Creative Ownership in the Age of AI*

Annie Liang† Jay Lu‡

July 29, 2026

Abstract

Copyright law asks whether a new work is "substantially similar" to an existing one, but generative AI can closely imitate a creator's style without copying any content—a capability now central to ongoing litigation. We argue that existing definitions of creative ownership are ill-suited to this setting and propose a new criterion: an AI-generated output *depends essentially* on an existing work if it could not have been generated without that work in the training corpus, and it is *permissible* if it does not depend essentially on any existing work. We establish structural properties of the set of permissible outputs and prove a sharp asymptotic dichotomy: when the creative process is light-tailed, dependence on individual works eventually vanishes, so in the long run the criterion places no limits on AI generation; on the other hand, when it is heavy-tailed, individual works remain pivotal and the criterion continues to bind even for arbitrarily large corpora.

* We thank Simon Board, Joshua Gans, and Ariel Procaccia for helpful comments.
† Department of Economics, Northwestern
‡ Department of Economics, UCLA

1 Introduction

Generative artificial intelligence systems now produce novels (Brodsky, 2024), comics (Schor, 2024), and even entire digital persons (Heritage, 2025). Because these systems are trained on large bodies of existing works, they can generate outputs that closely resemble a particular creator's style without reproducing any actual work. For instance, Figure 1 depicts a comic by the cartoonist Sarah Andersen alongside an AI-generated image created using her name as a prompt. The generated image reproduces recognizable features of her style (Andersen, 2022), yet need not constitute infringement under existing U.S. copyright doctrine.

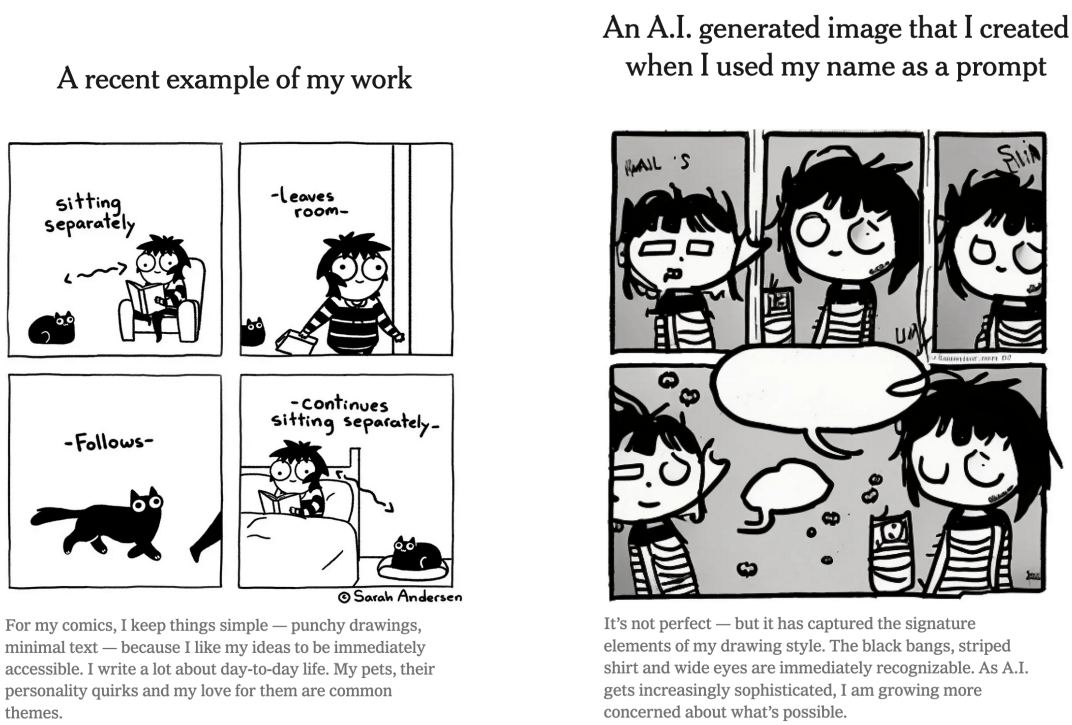

Figure 1: *Left:* an example of Andersen's own work. *Right:* an AI-generated image created when her name was used as a prompt. Images taken from Andersen (2022).

