Module 106

[The following is a transcript of the recorded lecture for Module 106 of the PatentX course. The recording of the lecture itself is available through https://ipxcourses.org. Stripped of the accompanying slides and other visual materials, the transcript will likely be hard to follow. It is not intended to be a free-standing document. Rather, its purpose is to assist students, who have already watched the lecture, when reviewing the material.]

A. Forbidden Conduct

Hello. I'm Terry Fisher.

This is the sixth in a series of recorded lectures on Patent Law. It addresses the kinds of activities that will give rise to liability for patent infringement.

The lecture is divided into two parts. In the first, I'll describe the kinds of activities that, unless licensed, would constitute infringement of a US utility patent.

In the second, I'll describe the defenses that sometimes permit potential defendants to engage in these various activities with impunity.

The key provision in the US patent statute is the same one we considered in the last lecture – namely, section 271(a), shown on your screen.

In the last lecture, I discussed the rules governing claim construction, which determine the scope of the "patented invention," and the rules that determine the term of a patent.

We now turn to the question of what you permitted to do to or with a patented invention during its term. As you can see, a wide variety of activities are forbidden.

Making
Using,
Selling,
Offering to sell,
And importing.

All of these prohibitions have an important geographic restriction: they only pertain to activities within or into the United States.

There's one final, less obvious aspect of this provision that bears emphasis. Unlike the copyright statute, which gives a copyright owner the exclusive rights to engage in particular activities with respect to the copyrighted work, the patent statute does not give the patentee any affirmative rights; rather it prohibits particular forms of conduct.

The reason this matters is that those prohibitions apply to everyone, including the patentee.

The inclusion of the patentee in the set of people is neutralized, in the overwhelming majority of circumstances, by the phase, "without authority." Typically, the patentee is permitted to engage in all these activities, because she gives herself permission. But in one important context, she cannot. That's where a product or process is subject to two or more patents, held by different people. I mentioned this doctrine once before in these lectures. Here's a bit more detail.

To illustrate, suppose that I invent an improvement to Mr. Gatewood's patented mousetrap. For example, I add Velcro strips to the inside of the lips of the bag, so that the bag stays closed once the string is pulled, and the operator never has to see the trapped mouse.

I apply for and obtain a patent on that improvement. (Assume, for the purpose of this illustration, that my improvement is both novel and nonobvious.). Am I permitted to manufacture and sell the improved version of the mousetrap? Not without Gatewood's permission, because to do so would violate his patent. Is he permitted to manufacture and sell the improved version? Not without my permission, because to do so would violate my patent.

In sum, neither of us can practice the improved invention without a license from the other. This is commonly known as a situation of "blocking patents." Ordinarily, the stalemate is resolved by contract. I'll get a license from Gatewood, or he'll get a license from me. In either case, the license will include a financial term that ensures that we split, in some ways, the potential income stream associated with the improvement mousetrap.

As I indicated when this topic came up before, Prof. Lemley has argued persuasively that, at least if transaction costs do not frustrate deals between the holders of patents on basic inventions and the holders of patents on improvements to those inventions, then this doctrinal structure will create a more economically efficient pattern of incentives for innovation than the doctrinal structure of copyright law. But the "if" is important. Keep it in mind when you encounter examples of improvements in the remainder of this course.

We turn now to the question of what sorts of behavior are encompassed by these terms. For this next portion of the lecture, I'll be using the map, rather than slides.

The meaning of "make" is pretty straightforward in most circumstances, but there are a few edge cases. Here's one: An invention covered by a so-called combination patent – in other words, one that requires combining two or more components – is not "made" until the components are actually brought together into one operable assembly.

That leads naturally to the question: What if a defendant manufactures in the US all of the components but does not assemble them in the US. Rather, the defendant ships them to customers outside the country, with instructions concerning how to put them together? In the Deepsouth Packing case, the Supreme Court took quite seriously the principle that patent law

stops at the water's edge – and held that this behavior did not constitute infringement. But that holding was subsequently overturned by Congress – in section 271(f)(1) of the current statute.

An analogous issue sometimes arises with respect to process patents. Suppose that the process in question includes three steps: A, B, and C. If a single entity performs all three, it has plainly engaged in infringement. But what if steps A and B are performed by one entity, and step C is performed by another. The current rule is that the patent is infringed only if one of the parties controls or directs the other – or they form a "joint enterprise." Those phrases are defined quite broadly, but not so broadly as to encompass truly autonomous entities.

As Rachel Sachs has shown, this seemingly technical rule has the effect of increasing opportunities for circumvention of patents on diagnostic procedures.

Here's another borderline issue: Suppose I buy a product that is covered by a product patent from an authorized manufacturer or retailer. I use the product heavily and it starts to wear out. So I fix it up. Have I "made" the product – and thus violated section 271? The answer depends upon how much fixing up was involved. The rule is that "repairing" a product is not "making" on the ground that a tacit term of the original sale is a license to repair, but that "reconstructing" the product does constitute "making." So what's the difference? Roughly speaking, the more extensive the rebuild, the more likely it is to be considered reconstruction, but that's an awfully rough guideline.

