

Sustained endogenous growth driven by structured and evolving general purpose technologies

Kenneth I. Carlaw · Richard G. Lipsey

Published online: 5 January 2011
© Springer-Verlag 2011

Abstract We address two interrelated issues: structured technology and non-stationary equilibrium growth. We do this by modelling multiple, co-existing, non-identical general purpose technologies (GPTs). Three sectors producing pure and applied research and consumption goods, employ different, evolving, technologies. Agents within each sector operate under conditions of Knightian uncertainty and path dependence, employing technologies that differ in specific parameter values. This behaviour produces a non-stationary (non-ergodic) growth process. Important characteristics of structured technology, previously only described historically, are successfully modelled, including co-existing GPTs some of which compete with each other while others complement each other in varying degrees. Because changes in technology are partial causes of, but not contemporaneous with, GDP changes, their separate evolutions can be studied.

Funding for this research was provided by the Royal Society of New Zealand's Marsden Fund Grant 04 UOC 020. Earlier versions were presented at the 2006 meetings of the Joseph Schumpeter Society and published as part of the Simon Fraser University Department of Economics working paper series, 2007. We are grateful for comments and criticisms from Clifford Bekar, Les Oxley and Alvero Pereira.

K. I. Carlaw (✉)

Department of Economics, University of British Columbia, Kelowna, BC, Canada
e-mail: kenneth.carlaw@ubc.ca

R. G. Lipsey

Department of Economics, Simon Fraser University, Vancouver, BC, Canada
e-mail: rlipsey@sfu.ca

Keywords Structured technology · General purpose technologies · Non-stationary growth · Technological evolution · Non-ergodic equilibria

JEL Classification O41 · O47 · O29 · N70

It is generally agreed that technological change is a necessary condition for sustained economic growth since growth based on capital accumulation with constant technology would sooner or later come to a halt. Empirical work on technology, as found in authors such as Rosenberg (1982, 1994), David (1990, 1997), David and Wright (1999, 2003) and Freeman and Louça (2001), shows technology to have a complex structure. For example, the technologies in use typically stand in a hierarchy to one another and require context specific complementary technologies in their application, causing them to differ among sectors, industries and even among firms in any one industry. Furthermore, these technologies typically evolve along path-dependent (i.e., non-ergodic) trajectories that differ from each other.¹ Empirical work also shows that growth in market economies does not proceed at an even, stationary rate. For example, Freeman and Louça (2001) argue persuasively that growth has proceeded over the last few centuries in a series of ‘speedups’ and ‘slowdowns.’ David (1990) investigated one of these slowdowns associated with the introduction of electricity.

In contrast to these empirical findings, most of the recently developed endogenous growth models follow Romer (1986) and Lucas (1988) in modelling technology as a scalar, pre-multiplier in an aggregate production function or as a single variable within that function. As a result, technical change is shown by a single growth rate and economic growth proceeds along a balanced (stationary) growth path that is contemporaneous with technological change.

By dividing technology into a general purpose technology (GPT) and supporting sub-technologies, the initial literature on GPTs went beyond the modelling of technology as a single variable by giving technology a rudimentary structure. By developing models in which the introduction of a new GPT produced a slowdown in the rate of growth of productivity and national income, it addressed at least one aspect of uneven growth.

The seminal article on GPTs was Bresnahan and Trajtenberg (1995), and the concept was further developed in the several models in Helpman (1998).² Although the models in that book were important first steps in the

¹Davidson (2002) provides a statement of the difference between the ergodic axiom that is central to the stationary equilibrium concept of New Classical and Neoclassical economics and the non-ergodic equilibrium concepts employed in Keynesian and evolutionary economics. We employ the latter here. Lipsey et al. (2005: Chapters 16 and 17) make the distinction critical for the development of what they call technology enhancement policy. Lipsey (2007) uses the same distinction extensively in his reconsideration of second best theory.

²The relevant models in Helpman (1998) are Helpman and Trajtenberg (1998a, b) and Aghion and Howitt (1998).

development of a full theory of GPTs, they all have a number of characteristics that are inconsistent with evidence accumulated by the historical research. One of the most important is that they all contain a single GPT whose behaviour dominates the behaviour of the whole economy. In spite of such limitations, however, the concept of a GPT has been applied to many issues in economic history, industrial organization and economic policy.³

The divergence between the existing one-GPT models and the many-GPT real world has caused at least three errors in GPT research. First, the attempt to test for the existence of a specific GPT by examining the economy's aggregate behaviour assumes incorrectly that every single GPT leaves detectable traces on the aggregate behaviour of the economy. Second, there is a widespread but erroneous belief that a basic hypothesis of GPT theory is that new GPTs invariably cause slowdowns in productivity and income growth. Third, the attempt to explain those productivity slowdowns that do occur by using models that contain only one class of GPT is inappropriate because different GPTs are observed to have different effects on productivity.

The considerations in the above paragraph strongly suggest the need for more empirically relevant models of structured and evolving technology in general and GPTs in particular, which is what we attempt in this paper. We first list those key characteristics of structured technologies (including GPTs) that have been established in the technology literature and that we are able to model here. We then compare what we do with the handling of those of the characteristics on our list to what has been modelled elsewhere.

1. Technologies are developed through endogenous Research and Development (R&D) activity.
2. The efficiency with which a technology delivers its services typically increases over time.
3. The use of any new technology including a new GPT spreads slowly through the economy and many decades are typically required for a GPTs full diffusion to many different sectors and uses.
4. Several non-identical GPTs exist at any one time.
5. General purpose technologies can be grouped into what in this paper we call “classes” of technology. Lipsey et al. (2005, hereafter LCB) distinguish five such classes for GPTs, materials, information and communication technologies (ICTs), power sources, transportation equipment, and organizational forms (e.g., the factory system), with at least one

³For examples in economic history see Crafts (2003), Greasley and Oxley (1997), Jovanovic and Rousseau (2005), Mokyr (2006), and Ruttan (2005); in applied microeconomics see Hornstein et al. (2005), Lindmark (2005), and Sapio and Thoma (2006); in assessments of policy see Forman et al. (2003), Safarian and Dobson (2002), van Zon and Kronenberg (2007), and Wehrli and Saxby (2006); in practical policy strategies that directly employ the concept, see New Zealand's development of two entire policy frameworks in *The Growth and Innovation Framework* (GIF) and the *Digital Strategy*. (see, <http://www.digitalstrategy.govt.nz/>) as well as the Australian Research Council (2006).

GPT in each class being used at any one time—but other groupings of technology may be convenient for other purposes.

6. Over time, many different GPTs within each class are invented. These often compete with each other and, as a result, there can be several such GPTs in simultaneous use. For example, in the early 1900s some textile factories were shifting to electricity as a power source, while most were steam powered, and not a few were still using water wheels—three GPTs all within the power delivery class of technologies.
7. In contrast, GPTs that are in different classes often complement each other, as when electricity enabled electronic computers and lasers.
8. The invention, innovation and diffusion of any major new technology, including GPTs, involves many uncertainties. In particular, uncertainty pervades the following: (i) how much potentially useful pure knowledge will be discovered by any given amount of research activity; (ii) the timing of the discovery of a new GPT; (iii) just how productive a newly innovated GPT will be over its lifetime; (iv) how well the new GPT will interact with GPTs of other classes that are also in use; (v) how long a new GPT will continue to evolve in usefulness; (vi) when it will begin to be replaced by a new superior GPT of the same class (vii) how long that displacement will take and (viii) if the displacement will be more or less complete (as were mechanical calculators by electronic computers) or if the older technology will remain entrenched in particular niches (as steam turbines remain an important source of power for generating electricity).⁴
9. Because of pervasive uncertainty, it is not feasible for agents to maximise over the lifetimes of new GPTs, let alone over the infinite horizon.
10. Individual GPTs can have very different life patterns. Some, such as the three masted sailing ship and the steam engine, disappear more or less completely from commercial use. Others, such as stone, bronze and the waterwheel, retreat into specialised niches ceasing to be GPTs. Others, such as writing and electricity, become fully assimilated into the fabric of the economy, continuing to be important economically but no longer having the often disruptive impacts on the economic structure that they had during their initial diffusion.
11. The introduction of a new GPT is sometimes, but not always, associated with a slowdown in the rates of growth of productivity and national income.

These 11 characteristics are discussed below and referred to by number when they are first modelled.

Characteristic 1 Despite the established endogeneity of R&D, the GPTs in all of the models in Helpman (1998) arrive exogenously. In van Zon et al. (2003,

⁴All of these sources of uncertainty combined with the endogenous behaviour of agents and the path dependant evolutions of each major new technology result in a non-stationary (non-ergodic) time series process.

hereafter ZFK), resources are arbitrarily divided once and for all between the production of the final good and research. The allocation decision is how many research resources to allocate to searching for a new GPT and how many to allocate to producing ancillary peripherals that augment the existing GPT. Carlaw and Lipsey (2006, hereafter C&L) presented a three sector model in which successive GPTs were developed endogenously in a pure research sector and then used by an applied R&D sector to produce technologies that were used in the consumption goods sector. Our present model uses the same three sectors.

Characteristic 2 The increase over time in the efficiency with which a GPT delivers its services has been modelled by Helpman and Trajtenberg (1998a hereafter H&Ta), Aghion and Howitt (1998 hereafter A&H) and ZFK as a consequence of the need to produce a number of complementary components that increase the GPT's potential efficiency. In H&Ta, after the GPT arrives, resources are diverted to produce these components and when a sufficient number have been produced, the GPT's efficiency passes a critical level at which it pays to discard the incumbent GPT and replace it with the new arrival. There is then a sudden discrete jump in productivity as the GPT comes into use in the economy's final goods sector. The GPT's efficiency may stay constant once it is adopted or it may rise linearly as further complementary components are developed. In A&H efficiency and diffusion are combined so we deal with them under Characteristic 3. In ZFK, each time a new GPT is invented, it creates its own CES production function with the GPT's efficiency as the zero element and all produced ancillary components as further elements. Thus all elements are substitutes for each other and their overall efficiency rises at a steadily diminishing rate as new components are added. LCB come closer to the historical evidence by separating the efficiency of the GPT from the efficiency of its applications. They also allow their endogenously developed GPT to improve its efficiency logistically as it is used over time.⁵ Our present model uses these LCB assumptions.

