10ISSN 1183-1057

SIMON FRASER UNIVERSITY

Department of Economics

Working Papers

16-15

"General Purpose Technologies in Theory, Applications and Controversy: A Review"

Clifford Bekar, Kenneth Carlaw,
& Richard Lipsey

Nov, 2016

SFU DEPARTMENT OF ECONOMICS

GENERAL PURPOSE TECHNOLOGIES IN THEORY, APPLICATIONS 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 GPTs are identified and definitions discussed. We propose a definition that includes many multipurpose and single-purpose technologies and that uses micro-technological characteristics not macro-economic effects. Identifying GPTs requires recognition that they evolve continually and that there are always boundary uncertainties concerning particular items. We consider existing 'tests' of whether particular technologies are GPTs, arguing that many of these are based on misunderstandings, either of what GPT theory predicts or of what such tests can establish. The evolution of formal GPT theories is outlined, showing that only the early theories predicted the inevitability of GPT-induced showdown and surges. More recent GPT models, that are designed to model GPT's characteristics, demonstrate that GPT theory does not imply the necessity of specific macro effects. We then show how GPTs can rejuvenate the growth process without causing slowdowns or surges. We conclude that the concept is helpful, while the criticisms can be resolved.

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

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

CONTENTS

GENERAL PURPOSE TECHNOLOGIES IN THEORY, APPLICATIONS AND CONTROVERSY: A REVIEW[1]

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—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 contributions of the measurable aggregate inputs.

With the publication of Bresnahan and Trajenberg's (1995) paper "General Purpose Technologies: 'Engines of Growth?'"[2] 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 work to remove anomalies, inconsistencies and other deficiencies; and, (ii) note instances where GPT theory has proven productive in the 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 provide our responses to his most important arguments throughout our paper where they are relevant to our discussions. The second comes from economists who have sought to test what they argued to be predictions of GPT theory, often reporting refutations. 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 and should not be rejected in response to such criticisms.

We begin by identifying six common characteristics of GPTs that in total serve to distinguish them from other technologies. We then discuss some alternative definitions of GPTs, arguing that GPTs should be defined in terms of their micro characteristics not their macro-

[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 for 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. Here after we reference the 1995 paper.

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

economic effects. We then argue for the usefulness of the broader of two definition of GPTs, the one that does not exclude single-generic-purpose technologies that are widely used (e, g., the printing press and the iron steam 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 a static technology but one that 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 the example of electricity. 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. Next we 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 (1995) generation of GPT theories predicted the inevitability of an aggregate showdown 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 are agnostic regarding the macro effects of GPTs, but are instead designed to model their characteristics. 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 has been helpful, and that many of the serious criticisms directed at it can be resolved.

I. DEFINING CHARACTERISTICS

Bresnahan and Trajtenberg (1995) isolate several 'regularities' associated with GPTs. These are derived from their study of historically transformative technologies – a study highlighting both qualitative and quantitative observations. We regard their set of stylised facts as the foundation of a well specified theory of GPTs and we term items 1-6 below as the defining characteristics of GPTs.[4]

1. "Most GPTs play the role of 'enabling technologies', opening up new opportunities rather than offering complete, final solutions" (Bresnahan and Trajtenberg, 1995, p. 84).

[4] Bresnahan and Trajtenberg (1995) suggest one characteristic that we reject: the provision of a generic function. Their illustrative example of the generic function of rotary motion being provided by the steam engine has been sharply criticised by historians, and rightly so in our view. Singling out any one property is not helpful in defining a GPT and it invites unproductive debate about what the generalised function really is. For example, the steam engine initially produced reciprocal motion, but knowledge concerning the translation of reciprocal motion into rotary motion had been around for centuries. Further, we might debate which function provided by steam—rotary motion or non-location-specific power—was the key generic function?

