

Intellectual Property and Creative Machines

Gaétan DE RASSENFOSSE
Adam JAFFE
Joel WALDFOGEL

June 2024

Innovation and Intellectual Property Policy
Working Paper series no. 27

Available at: <https://ideas.repec.org/p/iip/wpaper/27.html>



STiP lab

Intellectual Property and Creative Machines

Gaétan de Rassenfosse, École polytechnique fédérale de Lausanne

Adam B. Jaffe, Brandeis University

Joel Waldfogel, University of Minnesota

This version: June 24, 2024

Abstract. The arrival of creative machines—software capable of producing human-like creative content—has triggered a series of legal challenges about intellectual property. The outcome of these legal challenges will shape the future of the creative industry in ways that could enhance or jeopardize welfare. Policymakers are already tasked with creating regulations for a post-generative AI creative industry. Economics may offer valuable insights, and this paper is our attempt to contribute to the discussion. We identify the main economic issues and propose a framework and some tools for thinking about them.

Keywords: generative AI, machine learning, copyright, fair use

I. Introduction

Generative artificial intelligence (AI) can produce new inventions, images, musical works, poems, essays, novels, and other creative works. We call generative AI models ‘creative machines.’ Some of the outputs of these creative machines are indistinguishable from human creations. It is controversial whether creative machines will ever be capable of matching the range of creativity humans exhibit. Still, there is no question that significant portions of that range will be within the power of creative machines and that human creators will increasingly use such machines as part of their own creative processes.


Creative outputs are protected in most countries and by international treaties under legal regimes referred to broadly as ‘intellectual property’ (IP). IP comes in many forms, including copyrights, patents, design rights, trademarks, and trade secrets. It serves broadly to foster creation and enables the effective and efficient use of creative outputs by giving creators rights to control legally, to some extent, how others use their creations. New technology as powerful as these creative machines will inevitably have significant impacts on IP regimes. They pose challenges that may call for changes in IP rules.

Our IP system is under pressure. The stakes for copyright owners are higher than ever, with the scale of potential copyright infringement today dwarfing the music piracy disputes of the early 2000s. Uncertainties about copyright infringement pose significant legal risks for technology developers (and possibly users), potentially chilling investments in generative AI. High transaction costs to secure the relevant rights to training data may represent a barrier that only the wealthiest developers can overcome. Furthermore, the lack of copyright protection for AI-generated work may curb implementers’ appetite for such work. More generally, the copyright system was not designed for the generative AI world, and generative AI can potentially turn the copyright law upside down, as Lemley (2024) argued.

As with any new technology, it is impossible to foresee reliably how creative machines will develop or how they will interact with the economic, social, and institutional systems in which creativity plays out. In this paper, operating primarily from the perspective of economic analysis of social welfare, we analyze some of the issues that have arisen or are likely to arise as IP systems confront this new creative technology. We will focus primarily on matters for the copyright system, but where relevant, we also mention related issues connected to the patent system. Some of the patent-related issues are discussed in a separate paper by a partially overlapping set of authors (de Rassenfosse, Jaffe, and Wasserman 2024).

Figure 1 introduces a framework to help identify the main interactions at play. We consider two classes of creators, namely humans and machines. These creators produce intermediate outputs such as movie scripts or inventions, which then reach the market as final goods, for example, movies or innovative products. Creators and implementers may or may not be the same person. At the top of the Figure are generative AI models, which are trained on material produced by humans (at least initially) and which feed creative machines. The arrows depict the main interactions discussed in the present paper, for which IP plays—or has the potential to play—a pivotal role.

Figure 1. Key interactions in the creative industry



Note: The arrows depict the key interactions in which IP plays a role.

The rest of the paper is organized as follows: Section II explains the basic economics of IP to set the stage for the rest of the discussion. It focuses on both patent rights and copyrights. Section III extends the framework by considering relevant IP policy trade-offs and noting critical differences between patents and copyrights. Section IV briefly touches upon some non-economic considerations that we view as crucial for framing essential elements of the discussion. Section V discusses the potential effects of generative AI on creative activities. We introduce a series of tools to shed light on these effects. Section VI ventures into different compensation regimes for both copyright owners and developers of creative machines, taking the current situation as a starting point. Section VII considers long-run issues, notably regarding human-machine interactions. Finally, Section VIII provides some concluding remarks.

II. Basic Economics of IP

A. Intangible assets: costly to create but often easy to copy

Copyrights protect creative works such as text products (books, magazines, newspapers), video products (movies, television), music (recordings, publishing), software, and other original works of authorship. Patents protect inventions. Consumers—and producers—benefit from a continued flow of new creative products. Consumers enjoy the stream of services from these IP-protected products that they value beyond the price they pay, and creators (and their intermediaries) benefit from the revenue they receive in excess of their cost.

Creative works are generally seen as expensive to create but easy to copy. For example, an author must expend effort to write a book. In comparison, it is relatively easy for another person to create a copy of the book. The same also holds for technical inventions. Using survey data, Mansfield, Schwartz, and Wagner (1981) find that the cost of imitating an existing invention amounts to about two-thirds of the cost of developing the original invention. If competitors are allowed to offer copies for sale in competition with the original creator, the revenue diverted from the creator may

undermine the incentive to create in the first place. IP—the grant of an exclusion right to the creator—preserves the creator’s revenue and, therefore, their incentive to create.

In many cases, incentives for development and implementation are at least as important as incentives for the underlying creative act. Some people write novels they never publish or tinker with inventions in their garage that are never commercialized. The possibility of IP protection for these creations is typically not necessary to bring them forth. But suppose an inventor, artist, or other creator does wish to develop their idea into a commercial product. In that case, there may be significant investment necessary to get it to that point—think of testing a possible drug to prove it is safe or filming a movie. Regardless of whether the underlying creation would have been forthcoming without the possibility of IP protection, the investment in development and commercialization could be jeopardized if the resulting product could be easily copied and sold by others. The distinction between incentives for creation and incentives for development is crucial for the analysis of creative machines’ consequences for IP. Creative machines may significantly reduce the marginal cost of creation; all else equal, if creation becomes cheaper, less incentive is needed to bring it forth.¹

B. IP rights create a trade-off between incentives and costs

IP rights allow creators to restrict the extent to which others can compete with them (or their agents) in implementing the creations. The benefit of competition is one of the few things economists agree on. Monopolies raise prices and restrict output relative to efficient levels. The decision to create patents and copyrights reflects the view that the social benefit of protection—in preserving creative and development incentives—offsets the inefficiency arising from the exclusion right granted.

Said another way, monopolies create static inefficiency once products (creative works and inventions) exist. Their owners generally charge prices above marginal cost, preventing some efficient instances of potential usage (in which a user values the product above its marginal cost but below the price). The preservation of revenue for already-existing products is not the goal. Eliminating monopolies on already-existing products could give rise to more efficient allocation. However, removing protections would have a dynamic cost: future creators would have less incentive to create and develop new products. Discussions on a waiver of IP rights related to COVID-19 technologies offer a prominent recent example of such tensions. Some observers worried that an IP waiver would undermine the IP system’s ability to foster vaccines or other therapies for the next pandemic (*e.g.*, Hilty et al. 2021).

IP protection has another potential dynamic cost. The protection of existing creations may inhibit the creation of future creations that build in some way on the protected works; society loses out if fewer future works are forthcoming. Patents and copyrights deal with this tension in different ways. For patents, in exchange for the right to restrict others from using an invention, the patentee is required to disclose how the invention works, enabling others to use that knowledge in future inventions (in principle). A patentee has no right to limit the use of such future inventions unless they infringe the specific ‘claims’ that constitute the legal definition of the original patented

¹ The point about a drastic reduction in the cost of creation is particularly valid for (some) creative works. Regarding technological inventions, evidence suggests that ‘ideas’ are getting more expensive to produce (Bloom et al. 2020).

invention. In contrast, a copyright on a given work includes the right to control the production of ‘derivative’ works. What constitutes a derivative work has been the subject of case law, but it is inherently a difficult line to draw. An essential issue in the copyright treatment of creative machines is the extent to which their outputs might be legally ‘derivative’ of the copyrighted works used to train the machine. (See Section V.D.1 below.)

