Technological Innovation and Diffusion, Fluctuations and Growth (I): Modeling Technological Change and Productivity Growth*

Pier Giorgio Ardeni
University of Bologna

Mauro Gallegati
University of Pescara

June 1993

Classification JEL number: O30, O40

Abstract

In this work we study the relation between investment in R&D, technological innovation, diffusion, fluctuations and growth of output. Technological innovation is the result of a process (investment in R&D) whose final outcome is fundamentally uncertain. We model innovation as a Polya urn scheme, where the probability of an individual success increases with the number of (previous) successes. We also model technological diffusion as an epidemic birth-and-death stochastic process, tipically a non-linear process. The diffusion of a new technology is to a degree intrinsically stochastic: it depends on aggregate feedbacks (global environment) as well as on local feedbacks (imitation) or investment (output demand). Firm growth and aggregate growth are related: if we measure the former in terms of firm size, the latter depends both on firm growth and on growth in the number of innovators. Firm growth is due to productivity growth (technological change), whose incentive, from the firm point of view, is profit. Higher cash flow (profits) implies higher funds for investment, higher research effort, and thus potentially faster technological change. However, if innovations spread out, monopoly rents will be temporary, and when someone innovates sooner or later everybody will innovate (unless exiting the market). In Part I of the paper we describe the mechanics of technological innovation and diffusion and a model for an innovating firms. In Part II the deterministic and stochastic laws of motion which arise are analyzed in detail.

---

*This work is part of an ongoing research on technological innovation and growth, whose preliminary results have been presented at seminars held at the University of Trento, the Center for Economic Dynamics and Institutional Change of the University of Bologna, and at the University of Ancona, whose organizers Massimo Egidi and Pier Luigi Sacco, Gilberto Antonelli and Patrizio Bianchi, and Giorgio Fuà, respectively, and qualified audiences we thank wholeheartedly.

1. Introduction

Several contributions in growth theory have recently emphasized the positive relation between firm investment (maybe through self-financing), R&D expenditure, technological innovation and economic growth (e.g. Levin, Cohen, and Mowery (1985), Scott (1989), Cohen and Levinthal (1989), Aghion and Howitt (1993)). A great deal of recent models emphasizes R&D as an important source of productivity growth and hence of overall aggregate economic growth, as opposed to the traditional neoclassical framework where productivity growth was conceived as purely exogenous. As we know, the simple neoclassical model predicted a zero long-run growth rate. Thus, in order to explain observed aggregate growth, an exogenous productivity growth was incorporated into the neoclassical production function. Technological *progress*, in the form of *factor-augmenting* change in the scale of production was then introduced, giving rise to a composite taxonomy according to whether such progress was just capital-saving, labor-saving or both (see e.g. Burmeister and Dobell (1970)). Examples of such production-function concept of technological change are the well known works of Hicks (1932), Harrod (1948) and Solow (1956). Yet, technological advances as such were left unexplained: technological change was not only *exogenous* but also *disembodied*, i.e. not incorporated in new capital goods. Later models had *induced* technological change (e.g. Kennedy (1962)) or *embodied* technological change (Solow (1960) or Bliss (1966)), thus trying to endogenize the current productivity of factor inputs. Yet, vintage models have not succeded in separating investment in new technology from investment in capital replacement. Thus, the interest in technological change as an *endogenous* phenomenon has recently increased (see e.g. Gomulka (1990) for an overview).

As any theory of growth must have technological change as one of its central elements, a major task that the literature starting from the work of Arrow (1962) on the effects of *learning* and the accumulation of knowledge has tried to cope is the understanding of the sources of technological change. Several different strands of literature have studied the interactions between the acquisition of new technology and the accumulation of knowledge: from Griliches (1979) to Machlup (1980), to Rosemberg (1972), to David (1985). An explanation of the relation between endogenous innovation and economic growth has also been given in recent works in the *evolutionary* tradition (see for instance the volume edited by Freeman (1990)).

Among these contributions, one major source of productivity growth has been identified with investment, and therefore, whether directly or indirectly, with R&D expenditure. R&D expenditure may contribute to growth in different ways. It allows to introduce new types of goods, either capital or consumption goods. Also, it may generate spillovers on the aggregate "stock of knowledge". In both cases there will be constant returns to investment in R&D, as opposed to diminishing returns as in the traditional approach. Also, there may be temporary monopoly profits and externalities at the aggregate level: while the former will form an incentive to keep investing in R&D, the second will give rise to increasing aggregate long-run growth. These R&D models are the most recent offspring of the endogenous growth literature, whose interesting feature is that "endogenous" growth can be generated through the accumulation of knowledge alone, broadly defined. In general these are multi-sector models, based on product differentiation, monopolistically competitive markets and given rate of arrivals of new technologies: examples are the work of Romer (1987), Aghion and Howitt (1993), and Barro and Sala-i-Martin (1990).

In this paper we want to address the issue of long-run growth as due to productivity growth which stems from technological change induced by investment in R&D generating technological advances. By taking aggregate growth as the interaction between *productivity* growth and *firm-size* growth, we are able to focus on firm behavior and hence to its investment policy. We denote with *investment in R&D* that investment that is not devoted to the mere replacement of obsolete physical capital and which, for that matter, it could be just *plain investment*. Yet, we know that the standard neoclassical investment theory framework would predict and optimal investment rate just equal to depreciating capital (Lucas (1972)). Hence, we ought to distinguish investment in *new* capital or more generally in *new technology* from replacement investment. Also, we want to consider a key aspect of technological innovation, namely the fact that it is the result of a process (research and implementation) whose final outcome is fundamentally uncertain. And this is the reson why we think that technological change, whatever form it takes, ought to be modeled as a stochastic process.

In what follows, we present a model in which both technological innovation and innovation diffusion across firms are driven by stochastic processes. Our aim is to offer a general framework by which to explain the discontinuous character of the introduction and adoption of technological innovations as well as their diffusion and, ultimately, overall economic growth. While the recent growth literature has emphasized the endogeneity of technological advances that generate increases in factor productivity (e.g. Romer (1986), Lucas (1988)), few have been the contributions (most notably Aghion and Howitt (1993), Grossman and Helpman (1991)) that have pointed out one of the main feature of technical progress, namely that technological improvements are essentially stochastic, being the result of a research process whose outcome is uncertain.

