

MERIT-Infonomics Research Memorandum series

General Purpose Technologies and Energy Policy

Adriaan van Zon & Tobias Kronenberg

2005-012



*MERIT – Maastricht Economic Research
Institute on Innovation and Technology*

PO Box 616
6200 MD Maastricht
The Netherlands
T: +31 43 3883875
F: +31 43 3884905

<http://www.merit.unimaas.nl>
e-mail: secr-merit@merit.unimaas.nl



International Institute of Infonomics

c/o Maastricht University
PO Box 616
6200 MD Maastricht
The Netherlands
T: +31 43 388 3875
F: +31 45 388 4905

<http://www.infonomics.nl>
e-mail: secr@infonomics.nl

General Purpose Technologies and Energy Policy

by

Adriaan van Zon and Tobias Kronenberg

(MERIT, March 2005)

Abstract

We employ a general purpose technology model with endogenous stochastic growth to simulate the effects of different energy policy schemes. An R&D sector produces endogenous growth by developing radical and incremental technologies. These innovations result in blueprints for capital intermediates, which require raw capital and either carbon or non-carbon-based fuels. A carbon tax therefore affects not only the final production sector but also the R&D sector by making the development of non-carbon-based technologies more attractive. Due to path dependencies and possible lock-in situations, policy can have a significant long-term impact on the energy structure of the economy. Allowing for different elasticities of substitution between consumption and environmental quality, we examine the effects of different carbon policies on growth, environmental quality, and welfare. We find that an anti-carbon policy may reduce welfare initially, but in the long run there is a strong potential for a ‘double dividend’ due to faster growth and reduced pollution.

Keywords: general purpose technology, carbon tax, R&D, growth, carbon fuel consumption

1. Introduction

The current bad news about the adverse impact of human activity on the environment¹, stresses the need to reconcile further economic growth with environmental protection. Since the lion's share of environmental damage is caused by the consumption and distribution of energy, a decoupling of growth from pollution will require a massive reorganisation of the energy system, possibly leading towards a hydrogen-based economy (Rifkin, 2002). We are interested in finding out how economic policy may help to promote the transition from the current energy system to a sustainable alternative. There have been several energy transitions in the past, all of which were the result of technological breakthroughs. It thus seems reasonable to expect that technological progress will also be a *conditio sine qua non* for future energy transitions.

Technological progress has never been a smooth process. Schumpeter (1939) suggested that drastic innovations may give rise to technological waves, causing long-run cycles in GDP growth. A *drastic innovation* is a radically new idea which is reached after deliberate efforts at combining previously unrelated ideas². Such drastic innovations are a risky enterprise from an economic point of view, because they entail high risks due to the uncertainty about the actual usefulness of the idea. But from a technological point of view, the most important characteristic of a drastic innovation is that it opens up new opportunities for further development and expansion of the field of application of this radically new idea. The concept of a drastic innovation is, therefore, closely related to what Bresnahan and Trajtenberg (1995) and Helpman (1998) refer to as a "General Purpose Technology" (henceforth GPT). Similar notions of new ideas that are at least drastic enough to eliminate technological implementations of 'old' ideas through creative destruction, have been put forward by Aghion and Howitt (1992, 1998), who show how the uncertainty surrounding the research process itself may lead to cyclical growth patterns.

¹ See, for example, the February 2005 issue of the Scientific American, or the August 2004 issue of Physics Today, that both describe the detrimental effects of global warming on the Arctic region, including the North-pole and the Greenland ice-sheet. Especially the melting of the latter will pose severe problems for countries at an elevation slightly above or even partly below current sea-levels, like the Netherlands for example, possibly in a few decades from now.

² This phenomenon has also been described as "micro-innovations" in Mokyr (1990), "refinements" in Jovanovic and Rob (1990), and "secondary innovations" in Aghion and Howitt (1998, ch. 6).

Drastic innovations are usually contrasted with *incremental innovations*, which refine existing knowledge in a predictable fashion and which are generated when entrepreneurs combine older insights that are closely related from a technological point of view. The costs and risks associated with these incremental innovations activities are relatively low compared to those associated with drastic innovations. An important feature of incremental innovations is that they are highly path dependent, following specific technological trajectories (Dosi, 1988) that define, in our terminology, a so-called ‘technology family’ with the drastic innovation serving as its core.³ An endogenous growth model that is build around the expansion of such a technology family is the Romer (1990) model, in which innovations are effectively also of the incremental type, as technical change is embodied in the ensemble of intermediate goods rather than in just the latest new intermediate as in the Aghion and Howitt model.

In order to incorporate the crucial features of a GPT into one model, we have to allow for both vertical innovations (drastic innovations that form the core of a new GPT or technology family) in the sense of Aghion and Howitt (1992,1998) and horizontal innovations (incremental innovations that raise the productivity of an existing GPT) in the sense of Romer (1990). In both dimensions there will be Love of Variety effects, and researchers will have to decide to engage in either basic research or applied research. In our perception of R&D driven technical change, basic research is aimed at finding new cores that form the heart of a new GPT, while applied R&D is aimed at expanding the field of application of the core through the addition of ‘peripheral’ innovations. Thus, each successful basic R&D project gives rise to follow-up applied R&D projects.

It seems to us that this technology family perception is directly relevant in the context of an energy transition, as any conceivable energy system will be composed of a number of different technologies that together constitute the system as a whole. For the hydrogen economy as envisioned by Rifkin (2002), several existing technologies (most notably fuel cells and complex IT systems) must be combined to form a radically different energy distribution system. The hydrogen system can thus be seen as a radical innovation, and a number of incremental innovations is expected to occur once the system is in place. Therefore, if we want to analyze the transition towards a hydrogen economy, we need to adopt a technology family framework, as outlined above. In such a framework, different

³ Or, in the context of the family metaphor, its ‘founding father’.

technology families should be allowed to coexist at the same time, as it is the case in practice. This has the added bonus of allowing for a relatively smooth transition.

In order to implement this idea, we will use the GPT model by van Zon, Fortune, and Kronenberg (2003), further referred to as the ZFK-model, where growth occurs as the result of intentional basic research on the cores of new GPTs and applied research on new peripherals for the further expansion of existing GPTs. The model takes into account the uncertainty associated with drastic innovations by drawing the parameters that determine a technology's characteristics from a random distribution. Depending on these parameters, a new technology may result in a successful GPT with many practical applications, a complete failure, or anything in between these two extremes. Developing a new core of a GPT generates technological complementarities in the sense of Carlaw and Lipsey (2001), because the inventor of the core enables other researchers to begin applied R&D in the new technology.

In order to add an energy dimension to the ZFK-model, we extend it in three ways. First we allow for different types of fuels, a carbon-based one and a non-carbon-based one, and assume that carbon-based fuels generate adverse environmental externalities. As a second extension we introduce two different types of technical progress giving rise to the development of the cores and peripherals of new carbon-based GPTs or non-carbon-based GPTs. Third, we add a utility function describing the trade-off between growth and environmental quality. It should be stressed here from the outset that we do not claim the environmental features or the energy part of the model to be sufficiently realistic at this stage. Adding such realism makes sense to us only if the general behaviour of the model as that is driven by cycles in the allocation of R&D activity and the path dependencies that arise from the stochastic nature of the R&D process with respect to both the arrival times and the overall productivity characteristics of new GPTs, seems to be right, as we think it does. In order to illustrate this, we perform some simulation experiments using the extended ZFK-model in which we examine the impact of three different policy reforms on the allocation of R&D efforts, growth, and environmental quality.

