INTELLECTUAL PROPERTY:
LAW & THE INFORMATION SOCIETY

Supplement to the 6th Edition, July 2026

James Boyle
William Neal Reynolds Professor of Law
Duke Law School

Jennifer Jenkins
Clinical Professor of Law
Duke Law School

Copyright © 2026 James Boyle and Jennifer Jenkins
This work is made available under a Creative Commons Attribution, Non-Commercial, Share Alike License: https://creativecommons.org/licenses/by-nc-sa/3.0/

No claim is made to copyright on the public domain Federal cases or statutes. You do not have to comply with the license for those materials, or where your use is permitted by an applicable exception or limitation such as fair use.

The Sixth Edition of our casebook (2024) is the most recent edition. This supplement updates this edition with significant new legal developments and a broader focus on the ways that AI may impact intellectual property law and *vice versa*. It is current as of July 2026.

- **Introduction.** This supplement builds on the 6th edition's existing discussions of generative AI to investigate the larger question of how AI will affect creativity and innovation, both in creative expression and in the realms of science, invention, and discovery. It introduces students to some of the changes that AI may bring to the domains of both copyright and patent. There are vitally important legal questions in each sphere. Can an AI system be an author or co-author? An inventor or co-inventor of a patented invention? Is the training of a generative AI system on existing copyrighted works a fair use? But beyond those issues lie deeper possible transformations to the creative and innovative practices that those areas of law regulate, and the administrative systems that we currently depend on to make that law work. What if the cost of filing a patent were to fall dramatically because generative legal AI automated the process of patent prosecution, while scientific AI innovation labs churned out new inventions, both real and illusory? How would the system respond? The Introduction prompts students to think about those issues at the outset of the class, ready to meet the narrower individual questions as they come up in the context of the particular doctrines being discussed in later chapters.

- **Copyright law: Chapter 11.** A discussion of the copyrightability of AI-generated and AI-aided art is added to this chapter.

- **Copyright law: Chapter 13.** The supplement includes an analysis of two contrasting decisions from 2025 about whether training AI models with copyrighted material should qualify for fair use.

- **Copyright law: Chapter 14.** This chapter is updated to include the Supreme Court's recent decision regarding secondary liability for copyright infringement: *Cox Communications, Inc. v. Sony Music Entertainment* (2026).

- **Patent Law: Chapter 17.** A discussion of the significant impacts that AI might have on the patent system is added to this chapter.

- **Patent Law: Chapter 18.** This chapter is updated with a brief section on the patentability of AI-generated inventions and the USPTO's shifting guidance on the matter.

You can find free downloads of the casebook at https://law.duke.edu/cspd/openip/.

Table of Contents

INTRODUCTION, updated section on generative AI ……..…….. 1

Chapter Eleven, COPYRIGHTABLE SUBJECT MATTER, AI and authorship..…… 3

Chapter Thirteen, FAIR USE, contrasting early decisions on AI training and fair use……..…5

Chapter Fourteen, COPYRIGHT SECONDARY LIABILITY, *Cox v. Sony* (2026)……..... 7

Chapter Seventeen, INTRODUCTION TO PATENT LAW, AI and the patent system..…….... 14

Chapter Eighteen, PATENTABLE SUBJECT MATTER, AI and inventorship..……... 16

INTRODUCTION, page *xiii*, after the last full paragraph
This update about generative AI supplements the Introduction.

Throughout its history, intellectual property law has changed in response to disruptive technologies. Sometimes those changes were in the form of amendments to the respective statutes. Examples include the changes made to trademark law to deal with internet domain names, or the broadening of copyright law to address new methods of reproduction, whether audio recordings, television broadcasts, or digital downloads. Sometimes the changes were in the form of court decisions "adapting" the old pigeonholes of the law to deal with the new technologies being jammed into them. Courts' attempts to reconcile copyright law with computer software are a particularly good example.

As we write this supplement, our intellectual property system is attempting to respond to generative AI, a technology that arguably undermines some of its central doctrinal, economic, and even aesthetic assumptions. In fact, it is hard to come up with a major topic in intellectual property that *isn't* likely to face major questions as it confronts AI. Here are just a few.

- **How does copyright apply to AI-generated or AI-aided art?** Copyright only extends to original, creative expression emanating from an "author." Humans have long used technology to create—from the paintbrush to the camera to Photoshop or computer-generated imagery. In the past, copyright law has struggled with the question of how it should deal with these new, technologically-mediated creations. Initially, for example, courts said that photographs were not eligible for copyright because the work was created by a mere mechanical process, not a true act of authorship. Congress and the Supreme Court later brought photography within copyright's purview. How will we answer the same kinds of questions when the technology is generative AI? In *Thaler v. Perlmutter* (D.C. Cir. 2025), *cert. denied* (2026), the court said that purely AI-generated work can never receive a copyright, but what if there is a human in the loop? The Copyright Office has said that the answer depends on the degree of human involvement and creativity but that mere prompting does not satisfy copyright's authorship and creativity requirements. Can that position be maintained? We introduce you to these questions in Chapter 11 on Copyrightable Subject Matter.
- **Is the training of a generative AI system on existing copyrighted works a fair use?** This enormously consequential question is currently at the center of over 100 pending cases. It is introduced in Chapter 1 then explored in more detail in Chapter 13.
- **Can inventions that are solely AI-generated be patented?** As with copyright, the answer here is "no" because patents are reserved for human inventors. In *Thaler v. Vidal* (Fed. Cir. 2022) (yes, this is the same Thaler as in *Thaler v. Perlmutter* – computer scientist Stephen Thaler tried to obtain both a copyright and patents for AI-generated material), the Federal Circuit made clear that "there is no ambiguity: the Patent Act requires that inventors must be natural persons; that is, human beings." The Supreme Court denied certiorari in 2023. But again, as with copyright, the interesting questions are the more subtle ones. What of AI-aided inventions? And how could one ever prove the degree of AI contribution to a particular invention? The USPTO issued one set of guidelines on AI-assisted inventions in 2024, only to withdraw and completely revise them in 2025. *See* https://www.uspto.gov/subscription-center/2025/revised-inventorship-guidance-ai-assisted-inventions. We introduce you to these questions in Chapter 18 on Patentable Subject Matter.