In the past, producing such an image would have required a trained artist familiar with Andersen's work. Generative AI now makes it possible for untrained users to produce such outputs quickly and cheaply, with these outputs possibly competing

directly with the creator's own work.[1] Further complicating the issue, the system's capability to generate these substitutes may arise precisely from training on the works being displaced. Recent litigation—including *Andersen v. Stability AI* and *The New York Times v. OpenAI*—reflects creators' growing discontent not only with the protections they receive, but with the frameworks used to define creative ownership in the first place.

This paper makes three contributions. First, we propose a new criterion for creative ownership: a generated output *depends essentially* on an existing work if it could not have been generated without that work in the training corpus. Second, we establish structural properties of the set of *permissible* outputs—those that do not depend essentially on any existing work. Third, we characterize how the reach of the criterion evolves as the corpus of existing works grows large, and identify properties of the creative process that shape this.

In our framework, described in Section 2, a *creation*—a novel, image, song—is represented as a vector of features, i.e., a point in $\mathbb{R}^d$, and a corpus is a set of such points. We model an AI system in reduced form as a *generator*: a map that assigns to each corpus the set of creations that can be produced from it. This abstraction sets aside the internal mechanics of the system and isolates what our analysis requires, which is how the training corpus shapes which outputs can be generated. The *violation set* consists of the generable outputs that depend essentially on at least one existing work. The *permissible set* consists of the remaining generable outputs, which are robust to the removal of any single work from the corpus.

Section 4 establishes basic properties of the permissible set. We first provide a weak sufficient condition for nonemptiness of the permissible set based on the Radon number from convex geometry. We then show that the permissible set expands mono-

[1] Andersen writes: "The notion that someone could type my name into a generator and produce an image in my style immediately disturbed me. This was not a human creating fan art or even a malicious troll copying my style; this was a generator that could spit out several images in seconds. With some technical improvement, I could see how the process of imitating my work would soon become fast and streamlined." (Andersen, 2022)

tonically as the corpus grows, where incorporating a previous violation into the corpus strictly expands the permissible set, while incorporating a permissible output leaves it unchanged. Finally, we show that the permissible set is closed under generation, so that combining permissible outputs cannot produce a violation.

Section 5 characterizes the permissible set for large corpora. We model the creative process as follows: in each period, a creator selects a direction from a fixed finite set and innovates a random distance in that direction. For each period, we consider the corpus consisting of the works generated thus far. Our measure of interest is the *permissible ratio*, defined as the fraction of generable outputs that are permissible. We study the limiting behavior of the permissible ratio as the corpus grows large. We show that if innovation is light-tailed, then individual works eventually become replaceable. For a broad class of *convex-valued generators*, i.e. generators that produce convex generable sets, the permissible ratio converges to one (Theorem 1). By contrast, if innovation is heavy-tailed, large proportional breakthroughs recur forever, and the permissible ratio falls arbitrarily close to zero infinitely often (Theorem 2). Whether creative ownership (by our definition) retains force in the age of AI thus depends on the statistical character of human creativity: in fields that advance gradually, our criterion eventually imposes no constraints, while in fields that advance through outsized breakthroughs, it has substantial bite no matter how large the corpus grows.

The remainder of this introduction reviews the legal and technological context and the related literature. Section 2 introduces the framework, Section 3 defines our criterion, and Sections 4 and 5 present the main results. Section 6 extends the analysis to breakthroughs that open new creative dimensions and to the protection of groups of works.

1.1 Legal and Technological Context

Existing copyright law was not developed with AI systems in mind but nonetheless shapes how those systems are treated today. In most jurisdictions, copyright gives

authors exclusive rights over reproduction, distribution, and the creation of derivative works (United States Congress, 1976a). These rights potentially constrain both the inputs to generative AI systems (the works on which they train) and their outputs (the works they generate). We review the two constraints in Sections 1.1.1 and 1.1.2, and explain why each fits generative AI imperfectly, motivating consideration of a new criterion.