The extended ZFK-model has some features in common with the Matsuyama (1999) model, that also generates endogenous cycles in long-run growth. In Matsuyama's model one phase is characterized by capital accumulation, which runs into diminishing returns and thus cannot generate long-run growth. This Solowian phase is followed by a Schumpeterian innovation phase, and the interplay between the two phases generates long-run growth. The major achievement of Matsuyama's model is the integration of neoclassical capital

accumulation and Schumpeterian innovation-based growth. The ZFK model takes that integration a step further, but here it is the expansion of GPTs through applied R&D that eventually runs into diminishing returns, thus inducing a shift towards basic R&D, which in turn generates new possibilities for applied R&D. In short, the ZFK-model emphasizes the Schumpeterian aspects of growth, but also the role of R&D in both accumulation processes, as well as the uncertainties surrounding these processes, apart from disaggregating R&D activity itself into basic R&D and as many applied R&D processes as there are GPTs.

The extended ZFK-model also contributes its own specific features to the literature that are not present in the Matsuyama model. First, it introduces asymmetries in the intermediate goods market. The corresponding asymmetries in profit opportunities in the intermediate goods sectors that implement the ideas produced by the R&D sector, are at the heart of the interplay between basic and applied R&D, where, as Yetkiner (2003) has pointed out, falling profits provide the real incentives for R&D to find the next completely new technology. Secondly, we have cyclical growth that does not depend on the reallocation of homogeneous labour between production and R&D activities, as it is the case in the standard GPT approach by Bresnahan and Trajtenberg (1995), or in more recent work focusing on the clustering of R&D activity against a GPT-background that explains booms and busts in the business cycle (Francois and Lloyd-Ellis , 2003), or in (Matsuyama, 1999). Rather than this being a weakness of our model, we feel that it is actually a strong feature, as it is hard to imagine that the ‘eye-hand-coordination’ in the words of Romer (1990) is suited in practice to perform R&D tasks, and the other way around. Third, given the significant effects of technological change per se on economic growth, a better understanding of the reasons behind the cyclical evolution of output and technology is important from a policy perspective. Fourth, our model elaborates on the role of basic and applied R&D mechanisms in the growth process. It shows that the influence of these two R&D types on the long-run growth process is very different indeed. Fifth, the model shows that in a competitive equilibrium the allocation of researchers between basic and applied R&D will generally be inefficient. We will show that this inefficient allocation of R&D may give rise to situations in which the introduction of a carbon tax may yield a double dividend.

The latter findings are in line with Smulders and de Nooij (2002), who argue that if the allocation of R&D is inefficient, energy policy should be combined with a technology policy addressing the inefficiency in the R&D sector. We refine this analysis by distinguishing not only between two directions of technical progress (carbon-based versus non-carbon-based) but also between basic R&D and applied R&D. But in addition to this, our

simulations also show that the effect of any policy scheme depends crucially on the existing technology structure causing policy making to become highly path-dependent. If we are in a situation where carbon-based technologies form the dominant paradigm, then a carbon tax will – as a side-effect – move R&D resources away from applied R&D on the existing technologies towards basic R&D on new technologies, and this side-effect will reduce the market failure in the R&D sector. Thus the carbon tax alone tends to reduce carbon-based fuel consumption and increase growth. Tax recycling as a subsidy on non-carbon-based fuels will reinforce the carbon reduction, and tax recycling as a subsidy on basic R&D will reinforce the growth effect, and depending on preferences, any of these schemes could be preferable.

The organisation of the rest of this article is as follows. In section 2 we provide a short summary of the most important features of the ZFK-model, and our energy extensions of that model. Sections 3 and 4 show the results of some illustrative simulation exercises we have performed with the extended ZFK-model. Section 3 is devoted to the base run, and section 4 to some fiscal policy experiments. Finally, section 5 contains some concluding remarks.

2. The ZFK-model in outline

2.1 General overview

As the Romer (1990) model, the extended ZFK-model consists of a perfectly competitive final output sector that combines labour and effective capital to produce output. Effective capital is an aggregate of individual GPTs that in turn consist of intermediates supplied to the final output sector under imperfectly competitive conditions. Intermediates are built according to the specifications from the blue-prints obtained from the R&D-sector that uses the market-value of these blue-prints to pay for the resources used in producing them. The resources under consideration are R&D labour and R&D entrepreneurship.

The most important problem addressed in the ZFK-model is the ‘spontaneous’ arrival of new GPTs as the result of basic R&D, and the subsequent expansion of that GPT through applied R&D. Both types of R&D are profit-incentive driven, as is usual in new growth models a la Romer (1990) and Aghion and Howitt (1992, 1998). Both types of R&D activity alternate, as we assume that the very first component(s) of a GPT are the most important, and so they are invented first, while additional components invented through applied R&D grows

less and less ‘promising’ from a profitability point of view. This gives rise to decreasing returns to variety within a GPT that increases the relative attractiveness of finding the core of a new GPT rather than continuing the expansion of existing GPTs by finding still newer peripheral components. The decreasing returns to variety are a relative novelty that is also used in van Zon and Yetkiner (2003), and that hinges on the existence of asymmetries in the contribution of individual intermediates to their effective ‘ensemble’. In our case, this implies that the contribution of new peripherals to a GPT in productivity terms will be falling with the number of components of a GPT. Because of the latter, and assuming that profits generated by selling the peripheral components of a GPT are captured by the applied R&D sector, applied R&D becomes less and less attractive relative to basic R&D, resulting in a reallocation of R&D effort towards basic R&D. The latter in turn generates a higher probability that the core of a new GPT will actually arrive. Thus the expansion of a GPT will eventually lose momentum. But this loss in momentum also increases the incentive to find a new core. In addition to this, it may well be the case that the core of a new GPT is not promising at all. Then our set-up implies that the R&D sector will devote most of its resources trying to find a new core rather than producing peripheral inventions for a failing GPT. Thus we get cyclical growth patterns, implied by inherent technological expansion limits, rather than by reallocations of resources between the final output sector and the R&D sector as one usually finds in GPT models.

Growth in the ZFK-model then is caused by symmetric horizontal innovation which increases the number of GPTs and quasi-vertical⁴ innovation within each GPT, which effectively breaks the symmetry between GPTs again. Neither innovation is so drastic that it drives out older GPTs completely, however they can be asymptotically drastic.

2.2 The production structure

With respect to the production structure, we assume that a representative firm in the final goods sector produces output by using labour and effective capital that consists of a set of GPTs combined in an Ethier function (Ethier, 1982). Each GPT contributes symmetrically to the effective capital stock. Each GPT itself is a composite input made out of different components, i.e. its core and peripherals. These components contribute to the GPT in an

⁴ Because new peripherals contribute less to GPT productivity than ‘older’ peripherals by assumption, we actually have something looking like an inverted quality-ladder.

asymmetric fashion. As we assume that researchers roughly know what to look for and where to look for it, the core and then the most productive peripherals are developed first, while subsequent peripherals contribute less and less to the productivity of the technology as a whole. We therefore have:

$$Y = L_y^{1-\alpha} \cdot K_e^\alpha \quad 0 < \alpha < 1 \quad (1)$$

where Y represents final output, L_y is production labour, and K_e is the effective capital stock, that is itself defined in terms of individual GPTs as:

$$K_e = \left(\sum_{j=1}^A z_j^\alpha \right)^{1/\alpha} \quad (2)$$

where A is the number of GPTs currently active, and j indexes those active GPTs. z_j represents the “volume” of GPT j . Equation (2) is a CES function with a symmetric contribution of all factors z_j to the level of effective capital K_e .

An important aspect of the model is that inputs are not perfect substitutes: GPTs are better substitutes for each other than labour and the effective capital aggregate are. The fact that the elasticity of substitution between GPTs is bigger than 1, is responsible for the Love of Variety feature of the effective capital stock.⁵ It should be noted that equations (1) and (2) provide a very simple structure that can easily be generalised.⁶ But at present we are only interested in showing how the model works in principle, and learning about its qualitative policy implications.

GPTs in turn consist of a core and of peripherals that contribute less and less to the productivity of the GPT as a whole. An obvious candidate to describe the inner-structure of z_j is again a CES function:

$$z_j = \left(\sum_{i=0}^{A_j} c_{i,j} \cdot x_{i,j}^{\beta_j} \right)^{1/\beta_j} \quad (3)$$

⁵ This follows from the fact that the upper-level is a Cobb-Douglas function.