- **How might generative AI affect the patent system more broadly?** Our patent system depends on a series of fundamental assumptions, almost all of which might – or might not – be transformed by AI technologies. Everything could change: from the costs of innovation, and thus the need for patents in the first place, to the universe of prior art and costs of prior art patent searches, or the legal costs of filing a patent application (even an unworthy one). We could be facing anything from a scientific revolution – AI systems noticing "low hanging fruit" of innovations that human eyes have missed – to a scientific dead-end, fueled by corporate hype that leads to an enormously destructive financial bubble and collapse. We could face an explosion of AI-enabled patent trollery, as automated legal systems lower the cost of patent filing, while automated scientific systems generate a panoply of superficially patentable innovations. Each axis of change is significant on its own, and even more uncertain when combined with the others. How will the patent system accommodate these potential changes? All of these questions are explored in Chapter 17, our introduction to the hopes and fears behind the patent system.

The possibilities outlined above are speculation right now. It is hard to separate the hype and doomsaying from reality. One thing we can be certain about, however, is this: Whatever happens with AI and intellectual property law, there are lessons to be learned from the past, and its own technological revolutions and disruptions. For example, many of the readings in the book deal with an earlier technological transformation: the development of digital technologies and the World Wide Web and the ways that intellectual property law was forced, often unwillingly, to confront those changes. Themes from those readings will recur with generative AI.

Once again, we will see dramatic reallocations of the costs and benefits of the system. There will be conflicts pitting incumbents against newcomers. The former group will complain about rampant piracy. The latter group will say that the complacent incumbents are trying to hold back the tide of disruptive innovation – the lamp-oil sellers attempting to regulate electric lighting! There will be doomsayers trumpeting the death of our civilization and optimists talking about the coming Golden Age. There will be familiar questions surrounding judicial methodology. Can judges use the traditional incremental tools of the common law to adjust the law's formal categories to these new technologies, particularly given Congress's apparent acquiescence in this practice in the case of "common law statutes"? Or is that outright judicial usurpation of legislative power? We will debate institutional competence. Which institutions are best suited to deal with radical change? We will worry about path dependency, and argue about the benefits of private as opposed to public solutions when technology outpaces law. In each one of these conflicts, you will discover, again and again, the power of baselines, metaphors, and framing. Finally, there will be things that neither we nor ChatGPT can dream up. What a brave new world to have such problems in it! We hope this book will give you valuable tools with which to understand that world. It will shortly be yours.

Chapter Eleven, COPYRIGHTABLE SUBJECT MATTER, page 287
Insert after the end of the first paragraph

How does copyright apply to AI-generated or AI-aided art?

As we noted in the Introduction, humans have long used technology to create – from the paintbrush to the camera to Photoshop or computer-generated imagery. In the latter set of works, the technology obviously plays a central role. In the past, copyright law has struggled with the question of how it should deal with these new, technologically-mediated creations. Initially, for example, courts said that photographs were not eligible for copyright because the work was created by a mere mechanical process. Congress amended the Copyright Act in 1865 to include photographs, but doubt remained. The Congressional authority to pass copyright laws comes from the Constitution and specifies that it is aimed at the "writings" of "authors." Could a photographer be called an author? A photograph be called a writing? In a case involving a fetching photograph of Oscar Wilde, the Supreme Court found that the definition of writings could expand to fill new modes of artistic fixation and that, because the choices of lighting, framing, costume, and posing were made by the human photographer, photography could be a form of authorship. *Burrow-Giles Lithographic Co. v. Sarony* (1884). How will we answer the same kinds of questions when the technology is generative AI?

If a work is generated *entirely* by AI, it is ineligible for copyright. The D.C. Circuit confirmed this in 2025, holding that the copyright statute required human authorship. *Thaler v. Perlmutter* (D.C. Cir. 2025), *cert. denied* (2026). Notably, the court said it did not need to go further and consider the argument, which had been made by the Copyright Office, that the Constitution itself required human authorship. If that argument were to be accepted it would mean that no mere statutory revision could change that conclusion; Congress could *never* grant copyright to purely AI-generated works. Right now, that issue remains hypothetical because the courts have declined to reach it. It may not always remain that way.[1]

Thaler involved the claim that an AI system was the sole author of a work, so the court did not have to address the more challenging questions surrounding works that are the product of both human and machine creativity. Your carefully posed photograph, or a movie full of CGI effects, are both eligible for copyright despite the fact that they could not exist without the technological tool used to create them. By the same logic, if a person uses AI as a creative tool, then the human authorship requirement could be satisfied. But what degree of human input is sufficient?

The United States Copyright Office's response to that question is appropriately nuanced: "The answer will depend on the circumstances, particularly how the AI tool operates and how it was used to create the final work." (The Copyright Office is the administrative agency that grants copyright registrations and its reports and decisions have no precedential force, though courts sometimes use them for guidance.) In deciding whether works created using AI are copyrightable, the Office looks at whether it is the human or machine that provides the "traditional elements of authorship" and retains "creative control." It has granted registrations when a person selected, arranged, or iteratively edited AI-generated content—in these instances, only the human contributions are copyrightable.

