

GENERAL PURPOSE TECHNOLOGIES IN THEORY, APPLICATION AND CONTROVERSY: A REVIEW

by

Clifford Bekar, Kenneth Carlaw, Richard Lipsey

Contacts

Clifford Bekar: bekar@lclark.edu

Kenneth Carlaw: kenneth.carlaw@ubc.ca

Richard Lipsey: rlipsey@sfu.ca

Abstract

Distinguishing characteristics of General Purpose Technologies (GPTs) are identified and definitions discussed. Our definition includes multipurpose and single-purpose technologies, defining them according to their micro-technological characteristics, not their macro-economic effects. Identifying technologies as GPTs requires recognizing their evolutionary nature, and accepting possible uncertainties concerning marginal cases. Many of the existing ‘tests’ of whether particular technologies are GPTs are based on misunderstandings either of what GPT theory predicts or what such tests can establish. The development of formal GPT theories is outlined, showing that only the early theories predicted the inevitability of GPT-induced slowdown and surges. More recent GPT theories, designed to model the characteristics of GPTs, do not imply the necessity of specific macro effects. We show that GPTs can rejuvenate the growth process without causing slowdowns or surges. We conclude that existing criticisms of GPT theory can be resolved and that the concept remains useful for economic theory.

Key Words: General Purpose Technologies, technological change, patents, slowdowns, surges, growth theories, productivity

JEL Classes: N00, O30, O33, O40, O41.

GENERAL PURPOSE TECHNOLOGIES IN THEORY, APPLICATIONS AND CONTROVERSY: A REVIEW¹

Two main approaches to studying economic growth over the very long-run are found in the literature: the analysis of economic historians and the formal model building of economic theorists. Both literatures agree that technological change is a major driver of economic growth. Economic historians take a nuanced view of technology as a complex and multi-layered phenomenon, often focusing on historically significant signpost technologies such as the heavy plough, the printing press, the steam engine and their transformative effects on the economy, institutions, and social structures. Until recently, and in contrast, most formal theories of economic growth modelled technology as a scalar that is either an argument in, or multiplicative constant to, an aggregate production function. The contribution of technology to growth in this framework is usually estimated as the residual output after accounting for the contribution of measured aggregate inputs.

With the publication of Bresnahan and Trajtenberg's (1995) paper² "General Purpose Technologies: 'Engines of Growth?'" some theorists began to model technology and technological change in a more nuanced way than in most traditional macro growth models.³ Since that time, the General Purpose Technology (GPT) research program has been evolving piecemeal. In largely uncoordinated efforts, this research program has defined GPTs, modeled their evolution, tested their predicted and presumed effects, and criticised the very usefulness of the concept. As a result, the GPT literature contains many conflicting ideas and unresolved controversies. Given these developments, it seems appropriate to: (i) review the GPT literature to locate and remove anomalies, inconsistencies and other deficiencies; and, (ii) note instances where GPT theory has proven productive in economists' study of economic growth.

We find two main classes of criticism of GPTs. The first comes from economic historians who question the usefulness of GPTs as a way to think about technology and innovation. We take Field (2011) as representative of historians' criticisms and, where they are relevant to our discussions throughout our paper, we provide our responses to his most important arguments. The second comes from economists who have sought to test what they perceive to be predictions of GPT theory. We argue that both of these types of criticisms need to be met because the program of modelling technology as more than just a scalar is important for theoretical treatments of economic growth.

There is a significant overlap between the work of economists modeling economic growth and that of economic historians seeking to describe and understand the processes that lead to such growth. Given the distinct nature both of the questions motivating economic historians and economic theorists and their techniques for answering them, they may come to different conclusions regarding the potential usefulness of the concept of a GPT. It is our view that the program of modeling technology and technical change in a more structured manner than can be

¹ This is a revised version of a paper first presented at the International Schumpeter Society Conference, Montreal, Canada, July 2016. We are indebted to the participants in the session on General Purpose Technologies and to an anonymous referee for valuable comments and suggestions.

² This paper was first published in 1992 as an NBER working paper and later published in 1995 in *The Journal of Econometrics*. Hereafter we refer to the 1995 paper.

³ For a discussion of a range of precursors to this general concept see Cantner and Vannuccini (2012).

caught by a single variable is a critically important research program in growth theory, and it is to this issue that this paper's analysis is directed

We begin by identifying six technological characteristics that serve to distinguish GPTs from other technologies. We then discuss alternative definitions of GPTs, arguing that GPTs should be defined in terms of their micro-technological characteristics, not their macro-economic effects. We then argue for the usefulness of a broad definition of GPTs, one that does not exclude single-generic-purpose technologies that are widely used across the whole economy (e.g., the printing press and three-masted sailing ship). Next we discuss how to identify a GPT. This requires that we recognise two things. First, like any major technology, a GPT is not static but evolves continually over time. Second, as with any set of theoretical classes meant to stylise real observations, there is always uncertainty at the boundaries concerning whether a particular item is inside or outside of some given class. We illustrate this identification issue with two examples, electricity and the 3-masted sailing ship. This leads us to consider some existing 'tests' of whether particular technologies are GPTs. We argue that some of these tests are based on a misunderstanding of what are and are not predictions of GPT theory, while others are based on a misunderstanding of what such tests can actually establish. To deal with the latter point we address some methodological issues. Having rejected the allegations that these tests show that some particular technologies are not GPTs, we consider what can be learned from empirical work related to GPTs. Having done the required ground clearing with respect to definitions and tests, we move to assess the evolution of formal theories of GPTs, paying particular attention to what they do and do not predict. Most importantly we show that only the first post-Bresnahan-and-Trajtenberg generation of GPT theories predicted the inevitability of an aggregate slowdown and/or productivity surge after the introduction of a new GPT, and that these theories were designed to produce such predictions. More recent GPT models that are designed to model their characteristics are agnostic regarding their macro effects. These models clearly demonstrate that GPT theory does not imply the necessity of slowdowns and/or productivity surges. This leads us to consider how GPTs can rejuvenate the growth process without necessarily producing either slowdowns or surges. In the final section we conclude that the concept of a GPT is helpful for growth theory, and that many of the serious criticisms directed at it can be resolved.

I. CHARACTERISTICS AND DEFINITIONS

We first need to consider the general concept of technology before passing on to consider the specific technologies that are GPTs.

Technology

Technology is not always clearly defined in the literature, it is variously used to refer to: (i) the process of applying knowledge to practical uses; (ii) specific machines or processes themselves; and, (iii) the knowledge of how to make things of economic value. Lipsey, Carlaw and Bekar (2005: 58) define the concept as follows: "Technological knowledge, technology for short, is the set of ideas specifying all activities that create economic value." They argue that this knowledge is embodied in capital goods (e.g., machines, tools, structures, etc.), human capital, organisational forms (e.g., the limited liability corporation, the arrangement of machines on the factory floor, etc.) and institutions (e.g., rules, procedures, etc.). The definition separates

technology from pure science on the one hand, and its embodiment in capital goods and the economic structure on the other. While we use the standard shorthand of referring to artefacts as technologies, we strictly refer to the knowledge of how to make and use these artefacts. For example, when we refer to the technology of bronze we mean it as shorthand for the knowledge of how to produce and utilise bronze.

Six Key Characteristics

Our study of GPTs leads us to identify six characteristics that serve to identify a technology as a GPT. These characteristics are closely related to those of Bresnahan and Trajtenberg (1995), although these authors put more emphasis on productivity changes and we do not. We avoid this emphasis because we wish to make productivity effects possible consequences of GPTs rather than identifying characteristics. Also while we accept their point that GPTs often supply some single generic function, we do not make this one of our defining characteristics for two reasons. First, because we think that many non-GPTs also do this and, second, because we wish to avoid fruitless arguments as to just what the generic function is in the case of each GPT.⁴

The first three characteristics in the following list concern categories of technological complementarities which are defined by Carlaw and Lipsey (2002:1310) as, the actions of the initiating agents with respect to their own technologies that affect the value of other agents' technologies and/or their opportunities for further technological advances. The last three characteristics concern the scope of those complementarities and their inherently evolutionary nature.

1. *Complementarities with a cluster of technologies that define and support it.* The evolution of GPTs exhibit technological complementarities *within the evolving cluster of technologies* that define and support them. These complementarities are multi-directional among the elements of that cluster (For example, computers are used in the development and production of new chips, which are then used to increase the range of use and productivity of other digital devices, including computers.)
2. *Complementarities with a cluster of technologies that it enables.* GPTs exhibit technological complementarities *beyond the evolving cluster of technologies* that define and support them by enabling myriad new downstream inventions and innovations, many of which were technically impossible, or economically unfeasible, without the GPT and which are not identifiable as part of the GPT. These in turn enable further inventions and innovations as well as influencing the evolution of the GPT. (For example, electricity helped to enable the electronic computer, which enabled the Internet.)

⁴ For example Bresnahan and Trajtenberg (1995) suggest that the generic function provided by the steam engine was rotary motion. This view has been sharply criticised by historians, and rightly so in our view. The steam engine initially produced reciprocal motion, but knowledge concerning the translation of reciprocal motion into rotary motion and vice versa had been around for centuries. Further, we might debate which function provided by the engine—rotary motion or non-location-specific power—was the key generic function.