⁶ See van Zon, Fortune and Kronenberg (2003) for more details.

where x_{ij} for $i > 0$ represents the i -th peripheral of GPT j and x_{0j} represents the size of the core of GPT j .⁷ $c_{i,j}$ are the standard distribution parameters one normally uses with a CES function, while $1/(1 - \beta_j)$ is the elasticity of substitution between all the components of GPT j . A_j is the number of peripherals belonging to GPT j .

In the very early stages of a GPT, A_j can actually be equal to zero, and in fact, even in the medium and long run, A_j can remain equal to zero, thus underlining the ex post character of what we consider to be real GPTs i.e. technologies with a large number of peripherals. For GPTs with a small number of peripherals A_j , we will use the term “failed GPTs” from now on.

In order to simplify matters as much as possible, we will make the following assumptions:

$$\beta_j = \alpha \quad \forall j \tag{4.A}$$

$$c_{i,j} = c_{0,j} \cdot (\zeta_j)^i, \quad 0 \leq \zeta_j \leq 1 \tag{4.B}$$

Assumption (4.A) reduces the three-level organisation of the production process effectively to a two-level production function with asymmetric contributions of all components of all GPTs to final output, while (4.B) puts a technically useful structure on the distribution coefficients of the implied aggregate production function. This structure allows us to write the production function in terms of an aggregate of functions of the cores of the various GPTs only, so from a practical point of view we can essentially forget about individual peripherals. This is a big technical bonus since we have to deal with a number of “real GPTs” simultaneously. (4.B) states that the CES distribution coefficients are geometrically declining with the peripheral index.

The demand for GPTs is derived very much as in Romer (1990). Assumption (4.A) and equations (3) and (2) when substituted into (1), give rise to the following inverse demand

⁷ The inner-organisation of a GPT as a CES aggregate of core and peripherals with Love of Variety features is in our case responsible for the “innovational complementarity” character of GPTs as advanced by Bresnahan and Trajtenberg (1995), that is “the productivity of R&D in a downstream sector increases as a consequence of innovation in the GPT technology”. In fact the notion of complementarity hides the interpretation of productivity increases due to improved division of production tasks between intermediates in a setting where individual intermediates are direct substitutes for each other.

equations for each individual component, under the assumption of perfect competition in the final output market:

$$x_{i,j} = L_y \cdot \left(\frac{p_{i,j}}{c_{i,j} \cdot \alpha} \right)^{-1/(1-\alpha)} \quad (5)$$

Assuming, as in Romer (1990), that the components can be produced using raw capital only, where each component i of GPT j takes η_j units of raw capital k_{ij} to create one unit of the component $x_{i,j}$, the marginal production costs are equal to $\eta_j \cdot r$, where r is the interest rate and where we have ignored the depreciation of capital. Because each component has its own market niche (as described by (5)), the profit maximising rental price of each component is easily obtained as:

$$p_{i,j} = mc_{i,j} / \alpha = r \cdot \eta_j / \alpha \quad (6)$$

which is the familiar Amoroso Robinson condition for profit maximisation under imperfect competition, and where $mc_{i,j}$ is the marginal user cost of component i of GPT j . Using (6) to obtain total profits $\pi_{i,j}$ per component i of GPT j , we find:

$$\pi_{i,j} = L_y \cdot c_{0,j}^{1/(1-\alpha)} \cdot \zeta_j^{i/(1-\alpha)} \cdot (1-\alpha) \cdot \alpha^{(1+\alpha)/(1-\alpha)} \cdot (r \cdot \eta_j)^{-\alpha/(1-\alpha)} \quad (7)$$

Equation (7) has several interesting features. First, profits of a component rise with the overall level of final output production as proxied by L_y . Secondly, they fall with a rise in the production cost of a component (i.e. $r \cdot \eta_j$), while third they fall with the peripheral index i (since $0 < \pi < \zeta_j < 1$). The latter is one of the most important drivers of the overall behaviour of the model. Note too that for constant values of L_y and r , ex post profit flows are constant too. Under these assumptions, the present value of the profit stream associated with using component i of GPT j , i.e. $PV\pi_{i,j}$, would be given by:

$$PV\pi_{i,j} = \pi_{i,j} / r \quad (8)$$

As in Romer (1990) and Aghion and Howitt (1992), we assume that these flows are captured by the respective R&D sectors that created the designs for all individual components.

2.3 R&D activity

The blueprints that are needed to be able to produce the components of a GPT are obtained by the intermediate goods sectors from the R&D sector. We assume that each innovation, whether basic or applied, is the result of innovative activities by R&D labour. The determination of the kind of activity (i.e. basic versus applied R&D), which the research sector engages in, depends on the relative profitability of both types of R&D. Which type of R&D is done depends on the profitability of adding peripherals to an already existing technology (for which the core already exists) versus creating a completely new technology (for which a new core is required). As is clear from (7), the profitability of a peripheral falls the later it is introduced, so that there is a certain point in time at which pursuing basic R&D becomes increasingly more profitable than doing applied R&D, and research labour shifts from doing applied research to doing basic research. However, in order to avoid bang-bang reallocations of R&D activity, we assume that marginal R&D productivity is decreasing for both types of activities.

Our motivation for considering different R&D processes for peripherals and cores is our perception that the invention of a technology core requires “something more fundamental” than the further development of a technology by adding peripherals. We capture this difference by differentiating their contribution to total production. However, we also feel that finding a (core of a) potential GPT is subject to more uncertainty than finding a peripheral once a new technological “proto-paradigm” has arrived in the form of a potential GPT. We model this by assuming that the R&D process itself is uncertain, first because it is not able to predict the exact arrival time of the core of a new GPT, and secondly because it is not able to predict the actual characteristics of a potential GPT, that we assume to be known in actual fact only after the arrival of its core. These expectations pertain to the inherent productivity of the next GPT (i.e. $c_{0,j}$), the user costs of the peripherals (i.e. $r \cdot \eta_j$), the scope for extension (i.e. ζ_j) of the core, associated applied research productivity (i.e. δ_j , cf. equation (9.B) below) and research opportunities (i.e. μ_j , cf. equation (9.A) below).

In order to reflect the uncertainty associated with basic research, we assign random values drawn from a uniform distribution to the characteristics of each GPT, which are unknown until the core of that GPT is actually invented. Applied R&D can only begin after the core has been introduced. At this point in time the GPT characteristics become publicly known, and for reasons of simplicity, we assume that there is no uncertainty regarding the characteristics of new peripherals added to the GPT through applied R&D. Of course, there is still a random element because the actual arrival of an applied R&D invention depends on random draws from a Poisson distribution. However, because we assume risk neutrality, this random element has no effect on the outcomes.⁸

Even though we distinguish between basic and applied R&D, we do model both types along the same lines. R&D gives rise to innovations that arrive according to a Poisson probability distribution with arrival parameter λ . According to this distribution, the probability that an event (arrival of an innovation) occurs within T time units from now, is equal to $F(T) = 1 - e^{-\lambda T}$. As in Aghion and Howitt, we assume that the level of R&D activity directly and positively influences the arrival rate of innovations. But unlike Aghion and Howitt (1992), we assume that there are decreasing marginal returns to current R&D, giving rise to an effective arrival rate of innovations given by:

$$\lambda_j = \mu_j \cdot R_j^e \tag{9.A}$$

$$R_j^e = \delta_j \cdot R_j^\beta \tag{9.B}$$

where λ_j is the arrival rate of innovations associated with R&D process j (process 0 is associated with the basic R&D necessary to find the core of the next GPT, whereas $j > 0$ represent the processes necessary to find the peripherals of GPT j). μ_j is the arrival rate at a unit level of effective R&D, R_j^e . Equation (9.B) shows how effective R&D uses R&D labour R_j with a “productivity” parameter equal to δ_j , and where $0 < \beta < 1$ ensures that the marginal product of R&D labour falls with the level of R&D activity. The latter is another main feature of the model, since it ensures that in combination with (8), the marginal benefits of doing applied R&D and basic R&D are asymptotically falling to zero for increasing levels

⁸ Assuming risk aversion would make basic research less attractive compared to applied research. We will show, however, that a market failure will lead to an inefficient allocation between these two types of research anyway.

of R&D activity. But more importantly, they rise to infinity for levels of R&D activity that fall asymptotically to zero. The latter ensures that there will always be an incentive to employ a non-zero volume of R&D workers on any project, however bleak the prospects for success may be.

