expressed in Eq. (16):

$$s(t) = se^{-\delta(t-1)}e^{\gamma \sum_{i=1}^{t-1} n_i}. \tag{20}$$

The highest rate of growth for the sum series is γ. Then $s(t) \to 0$ whenever $\gamma < \delta$, and $p_t \to \infty$ based on Eq. (14). Consequently, $q_t = D(p_t) = \frac{a}{p_t^d} \to 0$. This proves case (1).

If $\gamma = \delta$ (case 2), the long run value of $s(t)$ depends on the convergence of the sum series $\sum_1^\infty (1 - n_i)$. A necessary condition for convergence is $\lim_{i\to\infty} n_i = 1$. Whenever this condition does not hold true, $s(t) \to 0$ (case $2a$). If $\lim_{i\to\infty} n_i = 1$ fast enough, then $\sum_1^\infty (1 - n_i)$ may converge to a positive value Σ, and $p_t \to p^* > 0$ (Eq. 14).

When $\gamma > \delta$, everything depends on the rate of $\gamma \sum_1^t n_i$ relative to the linear trend δt. If the rate of growth of the sum series is lower than $\frac{\delta}{\gamma}$, then $s(t) \to 0$ (case $3a$). If the sum series achieve a linear trend at a rate exactly equal to $\frac{\delta}{\gamma}$, then we have the convergence of $s(t)$ and p_t to positive values (case $3b$). Finally, if $\sum_1^t n_i$ grows faster than $\frac{\delta}{\gamma}$, we have $s(t) \to \infty$ from Eq. (20), and $p_t \to 0$ from Eq. (14).

The special case of Cases $2b$ and $3b$ imply a steady state $n^* = \frac{\delta}{\gamma}$. In this case, the argument of the sum series in $s(t)$ (Eq. 20) converges to $\frac{\delta}{\gamma}$ by assumption. The argument of the second exponential in Eq. (20)) becomes $\delta(t-1)$ in the long run, then, which exactly offset the argument of the first exponential. On the other hand, for $s(t)$ to converge to a finite value, the argument of the two joint sum series must converge to zero, which implies $n^* = \frac{\delta}{\gamma}$.

References

Acemoglu, D. 2002. Directed technical change. *Review of Economic Studies* **69**(4) 781–809.

Acemoglu, D. 2007. Equilibrium bias of technology. *Econometrica* **75**(5) 1371–1409.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, P. Howitt. 2005. Competition and innovation: an inverted-U relationship. *Quarterly Journal of Economics* **120**(2) 701–728.

Aghion, P., P. Howitt. 1992. A model of growth through creative destruction. *Econometrica* **60**(2) 323–351.

Aghion, P., P. Howitt. 1998. *Endogenous Growth Theory*. MIT Press, Cambridge, MA.

Alberth, S. 2008. Forecasting technology costs via the experience curve - myth or magic? *Technological Forecasting and Social Change* **75** 952–983.

Argote, L., D. Epple. 1990. Learning curves in manufacturing. *Science* **247** 920–924.

Arthur, B. 1989. Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal* **99** 116–131.

Bass, F. M. 1969. A new product growth for model consumer durables. *Management Science* **15**(5) 215–227.

Benoit, J. 1985. Innovation and imitation in a duopoly. *Review of Economic Studies* **52** 99–106.

Berndt, E. R. 1991. *The practice of econometrics: classic and contemporary. Ch. 3: Costs, Learning Curves, and Scale Economies*. Addison-Wesley, Reading, Massachussets.

Branch, W. A. 2004. The theory of rational heterogeneous expectations: evidence from survey data on inflation expectations. *Economic Journal* **114** 592–621.

Breschi, S., F. Malerba, L. Orsenigo. 2000. Technological regimes and Schumpeterian patterns of innovation. *Economic Journal* **110** 388–410.

Brock, W. A., C. Hommes. 1997. A rational route to randomness. *Econometrica* **65**(5) 1059–1095.

Bulow, J. I., J. D. Geanakoplos, P. D. Klemperer. 1985. Multimarket oligopoly: strategice substitutes and strategic complements. *Journal of Political Economy* **93** 488–511.

Cabral, L. M. B., M. H. Riordan. 1994. The learning curve, market dominance and predator pricing. *Econometrica* **62**(5) 1115–1140.

Ceccagnoli, M. 2005. Firm heterogeneity, imitation, and the incentives for cost reducing *R&D* effort. *Journal of Industrial Economics* **53**(1) 83–100.

Cefis, E., L. Orsenigo. 2001. The persistence of innovative activities: a cross-countries and cross-sectors comparative analysis. *Research Policy* **30** 1139–1158.

Conlisk, J. 1980. Costly optimizers versus cheap imitators. *Journal of Economic Behavior and Organization* **1** 275–293.

de Jong, E., W. F. C. Verschoor, R. C. J. Zwinkels. 2009. Behavioural heterogeneity and shift contagion: evidence from the asian crisis. *Journal Economic Dynamics and Control* **33** 1929–1944.

