Patent Thickets: Taxonomy, Theory, Tests, and Policy

UNABRIDGED EARLY DRAFT
DO NOT CITE OR DISTRIBUTE

Edward J. Egan and David J. Teece*

June 21, 2013

Abstract

We assemble the near population of 164 papers that define the term 'patent thicket'. We then ask three questions: 1) what is a patent thicket? We identify four distinct types of patent thicket definitions used in the literature – diversely-held complementary inputs, legitimate overlapping patents, spurious patents, and effectively saturated invention spaces – and describe the economic foundations of each in turn. We also identify a number of variants and sub-types that apply across and within these definitions, and so create the first taxonomy of patent thickets in the literature; 2) which measures and tests are appropriate to understanding patent thickets? Each type of thicket that we identify has dramatically different implications for measurement, tests of whether thickets exist and provide a hindrance to innovation, and appropriate policy responses. We articulate these implications and show how the measures, tests, and policy advice provided to date have been appropriate or inappropriate to each definition used and the context in which it was used; and 3) has the literature come to any well founded conclusions about patent thickets? Although we document a general and growing confusion over the meaning of the term 'patent thicket', and come to the unfortunate conclusion that many authors have implemented measures and tests, or given policy advice, that is at odds with the economic foundations of the types of patent thickets they have purported to study, we suggest that considerable progress has been made in the patent thicket literature. Overall it seems likely that two or three, if not all four, of our patent thicket types do exist but that generally any hindrance to innovation that they cause is a symptom of the technologically advanced and highly sophisticated innovation environment present in the early 21st century. The exception is a sub-type of the 'spurious patent' patent thicket, which relies on the issue of patents that fail to meet the requirements for novelty or non-obviousness. There is a growing concern that such patents are becoming more common, particularly in the U.S., but very little supporting empirical evidence. If these patents are being issued, then thicket problems seem almost inevitable, and a sensible policy response might be greater funding to support an increased effort in pre- and post-issue reviews of validity by the patent office.

*Dr. Egan is a post-doctoral fellow at the Sauder School of Business, University of British Columbia. Dr. Teece is the Thomas W. Tusher Chair in Global Business, Haas School of Business, University of California, Berkeley. We would like to thank Kai Rong Gan, Elizabeth Olson, Andrew Sang, and particularly John Davis, for their excellent research assistance. We would also like to thank Ed Sherry for his many useful comments, as well as for figure 2 on page 21. Please address all correspondence to: ed_egan@haas.berkeley.edu.

1 Introduction

A 'patent thicket' describes the situation where the relationships between existing patent rights have a potentially adverse effect on future innovation. Patents provide incentives for innovation but patent thickets, according to Shapiro (2001), are a "a dense web of overlapping intellectual property rights that a company must hack its way through in order to actually commercialize new technology". Patent thicket research is now just over a decade old. Around 50 papers have explored the nature and consequences of patent thickets. Hundreds more have done work on firms' strategic responses to patent thickets, private mechanisms to mitigate their supposed consequences, or have advocated (sometimes radical) policy responses to ameliorate their alleged negative effects. The European Patent Office (EPO) has recently hosted a workshop on patent thickets, the UK Intellectual Property Office (IPO) has commissioned a series of reports into patent thickets, and the last two major reports on intellectual property produced by the Federal Trade Commission (FTC) both devoted substantial space to patent thickets.[1] In short, patent thickets, and their perceived threat to the innovation ecosystem, are currently receiving a lot of attention in innovation research, and policy makers are poised to begin interceding in innovation markets to try to remedy the 'patent thicket problem'.

This paper examines the near-population of 164 papers that define the term patent thicket, including 50 papers directly on the topic of patent thickets, of which 20 have provided empirical analysis of their effect. We report a general confusion in this literature on many different dimensions. Specifically, we find that: there are actually many different kinds of patent thicket, each with a separate theoretical basis; authors frequently provide definitions of patent thickets that confound the separate theoretical arguments of the various underlying types; some papers mistakenly suggest that hold-up is an integral part of certain types of patent thicket; 'standard' measures of patent thickets suffer from construct validity issues, and it is far from clear that the measures represent just one, or in some cases any, of the separate types of patent thicket; papers estimating the impact of patent thickets have implemented experimental designs that are incapable of discriminating between the different types of patent thickets; and stylized facts concerning patent thickets are based upon findings that lack identification and could be caused by non-patent-thicket effects.

We uncover four main different types of patent thicket, which we label and name:

- A – diversely-held complementary input thickets
- B – legitimate overlapping patent thickets
- C – spurious-patent patent thickets
- D – effectively saturated invention-space thickets

A full taxonomy of patent thicket types, sub-types, and variants is provided in table 9 in the appendix. For now, we will ignore the sub-types, which are refinements of the main types, except for type B, which we will decompose into $B1$ – cumulative overlaps, and $B2$ – adjacent overlaps. Variants modify thicket types, for example adding transaction costs (variant T), search costs (variant S), or probabilistic patent (variant P) considerations. These, too, are discussed in the main body of the text. Figure 1, below, provides a simple overview of the economic fundamentals of the different types of patent thicket.

Innovation relationships based upon complementarities and substitution have different theoretical foundations. The modern definition of innovation consists of two parts: invention and commercialization. A relationship between a new invention and pre-existing patent rights and the relationship between commercialization of a new product and pre-existing patent rights, likewise has different theoretical foundations. Each of these four possibilities, as well as the possibility of transfer-seeking by patent applicants, leads to different type of patent thicket.

[1] See EPC and Board (2013), Team (2011), Hargreaves (2011), Commission (2003), and Commission (2011), respectively.

Figure 1: Flowchart for theory-based classification of patent thickets

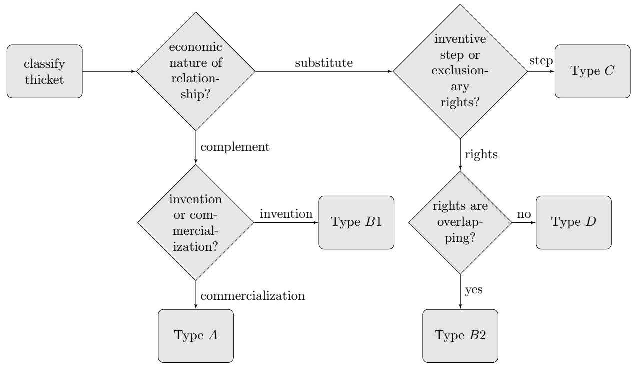

When complementarities pertain to commercialization, type A (diversely-held complementary input) patent thickets may arise. Envision a situation where the manufacturer requires many diversely-held complementary inputs that are patented to produce and sell a new product. If the holders of these patents are unable to coordinate, each will face a Prisoner's Dilemma when pricing their input. A higher input price will increase profits per unit while the decrease in profits from the reduction in demand that a higher price will induce, is shared with its rivals. The profit maximizing solution is for the N patent holders to share a single monopoly mark-up, but perverse incentives create N-fold marginalization instead. The more diversely-held the patented inputs (i.e., the greater the N), the greater the potential welfare loss from both lost rents to invention and deadweight-loss in product markets. This patent thicket problem was first articulated by Shapiro (2001). It does not involve hold-up.

A very different situation arises with when there are complementarities between existing patent-rights and a new invention. Now one invention needs another. This can happen when invention is cumulative and the exclusionary rights of the upstream invention 'overlap' with the exclusionary rights of the downstream patent. This is possible when each inventive step is different and so each patent is legitimate and correctly issued by the patent office. Examples include 'improvement patents', where a downstream patent adds to the existing functionality of the upstream patent in a novel, non-obvious, and useful manner, as well as patents on materials and research tools that will be used in subsequent inventions. This problem, which we name $B1$ – overlapping cumulative invention patent thickets, can cause economic inefficiencies when the allocation of rents to the downstream invention are shared with the upstream patent-holder sub-optimally. This was first described as a 'patent thicket' problem by Heller and Eisenberg (1998). As Reach -Through Licensing Agreements (RTLAs) are often used to distribute the rents to follow-on invention, the potential for hold-up is one underlying cause of market failure.

Yet more different types of patent thickets, again with different economic fundamentals arise from substitution. Substitution can happen either at the level of the inventive step or at the level of the exclusionary rights. When the inventive steps of two patents are substitutes, the second patent was not novel (or non-obvious given a combination of the patented prior-art) and should not have been issued. The issue of 'spurious patents' leads to type C thickets, which allow transfer-seeking and could have potentially devastating economic consequences for the innovation ecosystem. We argue that when a patent applicant exhibits guile in the filing of a spurious patent, the problem of spurious patents can be characterized as one of 'patent hold-up'.

Substitution at the level of exclusionary rights can be a problem either because it happens or because it can't happen. In the former case, if inventive steps are distinct and valid but rights are perceived as overlapping, more than one party may believe that they have exclusionary rights over the same domain of application. This belief can be rational when prosecution costs are prohibitively high for some firms or when the enforcement of rights is imperfect. We refer to this as a type $B2$ adjacent invention patent thicket; its economic loss arises from resources wasted in patent disputes.

On the other hand, if rights are discrete but the invention-space is effectively saturated with patents, further substitution may be impossible. We refer to this as a type D thicket. When different inventive steps can be used to achieve the same economic ends, firms may engage in 'ring-fencing' and patent the alternatives around their new invention. Perhaps surprisingly, we argue that this does not necessarily have negative welfare consequences.

The following example shows a patent thicket definition that confounds many different patent thicket types:

> "...complexity of a technology implies that patents are natural complements, and therefore hold-up arises easily if patent ownership is dispersed... patent proliferation is causing regulatory blockage in the form of 'thickets' of pre-existing patents and pending patents which impede genuine innovators wishing to enter markets.... A particular danger from increasing numbers of patents is the development of 'thickets' of patents with overlapping claims. The result of these is that businesses working at the leading edge of a particular technology may find it difficult or even impossible to know with whom they are in conflict, or whom they should approach for a licence.... As well as added transaction costs, patent thickets encourage strategic or defensive patenting behaviour, particularly where there is fragmentation of IPRs into the hands of multiple owners.... This would encourage the surrender of less valuable patents, reducing the density of thickets in a way consistent with achieving the maximum net economic benefit." – Harhoff et al. (2008)

The measurement and testing of patent thicket effects also suffers from issues. There are almost as many measures of patent thickets as there are papers measuring them. This would be helpful if the construct validity of measures was well understood, as then understanding could be reinforced or refined. Unfortunately, the construct validity of measures is exceedingly questionable, and has not been subject to systematic exploration. The confusion over which types of thicket are being analyzed is then amplified by a confusion over which measures might reflect each type. Two measures are particularly common. Ziedonis (2004)'s 'fragmentation' measure was intended to capture the extent to which complementary input patents are diversely-held (i.e., type A patent thickets). We show that it doesn't: instead it could capture the extent to which ownership of complementary inputs is fragmented holding the diversity of ownership of these inputs constant. In addition, Ziedonis (2004)'s measure has been used to estimate thicket types B through D, which have little if anything to do with diversity of ownership. Von Graevenitz et al. (2011)'s 'triples' measure is even more problematic. We discuss it in detail on page 42 and show that it could measure almost any type of patent thicket, or none of them.

The study of patent thickets also suffers from potential experimental design issues beyond those involved in measurement. Type A patent thicket provide a simple example. When products require a large number of diversely-held patents, their commercialization may become infeasible. N-fold marginalization issues may reduce the profits available to implementers below their opportunity costs, or, even more simply, the transaction costs involved in securing the rights to all of the complementary inputs could become prohibitive. In this case some products may not be commercialized and, knowing this, some inventors may decide not to pursue certain inventions because they would never be used. However, firms may anticipate this problem and engage in a 'defensive patenting' strategy. This means that firms may pursue inventions, and patents on inventions, that may be used as complementary inputs in their rival's products. This could provide a disciplining mechanism to mitigate the N-fold marginalization problem, and may also act as an inducement towards cross-licensing agreements that would mitigate the transaction cost problem. Papers in the literature have used measures of patent thickets to explain measures of patenting activity. Such an experimental design can not determine that patent thickets do not exist: A null finding would provide no evidence, and both a positive and a negative statistically significant correlation could be used as evidence to support the existence of patent thickets, depending on whether one pre-supposes the strategic response or not.

An experimental design based on thicket measures predicting patenting activity also can not differentiate between different types of patent thickets. Type A and type C (spurious patent) patent thickets both could be consistent with a finding that patenting activity is positively or negatively correlated with a thicket measure. For type C thickets, an increased thicket density should give rise to more spurious patents but fewer genuine patents. Type $B1$ (legitimate overlapping cumulative invention) patent thickets should reduce patenting activity. And even if the thicket measure did capture the extent to which complementary inputs are diversely-held, such an experimental design would fail to identify that the observed effect came from N-fold marginalization or transaction costs. Likewise, it is easy to conceive an omitted variable that might drive both an increase in the thicket measure and an increase in patenting activity. For example, the greater information storage, accessibility, and processing of the Internet age might facilitate both the creation of more patentable inventions and the assembly of more complex products.

This research sheds light on all of these issues, and attempts to put the foundations in place to address them for future research. In section 2 we discuss the economic foundations for innovation and the patent system in general, before turning to the economic nature of various patent relationships. A brief history of the tangled web of patent thicket definitions suggests that definition-by-analogy was probably responsible for much of the confusion over how, why and where patent relationships might hinder innovation.

In section 3, we establish our patent thicket taxonomy. For each type of patent thicket we trace its lineage in the literature; consider which industries, patent or product characteristics might be associated with it; discuss its economic foundations; where possible, decompose it into various sub-types and variants, summarizing relevant theory; and provide real-world examples.

In section 4 we conduct a meta-analysis on the near-population of papers that have defined the term patent thicket. This provides empirical evidence of a confusion in the literature, and suggests that this confusion is still growing. However, many points of agreement and understanding shine through. Type A thickets are associated with the semiconductor sector and other industries characterized by complex products, whereas type B thickets are more likely to be considered in biotechnology, academic research, and other sectors where inventions are likely to be close to basic science and patents may be considered 'broad'. Type C thickets have received less explicit attention, and researchers sometimes shy away from suggestion that the patent office may mistakenly issue a large number of spurious patents. There appears to be little to no evidence that this actually happens, but type C thickets are associated with software and business method patents, where patent quality issues are frequently alleged.

Section 5 examines the empirical foundations of the patent thicket literature. It analyses the approximate population of 20 papers that have reported the results of empirical research into patent thickets. We first consider the measures of patent thickets that have been used to date. Almost all of these measures are based on patent statistics, and the majority are based on patent citation counts. Patent citations are used extensively to measure the relationships between inventions, most likely because almost no other measures exist or are so easily obtainable. We question whether patent citations contain the information required to measure our different types of patent thicket. Citations exist to demonstrate that a patent is novel and non-obvious. It therefore seems unlikely that they will contain systematic information about which patented inputs might be assembled together in a future product. It seems more likely that they might contain systematic information about which patents have been (at least partially) replaced. When the measures are better understood, a reinterpretation of the literature to date might well provide a plethora of new stylized facts. We also discuss the tests used to estimate the effects of patent thickets, and suggest that more attention needs to be paid to developing experimental designs. The main effects for each dependent variable for each paper in our empirical papers sample are summarized in table 6. It is very difficult to draw any substantive conclusions from them in aggregate, which suggests the empirical underpinnings are far less developed than the theoretical ones. Table 7 summarizes whether empirical papers claim, rather than find, whether patent thickets exist and hinder innovation. Existence is widely claimed, but hindrance is not.

In section 6 we discuss policy recommendation that have been made in the hope of alleviating market failures for patent thickets, and in section 7 we summarize our work and conclude by providing our opinions on whether patent thickets exist and hinder innovation. We suggest that type A and B thickets do probably exist, at least in potential, but that private mechanisms most likely prevent any large scale hindrance to innovation and that policy interventions are unlikely to be more fruitful (and may be counterproductive). On the other hand, we have no evidence that type C of D thickets do exist. However, we remain extremely concerned over the potentially devastating effects type C thickets could have on the innovation economy.

2 Background

2.1 Innovation and Patents

The modern definition of innovation has two parts: Coming up with an invention and then commercializing it. Invention is the creation of something new and useful. Commercialization is the process of developing a product and bringing it to market. The novelty and usefulness of an invention is a pre-requisite for a utility patent, as are non-obviousness and a codifiable inventive step; a patent applicant is required to describe either the "advantages of the invention" or "how [the invention] solves problems previously existent in the prior art" to the patent office.[2] A utility patent is a transferable right to exclude granted to the inventor of something new and useful, providing that that something new and useful pertains to a "process, machine, article of manufacture, composition of matter, or [is an] improvement thereof".[3] This research is concerned only with utility patents and will use the term patent synonymously with 'utility patent'.

The two normative economic rationales for assigning patents – mitigating externalities and ameliorating information problems – are focused on invention more than its commercialization. Patents confer exclusive rights to their inventors to provide an incentive to invent. Without exclusive rights, others could free-ride on costly inventive activity, and Arrow's Paradox would apply.[4] And patents must be published for all to see.[5] Publishing a description of the invention puts that description into the public domain. Any interested party can see all of the extant protected inventions, use their ideas to conceive new ones, and avoid wasteful effort in duplication.

Commercializing an invention provides the return to inventive effort.[6] Commercialization might be undertaken by the inventor (or his/her employer) directly, or the patent might be transferred to another party. The transfer can take the form of an outright sale of rights, also called assignment, or a license, which might be exclusive or not, and which might have specific terms relating to follow-on invention

[2] See the Manual of Patent Examining Procedure (MPEP), section 608.01(a).

[3] This stands in contrast to a design patent which pertains to "ornamental design for an article of manufacture". Source: www.uspto.gov/main/faq

[4] Arrow's Paradox is where an inventor can't share, sell, or finance their idea, as to do so they must disclose it and risk appropriation.

[5] Granted patents have always been published by the USPTO. The USPTO began publishing patent applications, 18 months after the date of filing, on November 29th, 2000.

[6] It is also possible that a patent might be bought to be shelved. That is, a firm may purchase a patent for strategic reasons such as to prevent competition. This is discussed later.

attached to it.[7] Specialization and the efficient allocation of resources provides the economic rationale for the transferability of patent rights.

The Schumpeterian definition of innovation stresses commercialization over invention. Schumpeter defined innovation as the "carrying out of new combinations".[8] The components used in these 'new combinations' could already exist provided that their combination was new. Therefore patents play a slightly different role in Schumpeterian innovation; here it is their transferability that is key. The patent office stands as a central repository providing descriptions of components that innovators can license and assemble together.

In the last decade, however, there have been growing concerns that in certain specific situations patents might have the unintended effect of hindering innovation. These concerns have are based on the nature of certain relationships between patents – when patent rights form a patent thicket.

2.2 The Explosion of Patents and their Relationships

There has been an 'explosion' in patenting activity in the U.S. and E.U. since the 1980's. Many of the papers included in this review comment on this explosion, and Hall et al. (2012) provides an informative diagram.[9] In the U.S., for example, the number of patent applications ranged from around 50,000 to 100,000 per year from 1900 to 1980, and then rose from around 100,000 in 1985 to over 450,000 in 2006. It is possible that this rise has now ended, but in 2010 the backlog of patent applications at the USPTO was so large that the average time to prosecute a patent application was around three years.[10] The patent explosion might have important consequences for patent thickets. To see why, we must first understand how patents might relate to one another.

Patents might be economic substitutes for one another.[11] This substitution could be perfect, so that the pre-existing patent now have new competitors, or it could be imperfect, for example through the creation of incremental improvements, or superior or inferior alternative technologies that achieve roughly the same economic ends. Or patents might be complementary to one another, so that many patents can be used together to create a new product or improve an existing product. Or, perhaps, patents might be independent of one another, so that new patents indicate new things that did not exist before. In each case, positive effects for social welfare are apparent: Either something is being made better, something new is being created, or competition is being increased. But, except with independent patents, there are also hidden complexities that can prevent these potential gains to welfare. For example, a lack of cooperation between complementary input patent-holders can lead to over-pricing, or an inability to agree a 'fair' distribution of future rents might prevent a patent-holder from licensing their patent to an inventor who wanted to create an improvement or find a new usage for a technology. These hidden complexities, or equivalently the type of relationships between patents that gives rise to these complexities, are called patent thickets.

However, there is another possible relationship between patents, which also gives rise to another type of patent thicket that destroys existing welfare. If the patent office makes mistakes then it might issue more than one patent, and so provide more than one patent-holder with exclusionary rights on the same underlying technology. In this case, the latter patents do not embody a novel or non-obvious inventive

[7] Firms do license the rights to use inventions while their patents are still in application. Such licenses are typically contingent on the grant of the patent.

[8] See Schumpeter (1934), p. 65. Schumpeter also defined innovation as "the setting up of a new production function", which could be taken to mean the carrying out of existing combinations in new ways. See Schumpeter (1939), p. 87.

[9] See Hall et al. (2012), figure 1, page 12.

[10] See www.uspto.gov/news/pr/2011/11-12.jsp. Recent data on the population of patent applications is available but has not yet resulted in new reports.

[11] Patents do not confer their holders the right to a monopoly, they confer the right to exclude. This exclusion does not extend beyond the scope of the claims of the patent, and there can be many different technological solutions (and so different inventive steps that merit patents) that can be used to achieve the same economic purpose.

step given the prior-art, and so should not have been issued. These 'spurious' patents do not add to welfare – the invention they cover had already been invented – instead they can be used for economically wasteful transfer-seeking. The patent explosion might also enable a feedback loop for spurious patents.[12] When the number of patent applications increases but the resources of the patent office remain essentially unchanged, the patent office must spread its examination efforts ever thinner. This might make obtaining a spurious patent easier, and spawn more applications for spurious patents. The greatest problem with the patent explosion therefore pertains to spurious patents and the thickets they create for genuine innovators. In all other cases, thickets might hinder the manifestation of potential new welfare, but they do not destroy existing welfare.

2.3 A Tangled Web of Thicket Definitions

Shapiro (2001) appears to the dominant reference for a definition of patent thickets used in literature. Of the 164 papers that we uncovered that provided a definition for a patent thicket, 26 (16%) quoted a metaphor from Shapiro (2001) verbatim, and 103 (63%) referenced Shapiro (2001) in their own definition. Shapiro (2001) appears to have coined the term 'patent thicket' and defined it as:

> "[A] dense web of overlapping intellectual property rights that a company must hack its way through in order to actually commercialize new technology" – Shapiro (2001)

Before we attempt to hack our way down to the economic fundamentals behind this definition, and in particular what it means for patents to be 'overlapping', a little etymology is instructive. The path to this thorny use of definition by analogy appears to have begun with Kiley (1992) who said:

> "[T]he path toward public possession of real benefit is increasingly obscured by dense thickets of intersecting, overlapping, and cross-blocking patents." – Kiley (1992)

Thus the origin of "patent thickets" began with a notion of social welfare, rather than the private cost to firms engaging in commercialization activities. However, the metaphor became truly tangled by Merges (1996), who twisted its analogy back to the problems of private costs in bringing a product to market:

> "This Article is aimed at providing conceptual guidance for those who need to traverse the new thicket of intellectual property rights. Each vine, each plant, standing in one's path represents a distinct IPR owned by an individual. To pass through, one needs a license from each owner. Where a single right blocks the path, this is easy: a single licensing contract does the trick. Today, however, business people more often than not encounter a tangled, twisted mass of IPRs, which crisscross the established walkways of commerce." – Merges (1996)

Apparently independent of this line of thought, Heller (1997) noted that an 'anti-commons' problem could arise with patents. The problem was the opposite of that of the commons where the lack of property rights leads to 'over-grazing' – usage that exceeds the social optima. With too many exclusionary rights a resource may instead be underutilized. Heller and Eisenberg (1998) identified "an unintended and paradoxical consequence of biomedical privatization: A proliferation of intellectual property rights upstream may be stifling life-saving innovations further downstream in the course of research and product development". Heller and Eisenberg (1998) revert the focus back to social welfare, and enjoin us to be equally wary of the problems patent rights can cause with respect to both invention and its commercialization. As we will later see, this is because Heller and Eisenberg (1998) offered two distinct definitions of patent thickets.

[12] Defensive patenting can also act as a feedback loop with other kinds of patent thickets. This is discussed later.

Hall and Ziedonis (2001) mirrored this double-sided warning in their article on the 'patent paradox' – why surveyed firms that claim not to use patents to appropriate returns to R&D show a dramatic increase in their propensity to file for patents. In the context of firms in sectors where products are complex and require many patented inputs, Hall and Ziedonis (2001) say "[s]uch firms typically require access to a 'thicket' of external intellectual property to advance technology or to legally manufacture and sell products, elevating concerns about patent-related hold-up problems."

Other early sources of the definition of a patent thicket cited in the literature include Lessig (2001) – a book on the 'future of ideas' – and Muris (2001) – remarks to the American Bar Association made by the chairman of the Federal Trade Commission (FTC) – but both cite Shapiro (2001) and the former cites Heller and Eisenberg (1998) as well. Thus it seems likely that notion of a 'patent thicket' was present in the collective minds of economic and legal scholars in 1990's but that its formalization was generally vague. Patents were envisaged as 'intersecting', as well as overlapping, and this somehow led to 'cross-blocking'; either complex technologies or cumulative innovation were to blame; thickets might have restricted invention or commercialization or both; and somehow hold-up was considered an element.

In this research we argue that there are actually four main types of patent thicket, each with separate economic foundations and, as a result, each with different implications for measures, tests and potential policy interventions. Traces of three of these types – diversely-held complementary inputs, legitimately overlapping patents, and spurious patents – can be found in the quotations provided above, and the fourth – effectively saturated invention spaces – appears to have emerged in the early 2000's. We report evidence of a general confusion over these definitions in the literature, which appears to still be growing, and feel that the use of 'definition by analogy' is the probable culprit. In the next section we define each type in turn, taking care to make their economic fundamentals explicit.

3 Taxonomy of Patent Thickets

3.1 Shapiro (2001)'s definition

Shapiro (2001) is a paper that demands a careful reading. As the effective seminal source in the economics literature for the definition of patent thickets it deserves special attention. Unfortunately, many authors, particularly non-economists, appear to have either miss-read this work or failed to grasp some of its nuance. The first source of confusion appears to stem from the use of the term 'overlapping'. Shapiro (2001) provides two definitions of a patent thicket:

> "...an overlapping set of patent rights requiring that those seeking to commercialize new technology obtain licenses from multiple patentees." – Shapiro (2001), Executive Summary, page 119

> "a dense web of overlapping intellectual property rights that a company must hack its way through in order to actually commercialize new technology. With cumulative innovation and multiple blocking patents, stronger patent rights can have the perverse effect of stifling, not encouraging, innovation." – Shapiro (2001), section I "The Patent Thicket", page 120

The word 'overlapping' is used six times in Shapiro (2001): twice in section headers, once in each of the two definitions above, and twice when referring to these definitions. From the context, however, it is clear that 'overlapping' intellectual property rights should mean that multiple patents bear on the same product, not that multiple patents cover the same invention. Shapiro (2001)'s thicket definition is that there are diversely-held complementary patented inputs.

> "The generic problem inherent in the patent thicket is well understood as a matter of economic theory, at least in its static version. Consider, for example, a company seeking to manufacture

a new graphics chip... Suppose that the company's preferred design for this chip is likely to infringe on a number of patents; the process manufacturing methods used to actually produce the chip infringe on a number of additional patents. In order to produce the chip as designed, the company needs to obtain licenses from a number, call it N, of separate rights holders." – Shapiro (2001), section II "Market Responses to Overlapping Patents", page 122.

Although we object to the use of the term 'infringement' to include a situation where a firm can (and should) obtain all necessary usage rights prior to commercialization of its product and so never infringe (i.e., use without permission) an exclusionary right, the sentiment in the quote above is clear. Shapiro (2001) describes a situation where the commercialization of a product requires the usage rights to patented inputs – the inputs are complements to the product – from N separate rights holders: the patented complementary inputs are diversely-held. This is made fully explicit in the simple model provided in the appendix of that paper.

The model envisages a situation where N firms own patents on (perfectly) complementary inputs. Each input provider has a cost of producing a unit of input of c_i, and will charge an endogenously determined price p_i. The price of the product itself will be p, and assembling a unit of product will cost α. Shapiro assumes perfect competition at the assembly (and so product market) level, which ensures that $p = \alpha + \sum_i p_i$. Demand for the product is $D(p)$, and the price elasticity of demand is given by $\epsilon = -\frac{D'(p)\cdot p}{D(p)}$

The N input providers set their component prices independently and non-cooperatively. That is, the model assumes that each input provider is a monopolist so it sets price to maximize profits, given by $\pi_i = D(p)(p_i - c)$. The first-order condition for each input provider is then:

$$\frac{d\pi_i}{dp_i} = D(p) + D'(p)(p_i - c_i) = 0 \tag{1}$$

Summing first-order conditions across all input providers, i, and substituting in $\sum_i p_i = p - \alpha$ we then get:

$$\begin{aligned} D(p)N + D'(p)\sum_i(p_i - c_i) &= 0 \\ \therefore \frac{\sum_i(p_i - c_i)}{p} &= -\frac{D(p)N}{pD'(p)} \\ \therefore \frac{p - \overbrace{\alpha - \sum_i c_i}^{c}}{p} &= \frac{N}{\epsilon} \end{aligned} \tag{2}$$

With a single firm, $N = 1$, the Lerner index ($L = \frac{p-c}{p}$) is $\frac{1}{\epsilon}$, so with N firms each providing a complementary input the mark-up is N times the standard monopoly mark-up. Thus, in a static context at least, Shapiro's patent thicket definition is concerned with a specific hindrance to commercialization, which we refer to as the N-fold marginalization problem. This problem is akin to the double-marginalization problem inherent in the Cournot complements model, where an upstream monopolist provides an input to a downstream monopolist.

It is not the number of patents that causes N-fold marginalization but instead the number of input holders. Each input holder might own many complementary patents, but they each (independently) make a single profit maximizing decision.[13] And a patent on the final product itself is simply another complementary input. If the patent on the final product is held by the firm commercializing the product, it does not count towards the N, and if it isn't then it does. These points will be important we describe

[13] This is analogous to a problem in cooperative game theory where each input holder can choose to provide either all of their complementary patents or not, as so has a single veto.

the measures that have been operationalized for patent thickets later.

The second source of confusion regarding Shapiro (2001)'s definition is particularly prevalent among non-economists. A surprising number of papers assert that this N-fold marginalization problem is intentional 'royalty stacking'.[14] In fact, the input providers are each caught in a Prisoner's Dilemma. The profit maximizing prices of inputs (p_i^m) are ones that sum, in conjunction with the assembly cost, to give a single monopoly mark-up. However, each input provider has perverse incentives. When they increase their price above p_i^m they get an additional benefit from their increased component price and suffer an additional loss from the resulting decrease in demand for the final product. Unfortunately, they accrue all of the benefit and distribute a part of the loss (providing $N > 1$) as a negative externality on the other input providers.