The (expected) marginal benefits from doing R&D on the i -th peripheral of GPT j are now given by:

$$MB_{i,j} = \frac{\partial(PV\pi_{i,j} \cdot \mu_j \cdot \delta_j \cdot R_j^\beta)}{\partial R_j} \quad (10)$$

As the free mobility of R&D labour between its various uses implies that the marginal benefits (that are paid as wages to R&D workers) for different R&D activities should be the same, we obtain for the optimum ratio of applied R&D versus basic R&D workers:

$$R_j / R_0 = \left(\frac{PV\pi_{A_i+1,j} \cdot \mu_j \cdot \delta_j}{PV\pi_{0,A+1} \cdot \mu_0 \cdot \delta_0} \right)^{1/(1-\beta)} \equiv \varphi_j \quad (11)$$

where A_i is the number of peripherals of GPT j , and A is the total number of active GPTs. Equation (11) shows that higher (expected) profit flows on a peripheral in some GPT will divert R&D resources into further expanding that GPT. Such an expansion is promoted by an R&D process that is relatively efficient (high δ_j), or a process where innovations are relatively easy because of ample “fishing” opportunities (high value of the non R&D part of the arrival rate, i.e. the parameter μ_j). Because total R&D uses exhaust the available and exogenous supply of R&D workers R , we readily find:

$$R_0 = R / \left(1 + \sum_{j=1}^A \varphi_j\right) \quad (12.A)$$

$$R_j = \varphi_j \cdot R_0 \quad (12.B)$$

A fall in the present value of the expected profit flow associated with peripheral j will therefore decrease the corresponding φ_j and for given R , the corresponding value of R_j will

fall, thus leading to a decrease in the rate of expansion of existing GPTs and an acceleration of the arrival rate of new GPTs.

2.4 Love of Variety again

The organisation of capital in the form of different GPTs with differential impacts on effective capital ultimately results in the growth of capital productivity, hence of output itself, for a given amount of capital and labour. This is easy to see, since after some manipulation of the relation between effective capital and the size of the core of some GPT j , as well as the capital costs of building the core (and the corresponding peripherals), we get:

$$z_j = K_j \cdot c_{0,j}^{1/\alpha} \cdot \left\{ (1 - \zeta_j^{(1+A_j)/(1-\alpha)}) / (1 - \zeta_j^{1/(1-\alpha)}) \right\}^{1/\alpha-1} / \eta_j \quad (13)$$

The “GPT” productivity of “raw” capital K_j depends positively on the size of the contribution of the core (i.e. $c_{0,j}$), which we have already referred to as the intrinsic productivity of the GPT, positively on the scope for extension of the GPT (i.e. ζ_j), negatively on the unit raw capital cost (η_j), but most importantly, it depends positively on the number of peripherals of the GPT, (i.e. A_j). The latter is the Love of Variety effect implied by the concavity of the GPT in its individual components (cf. equation (3)). Love of Variety works at two different levels in this model, therefore, as opposed to Romer (1990). It works at the component level within each individual GPT, but also at the level of the GPTs that are used to produce final output.

2.5 Adding the energy dimension

We introduce energy into the ZFK-model by means of the following modifications. First we distinguish between two types of GPTs that use either carbon-based or non-carbon-based fuels. Secondly, the marginal (user-) cost of the components of these GPTs now also include carbon- and non-carbon-based fuel costs, next to capital cost. In fact, we assume that fuel-capital ratio's at the component level are fixed.⁹ However, these fuel-capital ratio's may

⁹ This adds a Leontieff layer at the component level to the existing multi-level structure.

differ between GPTs as this is another random-characteristic of a (potential) GPT. Third, we introduce a social welfare function that allows us to weigh growth against environmental degradation. To this end, we assume that a representative consumer derives utility from consuming final output and from enjoying the environment. Utility is then a function of consumption and environmental quality:

$$U_t = (\omega_c C_t^\rho + \omega_q Q_t^\rho)^{\frac{1}{\rho}} \quad (14)$$

In this formulation, C_t is consumption, Q_t is environmental quality, the ω 's are the respective weights, and $\sigma = 1/(1 + \rho)$ is the elasticity of substitution between C and Q . Naturally, if σ is relatively large, it is possible to substitute additional consumption for a less healthy environment, so economic growth can raise welfare even if it damages the environment. If σ is small or negative, economic growth will raise welfare only if it does not harm the environment excessively.

For environmental quality we adopt the simplest thinkable formulation:

$$Q_t = Q_t^{\max} - F_t^C \quad (15)$$

where Q^{\max} is the maximum attainable quality, which can only be reached if no carbon-based fuels are used. F^C is the total amount of carbon-based fuel that's being used in production. To simplify matters even more, we assume that pollution does not accumulate. This means that in the model, if carbon-based fuel consumption would be abolished altogether, the environment would return to its maximum quality immediately, which is obviously not a realistic assumption. Finally, we assume that the consumption of non-carbon-based fuels does not harm the environment at all.

2.6 Model closure

The model is closed by assuming that the interest rate is exogenously fixed. In the extended model, we also assume that fuel prices are exogenously determined.

3. Base run simulation results

Due to the complexity of the model a full analytical solution is simply not feasible. Therefore, we illustrate the working of the model by means of a simulation analysis. Since some of the GPT characteristics are relatively abstract concepts, it would be very difficult to find real-world data that could be used directly for calibration purpose. Instead, we choose a set of arbitrary parameter values. This approach allows us to draw only qualitative conclusions for policy analysis.

In order to illustrate the behaviour of the model, we first present the results of a base run. Figure 1 shows how the number of available GPTs rises over time. We assume that the economy starts with just one carbon-based and one non-carbon-based GPT. Figure 2 shows the number of peripherals that are developed for each GPT. Interestingly, some GPTs develop only very few peripherals or none at all. These are ‘failed’ GPTs’, consisting of only a core with few or no peripherals. But there are also ‘true’ GPTs such as C03 or N04, which develop dozens of peripherals.¹⁰ We see these ‘failed’ GPTs and ‘true’ GPTs in our ex post perspective, but the researchers who developed them were seeing them as potential GPTs, because they could only guess at their true characteristics ex ante.

The number of peripherals allows us to say something about the usefulness of a potential GPT, but it does not tell us very much about the economic impact of the GPT. But one important general feature of a GPT is that it affects a large share of the economy. In order to measure this feature we show the Eulerian contribution of each GPT to effective capital K_e in Figure 3 as a measure of the relative contribution to output (through the effective capital stock) of each individual GPT.¹¹ These curves show a cyclical pattern, which is consistent with some long wave views on economic development¹². Freeman and Perez (1988), for

¹⁰ The names C_i and N_j refer to carbon-based GPT i and non-carbon-based GPT j .

¹¹ Based on equation (2), the Eulerian contribution of GPT z to the aggregate stock of effective capital K_e is calculated as $z \cdot (\partial K_e / \partial z) / K_e$.

¹² The curves in Figure 3 are not diffusion curves in the usual sense, but they measure the share of effective capital which is linked to a specific technology, which is a closely related indicator. It should be noted that the expansion of a GPT can be viewed as the diffusion of the application of the core in the economy in a more standard diffusion sense.

example, identify five Kondratieffs in economic history since the late 18th century. Such Kondratieffs are characterised by the dominant GPT of their times, for instance the “steam power and railway Kondratieff” in the mid-19th century or the “information and communication Kondratieff” that started in the late 20th century.