1.1.1 Training on Copyrighted Works

Training an AI system on copyrighted works can constitute infringement under U.S. law, with exceptions governed by the fair use doctrine (United States Congress, 1976b). This doctrine considers factors such as whether the use is transformative, and whether it harms existing or potential markets for the original works (Supreme Court of the United States, 1994; United States Court of Appeals for the Second Circuit, 2015).[2]

There are several reasons why this doctrine may fit AI generation poorly (Zhang, 2025). First, critics argue that it unfairly constrains AI systems, since human innovation typically depends on learning from copyrighted material. Second, not all creations we would like to protect are copyrighted—for example, a generated character may appropriate an existing person's personality or likeness, neither of which copyright protects. Finally, an AI system need not train directly on a copyrighted work to capture its distinctive features; in many cases, training on summaries or adaptations is sufficient, placing the relevant activity outside the doctrine's reach.

1.1.2 Generation of "Substantially Similar" Outputs

Copyright law also restricts the generation of outputs that are "substantially similar" to an existing copyrighted work (United States Court of Appeals for the Second

[2] In prior cases, courts have found large-scale copying to be fair use when the copied material is not meaningfully disclosed to the public and the copying serves a distinct analytical function, such as search or indexing (Leval, 1990; United States Court of Appeals for the Second Circuit, 2015).

Circuit, 1946). The law has historically imposed a narrow definition for similarity, drawing a line between *expression*, which is protectable, and *creative style*, which is not (17 U.S.C. §102(b); Supreme Court of the United States, 1879; Hand, 1930). This choice reflects a longstanding policy objective to preserve creative freedom and downstream innovation (Leval, 1990). As a result, mimicking style without reproducing content does not necessarily amount to infringement, even when that style is distinctive or commercially valuable (United States Court of Appeals for the Ninth Circuit, 2003; Sag, 2024).

Because generative AI substantially reduces the difficulty of imitating style, critics argue that the expression–style distinction no longer offers creators sufficient protection. Defenders respond that protecting style would be difficult to enforce and could give creators control over broad creative domains (Leval, 1990; Sag, 2024). Our criterion offers a path forward by asking not whether a generated output *resembles* an existing work, but whether it could have been generated without that work.

1.2 Related Literature

Our paper contributes to a growing body of work in economics and computer science on generative AI and intellectual property. A central open question in this area is "the extent to which [the] outputs [of creative machines] might be legally 'derivative' of the copyrighted works used to train the machine" (de Rassenfosse et al., 2024). We propose an answer, where an output is derivative if it could not have been generated without the prior work in the training corpus.

More broadly, our paper relates to models of copyright policy for generative AI (see Lutes (2025) and de Rassenfosse et al. (2024) for excellent surveys). Early works include Gans (2024)—which studies how copyright policy should account for contracting frictions, distinguishing settings where training data can be licensed ex ante from larger-scale settings where such licensing is not feasible—and Yang and Zhang (2025)—which studies how fair-use rules for training data and copyright protection

for AI-generated outputs jointly shape incentives to create new data and improve model quality. Our contribution is complementary to these papers: instead of deriving optimal copyright policy in a model of creative production, we define when an AI output depends on a protected work. Most closely related, Dai et al. (2026) show that existing copyright regimes can create poor incentives for human creation, with the largest distortions falling on the most original creators. Although we do not model creative incentives explicitly, our search for a new definition of creative ownership is motivated by a similar concern with incentives for creation. The main difference in our analysis is that we consider the derivativeness of generated outputs, while Dai et al. (2026) focus on rights over works used as training inputs.

Our approach also relates to several ideas from machine learning. The Shapley value of a training example is its average marginal contribution across all subsets of the training data (Shapley, 1953; Ghorbani and Zou, 2019). Our definition differs in considering only counterfactuals in which the training example is removed from the full corpus (rather than from every set to which it belongs). Our criterion reflects our focus on whether a generation depends essentially on an existing work, rather than on how creative credit should be divided across multiple predecessors. Shapley-style attribution may still be useful for compensation or revenue sharing once dependence has been established.