2. “This phenomenon involves what we call ‘innovational complementarities’ (IC), that is, the productivity of R&D in a downstream sector increases as a consequence of innovations in the GPT” (Bresnahan and Trajtenberg, 1995, p. 84).
3. “As a GPT evolves and advances it spreads throughout the economy, bringing about and fostering generalized productivity gains” (Bresnahan and Trajtenberg, 1995, p. 4). We interpret “generalized productivity gains” to be those micro economic productivity gains that spread beyond the initial application of the technology in an individual firm’s specific activities as the technology evolves and advances in use.
4. GPTs enable many new downstream inventions and innovations that were technically impossible without the new GPT. These in turn enable further inventions and innovations as well as influencing the original GPT itself.
5. Since a critical characteristic of GPTs is pervasiveness, a mature GPT may have multiple uses, as does electricity, or a single generic use as do railways, as long as the use is widespread and shares the other characteristics listed here. (We say more on this later in the paper.)
6. In most of its uses the technology has no close substitutes. This distinguishes a widely used technology such as the screw, which has close substitutes in a range of fasteners, from those such as electricity for which there are no close substitutes in many of its uses such as powering the electronic computer.

The first three items are explicit in Bresnahan and Trajtenberg, the last three are implicit. For a technology to be a GPT it is necessary that it displays all of these defining characteristics. Other technologies may incorporate some of these characteristics without being a GPT. Importantly, and in contrast to some work in the GPT literature, all six characteristics defining a GPT are microeconomic technological characteristics and interactions.

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 potential, 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 does not interact with other technologies cannot become a GPT. The ability to interact depends on the internal nature of the technology itself and on what other technologies exist. On efficiency, having potential is not enough. If the technology cannot be made efficient enough to be adopted at an economic cost throughout much of the economy, its potential will not be realised. For example, hydrogen fuels cells, have many of the characteristics that provide potential to become GPTs. However, their efficiency could not (at least to date) be increased sufficiently for them to become pervasive and hence to realise that potential.

Two criticisms from historians are relevant at this point. First, some economic historians have questioned 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 those working in the literature constitutes a coherent set. We argue that our GPTs belong to a theoretically, and historically, coherent set because they all share the six defining characteristics. Other cases of potential 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 an 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, p. 218) puts it "There are many single-purpose innovations (for example, the cotton gin) with arguably large implications for TFP (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 immediate substitutes, they do not satisfy defining characteristic six.

We do not doubt that there are some "historically important innovations" that are not GPTs according to our list of characteristics, 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 set of theoretical classes that attempt to abstract usefully from the continuum of reality. 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: (i) their direct applications, (ii) their externalities, and (iii) 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 to divide individual items into sets of categories in order to theorise about any of these situations. In each case, the characteristics of the typical item will differ markedly from the characteristics of the typical items in other classes, although there will be individual items in each set that are close to boundary elements in adjacent sets. 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. Isolating an extreme class from a continuous set is not uncommon in economics.

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, p. 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".

II. DEFINITIONS

Not surprisingly there has been some debate and confusion about an appropriate definition of GPTs. We do not wish to fall into the essentialist trap of assuming that there is one correct 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 wide a group of multiple-use technologies should we include and second should we also include single-use technologies that are widely used and share all of the other defining characteristics of GPTs.

Considering the width 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.

We 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 pushes these interactions aside. 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, water wheels, 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 sense, as long as they are widely used over much of the economy and share the six 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.

Scholars are, of course, free to study any well-defined class of technologies. Both classes–those that only include multiple use technologies and that that admit single-use technologies that have all the other six characteristics–are of interest.-Thus we offer two definitions depending on which group one wishes to study.

> ***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 an input in 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 is an input in 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, for the reasons discussed above, more useful than the study of the narrower class (GPT-a). 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, Bekar and Carlaw (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 the list of defining characteristics. However, Jovanovic and Rosseau (2005, p.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 both the micro and the macro 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 them.[5]

A clear test case is provided by lasers. According to our list of defining characteristics, lasers are GPTs. 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 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 structural changes and so did not induce ‘protracted aggregate impacts’ in order to become effective.