Plainly, the notion that a tacit term of the original sale is a license to repair but not to rebuild is no help at all in sharpening up the distinction. If you encounter a case of this sort, your best bet is to poke around in the case law and try to ascertain whether the magnitude of the rebuild with which you are concerned more closely resembles the cases in which the courts have found activities to be repairs or the cases in which they have found them to be reconstructions.

OK, that's "make". What about "use"? As you might imagine, if you buy a product covered by a patent, and the product has been lawfully made by the patentee or his licensee, then you can use the product without violating the patent statute. But what if the product was not lawfully made – in other words, it was created by an infringer? In that case, using it would violate the statute. The effect of this rule is to extend liability beyond the primary infringer (namely the manufacturer) down the chain to include people who buy from him – in a fashion analogous to the criminalization of the receipt of stolen property.

Over the years, the courts have considered a variety of edge cases involving activities that might or might not constitute "use" of a infringing product and thus trigger this rule. A few of them are noted on your screen. So for example, merely possessing the product does not constitute using it, but testing it does constitute use. Importantly, employing the product for a purpose other than the one contemplated by patentee does constitute infringement.

What constitutes "selling" a product is usually obvious. The main thing to keep in mind is that the patentee's ability to control sales is qualified by the exhaustion doctrine, which we will consider in the third segment of this lecture.

What constitutes an offer to sell is also usually straightforward. The principal nuance here is that nowadays, if the offer is made in the US, the defendant cannot escape liability by pointing out that the sale itself would be consummated outside the US.

Importing things covered by US product patents clearly constitutes infringement. A less obvious but very important variation on this theme is that importing things made overseas using a process that is covered by a US process patent is also infringement – unless the product is subsequently materially changed or is a trivial component of a larger product.

Like the prohibition on "sales," the prohibition on importation is qualified by the exhaustion doctrine, which we will get to in due course.

OK, that concludes our survey of the kinds of activities in which one is not permitted to engage without permission from the patentee.

Up to this point, I've been speaking casually about products that are "covered" by patents. We now have to sharpen up that concept.

You will recall, from the previous lecture, that the dominant way of defining a patented invention in the US is the peripheral claiming method, in which the ambit of the patentee's rights is set by the claims, which the patentee is primarily responsible for drafting.

That means that, for the most part, whether a particular product (or process, for that matter) is covered by a patent is determined by whether it is covered by one or more of the claims in that patent.

That inquiry is governed by what is called the "element by element" approach. As its name suggests, to establish infringement the patentee must demonstrate that every one of the elements of the claim at issue are found in the defendant's product or process.

Here's a stylized illustration

Suppose that a claim has 5 elements.

An accused device infringes the claim only if It contains all 5 of the elements.

If one of the 5 features is different, then there is no literal infringement. (In just a minute, we'll consider the possibility of infringement under the doctrine of equivalents, but put that to one side for the times being.)

The same is true if one feature is not different, but missing altogether.

The method for determining whether each of the elements found in the claim are also found in the accused device Is the same as the method used, when applying the novelty doctrine, for determining whether all of the elements are contained in a single reference in the prior art.

Here's how these concepts would be applied in a simple, hypothetical case: Here, once again, is Gatewood's patented mousetrap.

Suppose that a defendant manufactured and sold mousetraps that, instead of a window, had a screen through which the trapped mouse could be viewed.

As you will recall, the principal claim in Gatewood's patent Had five elements – each designated here by the capital letters I have inserted.

To determine whether the defendant is infringing, we line up his device with those five elements

It definitely has A, the housing
And B, the trapping means – in other words, the glue
And C, the handle
And D, the opening in the top face
But not E. A screen is not a translucent window.
The result is that there is no literal infringement in this case.

Now, you may be thinking: But that's not fair. The defendant's trap is really the same as Gatewood's. It just uses a slightly different material for one component.

One response is that the misalignment that prevents a finding of infringement is Gatewood's own fault – or that of his lawyer. He could have written the claim more broadly. For example, instead of "a translucent window adhesively attached to the periphery of said viewing opening in said top face," he could have said, "a translucent or transparent material such as a window attached to the periphery of said viewing opening in said top face." Having opted for the narrow version, he should not complain if the defendant evades it. After all, a landowner cannot complain if someone walks just outside the edge of his property. So a patentee should not be able to complain if someone walks just outside the edge of his claim.

For better or worse, US patent law does not take such a harsh position. It contains a doctrine that sometimes allows a patentee to prevail in circumstances of this general sort. That doctrine is known as equivalents. Here's how it works.

To repeat, If a claim has 5 elements, And the accused device has the same elements, There's literal infringement.

If one or more of the elements does not show up in the accused device, but the accused device has similar analogues of each of those elements,

Then there is no literal infringement, but there may be liability under the doctrine of equivalents.

Well, when exactly?