3. Complementarities with a cluster of technologies that typically include those that are socially, politically and economically transformative. As the uses of a GPT expand in range and variety of their applications, the technologies involved have significant impacts on the economic structure and very often on the political and social structures as well, requiring changes in, or opening opportunities for, innovations in the technologies that are or will become embodied in these structures. (For example, as well as their major impacts on much of the economic structure, modern ICTs enabled the social media that are altering the techniques and content of communication in politics and society).⁵
4. *There are no close substitutes.* GPTs create complementarities with many or most of their applications that are: (i) unique — no other combination of technologies can produce the application — and (ii) Leontief — without it the whole system would not work. (For example, there is no alternative to electricity for most of the contemporary technologies in which it is the power component and without it most of these technologies could not function or even exist.)
5. *Have a wide array of applications.* The GPT's wide array of complementarities (points 1 – 3) result either in GPTs having multiple uses (e.g. electricity), or a single generic use that itself has many economic applications across much or all of the economy (e.g., railways).
6. *Initially crude but evolving in complexity.* Technologies that evolve to become GPTs typically begin in a crude and incomplete form, often with a narrow range of uses. This makes the identification of any one GPTs depend on a judgment as to when the technology has developed sufficient technological complementarities that it now has all of the other five characteristics in this list. We can know when a technology is clearly not a GPT and when it clearly is one, but the change from one to the other is a continuous process with no clear point in time. (For example, the electronic computer was clearly not a GPT in 1950 and clearly was one in 1990, but one can argue at what point in time the change occurred and what was *the* critical invention.)

Characteristics 2 and 3 are clearly related. Characteristic 2 is the more immediate and often more mundane of the complementarities, while 3 is the more downstream and often more dramatic. Most GPTs will have characteristics that vary more or less continually from one extreme to the other. Because this may cause some uncertainty as to whether some borderline characteristic should be classified under 2 or 3, one might wonder why we separate these two classes. The reason is that some non-GPTs will have many entries under class 2 but none under 3, so what matters in identifying a technology as a GPT is that there are items that are clearly in c-2 and others that are clearly in c-3.

⁵ Characteristic 3 preserves an important connection to work done on a class of enabling technologies before the formal modeling of Bresnahan and Trajtenberg (1995). Freeman and Perez (1988) refer to a constellation of key technologies and their support structure as a Technoeconomic Paradigm. For a full discussion of technological spillovers and related concepts see Carlaw and Lipsey (2002).

Three things determine whether a new technology evolves to possess these characteristics: (i) its internal nature (which helps to determine its potential); (ii) its external interaction with other technologies (which helps to determine the extent and magnitude of its effects); (iii) the efficiency of the technology in delivering its services (which helps to determine whether or not the technology is widely adopted).

On internal nature, many technologies have the potential to evolve into a GPT. Some are easily identified early on as having this potential. For example, it was clear early on that both nano-technology and bio-technology had the potential to become GPTs. (See for example Crandell (1996) for nano-technology and Grace (1997) for bio-technology). In other cases, it was a surprise that a new technology, often invented for specific uses, evolved to become a GPT. This was the case, for example, with computers and the Internet. On interaction, a technology that only interacts weakly with other technologies will not become a GPT. The extent of a technology's interaction depends on its internal nature and on the set of other existing technologies. On efficiency, having potential is not enough. If a technology cannot be made efficient enough to be widely adopted throughout much of the economy, its potential will not be realised. For example, hydrogen fuels cells, have many of the characteristics that provide the potential to become a GPT. However, their efficiency could not (at least to date) be increased sufficiently for them to become pervasive and hence to realise that potential.

Some Criticisms

Two criticisms from economic historians are relevant at this point. First, some historians question whether GPT theorists have added historical depth to their theorizing or simply misread their historical case studies; arguing that the properties listed above may have been derived from flawed histories and so may offer misleading lessons to theorists. A related issue is whether the set of technologies posited as GPTs by theorists constitutes a coherent set. We argue that our GPTs belong to a theoretically, and historically, coherent set because they all share the above six defining characteristics. Other cases of assumed historical errors are considered later in the paper. While we agree that it is important to get the history correct—and we wish to correct it where it has been in error—we believe the above list defines a set of characteristics found in a historically and theoretically important set of technologies.

Second, some historians have questioned whether or not GPTs can be meaningfully distinguished from a broader class of “historically important innovations?” As Field (2011: 218) puts it “There are many single-purpose innovations (for example, the cotton gin) with arguably large implications for TFP [total factor productivity] (and all sorts of other outcomes), and there can be “general-purpose” innovations that represent relatively marginal improvements over previously existing technology and have low social saving/contribution to TFP (for example, the felt-tipped pen).” On Field's first point, nothing in GPT theory precludes the existence of non-GPTs that have some of the same effects as GPTs. On his second point, the felt tipped-pen and other similar single-purpose technologies are ruled out as GPTs because, having many close substitutes, they do not satisfy defining characteristic 4.

We do not doubt that there are some “historically important innovations” that according to our list of characteristics are not GPTs, nor that there are some technologies that share most but not all of these characteristics. But these are the same boundary problems associated with any theoretical abstraction that attempts to categorise reality into discrete and mutually exclusive

boxes. In the present case, if we were to enumerate all of history's technologies, they would be seen to cover a continuum from the very specific to the very general in their direct applications and their technological complementarities. Such ranges of continuous variation are commonly encountered in economics: capital output ratios vary more or less continuously across firms and industries; the number of firms in an industry varies more or less continuously from 'one' to 'very many.' It is, however, demonstrably useful in each of these cases to divide individual items into sets in order to theorise. For example, we talk of capital-intensive industries and contrast their behaviour with labour-intensive industries in spite of the fact that industries with many different ratios exist. Also, we theorise about monopolies, duopolies, oligopolies, monopolistically competitive and perfectly competitive industries, although in fact there is no dividing line that clearly distinguishes the behaviours of firms with n , $n+1$, and $n-1$ competitors, for any but very small values of n . In each case, the characteristics of the typical item will differ markedly from the characteristics of the typical items in other sets, although there will be individual items in each set that are close to boundary elements in adjacent sets.

One example of this useful grouping of many items into a few classes is when one isolates an extreme class from a continuous set. The concept of a GPT is a case in point. We have defined such technologies to lie at one extreme of the characteristics of all technologies. There are technologies that might be just below our threshold criteria in any or all of these dimensions and others that are slightly less than that and so on. As Lipsey, Carlaw and Bekar (2005: 97) put it: "...what distinguishes GPTs from other technologies is a matter of degree. So there will always be technologies that on our definition are almost, but not quite GPTs".

Definitions

While the above six characteristics are the identifying features of GPTs, we need to compress them into a succinct definition for easy reference (always understanding that when in doubt it is the six characteristics that matter). Not surprisingly there has been some debate and confusion about an appropriate definition of a GPT. We do not wish to fall into the essentialist trap of assuming that there is one true definition. We are instead interested in ensuring that the definition of GPTs clearly delineates a class of technologies whose characteristics we have listed above.

There are two separate issues here. First, how inclusive should a definition of GPTs be? Second, should GPTs include single-use technologies that are widely used and share all of the other six defining characteristics?

Considering the inclusiveness of our definition Cantner and Vannuccini (2012) put it this way: there is a "knife-edge distinction...between those scholars who recognize only two or three GPTs since the industrial revolution (the steam engine, electrification and the more questioned ICTs) and see them as singularities or extreme cases of radical innovations ("epochal innovations" as Rosenberg and Trajtenberg (2004) term them), and those who expand the list to a much more wide range of technologies." Clearly Bresnahan and Trajtenberg (1995) are in this latter class when they observe that not just a single GPT but "a handful" populate every historical period.

Of course as economic historians we are free to study any class of technologies that interest us. But as theorists we wish to use in our models a class that captures the features that we think need studying. These considerations leads us to choose to study the broader class of GPTs

for a number of reasons. First, and perhaps most importantly, we believe an important application of GPT theory concerns how multiple overlapping GPTs interact. Too narrow a definition usually renders at most one GPT per ‘economic epoch.’ We believe that the historical evidence shows that interactions among GPTs (and of course among other technologies as well) are important determinants of how they evolve. For example, the three emerging GPTs of biotechnology, nanotechnology and robotisation (all of which are enabled by computers and electricity) are interacting with each other today and will do so increasingly in the near future. Second, and related to the first point, some important implications from GPT theory are derived from the dynamics of how existing GPTs compete and are replaced with newly emerging GPTs. Confining our study to a very small set of technologies forces us to ignore these interactions. Third, we believe that limiting the set of GPTs to only two or three in all of history overly limits the applicability of GPT theory. The concept of GPTs has been usefully employed to study the water wheel, the internal combustion engine, the Internet, and so on. It is not clear why one should wish to rule out such potential applications *ex ante*. So we include a wider class of multiple-use GPTs, as long as they display our six defining characteristics.

On the second issue, although most definitions stress the multiple uses of a GPT, some treatments include technologies that are single use, at least in a broadly defined, generic sense, as long as they are widely used over much of the economy and share the other defining characteristics of GPTs. This extends the definition to include such technologies as the movable-type printing press and transportation technologies such as the three-masted sailing ship, the railroad, the motor vehicle, and the airplane. Both classes are of interest, those that only include multiple use technologies and those that admit single-use technologies but have all the other characteristics. Thus we offer two definitions depending on which group one wishes to study, the difference between the two being only in the italicised portions of the definitions.⁶

GPT-a: A GPT is a single technology, or closely related group of technologies, *that has many uses across most of the economy*, is technologically dynamic in the sense that it evolves in efficiency and range of use in its own right, and is complementary with many downstream sectors where those uses enable a cascade of further inventions and innovations.