III. IDENTIFYING A GPT

Economic historians have been rightly critical of some of economists’ choices when they identify a specific artefact—such as the dynamo—as a GPT. As Field (2011, p. 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.” [6]To deal with this and some related issues, we need first to consider the concept of technology.

Technology

Technology is not always clearly defined in the literature and 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, p. 58) define the concept as follows: “Technological knowledge, technology for short, is the set of ideas specifying all activities that create economic value.” We 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 or the arrangement of machines on the factory floor) and institutions (e.g., rules, procedures, etc.). The definition thus 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.

GPTs as evolving technologies

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. Indeed, although Lipsey,

---

[5] This problem is illustrated by a parallel with growth theory. When aggregate growth models defined technology in such a way that changes in it had to cause contemporaneous changes in GDP and productivity, the ‘productivity paradox’, where one changed while the other did not, could not even be studied.

[6] 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 development of any specific GPT.

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, p. 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 a 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 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 that go to into the knowledge of how to do something fundamental 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 the whole bundle of how to do the job as the GPT: e.g., how to make commercially useful bronze, electricity, and genetically manipulated products.

An implication of the more expansive 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 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, p. 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 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 the list of six defining characteristics. The relevant question that remains is: Does the level of abstraction needed to define GPTs help in studying some aspects of economic growth and technical change? That GPTs prove useful in theoretical models of endogenous growth suggests our answer to this question.

An example from electricity

Not surprisingly with an evolving research program, there has been some confusion on identifying particular GPTs.[7] 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. We can illustrate this procedure with respect to electricity, among the most pervasive of GPTs. A short list of just some of the key effects of this GPT include the following, all of which are examples of our defining characteristics.

[7] It is not clear, however, how much disagreement would remain on what was and was not a GPT if GPT researchers were gathered together. We suspect little.

1. The replacement of the central drive shaft used in steam powered factories by unit-drive motors attached to each individual machine allowed for the rearrangement of the factory floor from machines that used most power being placed close to the shaft to machines being placed in the order in which the product was assembled. This in turn 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 (related to defining characteristics 1, 2, 3 and 4).
2. The removal of overhead steam-driven drive shafts allowed for improved illumination, ventilation, and cleanliness (related to defining characteristics 1, 2, 3 and 4).
3. The factory was redesigned from multi-story (to get machines as close as possible to the steam-driven drive shaft located near the ground-floor ceiling) to single-story structures that were safer and cheaper to build. (related to defining characteristic 2).
4. When repair or refitting of the power source was required, only the affected machine had to be shut down, rather than all of the many machines driven by one shaft (related to defining characteristics 1, 2, 3 and 4).
5. The size of production facilities was altered dramatically. On the one hand, the size of the assembly plant was no longer limited by the efficiency of a single drive shaft to deliver power throughout the plant. On the other hand, unit drive allowed the development of many small parts suppliers. The outcome was the breakup up of multi-purpose manufacturing firms into a small number of centralised assembly firms and a large number of smaller, parts manufacturing firms (related to defining characteristics 1, 2, 3 and 4).
6. The skill level of production workers was lowered drastically as they no longer had to service and repair their individual machines. Instead, the increased reliability of unit drive machines implied that these tasks could be delegated to a small number of specialists (related to defining characteristics 2, 3, 4, and 6 but in the reverse sense that any unit of labour became a close substitute for any other).
7. Electric lighting had many unexpected advantages over gas lighting: fewer acidic fumes, important in textile manufacturing; non-moisture producing compared with gas, important where control of humidity mattered; no gas leaks, important where explosions were a risk; no temperature increases, important where temperature control mattered; no flickering, important where even light mattered (related to defining characteristics 1, 2, 3, 4 and 5).
8. Glass blowing machines replaced skilled glass blowers (related to defining characteristics 1, 2, 3 and 4).
9. Skilled mule drivers were replaced in mines by electric locomotives (related to defining characteristics 1, 2, 3 and 4).
10. Machines that could regulate electrically generated heat by the push of a button replaced the skilled workers required to maintain heat in furnaces run by fossil fuels (related to defining characteristics 1, 2, 3, 4 and 6).
11. Electricity created a whole range of new household products and the rise of industries to produce them–washing machines and dryers, dishwashers, refrigerators, electric

stoves, toasters, mixers, can openers, blenders, automatic bread makers and so on (related to defining characteristics 1, 2, 3, 4, 5 and 6).