Coordination mechanisms are therefore key to prevent this type of patent thicket from causing sub-optimal commercialization of new products. Without coordination of some kind, the N-fold marginalization reduces the quantity of new goods provided to consumers, creating a large deadweight loss. Likewise, the rents to invention are also compromised – rather than earning the optimal rents for their inputs, the returns to input holders drop precipitously as N increases. Shapiro focuses on cross-licensing agreements, patent pools and cooperative standard setting as potentially effective mechanisms to mitigate the problems of this type of patent thicket, and stresses that anti-trust authorities should be cognizant that cooperation in this context will increase and not decrease consumer and total welfare.

> "Unfortunately, antitrust enforcement and antitrust law have a deep rooted suspicion of cooperative activities involving direct competitors. But such cooperation in one form or another may be precisely what is required to navigate the patent thicket. As a result, unless antitrust law and enforcement are quite sensitive to the problems posed by the patent thicket, they can have the perverse effect of slowing down the commercialization of new discoveries and ultimately retarding innovation, precisely the opposite of the intent of both the patent laws and the antitrust laws" – Shapiro (2001), section I (The Patent Thicket), page 122.

We concur. However, unilateral mechanisms should not be neglected. There are three standard unilateral approaches to resolving the Prisoner's Dilemma and each merits discussion. The first is credible commitments: firms make prior commitments that they are either unable to or would not want to renegotiate afterwards. Anecdotal evidence suggests that Fair/Reasonable And Non-Discriminatory (F/RAND) licensing commitments are particularly common in industries with complex products – those that require many complementary inputs. FRAND commitments are good examples of credible commitments. The second and third approaches require repeated games. With a positive chance of repetition of the game every time (or equivalently an infinitely repeated game with a non-zero discount factor), a firm can establish a reputation for playing the cooperative action and potentially induce cooperation by the other party; and punishment mechanisms, such as tit-for-tat or the grim trigger, can also be used to maintain cooperation. These mechanisms could work with each input provider solely as a technology developer but will likely be more effective when firms that provide inputs also commercialize products, as then the punishment for deviation can be higher. Thus, even absent overt coordination mechanisms like patent pools, it seems possible that firms might be able to solve the N-fold marginalization problem inherent in assembling diversely-held complementary inputs.

One way in which patent thickets are said to hinder innovation is by creating barriers to entry. It is entirely unclear how this relates to this definition of patent thickets: Shapiro's model supposes a fully competitive product market with no fixed costs where input holders are eager to license their patents and the sole hitch is in their coordination. Instead, we might be concerned that entrants might magnify the

[14] A small minority of papers, all written by non-economists, asserted that input providers want to charge as high a price as possible for their inputs, implicitly confusing high prices with high profits.

thicket problem. Aside from adding to the N, entrants might play a one-shot game while incumbents use reputations or punishment mechanisms to keep input prices optimal. We will return to issues regarding entrants, their role in the patent ecosystem, and whether a reduction in entry should be considered a reduction in innovative activity, in the measures and tests section later.

The third source of confusion regarding Shapiro (2001)'s definition concerns whether or not the patents are spurious in some fashion, particularly regarding whether all of the patents covering complementary inputs meet the requirements for novelty, usefulness and non-obviousness laid out by the patent office. We believe that Shapiro intended issues of validity to be separate from the definition of a patent thicket, consistent with the first of the two quotes below, but understand that the second quote may have given other researchers a different impression.

> "But concerns about a patent thicket, **and** excessively loose standards at the PTO, are hardly confined to e-commerce and business method patents." (emphasis added); and "This paper takes as given the flood of patents currently being issued by the PTO, and assumes that these patents are indeed creating a patent thicket in the sense that many new products would likely infringe on multiple patents." (emphasis added) – Shapiro (2001), section I (The Patent Thicket), page 121.

The final major source of confusion regarding Shapiro (2001)'s definition concerns hold-up. We will discuss hold-up in detail next, but first we assert that Shapiro did not intend hold-up as a part of his definition.

> "The patent thicket is especially thorny **when combined with** the risk of holdup, namely the danger that new products will inadvertently infringe on patents issued after these products were designed." (emphasis added) – Shapiro (2001), Executive Summary, page 119

3.2 Hold-up

The concept of hold-up originates with Williamson (1971).[15] The term 'hold-up' has a very precise definition in economics. Hold-up requires: 1) an incomplete contract between two-parties; 2) a transaction-specific investment by at least one party; and 3) opportunism (essentially moral hazard) the other party. Contracts are incomplete because it is either prohibitively expensive to contract over all of the possible states of the world or not all future contingencies can be anticipated. Once the contract has been made, one party to the contract undertakes an investment that is dependent on the contract. An example might be to invest in research and development of a product within a joint venture. Then, after this transaction-specific investment has been made, the other party engages in opportunism to appropriate some or all of the transaction-specific investment. Williamson (1979) defines opportunism as follows:

> "Opportunism is a variety of self-interest seeking but extends simple self-interest seeking to include self-interest seeking with guile. It is not necessary that all agents be regarded as opportunistic in identical degree. It suffices that those who are less opportunistic that others are difficult to ascertain ex ante and that, even among the less opportunistic, most have their price." – Williamson (1979), footnote 3, page 234.

The precise mechanism for hold-up is somewhat complicated. Imagine that the two parties can contract over states of the world $\omega \subset \Omega$, but must leave states $\omega' \subset \Omega$ uncontracted. If a state $s \in \omega'$ occurs, the contract is void and will be renegotiated afresh. With a contract over states ω in place, at least one firm makes a transaction-specific investment. This is an investment that depends upon the continuation of the contract terms. It must have a lower value, sometimes called its salvage value, outside

[15] Williamson (1971) describes hold-up but Goldberg (1976) apparently coined the term. See Klein et al. (1978) and Williamson (1979) for other seminal references.

of the contract than inside it. After this investment has been made, the other party to the contract then engages in opportunism. This party claims that a state s has occurred and that the contract is void. The occurrence of such a state s is unverifiable by the invested party, or by the courts.[16] The opportunistic party then renegotiates the contract, 'holding-up' the other for its transaction-specific investment. The held-up party can recover only its salvage value for the relation-specific investment if a new contract isn't created. Thus the new contract 'creates' the difference between the salvage value of the investment and its full value, and this difference can be (at least partially) appropriated by the opportunistic party.

An example from the literature is instructive. Iyer and Schoar (2010) purchased pens from 182 wholesalers in India. Some pens were plain but others were branded with custom designed logos. The branding of the pens constituted a transaction-specific investment in the contract by the wholesalers – branded pens were determined to have essentially no salvage value, whereas unbranded pens could be resold at their full price. The pen purchasing contract did not specify what would happen in the state of the world where the shoppers hired by the researchers couldn't (or at least claimed they couldn't) pay for the pens. The shoppers then claimed that they unexpectedly couldn't pay the full the amount for the pens, and so engaged in opportunism and held-up the wholesalers.[17]

The term 'patent hold-up' appears to have originated with Skitol (2005), Lemley and Shapiro (2006), Lemley (2007), Shapiro (2010), and Farrell et al. (2007).[18] According to these papers, the underlying requirements for patent hold-up, as opposed to Williamson hold-up, are: 1) an irreversible investment in the development of a product or invention; and 2) the 'surprise' revelation of pre-existing patent rights that bear on the product or invention. These papers also claim that the problems associated with patent hold-up are then amplified by probabilistic patents (or outright spurious patents), injunction threats (particularly in conjunction with costly and time-consuming redesign), patents on minor features, and the inclusion of the certain patent rights in standards.

Historically, the term 'transaction-specific' investment was used in the hold-up literature, but 'relation-specific' investment is now more frequently seen. This subtle shift in nomenclature helps when we try to reconcile patent hold-up with Williamson hold-up. The contractual component of Williamson hold-up is lacking from patent hold-up, but this could be addressed in two ways. First, one could take the view that a lack of a contract is an incomplete contract; making an investment that is dependent on a relationship with another party, without a contract with that other party, does intuitively expose one to their guile. Second, the patent office might make the contract through the issue of a patent, allowing the patent-holder to selfishly exert, or threaten to exert, their exclusionary right. Generally, though, there is no guile involved: in this second case, there was no agreement that was unilaterally changed through dishonesty, instead there is usually just the (honest) exertion of rational self-interest following the surprise revelation of the existence of a relationship. Submarine patents and guile within an RTLA provide specific exceptions to this general rule. 'Patent ambush' through the manipulation of an SSO and 'abuses' of FRAND commitments, might also constitute guile.[19] These are discussed shortly.

The market failure from hold-up comes not from the act of hold-up itself, but rather from the inefficiently low levels of investment in the shadow of potential hold-up. Hold-up itself is a transfer from the invested party to the opportunistic party. Transfers are not a source of economic inefficiency, except in so far as resources are wasted in their pursuit (as in the case of type C thickets). Economic activity forgone

[16] It can also suffice for it to be prohibitively costly to verify.

[17] It is worth noting that the wholesalers anticipated the hold-up, at least in a probabilistic sense, by demanding a deposit for branded pens but not for plain pens; that they appeared to understand that they were being held-up when it happened; and that they did accept a lower total price for the custom pens than the plain pens after they were held-up.

[18] In approximate order of origination, not publication. ABA (2004, pp.60–64) is the first published source to which the semantics of the term are 'patent hold-up' attributed. However, ABA (2004) is concerned with 'patent ambush' and says: "[In patent ambush cases] it has been alleged that a firm improperly induced [a Standard Setting Organization] to adopt a standard incorporating the firm's intellectual property without properly disclosing that the firm owned these intellectual property rights. It is alleged that such conduct [has conferred] substantial market power."

[19] See footnote 18 for the definition of patent ambush.

because of a risk of appropriation, on the other hand, leads to non-Pareto allocations of resources. This distinction sheds light on a fundamental misunderstanding in papers advocating a policy response to patent hold-up. Many of these papers redefine patent hold-up as the difference between an (usually hypothesized) ex-ante (to the relation-specific investment decision) price and the demanded ex-post (again to the relation-specific investment decision) price of a technology license.[20] Their suggested policy responses typically include compulsory licensing, restrictions on the use of injunctions, the implementation of 'incremental value tests', and so forth, in attempts to undo 'patent hold-up'.

> "'Hold-up' is used throughout this report to describe a patentees ability to extract a higher license fee after an accused infringer has sunk costs into implementing the patented technology than the patentee could have obtained at the time of design decisions." – Commission (2011)

In this paper we argue that patent hold-up, as currently defined in the literature, is lacking a key ingredient: guile.[21] The 'surprise' revelation of a patented input might create a similar market failure to hold-up, but does not merit the same policy responses. Absent guile, investment decisions under uncertainty can result in inefficient levels of investment: if the firms suffer from over-confidence bias, experience disproportionate disutility of losses relative to gains as in Kahneman and Tversky (1979)'s prospect theory, or otherwise exhibit behavioural economic deviations from rationality, then inefficiently high or low levels of investment could result.[22] The key intuition is simple. A rational firm makes estimates of the likelihood unsecured patent rights will bear on its product or invention, as well as the costs of securing or infringing these rights, and proceeds (optimally) accordingly.[23] Inefficient levels of investment can result from either deviations from rationality or guile by another party.

Information asymmetries and a lack of verifiability are necessary but not sufficient conditions for Williamson hold-up. The opportunistic party needs to be able to claim that a state of the world that nullifies the contract has occurred and the invested party must be unable to demonstrate that this isn't the case. However, it is the guile – essentially an act of moral hazard – that is the sufficient condition for Williamson hold-up, and this carries through to patent hold-up. Patent-based relationships can certainly be subject to both information asymmetries and a lack of verifiability, but absent guile this doesn't have anything to do with patent hold-up. For example, a potential licensee may know less about the quality of patent rights than the licensor, which might make reaching mutually agreeable terms difficult, but unless a surprise revelation about the quality of rights arises from an action taken by the licensor (i.e., guile) after the contract has been formed, and a relation-specific investment has been made, there is no patent hold-up.

Information asymmetries concerning the existence of patent rights and 'fair' terms for their use deserve particular attention. In 'patent ambush' a firm deliberately induces an SSO to include features protected by its patents without disclosing that it has such patents. An implementer then faces an information asymmetry regarding the existence of patent rights that was deliberately induced by the patent-holder's guile. This is possible because the SSO or implementers can't verify whether or not the failure to disclose the holding of relevant patents was deliberate. Given the nature of standard setting processes, and that the participants in standards are involved in a repeated game that allows effective punishment for this

[20] Numerous papers have pointed out grave flaws in this reasoning. Epstein et al. (2012), for example, very cogently warns against both the "non-standard and misguided definitions of economic terms of the art such as 'ex-ante' and 'hold-up'" and the imposition of a new regime of "government hold-up replacing private coordination".

[21] A full discussion of patent hold-up is beyond the scope of this paper.

[22] For over-confidence bias see, for example, Kahneman and Tversky (2000) and Moore and Healy (2008). A discussion of the foundations for behavioural economic deviations from rationality is provided in McFadden et al. (2000).

[23] When a petroleum firm decides to drill an oil well but comes up dry, the firm hasn't been held-up. Such a firm made a decision under uncertainty. Sometimes the outcome will be good and sometimes it will be bad but, absent deviations from rationality, making a decision to drill based on an assessment that expected benefits will out-weight expect costs, is still correct and economically efficient.

kind of behaviour, we do not think it likely that this actually occurs in practice, but we acknowledge that it is theoretically possible.

Fair/Reasonable And Non-Discriminatory (FRAND) licensing is also frequently associated with hold-up. With FRAND commitments it is possible that there is an information asymmetry regarding 'fair or reasonable' terms for the use of intellectual property rights. The usual argument is that a firm has made a FRAND commitment to license on some set of terms and then demanded an alleged non-FRAND set of terms when it discovers infringement, and so engaged in hold-up. The FRAND commitment might be thought to act as the contract and the alleged guile comes from the (unverifiable) change in terms.[24] But FRAND commitments should not entitle every licensee to the same terms, just terms that are 'fair or reasonable' and non-discriminatory. Non-discrimination means that a firm can not set licensing rates (or refuse to license) in an anti-competitive manner. That is, the firm can not unduly raise a rival's costs, create barriers to entry, or otherwise engage in actions that are illegal under anti-trust law. This is distinct from hold-up. Thus, in the context of a FRAND commitment, a charge of hold-up hinges on a change to FRAND terms that makes them unfair or unreasonable.

It can be fair and reasonable to price licenses differently at different times. Early adopters might be given preferential rates to encourage their usage, and once demand has been established rates might rise (or fall) for subsequent adopters. More importantly though, to an economist it is fair and reasonable to allow firms to set rates that punish those that abuse its intellectual property rights. Demonstrating willful infringement can be very difficult, so an infringed patent-holder might not be able to rely on punitive damages. The rational response may be then to price licenses after infringement has taken place at a multiple of the price that would have been offered prior to infringement, taking into account the probability of detection of infringement, the likelihood of success in prosecuting an infringement claim, and so forth.[25] This is would not be a violation of the FRAND terms, and would not constitute guile – a rational infringer would know this and act accordingly.

The lack of verifiability concerning the unlicensed use of a patent-holder's rights creates the potentially for both 'reverse hold-up' and for 'government hold-up'. Suppose that policy to 'correct patent hold-up' were put in place. For example suppose that injunctions were forbidden and that the same licensing terms would be imposed regardless of whether or not the 'infringing' firm undertook an investment to make a new product or not. Further assume that the transaction costs of licensing are approximately the same whether they are conducted privately or through the courts. The infringing firm then has no incentive to search for patented inputs or engage in private ex-ante (i.e., prior to infringement) licensing negotiations. It would rather infringe and hope that it isn't caught. In the best case it gets to use the patented input for free and in the worst case it is in the same situation as it would have been 'ex-ante' but without incurring the search costs. This problem is amplified if patents are probabilistic. Moreover, to avoid charges of willful infringement, the infringing firm must engage in guile! A relation-specific investment is being held-up – it is the R&D investment to create the patent in the first place. In the shadow of this potential reverse-hold-up, such R&D investments might not take place as the rents to invention are eroded.

The policy to 'correct patent hold-up' might also constitute government hold-up. The development of a patented invention constitutes a relation-specific investment with the government: The government has essentially entered into a contract that assures the future patent-holder the right to unilaterally dictate commercial terms, within the legal limits imposed by anti-trust law, for the usage of their invention. If

[24] In practice the courts can and do demand that firms provide evidence that the rates that they are demanding are comparable to rates already secured. It therefore seems questionable that a change in terms would be unverifiable and that patent hold-up really occurs with FRAND commitments.

[25] Absent 'exceptional circumstance', patent plaintiffs can not recover their legal costs. Furthermore, a judgement of willful infringement entitles the plaintiff to up to three times the assessed damages. This may not be sufficient to deter infringement if the likelihood of detection and prosecution are sufficiently low. See: 35 U.S.C. 285 and 284, respectively.

the government opportunistically changes the terms of the contract to extract a portion of the inventor's rents then this would be government hold-up.[26]

A material fraction of papers in the thicket literature either implicitly or explicitly claim that hold-up is a fundamental attribute of patent thickets (see table 2). Regardless of whether or not patent hold-up has a sound theoretical basis, the issue at hand is how it might relate to patent thickets. We will discuss this for each type in turn in the following sections.

3.3 Heller and Eisenberg (1998)'s first definition

Heller and Eisenberg (1998) provided the seminal definition of patent thickets for the law literature in much the same way as Shapiro (2001) provided it for the economics literature. Heller and Eisenberg (1998) is also somewhat prone to misstatement or confusion, perhaps because as legal scholars the authors were not concerned with providing a formal mathematical model that would render the semantics unambiguous, but more likely because they actually provided two explicitly different definitions. In this section we will focus on the first definition, which has greater synergy with the underlying theory of their paper. This definition is based on Heller (1997)'s observation that patents might be subject to an anti-commons problem – a situation where there are many rights to exclude that bear on a resource leading to under-utilization.

> "...a resource is prone to underuse in a 'tragedy of the anti-commons' when multiple owners each have a right to exclude others from a scarce resource and no one has an effective privilege of use." – Heller and Eisenberg (1998), page 698.

Heller and Eisenberg (1998) argue that an anti-commons (i.e., patent thicket) problem may arise as a result of too many 'concurrent fragments'. They say that:

> "A proliferation of patents on individual fragments held by different owners seems inevitably to require costly future transactions to bundle licenses together before a firm can have an effective right to develop these products" – Heller and Eisenberg (1998), page 699.

Thus, Heller and Eisenberg (1998)'s first definition appears to mirror Shapiro (2001)'s. Heller and Eisenberg (1998) also mention 'overlapping' patent rights, but it is clear from the paper that these are not intended to be a part of the 'concurrent fragments' definition, rather they "aggravate the problem of concurrent fragments".

3.4 Type A: Diversely-held Complementary Inputs

In this sub-section and the sub-sections that follow we will describe four different types of patent thicket that are consistent with definitions used in the literature. We name these types A. B, C and D. Where possible we will divide these types into sub-types and variants, creating a taxonomy of thicket definitions with inherited attributes. Our understanding is that Shapiro (2001)'s patent thickets and Heller and Eisenberg (1998)'s anti-commons ('concurrent fragments') problems are one and the same. Accordingly, we therefore declare patent thicket type A as:

Patent Thicket Type A: *Products require many diversely-held complementary inputs which are protected by discrete patents. Patents are valid and correctly issued by the patent office.*

By 'discrete patents' we mean that patents are not overlapping in the domain of their application, so that no two patents make claims on the same thing, and by 'valid and correctly issued' we mean that

[26]The change could be opportunistic if, for example, it was the result of lobbying.

patents cover novel, useful and non-obvious inventions and this has been correctly certified by the patent office.

Type A thickets are immediately classifiable into two sub-types:

Sub Type Aa: *Patented inputs are perfectly complementary, so that the full set of inputs must be assembled.*

Sub Type Ab: *Patented inputs are imperfectly complementary, so that varying degrees of functionality can be acheived with certain sub-sets of inputs.*

As we will later show, authors in the literature seldom make a distinction between these two sub-types but it remains useful. In multi-lateral bargaining theory there are two distinct situations: those which require unanimity and those where partial accords are possible. The second situation arises when the incentive condition for a complete agreement for at least one participant is less than the incentive condition for a partial accord, so that they would prefer to 'hold-out'. As such, Farrell (2009) notes that type Ab patent thickets can be subject to a hold-out problem.

3.5 Thicket Variants

Four variants to type A thickets are possible; they are also applicable to varying degrees to type B, C, and D thickets, so we discuss them now.

3.5.1 Variant C (complete information)

We begin with variant C – a hypothetical variant that we introduce as a theoretical benchmark.

Variant C (complete information): *All information about patented inputs is known to the producer. Search and transactions costs are not considered (or are assumed sufficiently low that they can be ignored).*

Even with variant C, where all information about patented inputs is known to the producer, there can still be information problems that prevent the producer from commercializing the product. Specifically it is remains possible that all of the bargaining parties (producer and input providers) might not know the demand for the product, resulting in different expectations of the gains and hence their fair division. Given that the product is innovative and new, its demand is surely difficult to accurate forecast.

Moving away from variant C, information problems begin to increase dramatically. Even if the producer knows of the existence and scope of all exclusionary rights, they are unlikely to know each patent-holder's marginal cost of providing their input. The marginal cost of a patent license is seldom zero. Transaction costs (the cost of enacting the licensing agreement) aside, the marginal cost of a patent license is the patent-holder's marginal opportunity cost for that license. As this is likely unobservable to the producer and other patent-holders, and as the patent-holder may not be able to convey it credibly, there is no way to enforce optimal pricing. Even agreeing to a $\frac{1}{N}$ share of the proceeds from the commercialization of the product does not resolve this problem – some products may still not be commercialized at all. As Farrell (2009) concisely puts it: "A symmetric draft agreement that divides the gains equally will command unanimous assent only if the gains per participant exceed the maximum among the N participants costs of participation. This order statistic is likely to be well above the average, and agreement may very well fail even though the total gains easily exceed N times the average cost of participation."

3.5.2 Variant T (transaction costs)

Variant T thickets, which consider transaction costs to be a part of the thicket definition, are very commonly used in the literature. Many of the early papers on patent thickets included a variant T

compliant definition, and this was particularly common in conjunction with a type A definition. Heller and Eisenberg (1998), Walsh et al. (2003), Clarkson (2004), Ziedonis (2004), and others, were all explicit that transaction costs were fundamental to a type A thicket definition, and Shapiro (2001) noted that transaction costs are a possible issue with type A thickets.

Variant T (transaction costs): *All patented inputs are known to the producer. Transactions costs are considered material and may prevent successful negotiation of usage rights for patented inputs.*

We do not know of any empirical work estimating the average transaction costs involved in licensing a patent, but it seems conceivable that the average legal fees and other costs of time might be on the order of $10,000. There are reports in the literature that some products require several thousand patented inputs – it therefore seems conceivable that transactions cost can become prohibitive, perhaps especially for small entrants.[27]

Transaction costs can be reduced through cross-licensing agreements, the creation of patent pools, and other mechanisms, and firms have incentives to minimize them. However, it is far from clear that we could create policy to further reduce them; the most frequently made applicable policy recommendation is for anti-trust authorities to adopt a 'hands-off' approach to the private approaches, which may involve coordination.

3.5.3 Variant S (search costs)

Variant S (search costs): *The number of patented inputs is sufficiently high that search costs become prohibitive or cognitive biases prevent complete identification. Some subset of input patent usage rights may be unsecured.*

In the context of information problems, Farrell (2009) makes the point that producers might not know with whom they must negotiate concerning patents (i.e., which patents actually bear on their product). His paper refers to this as a "potential patent thicket". This hints the larger problem of search costs. Setting aside cognitive limitations of the individuals doing the search for a moment, it does seem likely that searching for relevant patents will impose material costs. Duplication of search can create an economic inefficiency. Every firm that makes a competing product from the same set of inputs must search independently for those inputs. Again, mechanisms like cross-licensing agreements and patents pools, as well as industry standards, can all reduce search costs. However, this is one area that policy could be effective: the patent office is the central repository for the data used in searches, and enhancing this data with information about where patents are used, which other patents they are typically used with, and so forth, could materially reduce search costs.

A firm considering investment in the development of a new product has three basic considerations: 1) the cost of searching for and of securing required patent rights, such as usage rights to complementary input inventions covered by patents; 2) the costs in delays to production, particularly those resulting from injunctions, if not all necessary patent rights are secured and the infringement is discovered and the rights are upheld; 3) the likelihood that infringement will be discovered and the rights will be upheld. The cost of securing rights ex-ante of the investment as opposed to ex-post of the investment are complicated by whether or not the infringement could be proved as willful, as well as issues of reverse patent hold-up. Patent law makes a strict distinction between willful and non-willful infringement of patent rights; willful infringement is subject to up to treble damages.[28]

[27] See Callaway (2008) and Lemley and Shapiro (2006) for discussions of the volume of patented inputs required in complex products. Goodman and Myers (2005) reports that 7,796 patents were declared essential to two standard third-generation cellular phone technologies.

[28] See 35 U.S.C. 284. According to Lemley and Shapiro (2005), 'willful infringement' occurs only when an infringer is aware of the patent and believes the patent is valid and believes that its conduct infringes. Proving these three things seems a high hurdle to cross.

3.5.4 Variant P (probabilistic patents)

Variant P (probabilistic patents): *Patents do not have guaranteed validity and incontrovertible rights to exclude usage to the inventions detailed in their claims. Instead patents are always potentially subject to being ruled invalid in general or inapplicable in some specific application. The courts are the ultimate arbitrators of patent rights, and may decide rights on a case by case basis with imperfect judgment.*

Lemley and Shapiro (2005) argue that a patent "does not confer upon its owner the right to exclude but rather a right to *try* to exclude by asserting the patent in court" (emphasis in original). Set against the backdrop of the 'patent explosion' and in the context of (predominantly) jury-based trials where laymen will make imperfect decisions, Lemley and Shapiro (2005) claim that patents are tantamount to lottery tickets. Invalid patents (discussed under type C thickets) have some probability of being held valid, and a valid patent that carefully and clearly articulates even the most fundamentally sound of inventions might be held invalid. All that varies is the probability of successfully upholding a patent in a specific suit, famously labelled as $\theta \in [0, 1]$ in Lemley and Shapiro (2006).

The probabilistic nature of patents bears upon search costs in two ways: In enters the calculus faced by a firm directly; and the detectability of infringement is related to the likelihood that a patent suit would be upheld. To the extent that parties have different beliefs about the strength of patent rights, it could also bear upon transaction costs.

3.5.5 Defensive patenting

For type A patent thickets, many authors have claimed that the extent of the thicket and the danger of potential patent hold-up are linked. Their argument is that the more patented inputs that a product manufacturer requires, the more likely it is that one might come as a surprise that is somehow based on the patent-holder's guile. FRAND commitments might be used to prevent N-fold mark-up problems while at the same time facilitating patent hold-up. We think this unlikely.[29] But another commonly postulated strategic response, which is supported by some empirical evidence, called 'defensive patenting' can mitigate both the adverse effects of the thicket and the patent hold-up problem.

Defensive patenting is the intentional development of inventions, and the filing of patents to protect them, which are likely to be needed by rivals, in order to protect a firm's commercialization of its own products. The literature frequently refers to defensive patents as 'bargaining chips'. They act as potential punishment mechanisms to discipline the behaviour of the owners of patented complementary inputs. The desired behaviour could be either the correct pricing of inputs (i.e., avoiding the N-fold marginalization problem) or a lack of hold-up. This is evident in the two quotes below.

> "One way to cut through the patent thicket is for incumbents with extensive patent portfolios to enter into broad cross-licenses (that is, exchanges of roughly symmetric patent positions) to 'clear' the thicket. However, new entrants who lack large patent portfolios may be at a major disadvantage in this situation because they have no patents to trade. Without such cross-licenses, the result is [the N-fold marginalization problem]. Defensive patenting is a natural, even inevitable, strategy in industries with patent thickets, but defensive patenting itself can increase the density of the thicket." – Lemley and Shapiro (2005)

> "We examine whether firms most subject to 'holdup' responded strategically to the shift in the U.S. legal environment by patenting more aggressively... to safeguard against the threat of costly litigation and to negotiate access to external technologies on more favorable terms." – Hall and Ziedonis (2001) describing their "strategic response" (i.e., defensive patenting) hypothesis.

[29] We find the possibility of reverse hold-up and government hold-up in FRAND commitments far more compelling.

3.6 Welfare Considerations for type A thickets

Welfare might be increased or decreased by type A thickets depending on how firms respond to them. This will have consequences for measurement that we discuss later. In the simplest case, the more diversely-held the required patented inputs, the greater the N-fold marginalization problem, and the lower quantity of the resulting product provided to consumers. The deadweight-loss caused by the multiple mark-ups is proportionate to the extent that patented inputs are diversely-held. It is not affected by the number of inputs required directly. However, even this simplest case hides a problem of dynamic efficiency. It is superior, in the static analysis, to have all of the inputs held by a single firm so only a single monopoly mark-up will be imposed. But this would change the incentives to innovate. Aghion et al. (2005) provides evidence (and an accompanying theoretic model) that the relationship between competition and innovation takes the form of an inverted U.

Extending the simplest case slightly, we might suppose that for some products transaction and search costs render commercialization infeasible. We would therefore also expect an economic loss from the forgone products, though mitigating this loss is difficult and comparison to a first-best world without transaction costs is unrealistic. Likewise, if patents are probabilistic the commercializing firm would face greater uncertainty as well as potential legal costs in determining which patents will be held valid and infringed. When commercialization becomes potentially infeasible, a rational inventor may prefer not to invent. Invention and the patenting of invention are both costly activities. In the presence of a type A thicket, an inventor may anticipate that the compensation from commercialization will never materialize.

Next, firms may use mechanisms other than defensive patenting to respond to type A thickets. When patent-holders are also implementers, cross-licensing agreements are the likely the most common. These could resolve the N-fold marginalization problem, with comparatively low transaction costs, but could also act as a barrier to entry for new inventors wanting to contribute to the current generation of technology. Entrants could still undertake radical innovation that cleared away the thicket; and would have incentives to do so. The displacement of entrants from incremental to radical invention may or may not reduce welfare. When patent-holders aren't implementers, patent pools (and other technology market solutions) could also resolve N-fold marginalization problems. Entrants could then buy the necessary intellectual property rights from the patent pool, and any adverse welfare consequences of the thicket could be mitigated.

When firms use defensive patenting to respond to type A thickets, thickets could increase welfare. A patented input could have many different applications, and under Schumpeterian innovation, the more new things there are to combine, the more new combinations become possible. It may matter to welfare who invents something. Invention costs may differ between firms, and ideally firms with lowest costs would incur them, but firms will also have differing incentives to allow other firms to commercialize inventions. Perhaps more importantly though, when patent-holders are also implementers, they have incentives to invent around existing patents and create economic substitutes. When combined with defensive patenting, substitute inventions would lead to a competitive input market, lowering costs to non-patent-holding implementers as well.

3.7 Type B: Legitimate Overlapping Patents

Although we feel it likely that many authors who defined a patent thicket as being based upon 'overlapping' patents did so as a result of a misunderstanding of Shapiro (2001), 'overlapping' patents are indeed at the heart of a genuine patent thicket definition. Heller and Eisenberg (1998)'s second definition of a patent thicket was based upon 'stacking licenses'. They say:

> "The use of reach-through license agreements (RTLAs) on patented research tools illustrates another path by which an anti-commons [patent thicket] may emerge. As we use the term, an

RTLA gives the owner of a patented invention, used in upstream stages of research, rights in subsequent downstream discoveries... In practice, RTLAs may lead to an anti-commons as upstream owners stack overlapping and inconsistent claims on potential downstream products."
– Heller and Eisenberg (1998), page 699.