GPT-b: A GPT is a single technology, or closely related group of technologies, *that is widely used across most of the economy*, is technologically dynamic in the sense that it evolves in efficiency and range of use in its own right, and complementary with many downstream sectors where those uses enable a cascade of further inventions and innovations.

We believe that in terms of formal modeling and empirical applications, study of the wider class of technologies (GPT-b) is more useful than the study of the narrower class (GPT-a). It seems to us to be most fruitful to study together all of the technologies that share our six characteristics, whether they have multiple uses or a single generic use, as long as the latter are

⁶ For a fuller discussion of the concepts of being ‘widely used’ and “having many uses’ see Lipsey, Carlaw and Bekar (2005: Chapter 4)

widely used through most or all of the economy. So from here on when we speak of GPTs we refer to class-b type.

Another potential source of confusion needs to be addressed. For example, are calculus and Greek mechanics GPTs? They were both certainly important in helping to produce a wide range of valuable technologies. To deal with this issue we need to distinguish GPTs from what Lipsey, Carlaw and Bekar (2005) term General Purpose Principles (GPPs). The difference between GPPs and GPTs is that technologies (GPTs) represent the knowledge underlying some particular value-creating activity while principles (GPPs) are used to develop a range of technologies. For example, the Greek concept of mechanical advantage, underlying the technology of the lever, is a GPP; the knowledge of how to make a lever that exploits mechanical advantage is a technology.

Micro and macro effects

A major difference among definitions that exist in the literature is whether or not they include the effects of GPTs on a range of macro-economic indicators, as well as the broader economic, social and political transforming effects that often accompany their evolution. Neither of our two definitions include these possible macroeconomic effects, nor are they in our list of defining characteristics. However, Jovanovic and Rosseau (2005: 1182) include them stating: “A general purpose technology or GPT is a term coined to describe a new method of producing and inventing that is important enough to have a protracted aggregate impact.”

We argue that such possible macro effects should not be part of the definition because it is important to specify the defining characteristics of a GPTs and then ask: What are the implications of these characteristics? When modeled, do they always produce given changes and if not, under what circumstances do they do so? These are interesting questions but if we only admit to our universe of study those technologies that have these characteristics, we cannot ask such questions. A parallel with growth theory illustrates this crucial issue. Until recently, most aggregate growth models defined technology in such a way that changes in it had to cause contemporaneous changes in GDP and productivity.⁷ The ‘productivity paradox’ occurred because empirical observations showed that these did not always change contemporaneously in the same direction. Such an outcome could not occur in these models, and hence they could not be used to study it.

Lasers provide a clear illustration of why GPTs should not be defined in terms of their macro effects. According to our list of defining characteristics, lasers are a GPT. They are used across the whole economy from astronomy, to retail outlets, to hospitals, and they have made possible many developments that could not have occurred without them. But they have not had the disruptive macro effects found with the introduction of most new GPTs. Why? Our answer is that they fitted well into the current structure of the economy (what Lipsey, Carlaw and Bekar (2005) call the facilitating structure) and so, unlike most other GPTs, they did not require major

⁷ In Hicks-neutral growth models, technology is a constant multiple to the aggregate production function: $GPT = Af(K/L)$, while in Harrod-neutral models technology is defined as a constant multiple attached to the variable input labour: $GDP = f(K, AL)$. In both cases, changes in technology, A , have to coincide with changes in both GDP and TFP.

structural changes and so did not induce ‘protracted aggregate impacts’ in order to become effective.

II. IDENTIFYING A GPT

It can be tempting to think of a GPT as being embodied in a single generic artefact, such as a dynamo, an electronic computer or a biotechnology laboratory. Economic historians have been rightly critical of some authors’ identification of such artefacts—e.g., the dynamo—as GPTs. As Field (2011: 223) notes, “It matters that we get the history right. The critical innovation in making commercial electric power a possibility was not the magneto or the dynamo, which had been under development for decades, but the steam turbine.”⁸ Our response to this criticism starts by conceiving of GPTs as an evolving system of technologies.

GPTs as evolving technologies

While Lipsey, Carlaw and Bekar (2005) are clear that electricity had a long evolutionary history stretching over centuries, they appear to put inordinate emphasis on one invention and one date in that trajectory, the invention of the practical dynamo in 1867 (see pages 197 and 255). However, a critical point that has caused some confusion in the literature is that a GPT cannot be identified by any one artefact that was innovated at one point in time because the knowledge that is the GPT evolves through time as its own efficiency and range of applications change. What was an electronic computer in 1945 was very different from what was an electronic computer in 1950, 1955, 1960, and so on. This was also the case with electricity. We accept that the dynamo, the alternator, the rotary converter, and the transformer, which together allowed for the generation of commercially viable electricity, along with the transmission lines that facilitated its distribution, were some of the main sub-technologies that were important parts of the evolving GPT: the knowledge of how to generate and distribute electricity. No single one of these was *the* critical invention. Lipsey, Carlaw and Bekar (2005: 211) make this point clearly in their discussion of biotechnology:

“The GPP of understanding the genetic code was turned into a GPT by several inventions that made it a practical technology. First, was the discovery of a family of enzymes...that can recognise a particular sequence in DNA and cut it at the required point. Then came the technique that allowed the various fragments of DNA to be separated into homogeneous groups. After that was recombinant DNA in which fragments of DNA are joined....With these inventions, biotechnology became a true process GPT, which like any process, could be used in the manufacturing (or creating) of many different products.”

Because there is rarely if ever a single technological artefact that can be pointed to as an entire GPT, these technologies constitute an evolving knowledge set that is almost always made up of many cooperating artefacts. GPTs are sometimes referred to as a single technology and

⁸ Note, however, that the steam turbine was not a part of the GPT of electricity but an illustration of the importance of complementary technologies in the applications of any specific GPT.

there can be no harm in using this as a shorthand for the evolving process of how to do something, unless it obscures the long evolution and slow accretion of technological knowledge that is the typical story of a GPT. For example, Feldman and Yoon (2011) refer to rDNA as the GPT that we discussed above as biotechnology. Ours is a broader definition, covering the class of discoveries and inventions that make modern biotechnology possible. Theirs is a more specific case of knowledge about one of the key elements in that development. Defining one of these more specific elements as the GPT can give rise to unproductive arguments as to which of the various sub-technologies is the GPT. It can also obscure the fact that a GPT is an evolving body of knowledge. For these reasons, it seems essential to understand that the whole bundle of how to do the job is the GPT: e.g., how to make commercially useful bronze, electricity, or genetically manipulated entities.

An implication of the definition of a GPT as an evolving knowledge set is that there is rarely a definitive unique date for a GPT's emergence, or for its disappearance. Some, such as the three-masted sailing ship, disappear almost completely; some, such as the classic Victorian steam engine, retreat into specialised niches but no longer have the characteristics of a GPT; others, such as electricity, remain as GPTs for centuries, enabling countless new innovations. In theoretical models, in contrast, it is always clear when a new GPT arrives and when it is replaced. In empirical work a starting and ending date must be somewhat arbitrary. For example, in their comparison of electricity and IT, Jovanovic and Rosseau (2005) select the following dates: 1894-1930 for electricity and 1971 until the present for the "IT era." We will have more to say about this choice later on, but the issue should be clear: many of the technological developments since 1930 could not have occurred unless they were enabled by electricity; if electricity is to have an end date as a GPT that is still well in the future.

GPTs as a theoretical class

Field argues that—as a classificatory ordering—GPTs may obscure more than they clarify. Considering steam, electricity and ICT, Field (2011: 220) questions, “whether their histories share enough similarities to justify the distillation of their experiences into a common category with presumably broader applicability.” This raises an important question about abstraction in theory and practice.

Consider another applied abstraction as an illustration. Business school economists spend much time studying the internal working of firms, important in understanding much of their behaviour. Industrial Organization theorists typically suppress much of this detail and group firms into classes such as monopolies, oligopolies, and perfect competitors. Although this abstracts from much internal behaviour, it proves useful in studying some similarities in how firms behave in different market conditions. Micro theorists often abstract even further, typically treating all firms as simple maximising entities. Although this ignores even more, it concentrates on how firms tend to react to changes in such things as input and output prices. Macro theorists abstract even further, treating the production sector as a single unit distinct from the consumption sector. Because each level of theoretical abstraction is appropriate for dealing with a different class of questions, and although the more abstract theories are not helpful in answering finer-grained questions, it is not productive to argue which is the best level of abstraction.

To return to Field's critique, steam, electricity and computers reveal many differences in their individual evolutions and in their influences on the economy. By stressing their common

elements, GPT theory suppresses many of those differences. Our answer to Field's question posed at the outset of this section is that all three technologies share the characteristics laid out in our list of six defining characteristics. The relevant question that remains is: Does the level of abstraction needed to define GPTs help in modelling some aspects of economic growth and technical change? That GPTs prove useful in theoretical models of endogenous growth suggests that our answer to this question is "yes".

