

Economic theories of the non-obviousness requirement for patentability: A survey¹

Vincenzo Denicolò
University of Bologna

Abstract. This paper reviews several approaches to the non-obviousness requirement for patentability developed in the economics literature and discusses their limits. There have been four main approaches: option value, sequential innovation, error-cost, and complementary innovation. These approaches share the premise that patenting may impose negative externalities and so is not always socially beneficial. In particular, when innovation is sequential, granting patent protection to trifling improvements of a basic technology may lower the pace of technological progress. Similarly, with complementary innovation the overall incentive to innovate may be harmed if every minor component of a complex technology is separately patentable. In such circumstances, it may be desirable to deny patent protection even to genuine innovations.

¹ This article is a revised version of a paper presented at the conference on “Nonobviousness: The Shape of Things to Come”, held at the Lewis and Clark Law School in October 2007. I am grateful to the conference participants, and in particular my discussant Michael Katz, for helpful comments. Mailing address: Vincenzo Denicolò, Department of Economics, University of Bologna, Piazza Scaravilli, 2; I- 40126 Bologna, Italy; e-mail vincenzo.denicolo@unibo.it.

1. Introduction

What is the economic function of the non-obviousness requirement in patent law? Answering this question may help shed light on various tests of non-obviousness that have been proposed over the years, from “flash of creative genius” to the “teaching, suggestion, or motivation” approach recently struck down by the KSR Supreme Court ruling.² This paper reviews several approaches to non-obviousness developed in the economics literature and discusses their limits.

In assessing patentability, the most important issue faced by patent offices and the courts is, perhaps, which verifiable characteristics of the innovation should be considered to ascertain whether the non-obviousness requirement is met. Unfortunately, economists are ill equipped to address this question, since the way an innovation is conceptualized is a pre-analytical modelling choice about which economic theory provides little guidance. By and large, the economics literature has implicitly assumed that the relevant variable for assessing non-obviousness is the “size” of the innovation. The exact meaning of “size” depends on the specifics of the model: it may be the magnitude of the cost reduction for process innovations, it may be the extent of the quality improvement, it may be the size of the market for the new product, but in any case it is something positively correlated to the value of the innovation. The obvious alternative to “value” is “cost”: this option, however, has proven less popular, perhaps because R&D expenditures are notoriously hard to verify.³

Once the choice of how to describe an innovation is made, economists are much better equipped to analyze the optimal threshold for patentability, or at least to clarify the underlying trade-offs. There have been several approaches: option value, sequential innovation, error-cost, and complementary innovation. The option value approach is based on the notion of irreversible investment, and draws an analogy between the introduction of a new product or device and the exercise of an option. It argues that in a market economy there is a tendency to exercise the option too early, as today’s innovators do not internalize the externality they impose on future potential innovators by depriving them of the opportunity to develop a better innovation. The non-obviousness requirement serves to protect future innovators against such premature inventions.

In the sequential innovation approach, by contrast, the non-obviousness requirement is viewed as a tool for protecting early innovators against competition from subsequent improvements of their innovative technological knowledge. In a sense, this is the dual of the option value approach.

² See *KSR v. Teleflex*, 550 U.S. 127 S. Ct. 1727 (2007).

³ In addition, the U.S. Patent Act states that patentability “shall not be negated by the manner in which the invention was made” (35 U.S.C. § 103 (A)), which may be interpreted as a prohibition of taking R&D costs into account in assessing patentability.

The error-cost approach views non-obviousness simply as a strengthening of the novelty requirement, designed to cope with the risk of type II errors – granting a monopoly over a technology that is already in the public domain.

Finally, the complementary innovation approach focuses on the issue of the division of profit between different firms that jointly contributed to the discovery of a complex technology comprising several complementary components. It argues that in certain circumstances denying patentability to some innovative components may improve the overall incentives to innovate.

While all of these approaches offer useful insights, none provides compelling policy prescriptions. One must recognize that all that legal scholars and practitioners can take away from existing economic analyses of the non-obviousness requirement is a clarification of certain basic, abstract questions. Why should some minimal inventive step be required on top of novelty for an innovation to be patentable? How can it be optimal to grant either full patent protection or no protection at all? These are certainly important issues, and genuine progress has been made in this area. But much work remains to be done before economic analysis can provide practical guidelines for assessing non-obviousness.

I start my survey with a discussion of the “all-or-nothing” property of current patent policy, whereby an innovation may be eligible to full patent protection or no protection at all. Then, in Sections 3 to 6 I discuss each of the approaches outlined above. Section 7 concludes the paper suggesting some directions for future research.

2. Patent policy: an “all-or-nothing” solution?

Under current patent law, an obvious innovation (whatever this exactly means) cannot be patented, but as soon as the non-obviousness requirement is met the innovation enjoys a fixed level of protection, irrespective of how “non-obvious” the innovation is. How can such an “all-or-nothing” policy be optimal?