Figure 1

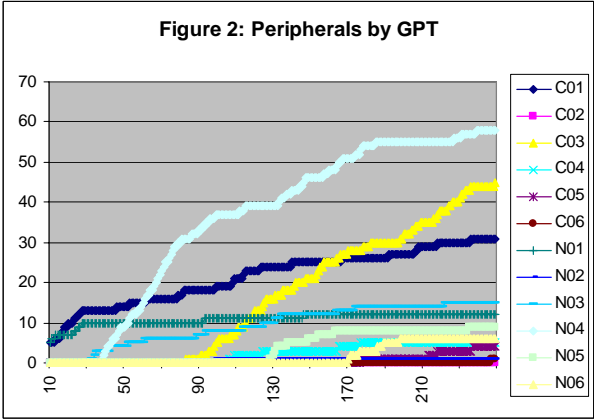


Figure 2




Figure 3




Figure 4

Note that in contrast to many long wave theories, there is nothing mechanistic in the coming and going of Kondratieffs in our model. Due to the structure of the model every GPT will at some time run out of further extension possibilities, and the search for a new GPT begins, but the length of these “long waves” is endogenously determined and highly variable. In the current simulation run, for example, the first Kondratieff is dominated by the initial technologies C01 and N01. It is quickly succeeded by another Kondratieff which is dominated by N04. Around the year 120, N04 is succeeded by C03, which dominates the next

Kondratieff lasting until the end of the simulation period. Other simulation runs show a similar picture. Usually, the length of a Kondratieff is in between 10 and 40 years, but sometimes an extremely successful GPT is invented which remains dominant for 100 years or more.

Figure 4 shows how the dominance of a certain technology determines the economy's fuel mix. During the N03 Kondratieff, non-carbon-based fuel consumption is much higher than carbon-based fuel consumption. During the transition to the C03, however, non-carbon-based fuel consumption levels off, while carbon-based fuel consumption rises quickly. At the end of the simulation period, the economy consumes a balanced mix of both fuels. There is no long-run tendency towards either fuel, because we assume that carbon-based fuel technologies are intrinsically just as productive as non-carbon-based fuel technologies. Thus, the R&D sector has no reason to concentrate on either fuel, and over the long run the economy can be expected to develop just as many carbon-based GPTs as non-carbon-based GPTs.

Figure 5 shows the allocation of R&D workers between basic and applied R&D. We have chosen the number of researchers – plotted on the vertical axis – to be equal to five. Just as in Figure 3, we clearly observe cycles in R&D. Whenever an attractive GPT is invented, researchers move into applied R&D on the new GPT. In year 31, for instance, the freshly invented N03 absorbs almost all of the R&D workers, who are busily developing peripheral applications based on that GPT. As the extension possibilities of N03 are being exploited, further R&D on N03 becomes less attractive, and researchers gradually move back towards basic R&D.

Figure 6 shows the amount of applied R&D on each GPT. Only successful GPTs are shown for the sake of visual clarity. We can see how researchers move away from applied R&D on N04 as that GPT runs out of extension possibilities. After the introduction of a new GPT, there is a “jump” in the R&D sector insofar as the massive movement from basic R&D to applied R&D occurs instantaneously¹³. The Figure can help to explain the continued dominance of C03. It is a GPT with a very large scope for extension, in model terms its ζ is very close to one. New technologies, such as N05 or N06, attract a certain amount of applied

¹³ In the real world, such a massive refocusing of R&D efforts would take more time, since R&D resources are not perfectly mobile as they are in the model, and information about new GPTs may take some time to diffuse. One might account for sluggish R&D labour movement in the model, but this would only complicate matters and divert attention away from the central issues.

R&D for a short time, but as these technologies offer a limited scope for extension, R&D workers move back into applied R&D on C03.




Figure 5




Figure 6

The economic intuition behind these R&D movements can be shown graphically in Figure 7. The horizontal axis represents the (given) total number of R&D workers, who are either working in basic R&D or applied R&D. Moving to the right, the ratio of basic R&D to applied R&D increases, and the other way around. The two curves depict the value marginal product (or marginal benefits (MB)) of employing basic and applied R&D. Since we have assumed diminishing returns ($\beta < 1$), the curves are downward sloping. The point of intersection of both MB-curves determines the profit maximising allocation of researchers, as both profits (and R&D wages if R&D workers are paid their value marginal product, as we assume) can not be increased by moving from the low MB R&D activity into the high MB R&D activity.

If, in this setting, applied R&D is successful, a new peripheral is developed, and we know that the next peripheral will be less productive (because $\zeta < 1$). Thus, the MB-curve associated with applied R&D shifts downward. At the new point of intersection, more researchers are working in basic R&D than before. The bottom half of the figure shows the arrival rate of new GPTs as a function of the number of active basic R&D workers. We can see that as more and more researchers move into basic R&D, the arrival rate increases, and the arrival of the next GPT becomes increasingly more likely. When the next GPT with a lot of development potential actually arrives, new possibilities for applied R&D arise. The VMP curve of applied R&D shifts up (cf. equations (7), (8) and (10)), and the allocation of R&D

workers changes again in favour of applied R&D. Thus, the alternating arrivals of GPT cores and peripherals generate cycles in the R&D sector, with researchers moving back and forth between the two R&D activities, thus generating growth.

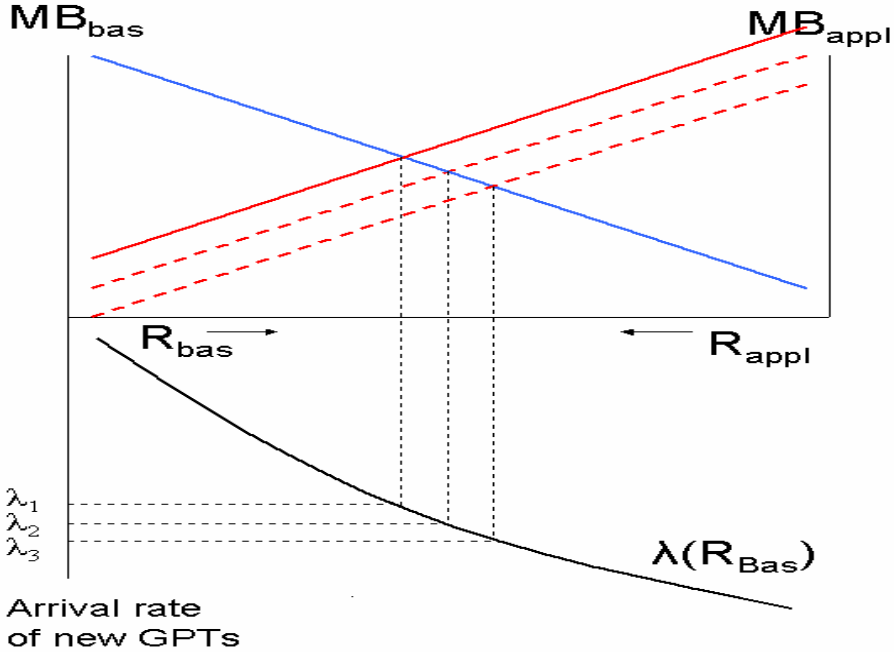


Figure 7

The arrival of successful GPTs generates cycles in the growth rate of the economy as we can see in Figure 8, which shows the growth rate of final output. Successful GPTs have two effects on growth. First, they raise output immediately in the period after their introduction, simply because of our Love of Variety structure. But secondly, and more importantly, they can raise the average growth rate over a period of several decades, as researchers are exploiting the new possibilities for applied R&D.

The actual growth impact of new GPTs, of course, depends on the intrinsic characteristics of the GPT, which in turn determine its pervasiveness, given the ‘general GPT environment’.¹⁴ In the extreme case, i.e. a total failure with no peripherals at all, we have only

¹⁴ This is because the pervasiveness of a GPT depends for a large part on its performance relative to other existing GPTs. If, for some historical reason, the latter are not very productive, then even a mediocre GPT may become pervasive, in accordance with the adagio ‘in the kingdom of the blind, the one-eyed man is king’.

a short-lived growth hike due to the Love of Variety effect. This uncertainty about future growth rates is also a realistic feature of the model: we know with almost complete certainty that a new successful GPT will arrive and that it will set off a period of high growth, but we can never know exactly when this will happen or whether it will actually be the next GPT that turns out to be a big success.