We adopt their definition as type $B1a$ below, and broaden their definition to provide a definition of type B patent thickets that is consistent with almost half of the definitions implemented in the literature.

Patent Thicket Type B: *The claims of two or more patents have overlapping domains of application. All patents are valid and correctly issued by the patent office. Patents may be characterized as 'broad'.*

Type B thickets pertain to invention, not commercialization. With type A thickets, a problem arose when a firm wanted to create a product that required many diversely-held complementary input patents, but with type B thickets, a potential problem can arise when a firm wants to create a new invention and patent it. The type B thicket problem is based on the notion of overlapping claims, not overlapping inventive steps. With type B thickets we assume that the patent office does not make mistakes and issue patents that aren't novel or obvious given the patented prior art. However, the same economic functionality, which may or may not be protected by patents' claims over a domain of application, can often be achieved with different underlying technologies. At the risk of introducing a new analogy, this is akin to the well-known phenomena of paradigms in science. Isaac Newton (or Gottfried Leibniz, if you prefer) provides us with classical physics that can be used to describe the motion of moving bodies, then Albert Einstein (and others, again according to taste) provides us with relativistic physics that can also be used to describe the motion of moving bodies. When the bodies are moving at non-relativistic speeds, Newton's and Einstein's theories provide predictions over the same domain of application, albeit with markedly different 'inventive steps'.

'Broad' patents – those that cover a large domain of application – have been a concern for many authors, particularly in areas that are close to basic research. While overlapping patents do not need to be broad, greater breadth makes overlaps more likely. However, awarding considerable breadth to a patent is, in and of itself, not necessarily a bad thing. Ideally a patent's breadth of protection should match its breadth of inventive step. This is illustrated in figure 2 below. Patent that fall above the $45°$ line are awarded more protection than they merit from the scope of their inventive step, and conversely those that fall below the line are awarded less. Thus unmerited breadth, or unmerited narrowness, should be a concern, and not absolutely breadth or narrowness.

The patent office might have good reason to deviate from affording a patent-holder the 'correct' level of protection depicted on the $45°$ line. Scotchmer (1991) argues that 'foundational' patents, those giants on whose shoulders many others will stand, are likely to be undercompensated. This is particularly true when a foundational patent has little commercial value in itself, but enables a large amount of highly valuable follow-on invention. One normative mechanism to rectify this situation would be to allow such patents a broader claimed domain of application.

We decompose type B patent thickets into two main sub-types: $B1$ and $B2$. In $B1$ patent thickets the subsequent invention is cumulative to the first invention, whereas in type $B2$ patent thickets, the subsequent invention is adjacent to the first invention. As we will shortly see, this distinction leads to diametrically opposite economic relationships between the two patents.

Patent Thicket Type B1: *Invention is cumulative and an upstream patent's claimed domain of application overlaps on a downstream patent's claimed domain of application.*

Cumulative invention is particularly common in areas close to basic research. This naturally includes patents that are generated as a result of academic research. A sub-literature on the anti-commons problem, epitomized by Walsh et al. (2003) and Murray and Stern (2007), has studied the extent to which

Figure 2: Unmerited breadth or narrowness of patent claims

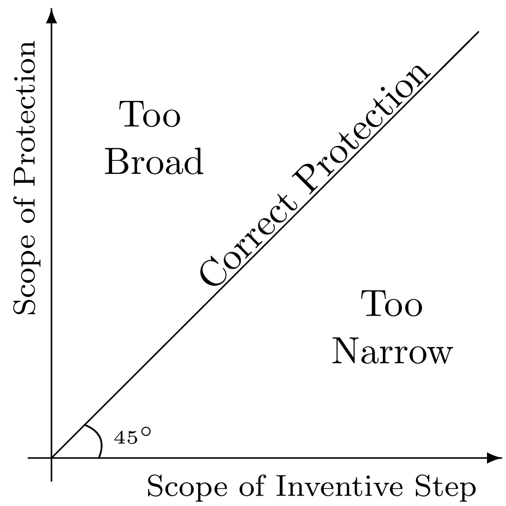

patent thickets influence academic research.[30] It also includes the biotechnology industry, particularly the genetics sector, where firms engage in science in order to create technology. However, cumulative invention happens in many sectors of the economy.

We further distinguish between two variants of type $B1$ thickets, according to the mechanism used to control the downstream invention. In type $B1a$, a Reach-Through Licensing Agreement is used. RTLAs allow upstream firms to capture some portion of the benefits of downstream research based on their upstream inventions. These can be highly controversial as the license may afford rights beyond the scope of the original upstream patent claim.[31]

Sub Type B1a: *Flow-through rights are made explicit in a Reach-Through Licensing Agreement (RTLA). The downstream party licenses the upstream patent prior to originating the downstream invention, and assigns rights back upstream accordingly.*

Sub Type B1b: *The downstream party licenses the upstream patent without a Reach-Through Licensing Agreement (RTLA), or does not license the upstream patent.*

Legitimate patents with overlapping claims in the context of cumulative invention create a complementary relationship. This can happen in three ways:

1. Improvement patents: The downstream patent refines the upstream patent in some manner.
2. Materials input: The upstream patent covers material or functionality that is embodied in the downstream patent.
3. Research tools: The upstream patent covers research tools that enable the creation of the invention embodied in the downstream patent.

In each case, an upstream patent is a complementary input into a downstream patent. However, arguments concerning whether or not an RTLA is merited are subtly different in each case. Improvement patents build off the inventive step of the upstream patent and so could stand alone. Yet they are making some invention or inventions better. The upstream patent holder was well placed, technologically, to make this improvement, but for some reason didn't. This suggests that the upstream patent-holder may not

[30] Both of these papers define the term patent thicket and so are included in the sample of this review.

[31] 'Post-expiration royalties' are particularly controversial and are sometimes seen as an example of patent misuse that is not protected by the statutory limitation of 35 U.S.C. 271(d).

have had an economic incentive to create the improvement, and so an RTLA might result in the invention being underprovided, and its welfare contribution being compromised. Material input patents differ from research tool input patents in that materials are used in the downstream patent, whereas research tools enable the invention behind the downstream patent. Research tools have little commercial value in themselves, so are most susceptible to the problems envisioned in Scotchmer (1991). Although in both cases RTLAs might be used to create fair allocations of rents based on the importance of the input, with research tool patents it might be very difficult to estimate the outcome of the research project.

The welfare loss from type $B1$ patent thickets has two sources. First there is the possibility of inefficient allocation of rents. Information problems, probabilistic patents, and behavioural deviations from rationality could each aggrevate this issue. Second, there is a concern that efficient investment might not take place. With type $B1$ thickets, Williamson hold-up is a genuine possibility once licenses are in place. Then there is a contract that will be subject to a relation-specific investment, which could be appropriated though opportunism. This hold-up is just one example of a broader market failure that arises from imperfections in property-rights. In the academic anti-commons literature, for example, the dominant concern is that researchers may not pursue projects because of the costs or difficulties in securing rights to necessary inputs on suitable terms. This concern should be mirrored in the commercial world. One counter argument to the academic anti-commons hypothesis is that researchers simply move to other projects. In the academic world, the frontier of science may be sufficiently broad that this does not have adverse welfare consequences. But in the commercial world, the difference in both private and social returns between first and second best projects may be more material.

In type $B2$ patent thickets, the subsequent patent has an inventive step that is adjacent to the pre-existing patent, rather than building upon it. To the extent that two adjacent patents have overlapping claims, they are economic substitutes and each patent-holder may believe that it has the right to exclude the other. This is complicated by the 'doctrine of equivalents', which allows a patent-holder to claim infringement for products or processes that do not fall within the literal scope of their claims, but are alleged to be equivalent to their claimed invention. Likewise, detectability and verifiability of the use of patent technology is almost always imperfect and frequently very difficult. However, as the inventive steps are genuinely different for $B2$ patent thicket patents, either probabilistic patents or potentially prohibitive legal costs are required for the patent-holder's conflicting beliefs to be rational.[32] The welfare loss from type $B2$ patent thickets comes from the resources wasted in attempting to exclude (welfare improving) substitution.

Patent Thicket Type B2: *Invention is adjacent. Although the inventive steps of two or more patents are different, they provide closely-related (or identical) economic functionality. As such the claims of the patents overlap in their domain of application. This is more likely to occur with broad patents that have large domains of application.*

We now suggest one final variant for patent thicket types, with the intention that it will primarily apply to type B thickets. It is a refinement to the concept of probabilistic patents, that we name probabilistic-breadth, or variant PB.

Variant PB: *One or more pre-existing patents are broad with a probabilistic diffusion of rights. Such a broad patent has a 'core' domain of application which it covers with validity (or a high-likelihood of validity), but as the distance from the core increases the likelihood that the patent would be deemed invalid or inapplicable increases.*

[32] Since the Supreme Court ruling in *Markman v. Westview Instruments Inc.* 517 US 370 (1996), the interpretation of the meaning of patent claims has been a matter of law, rather than a question of fact, and so determined by a judge, not a jury. 'Markman hearings', also known as 'claim construction hearings', are now held in pre-trial. The probabilistic patent framework requires that courts, not juries, are imperfect arbitrators of patent rights. However, it should be noted that juries are more imperfect than judges. See Moore (2000).

3.8 Type C: Dubious Patents

The patent office, perhaps particularly in the U.S., may engage in what Lemley (2000) described as 'rational ignorance'. As Lemley (2000) aptly put it, when it comes to whether the patent office spends enough time, effort and, ultimately, money, examining patents to ensure that each and every one is valid, "complaints are legion". However, given that many, if not most, patents will never be commercialized, licensed or used in any fashion, it can be efficient, in terms of social welfare, to allow the courts to make more detailed determinations when issues arise, rather than let the patent office incur the enormous costs required to make a proper determination for every application. The recent explosion in patent applications and introduction of new patent classes for genetics, software and business methods, may have altered the balance of the level ignorance that is rational – perhaps now some greater incurrence of cost by patent offices is merited – but it does not undermine the basic calculus. What might undermine the basic calculus is the strategic patenting of spurious patents and the feedback loop that this would cause.

Suppose that firms do set out to obtain patents on inventions that fail to meet the stated requirements for validity for strategic reasons, and rely on the enormous costs and difficulties involved in proving that a patent is invalid in court to render their strategy rational. We know of no evidence that firms do this, but there is considerable conjecture that in certain industries (i.e., business methods, software, and possibly some basic research areas like genetics and nanotechnology) this has become an active practice. The strategic reasons might be offensive or defensive: the firm might use the spurious patents to raise barriers to entry, raise rival's costs, discourage innovation in some area, or for naked transfer-seeking; or they might use them to ensure that they have bargaining chips against other firms, including those that pursue such offensive strategies. As more firms do this, the patent office receives more applications and it becomes easier to get a spurious patent granted as the examination effort per patent falls. Thus strategic use of spurious patents can create a feedback loop. More spurious patents likely beget more spurious patents. Moreover, strategic spurious patents are intended for use; as such they should be less likely than the average patent to be shelved.

> "The system is skewed toward the grant of patents of dubious objective validity, based on a brief, inconclusive process, which are then potentially subject to later disputes with other firms in which legal fees can easily run into millions of dollars for both sides." – Lemley and Shapiro (2005).

We therefore define type C patent thickets as those based on spurious patents. Type C thickets are likely to incur substantially greater search and transaction costs than type A or B thickets, as it will likely be harder to identify when a spurious patent bears on a new invention or product, and agreeing licensing terms is made more problematic by the (probable) mutual understanding that there are validity or informational issues.

Patent Thicket Type C: *Patents, or their claims, are spurious either because they fail to meet the stated requirements for validity or because they intentional undermine the rationale behind the patent system.*

Spurious patent-holders benefit when patents are probabilistic. But "probabilistic patents" are not a necessary condition for spurious patents. For example, it could be rational for a firm to pursue spurious patents if the cost of determining invalidity is so high that some firms could not afford to pay it.[33] Clearly though, if enforcement of patent rights is probabilistic this will favour the creation of spurious patents. The key distinction between spurious patents and probabilistic patents is guile. The filing of a spurious

[33] The legal mean costs in a patent suit were approximately $2.5m in 2010. See AIPLA (2011).

patent may be a willful act of dishonesty (firms may also mistakenly file spurious patents). Probabilistic patents merely require imperfect decisions by the courts.

Type C patent thickets are dividable into two distinct sub-types: those where the spuriousness of patents pertains to validity and those that intentionally undermine the rationale behind the patent system in another way. We will designate type $C1$ thickets as those based on patents which are incorrectly issued by the patent office because they do not embody a new or non-obvious inventive step. Patents that are issued on inventions that are 'useless' will not generate thickets as they have no use.

Patent Thicket Type C1: *At least a subset of patent claims do not embody a new or non-obvious inventive step. The patent office makes mistakes in the issue of some patents. Applicants take advantage of these mistakes to seek patents that cover unpatented or patented prior-art of historic, concurrent, or future inventions, creating apparent infringement.*

Type $C1$ thickets are further sub-dividable into those that do not embody a novel inventive step and those that do not meet the non-obviousness requirement.

Sub Type C1a: *Patents do not embody a novel inventive step in each and every claim. Thus multiple patents have claims, and have been granted exclusionary rights, covering the same domain of application.*

Sub Type C1b: *Firms seek patents that fail to meet the non-obviousness requirement given the prior art. Thus patents have claims, and have been granted exclusionary rights, over a domain of application which was obvious to a person skilled in the art at the time of the patent application.*

This subdivision is useful in its distinction between two kinds of prior art. In type $C1a$ patent thickets, patents are direct substitutes for each other. The prior art embedded in one patent has not been referenced or otherwise determined to prevent the issue of subsequent patents. Only the first patent should be valid. Type $C1a$ patent thickets are therefore more likely to come about when the inventive step is small and can be codified in many different ways. This might be particularly applicable to software and business method patents.

Type $C1b$ patent thickets come about because prior-art has been neglected. Such prior art would render the invention obvious to a person 'skilled in the art'. This is particularly likely with inventions that are close to basic research. Laws of nature, physical phenomena, and abstract ideas are not patentable subject matter.[34] They belong in the public domain. In biotechnology, nanotechnology, academic research, and other sectors, the distinction between such public domain knowledge and a patentable inventive step can be difficult to identify. Accordingly, we might expect that type $C1b$ thickets are associated with these areas.

From a firm's perspective, filing type a $C1$, and perhaps especially type $C1a$, patent is the creation of a real option. In economics a real option is an option – a freely abandonable right to something – that requires further investment to exercise. The penalty for filing an invalid patent is currently the loss of the patent, nothing more. Thus a firm can file for an invalid patent and later decide whether it will incur the costs to try to enforce the patent. If it does there is a chance that it will receive a payoff of zero (it will lose the patent) and a chance that it will secure some kind of return, whether through licensing compensation, or the ability to prevent entry, raise rival's costs, or otherwise increase (or prevent the decrease of) its rents from product markets. This opens up the intriguing possibility of a new policy response to type $C1$ patents: Legislation making the willful filing of invalid patents subject to fines and/or other punishment. Patent invalidity could be determined either by the patent office or by the courts. Assuming some of the revenue from fines accrued to the patent office, this could give the

[34] See the Manual of Patent Examining Procedure (MPEP), sections 2104-2106, and the decision of the Supreme Court in *Diamond v. Chakrabarty*, 447 U.S. 303, 206 USPQ 193 (1980) as well as *Association For Molecular Pathology v. Myriad Genetics*, 569 U.S. (2013).

patent office a rational incentive to increase its examination efforts. And perhaps more importantly, this would undermine notion that filing for spurious patents constitutes a real option: At any time an outside party could potentially request the court to declare a patent invalid and impose a fine.

Type $C2$ patent thickets do not depend on invalid patents, but rather are based upon patents that deliberately 'game' the patent system. Shapiro (2001) describes the problem of 'submarine patents' – patents that are deliberately kept in their application stage so that their claims can be adjusted to cover new inventions or products. Shapiro (2001) is explicit that submarine patents are not a part of his definition of a patent thicket, but at least one other author, Rubinfeld and Maness (2004), believed that they were. Submarine patents require an act of guile by the patent applicant. The applicant deliberately waits for a relation-specific investment to incur before inducing the 'contract' that enables the opportunistic extraction of rents by requesting the grant of the patent.

Patent Thicket Type C2: *Through the use of continuations, applicants can keep their patents in an application phase and adjust their claims. Prior to November 2000, patent applications were not disclosed. This made so called 'submarine patents' possible. An inventor might then incur costly research and development only to find that the resulting invention infringed another patent.*

Type C patents do not make positive contribution to economic welfare – nothing new (that wouldn't have been invented anyway in the case of type $C2$ patent thicket patents) has been created. Instead, much as in Tullock (1967), applicants seeking spurious patents are transfer-seeking. All resources involved in this transfer-seeking, including the resources devoted to the creation and prosecution of the patent, the search and transaction costs of firms that attempt to navigate around this patent, and the injunctions, suits and damages involved in 'infringement' and 'validity' challenges, are all wasted. Moreover, the willful filing of a spurious patent is an act of guile, and the patent office creates a contract between society and the spurious-patent's owner in the (erroneous) assignment of exclusionary rights. We therefore anticipate a market failure from hold-up. Genuine innovators may opt not to innovate for fear of appropriation of their rents by spurious patents. It is conceivable that many authors of papers in our sample had spurious patents in mind as the underlying mechanism for hold-up, even though this was never articulated.

3.9 Type D: Effectively Saturated Invention Spaces

> "Everything that can be invented has been invented." - Charles H. Duell, Commissioner, U.S. Office of Patents, 1899.[35]

Patent Thicket Type D1: *A number of patents effectively saturate the potential invention space. Patents may be small and cover 'marginal' inventions.*

Many authors have used the term 'marginal patents' apparently to mean spurious patents.[36] We take the term to mean that the marginal benefit of the patent just exceeds the marginal cost of its filing. As such we assume that a marginal patent has a small inventive step, just above the threshold necessary to establish patentability, and a correspondingly small claimed scope of exclusionary rights. The notion of a marginal patent implies that patents can not saturate an invention space; there will always inventions that are too small to be patented. Accordingly, we will refer to an invention space that is completely covered with patents, so that no further patenting is possible, as being effectively saturated.

[35]It appears that this well-known quote is apocryphal; see Sass (1989). However, Henry L. Ellsworth, Commissioner, U.S. Patent Office, 1843, included the following "rhetorical flourish" in his report to Congress: "The advancement of the arts, from year to year, taxes our credulity and seems to presage the arrival of that period when human improvement must end."

[36]These authors may have meant 'marginal validity', but this is an awkward concept. In our classification, the concept of marginal validity could be applied to variant PB patents, where the validity decreases as the distance from the core of the claimed domain of application increases.

Although we think it unlikely over the long term, it is possible that over short to medium terms that areas of invention space can become effectively saturated. That is, for a period, in certain areas, everything that can get patented gets patented.[37] Earlier patented inventions in such areas are probably larger and later inventions then probably fill in the holes, adding, removing, or refining functionality. Paradigm shifts and waves of creative destruction may then clear the invention space (or perhaps move it to a higher level) so that the process of patented invention can begin again.[38]

Firms may intentionally effectively saturate invention spaces with patents. When a new invention is discovered, many other ways of implementing the same economic functionality might become apparent. Firms may, therefore, adopt a strategy known as 'ring-fencing'. When a patent on an invention is susceptible to work-arounds that are themselves novel and non-obvious, the free-riding problem that patents are intended to solve extends beyond the original invention.

Type D patent thickets may be more common in pharmaceuticals and perhaps chemicals, than in the development of physical devices, which in turn are probably less susceptible than computer software or other more complex abstractions. When an invention is a composition of matter whose purpose is linked closely to nature of composition, variations on the composition are unlikely to serve the same ends. Many drugs, for example, are very specific compounds; altering the compound even slightly might render the variant of the drug ineffective. An instance of software, at the other extreme, represents one of many different ways of achieving the same technical functionality, and the technical functionality itself may have many economic substitutes. Furthermore, type D thickets are probably more common in what Cohen et al. (2000) categorized as discrete, rather than complex, technologies. When a patent is a product, as in pharmaceutical, it makes economic sense to protect all of the potential substitutes. But when a product is made up of many patents, each of which in turn might be made up of many patents (and so forth), as is common in hardware and software, protecting all of the variations can become infeasible.

There are no obvious welfare consequences from type D thickets on invention. One ownership structure of patents inside the thicket might have different welfare consequences from another – having dispersed ownership of substitutes would promote competition though likely with increased invention costs – but once the thicket is in place, there is no impact on future invention, as there is no future invention. Schumpeterian innovation would suggest that type D thickets are welfare improving for commercialization. With every technical way of doing something already created, codified, and displayed for all to see at the patent office, Schumpeterian innovators are still free to find new combinations and create new welfare in new ways.

Finally, we note that 'evergreening' – the process of making a minor adjustment to an patented invention and filing for a new patent shortly before the expiration of the previous patent – could be consistent with type B, C or D patent thickets. On one hand it appears much like ring-fencing. However, the time-lag between applications begs the question of why a competitor did not patent the modification themselves. Patent thicket possibilities include that: the original invention is a complementary input and the patent-holder declined to license the patent for cumulative invention to assure itself the right just prior to expiry (though this strategy would only work for a single generation); the new invention is adjacently overlapping and the competitors were deterred by the threat of litigations; or the new invention is actually not novel or non-obvious given the previous patent. Non-patent-thicket possibilities include information problems (perhaps the patent did not fully codify the invention) and a lack of competition (either because no competitors are present with the requisite knowledge to create the modification or

[37] The mouse trap provides an imperfect example. Dagg (2011) provides a fascinating history of the mouse trap that both supports and nullifies this argument. A Google Patents search for "Mouse Trap" returns about 42,900 results, with around 1,100 applications filed in 2012.

[38] Paradigm shifts restore the infinite possibilities that the human mind seems to demand. The contest is then between two infinite potentials for betterment: one linear and the other a spiral.

because competitors actively collude and allow each other to evergreen).

4 The Literature

4.1 Literature review process

The aim of our literature review process was locate, document and analyze the population of papers that explicitly focused on patent thickets and a near-population of papers that provide a definition of a patent thicket. The paper collection effort was undertaken in four waves. The first wave was a keyword based draw from journal repositories and search engines. Every result of a keyword search for variations on the term 'patent thicket', including variants on 'anti-commons', was retrieved from the Proquest, EBSCO, World of Science, and JSTOR journal repositories. In each case, papers were downloaded and given a cursory review for suitability for inclusion in one of three groups: a core group where papers are explicitly focused on the topic of patent thickets; a downstream group consisting of papers that used patent thickets as an input in some fashion and explicitly defined the term 'patent thicket'; and an upstream group of papers that could be used to provide underlying theory, stylized facts, useful definitions of other terms, or other material that would be relevant to understanding patent thickets. In addition, more than a thousand results from Google Scholar searches for patent thickets, variant terms, and other keywords that are used in the patent thicket papers such as 'blocking +patent', 'infringing +patent', 'network +patent', 'Herfindalh +patent', 'fragmentation +patent', were given a cursory review. This process yielded 251 papers of interest.

We then released the list of papers to a small number of academics and practitioners and solicited additions. Surprisingly, we received only 14 additional suggestions. The third wave of the collection effort was a two pronged convergence process. We identified 40 candidate core papers and extracted all of their references. These references were then matched against one another using custom-built software to create a list of all (313) referenced papers and their citation counts from within the candidate core. Of the 313 referenced papers, 112 were already in our set of papers of interest. Each of the remaining 201 papers were given a cursory review, and of these 46 were added to our list of papers of interest. Concurrent with this automated convergence process we also undertook a manual convergence process. We compiled a list of 12 papers that had themselves undertaken detailed reviews of the patent thicket literature and checked each of their references. We also compiled a list of papers from the 'early-period' of the patent thicket literature (prior to 2005) that we thought were likely to be used as the definitional reference for patent thickets in downstream papers. This list, in the approximate order of its yield, was Shapiro (2001), Heller and Eisenberg (1998), Ziedonis (2004), Hall and Ziedonis (2001), and Walsh et al. (2003). We then retrieved lists of each paper that cited these papers from Google Scholar and reviewed them for inclusion. This process added just another 13 papers to our list of papers of interest, which we took as a sign that we were approaching the population of papers that were suitable for our purposes.

Finally, we went through each paper in our list in some detail. This involved preliminary tagging of the paper's 'thicket stance' (the attitude of the paper towards whether or not patent thickets exist or hinder innovation), discipline (economics, law, etc.), research type or types (theory, empirical, etc.), and other attributes that we will discuss shortly. We also made a final determination as to whether papers should be categorized as core, upstream, downstream, or discarded. This resulted in 114 papers being assigned to the downstream group and exactly 50 papers assigned to the core group. For each of the 50 papers in the core group we wrote up a small review.[39] As we went through each paper we checked that referenced works that were likely to be useful for our current undertaking were included in our list of papers of interest. We found 3 further works, but each of them appears to be a contribution to an out

[39]These reviews are available online at www.edegan.com.

of print book or an unavailable conference proceeding. Aside from papers from recent years (essentially 2012 or 2013), we feel confident that no core papers have been missed and that our set of downstream paper closely approximates the population.

4.2 Patent thicket definitions

For the core and downstream paper groups we manually determined which patent thicket types, sub-types and variants listed in our taxonomy were consistent with the definition of a patent thicket used in the paper. Unfortunately many papers either did not provide explicit definitions or provided such sparse explicit definitions that in general the assignment of types to papers relied heavily on contextual statements made by authors. Core papers are explicitly about patent thickets. Some core papers were theory papers, such as Shapiro (2001) and Heller and Eisenberg (1998), that set out to define the mechanisms and potential consequences of patent thickets, and so the assignment of types was comparatively easy.[40] However, other core paper were empirical or survey papers and set out to provide stylized facts about thickets or evidence of their existence and consequences; or were discussion papers either for policy institutions or by legal or economic scholars for the general academic community; or were papers introducing new measures of patent thickets. Some of these papers did provide clearly articulated definitions, but many did not. Thus, although every attempt has been made to correctly assign types to the authors' definitions, we must stress that the assignment is subjective and may be incorrect in some instances. We provide a full list of which types were assigned to which papers in the appendix.[41] This task of type assignment was more problematic in the downstream group, where definitions were invariable sparser and more frequently implicit. To reduce misclassification errors as much as possible, this classification process was undertaken twice, independently by the authors and a group of research assistants. The research assistants were not given the actual types in our taxonomy but instead were asked to tag each paper with a series of keywords which were used to discriminate between types. The concordance between the research assistants' tagging and our definitions was very strong, suggesting that classification is sufficiently simple to be undertaken by non-experts and any classification errors in our samples are likely to be fairly small. We do not expect that the correction of classification errors (if such a task is possible) to materially affect the results we present below.

A paper's patent thicket definition might be consistent with multiple different patent thicket types, sub-types and variants. Our approach was to mark the definition with the finest set of sub-types and variants possible. Thus a definition that was consist with thicket type A–PST (main type A, variants P, S, and T) was marked A–PST, and not just type A, and not A–T, A–S, and A–P. However, when it was impossible to distinguish between sub-types, the definition was marked with a high-level type. For example, a paper with a definition that was consistent with both $B1$ and $B2$ sub-types would be marked B. Table 1 below provides detail on the incidence of thicket types across our full sample of 164 papers with codifiable definitions, made up of 50 core papers and 114 downstream papers.

As shown in table 1, type A (diversely-held complementary inputs) thickets are the most common. Around 60% of papers (99 out of 164) that provide a definition of a patent thicket provide one that is compatible with type A thickets. Almost half ($\frac{38+8+1}{90} \approxeq 48\%$) of these papers refer to variant T thickets – they mention the transaction cost problems inherent in type A thickets. Type B (legitimate overlapping patents) thickets, however, are almost as common. When a definition is clear about which sub-type of type B thickets they are referring to, they use the sub-type $B1$ (cumulative invention) definition 80% ($\frac{16}{16+4}$) of the time. Type C (spurious patent) and type D (effectively saturated invention space) thickets are substantially less common, being mentioned in just 16% and 10% of papers respectively. Type C

[40] A great deal of attention was still aid to exactly what the authors meant.

[41] Authors of referenced works are welcome to correct us where we are in error; please understand that we had to work with what you actually said and not with what you meant.

Table 1: Occurrence of variants and sub-types

The table describes the incidence of variants and sub-types of the high-level types for papers (which provided codifiable definitions) from the literature review. Main (without refinement) was used when a definition was consistent with the overarching high-level definition but not specific variants or sub-types. Note that variants are not exclusive with sub-types so the counts do not sum to the totals at the bottom.

	Type A	Type B	Type C	Type D
Main (without refinement)	50	39	2	15
S Variant	7	2	0	2
T Variant	34	17	1	0
ST Variant	8	5	1	0
PST Variant	1	0	0	0
Sub Type 1	-	16	26	-
Sub Type 2	-	4	2	-
Sub Type a	2	-	1	-
Sub Type b	0	-	0	-
All variants and sub-types	99	80	28	17

thickets are almost invariably sub-type $C1$ (failure to meet patenting requirements) when a sub-type is mentioned. Just one papers mentions sub-type $C2$ (submarine patent) thickets: Rubinfeld and Maness (2004) explicitly include them in their definition of a thicket. Sub-types a and b are almost always never differentiable for type A or type C thickets, where they are possible. Just two papers – Ayres and Parchomovsky (2007) and Lerner and Tirole (2008) – mention sub-type Aa (perfect complement) thickets, and only a single paper – Eisenstein (2010) – mentions type $C1a$ (novel inventive step) thickets. And over all definitions, 40% ($\frac{34+17+1+8+5+1+1}{164}$) and 16% ($\frac{7+2+2+8+5+1+1}{164}$) are consistent with variants T (transaction costs) and S thickets. This is concentrated almost entirely in type A and B thickets, despite being applicable to type C and D thickets as well. In fact, as has already been mentioned, transaction and search costs should be considered central to type D thickets and are highly relevant to type $C1$ thickets. Variant P thickets were mentioned (explicitly and clearly) in Ziedonis (2004).

We now move to an analysis of the co-occurrence of thicket types, and to simplify the analysis we focus on just high-level types. That is, we will ignore sub-types and variants, and, for example, treat type $B1$, $B1{-}ST$, $B2{-}T$ and other type B's as just type B. Table 2 below documents the co-occurrence between high-level thicket types. The table provide log-odds ratios for the co-occurrences of types by paper, with statistical significance determined using Fisher's Exact Test. Log-odds ratios are particularly nice for two reasons: odds ratios below one give rise to negative log-odds ratios so it is clear when one type is unlikely to occur with another; and log odds ratios are comparable in their absolute magnitudes so a log odds of -0.1 is the same distance from equal odds as a log odds ratio of 0.1. Fisher's Exact Tests are exact tests of statistical significance; they can determine significance with any sample size and do not just approximately significance only becoming exact in the limit ($N \to \infty$).[42]

Despite our comparatively small sample size (164 papers) for statistical analysis, we do observe a highly statistically significant negative correlations between type A and types C and D thicket definitions. Authors who define their thickets consistent with type A thickets are statistically unlikely to also include statements consistent with type C thickets at the 1 in a 100 level, and are statistically unlikely to also include statements consistent with type D thickets at the 5 in 100 level. No other statistically significant correlations are present. Type A thicket definitions are not statistically likely to refrain from include

[42] There is some concern in the literature that Fisher's Exact Test might be conservative in establishing significance. See Liddell (1976) and others.