The story many law and economics scholars have in mind is very simple. Innovations vary in size. The optimal level of patent protection, so the argument goes, is low for small innovations and increases smoothly with innovation size. Ideally, then, policy should tailor the degree of protection to each particular innovation. In practice, however, it is very difficult, if not impossible, to fine tune the policy – the well known “one-size-fits-all” problem. The ideal curve must then be approximated with a step-wise function, implying that patent protection is granted only above a threshold innovation size, which becomes the minimal “inventive step” for patentability.

This argument can be illustrated graphically with the help of Figure 1. The increasing curve depicts the hypothetically optimal level of patent protection as a function of the size of the innovation. As drawn, the optimal protection curve is smoothly increasing. The Figure also depicts a step-wise

approximation of the ideal curve, which accounts for the constraint that patent protection can be granted or denied, but cannot be fine tuned.

[insert Figure 1 around here]

Clearly, this entire argument rests on the ideal curve being increasing. But this is far from obvious. An increase in the size of an innovation generally increases the investment in research that society ought to make in that innovation. However, for any given level of patent protection, it also increases the inventor's prospective profits and so automatically increases equilibrium R&D investments. It is uncertain whether there is any need to adjust the strength of patent protection and, if so, in what direction.

Starting with Nordhaus (1969), economists have developed a formal economic analysis of the optimal strength of patent protection, which suggests that the optimal level of patent protection is the outcome of a trade-off between fast innovation and static allocative distortion. Stronger patent protection entails both social benefits (innovative activity is encouraged) and costs (society bears larger monopoly deadweight losses), and the optimal level of protection must equate the marginal social benefits and costs. As shown in Figure 2, both are larger for larger innovations, but it is uncertain whether the intersection point moves to the right or to the left.

Within this framework, Shapiro (2007) and Denicolò (2007) have shown that under reasonable simplifying assumptions the key determinant of the optimal degree of patent protection is the elasticity of the supply of inventions – i.e., the percentage increase in the number of innovations associated with a 1% increase in R&D expenditures. If the elasticity is constant the optimal level of protection is independent of innovation size; this is the case depicted in Figures 2 and 3. In Figure 2, the increasing curves represent the marginal social cost of patent protection, and the decreasing curves the marginal social benefit. The upper, solid curves represent a large innovation and the lower, dashed curves a small one. As drawn, the points of intersection of the two pairs of curves have the same abscissa.

[insert Figure 2 around here]

[insert Figure 3 around here]

When the elasticity is not constant, the size of the innovation affects the optimal level of protection indirectly, as it determines the point on the “innovation production function.” However, it is hard to tell whether the elasticity raises, falls or stays constant with the level of investment in research, and hence whether the optimal degree of patent protection increases, decreases or stays constant with the size

of the innovation.⁴ The size of the innovation may also affect the optimal degree of protection through other, more indirect channels, but again the effect is generally uncertain.⁵ In sum, economic analysis does not seem to support the claim that the ideal level of patent protection is a smoothly increasing function of innovation size.

Even if the optimal level of patent protection is constant or decreasing, however, the expected total net benefit from large innovations does exceed that from small ones. (This is illustrated in Figure 3, where the dashed curve again represents a small innovation and the upper, solid curve a large one.) Starting from this premise, a more promising approach focuses on the possibility that patents may entail fixed social costs, in addition to those, like monopoly deadweight losses, which vary with the level of patent protection and so are already included in the curves in Figure 3.

If there are no such fixed social costs, as it is implicitly assumed in Figure 3, all innovations should be patentable and receive the same level of protection. But if patenting entails social costs that are independent of the level of protection, this will shift down the social benefit curves (Figure 4), implying that not all innovations should now be protected. The optimal non-obviousness requirement would generally depend both on the size of the innovation (which determines the height of the gross social benefit curves depicted in Figure 3) and the magnitude of the externality.

[insert Figure 4 around here]

Notice that if those fixed costs were borne by the patent-holder, there would be no need of a non-obviousness requirement:⁶ if the costs exceeded the benefits, the inventor would choose not to seek patent protection and so the socially optimal decision would be implemented even if the innovation was patentable.⁷ To justify a non-obviousness requirement, the fixed social cost must be a negative externality that the innovator inflicts on somebody else. Economists have identified several possible externalities of this kind, which I now turn to review.

⁴ In the standard Poisson model of patent races, the elasticity is decreasing and so the optimal degree of protection decreases in the size of the innovation: see Denicolò (1999). More generally, Boldrin and Levine (2005) argue that the elasticity necessarily falls for extremely large values of the investment in research. Although theirs is an asymptotical result, it may be taken to mean that the elasticity is decreasing over the relevant range, suggesting that the optimal level of patent protection actually decreases with the size of the innovation.

⁵ See e.g. Berkowitz and Kotowitz (1982) and De Brock (1985).

⁶ Consider, for instance, the administrative and legal costs associated with the patent application and review process. To the extent that patent offices are funded by patent fees and renewal fees, most of the costs are eventually borne by patent-holders and would-be patent holders.

⁷ To the extent that the social benefit from the innovation exceeds the private benefit, the market equilibrium would be characterized by under-investment in research and under-patenting – a distortion that cannot be cured by making the innovation unpatentable.