Diks, C., C. Hommes, V. Panchenko, R. van der Weide. 2008. E&F Chaos: a user friendly software package for nonlinear economic dynamics. *Computational Economics* **32** 221–244.

Diks, C., R. van der Weide. 2005. Herding, a-synchronous updating and heterogeneity in memory in a cbs. *Journal Economic Dynamics and Control* **29** 741–763.

Dosi, G. 1988. Sources, procedures and microeconomic effects of innovation. *Journal of Economic Literature* **26** 1120–1171.

Dosi, G., L. Marengo, G. Fagiolo. 2005. Learning in evolutionary environments. K. Dopfer, ed., *The Evolutionary Foundation of Economics*, chap. 9. Cambridge University Press, Cambridge, England, 255–338.

Grossman, G. M., E. Helpman. 1991. *Innovation and growth in the global economy*. MIT Press, Cambridge, MA.

Grossman, S. G., J. E. Stiglitz. 1976. Information and competitive price systems. *American Economic Review Papers and Proceedings* **66**(2) 246–253.

Hartley, K. 1965. The learning curve and its application to the aircraft industry. *Journal of Industrial Economics* **13**(2) 122–128.

Hommes, C. H. 1994. Dynamics of the cobweb model with adaptive expectations and nonlinear supply and demand. *Journal of Economic Behavior and Organization* **24** 315–335.

Hommes, C. H. 2006. Heterogeneous agent models in economics and finance. Leigh Tesfatsion, Kenneth L. Judd, eds., *Agent-based Computational Economics*, *Handbook of Computational Economics*, vol. 2, chap. 23. North-Holland, Amsterdam, 1109–1186.

Hommes, C. H. 2011. The heterogeneous expectations hypothesis: some evidence from the lab. *Journal Economic Dynamics and Control* **35** 1–24.

Hommes, C. H., H. Huang, D. Wang. 2005. A robust rational route to randomness in a simple asset pricing model. *Journal Economic Dynamics and Control* **29** 1043–1072.

Iwai, K. 1984. Schumpeterian dynamics, part I: an evolutionary model of innovation and imitation. *Journal of Economic Behavior and Organization* **5**(2) 159–190.

Jovanovic, B., G. M. MacDonald. 1994. The life cycle of a competitive industry. *Journal of Political Economy* **102**(2) 322–347.

Kandori, M., G. J. Mailath, R. Rob. 1993. Learning, mutation, and and long run equilibria in games. *Econometrica* **61**(1) 29–56.

Koh, H., C. L. Magee. 2006. A functional approach for studying technological progress: application to information technology. *Technological Forecasting and Social Change* **73** 1061–1083.

Li, T. Y., J. Yorke. 1975. Period three implies chaos. *American Math. Monthly* **87** 985–992.

Lieberman, M. B. 1984. The learning curve and pricing in the chemical processing industry. *RAND Journal of Economics* **15**(2) 213–228.

Llerena, P., V. Oltra. 2002. Diversity of innovative strategy as a source of technological performance. *Structural Change and Economic Dynamics* **13** 179–201.

Malerba, F. 1992. Learning by firms and incremental technical change. *Economic Journal* **102**(413) 845–859.

Mansfield, E. 1961. Technical change and the rate of imitation. *Econometrica* **29**(4) 741–766.

McCabe, M. J. 1996. Principal, agents, and the learning curve: the case of steam-electric power plant design and costruction. *Journal of Industrial Economics* **44**(4) 357–375.

Nelson, R. R. 1988. Modelling the connections in the cross section between technical progress and *R&D* intensity. *RAND Journal of Economics* **15**(4) 546–554.

Nelson, R. R., S. G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.

Petrakis, E., E. Rasmusen, S. Roy. 1997. The learning curve in competitive industry. *RAND Journal of Economics* **28**(2) 248–268.

Romer, P. M. 1990. Endogenous technological change. *Journal of Political Economy* **98**(5) 71–102.

Schumpeter, Joseph A. 1942. *Capitalism, socialism and democracy*. Harper and Row, New York.

Sethi, R., R. Franke. 1995. Behavioural heterogeneity under evolutionary pressure: macroeconomic implications of costly optimization. *Economic Journal* **105** 583–600.

Simon, H. A. 1957. *Models of Man. Social and Rational*. Wiley, New York.

Spence, M. A. 1981. The learning curve and competition. *Bell Journal of Economics* **12**(1) 49–70.

Spence, M. A. 1984. Cost reduction, competition and industry performance. *Econometrica* **52**(1) 101–121.

Zeppini, P. 2011. *Behavioural Models of Technological Change (PhD Thesis), Chapter 4: A behavioural model of endogenous technological change*. Tinbergen Institute.