Table 2: Co-occurrence of patent thicket definitions

The table describes the co-occurrence of high-level patent thicket definitions, as well as with an identifier as to whether the definition included hold-up. Column 2 reports the counts of the occurrence of each definition. Note a paper's definition may be classified as consistent with multiple types (i.e., A and B) so the counts do not sum. The remaining columns report log odds ratios of co-occurrence, with Fisher exact tests used to calculate statistical significance. ***, **, and * denote significance at the 0.01, 0.05, and 0.1 levels, respectively.

	Count	Type A	Type B	Type C	Type D
Type A	99	-	-0.13	-1.62***	-1.15**
Type B	80	-0.13	-	-0.29	0.19
Type C	28	-1.62***	-0.29	-	0.81
Type D	17	-1.15**	0.19	0.81	-
Hold-up	26	0.06	0.24	-0.53	0.14

statements consistent with type B thickets, and type B thicket definitions can be confounded with any another thicket definition. During the thicket definition classification process we marked all mention of hold-up as a part of the thicket definition. It is interesting to note that hold-up is not statistically likely to be included or not included when a definition is consistent with any particular thicket type. We might that thought that hold-up would be correlated type $C2$ thickets in particular, and with type C in general, then perhaps with type B, and not with type D at all, but this is not the case. We take this as cursory evidence to support our claim that hold-up is poorly understood in the patent thicket literature.

Figure 3, below, shows the usage and number of thicket types supported by the papers in our sample over time. In figure 3a stack bars show the number of papers consistent with each type. The figure has a tail before 2002: when the literature was young, almost all papers where defining thickets somewhat independently for the first time, and all of these papers are in our core sub-sample. The main bulk of papers in our sample were published between 2003 and 2011, and the decline in 2012 and 2013 is almost surely due to our inability to find very recently published papers, rather than a decline in research. Type A definitions are clearly most common, and account for a little less than half all definition types used, more or less equally in every year. Type B definitions are a clear second, and again hold mostly constant in their usage over time. Type C and D definitions are much less frequent and start in earnest in 2003. Figure 3b shows the mean number of types that were consistent with a definition in a paper for each year from 2003 to 2012. The line of best fit has a substantial upward slope that is largely unchanged by exclusion of 2012 where the number of papers is small. This is concerning evidence of a history of growing confusion in the literature over the definition of a patent thicket.

EPC and Board (2013) documents the findings of the European Patent Office (EPO) Economic and Scientific Advisory Board's workshop on patent thickets. One of their major reported points was widespread agreement that a patent thicket "usually involves (1) multiple patents on (2) the same, similar, or complementary technologies, (3) held by different parties, making it difficult to negotiate intellectual property rights (for example, licensing agreements) to the point where some scholars feel it might be socially inefficient."[43] The 'same', 'similar', and 'complementary' technologies map into type C, B and A thickets rather nicely. Although we would argue that diverse ownership is only a necessary condition for type A thickets, and that multiple patents are only required either in so far as they are a sub-condition for diverse ownership or as an added onus towards transaction and search costs, we are heartened by the acknowledgement of definitional issues and the move to correct them.

[43] They were apparently unable to restrain themselves from including that "[a] patent thicket conjures up the image of a bramble, a large dense bush with thorns on the branches making it difficult to pass through without getting severely scratched."

Figure 3: Patent thicket definitions over time

(a) Frequency of high-level type usage

(b) Number of high-level types per paper

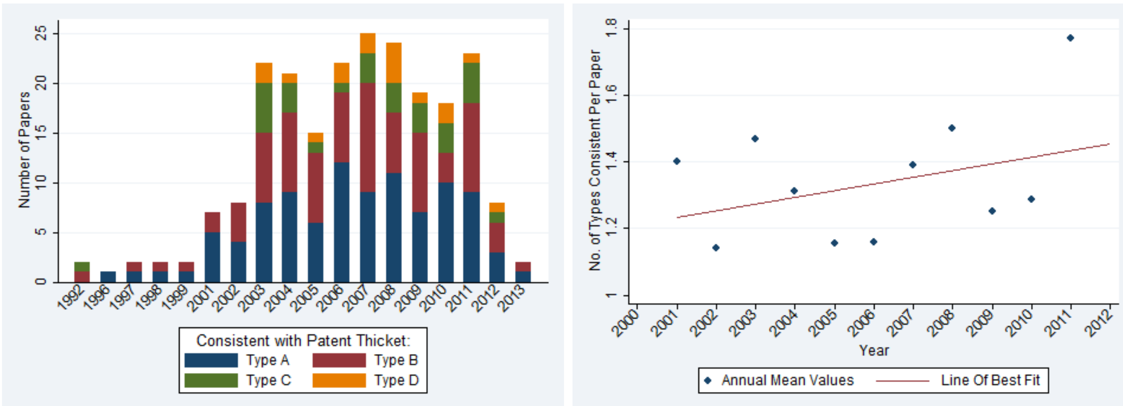

4.3 The 'geography' of thicket definitions

In this sub-section we consider the 'geography' of thicket definitions – where in the literature each high-level type is likely to be found. Each paper was categorized as belonging to a single discipline: Economics/Management, Law, General Science, and Policy Reports. The classification was based on the type of work undertaken in the paper, not by the (dominant) discipline of focus of the author(s) or by which journal it was published in, with the exception of policy reports which had to be commissioned by a recognizable and reputable entity like a patent office or the Federal Trade Commission (FTC). Each paper was also marked as containing, at least as a material component, work that could be described as theory, empirical, survey, measures or discussion. A single paper could be marked with several such designations. For example, Galasso and Schankerman (2010) contains both theory and empirics and was marked accordingly. The classification of and between theory and discussion was somewhat problematic for papers from the law literature, but we do not expect much controversy over our subjective judgments.

Likewise, downstream papers were coded for their topic. We identified five main topics and report results for each. Firm strategy papers provide strategic advice to firms regarding patent thickets; they discuss the strategic implications of blocking patents, pre-emptive patenting, secrecy, evergreening, willful infringement, engaging in Mexican standoffs or other mutually assured non-aggression or destruction, and other defensive or offensive patenting behaviors, as well the consequences of doings so on collaboration, industry structure including entry, and the value of firms. We might expect that firm strategy papers could use almost any definition type, though they might favor type C (spurious patents) which are particularly suited to pre-emptive patenting, certain types of evergreening, and some of the more obnoxious offensive or defensive patenting strategies.

Private mechanism papers discuss cross-licensing, patent pools, patent clearinghouses, patent collectives, Fair/Reasonable and Non-Discriminatory (FRAND) licensing, patent intermediaries including Non-Participating Entities (NPEs) of which patent trolls are a sub-species, shared platforms, and standards and Standard Setting Organizations (SSOs). As such we might expect that private mechanism papers will be associated with type A papers.

Intellectual Property Right (IPR) reform papers are more policy focused. They suggest reforms to the nature of IP rights, processes for granting patents at the patent office, and advocate approaches to patent-based transactions for anti-trust authorities. Reforms suggested include changes to renewals and duration limits, stricter patenting requirements, pre- and post-issue reviews of patent validity, industry specific patent policies and the creation of new patent classifications (sometimes with changes to the requirements for the new classes), as well as calls for compulsory licensing, research exemptions, and other changes to rights transfers that might mitigate the problems a thicket could cause for innovation. We might expect type C thicket definitions to be associated with IPR reform papers that are concerned with the quality of patents, and perhaps type $B2$ (adjacent overlapping patent thicket) definitions for papers that are concerned with the refusal to license to prevent substitution that would enhance welfare. To the extent that IPR reform papers focus on anti-trust considerations, they should use type A definitions. Anti-trust regards collusion mechanisms as potentially problematic, and type A thickets may need collusive mechanisms to prevent a hindrance of innovation.

The final two main topics we consider are 'academic research' and 'industry commentary'. As we have already mentioned there is a sub-literature on how patent thickets might affect academic research. We expect that this would focus on type $B1$ (cumulative invention) thickets, following Heller and Eisenberg (1998)'s and Walsh et al. (2003)'s seminal works on the anti-commons in biomedicine. We also found a small number of industry commentary papers that provided a codifiable thicket definition.[44] From our definitions of the types, we have no prior as to what papers that discuss patent thickets in various industry will use for their definitions above and beyond that which is best suited to each paper's focal industry.

For every paper in the full sample (i.e., both core and downstream groups) we recorded the industries that were given material attention. Generally in order for a paper to coded as belonging to an industry it had to have at least a section of its analysis devoted to that industry. Most papers were marked with a single industry and no papers were marked with more than four industries. We then aggregated the industries into four groups: industries that are likely to be associated with basic research, including academia, biotechnology, genetics, nanotechnology, and pharmaceuticals; industries characterized by complex products, including information and communications technology aside from software and the internet (associated with business method patents) and specifically including semiconductors, and manufacturing; the software and Internet industries; and 'other' industries, which included papers that were general in their analysis or did not reference any specific industries. We expect type B thickets to be associated with basic research industries, type A thickets to be associated with complex product industries, and type C patents to be associated with the software/business method industries, where there is a perception that the quality of patents might be lower despite some evidence to the contrary.

Table 3 provides a mix of reassurance and surprises. Type A thicket definitions are significantly correlated with theory, economics or management papers in the core. Given that Shapiro (2001)'s definition of a patent thicket as a diversely-held complementary inputs problem is more readily understood by economists than non-economists, and that it lends itself naturally to the development of other theory, this is no surprise. It is also unsurprising that type A thicket definitions are less likely to be found in papers discussing basic research or firm strategy. However, the statistically significant low likelihood of the use of a type A definition in IPR reform papers is troubling. Just under half (24 out of 49) of the IPR reform papers in our sample have a material focus on anti-trust considerations, where the use of definitions other than type A would suggest errors in analysis.

Type B patents aren't associated with economic theory. They are discussed in the core, but mostly by legal scholars. We suspect that the negative likelihood of a type B thicket definition being used in economic theory papers is simply due to the perception that type A thickets are the correct or dominant form in economics. Type B patents are highly statistically significantly associated with papers that focus on industries characterized by basic research, as we predicted that they would be. An association

[44] We found a very large number of industry commentary papers that did not provide a codifiable thicket definition. The vast majority of industry commentary papers, even of those that are explicitly discussing the ramifications of patent thickets, did not provide a definition for what a patent thicket is, and we found that we were generally unable to even hazard a guess of what their authors intended.

Table 3: The usage of thicket definitions by paper types

The table documents the occurrence of high-level patent thicket definition types with various attributes of papers (which provided codifiable definitions) from the literature review. The values are log odds ratios of co-occurrence, with Fisher Exact Tests used to calculate statistical significance. ***, **, and * denote significance at the 0.01, 0.05, and 0.1 levels, respectively. Note that for attributes that are not distinct (i.e., a paper may be both theory and empirical) the counts do not sum to the total.

	Count	Type A	Type B	Type C	Type D
Paper Type:					
Theory	57	0.65*	-0.86**	0.42	-0.27
Empirical	41	0.03	0	0	0.84
Survey	3	0.28	0.76	-	-
Measures	6	0.28	0.05	-	0.57
Discussion	79	-0.38	0.44	-0.26	-0.32
Paper Discipline:					
Economics/Management	83	0.71**	-0.84**	-0.38	0.37
Law	53	-0.58	0.58*	0.18	-0.88
Gen. Science	11	-1.05	0.25	0.08	-
Policy Report	17	0.21	0.73	0.46	1.15*
Core (Topic):					
Patent Thickets	50	1.32***	0.77**	-0.56	-0.79
Downstream (Topic):					
Academic Research	7	-0.74	0.35	-	-
Firm Strategy	26	-0.88**	-0.9*	0.96*	0.92
Industry Commentary	10	-1.92**	0.48	0.21	0.84
IPR Reform	49	-0.79**	-0.11	0.88**	-0.36
Private Mechanisms	69	0.04	-0.27	-0.51	-0.32
Industry Covered:					
Basic Research	59	-1.06***	0.89***	-0.85*	0.25
Complex Product	32	0.63	-1.28***	-0.44	0.27
Software/Business Methods	11	1.97*	0.25	-	-0.16

between type *B* patent thicekts and the 7 academic research papers in the downstream group failed to find significance.[45] Type *B* patents are also not associated with research on industries with complex products, again consistent with our priors.

Type *C* and type *D* consistent definitions are comparatively rare – only 45 papers (27% of our sample) used such definitions. The presence of type *C* thicket definitions is positively correlated with their papers being about firm strategy or IPR reform. Both correlations were expected. For firm strategy, some offensive or defensive patent strategies call for the use of spurious patents, and many IPR reform papers suggest mechanisms for dealing with spurious patents. It is also reassuring to note that type *C* definitions are not associated with basic research, as this is where spurious patents could do the most damage. Type *D* (effectively saturated invention space) thickets are associated with policy reports. We can not imagine what policy advice would be constructive in their regard, and consider this mild evidence of a lack of understanding of patent thickets by those advocating changes based on their alleged effects.

4.4 Definitions and stances

We now ask whether there is a systematic relationship between the thicket definitions used in papers and their 'stance' regarding patent thickets – whether or not they believe that patent thickets exist or hinder

[45]Because Fisher's Exact test is an exact test, issues of statistical power do not arise.

innovation. We coded each paper as having one of the following stances: anti, weakly anti, assumed anti, neutral, assumed pro, weakly pro, and pro. Papers were coded with assumed stances if they didn't make direct statements about whether or not patent thickets exist or hinder innovation but instead explicitly drew their stances from other work. Stances were coded as being weakly anti or weakly pro if their position was couched in moderating language. Table 4 below reports the incidence of stance codings using three classification scales.

Table 4: Paper thicket stance codings

	Assumed are neutral		Assumed are omitted		Assumed are ordered	
Value	Definition	Count	Definition	Count	Definition	Count
1	Anti	10	Anti	10	Anti	10
2	Weakly Anti	10	Weakly Anti	10	Weakly Anti	10
3	Neutral/Assumed	52	Neutral	15	Assumed Anti	1
4	Weakly Pro	35	Weakly Pro	35	Neutral	15
5	Pro	57	Pro	57	Assumed Pro	36
6	-	-	-	-	Weakly Pro	35
7	-	-	-	-	Pro	57
Total		164		127		164

In table 5, below, we report the relationship between the use of definitions consistent with the different types and the stances of the papers. The table reports Kendall τ_B rank correlation coefficients and their significance calculated using asymptotic τ_B standard errors between the occurrence of high-level thicket definition types and the stance scales described in table 4, for the core group of papers, the downstream group of papers, and the full sample. For the core papers, stances are generally based upon either empirical evidence or theoretical postulates – core papers are those that do work directly on patent thickets and many core papers rightly attempt to determine, deductively or inductively, whether or not thickets exist or hinder innovation. For the downstream papers, however, stances are based upon perception and beliefs. Authors of these works may make assertive stances about whether thickets exist or hinder innovation, but these statements are not supported by the research in their papers – downstream papers take patent thickets as inputs and study other things.

Neither type A nor type B thicket definitions are associated with an increased or decreased statement that thickets exist or hinder innovation. It appears that the jury is out for these types of thickets, both in terms of actual evidence and according to perception in the downstream literature. The number of observations is small for type C thickets, making it hard to find significance. However, it does appear that both in the core group of papers and overall in the full sample there is a weakly significant positive correlation between making a more 'pro' stance and using a thicket definition that is compatible with type C thickets. To the extent that we believe that spurious patents are being issued by the patent office and used strategically by firms, this is not surprising.

The most striking result in table 5 concerns type D (effectively saturated invention spaces) thickets. Type D thickets are by far the least commonly referenced; compliant with just 10% of definitions provided in our sample. Despite their infrequent use, it appears that type D thicket definitions are strongly significantly associated with theory, evidence, or statements in the core group of papers suggesting that thickets do exist and do hinder innovation. This sentiment is reversed in the downstream group of papers, who apparently believe that type D patents are (weakly significantly) not associated with thickets and their problems. The small size of the core and relatively large size of the downstream group makes type D definitions insignificant predictors of thicket stance overall.

Table 5: Thicket definitions and the paper stances

The table documents Kendall τ_B rank correlation coefficients, and their statistical significance, between the thicket definition types and the 'stance' of the paper. Stances were coded on either a 5-point scale (anti, weakly anti, neutral, weakly pro, pro) with papers that made it explicit that their stance was based on assumptions in the prior literature as neutral, or a 7-point scale with 'assumed anti' and 'assumed pro' at values 3 and 5. ***, **, and * denote significance at the 0.01, 0.05, and 0.1 levels, respectively.

Sample	Coding of Assumed	Count	Type A	Type B	Type C	Type D
Core	Stance:					
	Assumed are neutral	50	0.09	-0.13	0.16	0.23***
	Assumed are omitted	43	0.01	-0.17	0.13	0.22***
	Assumed are ordered	50	0.12	-0.12	0.17*	0.23***
Downstream	Stance:					
	Assumed are neutral	114	0.06	0.03	0.1	-0.15*
	Assumed are omitted	84	0.09	-0.01	0.14	-0.19*
	Assumed are ordered	114	0.05	0.04	0.1	-0.16*
Full	Stance:					
	Assumed are neutral	164	0.09	0	0.11*	-0.06
	Assumed are omitted	127	0.08	-0.06	0.14*	-0.07
	Assumed are ordered	164	0.09	-0.02	0.11	-0.06

5 Measures and Tests

In this section we describe the foundations of the measures and tests used in the 20 empirical papers of the core patent thicket literature. These are the near-population of papers that have, to date, contributed to our understanding of whether patent thickets exist and/or provide a hindrance to innovation. A list of these papers is provided in table 7. We then document the frequency with which these measures have used, the context of their use in terms of in which industrial area and with which outcome measures they have been applied, and their basic results. We conclude the section by showing which studies have claimed thickets exist and hinder innovation, and which types of thickets from our taxonomy their results might be applicable based upon the measures and tests that they have implemented.

5.1 Patent statistics

Patent statistics are any kind of codifiable metric concerning patents. The simplest patent statistic is the number of patents filed or granted in some area. But one of the joys of patents is their mandatory disclosure. A patent filing contains a wealth of information, including the identity of the inventor, the identity of the original assignee(s), a list of cited prior art, the technology class assigned to the primary claim by the patent examiner, the date of the application and grant, a full list of the claims made, a title and brief description for the invention, and figures and diagrams. Much of this information can be turned in to data and used in empirical analysis. The comparative paucity of other measures of innovation has made patent statistics hugely important to academic researchers. The next most common sets of measures pertain to R&D (i.e., R&D expenditure, the number of employees engaged in R&D, etc.) performed by publicly-traded firms or reported in surveys to census bureaus or other entities, and indirect innovation measures such as venture capital investment or those based on the production of basic science with innovative potential such as the number of academic publications in some field, their rates of academic citations, and so forth. However, there is a clear qualitative difference between these other measures and patent statistics. These other statistics are generally inputs, which may or may not result in invention, let alone innovation.[46] A patent is at least an output of an invention process, even if it is an input into a commercialization process (and perhaps an input into other inventions) and so an intermediate of an innovation process.

The major problem with patent statistics is one of construct validity – the degree to which a measure actually measures what is claimed it does or what its label implies it does. This arises in part because economists, legal scholars, and other academics have an imperfect understanding of patents; the theory of patents is well developed but sometimes contradicts practice. As examples, patent applications and renewals are costly, but around half of all patents are never used.[47] Patents confer exclusionary rights so that firms can appropriate rents to cover the cost of invention, but a survey of semiconductor firms reported that firms don't use patents to appropriate rents.[48] And attempts to explain observed inter-industry differences in the propensity to patent led authors to conclude that they were driven "largely by irrelevant fluctuations".[49] Academics (ourselves included) are, of course, working to reconcile the incongruity between theory and practice, but its persistence is a core point of this paper.

Construct validity problems also arise in part because patents are so wildly heterogeneous. As Jaffe and Lerner (2011) point out, there are patents on things ranging from a sealed and crustless peanut butter and jelly sandwich to the OncoMouse™.[50] Summing the counts of patents is therefore a questionable proposition. But this heterogeneity strains our confidence in other measures too. It isn't particularly reasonable to expect to be able to tell apart rodents that invariably get cancer from the signature product of J.M. Smucker's line of Uncrustables® based on a count of the number of claims (10 versus 12) or citations-made (7 versus 2).[51] An appeal to the law of large numbers is the primary recourse here.

Mostly though, the construct validity of patent statistics suffers as a result of what we will call Grilliches' Law. Griliches (1998) famously asked "Patents as indicators of what?" He went on point out that "What would we like [patent statistics] to measure?" is not the same as asking "What aspects of economic activity do patent statistics actually capture?" Although he identified numerous issues with patent statistics, advocated a process of validation through a combination of imposing structural relationships and determining supporting correlations, and was careful to suggest using other data whenever it was available, Grilliches' warnings appear largely unheeded. He concluded his 1998 paper by saying "In spite of all the difficulties, patent statistics remain a unique resource for the analysis of the process of technical change. Nothing else even comes close in the quantity of available data, accessibility, and the potential industrial, organizational, and technological detail... We should not be cursing the darkness, but rather, we should keep on lighting candles." This last sentiment was taken to heart by subsequent generations of researchers.

Grilliches' Law[52]: *If it relates to innovation, someone, somewhere, will claim that they can measure it with a patent statistic.*

In the next subsections we will discuss patent citations, the two dominant measures of patent thickets

[46]Entry into some sector or product class, changes to total factor productivity or firm value, and technology licensing are somewhat commonly used innovation output measures. Though technology licensing often involves patents and could be thought of as a patent statistic, and one firm's technology licensing output is another firm's technology licensing input.

[47]See Rossman and Sanders (1957).

[48]This is the 'Patent Paradox' discussed in Hall and Ziedonis (2001).

[49]The quote is from Griliches (1998) and the data is from Levin et al. (1987).

[50]U.S. Patent 6,004,596 and U.S. Patent 4,736,866, respectively.

[51]One might have more luck using the count of citations-received, which according to Google Patents stood at 9 to 332 in May of 2013. But we would caution that there are plenty of counter-examples for citations-received as well.

[52]Zvi Grilliches, 1930-1999, president of the Econometric Society 1975, president of the American Economics Association 1993, John Bates Clark Medal winner 1965, was an outstanding academic. Grilliches will be remembered as a founding father of innovation economics, both for his own contributions and as a mentor, supervisor and co-author to an entire field of researchers. The 'Tree of Zvi' is now more than 4 'generations' deep and covers many hundreds, if not a thousand, academics. Many, if not most, of the authors of core papers in this literature review are a part of Zvi's tree; as is one of the co-authors of this work. He deserves a better eponymous law.

– Ziedonis (2004)'s fragmentation measure and Von Graevenitz et al. (2011)'s triples measure – both of which are based on patent citations, and then other thicket measures.

5.2 Patent citations

Patents cite other patents. When a patent applicant submits a patent application to the patent office, they are required to disclose all relevant prior art. This is sometimes called 'Rule 56', after section 1.56 'Duty to disclose information material to patentability' in Appendix R (Consolidated Patent Rules) of Title 37 in the United States. This disclosure of prior art takes the form of citations to other patents as well as scientific publications and other documents. Some academic papers do use the non-patent citations, but the citations to other patents are by far the most commonly used patent statistic after the number of patents themselves. We will use the term 'patent citations' (or just citations) to refer to citations to and from other patented prior-art.

Patent citations are inherently directional – patents filed today reference historic patents from yesteryear. Using the notation $C_{i,j}$ to denote a citation from patent i to patent j: from the perspective of the patent-holder of patent i, citations $C_{i,j}$ are citations-made, which are also referred to as backwards citations or, in graph theory terminology, the out-degree of patent i; from the perspective of patent j though, citations $C_{i,j}$ are citations-received, or forward citations, or, again in graph theory terminology, the in-degree of patent j.[53] The notion of perspective on a citation is important because what it means to cite a patent is not necessarily the same as what it means to be cited by a patent.

Before we discuss what patent citations could mean, and what we feel that they are more or less likely to mean, we return briefly to the construct validity issues of patent statistics. All of the concerns that we voiced concerning the construct validity of patent statistics are amplified dramatically in citations. Whereas our theoretical understanding of patents is relatively advanced, out theoretical understanding of patent citations is almost non-existent. To the best of our knowledge there are no published formal economic models that inform our understanding of when and why patent citations are made.[54] Although patent citations are frequently used to control for the heterogeneity in patents, citation counts have been shown to be extremely non-linear. Count distributions have very long tails and most of the observed correlation between citation counts and other measures, such as firm value, is driven by these tails. Furthermore, citation counts themselves mask an enormous degree of heterogeneity. It appears that gross citation effects are the sum of a large number of reinforcing and opposing effects driven by underlying orthogonal meaningful variation. Counts of citations within patent classes convey different information from counts across patent classes; self-citations convey different information to citations to other firms; the concentration of citations across patent-holders may matter (see below); whether citations are made to active or expired patents can reverse results; citations accruing directly after application may mean something different from citations accruing just before expiry; and so forth.[55] And perhaps most importantly, there really are very few other conceivable, yet alone implementable, measures of the relationship between innovations.[56] This has led to the use of patent citations to represent any and every type of relationship between innovations.

The Patent Citations Addendum to Grilliches' Law: *If it relates to the relationship between innovations, someone, somewhere, will claim that they can measure it with patent citations.*

[53] Bronwyn Hall and many of her co-authors prefer to use the terms backwards and forwards citations. We prefer the terms citations-made and citations-received.

[54] There are graph theory models that have been used to explain the structure of patent citation networks. For example, Leskovec et al. (2005) show that a 'forest fire' graph theory models fits patent citations rather well. However, these models do not explain the economics of patent citations. Egan (2013) provides an economic model where patent citations represent substitution of system components.

[55] See Egan (????).

[56] A notable exception is the details of the interactions and characteristics of scientists working on related projects.

Patent citations have been used to measure knowledge spillovers (e.g., Jaffe et al. 1993), knowledge flows (e.g., Jaffe et al. 2000), how 'basic' or 'applied' the innovation is (e.g., Trajtenberg et al. 1997), technological importance (e.g., Albert et al. 1991), social value (e.g., Trajtenberg 1990), private value (e.g., Hall et al. 2005), substitution relationships (e.g., Lerner 1994), complementary relationships (e.g., Ziedonis 2004), license negotiation costs (e.g., Noel and Schankerman 2006), and many, many other things. An entire special issue of 'Economics of Innovation and New Technology' was devoted to the variation in uses of patent citations. Gay and Le Bas (2005) provides an overview.

The patent thicket literature largely treats patent citations as representing complementary relationships between patents. This seems one of the more questionable usages of patent citations. Patent citations exist as evidence of patentability. Citations might be added by the applicants themselves, by their legal counsel, during a patent search, or by patent examiners. But in each case, the citation should pertain to novelty or non-obviousness. Usefulness of invention is a requirement for patentability, but of all the requirements this gets the least focus; patenting is expensive and it would be pointless to apply for patent on something that was useless.[57] Novelty and non-obviousness are established by documenting the state of the art prior to the filing of the patent. The patent applicant has an incentive to include prior art which is rendered commercially obsolete by the patent, as then, when the patent issues, it will have a presumption of validity against this prior art. Thus patent citations should document the best that was, and the patent may be an economic substitute to some of it. But an applicant (or any other party to the process) has no obvious incentive to cite an example of a complementary invention to demonstrate patentability, much less one that will be used in a future product.

5.2.1 The Hirschman Herfindahl Index (HHI) 'fragmentation' measure

Ziedonis (2004) operationalized a Hirchman Herfindahl Index (HHI) for patent ownership that she named the 'fragmentation' index or F. This measure is one minus the sum of square of patent-ownership shares. It quickly became the dominant measure of type A (diversely-held complementary inputs) patent thickets and has remained so except for its replacement by the triples measure (which was introduced in 2011 and is discussed in detail below) for EPO data. The fragmentation index is used, either in its biased (F) or unbiased ($\hat{F}$) form, and counting either all citations or just X&Y citations (discussed in the next section). Papers studying patent thickets that have used Ziedonis (2004)'s fragmentation measure include Cockburn and MacGarvie (2009), Cockburn et al. (2010), Entezarkheir (2010), Galasso and Schankerman (2010), Harhoff et al. (2012), Huang and Murray (2009), Siebert and Von Graevenitz (2010a), Siebert and Von Graevenitz (2010b), Von Graevenitz et al. (2012), and, of course, Ziedonis (2004). The identity of patent assignees (i.e., the original owners of a patent) is available in most academic patent databases.[58] This allows researchers to calculate a proxy for ownership of patent rights and calculate the dispersion of ownership over citations-made. Specifically, for each patent or patent portfolio belonging to firm i, there are citations $C_{i,j}$ from firm i to other firms j, and a total of C_i citations-made to these j firms. F is then:

$$F = 1 - \sum_{j=1}^{J} \left(\frac{C_{i,j}}{C_i} \right)^2 \quad \text{where } i \neq j$$

The fragmentation measure was designed to capture the extent to which complementary inputs are diversely-held. Thus the driving assumption behind the measure is that patent citations represent privately-held complementary inputs. The inventions detailed in expired patents lie in the public

[57] This hasn't stopped many patent applicants. See, for example, US patents 4,429,685 (a method for growing unicorns) and 4,233,942 (dog ear tubes).

[58] Note that re-assignment, exclusive licensing, or other transfer of patent rights is generally not available in most academic patent databases, so the data provides only the identity of the original assignee.

domain and will not contribute to an N-fold marginalization problem. Patents can expire for a number of reasons, including reaching their statutory term (20 yrs from application after June 1995, 17 years from granting prior to this date), failure to pay renewal fees (due at 4, 8, and 12 years after granting), terminal disclaimers, and because they have been declared invalid either by the court or the patent office (typically in a post-grant opposition). So the fragmentation measure should be calculated using only active, privately-held patents, to avoid introducing biases from industries dominated by new patent classes, industries that have experienced dramatic turn-over in the number of operational firms, or industries that have short product and so patent life-cycles. Unfortunately this is simply never done, despite the availability of data.

But the more fundamental problem concerns whether patent citations really do represent complementary inputs. We are not aware of any evidence in the literature that they do, and in fact the rise of patent citation measures based upon the notion that patents are complementary inputs appears entirely due to the patent thicket literature. Whether or not citations can represent complementary relationships is a testable hypothesis. One could, for example, determine which patents are complementary inputs into a set a products and examine the hazard rate of citations between them. However, as it stands, this is entirely conjectured. The two main counter-arguments to this notion are that: 1) there isn't any incentive or requirement for patent applicants to cite complementary art; and 2) there is an incentive, if not a specific requirement, to cite substitute art. Even small numbers of substitute patents might undermine, or actually reverse, the meaning of the measure. Suppose that patent citations sometimes represent substitution relationships and sometimes represent noise (to an economist exploring economic relations between inventions), but seldom, if ever, represent complementary relationships. Then a sector of the economy that exhibited high levels of fragmentation would indicate that many firms hold substitute patents – this would surely be the antithesis of a patent thicket, and would instead represent a competitive innovation environment.