3. The option value approach

The option value approach is based on the notion of irreversible investment, and draws an analogy between the introduction of a new product or device and the exercise of an *option*. In their baseline model, Erkal and Scotchmer (2007) envision a population of potential innovators, to whom ideas to fill a market niche occur randomly. (There is a constant instantaneous probability that an idea occurs to somebody, which is a key parameter of the model.) Different ideas have different implementation costs, but are otherwise perfectly substitutable. Hence, if patent protection is sufficiently broad, or else product market competition is sufficiently intense, after the first idea is implemented and the innovation is patented the innovation process stops, as a second innovator could not recoup his innovation cost.

These assumptions may sound extreme, but it is undeniable that in addition to opening the way to subsequent improvements, early innovations may have a pre-emptive effect. In a market economy, when deciding whether or not to invest in an idea, private inventors do not internalize the externality they are imposing on future potential innovators by depriving them of the opportunity to develop a better innovation – i.e., one with a lower implementation cost. Hence, if inventors could appropriate the value of the innovation fully, the market equilibrium would exhibit over-investment in research; on average, that is to say, the innovation would occur too early at too high a cost.

As a result, there is room for policy, which should aim at impeding premature innovations. In the model, the first best solution can indeed be achieved by a policy that grants patent protection only when the implementation cost is below a critical threshold. Such a threshold is explicitly calculated by Erkal and Scotchmer, who also develop an instructive comparative statics analysis. Intuitively, if the probability that ideas occur is large, the option value of not investing is also large, and hence the threshold should be low. Conversely, the optimal threshold is higher (i.e., it should be easier to obtain patent protection) when ideas are *scarce*.

The analysis developed by Erkal and Scotchmer provides interesting insights for policy. First, the analysis implies that non-obviousness should be assessed in terms of the R&D costs. Quite counter-intuitively, however, in the model large R&D expenditures are a signal of obviousness. Second, patent offices and the courts should try and get a sense of the scarcity of ideas, and grant patent protection more generously when ideas are scarce.

One problem with these policy conclusions is that both R&D expenditures and the scarcity of ideas are difficult to observe, whereas policy should depend on verifiable variables.⁸ In addition, it must be said that the optimal policy in the Erkal and Scotchmer model could be implemented in various equivalent

⁸ Erkal and Scotchmer note that since in the model the arrival of ideas follows a Poisson process, the waiting time for the innovation is a proxy for the scarcity of ideas. However, the waiting time is also difficult to verify, since patent offices and the courts often do not know when the search for the innovation started.

ways. For example, the policymaker could use optimally chosen monetary prizes, or it could set an optimal level of patent protection (with no non-obviousness requirement at all), instead of using the non-obviousness requirement as a policy tool. A fully satisfactory explanation of the economic function of the non-obviousness requirement should not rest on artificial restrictions on the set of available policy tools. This will be a recurring theme of this survey.

4. The sequential innovation approach

Although today's innovation may to some extent pre-empt future innovative activity, as stressed by the option value approach, certainly it also stimulates certain subsequent technological improvements. The notion of cumulative research, which is familiar to inventors and scholars from all fields, has indeed been the focus of a large economics literature on sequential innovation. This literature has highlighted the fact that each innovation builds on the previous ones and in turn is the basis for subsequent improvements. In a market economy, however, the arrival of a new innovation exerts a negative externality on past innovators, whose profits are eroded by the competition from newer and superior technologies. As argued in Section 2, such an externality can explain why certain innovations may not deserve patent protection.

There are both analogies and the differences with the option value approach. In the option value approach, the negative externality is that the option of waiting for a better idea is lost; in other words, subsequent innovation is impeded. In the sequential innovation approach, by contrast, the negative externality of patenting is that a precedent innovation is discouraged. In both cases, however, there is a trade-off between present and future innovation, and the maximization of the pace of innovative activity calls for certain innovations not to be targeted. One way to achieve this goal is to make them not patentable.

To put it differently, when innovation is sequential innovators need *forward* patent protection in addition to *backward* protection, because otherwise future innovators could compete away the original innovators' profits; and because the first innovators should be rewarded for opening the way to the subsequent improvements (Scotchmer, 1991). Many authors, including Scotchmer and Green (1990), van Dijk (1995), Scotchmer (1996), O'Donoghue (1998), Denicolò (200) and Hunt (2004), have developed the analysis of the optimal non-obviousness requirement along these lines. Here I shall briefly review the model developed by Hunt (2004).

In Hunt's model, there is an infinitely long sequence of successive innovations. Innovative activity is sequential in that only after an innovation is achieved can the race for the next one begin. The timing of innovation is uncertain, being a probabilistic function of the amount invested in R&D. Each innovation is protected by a patent but when the next innovation arrives, the past incumbent's profits are driven to zero and the latest innovator becomes the new incumbent. The size of innovations is also uncertain. When they

invest in R&D, research firms do not know how large the innovation they target will turn out to be, and hence whether it will be patentable or not, nor its per-period level of profitability.