Characteristic 3 The diffusion of a new GPT over time, Helpman and Trajtenberg (1998b hereafter H&Tb) build on their earlier H&Ta model in which a new GPT arrives and complementary components begin to be developed. In this later model, there are many sectors, each of which must use the same GPT but which differs in its value to each of them. The sector that most values the new GPT starts the production of the complementary components. Then, as the potential efficiency of the new GPT rises, other sectors begin producing complementary components. When a sufficient number of components

⁵The two main empirically established sources of improvement after a major technology is introduced are 'learning by doing' and 'learning by using.' The former refers to improvements into the production technique of the GPT itself and the later to improvements in its design and application as experience with its use accumulates and is fed back to manufacturers and designers. On the former, see Arrow (1962) and on the later Rosenberg (1982: Chapter 6).

have been produced, the GPT is simultaneously adopted in all sectors. Thus the new GPT diffuses over time only in the sense that over time more and more sectors switch resources to creating its complementary components, but it comes into use all at once everywhere. In A&H, a new GPT arrives and a random draw determines when it becomes available for adoption in any of the identical sectors producing the final good. Once this happens, the sector diverts resources from producing the final good to producing a chance of a favourable draw that will produce the one necessary ancillary component to work with the GPT in that sector. When this happens, the sector adopts the new GPT whose efficiency remains constant after it is introduced. It is arbitrarily assumed that before a new GPT arrives, all sectors will have adopted the current GPT, thus avoiding the problem of the co-existence of multiple GPTs. In ZFK, there is no diffusion as each sector develops its own GPT and adopts it as soon as it is discovered. In C&L and LCB there is no diffusion of the GPT over time since the one applied research sector either adopts the new GPT or it does not. Our present model comes closer to the historical evidence in that whenever the GPT is efficient enough for use in any one 'facility' that conducts applied R&D, it is adopted and used in that facility, causing an immediate impact on output and productivity. Later, another applied facility may adopt it and a new impact is felt on productivity, and so on.

Characteristic 4 The only model published so far that contains more than one simultaneously existing GPT is ZFK.⁶ Their GPTs are never retired, although each one becomes less important over time as their market share diminishes because more and more GPTs accumulate. Final output is produced by a Cobb Douglas production function with two elements, labour and a composite index of capital. Since the aggregate capital value that is inserted into that production function is created by aggregating over all GPTs using a CES production function, each new composite GPT has diminishing returns and a decreased market share in terms of the final product (and the growth rate goes asymptotically to zero). While the ZFK model is a significant improvement over all previous models in having multiple co-existing GPTs, it has several counterfactual characteristics that are avoided in the present model. First, all GPTs cooperate with each other in jointly producing the composite final good. Second, each GPT has a smaller impact on output and growth than each previous GPT. Third, in the long run GPTs do not sustain growth, which goes asymptotically to zero. In the present model, each of the many simultaneously co-existing GPTs has its own distinctive characteristics, as outlined in Characteristics 5–7. Each new GPT may have a smaller or larger market share than any of its predecessors and the market impact of any given GPT is not predetermined to shrink as GPTs accumulate. Also growth is sustained indefinitely (but at varying rates) by a succession of new GPTs.

⁶H&Ta and b deal with only one GPT while A&H deal with a succession of them but each new arrival displaces the incumbent in all sectors before yet another GPT arrives.

Characteristics 5–7 While we think our modelling of characteristics 3 and 4 represent improvements over any existing models, there are no treatments of characteristics 5–7 in the literature since there are no models of multiple coexisting GPTs that sometimes complement and sometimes compete with each other. This modelling takes us significantly beyond what currently exists in the literature and significantly closer to the observed characteristics of technologically driven growth.

Characteristic 8 There is no randomness in either of the H&T papers, while there is randomness in A&H on the arrival date of new GPTs and in ZFK with respect to each new GPTs arrival date and its productivity. But both of the latter models employ the ergodic axiom under which agents know the distributions that are producing the randomness. This allows the use of dynamically stationary equilibrium concepts⁷ in which agents necessarily foresee the expected behaviour of a new GPT over its whole lifetime and maximize over the infinite horizon, a technique that is also used in the two H&T papers. All of the models that use this equilibrium concept soon become analytically intractable when more empirically relevant technological structures are introduced.

Characteristic 9 The use of the ergodic axiom and the dynamically stationary equilibrium concept that it permits implies that agents can foresee everything relevant about the GPT over its entire life cycle, which often spreads over more than a century. This unrealistic assumption was circumvented in C&L's dynamically non-stationary (i.e., non-ergodic) equilibrium model of GPT-driven growth and was used in LCB.⁸ In it, agents face a future that is genuinely uncertain because the environment in which they operate is non-ergodic and do the best they can with limited current knowledge. Not only does this approach avoid the totally unreal assumption of perfect long term foresight, it also allows the introduction of more complex and realistic aspects of the evolving structure of technology and the behaviour of agents. We use it in the present paper, which builds on the 'baseline model' presented in LCB's Chapter 14.

Characteristic 10 In ZFK's model, GPTs continue in existence forever and although each has a diminishing effect on aggregate productivity, this is only because more and more GPTs are added to the capital stock whose total amount encounters diminishing returns. In the present model, GPTs have

⁷Davidson argues (e.g., Davidson 1993, 2002, 2009.) that this stationary equilibrium concept is founded on what he calls the ergodic axiom whose introduction into mainstream neoclassical economics he largely attributes to Samuelson. The axiom allows for statistical inference from cross sectional and time series samples using standard empirical techniques and is the rationale for employing stationary equilibria in economic modelling.

⁸The dates in the study paper versions of C&L show that it was written before LCB, although published after.

various life patterns depending on when new competing GPTs of their own class are developed and on how long it takes for a new version to supplant the existing one in its various uses. Typically, any one GPT does eventually retire in the sense that it is no longer used as the current GPT in production. But it never fully disappears because the knowledge associated with past GPTs feeds into applications that are utilized to produce future GPTs. For example, knowledge of metallurgy first learned with bronze and iron is still used and although electricity and writing have passed into the fabric of the economy and are no longer going through a process of diffusion they are still vital to current GPTs. It is also possible that one GPT can exist indefinitely in some niche.

Characteristic 11 This Characteristic concerns the impact that a new GPT has on the behaviour of the economy as a whole, particularly on productivity. David (1990) provided a detailed analysis of the reasons why the introduction of electricity produced a slowdown. Freeman and Louça (2001) provided convincing evidence that changes in their techno-economic paradigms (TEP) were associated with what they call ‘structural crises of adjustment’ that typically cause slowdowns. But their TEP is a much more encompassing concept than GPTs.⁹ None of these authors formally modelled their theories. All of the models in Helpman (1998) produce slowdowns in the growth rate when any new GPT is introduced. However, LCB argued that the extent to which a GPT would cause a slowdown varied with the way in which it interacted with the economy’s existing structure. The more structural change required to give the new GPT effect, the more likely it was to cause a perceptible slowdown. But they only modelled this interaction in their one-GPT model.¹⁰ Although growth does slow as new GPTs are introduced in the ZFK model, this is a steady process caused by the use of a single aggregate production function that has diminishing returns to the composite index of capital. In contrast, the present paper develops a model of multiple, co-existing, non-identical structured technologies (some of which may become GPTs) that is consistent with much of the historical evidence and that fulfils a necessary condition for developing a satisfactory model of occasional slowdowns. Developing such a model will, however, require treatment in a full paper, which is the subject of our current research.

In summary, the present model has some attractive features that distinguish it from most other endogenous growth, and all other GPT, models. It produces long term, non-balanced growth and waves of activity associated with the life cycles of individual GPTs. The model also provides a tool that has many

⁹Although these two concepts are often treated as identical, a TEP includes one or more GPTs and their supporting technologies and economic structure. The distinction between the two is discussed in detail in LCB, pp 372–77.

¹⁰See LCB (2005: 482–88) “A Simple Model of Facilitating Structure and GPTs.”

other potential applications. Although each of these requires a separate paper, some are briefly outlined at the end of this paper. In addition, by permitting technologically driven slowdowns as possible but not necessary, the model allows specific technologies to cause either, or both, or neither of troughs or peaks.

We couch our discussion in terms of GPTs because we think they are an important concept in understanding Western economic growth over the past 10,000 years. However, this is only one possible application of our model whose key characteristics are that final goods are produced by two “upstream” sectors, one that produces generally applicable technologies and one that uses these to produce applications that can be used directly in the production of final goods, and that each of the three sectors contains many “firms” with production functions that may differ from each other. The pure research sector could be producing knowledge that leads to new GPTs, or it could be producing knowledge that is less universal as long as it provides fodder for the applied R&D sector to produce commercially viable uses. For example, there has been much research over the past decades to develop commercially viable hydrogen fuel cells. If these efforts ever succeed, further R&D will be required to develop their many possible applications, even if they never approach the general applicability of a typical GPT. In another application, the pure research sector could be producing knowledge that permits fundamental technologies developed abroad to be used to produce technologies useful in domestic production, allowing the model to be applied where countries are adopting and adapting technologies created elsewhere.¹¹ For example, countries, such as China, India, South Korea, and Singapore, have devoted research to successfully adapting the GPTs that drive their modern ICT, automobile and other manufacturing industries, although these GPTs were invented elsewhere. Other countries such as Kenya, Zimbabwe, and Uruguay, which lack the pure knowledge base required to adapt these GPTs, have been unsuccessful in using them to create significant industrial applications. What is specific to our model are the characteristics of the three sectors (e.g., technology is structured); what is general are the many applications of each of these sectors to real situations, as illustrated in the above examples.