C. IP and Creative Machines: Just say no?

A threshold question is whether the output of creative machines with little to no human input should be eligible for protection in the form of patents and copyrights. The answer to that legal question from both patent and copyright offices around the world has been essentially ‘no’ (e.g., Schwartz and Rogers 2021, Abbott and Rothman 2023). A machine cannot be named as an inventor on a patent, and it cannot be the creator of a work for which copyright registration is sought.² These conclusions derive directly from existing law that sees patents and copyrights as legal rights that are only awarded to natural persons or legal entities. Indeed, the fundamental purpose of IP laws for the last centuries was to reward human ingenuity, so there was no need to separate the subject matter from the human. But IP laws were written at a time when there was no artificial ingenuity.

From the economic perspective, the appropriate way to think about whether the output of creative machines should be eligible for IP protection is to balance the costs of granting such protection against the benefits that it would generate. A conclusion that such protection is never appropriate would make sense only if there are no or only very limited circumstances under which the benefits of such protection (in terms of incentives for socially desirable activities) exceed the cost. It is hard to see how or why this would be true. Creative machines are a powerful new technology that produces useful creative works at low marginal costs. It seems extremely unlikely that a blanket prohibition of IP protection is socially beneficial.

There is another reason to question the social value of a blanket prohibition on patents or copyrights for the outputs of creative machines. All such machines are owned and operated by humans or human institutions. If there is a blanket prohibition on the protection of machines’ output, those humans and institutions will still use the machines for creation, but perhaps in a more limited way that allows them to claim IP rights. From a social perspective, we are not indifferent as to *how* these machines are used. We would like them to be used in ways that maximize their social value. Put differently, the rules governing how and under what circumstances creative outputs associated with creative machines are entitled to IP protection (and how that protection might be structured) are a potential policy tool for shaping the use of this technology in socially beneficial ways.

A more subtle point relates to the fact that a blanket prohibition may force the machine users and owners to obfuscate the role of the machines. If IP rules induce everyone to pretend that AI plays little or no role in generating creations, it will become very difficult to learn about the different ways in which generative AI is used. This lack of information, in turn, may prevent the

² There are exceptions to this general observation in some countries. For example, computer-generated work can be copyrighted in the United Kingdom under Section 178 of the 1988 UK Copyright Designs and Patents Act. China also seems to recognize copyright protection for AI-generated work, as discussed in <<https://www.twobirds.com/en/insights/2024/china/copyright-protection-for-ai-generated-works-recent-developments>>, last accessed, June 22, 2024.

development of a nuanced IP policy that recognizes the role of generative AI and attempts to treat it in such a way as to maximize the social benefits it generates.

Therefore, notwithstanding the general legal conclusion that the output of creative machines cannot be protected by IP, our approach herein is to recognize that such protection *could* be granted. Consequently, we analyze factors that affect how different forms of and approaches to such protection are likely to affect the overall social value of the technology.

III. Analyzing IP policy trade-offs

A natural way to analyze the trade-offs implicit in IP is to maximize the present discounted value of consumer surplus and producer profit. Lesser IP protections (shorter, narrower, weaker, more expensive) benefit consumers for already-existing products but reduce rewards for creation. Broader protection creates stronger incentives for creation. This rubric generates an ‘optimal’ level of IP protection that depends on the costs of creation as well as the ease with which works can be copied. While we do not undertake any formal modeling of this kind herein (see, *e.g.*, Budish, Roin, and Williams 2016), it is still helpful as a framework to think about directional changes in IP policy that might be appropriate in response to the changes that creative machines introduce to the system.

The rubric of maximizing social surplus suggests a starting point for thinking about how a new technology might require changes in the IP system. Technological change affects both the costs of creation and creators’ ability to appropriate the economic returns to their creations. A look into recent history makes this clear. Consider the book industry and the evolution of copying technology. Historically, an author created a manuscript. If a publisher agreed to take it on, they invested in editing and design, then typeset and printed the book. Finally, the publisher invested in marketing and convinced retailers to stock it. These were costly activities, and many books failed. Once a book existed, it was relatively complex to copy it. A competing publisher could, in principle, copy it, but this sort of copying was relatively easy to detect and punish under copyright law.

With the invention of the photocopy machine, the cost of copying fell, and copying could be undertaken by decentralized individuals. With the exception of the copying of academic journal articles, the copy machine did not change effective protection. Most books were priced so that purchasing them was cheaper than copying them. Up to the 1980s, technological changes did not substantially change creation costs or appropriability. Academic publishers did change their library pricing, however (Liebowitz 1985).

The arrival of digital technologies, first computers and then the Internet, changed both the costs of creation and the difficulty of appropriation. Music and books were relatively easy to digitize, even with the technology available prior to 2000. Associated files were easy to share online without the permission or assent of the rights holder. Piracy flourished, especially in the music industry, reducing revenue significantly (Waldfoegel 2018). The situation began to look like the classic example of creators’ inability to appropriate the returns to their creation, threatening the long-run viability of the creative activity. On its own, a technological change facilitating piracy would require an offsetting increase in effective IP protection—or, at the very least, an increase in IP enforcement—to maintain creative incentives.

At the same time that the existence of the Internet threatened appropriability with easier copying, other aspects of digitization reduced the costs of creation, distribution, and even promotion. While it had traditionally cost roughly a million dollars to bring an album to market, an artist could now record music using inexpensive software (including GarageBand on their phone) and distribute it through iTunes and later Spotify without much investment. Given reduced costs of creation—and no other changes—pre-digitization creative incentives could be achieved with less protection.

The net effect of reduced appropriability and reduced costs delivered an explosion of creative activity in digitization's wake. The number of new books, songs, movies, and television programs increased sharply in the first decades of the millennium. Moreover, the appearance of music distribution platforms such as Apple Music and Spotify allowed for both a recovery of revenue generation via bundled sales, as well as a transformative product more valuable to consumers than the pre-existing, *à la carte* recorded music offerings.

The story of digitization's impact on copyright is perhaps a helpful prelude to thinking about generative AI. Generative AI can have many possible effects on the demand for and supply of creative content.

On the one hand, creative machines may serve as unauthorized distribution channels without the permission of—or compensation to—the creators. Consider a generative AI model trained on, say, New York Times articles. If the machine is able to express the same content in different words, it may directly compete with the New York Times, diverting demand from the newspaper. Moreover, creative machines using New York Times content may tarnish the newspaper's brand by offering hallucinatory recounting of their articles. To the extent that the technology diverts revenue from underlying creators whose continued output is needed, among other things, to train machines, the arrival of generative AI would call for a strengthening of effective IP protection.

On the other hand, generative AI is also a tool for creation, which may increase creators' productivity. An emerging literature seems to suggest that generative AI increases workers' productivity in both routine and knowledge-intensive tasks, although who stands to gain from increased productivity is context-dependent (Brynjolfsson, Li, and Raymond 2023, Dell'Acqua et al. 2023, Wang, Gao, and Agarwal, forthcoming).

A. Transaction costs

Markets involving creative outputs are typically characterized by a significant division of labor, such that the entities that create (*e.g.*, inventors, academic start-up firms, writers, photographers) are often different from the entities that develop and commercialize products based on these creations (Arora, Fosfuri, and Gambardella, 2004). This means that the social benefits that markets create may be sensitive to the nature and extent of transaction costs, as such costs can inhibit or prevent valuable division of labor. IP, such as patents and copyrights, often play a crucial role in facilitating these transactions (*e.g.*, Lamoreaux and Sokoloff 2001, Gans, Hsu, and Stern 2008, de Rassenfosse, Palangkaraya, and Webster 2016). On its own, this argument speaks in favor of maintaining IP rights for AI-generated content to prevent markets for creative works from unraveling.