However, in one area the Copyright Office offers a definite response. It maintains that entering prompts into an AI system is inadequate to establish human authorship, reasoning that "prompts function more like instructions to a commissioned artist—they identify what the

[1] For a more comprehensive discussion of the way that courts and legislatures may one day treat claims of legal personality on the part of AI systems, *see* James Boyle, THE LINE: AI & THE FUTURE OF PERSONHOOD (2024).

prompter wishes to have depicted, but the machine determines how those instructions are implemented in its output." This position is currently being challenged in the District of Colorado in *Allen v. Perlmutter*. The Copyright Office had denied a copyright registration for an artwork entitled "Théâtre D'Opéra Spatial," produced by Jason Allen using Midjourney. Allen engaged in "numerous revisions and text prompts at least 624 times" to make the image. His summary judgment motion analogizes his creative input to that of photographers whose works have been deemed copyrightable because (quoting circuit court precedent) "the photographer's decisions regarding pose, positioning, background, lighting, shading, and the like…can be said to 'owe their origins' to the photographer." This is true even though a photograph is "a mechanical reproduction of a natural phenomenon—it is the output of a machine." In its response the Copyright Office reiterates its view that "text prompting does not actually form the image or exercise creative control over the expressive elements…prompters do not determine the expression generated by an AI system like Midjourney and—like Mr. Allen—must keep trying until it produces an image they deem 'acceptable.' That is not authorship."

AI systems are currently the subject of an intense, and perhaps warranted, moral panic. Still, even if the plaintiff in *Allen* does not prevail, it is hard to see how the courts could accept a blanket assertion that prompting, no matter how complex and iterative, can *never* satisfy the creativity requirement for authorship. Are 624 rounds of iterative prompting less creative than the choices of angle, lighting, and even focus you made for your photograph of a beautiful dawn? As with prior technologies, there might be a danger in focusing only on the technology's role and not the user's choices. The ghost of Oscar Wilde would be amused.

Chapter Thirteen, LIMITATIONS ON EXCLUSIVE RIGHTS: FAIR USE, Page 502
Replace the final paragraph on this page with the following

These issues are heading to the courts in a blizzard of over 100 lawsuits (and counting). Two early decisions from 2025 illustrate contrasting perspectives on whether training AI models with copyrighted material is fair use. In *Bartz v. Anthropic* (N.D. Cal. June 23, 2025), Judge William Alsup held that when Anthropic trained Claude with lawfully obtained books, this qualified for fair use. (Training it with pirated books was a different story—this was copyright infringement.) Judge Alsup described the use as "quintessentially," "exceedingly," and "spectacularly" transformative: "The technology at issue was among the most transformative many of us will see in our lifetimes." He wrote:

> Like any reader aspiring to be a writer, Claude trained upon works not to race ahead and replicate or supplant them — but to turn a hard corner and create something different. If this training process reasonably required making copies within the LLM or otherwise, those copies were engaged in a transformative use.

Such a use was "'orthogonal' to anything that any copyright owner rightly could expect to control," and allowing it served copyright's goal of promoting progress, "*without* diminishing the incentive to create." In discussing market harm under factor four, Alsup reasoned that training AI systems was more akin to teaching children to write than unfair competition, and cited to the *Sega* case you read earlier in this chapter:

> Authors' complaint is no different than it would be if they complained that training schoolchildren to write well would result in an explosion of competing works. This is not the kind of competitive or creative displacement that concerns the Copyright Act. The Act seeks to advance original works of authorship, not to protect authors against competition. *Sega.*

Therefore the Authors' existing market was not displaced in a way cognizable under copyright law, and the emerging market for *licensing* books as training inputs was "not one the Copyright Act entitles Authors to exploit."

Compare this analysis to an opinion from Judge Vince Chhabria issued two days later in *Kadrey v. Meta* (N.D. Cal. June 25, 2025). On the evidence presented Judge Chhabria had "no choice but to grant summary judgment to Meta" on fair use, but he made clear that "this ruling does not stand for the proposition that Meta's use of copyrighted materials to train its language models is lawful. It stands only for the proposition that these plaintiffs made the wrong arguments and failed to develop a record in support of the right one." He provided a roadmap for finding Meta's training unlawful, beginning with this framing:

> "[C]ompanies have been unable to resist the temptation to feed copyright-protected materials into their models—without getting permission from the copyright holders or paying them for the right to use their works for this purpose. This case presents the question whether such conduct is illegal.
>
> Although the devil is in the details, in most cases the answer will likely be yes. What copyright law cares about, above all else, is preserving the incentive for human beings to create artistic and scientific works. Therefore, it is generally illegal to copy protected works without permission. And the doctrine of "fair use," which provides a defense to certain claims of copyright infringement, typically doesn't apply to copying that will significantly diminish the ability of copyright holders to make money from their works (thus significantly diminishing the incentive to create

> in the future). Generative AI has the potential to flood the market with endless amounts of images, songs, articles, books, and more. People can prompt generative AI models to produce these outputs using a tiny fraction of the time and creativity that would otherwise be required. So by training generative AI models with copyrighted works, companies are creating something that often will dramatically undermine the market for those works, and thus dramatically undermine the incentive for human beings to create things the old-fashioned way."