1 The sequential GPT model

Figure 1, summarises C&L’s model. One sector produces pure research that occasionally discovers a new GPT; one sector produces applied research that develops applications for the GPT that are embodied in either physical or human capital; one sector produces a consumption good using the capital

¹¹Cohen and Levinthal (1990) call this knowledge producing activity the building of “absorptive capacity”.

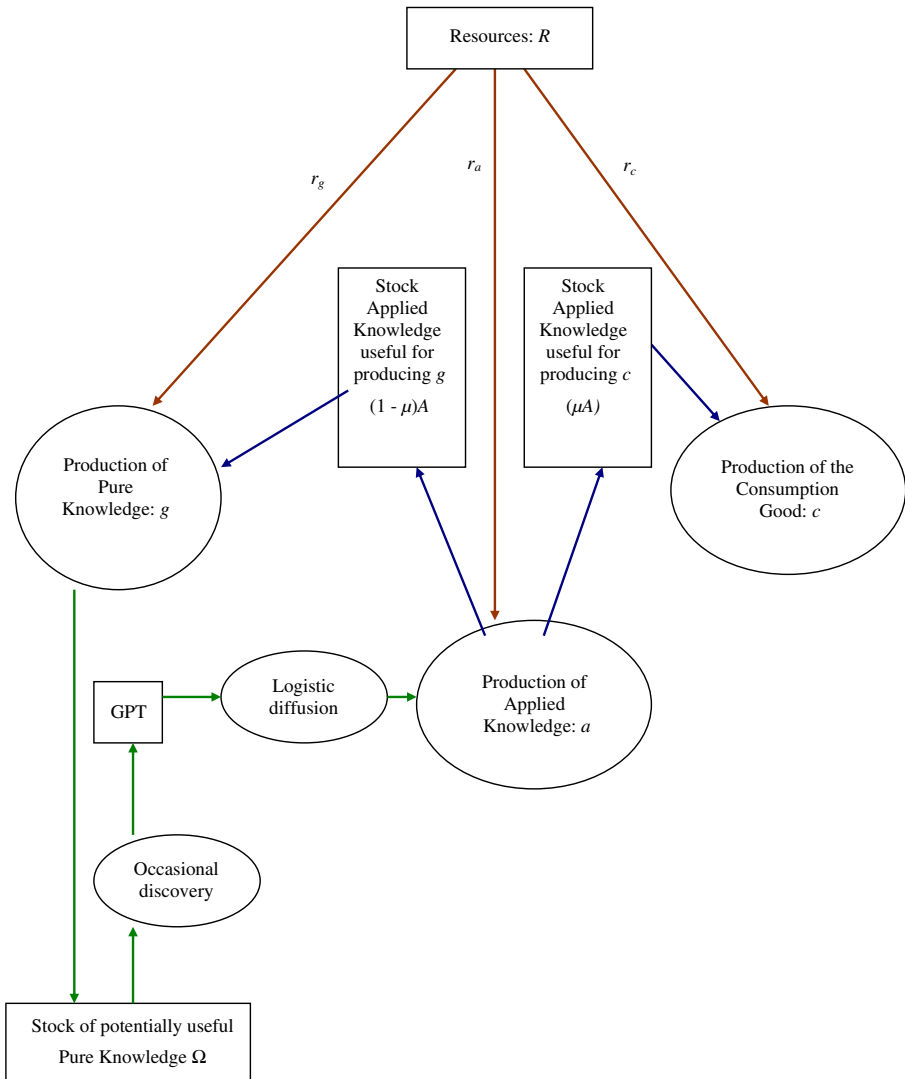


Fig. 1 The three sector model

produced by the applied research in its production function. Each sector is treated as a single unit with its own distinct aggregate production function, while a technology structure is imposed through the inter-sector relations among these three different production functions.

The economy has a fixed aggregate stock of a composite resource, R , which can be thought of as ‘land’ and ‘labour.’ The three aggregate production functions display diminishing marginal returns to this resource. The pure knowledge sector produces a flow of pure knowledge, g , which accumulates in a stock of potentially useful knowledge, Ω . Every once in a while a new GPT is

invented. The existing stock of potentially useful pure knowledge is embodied in it and then its efficiency slowly evolves according to a logistic function to become increasingly useful in applied research. The applied R&D sector produces practical knowledge which embodies the pure knowledge of the GPT in its applications that can take the form of physical or human capital or both. This embodied applied knowledge is useful in both the consumption and the pure research sectors, the latter being a feedback that is well established in the technology literature.¹² The stock of applied knowledge is *divided* between a proportion, μ , (where $\mu \in (0,1)$) that is useful in the consumption industry and, $1 - \mu$, that is useful in further pure research. For concreteness, the proportion μ may be regarded as physical capital that incorporates the applied knowledge, while the proportion $1 - \mu$ can be regarded as the human capital of those who work in the pure research sector.¹³ The proportion that is useful to the consumption industry enters its aggregate production function as a productivity coefficient with constant returns to scale. Thus the consumption sector's aggregate production function is of the form $C = \mu A (f_c)^\alpha$ where C is aggregate consumption output, μA is the amount of applied knowledge embodied in capital that is useful in the consumption sector, f_c is the amount of the composite resource that is allocated to the consumption sector, and $\alpha \in (0, 1)$.

Adopting the assumption that the capital that embodies the applied knowledge produced by the R&D sector is divided between what is useful in the consumption sector and in the pure knowledge sector ensures that positive externalities to accumulating applied knowledge are not introduced into the model. Among other things, this allows the production of a model of sustained endogenous growth without some of the characteristics that are often needed to sustain growth in other models of endogenous growth. Assuming that the whole stock of applied knowledge was embodied in physical capital which was useful in the consumption sector while some fraction of the knowledge also spilled over as an externality to increase the efficiency of pure research, would introduce increasing returns to knowledge. The no-externalities assumption also creates a model suited to studying the relation between technological change and standard measures of changes in total factor productivity (TFP), calculations of which are usually based on the assumption of constant returns to the accumulating factors.

The model produces endogenously determined sustained growth at a rate that varies over time.

¹²See, for example, Rosenberg (1982: Chapter 7).

¹³More sophisticated representations of the model would include explicit and formal embodiment processes. But at the current level of abstraction and aggregation these are unnecessary complications for our modelling exercise. When computing the total capital stock for this model as is done, for example, in Carlaw and Lipsey (2010) we value the stock of pure knowledge at its resource cost, as is usually done for R&D in official statistics.

2 The simultaneous GPT model outlined

We now introduce our alterations to the above model that are designed to incorporate characteristics 1–7 and 10 of the numbered characteristics listed in the introduction. The various sources of uncertainty outlined in Characteristic 8 are introduced in the course of the formal specification of our model in Section 3, while Characteristics 9 and 11 concern the behaviour of the model, not its specification characteristics. To allow for different classes of GPTs (Characteristic 5), we introduce multiple activities in each of the three sectors. The pure research sector has X laboratory complexes, which we call ‘labs’ for short, each producing a distinct class of pure technological knowledge and making use of all types of applied knowledge. Each lab occasionally invents a new GPT in its particular class (Characteristic 1). We number the GPTs in a particular class according to the order in which they were first adopted. The applied R&D sector has Y ‘facilities,’ each producing a distinct type of useful knowledge, and employing one GPT from each class. The consumption goods sector has I industries, each producing a different consumption good, and making use of the stocks of embodied applied knowledge. The efficiency of each GPT is denoted $(G_{n_x})_t$, where G is the stock of pure knowledge in the GPT, n is the position of the GPT in the sequence of GPTs developed in technology class $x \in [1, X]$ and t is a time index. Each GPT evolves according to a logistic function that is given in Eq. 8 below, until its full potential has been realized (Characteristic 2). This increasing productivity affects both the applied R&D facilities that are currently using the GPT and the GPT’s potential were it to be adopted by other facilities at some future time. In order to allow a new GPT to spread through the economy over time (Characteristic 3), the productivity of each GPT is allowed to differ in each of the applied R&D facilities. This is done by multiplying each GPT’s productivity by a parameter ν that is specific both to that GPT and the R&D facility in which it is operating.

Each applied R&D facility is initially seeded with one GPT of each class (Characteristic 4). When a new GPT is invented—we call it a ‘challenger’—each applied R&D facility must decide whether to adopt it or to stay with the GPT of that class that it is presently using—we call this the ‘incumbent.’ Because of the uncertainties listed in (Characteristic 8) above, the full evolution of the incumbent and challenging GPTs cannot be predicted in advance and then compared. So some more restricted choice criteria are needed (Characteristic 9). Of the many possibilities, we use a simple comparison of the challenger’s current level of productivity with the current productivity of the incumbent.¹⁴ If a new challenging GPT has a lower current productivity than the incumbent in all applied R&D facilities, it is sent back to the pure research sector for further development and it is reconsidered for adoption in every

¹⁴LCB (451–54) studied the effects of five different adoption criteria, any of which could be used here.

subsequent period. If it has a higher productivity in one or more facilities, it is adopted there. Thus, there may be several GPTs in any one class of technology in use at any one time (Characteristic 6). Each subsequent period, the facilities that have not adopted this new GPT compare its evolving productivity with the evolving productivity of their incumbent and switch when the former exceeds that of the latter. Over time, the use of a new GPT of any one class will spread through the economy, as one applied R&D facility after another adopts it and discards its incumbent (Characteristic 3). It is possible, however, that another GPT of that class will be invented before the first challenger has been taken up by all applied R&D facilities. If so, more than one comparison of the original incumbents to the possible challengers might be taking place for a period of time. Indeed, this could go on indefinitely as an early version of a GPT in some class persists in a niche where it remains superior on comparisons with all later GPTs in that class (Characteristic 10).

2.1 Some terminology

As noted earlier the variable x indicates a specific *class* of technology to which a GPT belongs. The index n_x identifies a GPT in the sequence of GPTs, $1 \dots n$, in technology class x that has been invented and then adopted by at least one applied R&D facility. At any one time, the latest GPT adopted in that class is denoted by n_x and the previously adopted version by $(n - 1)_x$. We refer to the productivity of the most recently invented GPT from class x at time t as $(G_{n_x})_t$ and the previously invented GPT in that class as $(G_{(n-1)_x})_t$. The variable t_{n_x} refers to the invention date of the latest GPT of the class x , while $t_{(n-1)_x}$ refers to the invention date of the previous GPT from class x , and $(t - 1)_{n_x}$ refers to the period just prior to the invention of GPT n_x . (In all cases, “invention date” means invented and adopted by at least one applied R&D facility.)