In this paper, technological innovations are conceived as Poisson process, whose arrival rate is variable as being generated by a Polya urn scheme, while the diffusion of innovations is conceived as a Poisson process of the kind of an epidemic birth processes. A major caveat is of order at this point. When we speak of innovation (or technological advances), we mean technologies whose *productivity* is higher then the existing ones. In this context, innovations (or even inventions) do not come out of nothing: in a way they are all already there, in an imaginary shelf, ranked according to their productivity (in ascending order)[1]. Two are the restrictions we impose on the acquisition process. In the first place, each firm can choose only between *two* technologies: an *existing* one and a *new* one. The latter is new because *nobody is already adopting it*. Firms cannot switch to a newer technology before they have switched to the new one (this is to avoid technological laps). Secondly, this choice is *stochastic*, meaning that technological change does not occur with certainty (albeit *on average* productivity does have a positive rate of growth), nor it does occur with certainty *at a given time*. The first assumption implies the *continuity* of technological change, as well as its *cumulativeness* (in line with Dosi (1988)). The latter encompasses overall *learning* for individual firms: before switching to a newer technology we have to learn what the new is. The second assumption implies that technological progress is not *certain* at the firm level (there might be delays of any kind in the learning process as well as in the adoption process) while it is at the aggregate level (*on average*). This also emphasizes that innovation is the uncertain outcome of investment in R&D, broadly defined.

Technological innovations are introduced randomly by firms, according to a Poisson law generated by a Polya urn drawing scheme which "simulates" the trial-and-error lying behind investment in R&D and R&D output and more generally technological research. As the overall chances of success increase with the number of draws (amount of investment), the positive relation between R&D expenditure and innovation is apparent. Here we want to

[1] This shelf resembles the *blueprint book* where all possible future blueprints have already been patented.

focus *on the mechanics* of technological innovation, by means of a scheme which emphasizes the stochasticity of the research effort. Such a general scheme simply assumes that a given amount of investment is devoted to technological change: obviously, this approach encompasses models that emphasize human capital as the fundamental input of the innovation process, as well as models that specify a research sector as distinct from the intermediate/consumer goods sector.

Yet, technological innovations take time to spread, innovative firms are often followed by imitators, and the introduction of an innovation and its diffusion by adoption appear to be distinct processes. Thus (differently from Aghion and Howitt (1993)), we model innovation and diffusion separately. The best way to model diffusion seems to be that of (discrete) stochastic diffusion models, namely *birth-and-death Poisson* processes. By adopting an epidemic model, we can model technological diffusion according to a non-homogeneous birth process, where the speed of adoption is non-linear in the number of adopters (from an initial slow speed to a faster one to a slower one in the end). The result is a growth path which is locally linear but globally non-linear, albeit convergent. Again, by focusing on the mechanics of the diffusion process, we abstract from a specification of the determinants, the latter being implicitly subsumed in the framework we adopt. Our approach, in fact, encompasses those models where innovations take the form of blueprints that are sold to the other sectors by the research sector (like Aghion and Howitt (1993) or Grossman and Helpman (1991)).

If the search for new technologies starts over again at any technological switch, the process keeps going forever (we have multiple technological paradigms), and the productivity increase is potentially unbounded. Thus, as technologies are ranked according to their productivity, we design a stochastic birth-and-death process where the *state* of the process is defined by the number of firms adopting a new (but well-established) technology. Firms will *enter* the state from the existing ("old") technology and will *exit* the state towards the newest (advanced) technology. This is a nonlinear stationary birth-and-death process. If multiple technological paradigms are allowed, stochastic trends in the growth path may show up at the aggregate level *in the long run* as the result of the superimposition of many different segmented locally linear trends. In a way, this resembles the "old" turnpike models, albeit in a different context. Here two elements contribute to the dynamics. There is a *short-run* component, given by the distribution across firms of technologies in any given time. That is, in any given time each firm will be adopting a given technology according to its own chances of success in innovating or in adopting an already existing new technology. This generates *fluctuations* around the short-run dynamic path which is convergent and locally linear. But there is also a *long-run* component. In the long-run, firms tend to stick to the new technology, which is established, not to the old nor to the newest (the largest proportion of firms, in fact, tends to concentrate on the new technology). But as newer and newer technologies are introduced, the shifting will go on. Trends that appear as linear at a local level (in the short run) become piecewise linear at the global level, and hence globally non-linear (in the long run). Yet, as the switching points are stochastic, even in the aggregate, they can well be considered as *stochastic trends*.

Thus, the stochastic approach we adopt allows a full description of the mechanics of innovation and diffusion, which can be given an economic content in a more structural model. A stochastic model typology is illustrated below in Section 2, while a theory of firm investment incorporating the dynamics of technological innovation is sketched in Section 3. The setup we adopt is fairly simple. By resorting to a standard investment model for an individual firm, we can express the growth of the market value of the firm (which is the choice variable and proxies for the firm *size*) as a nonlinear differential equation, where both the rate of growth of productivity and the number of *innovative* firms in the economy (or in an industry) appear as explanatory variables. All the assumptions lying behind this simple model are quite standard: all markets are perfectly competitive, all prices are given to the firm, its labor demand is also given by a production technology which is a function of labor only. The rate of growth of productivity is given by technological progress. The latter depends on firm investment, which we assume as proportional to total cash flow. As it is, the setup if fairly different from the recent "aggregate" representative agent optimizing models. Aggregate growth, in our context, stems from growth in firm size as well as in firm number.

The dynamics of the models is fully developed in Sections 4 and 5. As the stochstic processes described above all admit *deterministic equivalent* versions, we are able to analyze separately the behavior of a model where technological change follows deterministic laws of motion from a model where it follows stochastic laws. Both technological change and the firm value have their own dynamics, but it is their interaction which is the key aspect of our analysis. As it turns out, while a fully deterministic model predicts local saddle path stability along a balanced growth path, the stochastic model emphasizes the fluctuations thereby generated, as well as the nonconvergent asymptotic behavior of the model.

The approach we adopt is not an optimizing one, at least in the usual sense. In a way, we consider the behavior of the firm after the state variables of its maximization problem have already been "optimized out". The interaction between firm size and technological change is in fact the key aspect. Thus, the model is not even a model of aggregate growth, properly speaking. Further developments along the lines of a fully optimizing model are presently under study (Ardeni and Gallegati (1993)). Yet, many are the further developments we presently foresee of such an approach: from "segmented trends" to "stochastic endogenous dynamics", from the distinction between "generation effects" and "development effects" to the role of credit (i.e. availability of financial resources) in the growth process. One major result of this work is the attempt to combine the analysis of the growth process at the firm level with that at the aggregate level: there are obvious interactions which may potentially give raise to interesting dynamics that have been neglected in the literature on growth and investment. Both deterministic and stochastic laws of motion arise, generating growth paths about which fluctuations develop. The study of their behavior is a task which appears to have not been fully accomplished yet.