The role of transaction costs in markets for creative goods is also amplified by the fact that the other costs connected with the marginal use of a creative good are often very low. Particularly with digital technology, the cost of copying and transmitting music, text, or pictures to another user may be close to zero. This means that if the transaction costs associated with connecting a creator and a user are significant, they will be a substantial fraction of the entire cost of that user using that work. As discussed below, it is relatively easy for entities with appropriate funding and infrastructure to collect enormous amounts of existing material for training a generative AI model. But if they need to reach a contractual agreement with every creator of every one of the works in order to use them legally, the transaction costs associated with that set of permissions would be large relative to training costs—perhaps prohibitively large.^{3,4} Therefore, mechanisms that reduce the transaction costs of accessing training data will be fundamental to realizing generative AI’s full potential.

B. Differences between patents and copyrights

The theoretical frameworks for understanding patents and copyrights share apparent similarities. However, there are also notable differences. While it is beyond the scope of this paper to systematically analyze the differences between patents and copyrights, we will note a few that are crucial for their interaction with creative machines.⁵

The most crucial difference is that patents protect the concrete implementations of concepts and ideas to solve technical problems. In contrast, copyrights do not protect concepts or ideas but rather specific textual, visual, musical, or artistic expressions. It is controversial whether creative machines ‘understand’ or utilize concepts and ideas; their focus by construction is on the particular expressions (whether textual, visual, or other) on which they are trained. As a result, as discussed below, the training and use of creative machines raises immediate issues of infringement of existing copyrights. In contrast, because the focus of patents is on the concepts, the training and operation of creative machines do not raise significant issues of patent infringement; the issues they raise for patents have more to do with how the use of creative machines to generate inventions affects the operation of the patent system (de Rassenfosse, Jaffe, and Wasserman, 2024).

A second difference is that patents are in force generally for a maximum of ‘only’ 20 years from the time of application, whereas copyrights are in force for 70 years beyond the creator’s death. Furthermore, unlike patents, which require a formal application and approval process, copyrights do not require registration for the protection to be valid. As soon as an author creates a work and records it in some form, the work is immediately protected under copyright law. Thus, there is a much larger base and longer history of copyrights in force that might be infringed by the operation of creative machines. Also, unlike patents, copyrights remain in force without an explicit renewal system, making the identification of economically relevant copyrights and their owners particularly challenging.

³ Transaction costs also include non-monetary costs, such as the barrier of accessing the data, which might be encrypted or behind technological firewalls (Cuntz, Fink, and Stamm 2024).

⁴ Patent pools, in which technology owners decide to license their patents to one another or to third parties, exemplify the problem of transaction costs and one solution to deal with it (Lerner and Tirole 2007).

⁵ Much has been written on the differences between patents and copyrights, using more or less sophisticated arguments. For a basic overview of the differences, see, *e.g.*, Stim (2024).

Finally, U.S. copyrights are subject to a legal doctrine of ‘fair use,’ which deems specific categories of unlicensed use of copyrighted material non-infringing (*e.g.*, Samuelson 2015).⁶ The legal criteria for determining if a use qualifies as fair use are not entirely interpretable in economic terms. However, part of the idea is that fair use encompasses uses for which it would be complicated or inefficient to seek permission and uses that do not significantly diminish the economic returns of the copyright holder. As discussed below, part of the argument as to whether creative machines infringe existing copyrights revolves around whether the use of copyrighted materials in the training of the machines does or does not qualify as fair use. There is no corresponding issue with respect to creative machines and the patent system.⁷ In fact, the patent text is typically not subject to copyright restrictions. It can thus be used freely for training purposes.

IV. Non-economic considerations

The above discussion revolves around the role of copyrights and patents in shaping the economic incentives to create and implement or commercialize creations. However, (human) creators may care about controlling the subsequent use of their creations for reasons other than the desire to earn an economic return. The notion that creators should have a fundamental right to control how their works are used is sometimes referred to as creators’ ‘moral rights’ (especially in Europe). Some of copyright holders’ objections to using their works to train creative machines have this flavor—creators just do not like the idea that their works are helping build these machines, and not just because the machines might ultimately reduce their incomes (Peukert et al. 2024). In principle, patent holders might also feel this way about their inventions, but it does not seem to play a comparably significant role in patent policy discussions.

Another relevant consideration is ‘personality rights,’ which are the legal rights that protect an individual’s personal attributes from unauthorized commercial use. Consider *Heart on My Sleeve*, a song with AI-generated vocals made to sound like singers Drake and The Weeknd. The song was written by TikTok user ghostwriter977 and had garnered 15 million views on TikTok before its removal. While the song may be sufficiently different from other songs for there to be no copyright infringement, such uses ultimately affect the appropriability of creative works. (There is a finite demand for Drake and The Weeknd songs.)

The ‘optimal’ IP policy towards creative machines depends on how they develop and interact with markets and other institutions, a point we discuss in the next section. It also depends on the goals and values that we choose as a society and the weights we assign to potentially conflicting goals.

V. What will creative machines actually do?

At this point, we are like people with limited vision assessing an elephant. Blind people touching different parts of an elephant would each have different perceptions of the elephant. While limited

⁶ The situation is somewhat different in Europe, where the legislator has introduced text and data mining exceptions in the Copyright in the Digital Single Market Directive. This directive comes with its own set of issues (Ducato and Strowel, 2021).

⁷ There is no fair-use doctrine for patents, which makes sense given that the patent covers the underlying concept. There is a limited exception for using patented inventions for research purposes (*e.g.*, van Zeebroeck, van Pottelsberghe, and Guellec 2008), but it is not economically significant and has not arisen in discussions about inventions from creative machines.

perspective can lead to a fragmented or incorrect understanding of the issues at stake, the simple framework above suggests some possible conceptions of AI, with implications for how IP policy might adapt in response.

Broadly speaking, generative AI might deliver a) mass piracy machines, b) reductions in the cost (and possibly benefit) of delivering new creative products, c) transformative new products or services with value in excess of what they displace, and d) changed—increased or reduced—the quality of new creative products. We discuss these possibilities in turn, but we note at the outset that these possibilities are not mutually exclusive.

A. Mass-piracy machines

It is increasingly clear that many large language models (LLMs) are trained using copyright-protected material. At one extreme, imagine that LLMs did nothing but deliver chunks of copyrighted material (or minor variations thereof) taken from previous works. In that case, LLMs would not add value to society; they would simply displace revenue from creators to LLM-creating intermediaries. Short of this extreme hypothetical example is an intermediate case in which creative machines displace revenue from the IP they incorporate through a combination of piracy and the creation of new material, as ghostwriter977’s *Heart on My Sleeve* song illustrates.⁸

An obvious analogy is the availability of digital music files for ‘sharing’ via Napster and other unauthorized platforms at the turn of the 21st century. Presumably because the unpaid files were essentially perfect substitutes for the original work, this availability caused a precipitous decline in revenue for music copyright holders.⁹ (Note that, as with Napster, giving pirated content away for free is still piracy, so whether owners of creative machines charge for their output is irrelevant here.)

However, it is not a priori obvious that the use of underlying IP in the development of a creative machine diverts demand from the original IP-protected works. First, some of the users of the creative machines’ output would not have purchased it, so that no revenue from those users is displaced. Second, such ‘piracy’ could, under some circumstances, stimulate demand for the underlying work. The Google Books project digitized thousands of books, some still copyright-protected. These books were fully searchable, but search results delivered only relatively short text snippets. Nagaraj and Reimers (2023) show that the availability of books in the Google Books project stimulated demand for other uses of the books.

B. Cost-reducing technologies for creation

The existence of creative machines significantly lowers the barriers to creation, and this cost reduction takes two possible forms. First, creative machines allow music, drawings, and inventions

⁸ Note that even without piracy (explicitly delivering to LLM users blocks of copyrighted material), the LLMs may still infringe the copyrights on training materials. We discuss this point in Section V.D.1.

⁹ A notable difference between today’s situation and music piracy *à la* Napster is that Napster was created by two teenagers, whereas the training of generative AI models is performed by multi-billion-dollar companies.

to be generated automatically, which leads to a drastic decrease in the cost of creation (or an acceleration in the speed of creation).