2.2 Relations among GPTs

The productivity coefficient in each applied R&D facility is the geometric mean of the productivities of the GPTs that it uses (one GPT from each class of technology), each pre-multiplied by an associated parameter value, v . (See Eq. 2 below.) The v s model several key characteristics. First, they allow any new GPT to have different productivities in each activity and so not to be adopted simultaneously everywhere but instead to compete for adoption with all of the exiting incumbents of its own class of technology (Characteristic 3). Second, they allow GPTs of different classes that are working together in any one facility to complement each other so that their joint productivity exceeds the sum of their separate productivities, and sometimes to work less well together so that their joint productivity is less than the sum of their separate productivities (Characteristic 7). They can evolve as a random walk as we do in our present theoretical treatment, or they can be set so as to model any desired degree of competitiveness among different GPTs of any one class,

and different amounts of complementarity among GPTs of different classes operating together in any one R&D facility.

The ν s can be set out in a $Y \times X$ matrix where a row indicates the class of technology and a column the research facility. We call this the ‘operative ν matrix’. Initially all the ν s in this matrix are set at unity. Thus the initial GPT in each class has the same level of productivity in each of the applied R&D facilities and the GPTs work together in each of the facilities according to their unmodified productivities. When a new GPT of a particular class of GPT is invented, it brings with it its own matrix of potential new ν s, which we call its ‘potential ν matrix’. In our model these new ν s differ randomly over a small range centred on unity from the ν s in the operative matrix and they are not known to each agent until the new GPT is adopted by that agent. Therefore, when each R&D facility decides whether or not to adopt the challenger, it must form expectations over the potential ν s required to make the calculation of what its output would be if the challenger were used. We denote these expected values by $\bar{\nu}$. For simplicity we assume that the $\bar{\nu}$ for the challenging GPT being considered is the ν from its potential matrix while the $\bar{\nu}$ s that modify the other two GPTs in use by that facility are the current ν s from the operative matrix. In other words, the facility correctly perceives the direct productivity of the challenger but assumes, possibly incorrectly, that it will cooperate with the two other GPTs in use with the same degree of efficiency as does the incumbent.

Because the ν s in the challenging GPTs’ potential matrix vary across the relevant row, different facilities will evaluate its productivity differently and so some may adopt it while others do not. When a particular applied R&D facility adopts the challenger, the challenger’s potential ν s replace the existing ones in that facility’s column vector in the operative ν matrix. The resulting change in the ν s associated with GPTs in classes other than the newly adopted GPT indicates whether the new GPT cooperates better or worse with the existing GPTs of other classes than did the replaced incumbent (characteristic 7). As time goes by, the operative matrix of ν ’s changes and at any one time each column in the matrix is derived from the potential ν ’s of the latest GPT to be adopted by that R&D facility.

When calculating the output of each applied R&D facility in Eq. 2, all that is needed is to pre-multiply the productivity coefficient for each GPT being used by the corresponding ν s in that facility’s operative column vector. However, each period each facility must compare its incumbent GPT of each class with all the other GPTs of that class that are already in use by other facilities or have just been invented. To do this, we need a full identification of each ν , which requires four indexes, $\nu_{y,z}^{n_x}$. The superscript n_x tells us that the ν is modifying a GPT of class x , number n . The subscript tells us that the GPT is being used by research facility $y \in [1, Y]$ and replaced the previous ν when a new GPT of class $z \in [1, X]$ was adopted. Thus, z indicates the technology class of the GPT that was last adopted by that facility and is the source of all the ν s in that column, while n_x indicates the specific GPT of class- x to which the ν is being applied. Note that the x in the superscript and the y in the subscript do not change over time, the former indicating the class of technology from which the GPT came

Table 1 Illustrative distribution of ν s for a model with three GPTS and three applied R&D facilities

| | | $y=$ | | |
|------|---|-------------------|-------------------|-------------------|
| | | 1 | 2 | 3 |
| $x=$ | 1 | $\nu_{1,1}^{3_1}$ | $\nu_{2,1}^{4_1}$ | $\nu_{3,3}^{2_1}$ |
| | 2 | $\nu_{1,1}^{2_2}$ | $\nu_{2,1}^{2_2}$ | $\nu_{3,3}^{1_2}$ |
| | 3 | $\nu_{1,1}^{1_3}$ | $\nu_{2,1}^{1_3}$ | $\nu_{3,3}^{2_3}$ |

and the latter the R&D facility that is using it. In contrast, n and z do change over time, the n indicating the specific GPT of technology class- x being used by the facility in question and the z indicating the class of the last GPT adopted by the applied R&D facility where the n th GPT of class- x is being used (and hence the source of all the ν s in that facility’s column vector).

An example may help to illustrate how the ν s work. Let there be three classes of technology from which GPTs can be derived, developed by three pure research labs, $x = 1, 2, 3$, and three applied R&D facilities, $y = 1, 2, 3$. The nine individual ν s can be displayed in a 3×3 matrix. Each column modifies the productivities of each of the GPTs used in one of the applied R&D facilities. Each row modifies the productivities of the specific class- x GPT used in each of the applied R&D facilities.

As new GPTs are invented and adopted, the operative matrix of ν s evolves because the relevant column vector changes each time a new GPT of one class is adopted in a particular facility. Table 1 shows the fully identified ν s in an illustrative operative matrix that might have evolved after several inventions and adoptions of new GPTs.

The $z = 1$ subscript on all of the elements in the first column in the table tells us that the source of the ν s in the first R&D facility was a GPT of class 1 while the $n = 3$ superscript in row 1 column 1 tells us that the facility is using the third GPT of that class and that was the source of all the ν s in that column. The 2_2 and 1_3 superscripts on the second and third elements in that facility’s column vector tell us that the facility is using second GPT of class 2 GPT and the first GPT of class 3. The second column tells us that the source of the ν s in that column was the fourth GPT of class 1 ($z = 1$ in all subscripts and $n = 4$ in the first element in that faculty’s column vector). Evidently that GPT was rejected by facility 1, which is still using the third GPT invented in that class. This GPT is being used in combination with the second version of class 2 GPT (superscript 2_2) and the first version of the class 3 GPT (superscript 1_3). The third column tells us that the source of the ν s for the third applied R&D facility was the second version of class 3 GPT ($z = 3$ in all subscripts and $n = 2$ in the third element in that facility’s column vector), while the facility is using that GPT along with the second version of class 1 and the first version of class 2 GPTs.¹⁵

¹⁵ As a check, if a third version of GPT class 2 is invented and adopted by R&D facility 1, the first column in the operative matrix alters to read: $\nu_{1,2}^{3_1}; \nu_{1,2}^{3_2}; \nu_{1,2}^{1_3}$.

3 The model specified

The fixed supply of the composite resource of land and labour, R , is allocated by private price-taking agents in the consumption and applied R&D sectors and by a government that taxes the applied R&D and consumption sectors to fund pure research at an exogenously determined level. (In Section 6 below we change this allocation behaviour to analyse R&D support policy.) We make the assumption concerning the government because it reflects the modern reality that much pure research, whether funded through military and government procurement or conducted out of scientific curiosity, occurs in universities and public research laboratories such as CERN and other not-for-profit organizations.¹⁶ We make the following assumptions about the behaviour of private sector agents. (1) Agents are price takers. (2) Agents are operating in a non-ergodic environment and therefore under conditions of Knightian uncertainty for all of the reasons set out in Characteristic 8 in the introduction, which we model by assuming they do not know the probability distributions that are generating the disturbances on the outcomes. (3) Because agents cannot assign probabilities to alternative future consumption payoffs, they seek to maximize their profits on the basis of current prices, revising their behaviour as prices change period by period.

The constraint imposed by the composite resource is:

$$R_t = \sum_{i=1}^I r_t^i + \sum_{y=1}^Y r_t^y + \sum_{x=1}^X r_t^x, \quad (1)$$

where $r \in R$ is the amount of resource allocated to any one of the $I + Y + X$ possible activities in the economy and $i \in [1, I]$.

¹⁶Presented with an either-or choice between private and public financing of pure research, we think the better choice to reflect historical reality is public. It would be relatively easy to add some private behaviour allocating resources to the activity of creating pure knowledge and ultimately to the creation of GPTs. However, this is one of the many complications that we do not think would add as much to insight as it would to complexity so we have held with our either-or choice. Most major technologies developed in the twentieth century had significant amounts of public support at their early stages. (See, e.g. Ruttan (2001) As we go back in time before the twentieth century we do meet cases where private support is more important than public. But we also note that the historical case is mixed. For example, although it had rudimentary commercial beginnings in the West, writing was fully developed there by the state itself—in this case the Sumerian priests and their scribes (see Dudley 1991). Bronze and Iron seem to have enjoyed significant state support at least for key military applications, which may, as with so many twentieth century GPTs, have developed the production processes sufficiently for their uses to be extended fully to the civilian sector. Navigation that supported the invention and development of the three masted sailing ship certainly enjoyed state support, particularly in Portugal. Although British Red Brick universities were set up around 1900, US land grant colleges are 50 years earlier. Further back in time, much education was funded by the church (which we regard as part of the public sector) and by secular governments (both city states and more aggregated entities) in institutions that have existed in Europe since at least the thirteenth century. Many of the great inventions generated by China were entirely state financed, and so on.