12. The new electric trolley contributed to the development of modern suburbs. (related to defining characteristics 1, 2, 3, 4, 5 and 6)
13. Electric powered railways enabled subway systems that would have been impossible with steam power (related to defining characteristics 1, 2, 3, 4, 5 and 6).
14. Electricity enabled several later GPTs. Without electricity there could have been no electronic computers, no lasers, no internet, no industrial robots, no biotechnology, and no nanotechnology (related to defining characteristics 1, 2, 3, 4, 5 and 6).

Although this (partial) list of changes brought about by electricity is sufficient to demonstrate that electricity is a GPT, we can also consider a simple thought experiment. 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 carrier pigeons 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.

To repeat, 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 case of electricity, by an enumeration of observed characteristics. This usually requires extensive historical research, research that uses both qualitative and quantitative data.

IV. EMPIRICAL MEASUREMENTS AND TESTS

Rather than make the historical analyses referred to for electricity in the previous section, some writers have tried to determine whether or not a particular technology is a GPT through the analysis of quantitative empirical data. There are, however, two different ways in which these studies can be interpreted. They may test a prediction derived from GPT theory, or they may test the existential statement that the technology in question is a GPT. We argue that the studies considered in this section do not deal with what can be regarded as predictions from well-specified GPT theories and so are to be understood as testing the statement that technology X is a GPT. So we must consider how such existential statements can and cannot 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 just 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; 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. The concept of a GPT has been criticized by Field (2011, p. 218) for being unable to rule of such trivial technologies as felt-pens, which are widely used and for various purposes. Although the statement that felt-pens 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

In spite of what has just been said, some researchers have attempted to use empirical data to test whether or not a specific technology is a GPT. The two candidates for such data that are found in the literature are patent and macro data (typically related to productivity, GDP, and their growth rates). We postpone the consideration of the use of macro data until Section V where it is relevant. Here we consider the use of patent data. Is it possible that such data can test the statement that technology X is a GPT? Our answer is ‘no’ for several reasons.

First, none of the existing formal theories of GPTs make any predictions about patents. Although Moser & Nicholas (2004) state that “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”, 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. Second, as we show in detail below, the universe that contains all of the characteristics of a GPT is far broader in both time and space than the universe that contains the patents that have been studied. Furthermore, the relevant universe of all possible observations co-evolves with all technologies. So, while patent data may help to confirm that a particular technology is a GPT, it can never establish that it is not.

Let us illustrate with respect to electricity. Assume that we have been able to associate each of our six defining characteristics with specific types of patents and a search of the patent data finds no patents related to at least some of these characteristics. We would have failed to prove, using patent data that the technology was a GPT. But we would not have proved that it was not a GPT. Indeed, there are many reasons why we should not be surprised by this failure. Patents are an imperfect instrument for measuring a technology’s antecedents, externalities and technical complementarities, especially those attributable to a GPT. Many of the new patents, even those wholly reliant on electricity, did not cite either earlier patents, or earlier innovations related to electrification that were themselves not patented. Also, the technological complementarities cover a much bigger set of effects than can be captured by forward patent citations. To illustrate, consider the effects in the above list concerning the effects of electricity. Many, possibly most, of the effects mentioned in items 1 – 6 (concerning the layout of the factory, adoption of the unit drive, the skill level of production workers, etc.) in that list were not patented. While the entry under item 7 (electric lighting) no doubt gave rise to patents, the many important side effects sampled in the list did not. Items in points 8 -11 (the impact on the supply and demand of specific skills within electrified factories) obviously gave rise to patents as did those under 12 and 13 (rise of trolleys and railways). But in those cases most of the myriad side effects that resulted did not lead to patents that were attributed to electricity. Finally, the open ended item 14 (GPTs enabled by electricity) illustrates how when a GPT becomes a fully mature part of the economy, it often enables downstream GPTs whose myriad related technologies patents would not be covered in a survey of electricity patents. For example, few patent filings in contemporary Silicon Valley bother to cite electricity.