Two examples: electricity and the three-masted sailing ship

Not surprisingly with an evolving research program, there has been some confusion on identifying particular GPTs. Indeed, some historians have been critical of the differences among researchers that have resulted. We argue that a technology is identified as a GPT when it can be shown to have the defining characteristics listed in Section I. Here we illustrate this procedure with respect to two technologies, electricity a type-a GPT, that is among the most pervasive of all GPTs, and the three-masted sailing ship, a type-b GPT, that some readers have argued is not a GPT. We provide only enough detail in each of the 6 characteristics to illustrate how to establish that a technology is a GPT while understanding that there is much more that could be said in each case.⁹

Electricity

1. *Complementarities with a cluster of technologies that define and support it.* A number of related technologies show how to generate and distribute electric current efficiently (the knowledge of how to do so constitutes the technology of electricity). The practical generation of relatively low cost electricity provided the incentive for the development of transmission networks and generation facilities. Electricity has co-evolved with many supporting technologies including steam and hydroelectric generation methods. It has also co-evolved with many supporting applications in the pre-existing production structure.
2. *Complementarities with clusters of technologies that it enables.* Many of these downstream technologies were technically impossible without the GPT (e.g., electricity enabled computers). An example includes the replacement of the central drive shaft used in steam powered factories by unit-drive motors attached to each individual machine which allowed for, among many other important things, more efficient arrangement of machines on the factory floor. Those that used the most power no longer needed to be placed close to the drive shaft but could instead be placed in the order in which they were used in the production process. This transforming rearrangement involved (among many other things) the co-evolution of the electric motor, the unit-driven machine tools that it powered and the organizational technology of the flow of production within facilities. Along with the development of machine tools that could cut pre-

⁹ For a more detailed study of the characteristics of each of these GPTs see Lipsey, Carlaw and Bekar (2005), Chapter 6.

hardened steel, these technologies paved the way for Ford's assembly line and the vast productivity gains that came from its slow spread through the US economy over the next decades, and later to much of the rest of the world. Many of the other downstream technologies in turn often enabled further inventions and innovations. For example, a whole range of electrified household products were innovated, some of which were wholly new and some of which were mechanised versions of the hand or foot powered machines that they replaced: washing machines and dryers, dishwashers, refrigerators, electric stoves, vacuum cleaners, toasters, mixers, can openers, blenders, automatic bread makers and so on. Many of these products emerged in crude and incomplete form initially, but, eventually gave rise to wholly new industries to produce and service increasingly sophisticated versions.

3. *Complementarities with a cluster of technologies that typically include those that are socially, politically and economically transformative.* The Fordist production line was first used to mass produce the family car, which in its turn helped to create the suburb, dramatically altering the distribution of populations within many cities, first in North America and eventually worldwide. As an enabling component of the computer, the Internet, and, ultimately social media, electricity has also helped to dramatically alter modern political dynamics and social interactions.
4. *There are no possible close substitutes.* There exists no alternative to electricity for most contemporary technologies in which it is their power component. This is obvious with most of the jobs done by electricity, such as powering computers, lasers, electronic microscopes, radios, TV sets, GIS, GPS and so on *ad infinitum*. More generally, imagine removing electricity altogether, allowing agents to substitute into their next best alternative that either exists or is at least technically feasible with current knowledge. Telephones could be replaced by semaphore and fast inter- and intra-city mail service, but no technology could replace a host of products such as the internet, high speed electronic computers, the cloud, electronic microscopes, and most of the products of bio- and nano-technology. An estimate of the number of downstream technologies eliminated in such an exercise is a proxy for just how 'general purpose' electricity has become in today's world.
5. *Have a wide array of applications.* The myriad uses of electricity in our 'electronic age' are clear. For a slight variation on the argument in 4 above, note what stops working when a city-wide power failure occurs.
6. *Initially crude and evolving in nature.* The development trajectory of electricity and magnetism stretched back at least 200 years prior to its use in the telegraph early in the 19th century followed later by its emergence as a more generally applicable power source. Initially electricity was generated with the magneto and then the dynamo. While electricity solved some immediate problems, such as replacing gas lighting and powering small motors and streetcars, it also opened up

myriad new opportunities across the whole economy such as electrolysis and electroplating. These took decades to develop and exploit and are on-going today.

The Three-Masted (3M) Sailing Ship

1. *Complementarities with a cluster of technologies that define and support it.* The technology of the 3M sailing ship was knowledge concerning a cluster of technologies, including sail, rigging and hull design, which co-evolved along with a number of navigational and organizational technologies, including a general understanding of magnetism, techniques for time keeping, marine insurance and forms of limited liability investment. As the 16th century wore on, 3M ships improved in sea worthiness, navigation aids improved, and European knowledge of the world was filled in as more and more of it was mapped and charted. All of these developments gradually reduced the risks associated with long distance travel by sea.
2. *Complementarities with clusters of technologies that it enables.* There were many downstream inventions and innovations that would not have occurred without global trade on the scale that the 3M ship permitted. Although these were not technically impossible without the ship, they were economically unfeasible. The organization of firms and the mix of the products they produced and sold changed greatly. The development of sea route to the East undermined the established Venetian system of trade that had evolved over centuries in the Mediterranean basin. Shipping along the Atlantic required its own structural adaptations and innovations. Port cities, some of them previously only fishing towns, became the centres of trade, banking and shipbuilding. An elaborate apparatus was developed for production, distribution and sale of basic materials going from the rest of the world to Europe, manufactured goods from Europe to Africa and for slaves going from West Africa to the colonies. Large capital investments were required to provide and constantly improve the needed ships, port facilities, fortified bases, docks, storage warehouses, and a host of other facilities. As the Atlantic trade grew, many institutions were either invented, or greatly improved, as a result of the need to invest large amounts of capital under very risky conditions. Joint stock companies, and marine insurance were two of the most important. Italy's financial institutions arose around the large trading monopolies and the institutions and financial instruments were designed to accommodate these large players. Northern markets tapped new pools of capital by also catering to less wealthy investors. The innovation spilled over to have productivity enhancing effects on many commercial activities not directly related to transoceanic trade.¹⁰

¹⁰ One of the most interesting spillovers “concerned magnetism and the compass...when Columbus crossed the Atlantic, his navigators found their compasses continually changing their angle with true north. From then on, compass variation became a serious navigational problem [which was eventually solved by time keeping, time zones and the invention of longitude]....[I]n 1600, William Gilbert published *De Magnete* and....with his daring hypothesis that the earth was a gigantic magnet, he turned magnetism from a body of empirical observations into a science. [...] Since magnetism is closely related to electricity, Gilbert also took one of the

3. *Complementarities with a cluster of technologies that typically include those that are socially, politically and economically transformative.* The 3M ship enabled for the first time truly global trade, including North and South America, south and central Asia, as well as the pacific islands. Because of that ship, the central location of economic activity shifted to the Atlantic coast, while the Mediterranean region declined in relative terms. Overland trade within Europe, and between it and the East, became a less important and more costly alternative, sealing the fate of the Champagne fairs. More importantly the 3M ship enabled the institution of European colonialism of the Americas, Australia, many Pacific islands, parts of Africa and parts of Asia. These required many physical and institutional innovations that had social and political implications that persist to the present.
4. *There are no close possible substitutes.* Until the development of the iron steam ship in the 19th century, there were no substitutes for the 3M ship in its trans-oceanic uses.
5. *Have a wide array of applications.* Although 3M ships did not have the same vast array of multiple uses as did electricity, they did have more than a single use, as transport for freight and passengers, as fighting machines, as communication vehicles (carrying messages overseas), and as vehicles for exploration and scientific discovery. More importantly, as globalised trade expanded to become the main source of (Smithian) economic growth in the 16th and 17th centuries, their use spread over many sectors of the economies that required sea transport for both inputs and outputs.
6. *Initially crude and evolving in nature.* This technology emerged from long history of ship development stretching back centuries. The 15th Portuguese caravel is often regarded as one of the immediate precursors of the emerging GPT of the three-masted sailing ship (although competing versions emerge in northern Europe and existed in South-East Asia at or near this time). The fundamental breakthroughs in rigging were to combine three known technologies: the square sail, the lateen sail and the triangular foresail, but on a long forespar. The fundamental breakthrough in hull design was a move from overlapping planks to butted planks sealed with caulking. These innovations incrementally improved the efficiency of sailing in rough seas and tacking against the wind and thus vastly extended the possible range and safety of voyages.

The question of whether or not a technology has enough of the defining characteristics to classify it as a GPT can be answered, as in the above cases of electricity and the 3M ship, by an enumeration of observed characteristics. This usually requires extensive historical research,

first steps, and it was an enormous step, towards the invention of the dynamo and practical electrical power. [...] The development of the science of magnetism illustrates the often surprising complementarities between technologies—here the three-masted sailing ship on the one hand and magnetism and electricity on the other.” (Lipsey, Carlaw and Bekar, 2005: 171-2)

research that uses both qualitative and quantitative data. The most direct evidence that a technology may not be a GPT is failure to identify it as having our six defining characteristics. Of course as discussed in detail in the immediately preceding section, there will always be borderline cases for which the enumerations under these 6 items leaves uncertainty as to whether or not to call a specific technology a GPT. But as argued there, this is an inevitable result when the items in any continuously varying series are divided into a few discrete classes in order to theorise about them.

III. EMPIRICAL MEASUREMENTS AND TESTS

Rather than make the historical analyses referred to for electricity and the 3M ship in the previous section, some writers have tried to determine whether or not a particular technology is a GPT through the study of quantitative empirical data. The two candidates for such data that are found in the literature are patent data and macroeconomic data (typically measures of productivity and/or GDP growth). We consider the use of macro data in Section IV. Here we consider the use of patent data.