Figure 8

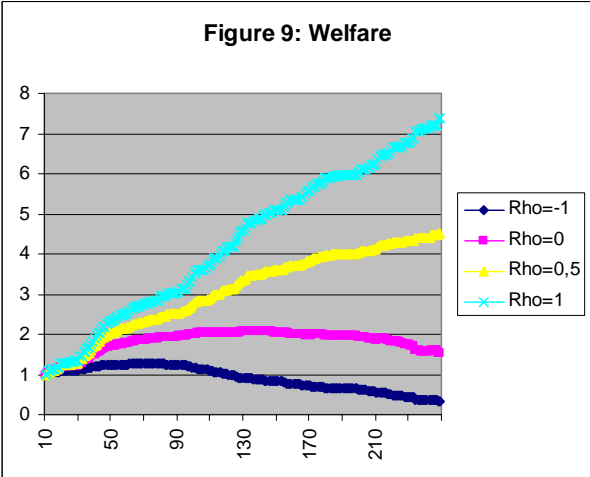


Figure 9

Figure 9 shows the representative consumer’s utility for different assumptions about the substitutability between consumption and environmental quality. For non-positive values of ρ , hence low values of the elasticity of substitution between consumption and environmental quality, the effect of environmental degradation is so large that utility actually starts falling at some point in time. In the extreme case of $\rho=1$, there is no peak in utility, because if consumption and environmental quality would be perfect substitutes, it is always possible to make up the utility loss from environmental degradation through higher consumption. If, however, ρ is smaller than one, the marginal utility of consumption diminishes vis-à-vis that of environmental quality, and it becomes increasingly difficult to raise welfare through ‘dirty’ growth. The lower the value of ρ , the sooner the utility peaks. If $\rho=0$, the utility peak occurs in the year 164, and for $\rho=-1$, it occurs already in the year 78. Thus, if environmental quality and consumption become less ‘easy’ to substitute, it becomes more difficult for increasing consumption to offset the effect of environmental degradation.

4. Fiscal policy experiments

4.1 Introduction

In most Western countries governments are concerned about carbon-based fuel consumption for at least two reasons. First, carbon-based fuels lead to unavoidable CO₂ emissions, which are damaging the environment. Improvements in energy efficiency and CO₂ sequestration may ameliorate this problem, but they cannot solve it in the very long run. Second, most oil deposits are situated in countries that are politically unstable. Therefore, almost all Western governments are implementing policies to reduce their dependence on carbon-based fuels.

However, it is not so much the *amount* of fuel consumption that worries governments and environmentalists; it is primarily the *mix* of fuels. If it were possible to substitute non-carbon-based fuels (e.g. hydrogen produced from non-carbon-based energy sources) for carbon-based fuels (e.g. oil and gasoline), the problem could be solved. One would simply have to discourage carbon-based fuel consumption by means of a carbon tax. The question then is what to do with the tax revenue. Some have proposed to use the tax revenue to subsidise the consumption of non-carbon-based fuels. This would alter the relative price of non-carbon-based fuels, and the resulting substitution effect might lead to the desired change in fuel consumption. Alternatively, one could subsidise R&D in the hope that new technologies based on non-carbon-based fuels are developed.

We now examine the impacts of two different energy policy schemes on growth, the environment, and ultimately welfare. First we show the effects of each policy on the production side, because these do not depend on our assumptions about the utility function. After that, we show how the outcomes of these policy schemes affect the utility of the representative consumer. The graphs associated with the various policies are presented in Figures 10-12 below. These Figures are discussed in subsections 4.2-4.3.

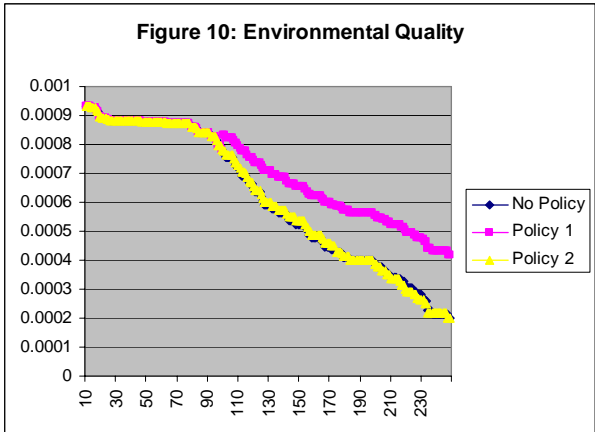


Figure 10




Figure 11




Figure 12

4.2 Policy 1: a carbon tax and non-carbon-based fuel cost subsidy

Our first policy experiment is going to be a tax on carbon-based fuel, which is to be recycled in the form of a subsidy on non-carbon-based fuel consumption. This is going to generate a substitution effect away from carbon-based fuel and towards non-carbon-based fuel. We introduce the policy in the year 100.¹⁵ Figure 10 shows that the policy is effective in achieving the goal of increasing environmental quality. The introduction of the tax/subsidy leads to an immediate fall in carbon-based fuel consumption, and due to the one-to-one

¹⁵ We use the first 100 ‘years’ to get rid of the influence of the initial values of the capital-stock variables in the model.

relationship, environmental quality is immediately improved. As the economy grows, carbon-based fuel consumption grows along, but it remains far below its base run value, and although environmental quality keeps falling, it remains way above the baseline value.

Of course, the government will not only be concerned about carbon-based fuel consumption, it will also be concerned about the policy's effects on output. The data show that the policy leads to an initial drop in output, which is too small to show up in a graph at any reasonable scale. The initial output loss is not surprising because the tax/subsidy creates a distortion in the final output sector. However, this output loss is eventually overcome: In the year 166 output under policy 1 surpasses the base run value. From that year on, output is actually higher than in the base run.

Thus, the government may be able to realize a double dividend in this scenario. Compared to the base run, the policy increases output in the long run and reduces carbon-based fuel consumption immediately. One would have to weigh the present value of the short run output loss against the present value of the long run output gain, and the outcome depends on a subjective discount rate.

Figure 12 shows why there is an increase in output in the long run. The number of GPTs grows faster under policy 1, so that at the end of the simulation period there are now 16 GPTs instead of 14. The reason for the faster arrival of GPTs is that although the policy targets only the fuel market, it also has an effect on the allocation of R&D resources. By lowering the relative price of non-carbon energy, the policy raises the value of a patent on a non-carbon technology. The present value of these patents rises relative to that of the carbon patents. This means, first of all, that there will be more applied R&D on non-carbon GPTs and less on carbon GPTs. Furthermore, since the dominant CO₂ offering the most applied research opportunities is carbon-based, there will be a larger incentive to do basic R&D on finding the next non-carbon based GPT. Thus, there is a shift from applied R&D to basic R&D, and we have already shown that due to the externality of basic research, there is generally too little basic research in equilibrium. Therefore, the policy-induced shift towards basic R&D leads to a higher long run growth rate.

4.3 Policy 2: a carbon tax and basic R&D subsidy

Realising that the increase in growth was only possible because the market failure in the R&D sector was partially overcome, one might be tempted to tackle the root of the

problem. If basic R&D is too low, why should we not subsidise basic R&D directly? Therefore, the next policy experiment is the same carbon tax as before, with the difference that we now use the tax revenue to finance a subsidy on basic R&D.

Figure 11 shows that this policy is indeed very successful in terms of growth. There is still the initial output loss, but it is overcome in the year 118, much sooner than under policy 1. Thus, in a net present value analysis, this policy has a much better chance of yielding a positive result than policy 1 did. The reason for this increase in growth is that the reallocation of R&D activity is much larger, and so the arrival rate of new GPTs is sped up considerably. Under policy 2, C05 is invented in the year 117, eight years earlier than under policy 1. Although this technology is a ‘failed’ GPT, it still contributes to output, and due to its earlier invention, output is higher under policy 2.