An example of the points just made is provided by Moser and Nicholas (2004). They analyze patent data in an attempt to determine whether or not electricity—which they argue is held to be among the quintessential GPTs—actually meets the theory's criteria. Specifically, they use patent citations in an attempt to determine how general the innovation was, and in how many downstream innovations it played a role. 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. They argue (Moser and Nicholas, 2004, p. 388-89), "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)."

The data do not sustain these strong conclusions. For example, their measure of originality covers 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 the major users of electricity, computers and communications, due to lack of data. Also, 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. Furthermore, many of the patents identified in chemicals actually required electricity-using innovations to make the technology in question work.

We conclude that although their work provides some interesting and informative data, it cannot sustain their strong conclusions. The absence of relevant patent data does not refute the proposition that any one technology, let alone electricity, 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 existence statement that electricity is a GPT.

Patent data corroboration

Patents have been successfully 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 derivative technologies that it has enabled. (But 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 GPTs, Trajtenberg, Henderson and Jaffe (1992) use patents to

> …explore the use of 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 Lacopette (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 based on the definition. In particular, they construct hypothesis 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 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 conclude that the technologies were not GPTs because they found insufficient evidence that they were.

V. 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 adequately 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;
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 productivity. 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 these theories have many assumptions that are not found in the observations stated by Bresnahan and Trajtenberg (1995) and repeated in our list of defining characteristics (or found in other well-known literature concerning transforming technologies). Some of these are merely for simplification. Others are patently counterfactual and important contributors to some of the results, in violation of the important methodological rule that assumptions that are patently counterfactual should be robust in the sense that they can be relaxed without seriously altering the theory's predictions. Key assumptions include the following (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. *In practice 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 first arrives 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 point 2 that in these theories there can be no uncertainty in any decision regarding GPTs. *In contrast, empirical evidence shows that uncertainties (in Knight's (1921) sense) are ubiquitous in the development of almost all new technologies, let alone GPTs.*
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 each other as did the assembly line, just-in-time inventories and robot-operated assembly lines, while others or are more or less independent of each other as were the vacuum cleaner and the electric washing machine—all of which depend on electricity.*
5. All the models that address the issue predict a slowdown of one macro sort or another, followed by a surge after 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). *The assumptions of the models were designed to produce such slowdowns.*

Although the heuristic value of these models is not at issue here, it is clear from these points, and the fuller discussion in Lipsey, Carlaw and Bekar (2005), that these models fall short of capturing our six defining characteristics. However, as Lipsey, Carlaw and Bekar (2005) also observe, 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" (p.384).

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 has several counterfactual characteristics that one would like to remove. First, 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 papers 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 transition to a new steady state will follow 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 and 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 theories

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 we consider Carlaw & Lipsey's, (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 output from 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 GPTs continued evolution			X
Uncertainty regarding GPTs replacement within its class			X
Uncertainty regarding time required to replace existing GPTs			X
Agents unable to maximize over life of GPT			X
Distinct evolutionary paths for each GPT			X

To incorporate all of these assumptions, the model cannot employ a stationary dynamic equilibrium concept. Instead, an unorthodox model is required. 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 spillovers 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 a variable rates of innovation.

Fortunately, it turns out that 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.'

One characteristic that is important but not included in any of the models mentioned so far is that a technology may emerge in a small 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, V. and I. Savin (2016).