3.1 The applied R&D and the consumption sectors

The output of applied knowledge at any point in time, a_t^y , from each applied R&D facility, y , depends on the amount of the composite resource, r_t^y , it uses and its productivity coefficient. This coefficient is the geometric mean of the set of GPTs (each denoted $(G_{n_x})_t$) each multiplied by its corresponding v term for the facility in question and then raised to the exponent β_x , as shown in Eq. 2.

$$a_t^y = \left[\prod_{x=1}^X \left(v_{y,z}^{n_x} (G_{n_x})_{t-1} \right)^{\beta_x} \right]^{\frac{1}{X}} (r_t^y)^{\beta_{X+1}}$$

$$\beta_x \in (0, 1] \forall x \in X, \beta_{X+1} \in (0, 1) \tag{2}$$

The stock of capital that embodies the applied knowledge generated from (and implicitly the pure knowledge used by) each facility accumulates according to:

$$A_t^y = a_t^y + (1 - \varepsilon) A_{t-1}^y, \tag{3}$$

where A_t^y is the stock of applied knowledge from facility y and $\varepsilon \in (0,1)$ is a depreciation parameter.

In the consumption sector, we make the simplifying assumptions (1) that there are the same number of applied R&D facilities and consumption industries, $Y = I$, and Eq. 2 that the knowledge produced in each of the facilities, y , is useful in only one corresponding consumption industry, i .¹⁷ The production function for each of the I industries in the consumption sector is expressed as follows:

$$c_t^i = (\mu A_{t-1}^y)^{\alpha_y} (r_t^i)^{\alpha_{Y+1}}, \tag{4}$$

where c_t^i is the flow of consumption output at time, t , from industry, i , and $\alpha_y \in (0,1] \forall y \in Y, \alpha_{Y+1} \in (0, 1)$, and $i = y$.

3.2 The pure knowledge sector

As already observed, there are X labs each producing one class of pure knowledge that leads to the occasional invention of a new GPT numbered n_x in the sequence of GPTs coming from that technology class. The productivity coefficient in each lab is the geometric mean of the various amounts of the Y different kinds of applied knowledge that are useful in further pure research

¹⁷This assumption greatly simplifies the calculations and is adequate for present purposes. In some specific applications, we need to make each consumption industry use more than one type of applied knowledge. At the extreme, if each uses all types, their production functions become:

$$c_t^i = \left[\prod_{y=1}^Y (\mu A_{t-1}^y)^{\alpha_y} \right]^{\frac{1}{Y}} (r_t^i)^{\alpha_{Y+1}}, \alpha_y \in (0,1] \forall y \in Y, \alpha_{Y+1} \in (0, 1).$$

(one for each applied R&D facility and each raised to a power σ_y).¹⁸ Each class of technology has a corresponding pure knowledge research lab. The output of pure knowledge, g_t^x , from lab x , at time, t , is a function of its productivity coefficient and the amount of the composite resource, r_t^x , devoted to that lab at time, t .

$$g_t^x = \left[\prod_{y=1}^Y ((1 - \mu) A_{t-1}^y)^{\sigma_y} \right]^{\frac{1}{\sigma_Y}} (\theta_t^x r_t^x)^{\sigma_{Y+1}}, \quad (5)$$

where $\sigma_y \in (0, 1]$, $\forall y \in Y$ and $\sigma_{Y+1} \in (0, 1)$. The term, θ_t^x , models the uncertainty surrounding the productivity of resources devoted to pure research as indicated in uncertainty source 8(i) listed in the introduction.

The stocks of *potentially useful* knowledge, Ω_t^x , produced by each of the X labs accumulate according to:

$$\Omega_t^x = g_t^x + (1 - \delta) \Omega_{t-1}^x, \quad (6)$$

where $\delta \in (0, 1)$ is a depreciation parameter.

New GPTs are invented infrequently in each of the X labs and their invention date is determined when the drawing of the random variable $\lambda_t^x \geq \lambda^{*x}$ (Characteristic 8 (ii)). For simplicity, we let the critical value of lambda for each of the X labs be the same: $\lambda^{*x} = \lambda^* \forall x \in X$. When at any time, t , $\lambda_t^x \geq \lambda^*$, indicating that a new version of class- x GPT is invented, the index t_{n_x} is reset to equal the current t , and n_x is augmented by one.

Because agents do not know how productive a new GPT will be over its lifetime (Characteristic 8 (iii)) they must make their adoption decisions with incomplete information. In this paper, as discussed in Section 1, we use the rule that the technology class- x challenger is adopted wherever its initial productivity is expected to exceed that of the class- x incumbent (determined according to Eq. 2, using the operative v matrix). The expected productivity from using the challenger is determined from its new $(G_{n_x})_t$ term defined in Eq. 8 below, the productivities of the other classes of GPT used by that facility and the challenger's expected \bar{v} s, which alter those productivities. As mentioned earlier, we assume that agents considering adopting the new GPT can correctly predict the potential v associated with the challenger, the minimum knowledge that they need to make some kind of evidence-based adoption decision, but that they assume that the challenger will cooperate with the other GPTs used by the facility with the same level of efficiency as did

¹⁸We maintain the assumption made by C&L and discussed in Section 1 above, that applied knowledge is divided into a proportion, μ , that is useful in the consumption industries and a proportion, $(1 - \mu)$, that is useful in pure knowledge production. LCB relax this assumption in their one at time GPT model and let the full stock of applied knowledge be useful in both the pure knowledge and consumption sectors. Doing so generates a knowledge spillover of the sort found in Lucas (1988) and Romer (1986) which could easily be incorporated into applications of the model developed here.

the incumbent, a prediction that may be falsified ex post since the potential v s brought in by the challenger may differ from those in the current operative matrix (Characteristic 8(iv)). Thus, agents assume that the other \bar{v} s in their facility’s column vector will be unchanged. For example, in the case illustrated in Table 1, if facility $y = 1$ is considering a challenger from class 2, it expects that the v in cell (1,2) will become the potential $\bar{v}_{1,2}^{3_2}$ associated with the challenger but that the v s in cells (1,1) and (1,3) of the existing operative matrix will remain unchanged (i.e., $\bar{v}_{1,2}^{3_1} = v_{1,1}^{3_1}$ and $\bar{v}_{1,2}^{1_3} = v_{1,1}^{1_3}$).

Since in each applied R&D facility the only v that agents expect to change is the one associated with the challenging x -class GPT, we can compare the productivities for any of the y facilities by simply comparing the $v_{y,z}^{(n-1)_x} (G_{(n-1)_x})_{t_{n_x}}$ that would result if the incumbent were left in place with the $\bar{v}_{y,z}^{n_x} (G_{n_x})_{t_{n_x}}$ that is expected to result if the challenger were adopted.¹⁹ This comparison is made in each of the Y applied R&D facilities at time $t = t_{n_x}$ so the test, stated generally for all applied R&D facilities, is:

$$\left[\bar{v}_{y,z}^{n_x} (G_{n_x})_{t_{n_x}} \right] \geq \left[v_{y,z}^{(n-1)_x} (G_{(n-1)_x})_{t_{n_x}} \right] \text{ for each } y \in [1, Y]. \tag{7}$$

If the test is passed, the new GPT is adopted in facility y .

If none of the Y applied R&D facilities adopts the GPT, it is returned to its pure knowledge industry. The indexes t_{n_x} and n_x are incremented back to their previous values—it is as if the favourable drawing of $\lambda_t > \lambda^*$ had not occurred. Pure research then continues to improve the new GPT and it is reconsidered every period until it is adopted or superseded in each applied R&D facility by a newer, superior version in its class.

The evolving efficiency with which the GPT delivers its services is shown in Eq. 8 below.

$$(G_{n_x})_t = (G_{(n-1)_x})_{(t-1)_{n_x}} + \left(\frac{e^{\tau+\gamma(t-t_{n_x})}}{1 + e^{\tau+\gamma(t-t_{n_x})}} \right) \left(\Omega_{t_{n_x}}^x - (G_{(n-1)_x})_{(t-1)_{n_x}} \right) \tag{8}$$

where e is the natural exponential.

The equation shows the efficiency of the GPT, $(G_{n_x})_t$, increasing logistically as the full potential of the GPT is slowly realized. t_{n_x} is the invention date of the version n_x , of the class- x GPT, $\Omega_{t_{n_x}}^x$ is the full *potential* productivity of the new version of GPT x , $(G_{(n-1)_x})_{(t-1)_{n_x}}$ is the *actual* productivity of the version that it replaced, evaluated at the time at which that earlier version was last

¹⁹If the \bar{v} s attached to the other GPTs operating in applied R&D facility, y , were also expected to diverge from those in the appropriate column of the operative matrix, we would have to calculate the complete value of the a_t^y term from Eq. 2 as it would be under the existing GPTs and the operative set of v s and then compare it with what that a term would be given the new column vector of expected \bar{v} s.

used, $t_{(n-1),x}$ and γ and τ are calibration parameters that control the rate of logistic diffusion.²⁰ The evolution of efficiency proceeds as follows. Initially, since $t_{n_x} = t$ (and because γ is very small, 0.07 in our simulations), the value of the efficiency coefficient is close to zero so that the initial productivity of the challenging GPT is close to that of the incumbent. As t increases over time the value of the efficiency coefficient approaches unity so that the GPT's productivity approaches its full potential.

In addition to modelling Characteristics 8(iii) and 8(iv), which we discussed above, the interaction of the G and the ν terms model uncertainty about when a new GPT will begin to be replaced by a challenger (Characteristic 8(vi)), how long it will take to displace the incumbent GPT (Characteristic 8(vii)), and whether or not a GPT will be completely displaced (Characteristic 8(viii)). Thus, they model the many aspects of the general observation that the applied potential of a GPT cannot be precisely predicted when it is originally being developed.

In the subsequent periods, the test in Eq. 7 is modified to note the productivity changes that occur over time:

$$\left[\bar{v}_{y,z}^{n_x} (G_{n_x})_t \right] \geq \left[v_{y,z}^{(n-1)_x} (G_{(n-1)_x})_t \right] \tag{7'}$$

for each $y \in [1, Y]$ that has not yet adopted GPT G_{n_x} .²¹

In our model, the economy's GDP is the current period's output of the consumption and applied R&D sectors plus the cost of the composite resource used in the pure knowledge sector.²² The knowledge produced by the applied R&D sector is an addition to the capital stock because it gets embodied in capital. In contrast, pure knowledge is an intermediate input into the applied research sector, being of no use in production until it is turned into applied knowledge. It is not, therefore, a part of the capital stock. Thus for our purposes, the stock of applied knowledge embodied in physical and human capital is the capital stock while the flow of the production from the applied sector plus the allocation of resources to the pure knowledge sector valued at their marginal products is gross investment.