The courts have articulated various tests for determining the necessary degree of similarity. One commonly invoked test is that each relevant component in the defendant's device must Perform the same function in substantially same way to obtain same result. This is known as the function-way-result test.

A vaguer guideline is that the differences between the relevant element in the claim and the corresponding component of the defendant's device must be "insubstantial."

Applying this analysis to our hypothetical problem,

As we saw, Four of the elements in Gatewood's claim map onto the defendant's trap

Because the fifth one does not, there's no literal infringement.

We would then ask, does a piece of screen perform the same function in substantially same way to obtain same result as a translucent window? Maybe not if used in a house – because one lets air through, while the other does not.

But in the context of a mousetrap, probably yes. Both, it seems, are designed to allow the user to ascertain that a mouse has been caught, but also partly to obscure the carcass.

The upshot is that the defendant is probably liable under the doctrine of equivalents.

Here are a couple of important nuances to be kept in mind when applying this doctrine.

Does the patentee have to show that the defendant behaved unscrupulously – i.e., knew of the patent and deliberately tried to skirt it? Once upon a time, maybe yes, but not anymore.

It helps the patentee if he can show that, at the time of the alleged infringement, a PHOSITA knew of the interchangeability of the technology described in the claim and the technology used by the defendant. It is not necessary, however, that the defendant's technology even have been in existence when the patent was granted. This last proposition is controversial, but there it is.

There are three important limitations on the doctrine of equivalents.. First, the patentee cannot use it to claim territory that he could not have claimed during patent prosecution because it

would have encroached on the prior art. A good illustration of this principle can be found in the Wilson case, which involved a patent on a dimpled golf balls.

Apparently, it's been known for a long time that smooth golf balls do not fly as far as dimpled golf balls – for reasons summarized in this diagram

Incidentally, if you are a golfer or a physicist, you would likely find interesting the 5-minute video the address of which appears at the top of your screen, which discusses the design of golf balls. If you are neither, you might still find the video amusing.

Although it's well established for this reason that dimpled balls are better than smooth balls, the optimal pattern of dimples remains uncertain – and manufacturers of golf balls are constantly trying out new patterns.

A particularly successful pattern was developed by the Wilson company, which used the pattern in its prostaff ball – and got a patent on it.

Here's the key claim.

Now, if you, like me, did not know what a icosahedron is, it's a polyhedron with 20 faces.

Here's one of the drawings that illustrate this claim

The Dunlop company, impressed by the sales of Wilson's balls, began selling a ball with a similar, but not identical pattern of dimples.

Wilson brought an infringement suit. But Wilson had a problem.

One of the elements of its claim was that its ball had six great circle paths that did not intersect any of the dimples in the pattern.

But in the Dunlop ball, the great circles did intersect some of the dimples.
That means that Wilson could not prevail on a theory of literal infringement.

So, in the diagram on the screen, the green zone represents the ambit of Wilson's claim. The Dunlop ball falls outside that zone.

Wilson's response was to try to expand the zone using the doctrine of equivalents. The effect of the doctrine of equivalents, when successfully invoked, is to create a buffer or penumbra, extending the reach of the patent to cover products or processes not covered by the literal terms of the claims.

Wilson might have succeeded in this case, but for one issue. The Federal Circuit ruled that, in order to reach the Dunlop ball, Wilson would need enough additional territory to be the functional equivalent of the hypothetical claim now set forth on your screen.

Unfortunately for Wilson, such a claim, had it been asserted by Wilson during patent prosecution, would have brought Wilson's patent fatally close to yet another dimpled ball, which had been previously developed – and patented – by Uniroyal.

Now, to be sure, there would still be a difference between the dimpling claimed in the Uniroyal patent and the dimpling claimed by Wilson using the doctrine of equivalents.

However, the court ruled, "these differences are so slight and relatively minor that the hypothetical claim--which permits twice as many intersecting dimples, but with slightly smaller intersections--viewed as a whole would have been obvious in view of the Uniroyal ball."
The Court's finding on this point could be depicted graphically as follows.

Surrounding uniroyal's actual claim is a buffer of obviousness. In other words, while an application for a patent filed after uniroyal that sought protection for a technology falling into the red zone would be rejected as not novel, an application for a patent that sought protection for a technology falling into the orange zone would be rejected as obvious.

For that reason, Wilson could not have gotten the hypothetical claim during patent prosecution.
And – here's the crucial point – Wilson cannot now use the doctrine of equivalents to claim territory that it would have been unable to get during prosecution.

So Dunlop prevailed.

The second of the three limits on the doctrine of equivalents is that a patentee cannot use the doctrine to claim material disclosed in the patent application – typically in the specification – but not claimed. One justification for that limitation is that purpose of the equivalents doctrine is to protect P against evasion through use of unanticipatable substitutes; it should not apply when P was aware of substitutes. In my view, that rationale is not particularly persuasive; the doctrine has a more general function that the rationale presumes. But that's where things stand.