There is another issue with the fragmentation measure that many authors have failed to take into account. Hall (2005) points out that this estimator is biased.[59] Specifically, assuming independent multinomial distributions for the count of citations to each assignee (subject to the constraint that citations across all assignees must sum to the count of all citations), the fragmentation measure will be positively correlated with the number of citations made. To rectify this situation, Hall (2005) provides a bias correction factor, which allows the calculation of the fragmentation measure $\hat{F}$ that is unbiased in expectation.[60]

$$\hat{F} = \frac{C_i}{C_i - 1} \cdot F$$

A simple numerical example is instructive. Using $F_{i,j}$ to denote the fragmentation measure for i citations-made spread across j assignees, and using equally-weighted averages on the permutations to construct the expectations in the cases where $i \neq j$, we provide the calculation of biased and unbiased, raw and expected fragmentation measures.[61] We do this for the cases of 2 citations across 2 assignees (i.e., $F_{2,2}$), 3 citations across 2 and 3 assignees (i.e., $F_{3,2}$ and $F_{3,3}$), and 4 citations across 2 and 4 assignees (i.e., $F_{4,2}$ and $F_{4,4}$), as follows:

[59] The key observations of Hall (2005) were included in an appendix of Jaffe and Trajtenberg (2002). Ziedonis (2004) implements the unbiased measure but a typographical error in the published version of her paper renders the bias correction factor incorrectly.

[60] Note that the bias correction factor isn't defined with only one citation.

[61] Using equally-weighted averages on the permutations to construct expectations is equivalent to assuming a uniform multinomial distribution.

$$F_{2,2} = 1 - 2\left(\frac{1}{2}^2\right) = \frac{1}{2} \text{ (2 ways)} \qquad \hat{F}_{2,2} = \frac{2}{1}\cdot\frac{1}{2} = 1 \qquad \mathbb{E}(F_{2,2}) = \frac{1}{2} \qquad \mathbb{E}(\hat{F}_{2,2}) = 1$$

$$F_{3,2} = \begin{cases} 1 - \left(\frac{1}{3}^2 + \frac{2}{3}^2\right) = \frac{4}{9} \text{ (3 ways)} \\ 1 - \left(\frac{2}{3}^2 + \frac{1}{3}^2\right) = \frac{4}{9} \text{ (3 ways)} \end{cases} \qquad \hat{F}_{3,2} = \begin{cases} \frac{3}{2}\cdot\frac{4}{9} = \frac{2}{3} \\ \frac{3}{2}\cdot\frac{4}{9} = \frac{2}{3} \end{cases} \qquad \mathbb{E}(F_{3,2}) = \frac{4}{9} \qquad \mathbb{E}(\hat{F}_{3,2}) = \frac{2}{3}$$

$$F_{3,3} = 1 - 3\left(\frac{1}{3}^2\right) = \frac{2}{3} \text{ (6 ways)} \qquad \hat{F}_{3,3} = \frac{3}{2}\cdot\frac{2}{3} = 1 \qquad \mathbb{E}(F_{3,3}) = \frac{2}{3} \qquad \mathbb{E}(\hat{F}_{3,3}) = 1$$

$$F_{4,2} = \begin{cases} 1 - \left(\frac{1}{4}^2 + \frac{3}{4}^2\right) = \frac{1}{4} \text{ (4 ways)} \\ 1 - \left(\frac{2}{4}^2 + \frac{2}{4}^2\right) = \frac{1}{2} \text{ (6 ways)} \\ 1 - \left(\frac{3}{4}^2 + \frac{1}{4}^2\right) = \frac{1}{4} \text{ (4 ways)} \end{cases} \qquad \hat{F}_{4,2} = \begin{cases} \frac{4}{3}\cdot\frac{1}{4} = \frac{1}{3} \\ \frac{4}{3}\cdot\frac{1}{2} = \frac{2}{3} \\ \frac{4}{3}\cdot\frac{1}{4} = \frac{1}{3} \end{cases} \qquad \mathbb{E}(F_{4,2}) = \frac{5}{14} \qquad \mathbb{E}(\hat{F}_{4,2}) = \frac{10}{21}$$

$$F_{4,4} = 1 - 4\left(\frac{1}{4}^2\right) = \frac{3}{4} \text{ (24 ways)} \qquad \hat{F}_{4,4} = \frac{4}{3}\cdot\frac{3}{4} = 1 \qquad \mathbb{E}(F_{4,4}) = \frac{3}{4} \qquad \mathbb{E}(\hat{F}_{4,4}) = 1$$

Using the fully fragmented measures, where $i = j$, we can see Hall (2005)'s concern, as well as her solution. $\mathbb{E}(F_{i,j=i})$ goes from $\frac{1}{2}$ when $i = 2$, to $\frac{2}{3}$ when $i = 3$, to $\frac{3}{4}$ when $i = 4$. As the number of citations-made, i, increases, the raw fragmentation measure asymptotically converges to 1 and its bias drops to zero. But for small counts of citations-made the raw fragmentation measure is obviously strongly positively correlated with the count itself. The bias correction factor equally clearly rectifies the problem: $\mathbb{E}(\hat{F}_{i,j=i})$ remains constant at 1 when the assignment of citations-made is fully fragmented irrespective of the count of citations-made.

However, there are two other problems with this measure. The first concerns the prima facie construct validity. The fragmentation measure is supposed to capture the extent to which cited patents (hopefully patents on complementary inputs) are diversely-held. It simply doesn't do this; instead, and just as its label implies, it measures the *fragmentation* of ownership. Suppose we hold the number of assignees constant at two – that is suppose that no matter how many citations a firm makes, all of those citations are always to patents held by just two firms. $\mathbb{E}(\hat{F}_{i,j=2})$ goes from 1 when $i = 2$, to $\frac{2}{3}$ when $i = 3$, to $\frac{10}{21}$ when $i = 4$.[62] Irrespective of whether we use the biased or unbiased fragmentation measure, the value drops steady while the economic situation remains the same – there are just two patent-holders to negotiate with.

Holding the number of citations-made constant (and putting aside issues of whether patent citations do actually measure economic relationships), the fragmentation of assignees can capture the competitiveness of the patent landscape. Thus we should always see the fragmentation measure implemented alongside the count of citations-made. But we never do – not in a single paper in our sample.

The second problem with this measure is much more subtle. Hall (2005) is explicit that her bias correction relies on the assumption of independence. She says that it must be the case that "there is no causal connection between the deviation of the observed outcome from the expected outcome in a particular cell and what happens in another cell" (other than through the "adding up constraint"). Put another way, she requires that assignees do not strategically respond to one another's holdings. When Ziedonis (2004) inaugurated the unbiased fragmentation measure, she ironically pointed out that when "the primary motives for patenting... are driven by concerns about strategic positioning, it is important

[62] With non-uniform multinomial distributions the expected values would be different because the permutation would have non-equal weights in the summation, but the same problem would persist as the expected values would still be convex functions of the potential outcome values.

to step back and question what is being captured in these [measures].”

5.2.2 The X&Y triples measure

Von Graevenitz et al. (2011) introduced a new measure of patent thickets that has quickly become the dominant measure for researchers working with European data. It is used in Hall et al. (2012), Harhoff et al. (2012) and Von Graevenitz et al. (2012). This measure is based upon what are called ‘X’ and ‘Y’ citations. Applications made to the European Patent Office (EPO) (including the filing of international patent applications filed under the Patent Cooperation Treaty (PCT) when the EPO acts as the designated or elected office) are subjected to a search process prior to their examination process. The Search Division of the patent office draws up a search report using both in-house and external collections of prior-art. It is accepted that this search is not exhaustive, but each and every claim in the patent application should be subjected to search against the prior-art.[63]

The search examiner may use citations to prior-art (whether patented or not) submitted by the applicant and may add their own citations. These cited documents are then appended to the search report. The search report specifies the relationship between the documents and the individual claims to which the document applies. Every document must bear upon on or more claims. Relationships between documents and claims are classified using one of the following codes:[64]

- A – the document defines the state of the art and does not prejudice novelty or the inventive step
- X – when taken alone the document provides evidence that the claimed invention cannot be novel or to involve an inventive step
- Y – combined with one or more other documents, the document indicates that the claimed invention is not non-obvious (i.e., in combination with the other documents, the invention is obvious)

The patent applicant is provided with a copy of the search report and is required to address all issues relating to X&Y citations. The applicant could do this in a number of ways. Specifically, they could: drop the claims involved from the application (providing the invention still met the requirements for ‘unity of invention’); provide evidence that the invention does indeed meet the novelty and non-obviousness requirements; modify the claims so that the invention met the novelty and non-obviousness requirements (if this modification required a substantive change then it would likely require a continuation application, and possibly a division of the patent into separate applications); or abandon the patent application. Regardless of which action they take, the citations marked as X&Y remain on the patent’s record with the EPO, and are included in the PATSTAT data distributed by the EPO.[65]

With this in mind, the triples measure is defined by an algorithm as follows:[66]

1. Let C^{XY} be the set of X & Y citations, such that $c^{XY}_{i,j} \in C^{XY}$ is an X or a Y citation from firm i to firm j
2. A reciprocating pair of X and/or Y citations R_{ij} is then a two element set: $R_{ij} = \{c^{XY}_{i,j}, c^{XY}_{j,i}\}$ where $i \neq j$

[63] When the claims of the invention “do not relate to one invention only, not to a group of inventions linked so as to form a single general inventive concept” the search examiner will normally be restricted to the first mentioned invention, but this restriction will be noted for the substantive examination. See EPO (2009), Part B, Section III-6.

[64] A relationship may also be marked with other letters, including ‘T’ (to indicate that the document relates to underlying theory), ‘D’ (to indicate the citation was mentioned in the patent application description), ‘E’ (when there is a potential priority date conflict that might undermine novelty), and so forth. See EPO (2009), Part B, Section X, p. 194.

[65] PATSTAT is the EPO Worldwide Patent Statistical Database, is update twice annually, and is one of the major sources of patent data used by academic researchers.

[66] Von Graevenitz et al. (2011) provide a written description of this algorithm. We reformulate it into mathematical notation to remove ambiguity. This reformulation deliberately does away the judgment-laden connotations implied by the authors, who are very strong on the notion that their measure captures ‘blocking’. It also makes it clearer that the ultimate measure is at the patent-holder (i.e., firm) level, despite the requirement for an analysis of the citations to and from individual patents within the patent-holders’ portfolios, which could be aggregated in a number of different ways.

3. A triple H_{ijk} is a set of patent-holders $\{i, j, k\}$ such that $\exists\{R_{ij}, R_{jk}, R_{ik}\}$ where $i \neq j \neq k$
4. The triples measure is the count of all *unique* H_{ijk} (i.e., each H_{ijk} has a different set of patent-holders with corresponding reciprocating pairs) in a technology area: $|!H_{ijk}|$

There are two ways to specific the technology area that the triples cover:

- In Von Graevenitz et al. (2011)'s original definition, C^{XY} in step 1 is the set of X & Y citations in a given technology area, which is defined using patent classes.
- An alternative definition could be implemented at the industry (of the patent-holders) level. In this case one should require that i, j, and k all operate in the same industry, for example defined using North American Industry Classification System (NAICS) or Standard Industry Classification (SIC) codes, in step 3.

Before we ask what these X&Y citations can mean, there are obvious mechanical considerations in the measure's assembly. All else equal, in particular supposing that the probability of a citation being of either type X or Y is constant across patents, the triples measure will be correlated with the size of firms' patent portfolios in a given technology area, the number of citations patents in the portfolio make, and the number of owners-cited. The size of a firm's patent portfolio might itself be a function of the number of firms in a technology area: a technology area with many patent-holders might have more patenting opportunity because of a selection-effect; likewise, stronger incentives to pursue patenting for rent appropriation, and stronger strategic incentives for patenting too, would suggest a treatment effect for these areas.

Patents that make one citation (or no citations) can't participate in triples. This perhaps isn't much of a concern, but patents that receive one or less citations is an issue that should be addressed. Citations take time to accumulate, triples require citations both made and received, and we should anticipate data-truncation issues with respect to citations-received. Generally though, there is a greater chance of finding a triple in areas where the typical patent makes (and receives) many citations. And this follows through for owners-cited: the greater the typical spread of citations-made across assignees in a technology area, the more likely it is that one would find triples. Taken together, this reasoning suggests that analysis using triples should take into account the size of the technology area, its existing density of patents (and perhaps its future patenting opportunity), and the characteristics of patents in the area in terms of the typical number of citations made and received as well as the typical dispersion of ownership across citations. This could be done by making adjustments to the measure itself, or through the use of control variables. Unfortunately, none of the papers in our sample that use the triples measure in an analysis of the effects of thickets does either.

The semantic content of an X&Y citation is open to debate. Von Graevenitz et al. (2012) report a correlation of greater than 0.3 between their triples measure and the HHI-based fragmentation measure in their sample, but it is clear from their regressions that both measures have orthogonal meaningful variation.[67] Von Graevenitz et al. (2011) argue that triples capture 'blocking'. They provide evidence that in areas identified as complex technologies by Cohen et al. (2000), the incidence of triples has increased steadily since the 1980s. Whereas in areas characterized by discrete technologies the incidence of triples has remained constant. They further claim that this is not driven by the number of patent applications in these areas and provide a figure as supporting evidence. Finally, they provide a "description validation" of their triples measure by using it to determine that some technology areas are characterized by thickets.[68] Tautologies aside, the foundation of Von Graevenitz et al. (2011)'s argument rest on an assertion. They

[67] When Von Graevenitz et al. (2012) include both measures in the same regression equation, they achieve statistical significance with opposite signs.

[68] Von Graevenitz et al. (2012) states that 9 out of 30 areas studied in Von Graevenitz et al. (2011) are characterized by thickets. This again raises the question of thresholds in measures, and at what level of triples (or fragmentation, or the number of patents, etc.) a technology area is declared a thicket.

say that they "provide a measure of the density of a patent thickets based on triples of firms that can mutually block some of each others' patents". Whether or not this could be true depends how the patent-holders that received search reports with X&Y citations responded to them, as well as whether the citations were X *or* Y.

If a patent applicant given an X citation on a particular claim is able to drop the claim with compromising the unity of their invention, and elects to do so, then it is possible that the cited patent is a complementary input into the citing patent. Thus the cited patent is a blocking patent with respect to invention (i.e., triples might measure type B thickets). It is also possible, as dropping the claim didn't compromise the unity of the invention, that the claim wasn't truly required for the invention. In this case an X citation may still indicate that both the citing and cited patents are likely to be complementary to one another and both act as inputs into a product. Thus the cited and citing patents could be both blocking patents with respect to the commercialization of some product (i.e., triples might measure type A thickets).

However, if the patent-holder takes any other action, or the applicant is given a Y rather than an X citation, the semantics of the measure are dramatically altered, if not reversed. Y citations indicate that the patent application's claim does not cover an invention because the 'invention' is obvious given the prior art (of which the cited patent is but a single piece). No matter the actions of the patent applicant, this does not confer blocking power to the cited patent. In the best case for establishing blocking, where the applicant withdraws the claim without compromising unity and the claim does represent a complementary input, the ownership of the input lies in the public domain as given the state of the prior art the 'invention' is obvious and so not patentable. That said, triples based on Y citations could still measure thickets. One possibility is that many reciprocal Y citations could be taken to indicate that invention spaces have become effectively saturated and that it is now difficult to conceive new inventions given the prior art (i.e., triples might measure type D thickets).

If a patent applicant given an X or Y citation successfully contests the classification of the prior art, then it is difficult to make definitive statements concerning the real relationship between the citing and cited patents. In such an instance, the assignment of the X or Y citation by the Search Division was presumably a mistake. One possibility for such a mistake is that although the claims genuinely do embody different inventive steps, their claimed domain of application is very similar (i.e., triples might measure type $B2$ thickets). And it is also possible that the Examination Division makes mistakes. Specifically, the patent office might (at least in some cases) grant a patent despite evidence that it isn't novel or non-obvious (i.e., triples might measure type $C1$ thickets).

And finally, if the patent applicant responds and is able to secure a patent with revised claims, this would provide fairly strong evidence of 'invent-around' and the creation of substitutes. Patent applicants given X&Y citations might well have the strongest incentives and best available capabilities to do this. They have already filed a patent application, presumably with some commercial purpose in mind, they know the pre-existing technology very well (as they have just tried to file a patent on it), and they now know which parts of their invention need fixing and what they have to differentiate their invention against. Legitimate substitute inventions, particular those that the patent office has effectively certified as substitutes through the granting of a patent despite X&Y citations, would undermine the existence, or innovation hindering problems, of type A, B, $C1$, and D thickets.

In conclusion, X&Y citations, and so the triples measure, could indicate that there are (type A, B, $C1$ or D) patent thickets, or that substitution is eliminating the problems of these patent thickets, or could measure something that has nothing to do with patent thickets. And the measure itself suffers from a host of mechanical issues in its construction that need to be addressed. Recent work on European data has adopted this measure whole-heartedly; we suggest that without an understanding of the measure's semantics and biases, this adoption has been premature and more work establishing its fundamentals should be undertaken.

5.2.3 Other measures of thickets

Despite the dominance of Ziedonis (2004)'s fragmentation measure overall, and Von Graevenitz et al. (2011)'s triples measures for EPO data in recent years, most papers use more than one measure of patent thickets. As a consequence, there are almost as many measures of patent thickets as there are papers about them. The three next most commonly used measures, in approximate order, are the number of patents, the CR4 (four firm concentration ratio) of assignees, and the number of owners-cited.

The number of patents is generally calculated in a technology area, defined by patent classes. Patent classes are extremely noise measures, so alternative bases of aggregation, particularly the industries of patent-holders or by competing product groups, are probably preferable.[69] One interesting exception is Reitzig (2004) who used questionnaires to determine approximately how many patents were protecting a particular invention.

Noel and Schankerman (2006), Geradin et al. (2007), and Galasso and Schankerman (2010) all use the CR4 of patent citations. The CR4 measures the share of patent citations held by the four largest patent assignees (in terms of citations). Specifically, let $s_j = \sum_j C_{i,j}$ for some set of patents i citing a firms j in a particular technology area, and then relabel these shares in order of descending magnitude, $s_1, s_2, \ldots, s_j$, so that s_1 is the largest share and s_j is the smallest. The CR4 is then $s_1 + s_2 + s_3 + s_4$. This measure is clearly analogous to the Ziedonis (2004)'s fragmentation measure, and shares the same conceptual basis. However, it suffers from some additional problems as it is unable to differentiate between sets of different concentrations among the top four patent assignees (i.e., one assignee owning all of the cited patents produces the same measure as any sharing of the cited patents among the four firms) and it doesn't capture the spread of cited patent ownership among patent-holders outside of the top four.

Noel and Schankerman (2006) and an earlier version of Cockburn et al. (2010) both used the total number of patent assignees cited. This measure has a very simple intuitive appeal. It still relies on the assumption that citations represent complementary relationships, but if we accept this proposition then the measure directly captures the extent to which complementary inputs are diversely-held. As such it could represent the size of the thicket for type A thickets, and could proxy (if one assumes that patents can always be licensed together in a single transaction from a single owner) for transactions costs in all thickets. We hope that this measure will received renewed attention in the future. However, we caution that its obvious mechanical correlation between owners-cited and citations-made should be corrected for in any analysis.

There are a large number of 'custom' measures based on patent citations used in the thicket literature. These range from simple weighting of patent counts by citations-received to exceedingly elaborate measures that have their foundations in network or graph theory. Siebert and Von Graevenitz (2010b) and Siebert and Von Graevenitz (2010a) provide examples of the latter. We are unable to disentangle their semantics, and generally take the position that combining many different patent statistics in complicated fashions is likely compound measurement issues and obscuricate, not clarify, construct validity. One possible exception to this objection is Nagaoka and Nishimura (2006). Nagaoka and Nishimura (2006) use the RIETI inventor survey to determine the average number of complementary patents needed for commercialization in each industry and use it weight Ziedonis (2004)'s fragmentation measure.

[69] Patents are assigned classes based on the nature of the technology and not the area of use of the technology. One exception is the International Patent Classification scheme used in Lerner (1994). A patent examiner determines the 'controlling claim' of the patent application and assignees a 'primary' or 'original' class to the patent accordingly. Although patents are assigned many classes for cross-reference purposes, generally only this original class is used by researchers. There are over 425 main classes in the US patent classification system, each with many sub-classes. Aggregating these classes and sub-classes into meaningful units of analysis is also hugely problematic. All of these features of the classification system, along with human error, make patent classes poor measures for analysis at the product or product-market level.

Claims-based measures offer a potential view into type C thickets. There has been a long-standing assumption in the literature that the number of claims proxies for the scope of the patent. This has been frequently extended so that claims-made proxy for the (implicitly unmerited) breadth of the patent, and so the extent to which a patent is spurious. As we discussed previously, we do not thing that a broad patent should necessarily be considered spurious. Unmerited breadth, that is the extent to which the scope of application exceeds the scope of the inventive step, would indicate some level of spuriousness, but it is far from clear that this could be measured with the number of claims relative to an industry or technology area average, let alone the number of claims itself. Hall and Ziedonis (2001), Huang and Murray (2009), and Cockburn and MacGarvie (2009) all implement claims-based measures.

A final patents statistic based measure of patent thickets is the non-patent share of prior art, used by both Cockburn and MacGarvie (2009) and Murray and Stern (2007). However, outside of patent statistics, there have also been some other measures of thickets. Hall and Ziedonis (2001) and Hall and Ziedonis (2007), as well Ziedonis (2004), took the view that patent thickets were related to a hold-up problem. They tell a story of 'defensive patenting' to secure rights that will be needed by other firms, so that a firm can protect itself (essentially through mutually-assured destruction) from hold-up by its rivals. Using Williamson-like reasoning, they argue that firms that make the greatest relation-specific investments face the greatest threat and have the greatest incentive to respond with defensive patenting. Accordingly, they use measures of the capital intensity of the firm as proxies. Hegde et al. (2009) likewise use the capital intensity of firms to distinguish between 'pioneering inventors' and firms that are more likely to be engaged in defensive patenting. And Reitzig (2004)'s unique questionnaire data allows him to include measures of whether patents are reported to be used to prevent copying or create revenues through royalties or licensing. This paper also uses Levin et al. (1987) to classify industries as being characterized by discrete or complex technologies.

5.3 Testing the effects of patent thickets

Now that we understand what the phenomenon is (e.g., type A, B, C, and D thickets, perhaps with variants T, S, or P), and that there is at least some hope of measuring the phenomenon, we turn to how one could estimate its impact. Researchers would like to answer two questions: 1) Do patent thicket exist? And 2) Do patent thicket *cause* an adverse effect on (i.e., hinder) innovation.

The language of patent thickets used in the literature, and particularly by policy makers, suggests binary states for patent thickets: Either a thicket exists or it doesn't. Our theoretical foundations, and the measures implemented, on the other hand, both generally suggest a continual phenomenon.[70] Patent thicket measures take the form of either counts of patent statistics or ratios of patent statistics. And using type A thickets as an example, a 'thicket' exists the moment that a single complementary input is required and increases in 'density' with the number of distinct patent-holders that own complementary inputs. A rephrasing of the first question to mix it with the second question would then be to ask: "At what point do patent thickets become onerous?"

Transaction costs and search costs are always economic losses, are privately incurred, and are applicable to every type of patent thicket in our taxonomy. We might try to create policy to minimize them but we don't expect to eliminate them (and we do expect firms to try to minimize them). So there is the possibility that when search and transaction reach a certain level, new inventions or new products become privately infeasible. These new innovations might well have social benefits in excess of their private returns. Nordhaus (2004) put the typical fraction of the return accruing to the inventor in the U.S. at 2.2%. This would suggest a very sizeable welfare loss, and provide normative grounds for subsidizing

[70] An exception is the appeal to cognitive limitations, or other artifacts of behavioral economics, in search costs. If, for example, there is an inviolable search threshold, beyond which firms will surely fail to find all of the required inputs for a type A thicket (diversely-held complementary inputs), then a binary world view is truly applicable.

search and transaction costs. However, although search and transaction costs may be material, each patent thicket type has its own unique welfare considerations that are likely far more important.

With type A thickets which characterize the problem of diversely-held complementary inputs and lead to a potential N-fold marginalization problem, one economic loss comes from the deadweight-loss arising from under-supply of a new product. Measuring type A thickets with dependent variables that capture invention and not commercialization then entirely misses this economic loss. Firms have incentives to mitigate type A thickets by inventing-around existing patents on complementary inputs, filing defensive patents that will be used by competitors producing other products, building reputations for fair pricing, engaging in credible commitments such as FRAND agreements, pooling patents, engaging in cross-licensing, and so forth. If they don't mitigate this problem, the rents to invention will be compromised, adding another economic loss from the reduced incentive to innovate. We can categorize their mitigation strategies into two distinct types: Mitigation strategies that rely upon more invention; and those that use other non-invention related strategic actions.

In a world of Schumpeterian innovation, the more inventions that are made available for new combinations, the healthier the innovation ecosystem, regardless of whether these patents were filed for strategic purposes or not. A positive correlation between thicket measures and invention outcomes, which might inform us about the extent to which strategic patenting occurs (assuming we could genuinely identify strategic patenting as a response to a thicket from other patenting activity), can say nothing about the point at which patent thickets become onerous. A negative correlation, however, could indicate that the rents to invention have been compromised. However, in a world where innovation equals invention plus commercialization, the dependent variable for a type A thicket should capture some aspect of commercialization. This might allow us to make statements along the lines of "beyond a certain threshold of type A thickets (however measured), we can observe a statistically significant drop in commercialization activity". In either world, dependent variables that capture non-invention related strategic actions have little to say about the welfare consequences of patent thickets. They can only provide suggestive evidence that firms are responding to thickets (though again we would stress the need for clean identification).

With type $B1$ thickets, which characterize the problems of legitimate overlapping patents in cumulative invention, one important economic loss comes from an inefficient allocation of rents. Scotchmer (1991) argues that there are grave risks of undercompensating inventors of foundational technologies, especially when the foundational technology has very little value on its own. With a single foundational patent, this perhaps is not too much of an issue. But with multiple foundational patents, all of which are (rightfully, according to Scotchmer 1991) demanding large royalties and complex Reach-Through Licensing Agreements (RTLAs), there is the possibility that the 'royalty stack', negotiation costs, or conflicting restrictions become prohibitive and new inventions are not undertaken. And on the other hand, cross-licensing agreements or other arrangements that do not accord foundational patents the rents they deserve may undercompensate these inventions.

A second important economic loss with type $B1$ thickets comes from the possibility of hold-up. Recall that we have argued that RTLAs are the one place that Williamson hold-up is truly possible. With an RTLA, there is a contract that may well be incomplete, is about to be subject to relation-specific investment, and is subject to opportunism. As such economists anticipate market failure in the form of inefficiently low levels of investment. Thus, for type $B1$ thickets, it is appropriate to measure the effects of thickets on invention, and a statistically significant reduction in invention associated would be grounds for concern.

With type $B2$ thickets, the welfare loss comes from inefficient litigation. Type $B2$ thickets concern overlapping adjacent inventions, and in this case two or more patents at least partially substitute for one another. However, the earlier patent-holder may believe that they have an exclusionary right. This is particularly possible with $B2$–P thickets, where courts are imperfect arbitrators of patent rights and the earlier patent-holder can rely on the probabilistic nature of their patent. Thus we might expect to see litigation to prevent the use of a new invention, which aside from the welfare loss involved in the litigation costs, may result in inefficient levels of invention and so commercialization.

Which leaves types C and D thickets. The former essentially *is in itself* a welfare loss and the later has no alterable welfare connotations. Type C thickets are those based on spurious patents. Type $C1$ are invalid and incorrectly issued by the patent office and type $C2$ are submarine patents. In either case the their presence is expected to cause the threat of litigation, and with variant P this threat is likely be to successful. Patent-holders of spurious patents have created no new invention themselves and so are not adding to economic welfare in any way. They can only detract from it. Type D thickets, on the other hand, represent effectively saturated invention spaces. Although we might expect these spaces to be cleared by radical inventions that alter the entire paradigm of invention, or through the eventual expiry of the patents, it is still possible that for periods of time there is simply nothing further to invent in an area. This means that no new welfare can be created, except through Schumpeterian innovation, but also that there is no new potential welfare to be lost.

5.3.1 Determining causality

Even with good measures, questions of causality are exceedingly difficult to address for patent thickets. To move beyond suggestive correlations, studies need to implement an instrument that shocks the explanatory variable without affecting the outcome variable directly. As we will shortly see, some of the largest and most well-known studies use patent statistics on both sides of their regression equations, making this immediately problematic. Generally these studies would require an instrument that shocked the relationship between innovations without shocking the volume of innovations. Appeals to regime changes, such as the creation of the Court of Appeals for the Federal Circuit (CAFC) in 1982 to hear patent appeals, or the announcement of decisions of important suits, such as Diamond v. Chakrabarty, 447 U.S. 303 (1980) or State Street Bank and Trust Company v. Signature Financial Group, Inc., 149 F.3d 1368 (1998), are (in our opinion) rightly claimed to affect the strength and enforceability of patent rights. However, it is hard to argue to these affect the relationships between patent rights and not the volume of patent application likely to be filed. This problem persists when authors use outcome measures that represent the value of patent rights, firms' strategic usages of their own patent rights, or firms' strategic responses to other patent-holders' rights.

A firm's decision to license, cross-license, litigate, or block usage can be a strategic use of their own patent rights. Likewise, the filing of continuations on their own patents or post-grant oppositions to a rivals patents can both represent a firm's strategic responses to other patent-holder's rights. The strengthening of patent rights alters the incentives to pursue these strategic actions irrespective of patent thickets, and so instruments based on changes to the strength of patent rights fail the exclusion restriction. Moreover, strategic responses to patent thickets, such as the pursuit of complementary inventions that might be used by rivals, are themselves a cause of patent thickets. We therefore face the classic problems: causation works in both directions, and both thickets and their response might be the effect of an underlying omitted variable.

5.4 Measures and tests from the literature

In table 6, below, we report some descriptives for the 20 papers that constitute the near-population of papers that have done empirical research on patent thickets, as well as some of their main-effect results. Included in these descriptives are the dependent variables studied. Because some papers report results for multiple dependent variables, a paper may appear in multiple rows of the table. The table is organized into three sections. The first section reports results for papers that study the effects of thickets on patenting or research. Dependent variables in this section include the propensity to patent, R&D expenditure, the rate of invention, and the accrual of academic citations (for the two papers that study whether patent thickets hinder research in academia). There are two countervailing possibilities: either thickets increase inputs to innovation, which would be consistent with a defensive patenting response to type A thickets; or they decrease inputs to innovation, which might be consistent with type B, C or D thickets.