A key to modelling this characteristic is that GPTs should not be *ex ante* identifiable. In all four of the above models (and in Carlaw and Lipsey (2011)) a candidate technology may fail to become a GPT.[8] To allow for this, 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.

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 GDPs 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 verbally in Lipsey Bekar and Carlaw (1998b, p. 207).
2. In the model's many simulations, changes in measured total factor productivity (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).[9] This also conforms to what we observe, as for

[8] 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 fail to become an actual GPT if it fails an expected efficiency criteria and is not eventually adopted in at least one applied sector of their model.

[9] See Carlaw and Lipsey (2002 and 2004) and for discussions of this point.

example when, in the late 1980s, there were large changes in ICTs but little change in the statistics for the 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 new 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 studied after these theories were presented.

Another important prediction of GPT theory is that GPTs transform the structures of the economy and society. Although this prediction cannot yet be derived from a formal model, there is ample historical evidence of its correctness. 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 created 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 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.

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 could cause 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 & Solomou (p. 229) assert 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."

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 prediction of GPT theory.

There are at least three serious problems with equating the arrival of GPTs with growth slowdowns and surges, all manifested in existing empirical work. First, different first generation theories predict different types of slowdowns. For example, Helpman and Trajtenberg (1998a) predict cycles of slowdowns and bonuses in both total factor productivity (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, p. 1314-15) conclude: "…there needs to be no observable impact of new technologies on rates return…" Also Lipsey, Bekar and Carlaw (1998b) argue (p. 207) "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: "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" (p. 212). This is what the Carlaw and Lipsey (2011) model predicts: slowdowns and surges associated with the introduction of a new GPT under some specified conditions but not others.

Third, it is clear 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). Furthermore, we observe that the electronic computer (a major component of IT) coincided with a secular boom in GDP and productivity over the first thirty years of its evolution (circa 1945-75) followed by a slowdown in productivity growth that lasted for many decades after that.

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, p. 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 (p. 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 the 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. For examples, Jovanovic and Rousseau (2005) argue that electricity was associated with a productivity slowdown as measured by labour productivity. 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 (interesting although not quite the slowdowns followed by bonuses predicted for the macro economy by the first generation theories). Finally, Rainer and Strohmaier (2014) 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" (p. 427).

VI. REJUVENATION OF GROWTH POTENTIAL WITHOUT THE NECESSITY OF SLOWDOWNS OR SURGES[10]

Regarding our list of predictions above, 1 - 3 agree closely with the evidence, but one might wonder how 1 and 4 can both be correct. How can a single new GPT rejuvenate the growth process without causing a productivity slowdown and/or 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 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 what how the growth rates of productivity and GDP will change when a new GPT and its derivative

[10] This section relies heavily on Carlaw and Lipsey (2002)

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.