The third of the limits on the doctrine of equivalents is the doctrine of prosecution history estoppel – sometimes known by the acronym PHE. The essence of this doctrine is that a patentee may not use the doctrine of equivalents to recover territory that he originally sought in the early stages of patent prosecution, but that he surrendered during prosecution, typically in response to a rejection by the examiner.

There are some subtleties lurking in that summary of the essence of PHE, which can be best understood through use of yet another diagram.

Suppose that the zones encompassed by the prior art could be represented graphically by the red zones.

We know from the golf-ball case that, surrounding the red zones are orange zones of obviousness.

Against this backdrop, suppose that, in his initial application, the applicant included the broad claim represented by this dark grey zone.

The examiner rejects the claim on the ground that some of the ground it covers is obvious in light of the prior art – ie., runs afoul of section 103.

Suppose that, in addition, the examiner also bases the rejection on the ground that some of the language in the claim is too vague and thus runs afoul the "claim definiteness" requirement in section 112(b)

In response, the applicant adds limitations to the claim, which have the effect to cutting back on the territory he originally requested.
The examiner allows the modified claim, and the patent issues with it.
The scope of the modified claim is indicated by shrunken grey zone in this diagram.
Now suppose that the patentee wants to bring an infringement suit against a competitor whose product comes close to the claim, but falls just outside it.

As we saw in the golf ball case, the patentee's strategy is likely to be that, even though the defendant's product did not fall within the claim read literally, it fell within the buffer zone surrounding the claim created by the doctrine of equivalents.

Prosecution history estoppel blocks the patentee from obtaining in this fashion the yellow zones – because those were zones he originally claimed, but then gave up.

Just to complete the story, you know from the golf ball case that the patentee also cannot use the doctrine of equivalents to grab the purple zones. Unlike the yellow zones, the purple zones are not areas the patentee claimed and then abandoned during prosecution. But allowing him to take them now would permit him to acquire territory that he could not have gained during prosecution because they would have been deemed obvious.

Nor is he permitted to claim territory that he revealed in his application but never claimed.

The net result is that the effective scope of his claim, augmented by permitted applications of the doctrine of equivalents, can be represented by this oddly shaped diagram.

OK. This diagram enables one to see one of the nuances I mentioned a few minutes ago

Recall that this was this was part of his original, ambitious claim.

It was rejected because it encroached in the prior art.

The applicant avoided the rejection by adding a limitation that had the effect of giving up this territory.

But notice that he gave up more territory than he had to.

Does estoppel deprive him now of the ability to recover any of this zone?
The Federal Circuit, in the Festo case, said yes, adopting what it described as the "complete bar" rule.

The Supreme Court rejected that position, in favor of what it called the "forseeable bar" rule.

In a bit more detail, here's the holding of the case…

To be honest, I doubt that the decision makes all that much difference in practice. The set of cases in which the foreseeable bar rule would allow a patentee a more generous zone of equivalents than the complete bar rule is probably small.

But the decision also contains some very interesting reflections on the nature of patents that have much broader implications. They are set forth on your screen. As you can see, the Court is remarkably pessimistic about the ability of inventors to capture the essence of their inventions in the language of the claims – and suggests that courts, using the doctrine of equivalents, can and should give effect to what they really invented. Language of this sort, I hope you see, exemplifies the central claiming approach, not the peripheral claiming approach. The Festo decision makes clear that, in the eyes of the Supreme Court, equivalents doctrine is a helpful tool in the service of central claiming. Whether the judges of the Federal Circuit share that view is much less clear.

We come finally, to the topic of secondary liability. Up to now, we've been considering ways in which a putative defendant can violate section 271 by himself or itself engaging in patent infringement – so-called direct infringement. But patent law also contains two doctrines that enable a patentee to reach persons who assist others to engage in activities that would constitute infringement.

The first of those doctrines is known as contributory infringement. It's captured in section 271(c) of the US statute, which appears on your screen.

This is a dense provision. Take a moment to read it carefully. Here are a couple of things worth emphasizing.

First, note the reference to the state of mind of the defendant. Direct patent infringement, as you know by now, is a strict liability offense. The intent of the defendant is irrelevant. By

contrast, contributory patent infringement occurs only if the defendant knew that the combination for which he or she was providing a part was both patented and infringing.

Second, note that liability does not arise when the pertinent piece, component, or supply has a substantial noninfringing use. The effect is to immunize against liability people who supply staple items of commerce. For those of you who know US copyright law, this is the origin of the aspect of the Sony case (reaffirmed in Grokster) that incorporates into the fair-use defense the concept of substantial noninfringing uses.

The second of the two forms of secondary liability in patent law is inducement, which finds expression in section 271b. This one is conceptually simpler. If you actively aid someone to engage in infringement, you too are liable.

Some examples appear on your screen: distributing brochures advertising infringing products; providing instruction in performing a patented process; buying articles made without permission through the use of a patented process; indemnifying an infringer; and persuading someone to breach a patent license agreement. This is a nonexhaustive list; other ways of engaging in inducement are readily imaginable.