There are also two ways to interpret these studies: (i) as a test of a prediction assumed to be derived from GPT theory; or (ii) as a test of the existential statement that the technology in question is a GPT. With regard to the first possibility, none of the existing formal theories of GPTs are designed to produce predictions about patents. While this does not preclude the use of patent data to establish specific things about GPTs, it does imply that those wishing to test predictions with such data should seek to derive testable implications from existing theories. However, no one seems to have accomplished this, or even tried to do it. The closest anyone gets to doing this is Moser and Nicholas (2004) in the paper discussed in more detail below when they state: “Bresnahan and Trajtenberg (1995 p. 3) argue that the range of later generations of inventions that benefit from an early patent can be measured as the range of industries that cite the early patent”. However, the only reference to patents in Bresnahan and Trajtenberg (1995) concerns using them to estimate variables in their reaction functions, not as predictions of their theory. So Moser and Nicholas’s, and other similar studies, study must be understood as testing an existential statement that electricity is a GPT rather than testing any specific predictions of the GPT theory. To consider this latter possibility we need first to ask how such existential statements can be tested.

Testing a general existence statement.

The existence statement that technology X is a GPT is neither a theory nor a prediction that is derived from a theory. It is instead a statement that technology X meets the requirements for being a GPT. The writers we are about to consider hold that such statements can be disproved for particular technologies using empirical data. But can such statements be proved or disproved? This question raises a more general methodological issue. Finding an example of X demonstrates the general statement that it exists, but not finding X does not prove it does not exist. In the latter case, one might not have looked hard enough, or in the right places, or any one of a multitude of other reasons. To take Karl Popper’s famous example, the statement “Angels exist” can be proven by finding what we all agree to be an angel, but cannot be falsified by failing to find an

angel. So we can prove a general existence statement; we can fail to prove it with some specific exercise or measurement; but we cannot disprove it.

We can prove that technology *X* is a GPT by showing that it has all of our six defining characteristics. If our historical studies fail to find convincing evidence that it has the necessary defining characteristics, we would have failed to prove that it was a GPT; but we would not have proved that it was not a GPT—further research might, for example reveal characteristics that we missed earlier.

Testing a qualified existence statement

For an existence statement to be capable of being disproved by evidence it must be closely confined by enough conditions as to make it refutable. To rephrase Popper, the statement “an angel exists in visible form and is in this room” can be falsified by not finding it in this room. One case in point has already been mentioned. Field (2011, p. 218) criticized the concept of a GPT for being unable to rule out such trivial technologies as felt-pens or screws, which are widely used and for various purposes. Although the statement that felt-pens or screws are a GPT is an existence statement, it can be refuted because the universe of their use is sufficiently circumscribed. We can enumerate all of their uses and show that these have sufficiently close technological substitutes to disqualify them as being GPTs.

Inappropriate tests of non-existence

As was argued above, to test the existential statement that a specified technology is not a GPT the statement must be sufficiently bounded in time and space so that the indicators of their existence can be exhaustively shown to be non-existent. But the technological characteristics associated with GPTs and their antecedents—technological complementarities and their evolution through time, number of uses, lack of technical substitutes, generality, etc.—do not map well or clearly onto the legal concepts underlying patent law. We note furthermore that other authors have been even more critical of the general value of patent data in economic studies.¹¹

As a case in point, we consider the study by Moser and Nicholas (2004) who analyze patent data in an attempt to determine whether or not electricity—which they argue is held to be among the list of quintessential GPTs—actually meets the criteria for being a GPT. Specifically, they use patent citations in an attempt to determine electricity’s generality, and how important it

¹¹ Von Hippel (1988) notes that it is often better for agents concerned with maximizing the value of rents from innovation to forgo a patent. He notes further that this strategy is often pursued in cases of rapid innovation for technologies with many potential uses. Derochers (1998, p. 72) notes that “Austrian economists have long been hostile to the use of economic statistics...but their epistemological claims have not been echoed in the mainstream of the profession which, for a number of reasons, has always been fond of patent statistics. The purpose of this article was therefore to demonstrate that, even on strictly empirical grounds, the drawbacks associated with patent statistics are nothing short of major and that even in the best circumstances, these data only give us a partial picture of the technological potential of a small number of innovations. Doing empirical studies of innovation is a worthy goal, but this does not justify the use of bad indicators on the grounds that they are the only ones that fit well with the dominant methodologies.”

was in downstream innovations. They argue that while electricity was an extremely novel innovation (patents on related innovations rarely cited earlier patents), electricity was no more general than a number of other innovations because their data show relatively few complementary innovations, “Our results contradict the hypothesis that electricity was a GPT according to conventional definitions such as those of Bresnahan and Trajtenberg (1995) or Lipsey, Bekar and Carlaw (1998). We find that ...electricity patents had lower generality scores, as measured by the distribution of their forward citations, fewer citations per patent (a measure of technological importance), and shorter citation lags (i.e., faster rates of knowledge depreciation)” (Moser and Nicholas 2004 388-89).

There are many reasons why their arguments do not sustain their strong conclusions.

1. Their measure of generality is a Herfindahl-Hirschman index that measures the range, or extent, of a patent’s generality. They argue that the index shows that electricity was not general in its application in the period 1976 - 2002. What the index actually shows is that those few patents that they observe in the USPTO industrial classification of electricity are no more general than the patents they observe in other classifications such as chemicals. Yet if one looks at the industrial and household activities of any developed economy in the period 1976-2002 one would be hard pressed to identify many activities that did not employ electricity somewhere in their production technologies.
2. They argue that patents for chemical technologies were found to be more ‘general’ than electrical technologies. However, in many cases, the chemical innovation that were cited relied on electricity as a cooperating or enabling technology but that relationship was a technological and not a legal one. They write (392) “The most striking fact about our list of the most general inventions is the relative paucity of electricity patents. Even Fuller’s invention of an electric contact is distinguished by the chemical discovery that silver-copper alloys are less sensitive to damage caused by heat and arcing.” But of course the electric contact is part of the cluster of electricity related technologies making up the GPT. That it was a chemical based innovation is immaterial. Similarly many of the other patents for innovations that were wholly reliant on electricity either did not cite earlier patents, or the earlier electrical innovations were themselves not patented.
3. Most of the vast range of products that were enabled by electricity from the array of household machines to radio, TV, electronic computers, the Internet, electronic microscopes, and GPS when they were patented did not typically mention electricity as a necessary condition for their existence. Also as any GPT becomes increasingly ubiquitous, as is electricity today, it tends to be cited relatively less than newer and more novel technologies with a more limited range of applications.
4. Electricity’s technological complementarities cover a much bigger set of electricity-dependent innovations than can be captured by forward patent citations. To illustrate, consider our case study of electricity examined earlier. The redesign of modern factories around the central drive shaft with unit drive motors was technologically and economically transformative but many of its elements were not patented.
5. Their measure of originality covers only a small part of electricity’s chronology (1920 to 1976) and range of effects (four industrial categories are surveyed: electricity, chemicals, mechanical, other). They exclude some major applications of electricity, computers and communications, due to lack of data.

We conclude that although their work provides some interesting and informative data, it cannot sustain their strong conclusions. The absence of relevant patent data cannot refute the proposition that any one technology is a GPT. Furthermore, the range of their study covers far less in time and space than the universe over which electricity was (and continues to be) a major enabler of inventions and innovations. Thus, it does not come near to the exhaustive study that would be needed to refute the existential statement that electricity is a GPT.¹²

So, while patent data may be useful for many purposes including estimating the parameters associated with a particular GPT, they cannot be used to establish a technology is not a GPT.

Patent data corroboration

Patents have been used in the literature to establish many things about technologies, which can include providing a minimum estimate of the existence of some of the defining characteristics of a GPT, such as the downstream technologies that it has enabled. However, for reasons already discussed these must be understood to be minimum estimates not exhaustive listings.

Here are some examples of the productive use of patent data on technologies. In a study that predates the introduction of the GPT concept, Trajtenberg, Henderson and Jaffe (1992) use

...patent citations to measure the "basicness" and appropriability of inventions. We propose that [1] the basicness of research underlying an invention can be characterized by the nature of the previous patents cited by an invention; [2] that the basicness of research outcomes relates to the subsequent patents that cite an invention; and [3] that the fraction of citing patents that are assigned to the same organization as the original invention is a measure of appropriability.

Other studies, for example, Feldman and Yoon (2011), and Stuart and Iacopetta (2014), use patent data in an attempt to establish that the technologies of rDNA and nanotechnology are GPTs. Both of these studies try to measure characteristics of GPT, in particular they construct hypotheses that use patent data to measure the pervasiveness and the complementarities of the technologies they consider. They present examples of empirical studies that find evidence

¹² One might ask; if not patent data, then what data would provide a test that a specific technology was a GPT? Our answer is: the data concerning our six characteristics. And we have demonstrated how one might do this using qualitative data in our examples of electricity and three-masted sailing ships. However, if the desire is to find quantitative data to perform such an empirical investigation, then such data has yet to be collected, as far as we know. If such qualitative and/or quantitative data cannot be found for all six, then the proposition that the technology in questions a GPT becomes dubious. But all that can be concluded is that the technology is question cannot currently be shown to be GPT, while for many technologies the possibility of finding new data that changes this conclusion cannot be ruled out definitively. (Of course, there will always be borderline cases of technologies that have some but not of all of our six characteristics, or that only do so in a small way for at least some of them.)

showing that certain technologies have some of the characteristics of GPTs. But what they could not do (nor do they try to do) is to prove that these technologies were not GPTs.

IV. THEORETICAL MODELLING OF GPTS

Cantner and Vannuccini (2012) provide an excellent summary of the development of GPT theories from Bresnahan and Trajtenberg (1995) to Carlaw and Lipsey (2011). So here we only mention a few salient features relevant to our present purposes.