Under factor four, Judge Chhabria took issue with Judge Alsup's learning analogy: "When it comes to market effects, using books to teach children to write is not remotely like using books to create a product that a single individual could employ to generate countless competing works with a minuscule fraction of the time and creativity it would otherwise take." Even if the copying was transformative, "harm to the market for the copyrighted work is more important than the purpose for which the copies are made." Notably, Judge Chhabria embraced the nascent theory that *diluting* the market for copyrighted works through "the rapid generation of countless works that compete with the originals, even if those works aren't themselves infringing" can demonstrate market harm. Even though this was "indirect" rather than "direct" substitution, "market dilution" was a "potentially winning argument." Do you find Judge Alsup or Chhabria's perspective more convincing?

The problem that follows puts you in the shoes of a lawyer trying to wrestle with the issues we have just discussed. As you work on it, imagine how you would make your arguments salient and convincing to a judge who might be unfamiliar with technology or with copyright itself. What analogies and precedents would you use? How would you use your skills in implicitly framing issues to build a convincing case on either side?

Chapter Fourteen, SECONDARY LIABILITY FOR COPYRIGHT INFRINGEMENT, Page 507
Replace Introduction *with this revision*

Introduction

Generally, in the common law, "secondary liability" is imposed on someone who does not commit the legal wrong directly, but is found responsible for encouraging, facilitating, or profiting from it. For there to be secondary liability, there needs to be underlying direct infringement. In copyright law, someone must directly violate the exclusive rights provided by § 106 before secondary liability can be imposed on any third party.

As you will see, the Copyright Act has no provisions imposing secondary liability. (Compare to § 271 of the Patent Act.) As a result, the secondary liability provisions of copyright law are judicially developed, without even an open-ended statutory basis like that given to fair use jurisprudence under § 107. *Sony v. Universal* (1984), the case you read in Chapter 13, was the Supreme Court's landmark decision regarding copyright secondary liability and new technology. Twenty-one years later, in 2005, the Supreme Court revisited secondary liability and technology in *MGM v. Grokster*. Then, another twenty-one years later, it significantly reinterpreted *Sony* and *Grokster* in *Cox Communications, Inc. v. Sony Music Entertainment* (2026). This latest episode in the progressive rewriting of secondary liability by the courts concludes this chapter.

As you learned in Chapter 13, *Sony* established the safe harbor for technologies that are "capable of substantial, commercially significant noninfringing uses" and discussed two forms of secondary liability—contributory infringement and vicarious liability—though it did not clearly delineate between them. *Grokster* then outlined another basis for secondary liability: inducement. It was unclear about whether inducement was a third form of secondary liability or a subset of contributory infringement, but it did distinguish inducement from the scenario in *Sony* ("*Sony* dealt with a claim of liability based solely on distributing a product with alternative lawful and unlawful uses, with knowledge that some users would follow the unlawful course...MGM's evidence in this case most obviously addresses a different basis of liability").

While Congress has not explicitly spelled out the bases for imposing secondary liability, it did implicitly address secondary liability in the 1998 Digital Millennium Copyright Act ("DMCA") when it created a safe harbor that shielded qualifying online service providers from both direct and secondary liability. However, this was a statutory *defense* from liability rather than a Congressional articulation of the underlying standard.

In *Cox*, Sony Music Entertainment had sued an internet service provider that had not met the requirements for the DMCA safe harbor. (In *Sony v. Universal*, Sony had been the technology provider; now Sony was the copyright owner suing Cox for the infringing activity of Cox's customers. Ironically, Sony ultimately lost the *Cox* case partly because of the case it had won in 1984.) The Fourth Circuit had affirmed a billion-dollar verdict against Cox for continuing to provide internet service to subscribers after being notified that those subscribers had illegally downloaded music. The Supreme Court reversed. In doing so, it held that there are two categories of secondary liability for copyright infringement: "vicarious" and "contributory" liability, and folded inducement liability within the latter category.

Earlier cases had defined vicarious liability as having 1) a right and ability to supervise the direct infringement and 2) profiting directly from the infringement. *Cox* only addressed

contributory liability because the appeals court had concluded that Cox was not vicariously liable. The Court provided this framework for contributory liability:

> A party is contributorily liable for a user's infringement if it "*intended* its service to be used for infringement," either by 1) affirmatively *inducing* the infringement, as in *Grokster*, or 2) selling a service "*tailored* to infringement…A service is tailored to infringement if it is "not capable of 'substantial' or 'commercially significant' noninfringing uses." (emphasis added)

While the result in the case was not entirely unexpected, the Court's rearticulation of secondary liability took copyright experts by surprise. Decades of appellate decisions – including one that you will read in this chapter – had understood *Sony* as laying out a standard for contributory liability that focused on culpable *knowledge*, not merely intent. Many thought that *Grokster*-style inducement was a third category of liability, not a subset of contributory liability. (The authors of this casebook were among them, largely because the Court had repeatedly indicated that its intent was to develop such a new, third, type of liability.) The "tailored to infringement" language is nowhere in *Sony* or *Grokster*. In introducing it, *Cox* ironically cites a Justice Ginsburg's concurring opinion from *Grokster.* We say ironically, because the Court is citing a concurrence that would have *expanded* secondary liability and *narrowed* the *Sony* safe harbor in order to support a standard that does exactly the opposite. *Cox*'s reinterpretation also appears to be in tension with the premises underlying the DMCA safe harbor for internet service providers.

The materials that follow review the *Sony* decision, introduce you to the *Napster, Grokster,* and *Cox* trinity of cases, along with readings on the DMCA safe harbor against both direct and secondary liability. As you read this material, focus on how the law has morphed and why, on the underlying interpretive methodology and visions of IP motivating the judges, and on what is at stake if the Court errs in either direction.

Chapter 14, SECONDARY LIABILITY FOR COPYRIGHT INFRINGEMENT, Page 546
Insert at the end of Chapter 14.