Although policy 1 is highly beneficial for growth, it is less effective in reducing carbon emissions, for two reasons: First, the change in the relative price of non-carbon-based fuel is smaller. Under policy 1 the price effect of the tax on carbon was reinforced by the subsidy on non-carbon. Now there is only the tax effect, because the subsidy on non-carbon-based fuel is replaced by another subsidy. With a smaller relative price change, the substitution effect is weaker. The second reason is that the increase in growth again raises the demand for both carbon and non-carbon-based fuels. As a result, from the year 174 on, carbon-based fuel consumption is actually higher than in the base run, and environmental quality is lower than in the base run.

4.3 Welfare effects

The welfare effects of any policy experiment depend crucially on the assumptions we make about the utility function, especially with respect to the elasticity of substitution between consumption and environmental quality. Therefore, we distinguish four cases, that of perfect substitutes ($\sigma = \infty$), imperfect but relatively good substitutes ($\sigma = 2$), the Cobb-Douglas case ($\sigma = 1$), and imperfect but relatively ‘bad’ substitutes ($\sigma = 1/2$).

Table 1 summarises the effects of the two policy schemes on welfare. Two observations spring to mind. First, in the long run both policy schemes are welfare-improving. In the case of policy 1, this positive effect comes mainly from the cleaner environment, but also from the increase in final output. In the case of policy 2, the positive welfare effect comes

from higher output, which offsets the increase in pollution. Second, with both policy schemes there may be an initial negative impact on welfare if substitutability is high. This is intuitively plausible, because the initial output loss itself is welfare-decreasing. Over time, the policy schemes increase economic growth, thus raising welfare again.

σ	Policy	1		2	
		Immediate effect	After 150 years	Immediate effect	After 150 years
$\sigma = \infty$		Negative	Positive	Negative	Positive
$\sigma = 2$		Positive	Positive	Negative	Positive
$\sigma = 1$		Positive	Positive	Positive	Positive
$\sigma = 1/2$		Positive	Positive	Positive	Positive

Table 1: Effects of policies 1 and 2 on welfare

4.4 Robustness of results

In order to get an impression of the robustness of the results we have run a number of different simulations. It turns out that policy 2, the subsidy on basic R&D, is always growth increasing, because it provides a remedy to the market failure in the R&D sector. Of course, the effect of the policy also depends on the size of the tax/subsidy. If we set the subsidy on basic R&D extremely high, we will at some time reach the point where basic R&D is actually too high, and growth is slowed down again. There must be an optimal subsidy that exactly equates the social marginal benefits of basic and applied R&D. The exact value of this optimal subsidy is difficult to determine in a model of this complexity, however. The ultimate reason for this is that the arrival of new GPTs and the subsequent expansion of these GPTs is governed by random events, which would allow us to equate these marginal social benefits only in an average sense.

In addition to this, the model generates path-dependencies which in combination with the randomness referred to above, makes it impossible to define generic policy prescriptions even in an average sense. To illustrate the point, in the scenario described above a carbon based GPT happened to be the dominant one for most of the simulation period. Under these

circumstances, policy 1 will induce researchers to move away from applied R&D into basic R&D, because they hope to discover a new GPT that is based on non-carbon-based fuels, and we have shown that re-allocating resources towards basic R&D speeds up growth. If, on the other hand, a non-carbon GPT happens to be dominant, the results are exactly reversed: resources shift from basic R&D towards applied R&D on the non-carbon GPT, and consequently growth slows down. Therefore, the growth effect of policy 1 depends crucially on the existing technology structure.

To check our intuition behind the effects of policy 2, in which tax revenues are used to subsidise basic R&D, we have run simulations with still another policy scheme in which applied R&D is subsidised. This policy 3 is never the optimal one, because it will yield the same substitution effect as policy 2 (and a smaller one than policy 1), but it exacerbates the misallocation between basic and applied R&D, and its effect on growth will generally be negative.

The choice between policy 1 and 2 depends on the preferences that are assumed. In the simulations shown above we have assumed that a representative consumer values both consumption and environmental quality. We have seen that if these two arguments of the utility function are regarded as good substitutes, policy 2 is likely to be optimal because it speeds up growth, albeit at the cost of higher pollution. If the two inputs are poor substitutes, policy 1 is likely to be preferable, because it is more effective at bringing down pollution.

The scenario we have chosen for the simulation analyses is one where a carbon-based GPT happens to be dominant. Therefore, the economy is highly polluted, and environmental quality is the ‘scarce’ good. If, however, the economy was in a situation where either output is low or ‘clean’ technologies dominate, the policies may have fundamentally different effects. It is important to note that the optimal policy depends on the existing technology structure. Therefore, there may not be a universally optimal policy. It might be optimal, for example, to employ policy 2 at low income levels to speed up growth, and then, at higher income levels, to move towards policy 1 in order to curb carbon-based fuel consumption. More simulation experiments may yield further insights into these matters.

Finally, it should be mentioned that the effects of policies in this model are usually not very long lasting. With this statement we mean that if a policy is removed – even after a very long time period – the economy tends to jump back towards a development path very close to the one it would have followed without the policy. This phenomenon is due to the ‘putty’

capital we have assumed, and that is a standard feature of many endogenous growth models.¹⁶ If a policy measure changes the relative price of any fuel, the capital stock is completely free to adjust instantaneously. In reality, capital is of a ‘clay ex-post’ nature. That is, if the price of one type of fuel rises, it is not possible to immediately retrofit all engines to use the other type of fuel. A future version of the model will incorporate some of these ‘clay’ capital features, and we expect the effects of policies to be more long-lasting than in the current model.

5. Concluding remarks

The basic GPT model has shown that a large number of GPTs can coexist in a model that is still not excessively complex. With the structure that we have assumed, it is possible to generate GPTs that arise from a highly successful core technology that is extended with a number of peripheral technologies. The model creates long run growth with a Schumpeterian flavour, because the introduction of a new core technology (a “radical innovation” in Schumpeterian terms) is followed by a quick succession of peripheral technologies (“incremental innovations”) that give rise to alternating phases of fast and slow growth. There is nothing mechanistic behind these “long waves”, although the demise of a GPT is certain in the long run because the possibilities of extension are limited, and new GPTs are sure to arise. These features give rise to a model economy capturing many of the stylised facts that we observe in actual economies over the long run.

Assuming that different GPTs are based on different energy sources, it is clear that this succession of GPTs will also generate long waves in the consumption of different energy sources. We have constructed a model with two energy sources, where each GPT is based on either carbon-based fuels or non-carbon-based fuels. Policies that discourage the use of carbon-based fuels can have a large impact on the allocation of R&D resources in such a framework.

We have examined the impact of different policy scenarios. A common element in all policies was a tax on the consumption of carbon-based fuels, but the tax recycling was different in each scenario. The results of each policy scenario differed considerably. Thus, when we discuss the usefulness of a carbon tax, it is very important to know how the tax is

¹⁶ However, we are now also working on an extension of the model that includes a simplified version of putty-clay capital. See van Zon (2005).

going to be recycled. In addition, we have seen that the effect of such a policy also depends on the existing technology structure. If the dominant GPT is carbon-based, a carbon tax will speed up the search for new GPTs, thereby stimulating growth. But if a non-carbon-based GPT dominates, the carbon tax might result in excessive applied R&D on that GPT, thus delaying the arrival of new GPTs and reducing growth.

We have shown that if the competitive allocation of R&D resources is inefficient, policy may increase the growth rate. In such a case, the carbon tax alone is a second-best solution and should be combined with a policy that addresses the inefficiency in R&D allocation. A similar argument has been made by Smulders and de Nooij (2002), who examined the effects of an energy conservation policy in a model with two directions of technical progress. In our model, however, the inefficiency results from a non-optimal allocation between basic and applied R&D. Therefore, if growth was the primary policy objective, it would be easiest to tax applied R&D and subsidise basic R&D, without creating any distortions on the energy market. If one wishes to promote a carbon tax on the grounds of R&D efficiency, one would have to show that R&D on carbon technologies is inefficiently high. In fact, our model provides potential support to such an argument.