However, inducement is not established by mere inaction. So, for example, when a parent company fails to prevent a subsidiary from engaging in infringing behavior, the parent is not on the hook.

Like contributory infringement, inducement has a scienter requirement. To prevail on this theory, the patentee must demonstrate that the defendant was both aware of the patent and aware that the actions the defendant helped or encouraged constituted infringement.

An important nuance: Can the defendant avoid liability by showing that he believed in good faith that the patent was invalid? Answer: No.

This concludes my summary of the kinds of activities that are forbidden under US patent law. In the second part of this lecture, I'll examine the defenses that a defendant can assert to avoid liability.

B. Defenses and Limitations

Hello I'm Terry Fisher.

This is the second part of the lecture on patent infringement. In the first part I described the kinds of activities that may be deemed to infringe a patent. In this lecture I'll describe the principal defenses that may be asserted by a person or firm that is engaged in one of these activities.

A full sense of the entitlements enjoyed by a patentee of course requires consideration of both parts. From a policy standpoint, the two parts are commonly said to work together. Specifically one of the functions of the defenses is to prevent patentees' assertions of their rights in ways that would undermine, rather than advance, social welfare as a whole. As I review the various defenses, ask yourself whether, if they are to perform this role, the defenses should be expanded, contracted, or left just the way they are.

The first and most obvious of the defenses is the one that we have considered repeatedly already in this course of lectures – namely, it is routine for a defendant sued for infringement to defend itself on the ground that the patent on which its liability apparently depends is invalid.

As you know by now, the primary way of asserting that defense in the US is through a challenge offered in the District Court where the infringement suit is initially brought -- although it is possible now for a defendant to make the same assertion collaterally by initiating an IPR review before the PTO.

A couple of things to emphasize about this universal strategy. First, it's done on a claim-by-claim basis. In other words, the defendant to escape liability altogether needs to establish that every one of the claims on which the plaintiff relies is invalid.

Next the defendant bears the burden of establishing invalidity by clear and convincing evidence. This standard is a little higher than the preponderance-of-the-evidence standard that governs many issues in civil litigation. As a practical matter, it doesn't make a great deal of difference to courts, but it does somewhat amplify the tendency of juries to give credit to the presumption of validity enjoyed by patents.

An important procedural aspect of assertions of invalidity is whether a patentee who loses a patent infringement suit on the ground that the patent is declared invalid can thereafter bring suit against a second defendant or instead is bound by the declaration of invalidity in the first suit.

Some time ago, the answer was that the patentee could go forward. In other words, the finding of invalidity in the first suit did not bind the patentee in a subsequent infringement suit against a different defendant.

But in 1971, in the Blonder Tongue case, the Supreme Court changed the position on this issue -- holding that a final judgment of invalidity does bind the patentee in a subsequent infringement suit against another defendant, so long as the patentee had a full and fair opportunity to litigate the issue.

This plainly raises the stakes for a patentee in the first suit in a planned chain of suits against several defendants.

The next defense – sovereign immunity – has an interesting and complex history in the US. But, because it does not appear in the law of any other country, I will ignore it for the purposes of this lecture.

We turn now to defenses that hinge upon behavior by the defendant.

The first of these is very rarely invoked nowadays but it's conceptually quite interesting – and, as you'll see, some scholars have made strong arguments that it ought to play more of a role in modern patent law than it currently does.

The doctrine is known as reverse equivalents. It was developed and is best illustrated by the Westinghouse case which was decided by the United States Supreme Court in 1898.

I'm going to describe the facts of the case in some detail, in part because they provide an excellent illustration of sequential innovation, a phenomenon that implicates many doctrines and policies in patent law.

Starting in the 1860s, transportation in the United States became increasingly dependent upon the transcontinental railroad lines. Growing amounts of freight, as well as growing numbers of passengers, were transported on the lines that you can see on your screen.

Unfortunately, the technology used in the braking systems of the trains failed to keep pace with technologies used to provide the engines power. In particular, the brakes did not work well when trying to slow or stop long trains. Among the crashes caused by the inability of brakes to restrain trains coming down hills was a catastrophic 1893 circus train wreck in Pennsylvania.

So, as you might expect, the late 19th c. saw a sustained effort by inventors, many of them employed by the railroads, to improve the braking systems.

A pioneering innovation was the air brake first developed by Geroge Westinghouse Jr., which relied upon compressed air generated in the locomotive to press brake pads against the wheels of the train cars. These early air brakes helped a great deal, but they had some significant disadvantages.

Specifically, it took a while for air pressure generated in the locomotive to travel through the various pipes to the cars. The longer the train, the more serious was the delay. In addition, the farther away from the engine was a particular car, the lower was the pressure of the air reaching it. That, in turn, created a risk that a braking train would fold up upon itself -- much the way a modern 18 wheeled truck sometimes jackknifes.

A final problem associated with the original air brakes is that a break in the chain of couplings between the cars would make it impossible to set the brakes on any of the cars.