FIGURE 1: CAPITAL ACCUMULATING AND TECHNOLOGICAL CHANGE

FIGURE 1

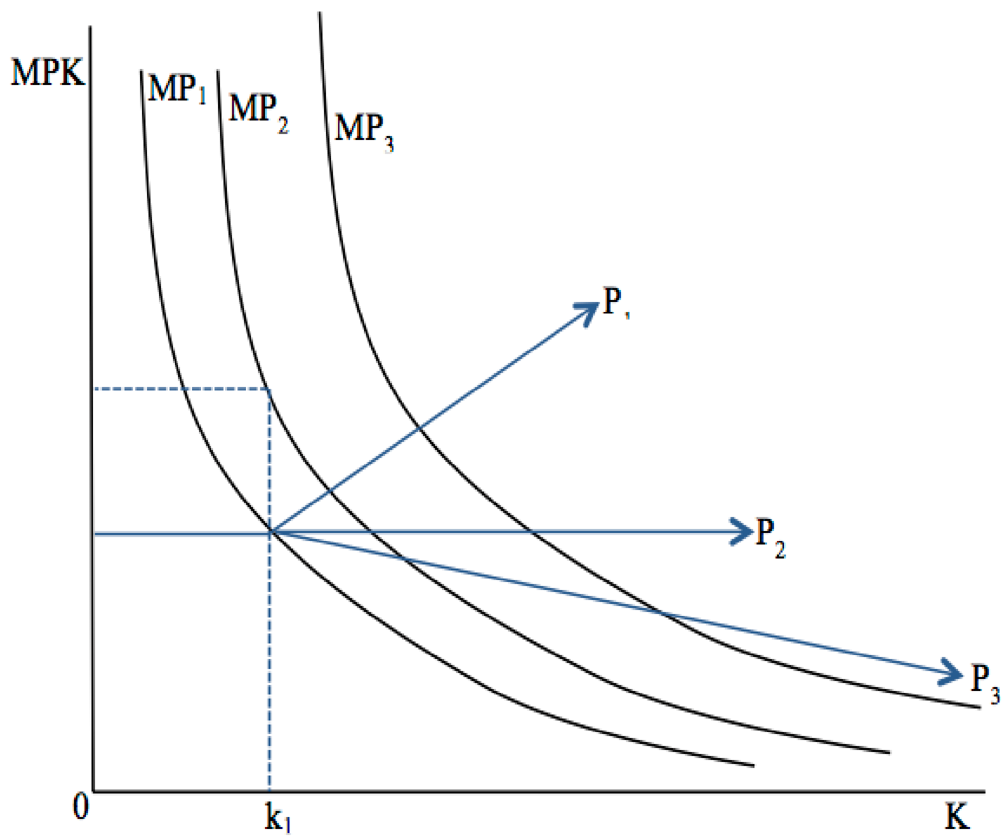

VII. CONCLUSION

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 that 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, a least for existing growth theory.)

They 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. The resulting concept of GPTs may or may not be useful to economic historians. However, if economists wish to continue these research programs 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.

END

REFERENCES

Aghion, P. and Durlauf, S.N. (2005). *Handbook of Economic Growth. Amsterdam*: Elsevier North-Holland.

Aghion, P. and Howitt, P. (1992) "A Model of Growth Through Creative Destruction." *Econometrica*. vol. 60 no. 2: 323-51.

Bresnahan, T. and Trajtenberg, M. (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 and Richard G. Lipsey (2002) "Externalities, Technological Complementarities and Sustained Economic Growth." *Research Policy*. 31, 1305-1315.

__________________________________ (2004) "Total Factor Productivity and the Measurement of Technological Change." *The Canadian Journal of Economics*. November 31(4), 1118-1150

___________________________________(2006) "GPT-Driven Endogenous Growth" (with Kenneth Carlaw). *Economic Journal*. January, 116, 155-74

___________________________________(2011) "Sustained Endogenous Growth Driven by Structured and Evolving General Purpose Technologies." *Journal of Evolutionary Economics*. 21(4), 563-593.

____________________________________(2012) "Does History Matter: empirical analysis of evolutionary versus stationary equilibrium views of the economy." *Journal of Evolutionary Economics*. 22 (4), 735-766.

Cantner, U. and Vannuccini, S. (2012) "A New View of General Purpose Technologies" in Heilemann U., Wagner A., *Empirische Makroökonomik und mehr*. Lucius et Lucius, Stugart.

___________________________ (2016) "Competition for the (Downstream) Market: Modeling Acquired Purposes." Paper presented to the International Schumpeter Society Conference, Montreal, Canada, July 2016.

Crafts, N.F.R. (2004) "Steam as a General Purpose Technology: A Growth Accounting Perspective." *Economic Journal*, Vol. 114, 17-29.

Crandall, B,C, (1966) 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*. Vol. 80, No. 2, 355-361.

_________ and Wright, G. (2003) “General Purpose Technologies and Surges in Productivity: Historical Reflections on the Future of the ICT Revolution.” In P.A. David and M. Thomas (eds.), *Economic Future in Historical Perspective*, Chapter 4. Oxford: Oxford University Press, 135-166.