The paper is divided in two parts. In Part I we describe the mechanics of technological innovation and diffusion and derive a class of stochastic processes for the number of innovating firms (Section 2), while discussing the economic rationale which justifies them and their consequences (Section 3). In Part II we describe the deterministic and stochastic laws of motion associated with such processes (Sections 4 and 5).

2. The mechanics of technological innovation and diffusion. A stochastic model typology

2.1. Innovation with no diffusion

We can conceive any technological improvement as the result of a research process whose final outcome is uncertain albeit positively related to the research effort. Consider a closed economy with no public spending, in which all technological advances come from firms who try to exploit an exogenously given scientific body of knowledge. Such aggregate stock of knowledge is a purely public good, which cannot be used directly and has to be *converted* in the form of technical knowledge. The existing body of knowledge is thus assumed to be already "incorporated" in the current prevaling technology. Suppose firms invest in R&D in order to *discover* (devise or invent) more efficient technologies. Since with the existing technology (shared by everybody) all firms operate at a normal profit level, there is an incentive due to the temporary monopoly rents which might be gained by adopting a more efficient technology. These monopoly rents will be slowly eroded as innovation diffusion will take place, thus driving down all temporary profits. However, as the timing of discovery and adoption is random, the whole process may take place in quite different ways, thus potentially generating fluctuations around the mean growth path.

The devise of a new technology can be modeled as an urn drawing scheme, where we can think of an urn containing two *patents* (blueprints) in different proportions, one representing a *new* technology and the other an *old* one. Assume every new technology is more productive than the older one (it produces a higher level of output given the same quantity of inputs). For simplicity, consider the case of just one urn, containing r red balls representing the new technology and b black balls for the old technology. Thus, the drawing of a red ball symbolizes the "discovery" of a new technology, i.e. a technological *innovation* by the firm who drew the ball. Firms partecipate to this "urn lottery" through their R&D expenditure: they pay a ticket to join the lottery so that the higher their R&D expenditure, the more tickets will be available to them, the higher the chances to draw a red ball.

To model this innovation process we might start with the urn model introduced by Polya (Feller, 1971). An urn contains b black balls and r red ones. When a ball is drawn, it is reintroduced together with c balls of the same color (the simplest process has $c=1$). Thus, if we draw a red ball from an urn with $b+r$ balls, at the next drawing we will have b black balls and $r+1$ red ones. In that case, while the probability to pick a red ball in the first draw was $r/(b+r)$, given that the first ball drawn was red, the probability to pick a red ball in the second draw will be $(r+1)/(b+r+1)$. Therefore, the drawing of one of the two colors increases the probability that the color drawn is drawn in the next draw. Thus, the probability of getting n_1 red balls (when n_2 are black ones) in $n=n_1+n_2$ draws is:

$$p_{n_1,n} = \frac{\binom{-r}{n_1}\binom{-b}{n_2}}{\binom{-(b+r)}{n}} \tag{2.1}$$

Now, if we let $p=r/(b+r)$, $q=b/(b+r)$, and $\gamma=c/(b+r)=1/(b+r)$, we get the so-called *Polya distribution* on the integers $\{0, 1, \dots, n\}$:

$$p_{n_1,n} = \frac{\binom{-p/\gamma}{n_1}\binom{-q/\gamma}{n_2}}{\binom{-1/\gamma}{n}} \tag{2.2}$$

such that $p+q=1$. It is well known that the limit form of the Polya distribution is the *negative binomial distribution*. In fact, if $n\to\infty$, $p\to 0$ and $\gamma\to 0$, we have that $np\to\lambda$ and $n\gamma\to\rho$, and, for a given n_1 :

$$p_{n_1,n} \to \binom{\lambda\rho+n_1-1}{n_1} \left(\frac{\rho}{1+\rho}\right)^{\lambda\rho} \left(\frac{1}{1+\rho}\right)^{n_1} \tag{2.3}$$

which sums to unity for fixed values of λ and ρ.

Now, if we let a binary random variable X_n be 1 or 0 depending on whether the n-th draw is red or black and Y_n be the number of red balls in n draws, i.e. $Y_n=X_1+X_2+\dots+X_n$, we then have that

$$E(Y_n) = \frac{nr}{(b+r)}$$
$$VAR(Y_n) = \frac{nbr(b+r+c)}{(b+r)^2(b+r+c)} \tag{2.4}$$

and the sequence $\{Y_n\}$ is an ordinary Markov chain, with constant transition probabilities. This allows us to derive the so-called *Polya stochastic process* which is nothing but a Markovian process in continuous time, a *pure non-stationary birth process*. Such a process can be obtained as the extension to the limit of the Polya urn scheme. If we define the *state* of the system as the number of red balls drawn, then the transition probability $E_k \rightarrow E_{k+1}$ at the $(n+1)$-th drawing[2] is given by:

$$p_{k,n} = \frac{r+kc}{b+r+nc} = \frac{p+k\gamma}{p+n\gamma} \tag{2.5}$$

where $p=r/(b+r)$ and $\gamma=c/(b+r)$. If one assumes that draws occur at a rate of one every h time units and $h \rightarrow 0$ as $n \rightarrow \infty$ so that $np \rightarrow t$ and $n\gamma \rightarrow at$ [3], we then have that, in the limit, (2.5) tends to:

$$\lambda_n(t) = \frac{1+an}{1+at} \tag{2.6}$$

where $\lambda_n(t)$ is the birth rate of the non-stationary Markovian birth process, which thus depends on time t [4]. We also have that, in the limit, (2.5) becomes:

$$P_0(t) = \left(\frac{1}{1+at}\right)^{1/a} \tag{2.7a}$$

---

[2] That is, the probability to go from k red balls drawn to $k+1$ red balls drawn at the $n+1$-th drawing.

[3] This implies that $nr/(b+r) \rightarrow t$ and $nc/(b+r) \rightarrow at$, that is $c/r \rightarrow a$.

[4] We can write (6) as $\lambda_n(t)=(1+nc/r)/(1+nc/(b+r))$.

$$P_0(t) = \frac{t^n \left[\prod_{i=1}^{n-1} (1 + ia) \right]}{n! \, (1 + at)^{\frac{n-1}{a}}} \qquad (2.7b)$$

that is, given the initial condition $P_0(0)=1$, (2.7) is a solution of (2.6).