Second, creators may be able to use machines to complement the creative process, for example, by creating a list of bullet points and asking an LLM to convert it to flowing text. This may, in turn, allow creators to do the sort of work they had already done more efficiently. Or it may allow creators to develop works of sorts they could not have created earlier, a topic we return to in our discussion of the effects of creative machines on the quality of creative output.

1. Change in the status quo between creators and implementers

It is likely that the skillset required to create artwork will evolve. Craftsmanship and creative abilities are poised to become less important, in contrast to technology proficiency. This observation has two broad consequences. First, as technology becomes easier and friendlier to use, we can expect a greater number of contributors and, therefore, greater fragmentation—perhaps in the form of a longer tail—of creative output. Second, as technology becomes more central in the creation process, technology firms will capture a greater share of the industry’s added value. The digitalization wave in the publishing industry provides a recent example of these trends. The appearance of on-demand book printing and direct publishing to e-readers such as Amazon’s Kindle has empowered a significant number of users to publish books. It has shifted some industry profits from traditional publishers to technology firms like Amazon (Waldfoegel and Reimers 2015). Similar trends occurred in the music industry. Spotify now plays the gatekeeper role, diminishing the influence of traditional major labels (Aguilar and Waldfoegel 2021). To some extent, we can expect dominant AI companies to enjoy significant market power, perhaps up to the point of becoming the creative industry’s new gatekeepers. This expected market power strengthens calls for fair compensation for the original content creators on which generative AI models are trained.

It is also possible that, in some cases, the frontier between intermediate and final output creators will sharpen, leading to a greater division of labor in the industry. For instance, while a larger number of users may be able to compose music or write a movie script, a more limited number of them may have the ability to be live-performing artists or to direct a movie. The contrast might be similarly pronounced for inventive activities, where ‘invention machines’ may allow a larger number of users to produce new inventions; but production, distribution, and marketing capabilities may remain in the hands of a few. Greater disintermediation would reinforce the role of markets for technologies, where inventive and creative content is exchanged, and, therefore, the role of IP underlying these transactions.

2. Reduction in the costs of creation

In theory, reducing the costs of creation can be compensated by lowering the incentives to create provided by the IP system. If the marginal cost of using the machine is close to zero, there would be no need to incentivize its use. Put differently, the cheaper the creation activity becomes, the weaker the case for incentivizing this activity.

However, lack of appropriation of the machine’s output may have some trickle-down effects on the development of the machine itself. Developing a machine is presumably costly, and a lack of IP protection may make the machine’s output hard to sell. This effect shifts the incentive problem upstream, as IP protection on a machine’s output may be needed to incentivize its development.

Just how much this effect will play out depends to some extent on the business model of the machines' creators. For example, delivering generative AI capabilities as a Software as a Service (SaaS) application may guarantee strong enough appropriability since the machine's owner will be able to set the price for its use freely.

Furthermore, as explained in Section II.A, IP rights serve functions other than a reward-based creation incentive. They also incentivize investments in follow-on development activities. Securing the patent right to an active ingredient or the copyright to a movie script enables follow-on investments by implementers, who may not be the original creators. This argument speaks in favor of some form of protection for AI-generated work, especially considering the division of labor in the creative and inventive industries.

3. Potential reduction in the benefits of creation

Although one can expect the cost of machine-made creations to be particularly low, the return may be similarly modest. An abundant supply of creative content may lead to its commoditization, similar to the effect of the Fourdrinier machine in the early 19th century on the cost of paper production (Clapperton 1967). Technological progress gradually transformed paper from a precious, labor-intensive good to a commodity. Expanding on the analogy, LLMs may make creative content worth no more than the paper on which it is printed.

While some creative content may be at risk of commoditization, bringing down its value, there will surely always be some creative content of exceptionally high value. How exactly creative machines will shift the value distribution is unknown, but this shift will have implications for the appropriate IP regime.

Another aspect concerning the benefits of creation relates to the creator's ability to appropriate the returns to its creation. As elaborated in Section V.A., the use of LLMs as a mass-piracy machine may displace revenues from original creators to 'pirates,' possibly as a zero-sum game, and significantly lower the benefits of creation.

C. Will creative machines deliver 'transformative' products and services?

The two mechanisms already discussed—piracy and reduced costs and benefits of creation—fit squarely into a traditional framework for thinking about IP. In order to maintain creative incentives, reduced benefits (especially reduced appropriability)—on their own—require a strengthening of effective protection. And, on their own, reduced costs allow the achievement of pre-existing creative incentives with reduced protection.

But creative machines raise additional possibilities. Much as the bundling of musical recordings into streaming services with value-added curation features created services with more value to consumers than previous musical offerings could deliver, contemporary creative machines may allow for products and services whose value creation exceeds the revenue diverted from existing IP. The question of whether creative machines can increase total social surplus is essential. If the answer is no, there would be no reason to encourage their development. For example, this would be the case if revenue to creative machines is simply diverted from traditional IP owners, with no additional benefit to consumers.

One can imagine many examples of possible surplus expansion via creative machines. Consider the following simple example. Today, a user seeking information issues a search query at Google. This delivers valuable leads, and further searches allow the searcher to find sources with information on what they seek. An LLM, by contrast, might allow the same searcher to get a coherent essay with an organized answer to the question they seek. That is, the LLM might deliver the equivalent of the output produced by a search engine user, along with an analyst and a writer. Or an LLM with access to many new sources could deliver authoritative summaries of the news topics of interest to a particular user.

Our analysis, and much economic discussion of IP, emphasizes that the ownership conveyed by IP is important for incentivizing creation. However, we should note that, particularly in the digital realm, there is a meaningful role for open-source works, where monopoly control is explicitly eschewed to varying degrees. Wikipedia has managed to attract more users than Encyclopedia Britannica despite no ownership of its text. Many LLMs are currently experimenting with varying degrees of free usage. Allowing space for experimentation with different ownership models is another consideration as IP policy towards creative machines evolves.

The transformative potential for creative machines may ultimately manifest in ways we cannot now conceive. Generative AI has the features of what economists call a ‘general purpose technology’ or GPT (Bresnahan and Trajtenberg 1995). GPTs can be used in many different sectors; this widespread use contributes to their continuous improvement and so makes them powerful sources of overall economic benefit. While the GPT nature of generative AI does not suggest obvious implications for IP treatment of creative machines, it does imply that policies that inhibit the development of these technologies may have significant unforeseen adverse consequences.

D. Creative machines and quality of work created

Creative works differ substantially in various quality measures, including both aesthetic quality measures (*e.g.*, acclaim from critics) and the economist’s notion of quality (consumer appeal). Most books sell very few copies, while a small number of books achieve substantial sales and acclaim. The same is true for music, movies, and other creative works. Inventions follow a similar value distribution, with most having a low economic impact and technical merit (Silverberg and Verspagen 2007, de Rassenfosse and Jaffe 2018).

How might creative machines affect the distribution of creative works’ qualities? One possibility is that creative machines function as substitutes for traditional human creations. Creative machines ‘on autopilot’ might produce works with little human input. These works, in turn, might have predictably modest quality. This would presumably have some value to consumers, but creative output of this type would have little chance of diverting substantial interest from human-created fare. Creative machines might also function as complements to human creation, allowing humans to create more quickly and create works of unpredictable—and sometimes high—value to consumers.

1. *Derived output and ‘derivative works’*

We noted above that LLMs engaging in piracy—finding chunks of text or images from their training data to return to users in response to their queries with nothing else added—create

negligible social value. The more prevalent—and more complex to analyze—case is where the creative machine produces new text, music, or images in response to a query. Since the machine is, by construction, an algorithm for recombining its training data, any output it produces is necessarily derived from its training data. There is simply nowhere else from which its output can come.