In this world of uncertainty, the incentive to invest depends on the *expected* profits of a successful innovator. In addition to the probability that the innovation is patentable and the expected size of the innovation, expected profits also depend on the expected duration of the innovator's monopoly. A stronger patentability requirement reduces the probability that the current innovation is patentable, but it also reduces the probability that the next generation's innovation will be patentable. Thus, it prolongs the expected duration of monopoly if the innovation happens to be patentable, meaning that a stronger patentability requirement has two opposite effects on the incentive to innovate.

To see why the optimal patentability requirement is generally positive, suppose that the only objective of the policy-maker is to maximize the pace of technological progress. (If the policy-maker is also concerned with static monopoly deadweight losses, the optimal non-obviousness requirement is likely to be even stricter). Starting from zero, let us consider the effect of increasing the patentability threshold. Such a policy move, which makes very small innovations unpatentable, has a second order negative effect on the innovator's prospective profits, since small innovations earn little profits even if they are patentable. The positive effect on the duration of monopoly, however, is first order, since it does not depend on the size of unpatentable innovations but only on their frequency. Thus, starting from zero a small increase in the patentability threshold unambiguously improves the pace of technological progress.

Hunt then proceeds to develop a comparative statics analysis of the optimal patentability requirement. The main result is that the minimal inventive step should be higher in highly innovative industries. The intuition is that in such industries the process of profit erosion is especially rapid and so there is more need of forward patent protection.

One problem posed by this literature is that the patent system can in fact provide forward protection in two ways: the non-obviousness requirement and *leading breadth*. That is, even patentable improvements may constitute infringement on the original patent, depending on the first-generation patent's leading breadth (Green and Scotchmer, 1995; Chang, 1995; O'Donoghue, Scotchmer and Thisse, 1998). This raises various questions. Under what conditions are the non-obviousness requirement and leading breadth perfect substitutes? And, when they are not, how should forward protection be provided: through the non-obviousness requirement, leading breadth, or a combination of the two?

Denicolò and Zanchettin (2002) argue that the non-obviousness requirement and leading breadth protect innovators in different ways. The fact that an innovation must satisfy non-obviousness requirements may block or impede subsequent improvements, thus lengthening the duration of the innovator's monopoly. On the other hand, by finding that an improvement infringes on a basic patent, the courts can force the patent holders to bargain over profit shares, thus allowing the original innovator to capture some

of the rents from the improvement. This means that the non-obviousness requirement has a “blocking effect,” and leading breadth a “sharing effect.”

Denicolò and Zanchettin further argue that sharing is generally better than blocking, because it does not stifle technical progress, and because the inter-temporal externality from the cumulative nature of innovation is at least partially internalized. As a consequence, leading breadth tends to dominate the novelty requirement as a way of providing forward protection.⁹

This means that existing economic justifications of the non-obviousness requirement as a tool to provide forward patent protection are at best incomplete. In addition to explaining why it may be desirable to provide forward protection, theory must also explain in what circumstances the non-obviousness requirement is the most efficient tool to achieve that goal. This issue is still largely open. The most obvious reason why the non-obviousness requirement could be preferable is that leading breadth creates fragmented intellectual property rights and so involves greater transaction costs (Merges and Nelson, 1990; Heller and Eisenberg, 1998; Shapiro, 2001). However, the inclusion of those transaction costs in formal economic analyses of the non-obviousness requirement still awaits future research.

5. The error-cost approach

As is well known, patents have some breadth, i.e., they confer a monopoly over an entire set of technologies. If, for instance, the technology space could be described as a line, a patent would typically cover an interval, not just a point. The reason is that if the patent-holder’s rights were confined literally to the device or product disclosed in the patent specification, the patent would be practically useless since it could be invented around easily. One of the major issues in the patent design debated is, indeed, what determines the optimal breadth of patent protection.

The mere fact that patents have breadth, however, may create problems: even with the best of the intentions, the ambiguity of human language engenders the risk that patent offices and the courts err in defining the boundaries of the patent-holder’s monopoly. There are two types of errors: denying protection over technologies actually disclosed by the innovator (a type I error), or granting a monopoly over something that is already known, and thus should stay in the public domain (a type II error).¹⁰ The error-cost approach maintains that to reduce the risk of type II errors, patents should be granted only when the

⁹ To be precise, leading breadth also has a blocking effect: if the second innovation is small and infringes, the profit left to the second innovator may be too low for investment to be profitable (R&D costs are sunk when bargaining between the patent holders takes place). This means that leading breadth inevitably prevents some second-generation improvements. In fact, Denicolò and Zanchettin show that while this effect is always at work, the sharing effect operates only if leading breadth is sufficiently great. Consequently, when forward protection is weak, the novelty requirement and leading breadth are perfect substitutes. By contrast, when forward protection is strong, it is preferable to use leading breadth only.

¹⁰ Symmetrically, there is a risk of inventors getting a monopoly over something not yet invented, thereby stifling subsequent innovation.

innovative technology lies sufficiently far away from the technological frontier, beyond some minimal inventive step.