Concerning the coming energy transition, it is clear that the hydrogen system is not even yet in its infancy. It is going to be a core technology, forming the basis for a potentially wide range of peripheral innovations. A lot of basic research will have to be done before the system can actually be implemented. The traditional, oil-based energy system, by contrast, is a very mature system. No more basic research is required, only a few more peripheral applications will be feasible. Since basic research generates positive externalities it is probably underprovided by markets, and a carbon tax combined with a subsidy on basic R&D would be a very effective tool to increase not only the speed of transition, but also the growth rate of the economy as a whole. A subsidy on non-carbon-based fuels would not be the most effective solution, because it would end up supporting current producers of non-carbon-based fuels, and they are not necessarily the ones who develop new non-carbon technologies. The most effective way to speed up the transition towards a non-carbon energy system will be to promote basic R&D on non-carbon energy technologies.

The current model has been specified in a fairly minimalistic way, since it was only meant to serve the purpose of illustrating the principles involved. But it could be extended along several lines. One interesting extension would be to include labour mobility between the final output sector and the research sector. This would enable us to assess the impacts of policy measures on wage differentials or unemployment. It will also make our results more

comparable to other GPT models, because in such models it is the movement of labour between the final output and the research sector that generates output booms and busts.

Furthermore, it has been argued that rising oil prices will contribute to the demise of the oil energy system and the rise of hydrogen economy. We may also incorporate this effect into the model. For the moment being, we have assumed constant real prices of capital and energy. We could allow energy prices to change exogenously over time, or include supply functions for carbon and non-carbon-based fuels.

Another promising route of further investigation would be to look at the optimal timing of policies. We have argued that a growth-promoting policy might be optimal at low income levels, whereas a pollution-curbing policy might be optimal at higher income levels. These issues depend on preferences, specifically the elasticity of substitution between consumption and environmental quality. Especially with regard to the timing of policies, it would be interesting to incorporate clay capital and different GPT vintages into the model. With clay capital the ex-post substitution possibilities between vintages are zero, and the timely execution of policies becomes much more crucial than in the current model. In this context, and considering the randomness and path-dependencies involved, Monte Carlo experiments may be helpful in finding robust policies that get the policy job done in almost all thinkable circumstances.

References

- Aghion, P., and P., Howitt, (1998b), "On the macroeconomic effects of major technological change", *Annales d'economie et de statistique*, 49/50, 50-75.
- Aghion, P., and P., Howitt, (1992), "A Model of Growth through Creative Destruction", *Econometrica*, 60, 323-351.
- Aghion, P., and P., Howitt, (1998a), *Endogenous Growth Theory*, MIT Press, Cambridge MA.
- Arrow, K.J., (1962), "The Economic Implications of Learning by Doing", *Review of Economic Studies*, 29(3), 155-173.
- Barro, R.J., Sala-I-Martin, X., (1995), *Economic Growth*, McGraw-Hill, New York.
- Bresnahan, T., and Trajtenberg, M., (1995), "General Purpose Technologies", *Journal of Econometrics*, 65, 83-108.
- Carlaw, K.I., and Lipsey, R.G., (2001), "Externalities versus Technological Complementarities: a model of GPT-driven, sustained growth", paper presented at the Conference in Honour of the 20th Anniversary of Nelson and Winter's Book *An Evolutionary Theory of Economic Change*, <http://www.sfu.ca/~rlipsey>
- Cheng, L. K., and E. Dinopoulos, (1992), "Schumpeterian Growth and Stochastic Economic Fluctuations", *University of Florida Mimeo*.
- David, P. (1990), "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox", *American Economic Review Papers and Proceedings*, 80(2), 355-361.
- Dosi, G., (1988), "Sources, Procedures, and Microeconomic Effects of Innovation", *Journal of Economic Literature*, 26 (September), 1120-1171.
- Ethier, W.J., (1982), "National and International Returns to Scale in the Modern Theory of International Trade", *American Economic Review*, 72(3) (June), 389-405.
- Francois, P. and H. Lloyd-Ellis, (2003), 'Animal Spirits Through Creative Destruction', *American Economic Review*, 93(3), 530-550.
- Freeman, C., and C. Perez, (1988), "Structural crises of adjustment, business cycles and investment behaviour", in Dosi et al. (eds.), *Technical Change and Economic Theory*, Pinter Publishers.
- Gordon, R.J., (2000), "Does the 'New Economy' Measure up to the Great Inventions of the Past?", NBER WP No.7833.

- Grossman, G.M., and E. Helpman (1991), *Innovation and Growth in the Global Economy*, Cambridge, MIT Press.
- Helpman, E., (1998), *General Purpose Technologies and Economic Growth*, MIT Press.
- Helpman, E., and M., Trajtenberg (1994), “A time to Sow a Time to Reap: Growth Based on General Purpose Technologies”, *Center for Economic Policy Research*, Discussion Paper No.1080 (December), 1-32.
- Jovanovic, B., and R., Rob, (1990), “Long Waves and Short Waves: Growth through Intensive and Extensive Search”, *Econometrica*, 58(6), 1391-1409.
- Lipsey, R., and C., Bekar, (1995), “A Structuralist View of Technical Change and Economic Growth”, In T. Courchene, ed., *Technology, Information and Public Policy*, Kingston, Ont: John Deutsch Institute for Study of Economic Policy.
- Matsuyama, K., (1999), ‘Growing through Cycles’, *Econometrica* 67, 335-347.
- Mensch, G., (1979), *Stalemate in Technology: Innovations Overcome Depression*, Cambridge MA, Ballinger Publishing Company.
- Mokyr, J., (1990), *The Lever of Riches*, Oxford University Press.
- Mukerji V. (1963), “Generalized SMAC Function with Constant Ratios of Elasticity of Substitution”, *Review of Economic Studies*, 30, 233-236.
- Olsson, O. (2001), “Why Does Technology Advance in Cycles?”, Available at: <http://swopec.hhs.se/gunwpe/abs/gunwpe0038.htm>.
- Rifkin, J. (2002), *The Hydrogen Economy*, Polity, Cambridge UK
- Romer, P. (1990), “Endogenous Technological Change”, *Journal of Political Economy*, 98, 71-102.
- Schumpeter, J., (1934), *The Theory of Economic Development*, Cambridge MA, Harvard University Press.
- Schumpeter, J., (1939), *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process*, 2 Vols, New York, McGraw-Hill.
- Segerstrom, P. S., Anant, T., and Donopoulos, E., (1990), “A Schumpeterian Model of the Product Life Cycle”, *American Economic Review*, 80, 1077-1092.
- Smulders, S., de Nooij, M., 2002, The impact of energy conservation on technology and economic growth, *Resource and Energy Economics* 25, 59-79.
- Solow, R., (1997), *Learning from “Learning by Doing” Lesson for Economic Growth*, Stanford University Press, Stanford California.
- Yetkiner, I. H., (2003), *Shocks and Growth: Four Essays*, PhD-Thesis, Groningen University, Groningen.

- van Zon, A., (2001), “A simple endogenous growth model with asymmetric employment opportunities by skill” *MERIT Research Memorandum*, 2001-029.
- van Zon, A., Fortune, E., and Kronenberg, T., (2003) “How to sow and reap as you go: a Simple Model of Cyclical Endogenous Growth”, *MERIT Research Memorandum*, 2003-026
- van Zon, A., and I. H., Yetkiner, (2003), “An Endogenous Growth Model à la Romer with Embodied Energy-Saving Technological Change”, *Resource and Energy Economics* 25(1), 81-103.
- van Zon, A., (2005) “Vintage modelling for dummies using the putty-practically-clay approach”, *MERIT Research Memorandum*, 2005-005.