Cox Communications, Inc. v. Sony Music Entertainment
607 U.S. ___ (2026)

THOMAS, J., delivered the opinion of the Court, in which ROBERTS, C. J., and ALITO, KAGAN, GORSUCH, KAVANAUGH, and BARRETT, JJ., joined. SOTOMAYOR, J., filed an opinion concurring in the judgment, in which JACKSON, J., joined.
JUSTICE THOMAS delivered the opinion of the Court.

Countless people use the Internet for legal activities, but some use it to illegally share copyrighted works, such as songs and movies. The Copyright Act authorizes copyright owners to sue these copyright infringers. In this case, however, instead of suing those infringers, the copyright owners sued petitioners, Cox Communications, Inc., and its subsidiary, who provided the Internet connections that the infringers used. They contended that Cox was itself liable for copyright infringement because it continued to provide known infringers with Internet access.

Based on this theory of infringement, respondents, Sony Music Entertainment and other

major copyright owners, secured a billion-dollar verdict against Cox. The United States Court of Appeals for the Fourth Circuit agreed that because Cox provided Internet service to known infringers, it was a willful infringer itself.

Under our precedents, a company is not liable as a copyright infringer for merely providing a service to the general public with knowledge that it will be used by some to infringe copyrights. Accordingly, we reverse.

I

A

…This Court has…recognized two categories of secondary copyright liability, which means liability for the copyright infringement of another. Those two categories are "contributory" liability and "vicarious" liability. *Metro-Goldwyn-Mayer Studios Inc.* v. *Grokster, Ltd.* (2005). This case concerns contributory liability.

The provider of a service is contributorily liable for a user's infringement if it intended its service to be used for infringement. To establish that a provider intended its service to be used for infringement, a copyright owner must show one of two things. First, it can show that a party affirmatively "induc[ed]" the infringement. *Ibid.* Or, second, it can show that the party sold a service tailored to infringement. *Id*. (Ginsburg, J., concurring)…

B

…Internet service providers, such as Cox, have limited knowledge about how their Internet services are used and who uses them. They do know which IP address corresponds to which subscriber's account, but they cannot distinguish one individual user from another. For instance, if an Internet service provider learns that someone illegally downloaded music from a coffee shop's IP address, the Internet service provider cannot determine which individual at the coffee shop infringed the copyright…

[Sony] enlisted the services of an entity called MarkMonitor to track infringement of its copyrights across the Internet…When MarkMonitor detects apparently infringing activity, it sends notices to the Internet service provider, identifying the IP address at which the infringement occurred. In the roughly 2-year period at issue here, MarkMonitor sent Cox 163,148 such notices.

Cox states that it took steps to limit copyright infringement by those using its Internet services. According to Cox, it created a system of responding to the notices that it received from MarkMonitor. After the second MarkMonitor notice for a subscriber's account, Cox sent a warning to that subscriber. After additional notices, Cox terminated Internet access to that subscriber's IP address until the subscriber responded to the warning. If it continued to receive notices for that IP address, Cox suspended service until the subscriber called and received a warning over the phone. After 13 notices, the subscriber was subject to termination of all Internet service…

The parties disagree about how to characterize Cox's efforts to protect Sony's copyrights. Sony points out that Cox terminated only 32 subscribers for infringement during the claim period, even as it terminated hundreds of thousands of subscribers for nonpayment during the same period. Sony also points to statements from Cox employees expressing frustration with the notices and an unwillingness to act on them in favor of protecting revenue from subscriber payments. Cox resists this characterization by pointing out that its warning and suspension system ended 98% of identified infringement.

C

…In the District Court, Sony prevailed as to both contributory and vicarious liability. The jury found in favor of Sony on both theories. It also found that Cox's infringement was willful,

awarding $1 billion in statutory damages[*]…The Fourth Circuit affirmed in part and reversed in part. It affirmed as to contributory liability because Cox continued to provide Internet service to known infringers. Applying Circuit precedent, it reasoned that "supplying a product with knowledge that the recipient will use it to infringe copyrights is exactly the sort of culpable conduct sufficient for contributory infringement." The Fourth Circuit reversed as to vicarious liability because it concluded that Cox did not "receiv[e] a direct financial benefit from its subscribers' infringement."…

II

A

"The Copyright Act does not expressly render anyone liable for infringement committed by another." *Sony*. Ordinarily, when Congress intends to impose secondary liability, it does so expressly. Although our precedents have recognized specific forms of secondary copyright liability that predated the Copyright Act, we are loath to expand such liability beyond those precedents.

B

The provider of a service is contributorily liable for the user's infringement only if it intended that the provided service be used for infringement. The intent required for contributory liability can be shown only if the party induced the infringement or the provided service is tailored to that infringement. *Grokster*; *Sony*.

A provider induces infringement if it actively encourages infringement through specific acts. *Grokster*…A service is tailored to infringement if it is "not capable of 'substantial' or 'commercially significant' noninfringing uses." *Grokster* (Ginsburg, J., concurring) (quoting *Sony*)…These two forms of contributory infringement track patent law. Under 35 U. S. C. §271(b), "[w]hoever actively induces infringement of a patent shall be liable as an infringer." Such liability requires that the party express "an affirmative intent that the product be used to infringe." Under §271(c), a party is liable when it sells a product used for infringement "knowing the same to be especially made or especially adapted for use in an infringement of such patent."