The derived nature of generative AI's output is not fundamentally different from human output. There are presumably few novelists who have never read a book, nor composers who have never heard a song. In complex ways that cannot be mapped but are nonetheless real, a novelist or composer must, to some extent at least, derive their works from their experience of previous works. Philosophers, neuroscientists, and AI experts can debate whether a human creating a new work using their experience of previous works is doing something that is fundamentally different from what a creative machine does based on its training data. But, on the outside, a human creating something new based at least in part on previous creations they have observed looks very similar to a creative machine producing a new output based on its learning from its training data.

The derived nature of generative AI's output has both economic and legal implications. Legally, a copyright owner has the right to control 'derivative works' based on the copyrighted work, where for that purpose, 'derivative work' is a legal term of art, with case law construing its meaning, but which is nonetheless inherently subjective (*e.g.*, Goldstein 1982, Reese 2008). One can ask a machine to write a novel about children who can do magic. If the characters in the result are named Harry and Hermione, it is probably legally 'derivative.' At the other extreme, there is a very large (if not infinite) number of possible characters and structures for such a novel without apparent relation to any specific work (Karjala 2006). However, many works of magical fiction involving children share common elements to some degree. The task of deciding how much overlap makes a work 'derivative' of a previous work is, in principle, no different for AI-generated novels than for human novels. The only difference is that the generative AI model has 'read' many more novels, and we might have a list of exactly which ones it has read. It is not obvious that this makes the task any easier—or more relevant.

From an economic perspective, one possibility is that the creative machine, overall or in combination, engages in more or less the same derivation process as humans, presumably not producing exactly the same products humans do or could do but perhaps producing the same distribution of products over relevant dimensions. In that case, the problem simplifies to the 'cost reduction' case discussed above, as it seems likely that machines can produce these derivations more cheaply than humans can. The other possibility is that collectively, creative machines produce a body of work that differs from what humans produce or could produce. The rest of this section explores that possibility.

2. *The role of predictability*

Historically, technologies reducing the cost of creation have had significant positive effects on welfare because of the unpredictability of quality at the time of creation. To see why this is so, suppose that the appeal of all book ideas was clear before authors and publishers made any investments. Then, all books with expected—and realized—value to consumers above their creation costs would be greenlighted. A technological change reducing the cost of creation would allow the greenlighting of new works, all of which would have expected—and realized—value

between the old, higher cost threshold and the new, lower one. This would be beneficial to society, but the benefits would be modest.

In reality, the quality of creative works is known to be highly unpredictable at the time of creation (Goldman 1983, Caves 2000). As a result, a reduction in the cost of creation delivers works with lower expected value; but because of unpredictability, it delivers works throughout the realized quality distribution. This is the reason why digitization delivered a digital renaissance (Aguilar and Waldfogel 2018, Waldfogel 2018). Shi and Evans (2023) have reported evidence of the beneficial role of ‘unpredictability’ on technological progress. Drawing on millions of research papers and patents, they show that surprise in terms of unexpected combinations of contents and contexts predicts outsized impact, as measured by the number of citations that these papers and patents attract.

The implications of these ideas for creative machines are clear. If creative machines allow the creation of predictably mediocre works, then the welfare benefits will be small, even putting aside the possibility of undermining revenue opportunities for conventional creators. If creative machines facilitate creative works of unpredictable quality—including high-quality works—then creative machines may have substantial benefits for society beyond mere cost reduction.

This line of reasoning considers that one can compare the quality of machines’ works to that of human creators. However, creative machines could also produce works that are fundamentally different from humans’ creations, *i.e.*, works humans cannot produce. We discuss this point in the next section.

3. *The nature of creations*

Some anecdotal evidence suggests that machine-made work may be as good as human creations. Jason Allen’s piece, *Théâtre d’Opéra Spatial*, was created by the Midjourney AI software and won first place in the digital art category at the 2022 Colorado State Fair’s annual art competition.¹⁰ AI-generated discoveries have already been published in prestigious scientific journals (*e.g.*, Merchant et al. 2023).

If machine-made output is undistinguishable from human-made output and vice versa, the market will not prefer one over the other; and we may as well direct human efforts toward more productive activities. In that case, today’s IP regime, devised for costly-to-create output, might be too strong (bearing in mind that generative AI may leave development costs unchanged). Having made this observation, the following logical questions are: what is the ‘optimal’ level of strength, and what are the features of the appropriate regime? There is also an ethical question behind this observation: Do we want to encourage machine-made creations over human creations as a society?

However, if machine-made output is in some way fundamentally different from human-made output, the right question to ask is how to devise an IP regime that supports both human and machine creations. Anecdotal evidence suggests that some of the machines’ creations have that

¹⁰ See <<https://www.artnews.com/art-news/news/ai-generated-artwork-colorado-state-fair-copyright-decision-1234679341/>>, last accessed on June 15, 2024.

flavor; for instance, Warner Music Group’s initiative to recreate the voice of the late French singer Edith Piaf to narrate the animated film *Edith*.¹¹

Given these unknowns, discussions about changes to the IP system are premature in some sense. This puts policymakers and IP scholars in the uncomfortable position of having to propose changes to the IP system in the face of pressure from lobby groups while being largely ignorant about the fundamental, long-term impact that technology will have on creative and inventive activities.

4. *Different kinds of creative machines*

We have been generically considering creative machines, including LLMs, but there are different types. Gans (2024) notes a distinction between ‘small’ and ‘large’ models; they might alternatively be labeled ‘specialized’ and ‘general’ models. Specialized or small models are designed for a specific purpose, so they are trained on a specific class of data related to that purpose. Examples include models for reading medical images for diagnosis or producing answers to consumer service inquiries. Large or general models are intended to field essentially any kind of inquiry and, hence, are trained on as large and diverse a set of training inputs as possible. Gans theoretically models negotiation between existing copyright holders and model developers over the use of the copyrighted material in model training. He finds that for small models, negotiation between model developers and existing copyright holders with the right to control the use of their works in training leads to socially desirable outcomes. However, for large models, the negotiation is complex and hence may not lead to socially efficient outcomes.

VI. Compensation regimes

The previous section illustrated a range of ways in which creative machines can be used and introduced a series of considerations regarding the nature and effects of their outputs. Given this considerable heterogeneity and uncertainty in outcomes, it becomes apparent that developing an IP regime that best supports a post-generative AI world is a highly challenging task. An important piece in this puzzle will be the compensation program between copyright holders and users, which generates significant tensions. The present section delves into this issue.

The starting point for these tensions is that almost all creative machines are based, at least in part, on training data that included copyrighted material, which was taken and copied without permission from, or compensation to, the copyright holders. Some proponents have suggested that any requirement for such compensation would kill the goose that could lay golden eggs. On the other side, it is unclear whether the owners of copyrighted works would be satisfied with some kind of generic licensing model in which compensation would be paid for training models on copyrighted works, but the owners of individual works might not control whether and in what ways their works are used.

At the outset, we should point out that a critical technological bottleneck limits the fair compensation of the use of copyrighted content. At this stage, it is technically challenging to identify in a systematic or general way the contribution of particular works in the training data to the model parameters. Further, outside of specific cases where a model is asked a question that

¹¹ See <<https://www.billboard.com/pro/ai-bring-back-dead-artists-musicians-estate-managers/>>, last accessed June 19, 2024.

essentially calls for a copy of a particular previous work, it is equally challenging to identify the contribution of particular works in the training data to model outputs. Although efforts are underway in the computer science community (*e.g.*, Haim et al. 2002, Georgiev et al. 2023, Grosse et al. 2023), we are far from industrial applicability.

While several solutions have been proposed (*e.g.*, Henderson et al. 2023, Samuelson 2024), to clarify the issues, we consider a few stark and simple approaches, which are intended collectively to span the range of possible policies. Although we note how each scenario might develop legally, we do not focus on legal or political feasibility because the purpose is to illustrate the issues rather than make actual policy proposals. We first consider approaches that are extremely favorable to either model builders or copyright owners and then consider compromises that might balance the competing interests.