As the Polya urn model, and the derived Polya stochastic process, can thus be treated as standard Markovian processes (of which we know the asymptotic behavior and the distribution), we can couch the problem in terms of well established theory of stochastic processes. In our case, the pure birth process is a non-stationary Markov process, and belongs to the class of time-dependent (non-homogeneous) Poisson processes. Now, the modeling of the technological innovation process as a non-homogeneous time-dependent Poisson process surely captures the basic notion of innovation as a search process with an uncertain outcome. In the evolutionary literature, Dosi (1988) has pointed to this aspect of the innovation process. Dosi makes a distinction between *weak* uncertainty (the probability distribution of an event is known) and *strong* uncertainty (not even the probability distribution is known) and argues that innovation involves a considerable degree of strong uncertainty. Of course, this latter notion is not captured in the model discussed here (where the probability distribution is known), as we do need some distributional results to model the process quantitatively. However, resorting to a Polya urn model for modeling technological innovation has a few advantages over other models in which technological advances are also stochastic but where the arrival rate of research success is constant over time. Aghion and Howitt (1993), for instance, assume that the research sector produces a random sequence of innovations: the Poisson arrival rate of innovations in the economy at any instant is λn, where n is the flow of labor used in research and λ a constant parameter given by the technology of research. The intensity of the research effort (which is R&D expenditure in our case) will thus depend on the amount of labor employed, but the chance of success will be constant. Within a Polya urn model as the one depicted above, the chance of success at each drawing is constant, too (and is given by the initial proportion of red balls), but the overall number of red balls drawn will increase over time linearly.

Now, we have assumed that there is just one urn, so describing two alternative technological paradigms, a new and an old one. If we let $p=r/(r+b)$ and $\gamma=c/(r+b)$ as before, as n becomes large, the stochastic process we have just described has mean equal to np (it is non-stationary in mean, albeit ergodic) and variance equal to $np(1+n\gamma)$ (it is also non-stationary in variance). Such a process can thus be thought of as evolving around a linear trend: the more firms draw balls from the urn (the more they spend in R&D), the larger the number of red balls in the urn[5]. Yet, the proportion of red balls will remain constant, but the number of firms "discovering" the new technology will increase linearly over time: in the limit, all firms will discover and adopt the new technology. Thus, the process, which evolves around a linearly increasing trend, is non-stationary in mean and has a variance that increases over time: this means that *fluctuations around the mean growth path increase over time*. An example is given in Figure 2.1.

Of course, this simple scheme does not allow for diffusion and adoption of the new technology other than by investing in R&D and successfully devising it. Namely, there is no imitation nor diffusion by knowledge spillovers or by imitation. In other words, we have implicitly assumed each firm must go through the same search process: *there is no external benefit in the R&D expenditure of others*. How then to model the diffusion of technological innovations, i.e. the increase in the number of firms that will adopt the new technology as

[5] Since $P_{n_1,n} > P_{n_1,m}$, $\forall\ n<m$, as the number of drawings increases, the probability of drawing at least one red ball will increase, as $P_{0,n} > P_{0,m}$, $n <m \Rightarrow (1 - P_{0,n}) < (1 - P_{0,m})$, $n <m$.

the latter becomes popular? The diffusion process typically takes place over time (or space, for that matter), possibly with diminishing marginal costs. Simple models for technological diffusion are illustrated in the following Sections 2.2, 2.3, and 2.4.

The simple Polya urn scheme depicted above is just one example, in a way the most general. We have in fact assumed that no matter how much firms invest, as far as they do not devise the more efficient technology, they will keep on drawing from the urn. Overall uncertainty regarding the probability of a discovery remains unchanged (asymptotically the probability converges to p). Yet, it is the timing of such discovery that changes, as it decreases the more draws are made. This simple process does not allow for *learning*, strictly speaking, and it concerns a very raw binary choice between two technologies, one of which is (maybe) just marginally better than the other. However, learning from past experience is implicit in the fact that *all previous technologies are subsumed in the existing one*. A more complex model would allow, for instance, three or more technologies among which to choose or a different ball replacement scheme. A typical learning process would imply, say, that *one ball of the color drawn is added only when a red ball is drawn*. Or, instead, that *two balls of the color drawn are added when a red ball is drawn*. All these extensions are possible and the implied stochastic process will be different. Here we will exploit this simple model, leaving all other extensions to future research.

2.2. Diffusion in a population of unlimited size

We are now going to consider a typical diffusion process of an innovation which is being introduced by one firm at some point in time (time 0). The diffusion process starts at the time a firm has discovered a new technology and the latter is made available to other firms independently from their R&D expenditure. This is basically the opposite of the previous case, where everybody invests in R&D and no diffusion is allowed. Here, the moment anyone finds a new technology, the diffusion process starts with no more investment in research (to make things simple we can assume that all R&D expenditure is then switched to the acquisition of the newly discovered and available tech).

The simplest model of diffusion could be of a kind of a pure birth process. With full appropriability of the new technology, we can think of the diffusion process as growing linearly with the number of firms adopting the new technology. This is equivalent to a birth process where each member of a population acts independently and gives birth at an exponential rate λ. If we suppose that technological adoption is irreversible and the overall population of firms is of unlimited size, then, if $X(t)$ represents *the size of the population of adopting firms* (the number of innovating firms) at time t, the sequence $\{X(t),\ t \geq 0\}$ is a pure birth process with rate:

$$\lambda_n = n\lambda, \qquad n \geq 0 \ . \tag{2.8}$$

This is a pure birth process called *Yule* process (Ross (1983, p. 144)) starting with a single individual (a firm devising a new technology) at time 0. Hence we know that the population size at time t will have a *geometric* distribution with mean $e^{\lambda t}$. If we let T_i be the time it takes for the population size to go from i to $i+1$, it follows that T_i is exponential with rate $i\lambda$.

Now, this is indeed a very simple diffusion process and can hardly be realistic. The size of the population of innovating firms increases linearly, i.e. the number of firms adopting the new technology grows unlimitedly. Also, even if we let the birth rate λ depend on the firm R&D expenditure, the total population is asymptotically unbounded.