A major step in addressing these problems was the development of auxiliary reservoirs. As their name suggests, these consisted of reservoirs where compressed air could be stored in individual train cars, thereby shortening the amount of time it required to exert pressure on the brake pads of the cars.

Westinghouse then came up with a very important refinement of this technology, the heart of which was what it called a triple valve. Its name derived from the fact that it had three positions. In the so-called "release position," air came from the engine through the brake pipe into the valve, passed through the "feed groove," and thus increased the pressure inside the auxiliary reservoir unitl it reached the pressure of the air inside the brake pipe. That prepared the brake for braking. When the conductor wanted to apply the brakes, he would *release" the pressure in the brake pipe. The slide valve in the center of the picture would move to the right, and air would flow from the reservoir into the brake cylinder, which in turn would press the brake block against the wheel. In the third position, known as the lap position, the valve shifts back to the left, isolating the reservoir and the cylinder, holding the brake pad against the wheel.

Clever. As you might imagine, it was patented by Westinghouse. The key claim in the patent is shown on your screen:

"In a brake mechanism, the combination of a main air-pipe, an auxiliary reservoir, a brake-cylinder, and a triple valve having a piston whose preliminary traverse admits air from the auxiliary reservoir to the brake-cylinder, and which by a further traverse admits air directly from the main air-pipe to the brake-cylinder, substantially as set forth"

Although the triple valve system was ingenious, it was not perfect. One of its imperfections is that if, in a long descent, the air pressure in the brake cylinder diminished because of leakage, it was impossible to augment the pressure without releasing the pad. In the early 1880s, George Boyden developed a modified version of the triple valve that made this possible. Schematically, the improvement looks like this. The actual valve looked like this.

Westinghouse brought suit against Boyden, arguing that the improved value infringed the patent on the unimproved triple valve. When the case eventually made its way up to the US Supreme Court, Boyden prevailed.

Here's the key passage in the opinion:

"But even if it be conceded that the Boyden device corresponds with the letter of the Westinghouse claims, that does not settle conclusively the question of infringement. We have repeatedly held that a charge of infringement is sometimes made out, though the letter of the claims be avoided. [That as you know by now is the regular doctrine of equivalence.] The converse is equally true. The patentee may bring the defendant within the letter of his claims, but if the latter has so far changed the principle of the device that the claims of the patent, literally construed, have ceased to represent his actual invention, he is as little subject to be

adjudged an infringer as one who has violated the letter of a statute has to be convicted, when he has done nothing in conflict with its spirit and intent."

That's the doctrine of reverse equivalents. It provides a way in which a defendant who has engaged in conduct that would be described as literal infringement can nevertheless escape liability.

Between 1898, the date of the Westinghouse case, and the establishment of the Federal Circuit in 1982, this doctrine was invoked with reasonable frequency by US courts. Somewhere between two and eight cases per decade resulted in findings of non-infringement on the basis of this doctrine of reverse equivalence.

However, the judges of the Federal Circuit concluded that the doctrine was misguided – and they have effectively overruled it.

Here's a representative case from the Federal Circuit. As you can see from the highlighted passages in the opinion the Federal Circuit declared: "never once has this court affirmed a decision finding an infringement based on the reverse doctrine of equivalents." And it refers to the Westinghouse case as "one anachronistic exception."

As you might imagine, this is highly unusual language for a court of appeals to use when describing a Supreme Court decision.

Some scholars have criticized this effective overruling of Westinghouse. In particular, Samuel Ernst and Dennis Crouch have both written articles that in my view are persuasive criticisms of the Federal Circuit on this issue.

But, for better or worse, the law is that the defense of reverse equivalents has almost no current salience.

We turn now to another defense that used to be more generous in the United States but has shrunken. In many countries, defendants can escape liability by establishing that they made use of patented inventions for experimental purposes. Examples include Germany, the United Kingdom, and to some extent Japan.

In the United States there was once a significant doctrine of this sort. Courts frequently mentioned it in dictum. Here's an early statement by Justice Story:

"it could never have been the intention of the legislature to punish a man, who constructed [an infringing device] merely for philosophical ex- periments, or for the purpose of ascertaining the sufficiency of the machine to produce its described effects."

In the hands of the Federal Circuit, this doctrine also has largely faded from view.

In the previous lecture, when talking about the pharmaceutical industry, I discussed the ruling by the Federal Circuit in the Roche Products case that experimentation on pharmaceutical products for the purpose of launching FDA review of what would become generic versions of those products could not be excused on the basis of an experimental used doctrine.

I then explained how, in that specific context, the Federal Circuit's stance was overturned by the Hatch Waxman Act.

But in all other settings, the skepticism of the Federal Circuit survives. The primary precedent is its 2002 ruling in Madey versus Duke University case. The key passage in the opinion is set forth on your screen.

"So long as the act is in furtherance of the alleged infringer's legitimate business and is not solely for amusement, to satisfy idle curiosity, or for strictly philosophical inquiry, the act does not qualify for the very narrow and strictly limited experimental use defense. Moreover, the profit or non-profit status of the user is not determinative."