Figure 5 illustrates this argument. Given that a patent must have some breadth, and in particular some lagging breadth, but such breadth cannot be controlled perfectly by patent offices and the courts, the only way to avoid type II errors is to deny patent protection to innovations that do not constitute a sufficiently significant improvement over the prior art. This policy may be desirable even if it inevitably increases the risk of type I errors.

[insert Figure 5 around here]

In the error-cost approach, the externality imposed by a patent is the risk of creating deadweight losses for which there is no *quid pro quo*. Clearly, the closer the disclosed innovative technology is to the technological frontier, the greater the risk. It follows that only sufficiently large innovations should be patentable.

Although the error-cost approach seems a natural candidate explanation for the non-obviousness requirement, to the best of my knowledge there has been no formal economic analysis along these lines so far. Many authors have noted that once a patentability threshold is somehow defined, errors may be made in the practical application of the threshold. While this is certainly true, it does not constitute a justification for the threshold. In a sense, the error-cost approach maintains that the non-obviousness requirement is needed to minimize errors in the application of the *novelty* requirement. The implementation of the non-obviousness requirement will in turn be subject to errors, but this is another story.

Needless to say, a fully worked out theory of the non-obviousness requirement based on the error-cost approach would have to address several open issues. What exactly determines the probability of errors? Why do not policymakers simply adjust the patent's lagging breadth when the innovation is small? How large are the social costs of type II errors, and how do they compare to the costs of forestalling small innovations? It is difficult to assess the potential contribution of the error-cost approach to our understanding of the non-obviousness requirement until these and other questions are answered.

6. The complementary innovation approach

The increasing complexity of modern technology and the surge in patenting have resulted in a proliferation and fragmentation of intellectual property rights. In many innovative industries, such as telecommunications, software, and biotechnology, production of new products often requires dozens, if not hundreds, of complementary innovative components, each of which may be protected by one or more patents. Since Merges and Nelson (1990), Heller and Heisenberg (1998), and Shapiro (2001), economists

have recognized that the fragmentation of patent rights increases transaction costs and may lead to pricing inefficiencies. This provides another potential source of externalities that an innovator may impose on others, and hence another possible motive for denying patentability to certain innovations.

Various models of complementary innovations have been proposed in the economics literature. Some posit specialist research firms, which are capable of achieving certain innovations but not others (e.g. Shapiro, 2007; Denicolò, 2007; Denicolò and Halmenschlager, 2008). Others assume that research firms are generalist and so in principle every single firm could develop all the components needed to operate the new technology (Scotchmer and Green, 1990; Ménière, 2005; Bessen and Maskin, 2008; Gilbert and Katz, 2007), although certain firms may have a comparative advantage in the search for certain components (Fershtman and Marcovich, 2006). In some models firms conduct the research sequentially, targeting one innovative component after another, while in others they invest simultaneously in all innovations.

These models differ also in the policy issues they address. Some ask very broad questions, such as whether patent protection is at all desirable. Only a few focus on the non-obviousness requirement. Specifically, Scotchmer and Green (1990) and Ménière (2005) ask whether each innovative component should be separately patentable or patent protection should be reserved for innovators that succeed in developing all (or most) the components needed to operate a new, complex technology. Denicolò and Halmenschlager (2008), by contrast, ask how strong the non-obviousness requirement should be for each separate innovative component.

Ménière (2005) builds on Scotchmer and Green (1990), differing mainly in the assumption of simultaneous rather than sequential R&D investments. Having already dealt with the sequential innovation approach in Section 4, here I focus on the Ménière paper. In his model, a strong non-obviousness requirement fosters innovation by reducing the coordination and transaction costs created by the fragmentation of intellectual property rights. However, it also reduces the incentive to disclose small innovations, thereby encouraging duplicative R&D. The optimal resolution to this trade-off requires that only those innovations that are sufficiently difficult to achieve should be separately patentable. If the probability of success is large, the innovative technology should instead be patented as a whole (that is, only if a single innovator developed all the innovative components). A crucial assumption in Ménière's approach is that R&D investment is a zero-one variable, so that firms can only choose whether to invest or not. It would be interesting to extend the analysis to the case in which the probability of success is determined endogenously as a function of R&D investment.

Denicolò and Halmenschlager (2008) distinguish between two patent policy tools: the probability that an innovative component is granted patent protection (the non-obviousness requirement) and the length and breadth of protection if a patent is granted (the strength of patent protection). These two policy tools jointly determine the prospective reward to innovators, and hence the incentive to innovate. But the patentability requirements also determine the probability that the new technology reads on several

patents or only one. Thus, taking the patentability requirements as a separate tool, they can explicitly address the role of policy in determining the level of fragmentation of intellectual property rights.

Their major finding is that with complementary innovations the non-obviousness requirement should be interpreted more strictly than in the benchmark case of stand-alone innovation; that is, it should be more difficult to obtain patent protection. However, if a patent is granted the strength of protection should be greater. The optimal expected prospective reward to innovators may be greater or lower than in the stand-alone innovation case, depending on the magnitude of the inefficiencies created by the fragmentation of intellectual property rights.