When looking at such theories we may ask what we would like to see in them. Among other things, a desirable list of features includes:

- (1) assumptions that, as far as possible, capture the empirically observed characteristics of GPTs;
- (2) although abstraction is necessary in any theory, *ceteris paribus*, the more relevant characteristics of the phenomena that are modelled in the assumptions the better;
- (3) the minimum of counter-factual assumptions and those that are used are robust in the technical sense that they can be removed without altering any of the theory's critical conclusions;
- (4) in particular, the absence of ad hoc, counterfactual assumptions that are employed merely to produce some desired result (e.g., one of the theory's key predictions);
- (5) the minimum of predictions (preferably none) that are clearly refuted by existing empirical evidence.

One of the main contributions of Bresnahan and Trajtenberg (1995) was to break the link between changes in technology and productivity that is inherent in models that use aggregate production functions and measure productivity as a residual (i.e., TFP). The way was then opened for theories that could explain the productivity paradox of changes in technology that were not systematically related to changes in productivity.

First-generation models

The first post-Bresnahan-and-Trajtenberg-generation theorists sought to explain an apparent association between the introduction of major new technologies and slowdowns and surges in GDP and/or productivity growth. All of their theories have assumptions not found in Bresnahan and Trajtenberg (1995) or in the literature on transforming technologies. Some of these assumptions merely simplify the authors' analysis. However, some others that are clearly counterfactual drive the model's results. This violates the important methodological rule that assumptions that are patently counterfactual should be robust in the sense defined above. The most important assumptions follow (with our comments in *Italic*).

1. GPTs arrive exogenously and there exists only one GPT in use at any one time so that the behaviour of the macro economy closely mirrors the behaviour of the GPT. *This implicitly uses the one-at-a-time definition of a GPT that we discussed in Section I and which we gave reasons for not using in our study. For one reason, there are usually several GPTs in active use at any one time.*

2. All agents are able to identify a new technology as a GPT when it first arrives and to predict its value over its entire lifetime. *This is necessary in order to apply conventional maximising theory to the models. In contrast, actual experience shows that it is not always clear if a new technology will develop into a GPT when it is first innovated and this uncertainty often extends over a long time span. Also, it is clearly impossible to predict the full future course of any GPT once it has been identified as such.*
3. It follows from assumption 2 that in these theories there can be no uncertainty in decisions regarding those GPTs. *In contrast, empirical evidence shows that uncertainties (in Knight's (1921) sense) are ubiquitous in the development of almost all new technologies.*
4. There is a single vertical complementarity between the GPT and its components and all of these components are horizontal substitutes for each other. *In fact downstream technologies often complement the development of the original GPT, and each other, as did the assembly line, just-in-time inventories and robot-operated assembly lines, while others are more or less independent of each other as were the vacuum cleaner and the electric washing machine—all of which depend on electricity.*

All of these first-generation models that address the issue predict a slowdown followed by a surge in a least one of the macro variables following on the arrival of a new GPT, immediately in Helpman and Trajtenberg (1998a) and with a lag in Helpman and Trajtenberg (1998b) and in Aghion and Howitt (1998). As we have noted, however, the original model in Bresnahan and Trajtenberg, (1995) was designed to produce such slowdowns. Although the heuristic value of these models is not at issue here, it is clear that these models fall short of capturing our six defining characteristics. However, such limitations are typical of the early stages of any research program for the development of some important new theory so that pointing them out "...should be understood as markers for further research rather than arguments for ignoring these first pioneering attempts at modelling the impact of GPTs" (Lipsey, Carlaw and Bekar, 2005: 384).

Later-Generation Models

In the first of the second-generation models, van Zon, Fortune and Kronenberg (2003) make two important improvements. First, they allow more than one GPT to exist at any one time and, second, they allow a potential GPT to fail if enough supporting components are not developed for it. While this model is a significant improvement over all previous models, it could be improved yet further by altering some of its assumptions that are not consistent with what we know about GPTs.

First, in their model all GPTs cooperate with each other in jointly producing the composite final good; they never compete with each other as electricity did with steam and steam with water wheels. Second, and counterfactually, each GPT has a smaller impact on output and growth than does each previous GPT. Third, in the long run GPTs do not sustain growth, which goes asymptotically to zero.

Most of the other second-generation models are concerned with the transitional dynamics that follow the exogenous introduction of a new GPT. Pestas (2003) alone makes the unequivocal prediction that the pattern of a slowdown and a surge in the transition to a new steady state will follow the introduction of all new GPTs. Both Harada (2009) and Rainer and Strohmaier (2014) make the pattern of a slowdown and/or surge conditional on the exogenously determined relative attractiveness of R&D devoted to GPTs and ancillary technologies.

All of the above models are technically complex in ways that make it extremely difficult to alter their assumptions to increase their empirical relevance. They all use dynamically stationary equilibrium concepts with agents that maximize over the whole life of the GPTs that are developed under conditions of either certainty or risk. This maximising condition requires the counterfactual assumption that agents can foresee the actual or probable consequences of a newly introduced GPT over a lifetime that often covers centuries and clearly involves many genuine uncertainties

Non-stationary models

The theories in Lipsey, Carlaw and Bekar (2005, Chapters 13-15) and in Carlaw and Lipsey (2006: 2011) differ in that they are not constructed to produce any set of macro effects; rather, they are constructed to capture many of the micro characteristics of GPTs, and then see what these imply with respect to growth slowdowns, surges and other aspects of GPT-induced behaviour.

As an example consider the model in Carlaw and Lipsey (2011). It has three sectors: a pure research sector that develops and produces GPTs, an applied research sector that develops applications of these GPTs that are suitable for use in production, and a goods-producing sector that uses the technologies developed by the applied research sector. The model's many assumptions concerning GPTs are all based on known facts related to technologies and technological change. Some of these are common to most technologies, some just to most major technologies, and a few are unique to GPTs. They are summarised in Table 1. To compare this model with those discussed in the previous section, the columns of the table show those assumptions that are found in all other GPT models, in some but not all other models, and in none of the others.

Table 1: Characteristics of GPTs in Carlaw and Lipsey (2011) compared with those in previous models			
Property of the GPTs in the Carlaw and Lipsey (2011) Model	Found in all other models	Found in some other models	Found in no other models
GPT is developed through endogenous R&D		X	
GPT efficiency increases over time	X		
GPT diffuses slowly through the economy			X
Several distinct GPTs exist simultaneously		X	
GPTs are grouped by class			X
GPTs may overlap within and between classes			X
Technologies in different classes may be complements			X
Uncertainty regarding the knowledge to be produced by a given amount of R&D		X	
Uncertainty regarding timing of new GPTs		X	
Uncertainty regarding new GPT's productivity			X
Uncertainty regarding GPTs relationship to existing GPTs			X
Uncertainty regarding if and when an incumbent GPT will be replaced by a new GPT within its class			X
Agents unable to maximize over life of GPT			X
No unique evolutionary path for each GPT because of random components on the outcomes of decisions			X

To incorporate all of these assumptions, the model cannot employ a stationary dynamic equilibrium concept. Instead, an unorthodox model is required — one that evolves in ways that are not wholly mechanistic and hence not wholly predictable. There is a transitional equilibrium in every time period, given the current expected marginal productivities of inputs in each sector. But because of technological advance and the fact that agents cannot foresee the consequence of the complementarities arising from their present resource allocation decisions, the marginal products of resources change in each line of activity from one period to the next. Thus the equilibrium outcome is different in each period and the economy never settles into a stationary equilibrium (neither a steady state nor a balanced growth path). This focuses the model on the historical, path-dependent process of knowledge accumulation and uneven patterns of growth driven by variable rates of innovation.

Fortunately, these alterations make the model technically much simpler than the earlier models of GPT-driven, stationary equilibrium growth. As a result, the model can be extended to include an increasing number of empirically relevant characteristics in its assumptions. Importantly, this theory is not constructed to produce preconceived results such as ‘protracted aggregate impacts’.

Predictions

A large number of simulation runs of the Carlaw and Lipsey (2011) model produce some interesting behaviours that can be regarded as the model's predictions.

1. There is no consistent correlation between the arrival of any one GPT and the slowdown or acceleration of either GDP or productivity. These may occur with the introduction of a new GPT, but only in some circumstances. First, the arrival of more than one new GPT at more or less the same time can combine to exert a measurable effect on GDP. Second, more than one GPT can be coevolving through similar stages of development at roughly the same time and so reinforce each other in their effects on the economy. Third, one GPT may be much more 'powerful' than the typical one and so have a noticeable effect on the economy by itself (as was probably the case with electricity). But more often than not, the effects of GPTs at various stages in their evolution cancel each other out in the macro data. This agrees with many researchers have found empirically and what was argued in Lipsey Bekar and Carlaw (1998: 207).
2. In the model's many simulations, changes in measured TFP do not correlate with changes in technology. This is interesting because of the long-standing debate as to whether or not TFP is a valid measure of technological change (as it is in Solow-type growth models).¹³ This also conforms to what we observe, as for example when, in the late 1980s, there were large changes in ICTs but little changes in the measured growth rates of either productivity or GDP.
3. The model produces the result that GPTs rejuvenate and sustain the growth process by creating an array of opportunities to develop new innovations applicable in the production of final goods. If no further GPTs are allowed to arrive, existing GPTs continue to improve in efficiency and the applied R&D sector continues to innovate. But both of these activities eventually encounter diminishing returns, causing the growth process to slow and eventually to stop. If new GPTs are again allowed to arrive, the endogenous growth process is rejuvenated and growth is sustained into the indefinite future.
4. The model's GDP-driven growth is sustained without any activity exhibiting increasing returns to the accumulating factors as is required in some endogenous growth models in the tradition of Romer (1986, 1990) and Lucas (1988)—returns that researchers seeking to measure R&D spillovers, were unable to agree that they had detected in the macro data that were studied after these theories were presented.