A. Unrestricted access

A first approach involves granting unrestricted access to copyrighted works, without permission or compensation, for the purpose of training algorithms (*e.g.*, Lemley and Casey 2021). Something close to this might result if a Court made a broad finding that model training constitutes ‘fair use’ of the copyrighted material and would resemble the legislation in Japan.¹² Such an approach would satisfy model builders’ desire for maximum flexibility at the expense of depriving historical rights holders of rents that they may have reasonably anticipated receiving and undermining long-held presumptions about how copyright operates.

Note that this approach would not eliminate all rights of existing copyright owners relative to creative machines, only the right relative to the use of training data. As discussed above, a specific output of a creative machine could be said to infringe on a pre-existing work if it were deemed to fall under the ‘derivative work’ doctrine.

B. Strict enforcement

A second approach involves strictly enforcing the right of copyright holders to compensation for copying their works without permission. Such an approach would prohibit both outright piracy, where the originals of works used in training were acquired without permission or compensation, and also would reject the possibility that the ‘fair use’ doctrine covers the reproduction of copyrighted works that occurs during the preparation of materials for training, or in the process of the training itself.

Without getting into subtle legal details, the existence of statutory damages (whereby damages are due for infringement without any requirement to prove actual harm) would create potentially substantial damage liabilities for those who train large models. These damages would potentially bankrupt some current players, and significantly increase development costs for those who could bear them.

This situation would incentivize the owners of creative machines to pursue some settlement, presumably offering significant amounts of money. However, the transaction costs would be high

¹² See <<https://www.privacyworld.blog/2024/03/japans-new-draft-guidelines-on-ai-and-copyright-is-it-really-ok-to-train-ai-using-pirated-materials/>>, last accessed, June 22, 2024.

because machine owners would, in principle, need agreement from many different rights holders. Some rights holders might demand disproportionately large compensation or refuse to grant permission at any price. It would certainly create a great deal of uncertainty while the situation awaited resolution.

C. Grace period

A third approach involves granting amnesty for use without permission up until a certain date, but strictly enforcing the right to control copying going into the future. While it is not possible to identify *ex-post* how particular works contributed to training, it should be technically feasible at some cost to screen *new* data being fed into the training process and include only works for which permission has been secured. If it is not practically feasible to sort out past infringement, perhaps the long-run integrity of the system and incentives for continued human creation could be maintained by pairing some kind of amnesty for past infringement with strict enforcement of the need for permission going forward.

Copyright holders are unlikely to favor this approach, but it is obviously better for them than losing all rights over their works being used in training. And because of the complexity and uncertainty created by strict liability, at least some rights holders would end up better off under this approach than under the strict liability approach. It has been suggested that realizing the full potential of existing models will require training them on significant amounts of new data. If this is correct, the willingness of creative machine owners to pay for the right to use new data could be significant.

This approach reduces uncertainty and transaction costs, which is socially beneficial, but does it in a way that is prejudicial to rights holders relative to historical expectations. This suggests that this approach might be linked to a requirement that to enjoy the benefits of amnesty for any historical infringement, each model developer would have to make some kind of one-time payment into a compensation fund, which would then be distributed in some way among historical rights holders.

The consequences of such an approach would be very heterogeneous across different historical rights holders, with active creators winning and inactive creators of older works losing. Some compensation could mitigate these differences but probably not eliminate them.

D. New statutory blanket license

Congress could create a new blanket license analogous to that administered by SoundExchange for streaming sound recordings and by the Mechanical Licensing Collective for reproductions of musical works in the streaming process (Priest 2021). Copyright holders might dislike this idea; and because there are so many of them, most would likely get very little money. But if the blanket license fee were set to collect significant revenue, popular and successful creators would get a significant new revenue stream.

Implementing such an approach requires the determination of (1) a formula for the royalty owed by model trainers and (2) a formula for the distribution of collected royalties to different rights holders. An obvious candidate structure for both sides of the problem is to use revenue as the scaling variable.

With respect to the setting of the model developers' obligation, basing the royalty on revenue earned implies that developers would only owe royalties once they start selling/licensing some kind of product. Whether this is a feature or a bug depends on one's point of view. Start-up firms generally do not have the luxury of paying for their inputs before earning revenue. Furthermore, because this particular input is intangible and not consumed by use, there is an efficiency argument for allowing developers to escape royalty obligation if they never commercialize.

Because the contribution of particular works to the parameters of the models is hard or impossible to determine, there does not seem to be any obviously 'correct' method for determining the share of royalties collected that is then distributed to each specific rights holder. However, the lack of traceability of the contribution of specific works means that a simple approach based on widely available information will likely be as good as any effort to target the money more accurately. An obvious candidate is to distribute the money based on the works' revenue from sales or licensing. While the contribution of individual works to model training cannot be determined, as a general tendency, more popular works will be more likely to be relevant in creative machines' output. Sales and subscription revenue seem a good proxy for popularity. However, the sheer number of copyright owners (virtually any person or legal entity who has ever published something online) may make this approach impractical. Furthermore, many works are available free of charge—hence, they would receive no compensation for training purposes—but copyright owners may object to receiving no compensation (*e.g.*, Peukert et al. 2024).

On a going-forward basis, it should be possible to exclude works from the blanket license if their owners do not want them to be used for model training, thus respecting the rights of those creators who simply do not want their works used by the machines. A digital watermark could indicate whether the rightsholder has opted out of the blanket license. Model builders would be expected—at least going forward—to exclude these works from training data (though enforcement is challenging), and any works with such watermarks would be ineligible for any share of the royalties distributed.

We note that the blanket licensing model has its own issues and problems. ASCAP and BMI, two of the largest performing rights organizations that license public performances of musical works on behalf of songwriters, have both been subject to antitrust litigation by the Department of Justice for anti-competitive behavior (Einhorn 2000). Much of the data needed for any kind of model for distributing the revenues is proprietary in various ways. Many practical details would have to be worked out, and any approach chosen will be highly imperfect and likely highly controversial. We point out this option only to illustrate that the problem of high transaction costs is inherent in the usage of vast volumes of copyrighted material to train large models, and some mechanisms can be considered to reach some kind of accommodation despite those costs.

E. [Could a solution emerge endogenously in the marketplace?](#)

At the moment, incentives do not seem to be pushing towards a solution. Model developers seem content to keep unclear exactly how copyrighted material is used in training because greater clarity might just increase their liability. If the Courts, Congress, or the Copyright Office moved toward the strict liability option described above, that might create enough pressure for new approaches to emerge (Merges 1996). However, the technical and transactional challenges to crafting a solution under current laws are considerable. Furthermore, there is little doubt that appropriate

regulation may act as a catalyst, much like the ‘safe harbor’ provisions of the U.S. Digital Millennium Copyright Act allowed online platforms like YouTube, Amazon, and Apple to flourish by reducing their liability for user-generated content that might infringe on copyrights (*cf.* Goldman 2015).

VII. Long-run issues

Some of the issues raised by creative machines involve the division of rents from the already-existing pie of creative products. For example, if creative machines create new net value and revenue, how might that revenue be split among parties? However, as with many IP issues, many of the critical questions raised by creative machines are dynamic: How might creative machines affect continued creative activity?

A. The interaction between machine and human creation

There are vigorous debates about the effects of recent technological changes on artists’ ability to support themselves and to continue creation. Musicians and songwriters raise concerns about streaming payments.¹³ Creators of television programs raise similar concerns about deteriorating pay and working conditions in the digital age.¹⁴ Despite concerns about compensation, creation has not ebbed, raising questions about the need for policy action to ensure a sufficient continued supply of new creative products.

New creative machines may be different. It is possible that generative AI will allow a smaller number of creative workers to create many creative products. And while growing numbers of distribution channels have allowed far more songs, movies, and television programs to generate revenue in the digital era, the appetite for more output is probably finite. It is entirely possible that generative AI will displace some creative workers.

Suppose, for example, that ‘autopilot’ creative works appeal to enough users to divert some of their spending from traditional work. It is possible that fewer creators can support themselves from continued creation. This, in turn, has two potential adverse effects. First, the users who prefer traditionally produced works would now face fewer appealing options. Second, there would be a reduced flow of new artistic experimentation, the unpredictable creation that leads to highly valuable new products—assuming experimentation remains the appenage of human creators.