Denicolò and Halmenschlager use the probability that an innovative component is patentable as a policy variable. However, in their model all innovative components are symmetric and so they are effectively indistinguishable. Thus, the model is silent on the crucial policy issue of which component should be granted patent protection and which should not. Once again, economic analysis may offer some insights on why it can be desirable to have a non-obviousness requirement for patentability in place, but provides little guidance on how such a requirement should be implemented.

7. Conclusion

Forty years ago economists used to describe patent policy in terms of a single variable: patent life. A lot of progress has since been made, and today there is a burgeoning literature dealing with “institutional details”, like the non-obviousness requirement for patentability. In spite of this increasing interest, however, economic analysis provides little practical guidelines for policy in this area. All that legal scholars and practitioners can take away from the economics literature is a clarification of the economic function served by the non-obviousness requirement. While this is crucial to a full understanding of non-obviousness, much work remains to be done before economic analysis can really succour in the assessment of patentability.

Even at the abstract theoretical level, there are several unresolved issues. One weakness of various approaches reviewed above is that they artificially restrict the set of available policy tools in order to highlight the role of non-obviousness. When all conceivable tools, or even only those that are currently available in practice, are taken into account it becomes much more difficult to justify the existence of a non-obviousness requirement. This problem is particularly acute in the option-value and the sequential-innovation approaches. Another weakness is that much of the economics literature focuses on variables that are difficult to observe, whereas policy must depend verifiable variables. Addressing these weaknesses is an urgent task for future research.

Having said all this, I would like to conclude with a more optimistic note. There is, indeed, at least one way in which existing economic analyses of the non-obviousness requirement may already have a

direct impact on policy. Many scholars and practitioners are concerned that in the last decades patenting activity has proliferated to dangerous levels, especially in the US. On the other hand, policy-makers seem to believe that there cannot be too much of a good thing – innovation. The contribution of economic theory to this ongoing debate is to show that even if society’s goal is to maximize the pace of technological progress, in our second-best world not all innovations (and, *a fortiori*, not all patents) should be welcome. This is an important message for policy, which might reduce the resistance to certain desirable policy reforms.

References

- Berkowitz, M. K. and Y. Kotowitz (1982). 'Patent policy in an open economy', *Canadian Journal of Economics* 15, 1–17.
- Boldrin, M. and D. Levine (2005). 'IP and Market Size,' mimeo, Department of Economics, UCLA.
- Bessen, J. and E. Maskin (2008). 'Sequential innovations, patents, and imitation', *RAND Journal of Economics*, forthcoming.
- Chang, H. (1995). 'Patent scope, antitrust policy, and cumulative innovation', *RAND Journal of Economics* 26, 34–57.
- De Brock, L.M. (1985). 'Market structure, innovations and optimal patent life', *Journal of Law and Economics* 28, 223–244.
- Denicolò, V. (1999), 'The optimal life of a patent when the timing of innovations is stochastic', *International Journal of Industrial Organization* 17, 827-846.
- Denicolò, V. (2000), 'Two-stage patent races and patent policy', *RAND Journal of Economics* 31, 488–501.
- Denicolò, V. (2007), 'Do patents overcompensate innovators?', *Economic Policy* 25, 679-729.
- Denicolò, V. and C. Halmenschlager (2008), 'Complementary innovations, fragmented property rights, and optimal patent policy', mimeo.
- Denicolò, V. and P. Zanchettin (2002), 'How should forward patent protection be provided?', *International Journal of Industrial Organization* 20, 801-827.
- Erkal, N. and S. Scotchmer (2007), 'Scarcity of ideas and option to invest in R&D', mimeo, University of Berkeley.
- Fershtman, C. and S. Marcovich (2006), 'Patents, licensing and imitation in an asymmetric dynamic R&D race', CEPR Discussion Paper No. 5481.
- Gilbert, R. and M. Katz (2007). 'Efficient Division of Profits from Complementary Innovations', University of Berkeley, Competition Policy Center, Paper CPC07-076.
- Green, J. and S. Scotchmer (1995), 'On the division of profit in sequential innovation', *RAND Journal of Economics* 26, 20–33.
- Heller, M.A. and R.S. Eisenberg (1998), 'Can patents deter innovation? The anticommunist in biomedical research', *Science* 280, 698–701.
- Hunt, R. (2004), 'Patentability, industry structure, and innovation', *Journal of Industrial Economics* 52, 401-425.
- Ménière, Y. (2005), 'Non-obviousness and complementary innovations', mimeo.