We include Hall and Ziedonis (2001) and Ziedonis (2004) in this section. Even though these papers phrased their arguments in terms of a Williamson-like hold-up story, exactly the same mechanism of defensive patenting could be used as a disciplining mechanism to mitigate the problems associated with N-fold marginalization. For example, suppose that each patent-holder is also a producer of a product. Each faces the N-fold marginalization problem when assembling their product, but each can use its patented inputs into its rivals products as a disciplining mechanism to force rivals to cooperative in their pricing to achieve a joint-profit maximizing mark-up instead.[71]

The second section reports results for papers that study the effect of patent thickets on either firm value or patent value. We do not have clear priors for the effects of type A, $B1$, or D thickets on value: Type A thickets might reduce firm value, as N-fold marginalization problems reduce the rent to invention, but defensive patenting could reverse this effect; type $B1$ thickets, based on overlapping patents in cumulative invention, might lead to under compensation for inventors of fundamental inventions and over compensation to inventors of follow-on inventions, or a market failure due to the threat of hold-up under Reach-Through Licensing Agreements (RTLAs) and so inefficiently low follow-on invention; as type D thickets increase, or invention spaces become more saturated, each additional invention could be more valuable or could be more marginal; and for thickets based on spurious patents (type C) value should, assuming the filing of such patents is rational, accrue to the spurious patent-holder. But such patents enable transfer-seeking, and we would expect the holders on legitimate patents (as well as the developers of new products) to suffer as a result.

The third and final section of the table reports the effects of patent thickets on strategic outcomes like licensing, litigation and blocking. We also included Cockburn and Macgarvie (2011) and Cockburn and MacGarvie (2009) in this section. These two papers study entry, the number of financings, the amount of investment, and the likelihood of undertaking an IPO for software firms in the presence of patent thickets. Of all of the dependent variables used in an analysis of the effects of patent thickets, we find these the most credible. New entrants, and their investors, presumably make their entry and investment decisions taking the thicket in an industry as given.

[71] Hall and Ziedonis (2001) and Ziedonis (2004) use the interaction between capital intensity of firms and a thicket measure. This does not identify their hold-up story. The disciplining mechanism we just described is more important to product producers who get rents from both their patents and their products (and who can be expected to have a higher capital intensity) than to non-participating entities. Increased levels of capital intensity could also be correlated with lower number of competing firms.

Table 6: Measures, tests, and main effects for empirical papers on patent thickets

The table provides a high-level review of the empirical results for papers that have, in some material way, provided evidence about the existence of patent thickets and/or whether or not they provide a hindrance to innovation. Note that only main effects are reported. Many of these papers use interaction effects to model the influence of patent thickets – these are indicated with a ✓ in the Interaction Effects column. For these papers the actual findings are much more nuanced than the simplified results presented; in many cases the interpretation of the interaction results by the authors leads them to the opposite conclusion from what one might expect from the main effects. $\uparrow$ and $\downarrow$ indicate statistically significant positive and negative correlations (or equivalently odds ratios of more than and less than one) respectively. ***, **, and * indicate statistical significance at the 0.01, 0.05, and 0.1 levels respectively. '~ 0' indicates results that are not statistically significant from zero. '?' indicates that a variable was used in an analysis but statistical significance (or standard errors) of the measure was not reported. Hall and Ziedonis (2001) use both forward citations (i.e., citations received) and the number of claims per patent to examine whether the quality of the average patent in the semiconductor industry declined. They say that these measures provide "weak evidence at best". When authors reported multiple results we chose the main result when one was clearly apparent and a typical result otherwise. The 'HHI' (Hirschman Herfindahl Index) over assignees, otherwise known as the 'fragmentation measure' introduced by Ziedonis (2004), is implemented using either just 'X&Y' citations (European data only), as discussed in the main text, or using all citations, and with either the bias correction factor introduced by Hall (2005) to make it unbiased, or not. These distinctions are included in brackets next to the results (i.e. (X, ✓) to indicate not just X&Y citations were used in a calculation of an unbiased HHI measure). Cockburn et al. (2010) has data from surveys for years 1993, 1995, 1996, 2000 and 2004, but are reported as covering the period 1993-2004 for simplicity.

						Interaction	Dependent	Patent Citation Based Measures				Non-Citation Pat. Stats			
	Paper	**Instrument**	**Unit**	**Period**	**N**	**Effects**	**Variable**	**Triples**	**HHI (X&Y, Unbiased?)**	**CR4**	**Other**	**No. Pats**	**Claims**	**Other**	**Other**
Patenting/Research	Hall and Ziedonis (2001)	1984-1988	Firm-Year	1975-1999	946	✓	Patenting				?		?		$\uparrow^{***}$
	Hall et al. (2012)		Firm-Class	2001-2009	~1m	✓	Patenting	$\downarrow^{***}$				$\uparrow^{***}$			
	Huang and Murray (2009)		Patent-paper pairs	1988-2005	1,279	X	Academic Citations		$\downarrow^{**}$ (X,X)			$\downarrow^{**}$	$\downarrow^{**}$		
	Nagaoka and Nishimura (2006)		Firm-year	2007-2010	~6k	X	Patenting				$\uparrow^{**}$				
	Nagaoka and Nishimura (2006)		Firm-year	2007-2011	~6k	X	Invention				$\downarrow^{***}$				
	Noel and Schankerman (2006)		Firm-year	1980-1999	~1k	X	Patenting & R&D			$\downarrow^{**}$					
	Murray and Stern (2007)		Patent-paper pairs	1997-1999	169	✓	Academic Citations					$\uparrow^{**}$		$\downarrow^{***}$	
	Von Graevenitz et al. (2012)		Patents	1980-2003	~175k	✓	Patenting	~ 0	$\uparrow^{***}$ (✓,✓)						
	Ziedonis (2004)		Firm-year	1980-1994	667	✓	Patenting		$\uparrow^{**}$ (X,✓)						
Value	Entezarkheir (2010)		Firm-Year	1979-2002	~68k	X	Tobin's Q		$\downarrow^{**}$ (X,X)						
	Geradin et al. (2007)		Firm-Year	1980-2005	371	X	Market Value			~ 0	~ 0				
	Noel and Schankerman (2006)		Firm-year	1980-1999	~1k	X	Firm Value			$\uparrow^{**}$					
	Reitzig (2004)		Patents	1994	612	✓	Patent Value								~ 0
Strategic Outcome	Cockburn and Macgarvie (2011)	1994-1998	Firm-Year	1990-2004	~150k	✓	Entry				$\downarrow^{**}$	$\downarrow^{***}$			
	Cockburn and MacGarvie (2009)	1994-1998	Market-Year	1990-2004	231	✓	Financings		$\uparrow^{***}$ (X,✓)			$\downarrow^{*}$	~ 0	~ 0	
	Cockburn and MacGarvie (2009)	1994-1998	Market-Year	1990-2005	232	✓	Investment $					$\uparrow^{*}$	~ 0	~ 0	
	Cockburn and MacGarvie (2009)	1994-1998	Market-Year	1990-2006	233	✓	IPO					~ 0	$\downarrow^{**}$	$\uparrow^{**}$	
	Cockburn et al. (2010)		Firm-Year	1993-2004	1,616	✓	Licensing		$\uparrow^{**}$ (✓,✓)						
	Galasso and Schankerman (2010)	CAFC (1982)	Cases	1975-2000	5,131	✓	Litigation Duration		$\uparrow^{**}$ (X,✓)	$\uparrow^{**}$	$\downarrow^{***}$				
	Hall and Ziedonis (2007)		Cases	1973-2001	547	X	Litigation								$\uparrow^{***}$
	Harhoff et al. (2012)		Patents	1980-2010	~1m	X	Post-grant Opposition	$\downarrow^{***}$	$\uparrow^{*}$ (X,✓)						
	Hegde et al. (2009)	ΔTerm (1995)	Patents	1981-2000	~360k	✓	CAP Continuation								$\uparrow^{**}$
	Hegde et al. (2009)	ΔTerm (1995)	Patents	1981-2001	~360k	✓	CIP Continuation								$\downarrow^{**}$
	Nagaoka and Nishimura (2006)		Firm-year	2007-2011	~6k	X	Cross-license				$\uparrow^{***}$				
	Nagaoka and Nishimura (2006)		Firm-year	2007-2011	~6k	X	Block Usage				$\downarrow^{**}$				
	Siebert and Von Graevenitz (2010a)		License	1989-1999	847	✓	Ex-ante Licensing				$\uparrow^{***}$				
	Siebert and Von Graevenitz (2010a)		License	1989-2000	848	✓	Ex-post Licensing				$\uparrow^{***}$				
	Siebert and Von Graevenitz (2010b)		License	1989-1999	847	✓	Ex-ante Licensing		$\downarrow^{*}$ (X,✓)		$\uparrow^{**}$	$\downarrow^{**}$			
	Siebert and Von Graevenitz (2010b)		License	1989-1999	847	✓	Ex-post Licensing		$\downarrow^{*}$ (X,✓)		$\uparrow^{***}$				

The second column of table 6 reports whether or not the paper used an instrument in some form, and whether or not the conclusions of the paper were driven by interaction effects. The main columns of table 6 reports the main effects of thickets variables. Although we felt that this was strictly better than simply marking which papers had used which measures, the main effects inevitably miss the nuance of some of the papers' results. Cockburn and Macgarvie (2011) provides a useful example. They use the changes in patent regime strength in the software sector from 1994-1998 as an instrument. They report results from before, during, and after this period, and make inferences based on the differences between the before and after period. We report their main result that entry decreases as measures of patent thickets, such as the number of patents in portfolios of incumbent firms, increase. This result is certainly informative. However, their conclusion is that "this may be mitigated by the stimulating effects of stronger patent protection".

It is hard to discern any meaningful patterns from the main effects reported in table 6. In every section of the table, there are approximately as many positive effects of patent thickets on dependent variables as there are negative effects. And within many papers the main effects of different patent thickets variables go in different directions. But the core thesis of this work is not that nothing has been learnt about patent thickets. Instead, we argue that a lack of clarity in the definition of a patent thicket, combined with poor construct validity in the measurement of thickets, has prevented clear conclusions from emerging. Table 7, below, attempts to summarize the problem we face in trying to disentangle the truth.

5.5 Claims of thicket existence and hindrance

Table 7: Empirical evidence regarding the existence and hindrance of patent thickets

The table provides summary data for the empirical papers that have, in some material way, provided evidence about the existence of patent thickets and whether or not they provide a hindrance to innovation. The Definition column indicates which thicket types from our taxonomy the paper's definition of a patent thicket is consistent with. The Region and Sector columns pertain to the paper's dataset's coverage. The Paper Claims columns represents our best understanding of the paper's claims regarding the existence of patent thickets (in the industries or certain sub-industries studied in their data) and whether or not said patent thickets caused a hindrance to innovation. Some authors' claims were implicit. The Implications column indicates whether the measures and tests in a paper might support an analysis of patent thickets of the various types. '?' denotes that an author didn't make an explicit or implicit claim on existence or hindrance. '-' indicates that hindrance is not applicable because a lack of existence is claimed. Reitzig (2004) has data on five industries: Chemicals, Drugs, Electronics, Machinery, and Transport Manufacturing. Von Graevenitz et al. (2012)'s 30 technology areas are discussed in the body of the text.

				Paper claims:		**Implications:**			
Paper	**Definition**	**Region**	**Sector**	**Exists?**	**Hinders?**	**A**	**B**	**C**	**D**
Cockburn and Macgarvie (2011)	A-T	US	Software	✓	✓	✓			✓
Cockburn and MacGarvie (2009)	A-ST,B-T	US	Software	✓	X	✓	✓	✓	✓
Cockburn et al. (2010)	A-T	DE	Mftr/Srvcs	✓	✓	✓			
Entezarkheir (2010)	A-T	US	Mftr	✓	?	✓			✓
Galasso and Schankerman (2010)	A,B	US	All	✓	X	✓	✓		
Geradin et al. (2007)	A	US	Cellular	X	-	✓	✓	✓	
Hall and Ziedonis (2001)	A,B	US	Semiconductor	✓	?		✓	✓	
Hall et al. (2012)	A,B	UK	All	✓	✓	✓	✓	✓	✓
Hall and Ziedonis (2007)	A	US	Semiconductor	?	?	✓			✓
Harhoff et al. (2012)	B	EU	All	✓	?	✓	✓	✓	✓
Hegde et al. (2009)	C	US	All	?	?		✓	✓	
Huang and Murray (2009)	A,B	US	Genetics	✓	✓	✓			✓
Murray and Stern (2007)	B-T	US	Biotech	✓	✓	✓	✓		✓
Nagaoka and Nishimura (2006)	A-T	JP	Mftr/Constr/ICT	✓	X	✓			✓
Noel and Schankerman (2006)	A-T	US	Software	✓	X	✓			
Reitzig (2004)	A	EU	5 Industries	✓	X	✓	✓		✓
Siebert and Von Graevenitz (2010a)	A-ST,B-ST,C1	US	Semiconductor	✓	X		✓		
Siebert and Von Graevenitz (2010b)	A-T,D	US	Semiconductor	✓	✓	✓	✓		✓
Von Graevenitz et al. (2012)	A,B	EU	30 Tech. Areas	✓	X	✓	✓	✓	
Ziedonis (2004)	A-PST	US	Semiconductor	✓	X	✓			

The jurisdiction of patent applications and the industry area studied might affect the type of patent thicket one would expect to find. Prior to the Leahy-Smith America Invents Act (enacted September 16, 2012), there was no post-grant opposition mechanism in the United States.[72] Whereas in the European Union, post-grant opposition has been possible since 1973. This suggests that there may be fewer issues relating to spurious patents (type C thickets) in the E.U. than in the U.S. Likewise, although we are aware of no supporting evidence, and Graham and Vishnubhakat (2013) and others provide rebutting evidence, there is a persistent believe that software and business method patents are of lower quality, and more likely to be spurious, than patents in other technology areas. Semiconductors, manufacturing, and Information Communications Technology (ICT), on the other hand, are all industries associated with complex products and type A patent thickets. Meanwhile, biotechnology, including genetics, and pharmaceuticals are associated with type B thickets, and particularly type $B1$ thickets that pertain to cumulative invention.

The measures, and to a lesser extent the dependent variables, that studies have used might provide implications for one or more types of thicket. And the definition of a thicket implemented in a particular paper might be (partially or fully) at odds with either the sample used or these implications. The derangement of four elements – sample predisposition, claimed definition, measured definition, and implemented test – makes drawing substantive conclusions essentially impossible.

We conclude this section by reporting the aggregate conclusions reached by the authors of empirical work on patent thickets. In table 7 we provide two columns under the heading of the 'paper claims'. We went through the results and conclusion sections of each of the empirical patent thickets papers and recorded their assertive statements regarding the existence of patent thickets and whether they claimed that patent thickets hindered innovation. We did not take into account whether or not these assertions were based upon evidence, the quality of measures, or any other factor. We denote cases where no assertion was made with a question mark, cases of positive assertion (i.e., thickets do exist or do hinder innovation) with a tick, and cases where no positive assertion was made with a cross. It is clear that almost all authors believe that patent thickets do exist: we record 17 claims for, 1 claim against, and 2 abstentions. Whether or not thickets hinder innovation, however, is much more mixed: we record 6 claims for, 8 claims against, 5 abstentions, and 1 not applicable.

Another way of determining whether or not patent thickets exist and hinder innovation is to ask practitioners. Mann (2004) did exactly that in his paper entitled "The Myth of the Software Patent Thicket". He says: "the idea of a 'thicket' or 'anti-commons' in the software industry is difficult to credit. When raised in my interviews, that thesis universally was rejected." And he is equally firm in his conclusions: "the patent system is not systematically preventing the initiation of product development. Beyond that, it is plain that the system is not obviously dysfunctional."

Taking the claims of papers that have actually done work on patent thickets at face value, the median message is not that patent thickets hinder innovation. Although thickets are claimed to exist, many authors claim that private mechanisms already in place remedy their issues. Moreover, a key point of this research is that these claims should not be taken at face value. The papers that policy makers might want to rely on have implemented tests that are at odds with their thicket definitions; used samples that are at odds with their measures, tests and definitions; and implemented measures that may simultaneously measure many different definitions or none of them. More work is needed on the foundations. Future research should establish measures that have strong construct validity and represent one and only one type of thicket from the taxonomy. Once we understand what has been measured, we may well be able learn a vast amount from the research already done.

[72] Patents could still be declared invalid by the courts or the patent office but there was no mechanism for a third-party to bring a validity dispute to the patent office.

6 Policy

In this section, we discuss what we have learnt about patent thickets through the lens of some suggested policy responses to the patent thicket problem. We focus on policy recommendations made by government bodies such as patent offices, government commissions, and regulatory agencies. They account for just 11 out of 49 articles in our full sample advocating policy changes to address the perceived problem of patent thickets, but they are likely to be the most influential.[73] A high-level summary of these reports are provided in table 8 below.

Table 8: Policy recommendations made in policy reports on patent thickets

Policy Report (By Year)	Types	Pages	Commissioning Organization	No policy recommendation	Improve Patent Quality	Improve Information	Improve Dispute Resolution	Regulate Licensing	Suggested Extent
Muris (2001)	A-S	8	F.T.C.	✓					–
Arundel and Patel (2003)	A-T, B	19	European Commission		✓		✓		Mild
Commission (2003)	B, C1	315	F.T.C.		✓	✓		✓	Strong
Attaran (2004)	A-T, D	7	W.H.O.	✓					–
Schacht (2006)	A,B,D	17	U.S. Congress	✓					–
Van Zimmeren et al. (2006)	A-T	8	W.H.O.					✓	Moderate
Cowin et al. (2007)	B1,D-S	71	EU Parliament		✓	✓		✓	Strong
Harhoff et al. (2007)	A-T	308	European Commission		✓	✓	✓	✓	Moderate
Competition (2008)	B,C1,D	426	European Commission	✓					–
Jacob (2009)	C1, D	11	European Commission		✓		✓		Mild
Commission (2011)	B, C1	309	F.T.C.					✓	Strong
Hargreaves (2011)	B-ST	130	U.K. I.P.O.		✓				Moderate
Regibeau and Rockett (2011)	A-ST,B-T	146	European Commission					✓	Strong
Team (2011)	A,B,C1	71	U.K. I.P.O.		✓	✓			Strong
Hall et al. (2012)	A,B	66	U.K. I.P.O.	✓					–
EPC and Board (2013)	A-T,B	28	E.P.O.		✓	✓	✓	✓	Moderate

Suggested policy responses to patent thickets fall into four basic categories[74]:

- Improve patent quality: suggestions include the establishment of post-grant opposition proceedings in the U.S. (implemented in the Leahy-Smith America Invents Act, September 16$^{\text{th}}$, 2012); peer-review of applications; financial incentives for opposition (Leahy-Smith granted the USPTO authority to set fees for oppositions, the proposed cost for a review of up to 20 claims is around \$30,000); increased renewal fees (Leahy-Smith slightly increased fees overall, while reducing fees for 'micro entities'); increased standards for patentability (particularly regarding obviousness and the required level of inventive step) and by preventing the inclusion of certain types of subject matter as patentable (or in some cases, a removal of the allowance for business method, genetics, software, and certain other controversial patent classes); and legislation to have validity determined on a preponderance of the evidence (rather than presumed). A general plea for increased resources for the USPTO, so that greater effort could be dedicated towards search and examination, was also common.
- Improve information about patents: for example through the publication of all patent applications shortly after filing (implemented in the Technical Amendments Act, November 29$^{\text{th}}$, 2000), the creation of centralized registrars of IP disputes, and improved clarity of patent filings.

[73] There are also 5 policy reports commissioned by government agencies that have not advocated policy changes. Attaran (2004) and Regibeau and Rockett (2011) conclude that it is simply too early to do so. They claim that not enough is known about the consequences of patent thickets in practice. Muris (2001) summarizes current practices, Schacht (2006) notes that changes considered by congress will differentially affect various industries, and Hall et al. (2012) conducts a large scale study in an attempt to establish stylized facts relevant to policy makers.

[74] This categorization holds for all 49 papers advocating policy changes, not just the 11 reviewed in table 8.

- Improve dispute resolution: for example through the establishment of a dedicated European Patent Court, creation of an arbitration process for patent disputes, or the facilitation of patent litigation insurance (so that small firms can participate in patent disputes).
- Regulation of licensing activity: Proposed changes varied from imposing limitations to patent damages and on the use of injunctions, to requirements that royalties be calculated on an 'ex-ante' (to infringement) basis, and onward to compulsory licensing schemes and the creation of national royalty collection and clearing houses.

Suggested policy responses among these 11 papers also varied in their extent of desired reform. In the last column of table 8 we characterize the suggestions as ranging from mild to moderate to strong. Jacob (2009), Hargreaves (2011), and Commission (2011), provide respective examples of the differing extents of desired reform:

> "Changes to the system should be viewed with great care – on the whole it works very very well. One should be very careful to avoid panic-driven or emotion-led changes which could damage an important and beneficial part of industry." – Jacob (2009)

> "The structure of patent renewal fees might be adjusted to encourage patentees to assess more carefully the value of maintaining lower value patents, so reducing the density of patent thickets." – Hargreaves (2011)

> "Concerns about punishing infringement, deterring infringement... should not inflate the reasonable royalty damage award [assessed by the courts] beyond what a willing licensee would have paid for a patent [ex-ante to the infringement]..." – Commission (2011)

We are not aware of any evidence that spurious patents are issued in material numbers or that they have a notable adverse effect on the innovation economy. However, as we said previously, there is considerable conjecture in the literature that spurious patents have been issued and that their number is growing. There is also, to the best of our knowledge, no recent data on the percentage of patents that are commercialized. A study from the 1950's, Rossman and Sanders (1957), estimated that around half of patents are commercialized. Lanjouw and Schankerman (2001) estimates that around 0.7% of U.S. corporate patentees are involved in patent litigation, and Moore (2000) shows that in a sample of 1,209 cases, 1,151 considered the issue of patent validity, of which 33% returned a judgement of invalidity. Theoretically, a single, well-placed, spurious patent might be sufficient to cause a type C 'thicket', as it could hinder innovation.

Policy to improve patent quality should decrease the likelihood of invalid patents. However, it might also remove patents that are currently 'marginal' – legitimate patents with a small inventive step and so a correspondingly small benefit to their applicants/holders (from their right to exclude on the claimed domain of application) that is only slightly greater than the cost of prosecuting a patent application. These patents may have social benefits that greatly exceed their private benefits. However, the inventions behind these patents, and the social benefits, may be placed in the public domain if patent protection is impossible, and so might not be lost. Furthermore, inventive effort may be rededicated to other projects. As such, the economic costs to refusing patents with small inventive steps are probably low, whereas, if type C patent thicket problems are material and policy successful reduces their incidence, the economic benefits could be substantial.

Lemley (2000)'s calculus of rational ignorance at the patent office would be undermined by a feedback loop induced by spurious patent applications. Moreover, the loss function to over or under funding the patent office is likely highly asymmetric. The first-best level of funding would match social benefits to social costs. Over funding the patent office would be an economically inefficient allocation of resources relative to this first best; a likely consequence of over-funding would be rent extraction by PTO staff.

However, patents are an input into many other activities, and under-funding is likely to lead to a host of problems relating to patent validity. Many downstream activities will generate increased social costs (such as patent litigation, forgone invention due to type C thickets, and so forth), and other downstream activities would experience reduced social benefits (such as the positive externalities of innovative activity). When making a decision under uncertainty with an asymmetric loss function, it is optimal to err towards the flat side of the loss function. Thus we argue that over-funding of the PTO should be preferred to under-funding.

Transaction costs and search costs (i.e., thicket variants T and S), as well as the coordination problem inherent in type A patent thickets, would benefit from increased information about patent-related activities. In the Internet Age, the marginal cost of distributing information is close to zero. The cost of assembling information, however, can be substantial. Not only could wasteful duplication of effort of information assembly be avoided, but in many cases the cost of assembly of information could be borne by the entity with the lowest cost – either the patent-holder who could report the information to patent office, or the patent office itself. As such, policy to increase information about patenting activities will generally have small costs and potentially large benefits. One important caveat concerns the period for which a patent application should be kept secret. On the one hand, secrecy enhances the possibility of guile (particularly through the creation of submarine patents) and on the other it prevents free-riding and gives the applicant an opportunity to file for protection on alternative implementations. Absent a period of secrecy, inventors may delay application for patents, which in turn may delay commercialization.

Improved dispute resolution processes could help mitigate the effects of type B thickets, where disputes over the allocation of rents are the driving consideration, as well as type A thickets, where coordination failures may result in disputes. However, policy to change a dispute resolution process is a likely target for lobbying activities – changing the rules for disputes between implementers and inventors (type A thickets) or pre-existing patent-holders and new patent-holders (type B thickets) could favour one party over the other – as it could enable transfers which might adversely affect incentives to innovate. That said, party-neutral changes to dispute resolution processes could be welfare improving. For example, making dispute resolution more efficient by enabling more accurate decisions, or making decisions more quickly and with fewer resources, would increase welfare. The creation of a European Patent Court could centralize expertise, create a level playing-field for EPO patent disputes, and increase the consistency of patent dispute decisions. Patent litigation insurance might afford smaller firms the same patent protection as their larger, more resource rich counterparts, but might also facilitate opportunistic rent-seeking. And arbitration processes, particularly those that are voluntarily agreed to both parties, could speed dispute resolution.[75]

Regulation of licensing activity is by far the most problematic of the proposed regulatory responses to patent thicket problems. Much like calls for improved dispute regulation, regulating licensing can favor the licensee over the licensor, or vice versa. Proposals like compulsory licensing, limitations on damages, imposition of 'ex-ante' licensing terms, the forbidding of injunctions, and caps on licensing rates all appear to favor the licensee. In the context of $B1$ patent thickets, Scotchmer (1991) argues that there are grave risks of undercompensating upstream inventors. And the removal of disciplining mechanisms, such as refusals to license or the charging of punitive licensing rates, might cause rather than cure type A thicket problems. Furthermore, many of these proposals might create patent hold-up problems. Using data on 1,209 patent suits tried from 1983 to 1999, Moore (2000) report that infringement was determined to have taken place in 66% of decisions, and that when willfulness of infringement was considered, infringement was found willful in 64% of decisions.[76] When a licensee is guaranteed a pre-infringement license price

[75] Imposing arbitration, especially arbitration with a limited set of outcomes or that forbids the use of legal instruments like injunctions, is much more likely to bias outcomes to favour one party over the over, and as such may not be welfare improving.

[76] It is interesting to note that infringement and willfulness of infringement were much more likely to be determined by judges than by juries. Judges found infringement in 71% of cases, and found it willful in another 71% of cases, as opposed to 59% and

after having been caught for infringement, then willful infringement, and so hold-up of the inventor for the rents to the invention, is a dominant strategy whenever the chance of detection and prosecution of infringement is sufficiently small to overcome any additional cost from having willfully infringed.

7 Summary and Conclusion

Just a little over ten years ago, the term 'patent thicket' was coined to describe one particular economic relationship between patent rights that might hinder innovation – the diversely-held complementary inputs problem that we name a type A thicket. Today the literature on patent thickets extends to hundreds of papers, and many different types of relationship between patent rights are described as causing a patent thicket. At least 164 papers, reviewed in a meta-analysis in this research, have provided a definition of the term. In this paper, we identified four main types of patent thicket (A, B, C, and D), with ten sub-types, and five variants. Each type and sub-type reflects a different relationship between patent rights and so embodies a different economic mechanism with different welfare consequences. Each different variant modifies these economic mechanisms and adds its own welfare issues. The four main types of patent thicket, together with brief descriptions of their economic foundations and welfare consequences, are as follows:

- Type A – Diversely-held complementary input patent thickets: When a new product requires many diversely-held complementary inputs that are protected by patents, an N-fold marginalization problem can arise from a lack of coordination or suitable disciplining mechanisms. As a result the price of the product creates a deadweight-loss and invention is under-compensated. The literature has confused this with hold-up. We argue that search problems and behavioural limitations may lead to commercialization without securing the necessary property rights, but that this a consequence of a decision made under uncertainty and that, absent guile by a patent holder, this does not constitute hold-up.
- Type B – Legitimate overlapping patent patent-thickets: When patents all have a novel and non-obvious inventive step but a new patent has a claimed domain of application that overlaps with a pre-existing patent or patents, a problem regarding the correct allocation or perceived allocation of rents can arise. In the context of cumulative invention, where the new patent builds upon pre-existing patent rights, the upstream invention is a complementary input into the downstream invention and may be under (or over) compensated. The economic foundation of this argument rests on inperfection in the assignment of property rights. One example mechanism is hold-up. For example, the fear of appropriation of rents through opportunism in Reach-Through Licensing Agreements (RTLA) may cause inefficiently low levels of investment in follow-on invention, and so under-compensation for upstream rights holders. On the other hand, in the context of adjacent invention, the new patent is a legitimate economic (not technical) substitute to one of more pre-existing patents, but pre-existing patent-holders may believe that the patent office has erred in the issue of the new patent. When patent enforcement is prohibitively costly or patents are probabilistic, this can lead to an over compensation of pre-existing patent rights.
- Type C – Spurious patent patent-thickets: When the patent office makes mistakes and issues patents that do not embody a new and non-obvious inventive step, firms may seek spurious patents in order to engage in transfer seeking. This can reduce the rents to genuine inventors, potentially reducing genuine invention, and wastes resources. When guile is used in the patent application, we argue that this is a form of 'patent hold-up'. Patent hold-up differs from Williamson hold-up in that the

53% respectively for juries.

'contract' is created by the grant of exclusionary rights by the patent office. Like Williamson hold-up, patent hold-up leads to inefficiently low levels of investment – this time in genuine innovative activity.

- Type D – Effectively saturated invention-space patent thickets: For a period of time, in a technological area, it is possible that "everything that can be invented, has been invented". When a small number of variants of an inventive step can achieve the same economic functionality, firms may engage in a 'ring-fencing' strategy to prevent free-riding and low-cost imitation. Type D patent thickets are welfare neutral as no new (patentable) invention is possible, through they may facilitate the creation of new combinations of patents and so further Schumpeterian innovation.

Thicket variants, such as T, S, and P, add notions like transaction costs, search costs and probabilistic patents (i.e., the enforcement of patent rights by the courts is imperfect) to patent thicket types. And refinement of types into sub-types, for example the decomposition of type B into type $B1$ for cumulative overlaps and type $B2$ for adjacent overlaps, creates a hierarchical taxonomy of patent thickets. This hierarchy, and the articulation of its economic foundations, is the core contribution of this paper.

We then documented a general confusion in the literature over the meaning and mechanisms behind patent thickets. The average paper in our full sample implemented a definition of a patent thicket that was consistent with 1.3 main thicket types. This confusion appears to be increasing over time, and in 2012 the average paper implemented a definition consistent with around 1.8 main thicket types. We found a frequent mismatch between the industry (i.e., semiconductors, biotech, software, etc.) and the type of thicket being discussed in papers. And we observed that papers that suggested changes to firm strategy, intellectual property right regime reform, or private mechanisms in response to thickets frequently fail to understand the economic fundamentals of the type of patent thicket that they purport to discuss. Misplaced references to hold-up problems were particularly common. We are optimistic that the taxonomy of patent thickets in this paper will help alleviate these problems in the future.