3.3 Resource allocation

As we have already noted, in the pure knowledge sector the government pays for and allocates a fixed amount of the generic resource, R , to each of the pure

²⁰ Agents in our model do not know the parameter values of the logistic function and these can be either fixed or determined randomly in the model each time a new GPT arrives. When determined randomly, they model uncertainty source 8(v).

²¹ This comparison must be made in every time period, for every applied R&D facility that has not yet adopted the latest version of every class of GPT in existence. It must also be made for all of the GPTs of any given class that are more recent than the one currently being used by a given facility.

²² Because the pure knowledge sector only periodically produces a useful GPT, we adopt the standard national accounting convention of valuing the output of that sector at its input costs in each period.

knowledge producing labs.²³ Producers in the applied R&D and consumption sectors maximize their profits each period taking prices as given. (We suppress time subscripts in Eqs. 9 through 16 because agents are not foresighted and are consequently performing a static maximization in each period.)²⁴ The prices for output from the I consumption industries are derived from the maximization of an aggregate utility function, which we assume is additively separable across the I consumption goods.

$$U = \sum_{i=1}^I (c^i)^{\phi^i} \text{ and } \phi^i = \phi^{i'} = 1, i \neq i' \forall i, i' \in I \quad (9)$$

where U is aggregate utility and c is defined as above.

The simplifying assumption that the exponents on each type of consumption are all equal to unity can, of course, be modified for specific applications. Maximizing this utility function and rearranging the first order conditions (FOCs) yields:

$$\frac{MU^{i=1}}{MU^{i \neq 1}} = \frac{P^{i=1}}{P^{i \neq 1}} = \frac{\phi^{i=1} (c^{i=1})^{\phi^{i=1}-1}}{\phi^{i \neq 1} (c^{i \neq 1})^{\phi^{i \neq 1}-1}} \quad (10)$$

Since $\phi^i = 1 \forall i \in I$ it follows that $P^{i=1} = P^{i \neq 1}$, i.e., the relative prices of all consumptions goods are unity.

We assume a competitive equilibrium in the market for the composite resource. This implies that it earns the same wage, w , regardless of where it is allocated.

²³For specific applications where the private/public investment mix in pure knowledge R&D is under investigation it is possible to accommodate a mix of private and government agents allocating resources to pure R&D. To do so requires that private agents maximize their expected current-period profits as is done in the consumption and applied R&D sectors of the current model. This in turn requires that we have an expected price associated with the output of the pure knowledge sector that is transacted when a GPT is invented and taken up by at least one applied R&D facility. The formation of expectations about these prices can be as simple or complex as the application requires.

²⁴Agents' lack of foresight has some implications for the allocation of resources in the model. We have run a number of simulations where agents have one or two periods of foresightedness in the sense that they are aware of the impact of more than one period's allocation decisions on the marginal product of resources in applied R&D. For the simulations run so far we find that with foresight agents allocate slightly more resources to applied R&D because they are aware of its future effects in raising consumption output but that this has a negligible impact on output over time. Depending on the relative allocations of resources to applied versus pure research and on the arrival rates of GPTs, the foresightedness that we allow sometimes causes the output growth to increase and at other times causes it to decrease relative to the simulation with no foresight. But in all cases, the overall impact is negligible. We do not allow agents to have infinite foresight because such an assumption would be meaningless in our model which incorporates Knightian uncertainty and because, even without uncertainty, having full foresight over the lifetime of a GPT that often extends for more than a century is quite impossible.

Each consumption industry maximizes its profits taking the price of its consumption output, P^i , and the prices of its inputs, composite resource, w , and applied knowledge, P^y , as given. Profits, π^i , are expressed as:

$$\pi^i = P^i c^i - w r^i - P^y A^y \tag{11}$$

Profit maximization yields the following FOCs in each of the I consumption industries:

$$\begin{aligned} P^i mp(r^i) - w &= 0 \\ P^i mp(A^y) - P^y &= 0 \end{aligned} \tag{12}$$

where mp represents marginal product. From the first FOC, the assumption the $P^i = 1$, and the definition of the production function for industry i we get:

$$r^{i*} = \left[\frac{\alpha_{Y+1}}{w} (\mu A^y)^{\alpha_y} \right]^{\frac{1}{1-\alpha_{Y+1}}} \tag{13}$$

This is the reduced form expression for the demand for the composite resource in each consumption industry, i .

From the combination of both FOCs from the profit function for consumption industry i and the definition of the production function we get:

$$\frac{w}{P^y} = \frac{\alpha_{Y+1}}{\alpha_y} \frac{A^y}{r^i}$$

which implies that:

$$P^{y*} = \frac{\alpha_y w}{\alpha_{Y+1} A^y} \left[\frac{\alpha_{Y+1}}{w} (\mu A^y)^{\alpha_y} \right]^{\frac{1}{1-\alpha_{Y+1}}} \tag{14}$$

Each applied R&D facility maximizes profits taking the price of its applied knowledge output, P^y , and the composite resource, w , as given. The pure knowledge input in the form the currently adopted set of X GPTs is provided freely to the applied R&D facilities by the government financed labs. Profits are expressed as:

$$\pi^y = P^y a^y - w r^y \tag{15}$$

Maximization of the profit function and algebraic manipulation yields the following FOC:

$$P^y mpr^y - w = 0$$

The demand for the composite resource from each of the Y applied R&D facilities is thus:

$$r^{y*} = \left[\beta_{X+1} \left[\prod_{x=1}^X (y_{y,z}^{n_x} (G_{n_x})_t)^{\beta_x} \right]^{\frac{1}{X}} \frac{P^{y*}}{w} \right]^{\frac{1}{1-\beta_{X+1}}} \tag{16}$$

With these resource demand equations, we now have a complete description of the allocation of the composite resource across the three sectors.

4 The model simulated

To simulate our model, we first make some simplifying assumptions, all of which can be altered for specific applications. We restrict the model to three industries within the consumption sector, three facilities in the applied R&D sector and three labs in the pure knowledge sector ($I = Y = X = 3$). We choose $\alpha_{Y+1} = \beta_{X+1} = \sigma_{Y+1}$ and $\alpha_y = \beta_x = \sigma_y$ to impose symmetry across sectors and specific activities within sectors (i.e., industries, facilities and labs). To simplify the utility function and make relative prices unity we set $\phi_i = 1 \quad \forall i \in [1, I]$.

We choose values of the parameters and initial conditions with the overall objective that the model will replicate the accepted stylised facts of economic growth. We set the values of most of the initial conditions at unity because none of the long run characteristics of the model are influenced by initial conditions.²⁵ Some values are chosen to ensure consistency with observed data in the following ways: $\alpha_{Y+1} = \beta_{X+1} = \sigma_{Y+1} = 0.3$ ensures diminishing returns to the composite resource in all lines of activity; $\varepsilon = \delta = 0.025$ produces an average annual growth rate between 1.5% and 2%;²⁶ $\lambda^* = 0.66$ allows a GPT within each class to arrive on average every 35 years, but with a large variance in arrival dates. We choose $\gamma = 0.07$ and $\tau = -7$ so that 90% of a GPT's potential is translated into its actual efficiency over the first 120 years of its life. We choose $\mu = 0.95$ in order to set the income shares of the labour and capital (physical and human) at approximately 0.3 and 0.7.²⁷ We set $\alpha_y = \beta_x = \sigma_y = 1$ to ensure that knowledge has constant returns.

Table 2 gives the parameter values and initial conditions used to simulate the results of the multi-GPT model as reported in the text and shown in the figures. The set of θ^x are random variables distributed uniformly with support $[0.9, 1.1]$, mean 1, and variance $(0.4)^2/12$, which sets modest bounds on the uncertainty concerning the productivity of pure research. The λ 's are derived from beta distributions, where each distribution is defined as $\text{beta}(x|\psi, \eta) = \frac{x^{(\psi-1)}(1-x)^{(\eta-1)}}{\text{Beta}(\psi, \eta)}$ with support $[0,1]$, mean $\left(\frac{\psi}{\psi+\eta}\right)$ and variance $\frac{\psi\eta}{(\psi+\eta)^2(\psi+\eta+1)}$. $\text{Beta}(\psi, \eta)$ is the Beta function, and ψ and η are parameters which take on positive integer values. We choose $\psi = 5$ and $\eta = 10$.

Next we need to determine the ν s in each challenger's potential matrix. They are calculated as the ν s in the current operative matrix modified by the addition of a random variable, ρ , drawn from a uniform distribution with support

²⁵The initial value of $t_{n_x} = 2$ is chosen because we have lagged variables indexed on it and MatLab does not allow zero as an index value.

²⁶We choose this range based on Angus Madison's historical data set (see Madison's web pages <http://www.ggd.net/maddison/> for the complete data set). We calculate the average annual growth rate of GDP per person from 1870 to 2003 for the USA to be 1.86% and for Canada to be 1.96%.

²⁷ μ can alter shares because it alters the marginal product of labour in consumption, and labour is paid the same wage everywhere.