- Merges, R. P. and R. Nelson (1990), 'On the Complex Economics of Patent Scope', ' *Columbia Law Review* 90, 839–916.
- Nordhaus, W. (1969), *Invention, Growth and Welfare*, Cambridge, Mass., MIT Press.
- O'Donoghue, T. (1998), 'A patentability requirement for sequential innovations', *RAND Journal of Economics* 29, 654–679.
- O'Donoghue, T., Scotchmer S. and J. F. Thisse (1998), 'Patent breadth, patent length, and the pace of technological progress', *Journal of Economics and Management Strategy* 7, 1–32.
- Scotchmer, S. and J. Green (1990), 'Novelty and disclosure in patent law', *RAND Journal of Economics* 21, 131-146.
- Scotchmer, S. (1991), 'Standing on the shoulders of giants: cumulative research and the patent law', *Journal of Economic Perspectives* 5, 29–41.
- Scotchmer, S. (1998), 'Protecting early innovators: should second-generations products be patentable?', *RAND Journal of Economics* 27, 322–331.
- Shapiro, C. (2001). 'Navigating the patent thicket: Cross licenses, patent pools, and standard-setting', in *Innovation Policy and the Economy*, edited by A. Jaffe, J. Lerner and S. Scott, National Bureau of Economics.
- Shapiro, C. (2007), 'Patent reform: Aligning reward and contribution', in *Innovation Policy and the Economy*, edited by A. Jaffe, J. Lerner and S. Scott, National Bureau of Economics.
- van Dijk, T. (1995), 'Patent height and competition in product improvements', *Journal of Industrial Economics* 44, 151–167.