We also discussed the measures and tests used in the near-population of 20 papers that have studied the effects of patent thickets. We drew attention to fundamental measurement issues. Specifically, we demonstrated that almost all measures of patent thickets either were unlikely to measure any of the relationships between patents that give rise to a particular type of thicket, or that they were likely to capture many different types of these relationships. We also argued that the majority of tests for the effects of patent thickets suffered from poor research design and were unable to differentiate between the effects of different types of thicket, or from positive or negative welfare consequences of thickets. As such, the empirical foundation for claims of patent thicket problems appears lacking. We hope that the renewed attention to the theoretical foundations provided in this work will allow the creation of new measures and the establishment of the semantics of existing measures, so that prior work can be given a meaningful interpretation in the future.

Finally, we reviewed the 11 policy articles published by government and non-government agencies on patent thickets, and examined their policy advice. The most common advice pertained to a perceived problem with patent quality. However, in our near-population of papers that provided a definition of patent thicket, only around 16% provided a definition consistent with a type C (spurious patent) patent thicket, for which such policy advice might be applicable. Around 60% of the papers in our sample implemented a definition consistent with type A (diversely-held complementary input) patent thickets. Type A thickets were the original thicket type discussed by Shapiro (2001), who noted that "antitrust law can potentially play such a counterproductive role, especially since antitrust jurisprudence starts with a hostility toward cooperation among horizontal rivals." Unfortunately, since the publication of Shapiro (2001), the calls to regulate licensing in a fashion that is likely to aggravate, rather than ameliorate, type A patent thicket problems have persisted if not increased. Some of the strongest calls have come from the F.T.C., despite research to suggest that their efforts are counter-productive.[77]

We now conclude by trying to answer the questions: Do patent thickets exist? Do patent thickets hinder innovation? And what sensible policy responses should be made to the problem posed by patent thickets?

There are products that require many diversely-held patented complementary inputs. A wide range of consumer electronic devices, from MP3 players to digital cameras to desktop computers, are all likely to meet this criteria. Smartphones, however, are the prototypical example. At the launch of the iPhone®, Steve Jobs famously claimed "and boy have we patented it!" Of course, they hadn't: the iPhone probably relies on thousands of patents held by hundreds other firms.[78] Thus smartphones, like the iPhone® do require many complementary inputs that are covered by patents and these patents are diversely-held, so type A patent thickets almost surely exist.

However, for a type A patent thicket to hinder innovation, there must be an N-fold marginalization problem that creates a deadweight-loss in product markets and, in anticipation of the reduced rents to invention, causes a drop in inventive activity in the complementary inputs themselves. We can observe some measures of deadweight-losses in product markets but we simply can not distinguish their causes.[79] We can also observe a rise in patenting activity in the components used in these products, but there is no way to know whether the size of inventive steps in these patents, or their rents, are declining, holding steady or increasing. And although the "smartphone wars" – a barrage of injunctions, infringement suits, validity contests, and other patent disputes – conjures the image of a sector in crisis, the outcomes have included broad-cross licensing agreements (often with one-off transfer payments to balance the exchanges) and other commitments to cooperation that could mitigate the N-fold commercialization problem.

Like all other researchers, we are therefore constrained to our best guess as to whether type A thickets hinder innovation. We suspect that private mechanisms are extremely effective in mitigating the adverse effect of an N-fold marginalization problem. Profits and social welfare are both aligned; the former does not come at the expense of the latter. As such, we expect firms to follow the advice given by Arrow (1973) and create informal institutions to solve their coordination problems. Cross-licensing agreements, participation in the creation of technology standards, patent pools, FRAND commitments, and self-regulation through reputations and disciplining mechanisms, all serve as examples. Furthermore, policy to force cooperation runs the risk of favoring implementers over inventors, or vice versa, and locking in a market failure, or worse yet introducing new market failures. Policy to prevent cooperation would imperil the mechanisms currently in use. As it stands, new and highly complex products are being invented, assembled, and brought to market at an astounding rate. It is hard to characterize this extraordinary innovation ecosystem as broken.

There are also many examples of patents with overlapping patent rights but different (and legitimate) inventive steps, so we feel comfortable in proclaiming the existence of type $B2$ patent thickets. When these patents are adjacent (type $B1$), disputes over property rights do arise and are settled by the courts.

[77] Shapiro (2001) concluded by saying "the Federal Trade Commission has exhibited less restraint, and arguably is making it more difficult for firms to engage in [mechanisms to prevent type A thicket problems]". See also Epstein et al. (2012) for an excellent critique of Commission (2011).

[78] The exact number of patents that pertain to the iPhone® is unknown but Apple's percentage is likely small. The number of other patent-holders – the N in the N-fold marginalization problem – is likely several hundred. According to Gaze and Roderick (2012), Apple has filed 1,298 iPhone related patents (416 of which are on core smartphone technology, the rest of which are on cameras, user interfaces, batteries, antennas, and other components used in smartphones) since 2000. Apple was also involved in 479 patent cases pertaining to the iPhone from 2008 to 2012; in 263 cases Apple was the defendant accused of infringing smartphone patents belong to a wide array of other patent-holders. Estimates of the total number of patents needed to create a smartphone vary wildly.

[79] Unsealed statements from *Apple Inc. v. Samsung Electronics Co. Ltd. et al.*, Case No. C 11-1846 LHK, revealed that Apple's gross margin on the iPhone is around 55%. Apple makes (presumably profit maximizing) pricing decisions in the context of the consumer market's taste for new and differentiated personal electronics goods and does not price at marginal cost. But we do not know the counter-factual of what Apple's marginal cost (and so pricing and deadweight-loss) would have been if it held all of the patents on an iPhone.

It would be better if these disputes did not take place, but we live in imperfect world. It is far from clear that policy could solve this problem but better information on patenting disputes and improved dispute resolution might reduce the economic losses incurred. We are not aware of research that provides a decomposition of patent suits suitable for an analysis of whether or not this is a growing or large-scale phenomena. Future research should endeavour to quantify this issue. For the moment, we suggest that type $B1$ thickets do not pose a material hindrance to innovation.

With cumulative overlapping patent rights (type $B1$) we are more concerned with whether the allocation of rents between upstream and downstream parties is optimal. One problem is that patents are somewhat "one size fits all". Although the scale of exclusionary rights should match the scale of the inventive step, some inventive steps do not have (or have few profitable) commercial applications in and of themselves but instead realize their commercial gains through follow-on inventions. Patents should not be allowed on laws of nature or products of nature, but this does not preclude the patenting of research tools or materials (new compositions of matter) that can be used in subsequent inventions, or the possibility that a downstream patent refines or repurposes an upstream patent to make it more profitable.[80] The question at hand is whether these rents are currently so systematically sub-optimal that they are preventing the creation of either upstream patents or downstream patents. As downstream patents can not occur without upstream patents and at least some welfare can potentially be realized through upstream patents alone, the under-compensation of upstream patents is the greater threat.

Since the implementation of the Bayh-Dole Act (1980), universities in the U.S. have engaged in extensive patenting.[81] The nature of university research favours the creation of upstream patents that are closer to basic science and more likely to cover research tools or materials. Universities are also less profit motivated, and are likely willing to accept smaller shares of rents. Fears of an under provision of upstream patenting by industry could therefore be off-set by increasing research funding to universities.

The question of existence of type C patent thickets is not whether spurious patents exist – some surely do – but whether legitimate innovators opt or are forced to form relationships with them. We do observe patent suits when claims of infringement are rejected on the grounds of invalidity. But this is evidence of rejection of a relationship; the observation of an invalid patent and the observation of a lack of (an ongoing) relationship is one and the same. However, if spurious patent patent-thickets do exist, then they will surely provide a hindrance to innovation. Spurious patents do not generate new economic welfare; they enable transfer seeking. We have argued that the loss function from over or under funding the patent office is likely to be highly asymmetric. Patents are inputs into commercialization and further invention, and under funding of the patent office has the potential for dire follow-on consequences. Over funding of the patent office may result in some bureaucratic wastage but this is unlikely to have follow-on consequences. Thus, given the uncertainty concerning first-best funding levels, we favor increased funding to the patent office and erring on the side of caution.

The existence of type D thickets is pure conjecture. We have no way of measuring the potential invention space at any given time. To do so would be to know the future. We also struggle to measure whether, when, and to what extent paradigm shifts have happened and displaced previous generations of technology. Thus we can not say whether invention spaces have ever become effectively saturated, or whether (let alone at what intervals) this saturation has been wiped away by radical innovation. However, type D patent thickets are welfare neutral, so policy to address them is largely irrelevant.

Transaction costs and search costs need to be quantified and their roles in patent thickets need to be better understood. Both certainly exist, but it is far from clear whether they are material to licensing decisions or if, perhaps at some threshold, they preclude future innovation. However, even without a

[80] That products of nature, particular DNA sequences, can not be patented was decided in the recent Supreme Court case *Association For Molecular Pathology v. Myriad Genetics*, 569 U.S. (2013).

[81] See Mowery (2004).

detailed understanding of transaction and search costs, policy to increase information about patenting activity should be considered. The marginal cost of distributing information is close to zero, and although the collection of information could add a regulatory burden, this is likely to be small compared with the benefits from better decisions – both commercial and with regards to policy towards to patent-related activities – that could be enabled.

Finally, we want to put the patent thicket 'problem' in context. The start of the 21st century is witness to a technologically advanced and highly sophisticated innovation environment. We are able to assemble products made up of thousands, if not millions, of components. Some of these components are covered by patents and are owned by different firms that must coordinate their activities on many different levels – the components must fit together, work together, and be priced together. We also have products that are the result of towering spirals of cumulative invention. Patents on research tools enable the discovery of new materials; patents on new materials enable the construction of new structures; patents on new structures enable the creation of new apparatus; patents on new apparatus enable both the creation of new research tools and new products, and so the cycle continues. Of course, there is conflict: the patent office had a 35 month backlog in 2010; in 2012, 5,836 court cases relating to intellectual property were recorded as filed with the district and federal courts in PACER; almost 200 products have been removed from shelves due to injunctions filed at the ITC over a recent five year period; the FTC has initiated actions over alleged anti-trust violations involving patent licensing agreements, patent pools, and standard-essential patents; and the Supreme Court has heard over 50 patent cases in the last 100 years to clarify patent law.[82] But this conflict needs to viewed in context: In 2012, 542,815 utility patent applications were made to the USPTO; 1,203 high-tech firms secured more than $4.2b in their first rounds of venture capital investment in 2012; U.S. export of intangible assets stood at $83b in 2007, and was growing at a rate of 14% a year; and around 38% of US GDP came from 'Knowledge and technology intensive industries' in 2007, growing at a rate of around 0.3% per year.[83] There is no data available to measure the number of new products and services that rely on patents that were created in the last year, but common sense suggests that the number is huge.

So, to our minds, there is no patent thicket 'crisis'. Instead the problems of patent thickets are generally symptoms of the enormous complexity of the innovation ecosystem. We might, rightly, worry about trying to reduce the frictions of patent thickets in the hope of an ever-more productive innovation ecosystem. But claims of a pressing problem are belied by the low levels of conflict and high levels of achievement. Moreover, any response to patent thickets must be cautious. Patent thickets arise because of certain specific relationships between innovations and pre-existing inventions. They are also often mitigated by other relationships between innovators and the patent-holders who own these pre-existing inventions. Interfering in these relationships, for example by regulating patent licensing, risks poisoning the cure while trying to fight the disease. Policy to reduce the incidence of spurious patents is the exception. In this case, any relationship between innovators and pre-existing spurious patent-holders can only be bad. In the process of conducting the literature review at the heart of this paper, we have seen hundreds of pleas for more funding for the USPTO, and no suggestions that their funding should be cut. We happily board this bandwagon. The recent Leahy-Smith America Invents Act (2011) is only a small step (generally) in the right direction: more should be done.

[82]Chien and Lemley (2012) claim 191 injunctions requests were heard by the I.T.C. from 2006 from 2011 to 2011, with a grant rate of around 90%. FTC actions are listed on www.ftc.gov. For a list of Supreme court cases see the "List of United States Patent Law Cases on wikipedia.org.

[83]Count of patent applications from the USPTO. Venture capital data from PWC Moneytree. Knowledge and technology intensive industry share of GDP and U.S. trade in intangible asset data is from the NSF and is the most recent data available.

8 References

ABA (2004). Handbook on the antitrust aspects of standards setting. American Bar Association (ABA) Section of Antitrust Law.

Aggarwal, Vikas A. and Hsu, David H. (2009). Modes of cooperative r&d commercialization by start-ups. Strategic Management Journal, vol. 30 (8), pp. pp. 835–864.

Aghion, Philippe, Bloom, Nick, Blundell, Richard, Griffith, Rachel, and Howitt, Peter (2005). Competition and innovation: An inverted-u relationship. The Quarterly Journal of Economics, vol. 120 (2), pp. 701–728.

AIPLA (2011). Report of the economic survey.

Albert, M.B., Avery, D., Narin, F., and McAllister, P. (1991). Direct validation of citation counts as indicators of industrially important patents. Research Policy, vol. 20 (3), pp. 251–259.

Allison, J.R. and Tiller, E.H. (2003). The business method patent myth. Berkeley Tech. LJ, vol. 18, p. 987.

Andrews, L.B. (2002). Genes and patent policy: Rethinking intellectual property rights. Nature Reviews Genetics, vol. 3 (10), pp. 803–808.

Aoki, R. and Schiff, A. (2008). Promoting access to intellectual property: Patent pools, copyright collectives, and clearinghouses. R&D Management, vol. 38 (2), pp. 189–204.

Arrow, Kenneth J (1973). Social responsibility and economic efficiency. Public Policy, vol. 21 (3), pp. 300–317.

Arundel, A. and Patel, P. (2003). Strategic patenting. In Background Report for the Trend Chart Policy Benchmarking Workshop" New Trends in IPR Policy.

Attaran, A. (2004). Patents do not strangle innovation, but their quality must be improved. Bulletin of the World Health Organization, vol. 82 (10), p. 788.

Ayres, Ian and Parchomovsky, Gideon (2007). Tradable patent rights. Stanford Law Review, vol. 60 (3), pp. 863–894.

Baluch, A.S., Radomsky, L., and Maebius, S.B. (2005). In re kumar: The first nanotech patent case in the federal circuit. Nanotech. L. & Bus., vol. 2, p. 344.

Baron, J. and Delcamp, H. (2010). Strategic inputs into patent pools. Cerna working paper.

Baron, J. and Pohlmann, T. (2011). Patent pools and patent inflation. In Conference Proceedings: 4th ZEW Conference on the Economics of Innovation and Patenting.

Barpujari, I. (2010). The patent regime and nanotechnology: Issues and challenges. Journal of Intellectual Property Rights, vol. 15 (3), pp. 206–213.

Barton, John H. (2002). Antitrust treatment of oligopolies with mutually blocking patent portfolios. Antitrust Law Journal, vol. 69 (3), pp. pp. 851–882.

Baumol, William J. (2004). Entrepreneurial enterprises, large established firms and other components of the free-market growth machine. Small Business Economics, vol. 23 (1), pp. pp. 9–21.

Bawa, R. (2007). Nanotechnology patent proliferation and the crisis at the us patent office. Alb. LJ Sci. & Tech., vol. 17, p. 699.

Bawa, R., Bawa, SR, and Maebius, S.B. (2005). The nanotechnology patent gold rush. Journal of Intellectual Property Rights, vol. 10 (5), pp. 426–433.

Bawa, Raj (2005). Will the nanomedicine patent land grab thwart commercialization? Nanomedicine: nanotechnology, biology, and medicine, vol. 1, pp. 346–350.

Beard, T.R. and Kaserman, D.L. (2002). Patent thickets, cross-licensing, and antitrust. Antitrust Bull., vol. 47, p. 345.

Bergman, K. and Graff, G.D. (2007). The global stem cell patent landscape: Implications for efficient technology transfer and commercial development. Nature biotechnology, vol. 25 (4), pp. 419–424.

Bessen, J. (2003). Patent thickets: Strategic patenting of complex technologies. Available at SSRN 327760.

Braun, V. and Herstatt, C. (2007). Barriers to user innovation: Moving towards a paradigm of 'licence to innovate'? International Journal of Technology, Policy and Management, vol. 7 (3), pp. 292–303.

Burk, D.L. and Lemley, M.A. (2003). Policy levers in patent law. Va. L. Rev., vol. 89, pp. 1575–1696.

Calderini, M. and Giannaccari, A. (2006). Standardisation in the ict sector: The (complex) interface between antitrust and intellectual property. Econ. Innov. New Techn., vol. 15 (6), pp. 543–567.

Callaway, D. (2008). Patent incentives in the semiconductor industry. Hastings Bus. LJ, vol. 4, p. 135.

Carrier, M.A. (2002). Why antitrust should defer to the intellectual property rules of standard-setting organizations: A commentary on teece & (and) sherry. Minn. L. Rev., vol. 87, pp. 2017–2034.

Carrier, M.A. (2003). Resolving the patent-antitrust paradox through tripartite innovation. Vand. L. Rev., vol. 56, p. 1047.

Carrier, Michael A. (2004). Cabining intellectual property through a property paradigm. Duke Law Journal, vol. 54 (1), pp. pp. 1–145.

Chien, Colleen V and Lemley, Mark A (2012). Patent holdup, the itc, and the public interest. Cornell L. Rev., vol. 98, pp. 1–209.

Choi, J.P. (2005). Live and let live: A tale of weak patents. Journal of the European Economic Association, vol. 3 (2-3), pp. 724–733.

Clarkson, G. (2004). Objective identification of patent thickets: A network analytic approach. Harvard Business School Doctoral Thesis.

Clarkson, G. (2005). Patent informatics for patent thicket detection: A network analytic approach for measuring the density of patent space. Academy of Management, Honolulu.

Clarkson, G. and DeKorte, D. (2006). The problem of patent thickets in convergent technologies. Annals of the New York Academy of Sciences, vol. 1093 (1), pp. 180–200.

Cockburn, I.M. and MacGarvie, M.J. (2009). Patents, thickets and the financing of early-stage firms: Evidence from the software industry. Journal of Economics & Management Strategy, vol. 18 (3), pp. 729–773.

Cockburn, I.M. and Macgarvie, M.J. (2011). Entry and patenting in the software industry. Management science, vol. 57 (5), pp. 915–933.

Cockburn, I.M., MacGarvie, M.J., and Müller, E. (2010). Patent thickets, licensing and innovative performance. Industrial and Corporate Change, vol. 19 (3), pp. 899–925.

Cohen, Wesley M, Nelson, Richard R, and Walsh, John P (2000). Protecting their intellectual assets: Appropriability conditions and why us manufacturing firms patent (or not). Tech. rep., National Bureau of Economic Research.

Cohen, W.M. and Walsh, J.P. (2008). Real impediments to academic biomedical research. In Innovation Policy and the Economy, Volume 8. University of Chicago Press, pp. 1–30.

Commission, Federal Trade (2003). To promote innovation: The proper balance of competition and patent law and policy. Washington, DC.

Commission, Federal Trade (2011). The evolving ip marketplace: Aligning patent notice and remedies with competition.. March, available at http://www.ftc.gov/os/2011/03/110307patentreport.pdf.

Competition, DG (2008). Pharmaceutical sector inquiry-preliminary report.

Cowin, R., Van der Eijck, W., Lissoni, F., Lotz, P., Van Overwalle, G., and Schovsbo, J. (2007). Policy options for the improvement of the european patent system. Scientific Technology Options Assessment (STOA) of the European Parliament.

Dagg, Joachim L (2011). Exploring mouse trap history. Evolution: Education and Outreach, vol. 4 (3), pp. 397–414.

Devlin, Alan (2009). Indeterminism and the property-patent equation. Yale Law & Policy Review, vol. 28 (1), pp. pp. 61–106.

Dhar, T. and Foltz, J. (2007). The impact of intellectual property rights in the plant and seed industry. Agricultural Biotechnolgy and Intellectual Property Protection: Seeds of Change, p. 161.

D'Silva, J. (2009). Pools, thickets and open source nanotechnology. European Intellectual Property Review, vol. 31 (6), pp. 300–306.

Egan, Edward J (????). The economics of patent citations: Measures, meanings and implications for research. Ph.D. Thesis, U.C. Berkeley. UMI Dissertations Publishing. DOI 3527121.

Egan, Edward J. (2013). Technology and commercialization strategy: Cooperate or compete. Working Paper.

Eisenmann, Thomas R. (2008). Managing proprietary and shared platforms. California Management Review, vol. 50 (4), pp. pp. 31–53.

Eisenstein, M. (2010). Up for grabs. Nature Biotechnology, vol. 28 (6), pp. 544–546.

Entezarkheir, M. (2010). Patent thickets and market value: An empirical analysis.

EPC, Economic and Board, Scientific Advisory (2013). Workshop on patent thickets. European Patent Office Recommendations, Available at http://www.epo.org/news-issues/news/2013/20130313.html.

EPO (2009). Guidelines for Examination in the European Patent Office. European Patent Office, Directorate Patent Law 5.2.1.

Epstein, Richard A, Kieff, F Scott, and Spulber, Daniel F (2012). The ftc, ip, and ssos: Government hold-up replacing private coordination. Journal of Competition Law and Economics, vol. 8 (1), pp. 1–46.

Epstein, Richard A and Kuhlik, Bruce N (2004). Is there a biomedical anticommons. Regulation, vol. 27, p. 54.

Evans, D.S. and Layne-Farrar, A. (2004). Software patents and open source: the battle over intellectual property rights. Va. JL & Tech., vol. 9, pp. 10–13.

Farrell, Joseph (2009). Intellectual property as a bargaining environment. Innovation Policy and the Economy, vol. 9 (1), pp. pp. 39–53.

Farrell, Joseph, Hayes, John, Shapiro, Carl, and Sullivan, Theresa (2007). Standard setting, patents, and hold-up. Antitrust Law Journal, vol. 74 (3), pp. 603–670.

Feldman, R. (2004). The open source biotechnology movement: Is it patent misuse? Minnesota Journal of Law, Science & Technology, vol. 6.

Feldman, R. and Nelson, K. (2008). Open source, open access, open transfer: Market approaches to research bottlenecks.

Galasso, A. (2007). Broad cross-license agreements and persuasive patent litigation: Theory and evidence from the semiconductor industry. LSE STICERD Research Paper No. EI45.

Galasso, A. and Schankerman, M. (2010). Patent thickets, courts, and the market for innovation. The RAND journal of economics, vol. 41 (3), pp. 472–503.

Gallini, N. (2011). Private agreements for coordinating patent rights: The case of patent pools. Economia e Politica Industriale.

Ganslandt, M. (2009). Intellectual property rights and competition policy.

Gaulé, P. (2006). Towards patent pools in biotechnology? Innovation Strategy Today, vol. 2 (2), pp. 123–143.

Gay, C. and Le Bas, C. (2005). Uses without too many abuses of patent citations or the simple economics of patent citations as a measure of value and flows of knowledge. Economics of Innovation and New Technology, vol. 14 (5), pp. 333–338.

Gaze, Laura and Roderick, John (2012). Inside the iphone patent portfolio. Thomson Reuters, IP Market Report.

George, G.D. (2006). What is hiding in the bushes-ebay's effect on holdout behavior in patent thickets. Mich. Telecomm. & Tech. L. Rev., vol. 13, p. 557.

Geradin, D., Layne-Farrar, A., and Padilla Blanco, A. (2007). Royalty stacking in high tech industries: separating myth from reality.

Geradin, D., Layne-Farrar, A., and Padilla Blanco, A. (2008). The complements problem within standard setting: assessing the evidence on royalty stacking. Boston University Journal of Science and Technology Law, Vol. 14, No. 2, 2008.

Gilbert, R.J. (2010). Ties that bind: Policies to promote (good) patent pools. Antitrust Law Journal.

Glover, Gregory J (2002). Patent thickets and innovation markets reviewed. National Law Journal, vol. 24 (56), p. C10.

Goldberg, Victor P (1976). Regulation and administered contracts. the Bell journal of economics, pp. 426–448.

Goodman, David J. and Myers, Robert A. (2005). 3g cellular standards and patents. Proceedings of IEEE WirelesCom, June 13 2005.

Goozner, M. (2006). Innovation in biomedicine: Can stem cell research lead the way to affordability? PLoS medicine, vol. 3 (5), p. e126.

Graham, Stuart and Vishnubhakat, Saurabh (2013). Of smart phone wars and software patents. The Journal of Economic Perspectives, vol. 27 (1), pp. 67–85.

Griliches, Zvi (1998). R & D and productivity. University of Chicago Press.

Hall, B.H. (2005). A note on the bias in herfindahl-type measures based on count data. Revue D' Economie Industrielle, Paris Editions, Techniques Et Economiques, vol. 110, p. 149.

Hall, B.H. (2007). Patents and patent policy. Oxford Review of Economic Policy, vol. 23 (4), pp. 568–587.

Hall, B.H., Helmers, C., von Graevenitz, G., and Rosazza-Bondibene, C. (2012). A study of patent thickets. Draft Report to the UK IPO, pp. 1–66.

Hall, B.H. and Ziedonis, R.H. (2001). The patent paradox revisited: an empirical study of patenting in the us semiconductor industry, 1979-1995. RAND Journal of Economics, pp. 101–128.

Hall, Bronwyn H., Jaffe, Adam, and Trajtenberg, Manuel (2005). Market value and patent citations. The RAND Journal of Economics, vol. 36 (1), pp. pp. 16–38.

Hall, Bronwyn H and Ziedonis, Rosemarie (2007). An empirical analysis of patent litigation in the semiconductor industry. University of California at Berkeley working paper.

Hargreaves, I. (2011). Digital opportunity: a review of intellectual property and growth: an independent report.

Harhoff, D., Von Graevenitz, G., and Wagner, S. (2008). Incidence and growth of patent thickets-the impact of technological opportunities and complexity. CEPR Discussion Papers, vol. 6900.

Harhoff, D., Von Graevenitz, G., and Wagner, S. (2012). Conflict resolution, public goods and patent thickets. Public Goods and Patent Thickets (April 15, 2012).

Harhoff, Dietmar, Hall, Bronwyn H, von Graevenitz, Georg, Hoisl, Karin, Wagner, Stefan, Gambardella, Alfonso, and Giuri, Paola (2007). The strategic use of patents and its implications for enterprise and competition policies. Report commissioned by European Commission.

Hegde, D., Mowery, D.C., and Graham, S.J.H. (2009). Pioneering inventors or thicket builders: Which us firms use continuations in patenting? Management Science, vol. 55 (7), pp. 1214–1226.

Heller, M.A. and Eisenberg, R.S. (1998). Can patents deter innovation? the anticommons in biomedical research. Science, vol. 280 (5364), pp. 698–701.

Heller, Michael A (1997). Tragedy of the anticommons: Property in the transition from marx to markets, the. HARV. l. Rev., vol. 111, p. 621.

Hemphill, T.A. (2003). Preemptive patenting, human genomics, and the us biotechnology sector: Balancing intellectual property rights with societal welfare. Technology in Society, vol. 25 (3), pp. 337–349.

Holman, C. (2006). Clearing a path through the patent thicket. Cell, vol. 125 (4), pp. 629–633.

Holman, C.M. (2005). Biotechnology's prescription for patent reform. J. Marshall Rev. Intell. Prop. L., vol. 5, p. i.

Holman, C.M. (2008). Trends in human gene patent litigation. Science, vol. 322 (5899), pp. 198–199.

Horn, L. (2003). Alternative approaches to ip management: One-stop technology platform licensing. Journal of commercial biotechnology, vol. 9 (2), pp. 119–127.

Huang, K.G. and Murray, F.E. (2009). Does patent strategy shape the long-run supply of public knowledge? evidence from human genetics. Academy of Management Journal, vol. 52 (6), pp. 1193–1221.

Hussinger, K. (2006). Is silence golden? patents versus secrecy at the firm level. Economics of Innovation and New Technology, vol. 15 (8), pp. 735–752.

Iyama, S. (2005). The uspto's proposal of a biological research tool patent pool doesn't hold water. Stanford Law Review, pp. 1223–1241.

Iyer, Rajkamal and Schoar, Antoinette (2010). Incomplete contracts and renegotiation: Evidence from a field audit. Tech. rep., MIT LFE Working Paper.

Jacob, Robin (2009). Patents and pharmaceuticals. A paper given on 29th November at the Presentation of the Directorate-General of Competitions Preliminary Report of the Pharma-sector inquiry.

Jaffe, A.B., Trajtenberg, M., and Fogarty, M.S. (2000). Knowledge spillovers and patent citations: Evidence from a survey of inventors. American Economic Review, vol. 90 (2), pp. 215–218.

Jaffe, A.B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. Quarterly Journal of Economics, vol. 108 (3), p. 577.

Jaffe, Adam B and Lerner, Josh (2011). Innovation and its discontents: How our broken patent system is endangering innovation and progress, and what to do about it. Princeton University Press.

Jaffe, Adam B and Trajtenberg, Manuel (2002). Patents, citations, and innovations: A window on the knowledge economy. MIT press.

Jensen, P.H. and Webster, E. (2004). Achieving the optimal power of patent rights. Australian Economic Review, vol. 37 (4), pp. 419–426.

Kahneman, Daniel and Tversky, Amos (1979). Prospect theory: An analysis of decision under risk. Econometrica: Journal of the Econometric Society, pp. 263–291.

Kahneman, Daniel and Tversky, Amos (2000). Choices, values, and frames. Cambridge University Press.

Kato, A. (2004). Patent pool enhances market competition. International Review of Law and Economics, vol. 24 (2), pp. 255–268.

Kesselheim, A.S. and Avorn, J. (2005). University-based science and biotechnology products. JAMA: the journal of the American Medical Association, vol. 293 (7), pp. 850–854.

Kiley, Thomas D (1992). Patents on random complementary dna fragments? Science, vol. 257 (5072), pp. 915–918.

Kim, S.H. (2004). Vertical structure and patent pools. Review of Industrial Organization, vol. 25 (3), pp. 231–250.

King, S.M. (2007). Clearing the patent thicket: The supreme court and congress undertake patent reform. Intell. Prop. & Tech. LJ, vol. 9, pp. 13–13.

Klein, Benjamin, Crawford, Robert G, and Alchian, Armen A (1978). Vertical integration, appropriable rents, and the competitive contracting process. Journal of law and economics, vol. 21 (2), pp. 297–326.

Kwon, I. (2012). Patent thicket, secrecy, and licensing. The Korean Economic Review, vol. 28 (1), pp. 27–49.

Lampe, R.L. and Moser, P. (2009). Do patent pools encourage innovation? evidence from the 19th-century sewing machine industry. Tech. rep., National Bureau of Economic Research.

Lampe, R.L. and Moser, P. (2012). Do patent pools encourage innovation? evidence from 20 us industries under the new deal. Tech. rep., National Bureau of Economic Research.

Lanjouw, Jean O and Schankerman, Mark (2001). Characteristics of patent litigation: A window on competition. RAND journal of economics, pp. 129–151.

Lanjouw, Jean O. and Schankerman, Mark (2004). Protecting intellectual property rights: Are small firms handicapped? Journal of Law and Economics, vol. 47 (1), pp. pp. 45–74.