This Court has repeatedly made clear that mere knowledge that a service will be used to infringe is insufficient to establish the required intent to infringe…In *Sony*, the Court explained that "[t]here is no precedent in the law of copyright" for liability based only "on the fact that [the defendant] has sold equipment with constructive knowledge of the fact that its customers may use that equipment to make unauthorized copies of copyrighted material." And, in *Grokster*, the Court confirmed that "a court would be unable to find contributory infringement liability merely based on a failure to take affirmative steps to prevent infringement."

III

Thus, Cox is not contributorily liable for the infringement of Sony's copyrights. Cox provided Internet service to its subscribers, but it did not intend for that service to be used to commit copyright infringement. Holding Cox liable merely for failing to terminate Internet service to infringing accounts would expand secondary copyright liability beyond our precedents.

Cox neither induced its users' infringement nor provided a service tailored to infringement. As for inducement, Cox did not "induce" or "encourage" its subscribers to infringe in any manner. Sony provided no "evidence of express promotion, marketing, and intent to promote"

[*] Cox could not invoke the DMCA safe-harbor defense based on its efforts to reduce infringement because an earlier decision had foreclosed that defense for the relevant period. [Editors' note: This earlier decision held that Cox did not qualify for the §512(a) safe harbor because its repeat infringer termination policy was not "reasonably implemented" in violation of §512(i). Its enforcement efforts fell short because Cox "failed to implement its policy in any consistent or meaningful way—leaving it essentially with no policy."]

infringement. And, Cox repeatedly discouraged copyright infringement by sending warnings, suspending services, and terminating accounts. As for providing a service tailored to infringement, Cox's Internet service was clearly "capable of 'substantial' or 'commercially significant' noninfringing uses." Cox did not tailor its service to make copyright infringement easier. Cox simply provided Internet access, which is used for many purposes other than copyright infringement.

The Fourth Circuit found otherwise based only on its Circuit precedent establishing a new form of contributory liability. The court did not suggest that Cox induced its users to infringe. And, it did not deny that Cox's service was "capable of substantial lawful use and not designed to promote infringement." Rather, the court held that "supplying a product with knowledge that the recipient will use it to infringe copyrights is . . . sufficient for contributory infringement." The Fourth Circuit's holding thus went beyond the two forms of liability recognized in *Grokster* and *Sony*. It also conflicted with this Court's repeated admonition that contributory liability cannot rest only on a provider's knowledge of infringement and insufficient action to prevent it.

IV

Finally, Sony argues that the DMCA safe harbor would have no effect if Internet service providers are not liable for providing Internet service to known infringers. The DMCA safe harbor protects Internet service providers that terminate repeat infringers "in appropriate circumstances." 17 U. S. C. §512(i)(1)(A). Sony argues that Congress must have enacted the DMCA on the presumption that Internet service providers could be held liable in cases such as these.

Sony overreads the DMCA. Sony does not contend that the DMCA expressly imposes liability for Internet service providers who serve known infringers. It does not. The DMCA merely creates new *defenses* from liability for such providers. And, the DMCA made clear that failure to comply with the safe-harbor rules "shall not bear adversely upon . . . a defense by the service provider that the service provider's conduct is not infringing." §512(l).

V

The judgment of the Court of Appeals for the Fourth Circuit is reversed, and the case is remanded for further proceedings consistent with this opinion.

It is so ordered.

JUSTICE SOTOMAYOR, with whom JUSTICE JACKSON joins, concurring in the judgment.

…The majority holds that Cox is not liable solely because its conduct does not fit within the two theories of secondary liability previously applied by this Court. In so doing, the majority, without any meaningful explanation, unnecessarily limits secondary liability even though this Court's precedents have left open the possibility that other common-law theories of such liability, like aiding and abetting, could apply in the copyright context. By ignoring those past decisions, the majority also upends the statutory incentive structure that Congress created.

I nonetheless agree with the majority that Cox cannot be held liable here for a different reason. Plaintiffs cannot prove that Cox had the requisite intent to aid copyright infringement for Cox to be liable on a common-law aiding-and-abetting theory. I therefore concur in the judgment…

II

A

To determine whether Cox may be held liable, the majority starts by correctly explaining that "[t]he provider of a service is contributorily liable for the user's infringement only if it intended that the provided service be used for infringement." The majority, however, errs in the very next sentence. It asserts, with no meaningful explanation, that "[t]he intent required for contributory liability can be shown only if the party induced the infringement or the provided service is tailored to that infringement." Because plaintiffs do not satisfy those theories of contributory liability, according to the majority, their claims cannot succeed.

The inflexible limit the majority imposes is nowhere to be found in either *Sony* or *Grokster*, the only authorities that the majority cites…Properly understood, *Sony* and *Grokster* preserved other forms of secondary liability derived from the common law…

B

The majority's limiting of secondary liability here dismantles the statutory incentive structure that Congress created. Congress passed the Digital Millennium Copy-right Act (DMCA) in 1998, 14 years after this Court held in *Sony* that the Copyright Act impliedly provided for secondary liability...Importantly, Congress did not provide that ISPs could never be secondarily liable for copyright infringement. Instead, it struck a balance by creating incentives for ISPs to take reasonable steps to prevent copyright infringement on their networks, while also assuring ISPs that they do not need to take on the impossible task of responding to every instance of infringement on their networks.

The majority's new rule completely upends that balance and consigns the safe harbor provision to obsolescence. Typically, this Court tries " ' "to give effect, if possible, to every clause and word of a statute."'" After today, however, ISPs no longer face any realistic probability of secondary liability for copyright infringement, regardless of whether they take steps to address infringement on their networks and regardless of what they know about their users' activity. For example, under the majority's rule, an ISP faces no liability if it sells an internet connection to a company that the ISP knows runs a website that exclusively hosts illegally obtained copyrighted material. That ISP also faces no liability even if it sells a connection to a customer who walks into the store and says that he needs a new internet connection because the other, more scrupulous ISP in town cut his connection after years of unabated piracy.