2.3. Diffusion in a population of limited size

In order to model the diffusion process more satisfactorily we can resort to a simple *epidemic* model of the kind of a stationary pure birth process. As before, by diffusion we mean the adoption of an existing newly discovered technology. Suppose we have a population of m firms that at time 0 consists of *one* "infected" and m-*1* "susceptibles" (time 0 is basically the moment any of the m firms draws a red ball). Once infected, i.e. once a firm has discovered a new technology, an individual firm remains in that state forever (the *truly innovative* is the first firm infected, while all the subsequent infected are the *imitating* ones) and we suppose that in any time interval h any given infected firm will cause, with probability $\lambda h+o(h)$, any given susceptible to become infected. Put another way, in any time interval any given susceptible firm will adopt (discover, i.e. devise, or imitate, or copy, or buy), with probability $\lambda h+o(h)$, the same new technology. If we let $X(t)$ denote the number of infected firms in the population at time t, the sequence $\{X(t),\ t \geq 0\}$ is a pure birth process with rate:

$$\lambda_n = (m - n)n\lambda, \qquad\qquad n = 1, \ldots, m-1 . \tag{2.9}$$

This is a non-homogeneous Poisson process, with a state-dependent arrival rate. Equation (2.9) follows since when there are n firms adopting a new technology, then each of the m-n susceptible will switch to the new technology at the increasing rate $n\lambda$: the more the new technology is spread over firms, the more likely other firms will adopt it. Clearly, if we let the total population size m go to infinity, we are back to the previous case. Yet in that model we did not have a *contagion* effect, which now comes in through a multiplicative effect.

If we let T denote the time until the total population adopts the new technology, then T can be represented as:

$$T = \sum_{i=1}^{m-1} T_i \ , \qquad (2.10)$$

where T_i is the time to go from i innovative firms (infectives) to $i+1$ infectives. As the T_i are independent exponential random variables (as $X(t)$ is Poisson) with rates $\lambda_i = (m\text{-}1)\ i\lambda$, with $i=1,\ldots,m\text{-}1$, we have that:

$$E[T] = \frac{1}{\lambda} \sum_{i=1}^{m-1} \frac{1}{i\ (m - i)} \qquad (2.11)$$

and:

$$Var[T] = \frac{1}{\lambda^2} \sum_{i=1}^{m-1} \left[\frac{1}{i\ (m - i)}\right]^2 . \qquad (2.12)$$

For a population of firms of reasonably large size, $E[T]$ can be approximated as (cfr. Ross (1983, p. 147):

$$E[T] = \frac{1}{m\lambda} \sum_{i=1}^{m-1} \left(\frac{1}{m - i} + \frac{1}{i}\right) \approx \frac{2\ \log(m - 1)}{m\lambda} . \qquad (2.13)$$

In order to give a better understanding of the actual implications of this diffusion model consider the following example. We have a finite number of firms m, each of which has access to an urn with $r=1$ red balls and $b=m\text{-}1$ black balls. Thus, a new technology can be discovered with a probability $\lambda=1/m$ at each draw. When any of the firms draws the red ball, the diffusion process begins (that is time 0). The number of adopting firms will then increase at a rate λ_i, where i is the number of adopters (notice, however, that the rate of *discovery* of a new technology for a single firm at any draw is still given by λ, but now the overall rate of *adoption* of that new technology is higher, and given by λ_i, since someone else has already "discovered" it). Now, suppose there are only ten firms[6], i.e. $m=10$. Thus, we will have (see also Figure 2.2):

$$\lambda_1 = \frac{9}{10}; \quad \lambda_2 = \frac{16}{10}; \quad \lambda_3 = \frac{21}{10}; \quad \lambda_4 = \frac{24}{10}; \quad \lambda_5 = \frac{25}{10};$$
$$\lambda_6 = \frac{24}{10}; \quad \lambda_7 = \frac{21}{10}; \quad \lambda_8 = \frac{16}{10}; \quad \lambda_9 = \frac{9}{10} .$$

As we can see, the diffusion rate increases from 0.9 up to 2.5 (corresponding to $n=m/2$) and then decreases. This is very interesting, as it captures quite well the essential feature of actual innovation diffusion processes: slow at the beginning (few adopters for a technology not yet established), then faster, then slower again (last come those firms who find more difficult to invest in new technologies). We can also compute the time it takes until the total population has switched to the new technology. From (2.10) and (2.11) above we have that it will take, on average, 1.11 time units to have one adopter, 1.74 time units to have two adopters, and so on, and 5.658 time units for the whole population of ten firms to follow the innovative leader. From (2.13) we can also see that, if the number of firm is large, say $m=10000$, it will take 18.42 time units, on average, for the whole population to adopt a new technology introduced by a single firm[7].

The diffusion process we have just described has several interesting features. It is a stationary process, albeit of finite length, as opposed to the model sketched in Section 2.2 above, which was nonstationary and unbounded. It is a non-linear process, as the size of the

---

[6] This implies that for each one of the ten firms, the probability of a successful research effort is 1/10. Now, this is a really high probability, if we consider that almost surely at least one out of ten drawings will be the right one. In fact, with $n=10$, $P_{0,n}=0.2105$, with $n=20$, $P_{0,n}=0.088$, and with $n=30$, $P_{0,n}=0.048$, i.e. the probability of no success in n trial decreases as n increases. If each firm pays a ticket corresponding to three drawings, the probability that nobody draw a red ball is less than 5%.

[7] It is important to point out that, because of the way this mechanism is designed, everybody actually benefits from the research effort of others. If in fact only one firm invests in R&D, sooner or later it will be successful, even though it will take much longer for the whole system to gain from the new technology.

population increases first and then decreases, whereas the former was linear. It allows for interaction between innovators and followers (imitators), which certainly ought to be capture by any reasonable model of technological diffusion.