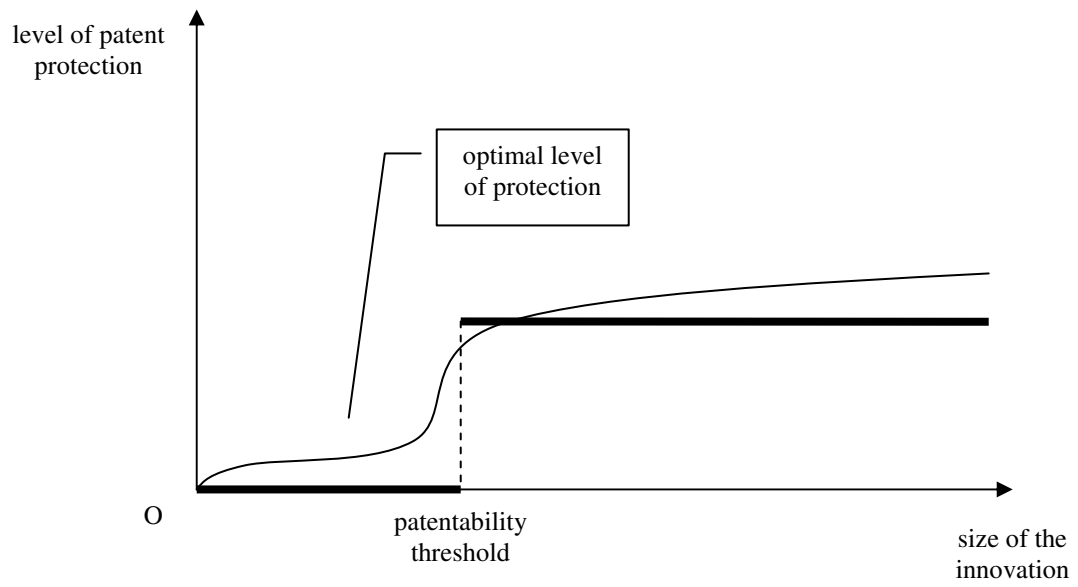


Figure 1

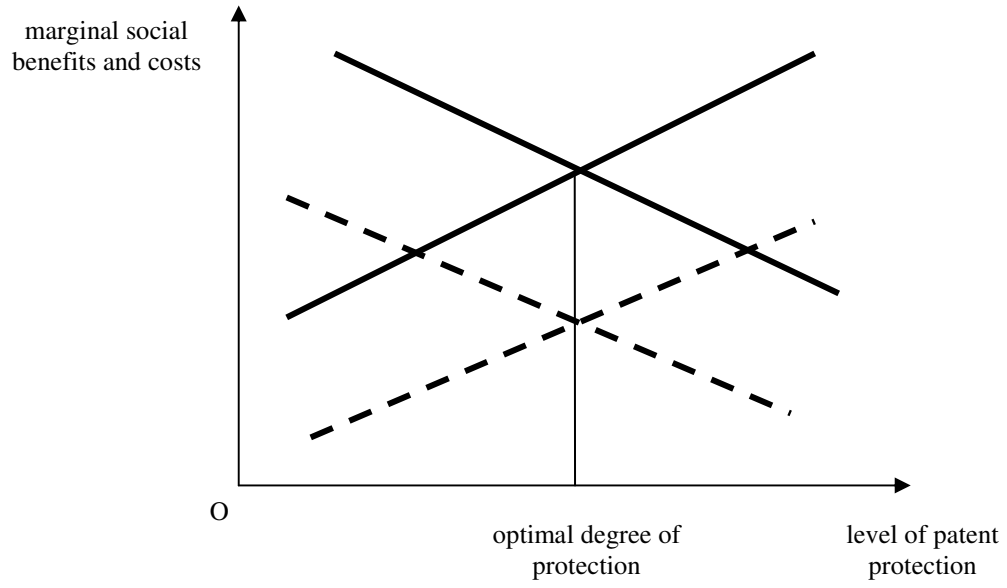


Figure 2

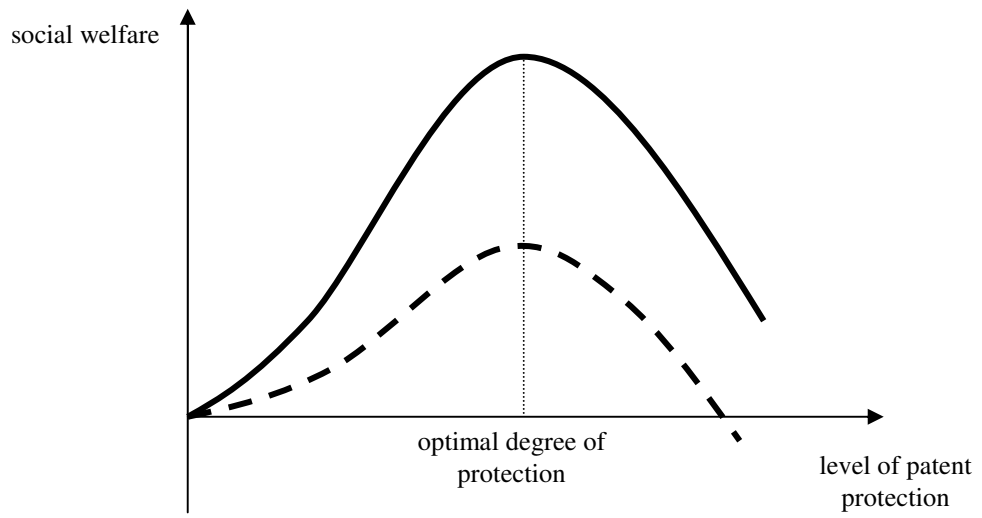


Figure 3

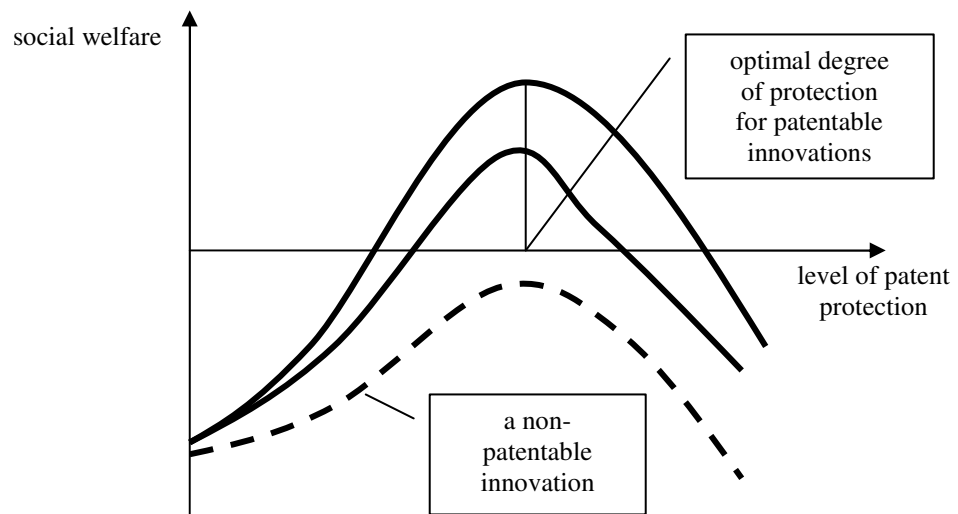


Figure 4

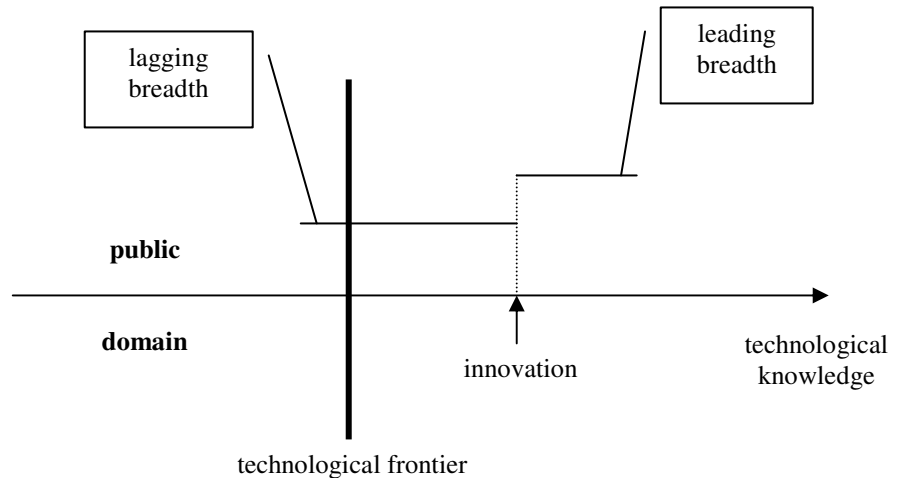


Figure 5