This doesn't leave much room for defendants to excuse behavior on the grounds that they were experimenting. There's been a fair amount of criticism of this position but thus far no effort to overturn it legislatively.

The third and last of the defenses that hinge upon conduct by the defendant has evolved in a different direction: over time, it has grown rather than shrunken. This is known as a prior user right.

Once again in other countries there are reasonably generous versions of this defense. In the United States, a less generous but not trivial version has been in effect since 2016.

The details appear on your screen. As a practical matter, this provision enables a defendant who uses a process -- or a product in a process – commercially, in the United States, more than one year prior to the filing of a patent application or disclosure by the a person who eventually receives a patent on the relevant technology, to keep on using the technology without permission indefinitely.

Since 1999 such a prior-user defense has existed with respect to the narrow category of business-method patents. What the new version of 273 a did was generalize that principle to all patent

It's a little early to tell, but the likely impact of this provision will be to make it a little less risky for inventors of new processes to choose trade secret protection instead of patent protection to deprive competitors of access to their inventions. Formerly, the creator of a new process who hid it rather than apply for a patent, ran a risk that someone else would invent the same process and would patent it and then use the patent to force the first innovator to stop using it. The reason this was a hazard is that, by keeping the process secret, the first innovator would

not place the new process into the prior art and therefore would not prevent the acquisition of a patent by the second innovator.

Section 273(a) reduces that risk – although the requirement that the commercial use commence at least a year before filing or disclosure by the second innovator means that the hazard has not been eliminated entirely.

We turn now to a set of defenses that hinge upon conduct by the patentee, rather than the defendant. The first pertains to conduct by the patent during the course of patent prosecution. The basic principle underlying the inequitable-conduct defense is clear enough. Applicants for patents have a duty to disclose to the Patent and Trademark Office all information known to them pertaining to patentability. Breach of this duty renders the entire patent unenforceable.

This used to be a very important defense. But here too the Federal Circuit has reduced its availability. Specifically in two cases in 2009 and in 2011, the Federal Circuit tightened significantly the requirements for establishing inequitable conduct. The current state of the law is that the defendant has to clear three hurdles: first, plead inequitable conduct with particularity. Second prove by clear and convincing evidence that the information in question was material – in other words, that, if the information withheld by the patentee had been revealed to the examiner or the false statements made by the patentee not been made, the examiner would not have allowed the claim in question. And finally, that the patentee intended to deceive the Patent and Trademark Office.

As you might imagine these requirements have made it much harder to establish inequitable conduct. Data gathered by scholars in the wake of the two decisions bear out that prediction. As you can see from this graph, in the three decades before 2009 the frequency with which inequitable conduct came up in district court patent decisions was rapidly growing. Between 2000 and 2009, the percentage of suits in which the defendant asserted inequitable conduct was also rising fast -- up to around 40 percent.

This table from the study by Rantanen & Petheridge shows that, since the Federal Cricuit's rulings in 2009 and 2011, the percentages of cases in which an inequitable conduct is asserted as a defense have been dropping fast.

Finally, in the cases in which inequitable conduct is asserted, the number of successful assertions has been dropping extremely fast. Interestingly, the most common reason for the failure of the defense is the inability of the defendant to prove that the patentee intended to deceive.

The remaining three of the defenses based upon conduct by the patentee all involve conduct after the patent is issued rather than during patent prosecution. They are closely related but distinct.

The first, like inequitable conduct, has shrunken recently. The scopes of the other two have not changed.

Delay involves situations in which the patentee could have brought suit promptly against an infringer, but waited. Section 286 provides that, in such situations, no recovery shall be had for any infringement committed more than six years prior to the filing of the complaint or counterclaim. How exactly that affects the recovery of damages I will discuss in the next lecture in this series. But the operation of 286 is clear enough.

Laches is a traditional equitable doctrine, not dependent upon 286 or any other statute, that bars stale claims. Until quite recently it was employed in patent cases, where it was employed to cut off claims both for damages and for equitable relief. To prevail on this theory, the defendant had to show two things: that the patentee delayed inexcusably and unreasonably and that the defendant thereby suffered material prejudice attributable to the delay. Although a 6-year delay created a presumption of laches, a defendant sometimes was able to invoke the defense when the delay was shorter.

In SCA Hygiene case, the Supreme Court overturned this body of law, ruling that that laches cannot be interposed as a defense against damages where the infringement occurred within the period prescribed by section 286. This leaves the statute as the primary defense based on delay pure and simple.

But what if the patentee not only waits, but also engages in conduct that misleads the defendant into reasonably inferring that the patentee does not intend to enforce the patent. The defendant then relies on that representation, and because of that reliance the defendant would be materially prejudiced if the patentee were allowed to renege.

Traditionally, under such circumstance, the defendant could avoid liability under the doctrine of equitable estoppel. Indeed, the effect was that the entire patent was unenforceable against the defendant thereby prejudiced.