2.4. Multiple technological paradigms: the diffusion process

Let us now consider the case of multiple successive technological paradigms. That is, we suppose that once a new technology has been devised by a firm, the latter then begins searching for a newer one. The stochastic process gets unavoidably more complex, but is still tractable, if we treat it as a *birth-death process*. We can define the *state* of the process by the number of firms adopting a given technology. Suppose we have just three possibilities: there is a standard technology (tech A), a new technology which by the end will overtake the old one (tech B), and a newer one (tech C). We focus on tech B and we define the *state* of the process by the number of firms adopting that new technology (i.e. firms that have switched from the old A to B) who have not found a *newer* technology yet. As before, the adoption of a new technology is irreversible, so that once a firm has switched to a new tech it sticks to that forever. Also, technological laps are not allowed, i.e. to go from A to C a firms has to go through B. Hence, we can denote as birth rate the rate at which firms switch to tech B (from tech A) and as death rate the rate at which firms switch to tech C (from tech B). Thus, we start at time 0 with a population of m firms consisting of *one* innovator ("infected") and m-1 potential adopters ("susceptibles"). This means that the diffusion process begins when anyone out of the m individual firms has discovered the new technology (after drawing the red ball from the urn). Once a firm has adopted the new technology, it will start drawing from a different urn, where the choice is between the new technology and a newer one (we assume the new urn has, say, g green balls for tech C and r red balls, so that the probability of getting a green ball for tech C is equal to μ). Thus, of the original population, only a fraction of those who have become "infected" will have a chance to access to the newer technology (this is also to avoid technological laps). Once the newer technology is discovered, those firms will "exit" the process, which thus concerns *only* the number of firms adopting the new technology (tech B) and *not* the newest (tech C). Hence, if we let $X(t)$ denote the number of firms adopting tech B at time t, the sequence $\{X(t),\ t \geq 0\}$ is a pure birth and death process with birth and death rates λ_n and μ_n, respectively:

$$\begin{aligned}
&\lambda_n = (m - n)\, n\lambda\,, && \mu_n = 0\,, && n = 1; \\
&\lambda_n = (m - n)\, n\lambda\,, && \mu_n = (n - 1)\,\mu\,, && n = 2, \ldots, m-1; \\
&\lambda_n = 0\,, && \mu_n = (n - 1)\,\mu\,, && n = m\,.
\end{aligned} \tag{2.14}$$

This is a birth and death process[8] with n states, $n=1,\ldots,m$, for which it is possible to determine the limiting probabilities of each state, which give the long-run proportion of time the process is in state n, meaning the long-run proportion of time we have n firms adopting tech B. By solving the balance equations (see Ross (1983, p. 152)) we get the following:

$$P_1 = \frac{1}{1 + \sum_{n=1}^{m} \left[\dfrac{\prod_{i=1}^{n-1} \lambda_i}{\prod_{i=2}^{n} \mu_i} \right]} = \frac{1}{1 + \sum_{n=1}^{m-1} \left[\dfrac{(m-1)!}{(m-i-1)!} \left(\dfrac{\lambda}{\mu} \right)^{i} \right]}\,; \tag{2.15a}$$

for $n = 1$;

[8] The rationale for such a death rate specification is that only the n firms that have already switched to tech B can go for trying to switch to tech C, but by innovating only and not by diffusion (as in Section 2.1), and they will do so at the individual rate μ. The way the death rate is specified thus corresponds to a Yule process, albeit with a different rate, i.e. firms exit the process according to the birth rate of a non-stationary process. This might seem contradictory, but the reason is that the whole process, and hence its birth and death rates, must depend only on n to be modeled as such in a closed form. Otherwise, it would become a branching process with a potentially infinite number of branches.

$$P_n = \frac{\prod_{i=1}^{n-1} \lambda_i}{\prod_{i=2}^{n} \mu_i} P_1 = \frac{\dfrac{(m-1)!}{(m-i-1)!}}{1 + \sum_{n=1}^{m-1} \left[\dfrac{(m-1)!}{(m-i-1)!} \left(\dfrac{\lambda}{\mu} \right)^i \right]} ; \qquad (2.15b)$$

$$\text{for} \quad n = 2, \ldots, m-1.$$

Consider again the case of $m=10$ and $\lambda=\mu=0.1$ of the example above. In this case:

$P_1 = 1.013 \cdot 10^{-5}$; $P_2 = 7.299 \cdot 10^{-5}$; $P_3 = 51.09 \cdot 10^{-5}$; $P_4 = 306.5 \cdot 10^{-5}$;
$P_5 = 0.015$; $P_6 = 0.061$; $P_7 = 0.184$; $P_8 = 0.368$; $P_9 = 0.368$.

That is, the limiting probability that n firms out of ten will be adopting technology B is clearly increasing in n, and it implies that 99.6% of the time there will be more than 5 firms adopting technology B, 98% of the time there will be more than 6, and 91.9% of the time there will be more than 7 firms. Only 0.00001% of the time there will be just one innovator alone. This simple example is illustrated in Figure 2.3. Hence, the *average* number of firms adopting technology B will be:

$$\sum_{n=1}^{m-1} n \, P_n = 7.99 \approx 8 \, .$$

Also, if we want to know the long-run proportion of time that a given firm sticks to the old technology A before adopting technology B, we can compute the equivalent limiting probability of its adopting technology A:

$$\begin{aligned} P[\textit{firm adopts } A] &= \sum_{n=1}^{m-1} P\big[\textit{firm adopts } A \mid n \textit{ firms adopt } B\big] \, P_n \\ &= \sum_{n=1}^{m-1} \frac{m-1-n}{m-1} P_n = \sum_{n=1}^{m-1} P_n - \frac{1}{m-1} \sum_{n=1}^{m} n \, P_n \\ &= 1 - \frac{1}{m-1} \sum_{i=1}^{m-1} n \, P_n = 0.11 \, . \end{aligned}$$

3. Technological change, production, and investment in R&D

To cast the analysis of technological innovation and innovation diffusion in terms of a theory of firm investment we need a few introductory considerations. We begin by assuming that there exists a large finite number of firms of the same kind, starting with a given initial technology. Output is produced according to a very simple production function linear in the labour input, which is the only input, whose wage rate is given to firms. This implies that firms use only circulating capital: technological progress, which is due to firms investment (R&D expenditure), is *disembodied* in this model, and comes in the form of productivity increases of the labour input. A single output good, whose overall demand is given, is produced and sold at a given price. If let $A(t)$, and $V(t)$ denote, respectively, the firm's nominal equity (i.e. the firm market value) and cash flow at time t, we then have that the firm's value rate of return at t is:

$$R(t) = \frac{V(t)}{A(t)} + \frac{\dot{A}(t)}{A(t)} \tag{3.1}$$

from which:

$$R(t)A(t) = V(t) + \dot{A}(t) \tag{3.2}$$

or:

$$\dot{A}(t) = R(t)A(t) - V(t) \; . \tag{3.3}$$

By solving (3.3) forward we would get the standard firm maximization problem, which states that, at any time t, firms maximize the present value of the stream of future cash flows subject to a production function constraint and a technological accumulation constraint. In the simplest setting, we take the rate of return on equity capital as given, i.e. $R(t)=R$ for all t. Now, cash flow at time t is given by:

$$V(t) = Pq(t) - Wl(t) - I(t), \tag{3.4}$$

where P, $q(t)$, W, $l(t)$, and $I(t)$ are, respectively, the good price, the output good, the nominal wage rate, the quantity of labour employed, and R&D expenditure (in nominal terms) at time t. If we let R&D expenditure be proportional to the firm's current cash receipts (gross profits), $I(t)=\psi V(t)$, where $0 \leq \psi \leq 1$, then:

$$I(t) = \psi V(t) \quad \Rightarrow \quad V(t) = (1-\psi)[Pq(t) - Wl(t)] \tag{3.5}$$

that is, profits are gross cash flows net of R&D expenditure by a factor ψ[9]. Therefore, after substituing into (3.3) we get:

$$\dot{A}(t) = RA(t) - (1-\psi)[Pq(t) - Wl(t)] \ . \tag{3.6}$$

We ought to point out, at this point, that, for the sake of uncovering the effect of technological change on output dynamics, we can not cast the problem in terms of a standard optimal control program. The reason is that not only the technological accumulation constraint implies a recasting of the optimal program in an different form, but it is the intrinsic stochasticity of the accumulation equation which would make it *stochastic* in the first place (hence, at best, a stochastic optimal control problem). As it will be clear from the transitional dynamics to the steady-state (Section 5, Part II), the endogenous stochastic dynamics thereby generated from the macro effects of small (random) perturbations at the micro level do alter the definition of a steady-state as a long-run equilibrium position. *The*

---

[9] Since in this model we do not have physical capital, R&D expenditure is basically investment in new "equipment", i.e. it is not technological "maintainance" of the labour input but the equivalent of an increase in the "capital stock" by means of productivity improvements. Yet, this form of investment is not *directly* for production purposes (as production takes place anyway, even if R&D expenditure is zero) and thus it is not a *production cost*: we assume it is proportional to current cash flows (which can be thought of as approximating the stream of future cash receipts).

stochasticity makes the equilibrium as a long-run concept meaningless, as the system tends to stay off the equilibrium most of the time independently of its saddle path stability. Two factors will be responsible of these features: the shift in productivity, and hence in the production function, and its randomness, due to the unpredictability of the research effort. Even though firms are posited equal to each other, in that they share the same technological endowment at time 0, they do invest in R&D, hoping to access to a better technology. No matter how we model the devising of a new technology, we know that in the long run the latter will be available to everybody: it is the very fact that at the beginning of the process, i.e. when the technology is first implemented and only one or few firms have access to it, that generates temporary profits. However, since technological innovations will arrive at a random rate, this will generate different adjustment paths for each firm and a fluctuating growth path at the aggregate level. Thus, there is a *technological uncertainty* which will make the growth path a stochastic one: firms only know their *expected* productivity level in advance, but not the actual one (see Figure 3.1).

Once we knew the conditions for (global or local) stability, we would look after the dynamic paths generated by technological advances, whether stochastic or not, whose effect will be that of a continuous sequence of impulses, like those of an external force pushing a pendulum away from its equilibrium. Obviously, under such conditions, it is more interesting to look at the motion of the pendulum than at its stationary state. Cyclical oscillations around the steady state growth level due to technological shifts are such fluctuations around the stationary equilibrium levels, with a difference: shifts in productivity are endogenously generated, albeit randomly, in this context (see Figure 3.2).

Consider, for instance, the technological urn model depicted in Section 2.1 above. Firms draw from the urn according to their R&D expenditure. Thus, each firm will produce

q units of output using l_a units of labour by means of the *old* technology (tech A) and using l_b units of labour by means of the *new* technology (tech B), where $l_a > l_b$. Hence, if:

$$q(t) = \phi_a \, l_a(t) \, , \tag{3.7}$$

where ϕ_a is the amount of labour required to produce q units of output with technology A, and:

$$q(t) = \phi_b \, l_b(t) \, , \tag{3.8}$$

where ϕ_b is the amount of labour required to produce q units of output with technology B, then, for a given level of labour (e.g. the "old"), we will have an increase in the level of output equal to ϕ_b/ϕ_a. Now, if we assume that the increase in productivity is constant, we will have that:

$$\frac{\phi_b}{\phi_a} = \frac{\pi \phi_a}{\phi_a} = \pi = constant, \qquad \pi > 1 \, . \tag{3.9}$$

If we let p denote the probability of switching to a new technology (e.g. the probability of drawing a red ball), then the *expected* level of output at time t will be:

$$\begin{aligned} E[q(t)] &= pq_b(t) + (1-p)q_a(t) = [p\pi + (1-p)]q_a(t) \\ &= [1 + (\pi-1)p] \, q_a(t) = q_a(t) + (\pi-1)p \, q_a(t) \, . \end{aligned} \tag{3.10}$$

Hence, at any time t, the firm level of output is given by the sum of the output attainable with the current technology and the output attainable with the more efficient technology, weighted by the probability of having access to it. Obviously, we are assuming that both the productivity increase (π is a constant) and the probability of success (p) are known in advance. The relevant uncertainty lies in the fact that it is the time of inception of the new technology that is not known *to the individual firm*. Thus, as all individual paths will be different, there will fluctuations around the *average* growth path.

If ψ is fixed, the expected value of output is uniquely determined, for any given π and p (there is basically a once-for-all jump). Conversely, if ψ is a choice variable, for the investment structure depicted above to make sense, we might posit a positive relation between the probability of a successful drawing p and R&D expenditure. Thus, if p is a function of ψ asymptotically bounded from above, in the sense that:

$$0 < p(\psi) < 1; \qquad \frac{\partial p(\psi)}{\partial \psi} > 0; \qquad \frac{\partial p(\psi)^2}{\partial^2 \psi} < 0, \qquad (3.11)$$

then we would choose a ψ value such that $\partial V^*/\partial \psi$ is zero, i.e. V^* is maximized with respect to ψ also (see Figure 3.3).