Table 2 Numerical simulation of the multi-GPT model

| | | | |
|---|----------------------------------|------------------------------------|---------------------|
| Parameter values: | | | |
| $\alpha_y = 1 \forall i \in [1, Y]$ | $\alpha_{Y+1} = 0.3$ | $\beta_x = 1 \forall x \in [1, X]$ | $\beta_{X+1} = 0.3$ |
| $\sigma_y = 1 \forall y \in [1, Y]$ | $\sigma_{Y+1} = 0.3$ | $\gamma = 0.07$ | $\tau = -7$ |
| $\phi_i = 1 \forall i \in [1, I]$ | $\varepsilon = 0.025$ | $\delta = 0.025$ | $\mu = 0.95$ |
| $I = 3$ | $Y = 3$ | $X = 3$ | $\lambda^* = 0.66$ |
| $R = \sum_{i=1}^I r_i^j + \sum_{y=1}^Y r_i^y + \sum_{x=1}^X r_i^x = 1000$ | | | |
| Initial conditions: | | | |
| $(G_{n_x})_0 = 1 \forall x \in (1, X)$ | $A_0^y = 1 \forall y \in [1, Y]$ | $n_x = 1$ | $\theta_0^x = 1$ |
| $t_{n_x} = 2$ | $P_0^i = 1 \forall i \in [1, I]$ | | $d_{y,x} = 1$ |
| $w_0 = 1$ | | | |

$[-0.05, 0.05]$. Thus the ability of a class- x challenging GPT to cooperate with the other classes of incumbent GPTs used in a given research facility varies randomly around the ability of the class- x incumbent to do the same thing. The support for ρ is chosen so that deviations from operative to potential v_s are not large. For the example in Table 1, the potential v_s associated with a challenger GPT $n + 1$ from class $x = 2$ are derived from the incumbent's v_s as follows:

$$v_{y,2}^{(n+1)_2} = \begin{cases} v_{y,2}^{n_2} + \rho & \text{if } (v_{y,2}^{n_2} + \rho) \in [0.5, 1.5] \\ 0.5 & \text{if } v_{y,2}^{n_2} + \rho < 0.5 \\ 1.5 & \text{if } v_{y,2}^{n_2} + \rho > 1.5 \end{cases} \quad y [1, 3] \quad (17)$$

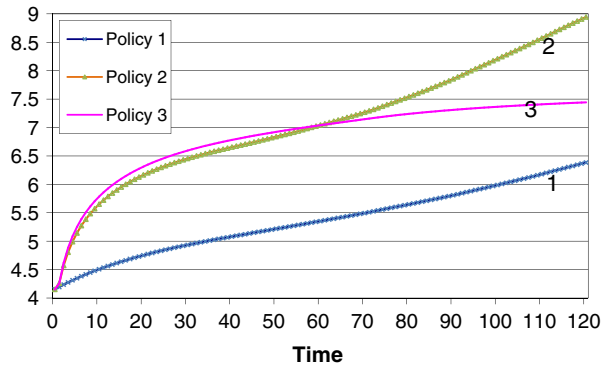
5 The model applied

To demonstrate some important properties of our model and show its applicability to practical issues, we now provide three illustrations. We also note that in the study paper version of this paper, Carlaw and Lipsey (2007), we show that the model can duplicate the accepted stylised facts of economic growth. This is not a strong test for a model as flexible as ours, but it is a necessary test in the sense that, if it could not be passed, the model would be of little use.

5.1 Policy: the allocation of funds between pure and applied R&D

We first illustrate the potential of the model to address policy issues by considering the appropriate mix of public funding to support pure and applied research. We make a number of simulations in which pure research is solely financed by the government while applied R&D is financed by a mix of government and private funds. In total, the government allocates 3.78% of

Fig. 2 Natural log of output under different R&D support policies



GDP to knowledge production.²⁸ In the first simulation, all of the government support goes to pure knowledge production (Policy 1). In the second, after several iterations, the government reallocates its funds to split them evenly between pure and applied R&D (Policy 2). In the third, again after the same number of iterations as Policy (2), the government allocates all funding to applied research (Policy 3). To make these three simulations display solely the effects of the different research allocation policies, we ran the first simulation then imposed the realized random variables from that simulation on the next two.

Figure 2 gives the natural logs of the GDP paths, starting at the date of these alternative policy changes. It shows that there is a large gain over any time period in having the government support some applied R&D (Policy 2) rather than solely supporting pure research (Policy 1). The best mix would have to be discovered iteratively and some experimentation shows that it is closer to 80% for applied R&D, rather than the 50% shown in the figure. The government must not, however, go too far. The figure shows that there is an early gain to allocating all public support to applied R&D (Policy 2)—provided, of course, that some amount of pure knowledge had already been accumulated and that existing GPTs have some unexploited potential. This gain is eventually more than offset by the loss from not maintaining support for the pure knowledge production that creates new GPTs. As the figure shows, Policy 3 ceases to out-perform Policy 2 after about 50 iterations and, since it eventually leads

²⁸We choose 3.78% because the United States allocates about 2.63% of total GDP to R&D and 2.3% of GDP to post-secondary education and research making 4.93% in total. We choose the figure 3.78% because we make the crude assumption that half of post-secondary expenditure goes into R&D activities. But since this ignores the significant fraction of total government expenditure that goes to R&D, particularly from the Department of Defense, this number would need to be significantly increased for practical applications. The upper limit would be the total public expenditure (not including transfers) which is 19.1% of GDP.

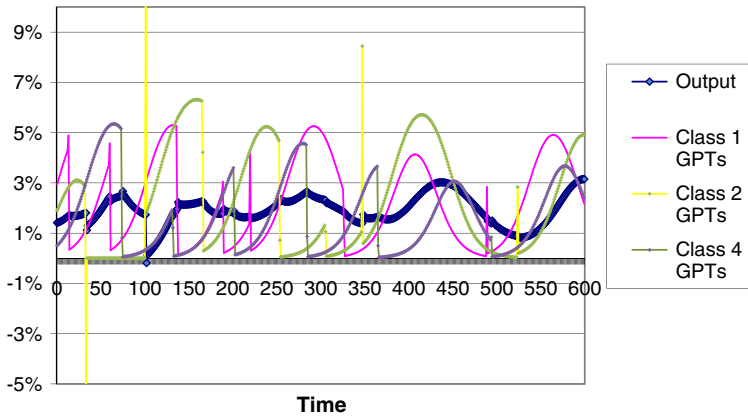


Fig. 3 Growth rates of output and useful pure knowledge

to a constant income (zero growth), it is eventually also out-performed by Policy 1.

Although this analysis is crude, the qualitative results suggest some interesting possibilities for further study in more sophisticated, context-specific versions of our model.

1. Given that the public sector is financing pure research, the private sector may under-produce applied R&D making some mix of private and publicly supported applied R&D preferable.²⁹
2. The highest possible rate of growth can be attained, and sustained for some time, by diverting all research resources to applied R&D. However, the absence of new GPTs created by pure research is sooner or later felt in a falling growth rate that eventually reaches zero, making this the worst allocation policy in the very long term.

5.2 GPTs and the whole economy

Because most other GPT models allow only one GPT to exist at any one time, the performance of the whole economy closely mirrors that of the existing GPT. This has led some researchers to look to the economy's performance to

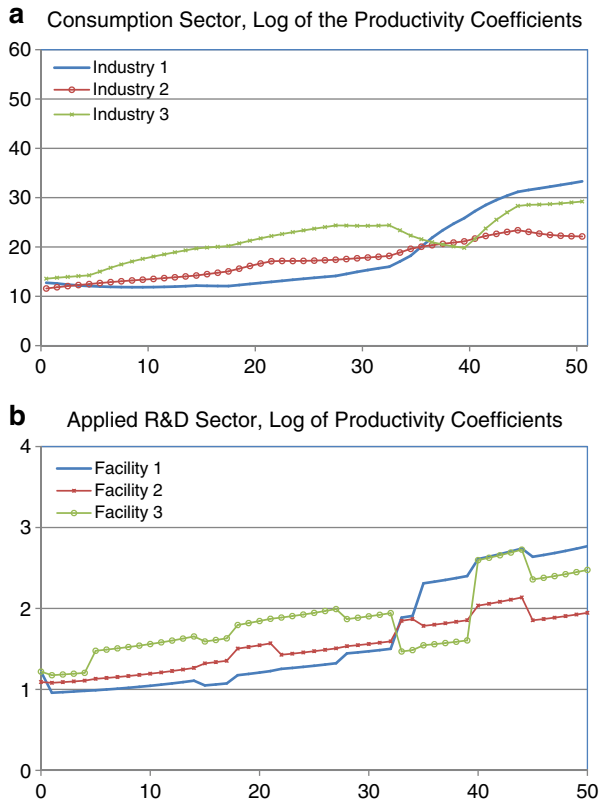
²⁹In our model, private agents allocate resources to applied R&D based on the effect of such research on the near-future productivity of the consumption goods industries. If private sector agents looked further into the future, there would be an incentive to allocate more to current R&D. However as long as they only look a few periods ahead when calculating their expected returns (as is often alleged) the public return from such R&D will exceed the private return, providing a case for some diversion of public funds to encourage private R&D (as governments do in many countries by policies such as tax relief or direct subsidies for R&D).

deduce something about the arrival and impact of an alleged GPT. It has also lead to the commonly held but erroneous belief that the early stages of the introduction of every new GPT will be associated with a slowdown in the rates of growth of productivity and national income. We might suspect, however, that the relation between an individual GPT and the whole economy that was present in a one-GPT-at-a-time model would not show up when several GPTs all at different stages in their evolution exist at any one time.

Figure 3 illustrates that this suspicion is correct in our model. The heavy line is the growth rate of aggregate output while the lighter lines are the growth rates of the efficiency of each class of GPT as determined from Eq. 8 in the simulation.

The figure shows both that the arrival and evolution of any one GPT does not typically determine the pattern of growth of aggregate output and that the growth of aggregate output cannot be used to infer the arrival and evolution of individual GPTs in a multi-GPT world. The spikes in the growth rate do not just signal the arrival of a new GPT. They occur whenever a new GPT is adopted by a particular applied facility, which may occur long after the

Fig. 4 The evolution of productivity in applied R&D and consumption



original arrival of a new GPT. Sometimes the spikes are negative because of a bad realization of the v s for GPTs of other classes that cooperate with the new GPT. Eventually however the new GPT does pay off in enhanced productivity.³⁰

5.3 The separate evolution of productivity in firms and sectors

As mentioned in the introduction, productivity is a single scalar variable in most existing models of both exogenous and endogenous growth. In our model, productivity has the potential of evolving differently for each activity (i.e., industry, facility or lab) in each sector, although in the present version we make the pure research laboratories all have the same productivity. Figure 4 shows the evolution of productivity for one run of 50 periods for each of the three industries and facilities in each of the two sectors, consumption and applied R&D. Although these are only illustrative, they show the potential to model what is observed in fact: the productivity of different sectors, and of different agents within any one sector, evolves differently but not totally independently of each other.