Concerns about keeping human creative workers engaged in creative activity may extend beyond the well-being of current creators and consumers; these concerns may also affect the future value of machine-created or machine-assisted works. Does continued improvement of the machines require continuous training on new material going forward? To some extent, the answer depends on the time-value of data, for which we have little evidence (Valavi et al. 2020). Intuition suggests that the time-value of data for inventive activities is relatively high, meaning that recent inventions are particularly more valuable than past inventions for producing new inventions. We can

¹³ See, *e.g.*, <<https://thetrichordist.com/2012/04/15/meet-the-new-boss-worse-than-the-old-boss-full-post/>> and <<https://stringsmagazine.com/will-streaming-ever-pay-for-musicians/>>, last accessed on June 15, 2024.

¹⁴ See, *e.g.*, <<https://www.nytimes.com/article/wga-writers-strike-hollywood.html>>, last accessed on June 15, 2024.

conjecture that the time-value of data is probably slightly lower in the creative industry (after all, romance novels have been around for a long time). However, to the extent that the creation process is cumulative, there is a need to continuously integrate new data into the training. This need calls for a mechanism to encourage new work to be used for training purposes. The next logical question is the extent to which this training material needs to be created by humans to ensure an original infusion of data into the training or whether machine-made content is equally fit for purpose. Research in computer science suggests that using model-generated content in training causes irreversible defects in the resulting models, a phenomenon known as ‘model collapse’ (Shumailov et al. 2023). Thus, regulation must maintain humans’ creation incentives high enough.

This discussion has focused on the possibility that machine creation will significantly displace human creation. However, humans will interact with machines in ways that are not yet entirely clear. Machines may either substitute for human creators or complement them. Perfect substitution arises if machines can create creative content without human input. Perfect complementarity arises if machines augment human output by reducing costs or increasing quality. An example of the latter is the Vocaloid software, which analyzes vocal data to replicate a singing style and voice in a different language.¹⁵ Such technology allows singers to reach more markets, which we view as a quality improvement.

Where the equilibrium settles on this continuum will affect the degree of commoditization of creative content and, therefore, the nature of the IP regime that best supports the creative industry. Furthermore, the nature of the IP regime itself could affect the degree to which humans are substituted for or complemented by machines. Consider a business facing the decision to distribute tasks among human pharmacologists or channel resources into an AI system for vaccine development. Ideally, the allocation of resources should be guided by efficiency only, with IP rights being neutral to this choice.

B. How will creative machines affect market structure and competition?

Few could have predicted the structure of today’s creative industry at the start of the digitization wave in the 2000s. Predicting changes in the organization of the creative industry brought by generative AI is similarly challenging. We can, however, provide some high-level insights guided by economic theory and observations from the recent past.

Creative machines are a high-fixed-cost and low-marginal-cost technology. Training large generative AI models is incredibly expensive, but running them is comparatively cheap. Industries dominated by such technologies often become highly concentrated. Only a limited number of firms can bear the required investment costs, and recovering these costs necessitates scale to distribute the output to a large number of consumers. This observation suggests that market dynamics may lead to a handful of profitable creative machine producers, leading to a concentrated industry.¹⁶ Perhaps in anticipation of this threat, but certainly also for reasons related to ‘cultural sovereignty,’ many national governments are backing locally-brewed generative AI models that are often open source.¹⁷ How exactly these government-sponsored models will affect market structure is an open

¹⁵ See <<https://www.vocaloid.com/>>, last accessed on June 15, 2024.

¹⁶ See Azoulay et al. (2024) for an in-depth analysis of the impact of generative AI on market structure.

¹⁷ See <<https://www.politico.eu/article/europeans-race-create-artificial-intelligence-chatbots-counter-english-ai/>>, last accessed on June 15, 2024.

question. However, we note that the extent of subsidies and the open-source nature of many models may tilt the balance towards more fragmentation as it lowers the barriers to entry (see Azoulay et al. 2024 for a similar argument). Furthermore, the market needs for small, specialized models versus large, general models is still unclear. However, one can expect that ease of access to training data will affect the supply of small and large models in different ways and, therefore, the market structure.

It is tempting to believe that technological progress and, with it, falling costs of creating the machines will eventually erode the dominant positions of the incumbents. After all, the music industry itself was controlled by a few Major Labels and technological progress has allowed independent artists to make substantial inroads. Two factors may dampen this optimism. First, technological progress has only led to a greater fragmentation of players in the downstream market. The emergence of platforms has kept the upstream market heavily concentrated. Generative AI is a candidate technology for becoming a platform industry. The machine producers would provide the core technology on which a host of developers could build applications like on Apple’s and Google’s app stores. The possibility of machine creators evolving into platforms reinforces the risk of market concentration. Second, the greater the barriers to accessing training data, the more difficult it will be for new entrants to break in.¹⁸ This calls for an IP right model ensuring that the training data are accessible.

VIII. Conclusion

As many current lawsuits attest, creative machines are posing substantial challenges to creators, media companies, and technology companies.¹⁹ Much is at stake, including the income and moral rights of existing creators, the possibility of undermining future creation, and the hastening or slowing of possible benefits from creative machines.

Despite the urgency, we note that there are costs to acting either too late or too soon. On the one hand, clarity about rules may facilitate necessary investment (or inhibit ultimately uneconomic investment). Lack of clarity invites actors ready to “move fast and break things,” which may, in turn—and possibly unfairly—tilt decision-makers in favor of allowing transgressions ex-post. Both concerns favor swift rulemaking. But once made, rules may be difficult to change, and formulating the right rules when the technology trajectory is highly uncertain may lead to rules that turn out to fit the technology poorly.

We are not in a position to advocate the swift adoption of any particular rules. Instead, the framework and considerations we have raised point to important questions that scholars and policymakers must address as they formulate policy.

These questions include:

1. How much do LLMs incorporating IP divert revenue from existing rights holders?
2. How do creative machines affect the cost of developing new creative works?

¹⁸ Besides data access, training costs are another major barrier limiting the development of LLMs. This is another reason why government-sponsored LLMs and national data science infrastructures are gaining traction. See, *e.g.*, the Swiss Data Science Center initiative at <<https://www.datascience.ch/>>, last accessed on June 15, 2024.

¹⁹ See a list at <<https://www.bakerlaw.com/services/artificial-intelligence-ai/case-tracker-artificial-intelligence-copyrights-and-class-actions/>>, last accessed on June 15, 2024.

3. How does creative machine-empowered creation affect the sorts of works produced and their distribution of quality?
4. Do creative machines deliver services that create more value than they divert?
5. Can we design a licensing system that allows the creation of socially valuable creative machines?

These are interesting—and challenging—times for those engaged in creative work, those developing creative machines, the scholars trying to come to grips with these developments, and policymakers. IP policy will shape how technology evolves and how markets develop in response to that evolution. Our current knowledge makes it difficult to confidently choose ‘ideal’ or ‘optimal’ policies to address these issues. Transaction costs and other market imperfections squarely put any analysis of IP and technological change squarely in the realm of the ‘second best.’ Uncertainties about how the technology will evolve, combined with the need for a somewhat forward-looking policy, mean that decisions must be made with limited ability to foresee their consequences. However, systematic efforts to identify costs and tradeoffs can foster sound policy even if they cannot promise optimal policy.