This simple setting concerns a unique and once-for-ever switch in technology. A richer and more interesting case is that of a continuously changing technological environment, where new technologies arrive at a given rate, which may be changing too. Consider a fixed large number m of firms, each having a positive small probability p of devising a new technology i in any given short time interval[10]. Then the arrival of new technologies is a *Poisson process* with rate $\lambda = mp$. Hence, technology evolves according to a Poisson law, and so does output. But then, random shifts from technology i-1 to technology i will occur at times exponentially distributed with mean $1/\lambda$. Now, if we keep assuming that each *new* technology implies a constant percentage productivity increase π with respect to the *previous* one, i.e. $\phi_i = \pi\phi_{i-1}$ then productivity changes according to the same Poisson law, that is, the probability that an increase in productivity takes place during a sufficiently short time interval equals $\pi\lambda$ for all i. Thus, the average growth rate of productivity is constant:

---

[10] Assuming each technology i has the same probability of being discovered.

this means that, if it would not be subject to random fluctuations, it would vary in accordance with the deterministic differential equation:

$$\dot{\phi}(t) \equiv \frac{d\phi(t)}{dt} = \pi\lambda\phi(t) \quad \Rightarrow \quad \frac{\dot{\phi}(t)}{\phi(t)} = \pi\lambda, \tag{3.12}$$

implying that:

$$\phi(t) = \phi(0)e^{\pi\lambda t}, \tag{3.13}$$

where $\phi(0)$ is the initial productivity value (at time 0). It is readily seen that the expectation of the Poisson process coincides with $\phi(t)$, and thus $\phi(t)$ describes not only a deterministic growth process, but also the expected productivity value at time t. Also, if λ is a function of ψ, $\lambda(\psi)$, and the actual discovery of a new technology for any given firm is, for instance, proportional to R&D expenditure, then new technologies at time t arrive at rate $\lambda I(t)$, that is, $\lambda\psi V(t)$, i.e. a time-dependent rate. Again, technology evolves according to a Poisson law, and so does output. But then, random shifts from technology i-1 to technology i will occur at times exponentially distributed with mean $[\lambda\psi V(t)]^{-1}$.

Once a new technology is implemented, in any case, output shifts from a trend level to a higher level and, for a given price, positive profits will be generated. However, it is easily seen that these profits can only be temporary, as innovation slowly spreads out across firms. Just write the price of output P at t as an inverse demand function:

$$P(t) \equiv P(E[Q(t)]); \quad \frac{dP}{dQ} < 0, \tag{3.14}$$

where E[Q(t)] is expected *aggregate* demand[11], given by:

[11] In the simplest case of just two technologies a and b.

$$E[Q(t)] = \sum_{i=1}^{m-n} q_{i,a}(t) + \sum_{i=m-n+1}^{m} q_{i,b}(t)$$
$$= (m-n)\ q_a(t) + n\ q_b(t)$$
$$= (m-n)\ q_a(t) + n\ \pi\ q_a(t) \tag{3.15}$$
$$= m\ q_a(t) + n\ (\pi-1)\ q_a(t)$$
$$= Q_a(t) + (\pi-1)\ n\ q_a(t)$$

namely the sum of current aggregate output (i.e. output produced with the existing technology *a*) and the extra amount of output produced with the new technology *b*. With *m* firms, aggregate output is the sum of outputs of individual firms, i.e. *m* times individual output, while the extra amount of output is *n* (the number of innovative firms) times the individual increase in output. In equilibrium (when all firms adopt the same technology), the single firm's share of aggregate output is just given by:

$$\frac{q(t)}{Q(t)} \equiv \frac{q(t)}{mq(t)} = \frac{1}{m}, \tag{3.16}$$

while when adopting a new technology it is given by:

$$\frac{q(t)}{Q(t)} = \frac{\pi\ q_a(t)}{m\ q_a(t) + n\ (\pi-1)\ q_a(t)}$$
$$= \frac{\pi}{m+(\pi-1)n} = \frac{\pi}{m-n+\pi n}\ . \tag{3.17}$$

As $\pi > 1$, $q(t)/Q(t) > 1/m$ for all $n < m$ and for all $n_1 < n_2$. Hence, the smaller the number of innovative firms, the highest potential profits for the single innovative firm. If the aggregate demand curve has standard regular behavior and m is large, the price of output will decrease in the long run. In particular, as $Q \equiv (m\text{-}n+n\pi)q$:

$$\frac{dP}{dQ(q)} \equiv \frac{dP}{dQ}\frac{dQ}{dq} = \frac{dP}{dQ}\frac{d((m-n)q+n\pi q)}{dq} = \frac{dP}{dQ}[m+(\pi-1)n] \tag{3.18}$$

and positive temporary profits are guaranteed as far as $n<m$ at *any level* of prices. For the sake of simplicity, assume a linear inverse demand function like:

$$\begin{aligned} P(Q(t)) &= a - bQ(t); \quad a,b>0 \\ \Rightarrow P(q(t)) &= a - b(m-n+\pi n)q(t) \end{aligned} \tag{3.19}$$

so that (3.18) reduces to (see Figure 3.4):

$$\frac{dP}{dQ(q)} = \frac{dP}{dQ}[m+(\pi-1)n] = -b[m+(\pi-1)n] \ . \tag{3.20}$$

A word of caution is of order at this point. Here we keep assuming that firms are sufficiently small as compared to the total size of the market and they are present in a large number (they are "price takers"), although they do affect prices (through output). A better model would admit monopolistically competitive markets, a finite (maybe not large) number of firms and an active pricing policy. But then, entry and exit from the market (relatively to the possibility of positive profits) and even bankruptcy should be taken into account. We leave such model developments to future research (see Ardeni and Gallegati (1993)).

Before we get into the description of the dynamics, which will be illustrated in Part II, it must be pointed out that, as we will see below, the processes we have to tackle are processes where typically *macroscopic deteministic* laws of motion arise, about which the random nature of the technological "shocks" generates a fluctuating part. Most importantly, the deterministic motion and the fluctuations arise directly out of the same description in terms of individual *jumps*, or transitions. It is in this respect that the descriptions in terms of jump processes, as the one depicted above in Section 2, and their corresponding master equations can be very satisfactory, as will be shown in Section 5 below, in modelling growth and fluctuations around growth paths simultaneously. We must have un mind that there exist an intrinsic interaction between output growth, firm growth, and investment, whose effect (and actually whose purpose) is technological change. Technological change implies higher output, which generates higher profits, an increasing number of competitors sharing the same technological level and, in due course, lower and lower marginal profits as the new technology become popular. Such a process can be thought of as following a deterministic law of motion, but random in nature in its timing and occurrence: successful outcomes arrive randomly at given rates, so that fluctuations arise about the deterministic growth path. It is this circular link between the firm growth engine (incentive to innovate) and aggregate growth (the diffusion of innovations) that we will try to exploit in the description of the dynamics in the next two Sections.