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*. Vol. 21, No. 2, pp. 249-275.

Field, A. (2011) *A Great Leap Forward: 1930s Depression and U.S. Economic Growth*. Yale University Press: New Haven.

Goldfarb, B. (2005) “Diffusion of General Purpose Technologies: Understanding Patterns in the Electrification of U.S. Manufacturing 1880-1930.” *Industrial and Corporate Change*. Vol. 14, 745-773.

Gordon, R. (2016) *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton University Press: Princeton.

Grace, C. S. (1997), *Biotechnology Unzipped: Promises and Realities*. Toronto: Trifolium Books Inc.

Hall, B. H., A. B. Jaffe and M. Trajtenberg (2001) ‘The NBER patent citations data file: lessons, insights and methodology tools.’ N*ational Bureau of Economic Research*, Working Paper No. 8498, Cambridge, MA.

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 Trajtenberg, M. (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.

Helpman, E. (ed) (1998) *General Purpose Technologies and Economic Growth*. Cambridge: Massachusetts Institute of Technology Press.

Jovanovic, B. and Rousseau, P.L. (2005) "General Purpose Technologies." in Aghion and Durlauf (eds). *Handbook of Economic Growth*.

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, Montreal, Canada, July 2016.

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.

Lipsey, Richard.G., C. Bekar, and K. Carlaw (1998a) "What Requires Explanation?" in Helpman (ed), *General Purpose Technologies and Economic Growth.*

____________________________________(1998b) "The Consequences of Changes in GPTs" in Helpman (ed), *General Purpose Technologies and Economic Growth*.

Lipsey R.G., K.. Carlaw and C. 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. (1992) *The Lever of Riches: Technological Creativity and Economic Progress*, Oxford University Press.

_________ (2003) *The Gifts of Athena: Historical Origins of the Knowledge Economy*, Princeton: Princeton University Press.

Moser, P and Nicholas, T. (2004) "Was Electricity a General Purpose Technology? Evidence from Historical Patent Citations." *American Economic Review*, Vol. 94, No. 2. 388-94.

Nelson and Winter (1982) *An Evolutionary Theory of Economic Change*. The Belknap Press of Harvard University Press, Cambridge, Massachusetts and London England.

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.

________________________(2016) "Studying general purpose technologies in a multi-sector framework: The case of ICT in Denmark," *Structural Change and Economic Dynamics*. 36, 34-49.

Ristuccia, C. and Solomou, S. (2014) "Can General Purpose Technology Theory Explain Economic Growth? Electrical Power as a Case Study." *European Review of Economic History*, Vol. 18, No. 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, s71-102.

Rosenberg, Nathan (1963) "Technological Change in the Machine Tool Industry, 1840-1910." *The Journal of Economic History*, 23-4, 414-443.

________________ (1976) *Perspectives on Technology*, Cambridge: Cambridge University Press.

________________ (1982) *Inside The Black Box: Technology and Economics*, (Cambridge: Cambridge University Press).

________________ (1994) *Exploring The Black Box: Technology, Economics And History*. Cambridge: Cambridge University Press.

________________ (1996) "Uncertainty and Technological Progress" in *The Mosaic of Economic Growth*, R. Landau, T. Taylor and G. Wright (eds). Stanford: Stanford University Press.

________________(1998) "Chemical Engineering as a General Purpose Technology." In Helpman (ed). *General Purpose Technologies and Economic Growth*.

_________________and Trajtenberg, M. (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*. No. 115/116

Solow, R. (1957), "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics*. XXXIX, 312-20.

Trajtenberg, Emmanuel, Rebecca Henderson and Adam Jaffe. (1992) "Ivory Tower Versus Corporate Lab: An Empirical Study of Basic Research and Appropriability" No 4146, *NBER Working Papers* No 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." in *Research Memorandum Series*, 029. Maastricht: Maastricht Economic Research Institute on Innovation and Technology (MERIT).