Layne-Farrar, A., Padilla, A.J., and Schmalensee, R. (2007). Pricing patents for licensing in standard-setting organizations: Making sense of frand commitments. Antitrust LJ, vol. 74, p. 671.

Layne-Farrar, Anne and Lerner, Josh (2011). To join or not to join: Examining patent pool participation and rent sharing rules. International Journal of Industrial Organization, vol. 29 (2), pp. 294–303.

Leaffer, M. (2009). Patent misuse and innovation. J. High Tech. L., vol. 10, p. 142.

Lee, A. (2006). Examining the viability of patent pools for the growing nanotechnology patent thicket. Nanotech. L. & Bus., vol. 3, p. 317.

Lei, Z., Juneja, R., and Wright, B.D. (2009). Patents versus patenting: Implications of intellectual property protection for biological research.

Lemley, M. and Shapiro, C. (2006). Patent holdup and royalty stacking.

Lemley, M.A. (2005). Patenting nanotechnology. Stanford Law Review, pp. 601–630.

Lemley, M.A. and Shapiro, C. (2005). Probabilistic patents. The Journal of Economic Perspectives, vol. 19 (2), pp. 75–98.

Lemley, Mark A (2000). Rational ignorance at the patent office. Nw. UL Rev., vol. 95, p. 1495.

Lemley, Mark A (2007). Ten things to do about patent holdup of standards (and one not to). BCL Rev., vol. 48, p. 149.

Lerner, J. (1994). The importance of patent scope: an empirical analysis. The RAND Journal of Economics, pp. 319–333.

Lerner, J., Strojwas, M., and Tirole, J. (2003). The structure and performance of patent pools: Empirical evidence. Working paper.

Lerner, J. and Tirole, J. (2002). Efficient patent pools. Tech. rep., National Bureau of Economic Research.

Lerner, J. and Tirole, J. (2008). Public policy toward patent pools. In Innovation Policy and the Economy, Volume 8. University of Chicago Press, pp. 157–186.

Lerner, J., Tirole, J., and Strojwas, M. (2007). The design of patent pools: The determinants of licensing rules. The RAND Journal of Economics, vol. 38 (3), pp. 610–625.

Lerner, J. and Zhu, F. (2007). What is the impact of software patent shifts? evidence from lotus v. borland. International Journal of Industrial Organization, vol. 25 (3), pp. 511–529.

Lerner, Josh and Tirole, Jean (2005). The economics of technology sharing: Open source and beyond. The Journal of Economic Perspectives, vol. 19 (2), pp. pp. 99–120.

Leskovec, Jure, Kleinberg, Jon, and Faloutsos, Christos (2005). Graphs over time: densification laws, shrinking diameters and possible explanations. In Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining. ACM, pp. 177–187.

Lessig, Lawrence (2001). The future of ideas: The fate of the commons in a connected world. Vintage.

Levin, Richard C, Klevorick, Alvin K, Nelson, Richard R, Winter, Sidney G, Gilbert, Richard, and Griliches, Zvi (1987). Appropriating the returns from industrial research and development. Brookings papers on economic activity, vol. 1987 (3), pp. 783–831.

Liddell, Douglas (1976). Practical tests of 2 2 contingency tables. Journal of the Royal Statistical Society. Series D (The Statistician), vol. 25 (4), pp. pp. 295–304.

URL: *http://www.jstor.org/stable/2988087*

Lin, D. (2001). Research versus development: Patent pooling, innovation and standardization in the software industry. J. Marshall Rev. Intell. Prop. L., vol. 1, pp. 274–309.

Lin, L. (2011). Licensing strategies in the presence of patent thickets. Journal of Product Innovation Management, vol. 28 (5), pp. 698–725.

Liu, K., Arthurs, J., Cullen, J., and Alexander, R. (2008). Internal sequential innovations: How does interrelatedness affect patent renewal? Research Policy, vol. 37 (5), pp. 946–953.

Llanes, G. and Trento, S. (2009). Anticommons and optimal patent policy in a model of sequential innovation. Tech. rep.

Macdonald, S. (2004). When means become ends: Considering the impact of patent strategy on innovation. Information Economics and Policy, vol. 16 (1), pp. 135–158.

Mallo, L., Roox, K., Pike, J., Brown, A., Becker, S., and Thaler, G. (2008). Patent-related barriers to market entry for generic medicines in the european union: A review of weaknesses in the current european patent system and their impact on market access of generic medicines. Journal of Generic Medicines: The Business Journal for the Generic Medicines Sector, vol. 5 (4), pp. 255–280.

Mann, R.J. (2004). The myth of the software patent thicket. bepress Legal Series, p. 183.

Mann, Ronald J (2005). Do patents facilitate financing in the software industry? Texas Law Review, vol. 83, pp. 961–1009.

Maskus, K.E. (2006). Reforming us patent policy: Getting the incentives right. Innovations: Technology, Governance, Globalization, vol. 1 (4), pp. 127–153.

Masur, J.S. (2010). Costly screens and patent examination. Journal of Legal Analysis, vol. 2 (2), pp. 687–734.

McFadden, Daniel, Machina, Mark J, and Baron, Jonathan (2000). Rationality for economists? In Elicitation of Preferences. Springer, pp. 73–110.

Ménière, Y. (2008). Patent law and complementary innovations. European Economic Review, vol. 52 (7), pp. 1125–1139.

Merges, Robert P (1996). Contracting into liability rules: Intellectual property rights and collective rights organizations. Cal L. Rev., vol. 84, p. 1293.

Merges, Robert P (1999). Institutions for intellectual property transactions: the case of patent pools. University of California at Berkeley Working Paper.

Merges, R.P. (2006). Introductory note to brief of amicus curiae in ebay v. mercexchange. Berkeley Tech. LJ, vol. 21, pp. 997–1016.

Mertes, M.M.M. and Stötter, G. (2010). Managing the patent thicket and maximizing patent lifetime in vaccine technology. Human Vaccines, vol. 6 (10), pp. 860–863.

Meurer, M.J. (2002). Business method patents and patent floods. Wash. UJL & Pol'y, vol. 8, pp. 309–342.

Moore, Don A and Healy, Paul J (2008). The trouble with overconfidence. Psychological Review, vol. 115 (2), pp. 502–517.

Moore, Kimberly A (2000). Judges, juries, and patent cases-an empirical peek inside the black box. Mich. L. Rev., vol. 99, p. 365.

Mossoff, A. (2009). A stitch in time: The rise and fall of the sewing machine patent thicket. George Mason University School of Law.

Mossoff, A. (2011). Rise and fall of the first american patent thicket: The sewing machine war of the 1850s, the. Ariz. L. Rev., vol. 53, p. 165.

Mowery, David C (2004). Ivory tower and industrial innovation: university-industry technology transfer before and after the Bayh-Dole act in the United States. Stanford University Press.

Muris, T.J. (2001). Competition and Intellectual Property Policy: The Way Ahead. US FTC.

Murray, Fiona and Stern, Scott (2007). Do formal intellectual property rights hinder the free flow of scientific knowledge?: An empirical test of the anti-commons hypothesis. Journal of Economic Behavior & Organization, vol. 63 (4), pp. 648–687.

Nagaoka, S. and Nishimura, Y. (2006). An empirical assessment of the effects of patent thickets. unpublished, Hitotsubashi University.

Napoleon, V.J. (2009). Impact of global patent and regulatory reform on patent strategies for biotechnology. Pitt. J. Tech. L. & Pol'y, vol. 9, pp. 1–31.

Nielsen, C.M. and Samardzija, M.R. (2006). Compulsory patent licensing: Is it a viable solution in the united states. Mich. Telecomm. & Tech. L. Rev., vol. 13, p. 509.

Noel, M.D. and Schankerman, M. (2006). Strategic patenting and software innovation.

Nordhaus, William D (2004). Schumpeterian profits in the american economy: Theory and measurement. Tech. rep., National Bureau of Economic Research.

Palangkaraya, Alfons, Webster, Elizabeth, and Jensen, Paul H. (2011). Misclassification between patent offices: Evidence from a matched sample of patent applications. The Review of Economics and Statistics, vol. 93 (3), pp. pp. 1063–1075.

Paredes, J.P. (2006). Written description requirement in nanotechnology: Clearing a patent thicket. J. Pat. & Trademark Off. Soc'y, vol. 88, p. 489.

Rai, Arti K. (2003). Engaging facts and policy: A multi-institutional approach to patent system reform. Columbia Law Review, vol. 103 (5), pp. pp. 1035–1135.

Regibeau, Pierre and Rockett, Katharine (2011). Assessment of potential anticompetitive conduct in the field of intellectual property rights and assessment of the interplay between competition policy and ipr protection. In Prepared for the European Commission, Publications Office of the European Union, ISBN 978-92-79-22210-8.

Reitzig, M. (2004). The private values of ?thickets? and ?fences?: towards an updated picture of the use of patents across industries. Economics of Innovation and New Technology, vol. 13 (5), pp. 457–476.

Rey, P. and Salant, D. (2012). Abuse of dominance and licensing of intellectual property. International Journal of Industrial Organization.

Rossman, Joseph and Sanders, Barkev S (1957). Patent utilization study, the. Pat. Trademark & Copy. J. Res. & Ed., vol. 1, p. 74.

Rubinfeld, D.L. and Maness, R. (2004). The strategic use of patents: Implications for antitrust. In Antitrust, Patent and Copyright Conference.

Sabety, T. (2004). Nanotechnology innovation and the patent thicket: Which ip policies promote growth. Alb. LJ Sci. & Tech., vol. 15, p. 477.

Santore, Rudy, McKee, Michael, and Bjornstad, David (2010). Patent pools as a solution to efficient licensing of complementary patents? some experimental evidence. Journal of Law and Economics, vol. 53 (1), pp. pp. 167–183.

Sass, Samuel (1989). A patently false patent myth. The Skeptical Inquirer, vol. 13, pp. 310–313.

Schacht, W.H. (2006). Patent Reform: Issues in the Biomedical and Software Industries. Congressional Research Service (7 April 2006), available at: http://sharp.sefora.org/wp-content/uploads/2007/12/rl33367.pdf.

Schmalensee, R. (2009). Standard-setting, innovation specialists and competition policy. The Journal of Industrial Economics, vol. 57 (3), pp. 526–552.

Schmidt, K. (2008). Complementary patents and market structure.

Schneider, C. (2008). Fences and competition in patent races. International Journal of Industrial Organization, vol. 26 (6), pp. 1348–1364.

Schumpeter, Joseph Alois (1934). The Theory of Economic Development: An Inquiry Into Profits, Capital, Credit, Interest, and the Business Cycle, vol. 46. Transaction Publishers.

Schumpeter, Joseph Alois (1939). Business cycles: a theoretical and statistical analysis of the capitalist process. McGraw-Hill.

Scotchmer, Suzanne (1991). Standing on the shoulders of giants: cumulative research and the patent law. The Journal of Economic Perspectives, vol. 5 (1), pp. 29–41.

Shand, H. and Wetter, K.J. (2007). Trends in intellectual property and nanotechnology: Implications for the global south. Journal of Intellectual Property Rights, vol. 12, pp. 111–117.

Shapiro, C. (2001). Navigating the patent thicket: Cross licenses, patent pools, and standard setting. In Innovation Policy and the Economy, Volume 1. pp. 119–150.

Shapiro, C. (2003). Antitrust limits to patent settlements. RAND Journal of Economics, pp. 391–411.

Shapiro, Carl (2010). Injunctions, hold-up, and patent royalties. American Law and Economics Review, vol. 12 (2), pp. 280–318.

Siebert, R. and Von Graevenitz, G. (2010a). Jostling for advantage or not: Choosing between patent portfolio races and ex ante licensing. Journal of economic behavior & organization, vol. 73 (2), pp. 225–245.

Siebert, R.B. and Von Graevenitz, G. (2010b). Licensing in the patent thicket-timing and benefits.

Skitol, Robert A (2005). Concerted buying power: Its potential for addressing the patent holdup problem in standard setting. Antitrust Law Journal, vol. 72 (2), pp. 727–744.

Somaya, D., Teece, D., and Wakeman, S. (2011). Innovation in multi-invention contexts: Mapping solutions to technological and intellectual property complexity. California Management Review, vol. 53 (4), pp. 47–79.

Somaya, Deepak (2003). Strategic determinants of decisions not to settle patent litigation. Strategic Management Journal, vol. 24 (1), pp. pp. 17–38.

Sternitzke, C., Bartkowski, A., and Schramm, R. (2008). Visualizing patent statistics by means of social network analysis tools. World Patent Information, vol. 30 (2), pp. 115–131.

Strandburg, K.J. (2006). Law and the science of networks: An overview and an application to the "patent explosion". bepress Legal Series, p. 1617.

Taylor, M.R. and Cayford, J. (2003). American patent policy, biotechnology, and african agriculture: The case for policy change. Harv. JL & Tech., vol. 17, p. 321.

Team, Intellectual Property Office Patent Informatics (2011). Patent thickets: An overview. UK Intellectual Property Office.

Trajtenberg, M. (1990). A penny for your quotes: patent citations and the value of innovations. The RAND Journal of Economics, pp. 172–187.

Trajtenberg, M., Henderson, R., and Jaffe, A. (1997). University versus corporate patents: A window on the basicness of invention. Economics of Innovation and New Technology, vol. 5 (1), pp. 19–50.

Tullis, T.K. (2005). Application of the government license defense to federally funded nanotechnology research: The case for a limited patent compulsory licensing regime. UCLA L. Rev., vol. 53, p. 279.

Tullock, Gordon (1967). The welfare costs of tariffs, monopolies, and theft. Economic Inquiry, vol. 5 (3), pp. 224–232.

Van Overwalle, Geertrui (2010). Turning patent swords into shares. Science, vol. 330 (6011), pp. pp. 1630–1631.

Van Zimmeren, E., Verbeure, B., Matthijs, G., and Van Overwalle, G. (2006). A clearing house for diagnostic testing: the solution to ensure access to and use of patented genetic inventions? Bulletin of the World Health Organization, vol. 84 (5), pp. 352–359.

Verbeure, B., van Zimmeren, E., Matthijs, G., and Van Overwalle, G. (2006). Patent pools and diagnostic testing. TRENDS in Biotechnology, vol. 24 (3), pp. 115–120.

Von Graevenitz, G., Wagner, S., and Harhoff, D. (2011). How to measure patent thickets – a novel approach. Economics Letters, vol. 111 (1), pp. 6–9.

Von Graevenitz, G., Wagner, S., and Harhoff, D. (2012). Incidence and growth of patent thickets-the impact of technological opportunities and complexity. Journal of Industrial Economics.

Wagner, R. Polk (2003). Information wants to be free: Intellectual property and the mythologies of control. Columbia Law Review, vol. 103 (4), pp. pp. 995–1034.

Walsh, John P, Arora, Ashish, and Cohen, Wesley M (2003). Effects of research tool patents and licensing on biomedical innovation. Patents in the Knowledge-based Economy, vol. 285, p. 286.

Wang, A.W. (2010). Rise of the patent intermediaries. Berkeley Tech. LJ, vol. 25, p. 159.

Williamson, O.E. (1971). The vertical integration of production: market failure considerations. The American Economic Review, vol. 61 (2), pp. 112–123.

Williamson, O.E. (1979). Transaction-cost economics: the governance of contractual relations. Journal of law and economics, vol. 22 (2), pp. 233–261.

Ziedonis, R.H. (2004). Don't fence me in: Fragmented markets for technology and the patent acquisition strategies of firms. Management Science, vol. 50 (6), pp. 804–820.

9 Appendix

Table 9: Thicket Type and Variant Definitions

Type	Definition
A	**Products require many diversely-held complementary inputs which are protected by discrete patents. Patents are valid and correctly issued by the patent office.**
−Aa	Patented inputs are perfectly complementary, so that the full set of inputs must be assembled.
−Ab	Patented inputs are imperfectly complementary, so that varying degrees of functionality can be acheived with certain sub-sets of inputs.
B	**The claims of two or more patents have overlapping domains of application. All patents are valid and correctly issued by the patent office. Patents may be characterized as 'broad'.**
−B1	Invention is cumulative and an upstream patent's claimed domain of application overlaps on a downstream patent's claimed domain of application.
−−B1a	Flow-through rights are made explicit in a Reach-Through Licensing Agreement (RTLA). The downstream party licenses the upstream patent prior to originating the downstream invention, and assigns rights back upstream accordingly.
−−B1b	The downstream party licenses the upstream patent without a Reach-Through Licensing Agreement (RTLA), or does not license the upstream patent.
−B2	Invention is adjacent. Although the inventive steps of two or more patents are different, they provide closely-related (or identical) economic functionality. As such the claims of the patents overlap in their domain of application. This is more likely to occur with broad patents that have large domains of application.
C	**Patents, or their claims, are spurious either because they fail to meet the stated requirements for validity or because they intentional undermine the rationale behind the patent system.**
−C1	At least a subset of patent claims do not embody a new or non-obvious inventive step. The patent office makes mistakes in the issue of some patents. Applicants take advantage of these mistakes to seek patents that cover unpatented or patented prior-art of historic, concurrent, or future inventions, creating apparent infringement.
−−C1a	Patents do not embody a novel inventive step in each and every claim. Thus multiple patents have claims, and have been granted exclusionary rights, covering the same domain of application.
−−C1b	Firms seek patents that fail to meet the non-obviousness requirement given the prior art. Thus patents have claims, and have been granted exclusionary rights, over a domain of application which was obvious to a person skilled in the art at the time of the patent application.
−C2	Through the use of continuations, applicants can keep their patents in an application phase and adjust their claims. Prior to November 2000, patent applications were not disclosed. This made so called 'submarine patents' possible. An inventor might then incur costly research and development only to find that the resulting invention infringed another patent.
D	**A number of patents effectively saturate the potential invention space. Patents may be small and cover 'marginal' inventions.**

Variant	Definition
C	All information about patented inputs is known to the producer. Search and transactions costs are not considered (or are assumed sufficiently low that they can be ignored).
T	All patented inputs are known to the producer. Transactions costs are considered material and may prevent successful negotiation of usage rights for patented inputs.
S	The number of patented inputs is sufficiently high that search costs become prohibitive or cognitive biases prevent complete identification. Some subset of input patent usage rights may be unsecured.
P	Patents do not have guaranteed validity and incontrovertible rights to exclude usage to the inventions detailed in their claims. Instead patents are always potentially subject to being ruled invalid in general or inapplicable in some specific application. The courts are the ultimate arbitrators of patent rights, and may decide rights on a case by case basis with imperfect judgment.
PB	One or more pre-existing patents are broad with a probabilistic diffusion of rights. Such a broad patent has a 'core' domain of application which it covers with validity (or a high-likelihood of validity), but as the distance from the core increases the likelihood that the patent would be deemed invalid or inapplicable increases.

Note that variant C are hypothetical, existing solely in the minds of economists, and are not found in the wild.

Table 10: Core papers

Paper	Discipline	Stance	Definitions	Theory	Empirical	Survey	Measures	Discussion
Bessen (2003)	Econ	Pro	A,C1-T	✓				
Clarkson (2004)	Econ	Pro	A-T,B-T				✓	
Clarkson (2005)	Econ	Pro	A-T,B-T				✓	
Clarkson and DeKorte (2006)	Mgmt	Pro	A-T,B-T					✓
Cockburn and Macgarvie (2011)	Mgmt	Pro	A-T		✓			
Cockburn and MacGarvie (2009)	Econ	Weak Pro	A-ST,B1-ST		✓			
Cockburn et al. (2010)	Econ	Weak Anti	A-T		✓			
Entezarkheir (2010)	Mgmt	Weak Pro	A-T		✓			
EPC and Board (2013)	Plcy Rpt	Weak Anti	A-T,B					✓
Epstein and Kuhlik (2004)	Law	Anti	A,B					✓
Evans and Layne-Farrar (2004)	Mgmt	Weak Anti	A,B					✓
Farrell (2009)	Econ	Assumed Pro						✓
Galasso (2007)	Econ	Assumed Pro	A-S	✓	✓			
Galasso and Schankerman (2010)	Econ	Anti	A,B		✓			
George (2006)	Law	Assumed Pro						✓
Geradin et al. (2007)	Econ	Neutral	A-T,B-T		✓			
Geradin et al. (2008)	Econ	Assumed Pro	A					✓
Hall and Ziedonis (2001)	Econ	Assumed Pro	A,B	✓	✓			
Hall et al. (2012)	Plcy Rpt	Pro	A,B		✓			
Hall and Ziedonis (2007)	Econ	Assumed Pro	A		✓			✓
Hargreaves (2011)	Plcy Rpt	Assumed Pro	B-ST					✓
Harhoff et al. (2012)	Econ	Assumed Pro	B		✓			
Harhoff et al. (2008)	Econ	Pro	A-ST,B-ST,C1,D		✓			
Harhoff et al. (2007)	Plcy Rpt	Pro	A-T	✓				✓
Hegde et al. (2009)	Econ	Pro	C		✓			
Heller and Eisenberg (1998)	Law	Pro	A,B	✓				
Heller (1997)	Law	Pro	A,B	✓				
Huang and Murray (2009)	Mgmt	Pro	A,B		✓			
Kiley (1992)	Law	Pro	B,C1					✓
Lessig (2001)	Law	Pro	A,B					✓
Mann (2004)	Law	Anti	A,B			✓		
Mann (2005)	Law	Anti	A-ST,B-ST	✓		✓		
Merges (1996)	Law	Pro	A-T	✓				✓
Merges (1999)	Law	Neutral	A-T,B-T					✓
Mossoff (2011)	Econ	Pro	A-T,B-T					✓
Mossoff (2009)	Econ	Pro	A-T,B-T					✓
Murray and Stern (2007)	Econ	Weak Pro	B-T		✓			✓
Nagaoka and Nishimura (2006)	Mgmt	Weak Pro	A-T		✓			
Noel and Schankerman (2006)	Econ	Weak Pro	A-T		✓			
Regibeau and Rockett (2011)	Plcy Rpt	Neutral	A-ST,B-T					✓
Reitzig (2004)	Econ	Weak Pro	A	✓				
Shapiro (2001)	Econ	Pro	A	✓				✓
Siebert and Von Graevenitz (2010a)	Econ	Weak Pro	A-ST,B-ST,C1	✓	✓			
Siebert and Von Graevenitz (2010b)	Econ	Pro	A-T,D	✓	✓			
Sternitzke et al. (2008)	Econ	Pro	D		✓		✓	
Strandburg (2006)	Law	Neutral	A-T				✓	✓
Team (2011)	Plcy Rpt	Neutral	A,B,C1					✓
Von Graevenitz et al. (2012)	Econ	Pro	A,B		✓			
Von Graevenitz et al. (2011)	Econ	Pro	A,B				✓	
Walsh et al. (2003)	Law	Weak Anti	A-T,B-T					✓
Ziedonis (2004)	Econ	Pro	A-PST		✓		✓	

Table 11: Downstream papers (Aggarwal to Kim)

Paper	Discipline	Stance	Definitions	Theory	Empirical	Survey	Measures	Discussion
Aggarwal and Hsu (2009)	Econ	Weak Anti	A, B		✓			
Allison and Tiller (2003)	Law	Assumed Pro	A-S,B,C1		✓			
Andrews (2002)	Gen Sci	Weak Pro	A					✓
Aoki and Schiff (2008)	Econ	Assumed Pro	A	✓				
Arundel and Patel (2003)	Plcy Rpt	Neutral	A-T,B					✓
Attaran (2004)	Plcy Rpt	Weak Anti	A-T,D					✓
Ayres and Parchomovsky (2007)	Law	Pro	Aa-T	✓				
Baluch et al. (2005)	Law	Anti	B					✓
Baron and Delcamp (2010)	Econ	Pro	A-T		✓			
Baron and Pohlmann (2011)	Econ	Pro	A-T,C1	✓	✓			
Barpujari (2010)	Mgmt	Weak Pro	B1					✓
Barton (2002)	Law	Weak Pro	A	✓				
Baumol (2004)	Econ	Assumed Pro	A					✓
Bawa (2007)	Law	Pro	B,C1					✓
Bawa et al. (2005)	Law	Pro	B	✓				
Bawa (2005)	Gen Sci	Assumed Pro	B					✓
Beard and Kaserman (2002)	Law	Pro	A-S,B-S	✓				
Bergman and Graff (2007)	Gen Sci	Weak Pro	B-T					✓
Braun and Herstatt (2007)	Econ	Pro	B	✓				
Burk and Lemley (2003)	Law	Pro	B2,C1	✓				
Calderini and Giannaccari (2006)	Econ	Pro	A-T					✓
Callaway (2008)	Law	Weak Pro	A-S	✓				
Carrier (2003)	Law	Assumed Pro	A-S	✓				
Carrier (2002)	Law	Pro	B1					✓
Carrier (2004)	Law	Pro	B,C	✓				
Choi (2005)	Econ	Weak Pro	C1	✓				
Cohen and Walsh (2008)	Econ	Weak Anti	A-T		✓			
Commission (2011)	Plcy Rpt	Pro	B,C1					✓
Commission (2003)	Plcy Rpt	Pro	B,C1					✓
Competition (2008)	Plcy Rpt	Pro	B,C1,D		✓			
Cowin et al. (2007)	Plcy Rpt	Weak Pro	B1,D-S					✓
D'Silva (2009)	Law	Weak Pro	B2-T					✓
Devlin (2009)	Law	Pro	A-S,B-S	✓				
Dhar and Foltz (2007)	Econ	Weak Pro	B1	✓	✓			
Eisenmann (2008)	Mgmt	Pro	A					✓
Eisenstein (2010)	Gen Sci	Assumed Pro	C1a					✓
Feldman (2004)	Law	Neutral	B-T	✓				✓
Feldman and Nelson (2008)	Law	Weak Pro	A-T,B-T	✓				
Gallini (2011)	Econ	Pro	A,B	✓				
Ganslandt (2009)	Econ	Weak Pro	A-T					✓
Gaulé (2006)	Mgmt	Weak Pro	A-T,B-T					✓
Gilbert (2010)	Law	Pro	A-ST					✓
Glover (2002)	Law	Neutral	B1					✓
Goozner (2006)	Gen Sci	Pro	A-T					✓
Hall (2007)	Econ	Assumed Pro	A-S	✓				✓
Hemphill (2003)	Econ	Pro	D					✓
Holman (2006)	Law	Weak Anti	B1-S					✓
Holman (2008)	Gen Sci	Anti	B1					✓
Holman (2005)	Law	Anti	B1,D					✓
Horn (2003)	Law	Pro	B					✓
Hussinger (2006)	Econ	Neutral	D		✓			
Iyama (2005)	Law	Pro	B					✓
Jacob (2009)	Plcy Rpt	Neutral	C1, D					✓
Jensen and Webster (2004)	Econ	Pro	C1	✓				
Kato (2004)	Econ	Assumed Pro	A, C1	✓				
Kesselheim and Avorn (2005)	Gen Sci	Weak Pro	B1					✓
Kim (2004)	Econ	Pro	A-T	✓				

Table 12: Downstream papers (King to Wang)

Paper	Discipline	Stance	Definitions	Theory	Empirical	Survey	Measures	Discussion
King (2007)	Law	Assumed Pro	B,C1					✓
Kwon (2012)	Econ	Weak Pro	A	✓				
Lampe and Moser (2012)	Econ	Weak Pro	C1,D		✓			
Lampe and Moser (2009)	Econ	Pro	B		✓			
Lanjouw and Schankerman (2004)	Econ	Assumed Pro	A		✓			
Layne-Farrar et al. (2007)	Econ	Assumed Pro	C1,D	✓				
Layne-Farrar and Lerner (2011)	Econ	Assumed Pro	A,B		✓			
Leaffer (2009)	Law	Weak Pro	B,C1	✓				
Lee (2006)	Law	Assumed Pro	B					✓
Lei et al. (2009)	Gen Sci	Anti	A, B			✓		
Lemley and Shapiro (2006)	Econ	Pro	A	✓				
Lemley (2005)	Law	Pro	A,B					✓
Lemley and Shapiro (2005)	Econ	Pro	A	✓				
Lerner and Tirole (2008)	Econ	Pro	Aa-T	✓				
Lerner and Tirole (2002)	Econ	Weak Pro	A,B	✓				
Lerner and Zhu (2007)	Econ	Weak Pro	A,B		✓			
Lerner et al. (2003)	Econ	Assumed Pro	A,B		✓			
Lerner et al. (2007)	Econ	Assumed Pro	A,B	✓	✓			
Lerner and Tirole (2005)	Econ	Pro	A	✓				
Lin (2001)	Law	Assumed Pro	A-T	✓				
Lin (2011)	Econ	Assumed Pro	B1,D	✓				
Liu et al. (2008)	Econ	Assumed Anti	B1		✓			
Llanes and Trento (2009)	Econ	Neutral	A	✓				
Macdonald (2004)	Mgmt	Weak Pro	A-T,B-T					✓
Mallo et al. (2008)	Gen Sci	Weak Pro	C1					✓
Maskus (2006)	Econ	Weak Pro	A					✓
Masur (2010)	Law	Assumed Pro	C1,D-S					✓
Merges (2006)	Law	Pro	A,C1					✓
Mertes and Stötter (2010)	Gen Sci	Weak Pro	A					✓
Meurer (2002)	Law	Weak Pro	A					✓
Muris (2001)	Plcy Rpt	Weak Anti	A-S					✓
Ménière (2008)	Econ	Assumed Pro	A-T	✓				
Napoleon (2009)	Law	Weak Pro	A,B1,D					✓
Nielsen and Samardzija (2006)	Law	Assumed Pro	A-T					✓
Palangkaraya et al. (2011)	Econ	Assumed Pro	C1		✓			
Paredes (2006)	Law	Assumed Pro	B1					✓
Rai (2003)	Law	Neutral	A,C1	✓				✓
Rey and Salant (2012)	Econ	Assumed Pro	A	✓				
Rubinfeld and Maness (2004)	Law	Neutral	C1,C2					✓
Sabety (2004)	Law	Neutral	A,B1					✓
Santore et al. (2010)	Econ	Assumed Pro	A	✓	✓			
Schacht (2006)	Plcy Rpt	Neutral	A,B,D					✓
Schmalensee (2009)	Econ	Weak Anti	C1-ST	✓				
Schmidt (2008)	Econ	Pro	A	✓				
Schneider (2008)	Econ	Assumed Pro	A,B2,D	✓				
Shand and Wetter (2007)	Law	Weak Pro	B					✓
Shapiro (2003)	Econ	Assumed Pro	A,C1	✓				
Somaya et al. (2011)	Mgmt	Pro	A-ST	✓				
Somaya (2003)	Mgmt	Anti	B2,D	✓	✓			
Taylor and Cayford (2003)	Law	Pro	B1	✓				
Tullis (2005)	Law	Assumed Pro	A,B					✓
Van Overwalle (2010)	Gen Sci	Weak Pro	A,B1					✓
Van Zimmeren et al. (2006)	Plcy Rpt	Weak Pro	A-T					✓
Verbeure et al. (2006)	Gen Sci	Weak Pro	A-T,B1-T					✓
Wagner (2003)	Law	Anti	A	✓				✓
Wang (2010)	Law	Assumed Pro	A-ST					✓