The majority's decision thus permits ISPs to sell an internet connection to every single infringer who wants one without fear of liability and without lifting a finger to prevent infringement. It also means that Cox is free to abandon its current policy of responding to copyright infringement. As Cox's counsel conceded at oral argument, under the rule the majority adopts today, the safe harbor provision will not "d[o] anything at all" going forward. Congress did not enact the safe harbor just so that this Court could eviscerate it…

Questions:

1.) *Cox* shifts the basis of contributory liability from culpable knowledge to culpable *intent*, even though both *Sony* and its corollary from the Patent Act, §271(c), explicitly focus on knowledge. In doing so, *Cox* supplanted the definition of contributory infringement applied in decades of appellate cases, including *Napster* ("one who, with knowledge of the infringing activity, induces, causes or materially contributes to the infringing conduct of another, may be held liable as a 'contributory' infringer"). The Court cites to *Sony* and *Grokster* for its new framework. You have read these cases; is *Cox* consistent with this precedent? Would Napster still lose under *Cox*'s new definition of contributory infringement? When might culpable knowledge and intent overlap:

should actual knowledge, or knowledge with substantial certainty that copyright infringement will occur, amount to culpable intent? Does this help harmonize *Cox* with precedent?

2.) Under *Cox* there are two forms of contributory liability: 1) *Grokster*-style inducement and 2) selling a service "tailored to infringement," defined as one that is "not capable of 'substantial' or 'commercially significant' noninfringing uses." As mentioned in this chapter's introduction, the "tailored to infringement" formulation is entirely new. It is defined as anything outside of the *Sony* safe harbor, essentially collapsing a ground for liability with a limitation on it. In his *Grokster* concurrence, Justice Breyer extolled the technology-protective merits of the *Sony* safe harbor. Is this a positive development, then, or a non-standard that says "you're only on the hook if you're not off the hook"?

3.) Sotomayor's concurrence raises concerns about the inconsistency between *Cox* and the DMCA 512 safe harbor. Has the majority opinion "consign[ed] the safe harbor provision to obsolescence" and "eviscerate[d]" Congressional intent? Think about the service providers that would fall within the 512(a), 512(c), and 512(d) categories. Can ISPs such as Cox now ignore the safe harbor's requirements? Do 512(c) entities such as YouTube still have incentives to comply with DMCA takedown notices? How about 512(d) search engines? (Hint: think about the difference between storing copyrighted content at the direction of users and linking to it.)

Chapter Seventeen, INTRODUCTION TO PATENT LAW: HOPES AND FEARS, page 627
Insert after the end of the second paragraph

2 b.) AI & the Patent System

As we noted in the Introduction, the patent system depends on a series of fundamental assumptions, almost all of which might – or might not – be transformed by AI technologies. Here are just a few of them.

- *The costs of innovation*: Patents are designed to get us inventions that the normal workings and incentives of the market would not generate. We are willing to pay the monopoly costs of a patent to get the, otherwise unattainable, innovation. That is one reason why we do not hand out patents for innovations that are "obvious"; no special incentive is needed to take the obvious next step in a ladder of innovation, even if no one has done it before. Some members of the patent bar fear that a glut of AI-produced "prior art" will render many inventions obvious and therefore unpatentable. When AI-aided search for prior art is added to the equation, this might make almost all inventions seem obvious with the benefit of hindsight, meaning that worthy inventions will be denied patents. The AI system's superhuman ability to scan and synthesize documents means that it will always be able to find existing prior art to challenge any innovation. This is a very real concern, but it is also a narrow one, given the magnitude of the potential changes to our world. What if AI systems dramatically lower the cost of innovation itself, not merely the cost of running a prior art search? Will we need the patent incentive less, or for a different and smaller class of inventions? Will we need it at all? Do things that were previously "non-obvious" – and thus patentable – become obvious to AI systems that can uncover *and implement, not merely document,* potential technological solutions that human eyes might miss? At the moment, we simply don't know whether AI-aided science will be a revolutionary change, a bubble built on hype that collapses with devastating effects on the larger economy, or something in between, a marginal improvement that leaves the basics unchanged. As a result, we don't know what challenges the patent system will face and how its assumptions might need to be rethought.
- *The patent examination process:* Does AI transform not only the process of invention, but of patent *examination*, as patent applications on AI-assisted inventions, filed by AI-assisted lawyers, are assessed by AI-assisted patent examiners?
- *The costs of obtaining patents:* In our current world, it is expensive to get a patent. On average, costs range from $15,000 to more than $40,000. The price depends very much on the type of technology and the complexity of the legal work involved – the legal work making up a large percentage of the cost. Software, business methods, and genetic technologies are among the more expensive areas in which to gain a patent, partly because of the complexity of the technology and partly because of the sheer number of billable hours it will take to complete the patent process for such an innovation. (Compare this to the copyright system. Copyright applies instantly upon the fixation of a creative expressive work – no registration is needed, though one has to register a copyright in order to sue. But even registration is cheap.) The price tag of patents is a transaction cost that reduces overall efficiency, of course, a theoretical drag on innovation. But it can also play a positive role in the system. We know that there are entities misusing the patent system, gaining "junk patents" over innovations that are actually obvious, or have already been made. Some

of those patents make it through the system, sadly – particularly in certain areas such as software, business methods, and certain forms of genetic technology. But the cost of filing acts as a restraint on that misuse. A rational abuser of the system will only file when the expected benefits of the junk patent exceed the costs they had to pay up front. What if AI-assisted labs can churn out thousands of purported technological innovations, both real and fake, while AI-assisted law offices automate the process of prosecuting patents, lowering the cost substantially? We would expect an explosion of patent misuse, high-tech patent trollery, AI-style. And that might require major adjustments to the patent system.