This traditional defense survived the SGA Hygiene decision. The conditions necessary to establish it don't arise very often, but it exists.

Finally, even more rarely, a court will infer from the patentee's behavior an implied license to engage in particular otherwise infringing behavior.

We come finally to the defense of exhaustion. The basic idea here is that when a patentee places an item embodying his or her innovation into commerce, the patentee gives up some or all rights subsequently to control uses of, resales of, or transfers across international boundaries of, that product. The reason why this is called exhaustion is because it is said that the patent use entitlements are exhausted by the first sale.

There are two main settings in which this theory has been invoked. In the domestic context, exhaustion is sometimes known as the first sale doctrine. In the international context, the category of activities to which it is applied is known as parallel importation.

The history of the exhaustion doctrine in the United States is complicated. Traditionally, this defense was a broad. In other words, defendants were able to invoke it often -- in both the domestic situation and the international situation.

With respect to the domestic variant, in a series of cases, the Supreme Court ruled that authorized sale of a patented product to consumers carries with it the rights both to use and to resell that product. Similarly in the international context the traditional rule was that exhaustion applies to sales of goods overseas. In other words, the patentee cannot prevent reimportation of goods that are sold overseas -- unless the patent expressly forbade reimportation when the products were first sold.

In other words, this traditional version of the exhaustion doctrine could be overridden by the patentee by informing the first purchaser that reimportation is forbidden.

That was where things stood between 1996 and 2008.

Yet again, the Federal Circuit constricted the circumstances in which these defenses could be invoked. The primary way in which it did so was by enabling patentees to control more of the activities that purchasers of the products embodying the patented technology wanted to engage in.

Here are some examples: In the domestic context, the Federal Circuit held in an important pair of cases that exhaustion could be overridden by contract or notice. For instance, if the patentee stamped on the outside of a product, a notice saying that it could only be used once, the result was to override the otherwise general privilege of the purchaser to make use of the product repeatedly or to resell it.

In the international context, the Federal Circuit held that, even in the absence of an explicit notice, an overseas sale of a product embodying a U.S. patent did not entitle a purchaser to re-import the product into the United States.

The effect of both of these rulings was to augment the ability of patentees to engage in differential pricing of products embodying their inventions – in other words, to charge different consumers different prices for the same good or service.

For example, in the domestic setting, the ability of a patentee to override exhaustion through contract meant that the patent could charge a customer who wanted to use a device only once, less money than the patentee charged a customer wanted to use it multiple times – and need not worry that purchaser of the first type would resell their products to the purchasers of the second type.

In the international context, patentees could charge less for goods sold outside the United States than it charged for the same goods sold inside the United States – and need not fear that arbitrageurs would buy the goods for low prices in other countries and then re-import them into the United States thereby undercutting the high US prices.

Since 2008 the trend toward enhancement of patentees' rights in this respect has been reversed. In other words, the defense of exhaustion has been expanded.

Several steps have contributed to this expansion. Some of them made were by the Supreme Court in the 2008 Quanta case. Some were made by lower courts trying to figure out what the Supreme Court had in mind in Quanta.

But for our purposes, the intermediate steps are less important than the final outcome: the Lexmark decision, issued by the Supreme Court in 2017. The Court there ruled broadly that sale of a product embodying a patented technology exhausts all of the patentees' rights in that item, regardless of any restrictions the patentee purports to impose and regardless of the location of the sale. And this applies even if the sales were made by a licensee, rather than by the patentee itself.

The net effect has been to constrict the ability of patentees to engaged in differential pricing.

There remains, as you can see from this list, one more set of defenses to patent infringement – namely those rooted in the separate body of antitrust law. Those defenses are sufficiently complex and specialized that I won't examine them here. If you are very curious, the most important of them are taken up in the recorded lecture by Jorge Contreras on Standard Essential Patents, a link to which can be found on the IPXcourses website.

Here's an overview of the ground covered in Part B of this lecture. In the US, a defendant accused of patent infringement might be able to assert a variety of defenses.

The most common is invalidity.

A defense peculiar to the US (and thus of little interest in this course) is sovereign immunity. Then there are three defenses involving conduct by the defendant – reverse equivalents, experimental use, and prior use.

Four pertain to conduct by the patentee – inequitable conduct, delay, equitable estoppel, and implied license.

One of the most important defenses is exhaustion.

Finally, there are three defenses involving anti-competitive behavior that we have not yet discussed.

A highly impressionistic view of the relative importance of these defenses in 1980 looks like this. Very roughly speaking, the height of the columns reflects the frequency with which the defenses were invoked and the economic importance of their availability.

By 2000, their relative status had shifted.

Today, their relative status looks roughly like this.

Juxtaposing the three dates, you get this.

I hope that this lecture has provided you a sense both of the activities that, unless authorized, constitute infringement of a patent -- and the defenses that a defendant can assert to avoid liability for engaging in such activities.

In the next lecture we'll take up the remedies are available to a patentee who prevails in an infringement suit.