6 Summary and outlook

The model presented here foregoes the use of dynamically stationary, perfectly foresighted, equilibrium techniques where agents maximize over the infinite horizon (i.e., it avoids the ergodic-axiom). Instead, agent behaviour is modelled as profit seeking under several different types of uncertainty. This allows us to incorporate many of the relevant historical facts of technologically driven growth. We do this by modelling technology as structured and economic growth as a dynamically non-stationary process driven by a set of endogenously created and evolving GPTs, many of which exist simultaneously and interact. Across classes of technology, our GPTs often exhibit complementarities; within classes of technology, different GPTs often compete with each other.

Having incorporated many of the historical facts concerning technologically driven growth, we also observe that flexibility of our modelling technique allows it to be applied to a number of specific growth issues. To illustrate, we applied the model to the issue of publicly funded pure and applied R&D, to the issue of whether one can test of the existence of GPTs by observing aggregate output, and to the issue of how micro behaviour aggregates to economic output is a set of complicated and interdependent interactions. Many further applications of the model are foreseeable. Here are just three of the

³⁰The negative shock followed by a productivity gain is similar to what we seen in the 1970s–1990s with the ICT revolution.

many possibilities, each of which requires a full paper that will build on and apply the present model.

1. LCB hypothesise that slowdowns depend on both the nature of individual GPTs and how each interacts with the entire structure of the economy, which they call its ‘facilitating structure’.³¹ The present model meets one of the necessary conditions for testing this hypothesis and for developing a model of Schumpeterian long wave cycles of the sort studied empirically by Freeman and Louça (2001) in that it has multiple co-existing GPTs that can stand in various relations of competitiveness and complementarity with each other. What is now required is to model the facilitating structure and then study how individual GPTs impact on it.³² This not insignificant job is the subject of our current research.
2. The model has been used by Carlaw and Lipsey (2010) to test the real business cycle and time series econometric “proofs” of the stationarity of the underlying process generating GDP, which allows long run growth and short run cycles to be separated. Since our model is non-stationary, the GDP series that it generates can be used to check on the discriminatory power of the RBC and time series empirical tests for stationary. Carlaw and Lipsey (2010) find that the series generated by our model typically pass the stationarity test and also match the standard RBC growth facts. Their findings call into question the “proof” of the stationarity of real data and the discriminatory power of the time series tests.³³
3. In spite of many criticisms, total factor productivity (TFP) is commonly used to measure technological change. Lipsey and Carlaw (2004) and Felipe and McCombie (2006) have criticised this procedure from a theoretical standpoint and LCB (2005) used the output from their single-GPT models, where technological change is known, to test the efficacy of the TFP measure. They concluded that it was not even a rough measure technological change, at least when applied to the data generated by their models. Because the present model has many GPTs, it is much closer to what is seen in the real world. So the data it generates can be used as a stronger test of the efficacy of TFP to measure the technological change that is known from within the model.

References

- Aghion P, Howitt P (1998) On the macroeconomic effects of major technological change. In: Helpman E (ed) *General purpose technologies and economic growth*. MIT, Cambridge, pp 21–144

³¹This concept is defined and its significance discussed in LCB 55–63.

³²LCB make a start at modelling the interaction between a GPT and the facilitating structure in their one-GPT model on pp 482–496.

³³These results were presented in Carlaw and Lipsey (2010) at the 2010 International Schumpeter Society meetings on Innovation, Organization, Sustainability and Crisis in Aalborg Denmark.

- Arrow KJ (1962) Economic implication of learning by doing. *Rev Econ Stud* 29:155–173
- Australian Research Council (2006) Consistent national policies for converging technologies: some preliminary conclusions. Centre for strategic economic studies. Victoria University, Melbourne
- Bresnahan TF, Trajtenberg E (1995) General purpose technologies: engines of growth? *J Econom* 65(1):83–108
- Carlaw KI, Lipsey RG (2006) GPT-Driven, endogenous growth. *Econ J* 116:155–74
- Carlaw KI, Lipsey RG (2007) Sustained growth driven by multiple co-existing GPTs. Simon Fraser University Department of Economics Working Paper Series, dp07–17
- Carlaw KI, Lipsey RG (2010) Darwinian versus Newtonian views of the economy: empirical tests of Schumpeterian versus new classical theories. Presented at the 13th international Schumpeter society conference. Aalborg, Denmark, <http://www.schumpeter2010.dk/index.php/schumpeter/schumpeter2010/paper/view/466/89>
- Cohen WM, Levinthal DA (1990) Absorptive capacity: a new perspective on learning and innovation. *Adm Sci Q* 35(1):128–152
- Crafts N (2003) Steam as a general purpose technology: a growth accounting perspective. *Econ J* 114(495):338–351
- David PA (1990) The dynamo and computer: an historical perspective on the modern productivity paradox. *Am Econ Rev* 80:355–361
- David PA (1997) Path dependence and the quest for historical economics: one more chorus. Discussion Papers in Economic and Social History, no 20. University of Oxford
- David PA, Wright G (1999) Early twentieth century productivity and growth dynamics: an inquiry into the economic history of 'our ignorance. Discussion Papers in Social and Economic History, no 33. University of Oxford
- David PA, Wright G (2003) General purpose technologies and productivity surges: historical reflections on the future of the ICT revolution. In: David PA, Thomas MP (eds) *The economic future in historical perspective*. Oxford University Press, Oxford
- Davidson P (1993) Austrians and post Keynesians on economic reality: rejoinder to critics. *Crit Rev* 7(2&3):423–444
- Davidson P (2002) *Financial markets, money and the real world*. Edward Elgar, Cheltenham
- Davidson P (2009) Risk and uncertainty in economics. Manuscript presented to the conference on The economic recession and the state of economics. Westminster, London, February 6, 2009
- Dudley L (1991) *The word and the sword: how techniques of information and violence have shaped our world*. Cambridge University Press, Cambridge
- Felipe J, McCombie JSL (2006) The tyranny of the identity: growth accounting revisited. *Int Rev Appl Econ* 20(3):283–299
- Forman C, Goldfarb A, Greenstein S (2003) The geographic dispersion of commercial internet use. In: Cranor LF, Wildman SS (eds) *Rethinking rights and regulations: institutional responses to new communication technologies*. MIT, Cambridge, pp 113–145
- Freeman C, Louça F (2001) *As time goes by: from the industrial revolutions to the information revolution*. Oxford University Press, Oxford
- Greasley D, Oxley L (1997) Endogenous growth or 'Big Bang': two views of the first industrial revolution. *J Econ Hist* 57:935–49
- Helpman E ed (1998) *General purpose technologies and economic growth*. MIT, Cambridge
- Helpman E, Trajtenberg M (1998a) A time to sow and a time to reap. in Helpman, pp 55–84
- Helpman E, Trajtenberg M (1998b). Diffusion of general purpose technologies, in Helpman, pp 85–119
- Hornstein A, Krusell P, Violante GL (2005) The effects of technical change on labor market inequalities. In: Aghion P, Durlauf S (eds) *Handbook of economic growth*, vol 1. Elsevier, Amsterdam, pp 1275–1370
- Jovanovic B, Rousseau PL (2005) *General purpose technologies*. NBER Working Paper 11093. NBER, Cambridge Mass
- Lindmark S (2005) Coordinating the early commercialization of general purpose technologies: the case of mobile data communications. *Biotechnology & telecommunications innovation: conditions for emerging technologies* 7(1):30–60

- Lipsev RG (2007) Reflections on the general theory of second best at its golden jubilee. *Int Tax Public Finance* 14:349–364
- Lipsev RG, Carlaw KI (2004) Total factor productivity and the measurement of technological change. *Can J Econ* 37–4:1118–1150
- Lipsev RG, Carlaw KI, Bekar CT (2005) *Economic transformations: general purpose technologies and long-run economic growth*. Oxford University Press, Oxford
- Lucas RE (1988) On the mechanics of economic development. *J Monet Econ* 22:3–42
- Mokyr J (2006) The great synergy: the European enlightenment as a factor in modern economic growth. In: Wilfred D (ed) *Understanding the dynamics of a knowledge economy*. Edward Elgar, Cheltenham, pp 7–41
- Romer PM (1986) Increasing returns and long run growth. *J Polit Econ* 94:1002–1037
- Rosenberg N (1982) Inside the black box: technology and economics, chapter 6: learning by using. Cambridge University Press, Cambridge
- Rosenberg N (1994) *Exploring the black box: technology, economics, and history*. Cambridge University Press, Cambridge
- Ruttan VW (2001) *Technology, growth, and development: an induced innovation perspective*. Oxford University Press, New York
- Ruttan VW (2005) *Is war necessary for economic growth: military procurement and technology development*. Oxford University Press, New York
- Safarian AE, Dobson W (2002) Do institutions matter for economic growth. East Asia in transition: economic and security challenges. In: Safarian AE, Dobson W (eds) *University of Toronto Press, Toronto*, pp 1–20
- Sapio S, Thoma GLD (2006) The growth of industrial sectors: theoretical insights and empirical evidence from U.S. manufacturing. *Working Paper Series, 2006/9*. Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy
- van Zon A, Kronenberg T (2007) General purpose technology and energy policy. In: Bretschger L, Smulders S (eds) *Sustainable resource use and economic dynamics*, pp 71–99
- van Zon A, Fortune E, Kronenberg T (2003) How to sow and reap as you go: a simple model of cyclical endogenous growth. In: *Research memorandum series, 029*. Maastricht Economic Research Institute on Innovation and Technology (MERIT), Maastricht
- Wehrli R, Saxby G (2006) The importance of technological change for climate policy assessment. *Work in progress Institute of Economic Research, ETH*