Right now we are in a state of fundamental ignorance about which, if any, of these possible futures will come to pass. But the theme of this book is that one can nevertheless "future-proof" one's knowledge and skills, even in the face of radical technological uncertainty. We believe it will be borne out here. We think it is worth quoting a paragraph from the Introduction, noting the similarities that we can expect between this revolution and prior disruptive technologies.

> Once again, we will see dramatic reallocations of the costs and benefits of the system. There will be conflicts pitting incumbents against newcomers. The former group will complain about rampant piracy. The latter group will say that the complacent incumbents are trying to hold back the tide of disruptive innovation – the lamp-oil sellers attempting to regulate electric lighting. There will be doomsayers trumpeting the death of our civilization and optimists talking about the coming Golden Age. There will be familiar questions surrounding judicial methodology. Can judges use the traditional incremental tools of the common law to adjust the law's formal categories to these new technologies, particularly given Congress's apparent acquiescence in this practice in the case of "common law statutes"? Or is that outright judicial usurpation of legislative power? We will debate institutional competence; which institutions are best suited to deal with radical change? We will worry about path dependency, and argue about the benefits of private as opposed to public solutions when technology outpaces law. In each one of these conflicts, you will discover, again and again, the power of baselines, metaphors, and framing. Finally, there will be things that neither we nor ChatGPT can dream up. What a brave new world to have such problems in it! We hope this book will give you valuable tools with which to understand that world. It will shortly be yours.

Chapter Eighteen, PATENTABLE SUBJECT MATTER, page 683
Insert before Conclusion

Can AI-generated inventions be patented?

As we mentioned in the book's Introduction, the question is easy if the invention is solely generated by the AI. As with copyright, the answer here is "no" because patents are reserved for human inventors. In *Thaler v. Vidal* (Fed. Cir. 2022) the Federal Circuit made clear that "there is no ambiguity: the Patent Act requires that inventors must be natural persons; that is, human beings." The Supreme Court denied *certiorari* in 2023. But again, as with copyright, the interesting questions are the more subtle ones. What of AI-aided inventions? And how could one ever prove the degree of AI contribution to a particular invention? The answers here are considerably more murky.

The USPTO issued one set of guidelines on AI-assisted inventions in 2024, only to withdraw and completely revise them in 2025. *See* https://www.uspto.gov/subscription-center/2025/revised-inventorship-guidance-ai-assisted-inventions.

The 2024 Guidelines had applied to AI-assisted inventions the so-called *Pannu* test that is used in human-human collaborations on inventions. To be eligible for a patent:

- "the human must contribute in some significant manner to the conception or reduction to practice of the invention,
- make a contribution to the claimed invention that is not insignificant in quality, when that contribution is measured against the dimension of the full invention, and
- do more than merely explain to the real inventors well-known concepts and/or the current state of the art."

In other words, the AI was treated *almost like a co-inventor*. To be patent eligible at all, the human contribution had to be large enough that, had it been the contribution of a human inventor to a joint (human) research project, it would have entitled that co-inventor to be listed on the patent.

The 2025 Guidelines scrapped this guidance and substituted different rules. Far from being treated as a species of co-inventor, the AI was now treated as a mere mechanical tool, like laboratory equipment or computer software. The *Pannu* factors are once again reserved for purely human collaborations. The Guidance announced that the "same legal standard for determining inventorship applies to all innovations" regardless of whether they are AI-aided. In theory, this should make it easier for AI-assisted innovations to be patent-eligible.

Some have described the new Guidelines as focusing on whether the human had the "big idea" or central conception behind the innovation, using the AI merely to assist with the fiddly details, but the actual test as *stated* seems more stringent:

> The Federal Circuit has centered its inventorship inquiry around "conception," characterizing conception as "the touchstone of inventorship." Conception is "the formation in the mind of the inventor, of a definite and permanent idea of the complete and operative invention, as it is hereafter to be applied in practice." Conception is complete when "the inventor has a specific, settled idea, a particular solution to the problem at hand, not just a general goal or research plan." ... The question is whether the natural person possessed knowledge of all the limitations of the claimed invention such that it is so "clearly defined in the inventor's mind that only ordinary skill would be necessary to reduce the invention to practice, without extensive research or experimentation." Analysis of conception turns on

> the ability of an inventor to describe an invention with particularity. Absent such a description, an inventor cannot objectively prove possession of a complete mental picture of the invention at a later time.

So stated, the test would appear to require a lot of the human inventor. It would not be enough merely to collaborate on innovation research in which the AI system saw something in the data, proposed an innovation around it, and had the human tweak the details. On the other hand, how will these rules work in practice? The Guidelines also state that "the USPTO presumes those inventors named on the application data sheet or oath/declaration are the actual inventor or joint inventors of the application." In other words, the PTO's system will largely depend on honest disclosure by innovators of the precise roles they played in a complex inventive process. How many scientists are happily going to describe themselves as playing second fiddle to an AI? Even if we leave aside entirely the possibility of fraud and "creative patent drafting," behavioral psychology tells us that individuals often overstate their centrality to research projects. Challenges to such patents will probably be highly fact-intensive and depend on exhaustive discovery proceedings around the records of the inventive process. At the end of the day, then, so long as there is a human who can credibly assert that they had the big idea, the invention will probably be treated as patent-eligible by the PTO, regardless of what the facts are on the ground.