IX. References

- Abbott, R., & Rothman, E. (2023). Disrupting creativity: Copyright law in the age of generative artificial intelligence. *Florida Law Review* 75: 1141–1201.
- Azoulay, P., Krieger, J. L., & Nagaraj, A. (2024). Old moats for new models: Openness, control, and competition in generative AI. *National Bureau of Economic Research Working Paper* No. w32474.
- Peukert, C., Abeillon, F., Haese, J., Kaiser, F., Kiouka, A., & Staub, A. (2024). Strategic Behavior and AI Training Data. Available at <https://doi.org/10.48550/arXiv.2404.18445>.
- Aguiar, L. & Waldfogel, J. (2018). Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music. *Journal of Political Economy* 126(2): 492–524.
- Aguiar, L. and Waldfogel, J., (2021). Platforms, power, and promotion: Evidence from Spotify playlists. *The Journal of Industrial Economics*, 69(3), pp.653-691.
- Arora, A., Fosfuri, A., & Gambardella, A. (2004). *Markets for technology: The economics of innovation and corporate strategy*. MIT Press.
- Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020). Are ideas getting harder to find? *American Economic Review* 110(4): 1104–1144.
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies ‘Engines of growth’? *Journal of Econometrics* 65(1): 83–108.
- Budish, E., Roin, B. N., & Williams, H. (2016). Patents and research investments: Assessing the empirical evidence. *American Economic Review* 106(5): 183–187.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. *National Bureau of Economic Research Working Paper* No. w31161.
- Caves, R. E. (2000). *Creative Industries: Contracts between Art and Commerce*. Harvard University Press, 376 p.
- Clapperton, R. H. (1967). *The Paper-making Machine: Its Invention, Evolution, and Development*. London and New York: Pergamon Press, 365 p.
- Cuntz, A., Fink, C., & Stamm, H. (2024). Artificial Intelligence and Intellectual Property: An Economic Perspective. *World Intellectual Property Organization Economic Research Working Paper Series* No. 77/2024.
- Ducato, R., & Strowel, A. M. (2021). Ensuring text and data mining: remaining issues with the EU copyright exceptions and possible ways out. *European Intellectual Property Review* 43(5): 322–337.
- de Rassenfosse, G., & Jaffe, A. B. (2018). Are patent fees effective at weeding out low-quality patents? *Journal of Economics & Management Strategy* 27(1): 134–148.
- de Rassenfosse, G, Jaffe, A. & Wasserman, M. (2024). AI-Generated Inventions: Implications for the Patent System. *Southern California Law Review* 96:101–126.

- de Rassenfosse, G., Palangkaraya, A., & Webster, E. (2016). Why do patents facilitate trade in technology? Testing the disclosure and appropriation effects. *Research Policy* 45(7): 1326–1336.
- Einhorn, M. A. (2000). Intellectual property and antitrust: Music performing rights in broadcasting. *Columbia Journal of Law & the Arts* 24: 349–375.
- Dell'Acqua, F., McFowland, E., Mollick, E. R., et al. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Management Unit Working Paper* No. 24-013.
- Gans, J. S. (2024). Copyright Policy Options for Generative Artificial Intelligence. Mimeo.
- Gans, J. S., Hsu, D. H., & Stern, S. (2008). The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays. *Management Science* 54(5): 982–997.
- Georgiev, K., Vendrow, J., Salman, H., Park, S. M., & Madry, A. (2023). The journey, not the destination: How data guides diffusion models. arXiv preprint arXiv:2312.06205.
- Goldman, W., (1983). Adventures in the screen trade: A personal view of Hollywood and screenwriting.
- Goldman, E. (2015). How the DMCA's Online Copyright Safe Harbor Failed. *NTUT Journal of Intellectual Property Law and Management* 3(2): 195–198.
- Goldstein, P. (1982). Derivative rights and derivative works in copyright. *Journal of the Copyright Society* 30: 209–252.
- Grosse, R., Bae, J., Anil, C., Elhage, N., Tamkin, A., Tajdini, A., et al. (2023). Studying large language model generalization with influence functions. arXiv preprint arXiv:2308.03296.
- Haim, N., Vardi, G., Yehudai, G., Shamir, O., & Irani, M. (2022). Reconstructing training data from trained neural networks. *Advances in Neural Information Processing Systems* 35: 22911–22924.
- Henderson, P., Li, X., Jurafsky, D., Hashimoto, T., Lemley, M. A., & Liang, P. (2023). Foundation models and fair use. arXiv preprint arXiv:2303.15715.
- Hilty, R. M., Batista, P. H. D., Carls, S., Kim, D., Lamping, M., & Slowinski, P. R. (2021). COVID-19 and the Role of Intellectual Property. *Position Statement of the Max Planck Institute for Innovation and Competition of 7 May 2021*.
- Karjala, D. S. (2006). Harry Potter, Tanya Grotter, and the copyright derivative work. *Arizona State Law Journal* 38: 17–40.
- Lamoreaux, N. R., & Sokoloff, K. L. (2001). Market trade in patents and the rise of a class of specialized inventors in the 19th-century United States. *American Economic Review* 91(2): 39–44.
- Lemley, M. A., & Casey, B. (2021). Fair learning. *Texas Law Review* 99: 743–785.
- Lemley, M. (2024). How Generative AI Turns Copyright Law Upside Down. *Science and Technology Law Review*, 25(2): 21–44.

- Lerner, J., & Tirole, J. (2007). Public policy toward patent pools. *Innovation Policy and the Economy* 8: 157–186.
- Liebowitz, S. J. (1985). Copying and indirect appropriability: Photocopying of journals. *Journal of Political Economy* 93(5): 945–957.
- Mansfield, E., Schwartz, M., & Wagner, S. (1981). Imitation costs and patents: an empirical study. *The Economic Journal* 91(364): 907–918.
- Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., & Cubuk, E. D. (2023). Scaling deep learning for materials discovery. *Nature* 624(7990): 80–85.
- Merges, R. P. (1996). Contracting into liability rules: Intellectual property rights and collective rights organizations. *California Law Review* 84(5): 1293–1393.
- Nagaraj, A. & Reimers, I., (2023). Digitization and the market for physical works: Evidence from the Google Books project. *American Economic Journal: Economic Policy* 15(4): 428–458.
- Priest, E. (2021). The Future of Music Copyright Collectives in the Digital Streaming Age. *Columbia Journal of Law & the Arts* 45: 1–46.
- Reese, R. A. (2008). Transformativeness and the Derivative Work Right. *Columbia Journal of Law & the Arts* 31(4): 2012–2054.
- Samuelson, P. (2015). Possible futures of fair use. *Washington Law Review* 90(2): 815–868.
- Samuelson, P. (2024). Thinking About Possible Remedies in the Generative AI Copyright Cases. *Communications of the ACM*, DOI: 10.1145/3654699.
- Schwartz, D. L., & Rogers, M. (2021). ‘Inventorless’ Inventions? The Constitutional Conundrum of AI-Produced Inventions. *Harvard Journal of Law & Technology* 35(2): 531–579.
- Shi, F., & Evans, J. (2023). Surprising combinations of research contents and contexts are related to impact and emerge with scientific outsiders from distant disciplines. *Nature Communications* 14(1): 1641.
- Shumailov, I., Shumaylov, Z., Zhao, Y., Gal, Y., Papernot, N., & Anderson, R. (2023). The curse of recursion: Training on generated data makes models forget. arXiv preprint arXiv:2305.17493.
- Silverberg, G., & Verspagen, B. (2007). The size distribution of innovations revisited: an application of extreme value statistics to citation and value measures of patent significance. *Journal of Econometrics* 139(2): 318–339.
- Stim, R. (2024). *Patent, copyright & trademark: An intellectual property desk reference*. 18th Edition. Nolo, 704 p.
- Valavi, E., Hestness, J., Ardalani, N., & Iansiti, M. (2020). Time and the Value of Data. *Harvard Business School Working Paper* No. 21-016.
- van Zeebroeck, N., van Pottelsberghe de la Potterie, B., & Guellec, D. (2008). Patents and academic research: a state of the art. *Journal of Intellectual Capital* 9(2): 246–263.
- Waldfoegel, J. (2018). *Digital Renaissance: What data and economics tell us about the future of popular culture*. Princeton University Press, 320 p.

Waldfogel, J. and Reimers, I., (2015). Storming the gatekeepers: Digital disintermediation in the market for books. *Information economics and policy*, 31, pp.47-58.

Wang, W., Gao, G., & Agarwal, R. (forthcoming). Friend or foe? Teaming between artificial intelligence and workers with variation in experience. *Management Science*, available online.