Another important historical observation is that GPTs transform the structures of the economy and society. Although this result cannot yet be derived from a formal model, there is ample historical evidence supporting it. Writing allowed the first development of complex societies and cities in Mesopotamia. Bronze allowed the development of organised warfare, the development of multi-city empires, and the transfer of political authority from temple to crown. Movable-type printing contributed to a knowledge society unimaginable when all manuscripts had to be laboriously copied by hand. The steam engine, combined with automated textile machinery, led to the urbanisation first of England and then of many continental countries as

¹³ See Carlaw and Lipsey (2002) and Lipsey and Carlaw (2004) and for discussions of this point.

factories moved from dispersed sites with fast running water to the new vast Victorian industrial cities. As well as permitting countless new products and processes, the computer and related technologies, such as the cell phone and social media, has transformed the economic, social and political structures in profound ways. And so on through every major GPT and its set of derivative technologies.

A way ahead

One characteristic that is important but not included in any of the models mentioned so far is that a technology may emerge in a narrow use and, through interaction with other technologies and research processes, evolve into a full GPT, as for example did the steam engine, the electronic computer and the internet. The beginnings of an approach to modelling this characteristic are found in van Zon, Fortune and Kronenberg (2003), Cantner and Vannuccini (2012 and 2016), and Korzinovy and Savin (2016).

A key to modelling this characteristic is that GPTs should not be *ex ante* identifiable. In all of the above-mentioned models (and in Carlaw and Lipsey (2011)) a candidate technology may fail to become a GPT.¹⁴ To allow for this possibility, the four theories model GPTs as an endogenous emergent property of an evolutionary interaction between a sector that produces knowledge for the creation of GPTs and an applied sector that produces complementary technologies useful in the goods-producing sector. In this modelling, some part of the knowledge succeeds in creating actual GPTs while other parts fail to do so.

Along with the possibility of modeling the interaction of a new GPT with the economy's existing structure, which may need substantial alteration to accommodate the new GPT (as it did with steam engines and the computer), this incorporation of the GPT as an emerging property, and the feedback from the applied sector to the pure research sector, are all on the table for further important theoretical developments.

Slowdowns and Surges

One serious problem that contributed to the origins of GPT theory was concern over the 'productivity paradox' of the 1980s. Paul David's 1990 article provided detailed evidence as to why the introduction of electricity caused lagged responses—both slowdowns and accelerations—in the growth of productivity. This issue was then taken up by the post Bresnahan and Trajtenberg (1995) first-generation GPT theories. These were constructed to generate aggregate slowdowns followed by surges after the introduction of a new GPT. David's work on electricity in combination with these theoretical models have been so influential that many subsequent authors of empirical studies have assumed that all GPT theories predicted slowdowns and/or surges. For example, after following Lipsey, Carlaw and Bekar (2005) in dismissing slowdowns as a necessary characteristic of all GPTs, Ristuccia and Solomou (2014: 229) assert

¹⁴ Cantner and Vannuccini (2012) argue that Carlaw and Lipsey (2011) revert to a stochastic GPT arrival process that makes the GPT *ex ante* identifiable. But, in fact, while the knowledge for potential GPTs does stochastically arrive in the Carlaw and Lipsey (2011) model, that potential GPT may not become an actual GPT if it fails an expected efficiency criteria test and thus is not eventually adopted any of the model's applied sectors.

that: “The only thing that seems common to all GPT models ...is that the introduction of a new GPT should eventually determine a productivity surge”.

As already observed, this and other similar empirical work can be seen either as testing an assumed prediction of all GPT theories or as testing whether or not a particular technology was a GPT by assuming that this macro behaviour was part of the definition of GPTs. We have already examined in the case of electricity why the existence statement that a particular technology was GPT cannot be refuted by empirical data so we concentrate here on the interpretation that the empirical work is testing an assumed macro prediction of GPT theory.

There are at least three serious problems with equating the arrival of GPTs with growth slowdowns and surges. First, different first generation theories predict different types of slowdowns. For example, Helpman and Trajtenberg (1998a) predict cycles of slowdowns and bonuses in both TFP and GDP, while Helpman and Trajtenberg (1998b) and Aghion and Howitt (1998) predict this pattern for real wages and real output.

Second, as already argued, later theories that were not designed to produce slowdowns and surges of any sort do not predict their inevitability. Indeed, after an analysis of this issue, Carlaw and Lipsey (2002: 1314-15) conclude: “...there needs to be no observable impact of new technologies on rates of return...”. Also Lipsey, Bekar and Carlaw (1998: 207) argue “The main problem with any hypothesis concerning transitional effects is to establish a link between the evolutionary paths of individual GPTs and the observed macro behaviour of the economy....It is quite possible for all individual technologies in the economy to follow a logistic pattern of productivity growth, while macro growth follows random variations around a more or less stable trend.” After the long analysis that follows this statement, they conclude (212): “We argue only that GPTs do not always cause slowdowns. Further, when a GPT is the source of an observable macro slowdown, we would expect to find certain characteristic patterns within the development trajectory of the GPT itself, and interactions with other technologies, and the existing facilitating structure and policy institutions”. This is also what the Carlaw and Lipsey (2011) model predicts: slowdowns and surges associated with the introduction of a new GPT under some identifiable conditions but not others.

Third, empirical evidence shows that some GPTs have been associated with various types of slowdowns while others have not—an ambiguous result that is again predicted by the Carlaw and Lipsey (2011) model. For example, railways, the iron steamship, the laser, and biotechnology were not clearly associated with slowdowns followed by surges. In contrast, Jovanovich and Rouseau (2005) argue that electricity and Information Technology did exhibit this pattern, although the period they chose for electricity, 1894-1930, is clearly arbitrary and omits later times when electricity was being used for many new products and processes. In contrast, electricity is argued not to have a productivity bonus by Ristuccia and Solomou (2014). Castellacci (2010) finds that industries taken to be more closely related to ICT production and use experienced a productivity bonus with the introduction of this GPT. Those industries associated with the old (Fordist mass production) GPT or further away from the production and use of the new GPT experienced productivity slowdowns. This is interesting although it is not quite the slowdowns followed by bonuses predicted for the macro economy by the first-generation theories. Finally, Rainer and Strohmaier (2014: 427) using Danish industrial level data spanning the years 1966 to 2007 find that ICTs had what they call the empirical

characteristics of a GPT in that among other things “...persistent improvements in the technology have led to economy-wide productivity gains”.

Our criticism of the use of productivity data here is its use both to test GPT theory in general and whether a particular technology is a GPT. For example, Ristuccia and Solomou (2014: 227) argue that their finding of only modest productivity gains associated with electricity “challenges the usefulness of general purpose technology (GPT) theory as a way of conceptualizing the relationship between technological advances and long-term economic growth.” They conclude (244) by arguing that GPT theory either: (i) theorizes about a class of technologies that do not exist; and/or, (ii) is a set of truisms about the impact of big innovations on growth in GDP and productivity trends that have been long known and understood. The many points made in the above text provide our rebuttal to these contentions.

Nothing said here is an argument against determining whether or not the introduction of a particular technology, seen to be a GPT by its conformity with our six defining characteristics, was followed by the macro pattern of a slowdown and/or a surge—or indeed by such a pattern in one sector of the economy. The circumstances in which this will and will not happen are of interest.

Rejuvenation of Growth Potential Without the Necessity of Slowdowns or Surges¹⁵

Regarding our list of predictions above, 1, 2 and 4 agree closely with the evidence, but one might wonder how 1 and 3 can both be correct. How can a single new GPT rejuvenate the growth process without sooner or later causing a productivity surge?

Consider the stylised example illustrated in Figure 1. This is not meant to be a realistic representation of the whole growth process but is used only to illustrate the difference between rejuvenating growth and producing a productivity bonus. The figure relates the capital stock (K) to the marginal productivity of capital (MPK). Each curve, MP_1 , MP_2 and MP_3 , refers to a given state of technology. As the Classical economists long ago observed, since each curve is assumed to be negatively sloped by virtue of the law of diminishing returns, capital accumulation with no technical change will cause the marginal product to fall and eventually approach zero (the onset of the stationary state). In contrast, technological change without capital accumulation raises the marginal product of capital along a vertical line, such as the dashed line starting at k_1 . The combination of technological change and capital accumulation—and they do typically go together—moves the economy along some non-vertical trajectory. If capital accumulation is slow relative to technological change, the trajectory of productivity rises along the path such as P_1 ; if it is high relative to technological change, the path will be a falling one such as in P_3 ; while if it is just the right combination of technical change and capital accumulation, the path is the horizontal one, P_2 . So all that we can say in general about the effect of major new GPTs is that they rejuvenate the growth process by making productivity higher than it would have been if there had been the same amount of capital accumulation but little or no technological change. Nothing can be said about how the growth rates of productivity and GDP will change when a


¹⁵ This section relies heavily on Carlaw and Lipsey (2002)

new GPT and its derivative technologies first begin to diffuse throughout the economy, shifting the MPK curve to the right while the capital stock increases.

Care must be taken in using this analysis further as it is fully aggregated while the effects of an individual technology, GPT or otherwise, depends on complex interactions with the other technologies that are in use and the structure of the economy.

CAPITAL ACCUMULATING AND TECHNOLOGICAL CHANGE

Figure 1



V. CONCLUSION

The class of technologies that we call GPTs is important because GPTs spawn a host of new technologies and improve the efficiency of existing ones. They do this both by lowering the cost of broadly defined products or services, such as power or transportation, and by making possible a range of new products and processes many of which were technically impossible without it. This is one sense in which they rejuvenate the growth process. They may or may not cause a surge in the growth of productivity or GDP. But they do make future productivity and GDP higher than they would have been in their absence and so avoid the onset of the Classical stationary state caused by diminishing returns to capital. (They also have other more general transformative effects but these are less important, at least for existing growth theory.)

GPTs are not the only cause of such effects. Other important technologies or groups of technologies can have similar effects. This is not a problem for GPT theory. It only holds that there is an identifiable group of technologies that share common characteristics but not that these characteristics are each unique to GPTs.

Along with most economic historians, some economists, such as Nathan Rosenberg, Richard Nelson and Sidney Winter, have argued that a nation's technology is a complex, hierarchical and constantly evolving system. Yet the majority of models of economic growth continue to contain only one sector described by an aggregate production function and a scalar

value of technology, the latter determined either endogenously or exogenously. Following in the footsteps of Joseph Schumpeter, Nelson and Winter (1982) began the modern study of the economy as an evolving system while Bresnahan and Trajtenberg (1995) began the modeling of technology as having more characteristics than can be captured by a single number. This paper has followed in this tradition, arguing for the usefulness of GPTs as the central organizing framework for analyzing technological change and long-run economic growth. We adopt this concept over other possible frameworks, including Mokyr's macro inventions (Mokyr 1990) and Freeman and Perez's techno-economic paradigm (Freeman and Perez 1988). Briefly the techno-economic paradigm, although useful in historical studies (see e.g., Freeman and Louça 2001), includes far too much to be useful as a foundation for analytical models of technology-driven growth. Also Mokyr's class of macro inventions is too broad for our purposes, including as it does such things as spectacles, the weight-driven mechanical clock, the casting of iron, coke as a power source, the hot air balloon, the bicycle, and the screw propeller.

The concept of GPTs may or may not be useful to economic historians. However, if economists wish to continue with the research program of more sophisticated modeling of technology and technical change while also following the advice of some critics by abandoning the concept of GPTs, the obvious question is: What better theoretical concept of structured technology, capable of being used in formal models, would be put in its place and how different would it look? We know of none on the horizon. So we advocate proceeding on several fronts: learning from criticisms of GPT theory whether intended by their authors as constructive or destructive, making GPTs and related concepts more precise, modeling more of their known characteristics than are captured in present models, and testing the predictions of these models against data.

REFERENCES

- Aghion, P. and Howitt, P. 1998. "On the Macro Effects of Major Technological Change." in Helpman (ed), *General Purpose Technologies and Economic Growth*. Cambridge: Massachusetts Institute of Technology Press.
- Bresnahan, T. and M. Trajtenberg. 1992. "General Purpose Technologies 'Engines of Growth'?" *NBER Working paper* No. 4148.
- _____. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65: 83-108.
- Carlaw, Kenneth I. and Richard G. Lipsey. 2002. "Externalities, Technological Complementarities and Sustained Economic Growth." *Research Policy* 31: 1305-1315.
- _____. 2006. "GPT-Driven Endogenous Growth." *Economic Journal* 116: 155-74.
- _____. 2011. "Sustained Endogenous Growth Driven by Structured and Evolving General Purpose Technologies." *Journal of Evolutionary Economics* 21(4): 563-593.
- Cantner, U. and S. Vannuccini. 2012. "A New View of General Purpose Technologies" in U. Heilemann U and A. Wagner. *Empirische Makroökonomik und mehr*. Stugart: Lucius et Lucius.
- _____. 2016. "Competition for the (Downstream) Market: Modeling Acquired Purposes." *Paper presented to the International Schumpeter Society Conference* July 2016. Montreal: Canada.
- Castellacci, F. 2010. "Structural Change and the Growth of Industrial Sectors: Empirical Test of a GPT Model." *Review of Income and Wealth* 56(3): 449-482.
- Crandall, B.C. 1996. *Nano Technology Molecular Speculation on Global Abundance*. Cambridge Mass.:MT press.
- David, P. A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *The American Economic Review* 80(2): 355-361.
- Desrochers, Pierre. 1998. "On the Abuse of Patents and Economic Indicators." *Quarterly Journal of Austrian Economics* 1(4): 51-74.

- Feldman, M.P. and J.W. Yoon. 2011. "An Empirical Test for General Purpose Technology: An examination of the Cohen-Boyer's rDNA technology." *Industrial and Corporate Change* 21(2): 249-275.
- Field, A. 2011. *A Great Leap Forward: 1930s Depression and U.S. Economic Growth*. New Haven: Yale University Press.
- Freeman, Christopher and Carlota Perez. 1988. "Structural Crises of Adjustment: Business Cycles and Investment Behaviour.", in Dosi, Freeman, Nelson, Silverberg and Soete (eds). *Technical Change and Economic Theory*. London: Pinter.
- _____ and Louça. 2001. *As Time Goes By: From the Industrial Revolutions to the Information Revolutions*. Oxford: Oxford University Press.
- Grace, C. S. 1997. *Biotechnology Unzipped: Promises and Realities*. Toronto: Trifolium Books Inc.
- Harada, T. 2009. "The division of labor in innovation between general purpose technology and special purpose technology." *Journal of Evolutionary Economics* 20: 741-64.
- Helpman, E., and M. Trajtenberg. 1998a. "A Time to Sow and a Time to Reap: Growth based on General Purpose Technologies." in Helpman (ed), *General Purpose Technologies and Economic Growth*. Cambridge: Massachusetts Institute of Technology Press.
- _____. 1998b. "Diffusion of General Purpose Technologies." in Helpman (ed). *General Purpose Technologies and Economic Growth*. Cambridge: Massachusetts Institute of Technology Press.
- Jovanovic, B. and P.L. Rousseau. 2005. "General Purpose Technologies." in Philippe Aghion and Steven Durlauf (eds). *Handbook of Economic Growth*. New York: Elsevier, North Holland.
- Knight, F. H. 1921. *Risk, Uncertainty and Profit*. New York: Houghton Mifflin.
- Korzinovy, V. and I. Savin. 2016. "Pervasive Enough? General Purpose Technologies as an Emergent Property." Paper presented to the *International Schumpeter Society Conference*, July 2016 Montreal, Canada.
- Lipsey, Richard G. and Kenneth I. Carlaw. 2004. "Total Factor Productivity and the Measurement Of Technological Change." *Canadian Journal of Economics* 37(4): 1118-1150.

-
- _____. 1998. "The Consequences of Changes in GPTs" in Helpman (ed), *General Purpose Technologies and Economic Growth*. Cambridge: Massachusetts Institute of Technology Press.
- Lipsey, Richard G., Kenneth I. Carlaw and Clifford T. Bekar. 2005. *Economic Transformations: General Purpose Technologies and Long Term Economic Growth*. Oxford: Oxford University Press.
- Lucas, R. E. 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics* 22: 3-42.
- Mokyr, J. 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford: Oxford University Press.
- Moser, P and T. Nicholas. 2004. "Was Electricity a General Purpose Technology? Evidence From Historical Patent Citations." *American Economic Review* 94(2): 388-94.
- Nelson and Winter. 1982. *An Evolutionary Theory of Economic Change*. Cambridge, Massachusetts and London England: The Belknap Press of Harvard University Press.
- Pestas, I. 2003. "The Dynamic Effects of General Purpose Technologies on Schumpeterian Growth." *Journal of Evolutionary Economics* 13: 577-605.
- Rainer A. and R. Strohmaier. 2014. "Modeling the diffusion of general purpose technologies in an evolutionary multi-sector framework." *Empirica* 41: 425-444.
- Ristuccia, C. and S. Solomou. 2014. "Can General Purpose Technology Theory Explain Economic Growth? Electrical Power as a Case Study." *European Review of Economic History* 18(3): 227-247.
- Romer, P. 1986. "Increasing Returns and Long-Run Growth." *Journal of Political Economy* 94: 1002-37.
- _____. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98: 71-102.
- Rosenberg, Nathan and M. Trajtenberg. 2004. "A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-nineteenth-century United States." *Journal of Economic History*. Vol. 64, 61-99.
- Stuart J. H. G. and M. Iacopetta. 2014. "Nanotechnology and the Emergence of a General Purpose Technology." *Annals of Economics and Statistics* 115/116.

Trajtenberg, Emmanuel, Rebecca Henderson and Adam Jaffe. 1992. "Ivory Tower Versus Corporate Lab: An Empirical Study of Basic Research and Appropriability." *NBER Working Papers* 4146.

van Zon, A., E. Fortune and T. Kronenberg. 2003. "How to Sow and Reap as You Go: a Simple Model of Cyclical Endogenous Growth." *Research Memorandum Series* 029. Maastricht: Maastricht Economic Research Institute on Innovation and Technology (MERIT).

von Hippel, Eric. 1988. *The Sources of Innovation*. New York: Oxford University Press