A related notion is differential privacy, which requires an algorithm's output distribution to change only slightly when any single input is removed or replaced, *uniformly* over possible datasets (Dwork et al., 2006). Counterfactual credit attribution applies a similar test for attribution rather than privacy (Livni et al., 2024; Cohen and Zhang, 2026). Our criterion differs in asking whether a particular work is pivotal within a particular corpus, rather than evaluating whether this pivotality holds uniformly across datasets. Interestingly, the sensitivity to individual inputs that differential privacy seeks to minimize is precisely what makes a work protected in our framework.

Our long-run results connect to a broader literature on creative processes. Em-

pirically, the returns to creative work often appear heavy-tailed: a small number of creative products account for a large share of revenue, citations, and attention (Rosen, 1981; Chung and Cox, 1994; De Vany and Walls, 1999).[3] Several theories also explain why creative possibilities may expand in a heavy-tailed or high-dimensional way. For example, Weitzman (1998) models idea generation as recombinant expansion, which can produce growth exceeding exponential rates, while Lutes (2025) argues that AI tools may amplify variance in human creative output.

Finally, our paper relates to work on how AI systems acquire and displace human expertise. Garicano and Rayo (2025) study apprenticeship when AI can perform some of the routine work traditionally done by junior workers. In their model, AI can undermine training by reducing the value of the tasks through which apprentices help pay for learning. Cullen et al. (2026) study workers whose behavior and documentation become training data for AI systems that can later perform similar work. In their model, workers may withhold knowledge because supplying it improves the firm's future outside option. Creative works raise a similar problem: like worker expertise, a work that enters the training corpus enables the system to produce substitutes for it afterward. Our focus in the present paper, however, is on how to define dependence rather than on the strategic incentives this creates.

2 Model

2.1 Setup

A *creation* is a vector $x \in \mathbb{R}^d$. A *corpus* is any set of creations $C \subset \mathbb{R}^d$, representing a body of existing works. Examples include:

Example 1 (Novels)**.** Large language models trained on corpora of published books can generate new books. A recent surge in AI-generated books led Amazon's Kindle

[3] Drugov and Ryvkin (2020) show that whether innovation is light- or heavy-tailed has sharp implications for the design of contests and incentives.

store to impose limits on the number of books that could be self-published on the platform each week (Brodsky, 2024).

Example 2 (Cartoons)**.** Image generators such as Stable Diffusion are trained on billions of images scraped from the web (assembled in datasets such as LAION-5B). In *Andersen v. Stability AI* (2023), artists allege that these models can reproduce their distinctive styles without authorization.

Example 3 (Actors)**.** Companies such as Synthesia and HeyGen train models on video recordings of human performers to generate realistic digital avatars. The use of such likenesses was a central issue in the 2023 SAG-AFTRA strike and subsequent negotiations over AI provisions in performer contracts (SAG-AFTRA, 2024). The controversy has intensified with the emergence of fully AI-generated performers such as Tilly Norwood, an actress created by the AI talent studio Xicoia (Heritage, 2025).

We model AI systems in reduced form as a technology for producing new creations from existing corpora.[4] Formally, a *generator* is a map from the set of all possible corpora $\mathcal{C}$ into itself, which satisfies three properties.

Definition 1. The map $g : \mathcal{C} \to \mathcal{C}$ is a *generator* if for all $C, D \in \mathcal{C}$ it satisfies:

(1) Preservation: $C \subseteq g(C)$

(2) Monotonicity: $C \subseteq D$ implies $g(C) \subseteq g(D)$

(3) Idempotence: $g\left(g\left(C\right)\right) = g\left(C\right)$

We interpret $g(C)$ as the capability frontier induced by corpus C, i.e., the set of creations the technology can produce from that corpus. Preservation says that the original works themselves remain feasible.[5] Monotonicity says that adding works may

[4] We abstract from other details, including the architecture of the system, the prompts used to elicit outputs, and the stochastic process by which outputs are sampled. This abstraction is useful because our central object is not the internal mechanism of generation, but the dependence of generable outputs on the underlying corpus.

[5] We have in mind the system's raw capability, prior to any explicit restrictions on reproducing existing works, such as filters imposed to comply with current copyright law.

expand the frontier, but cannot eliminate previously feasible outputs.[6] (If $C \subseteq D$, then a technology with access to D can reproduce any use of C by ignoring the additional works.) Idempotence says that running g on its own outputs does not produce new creations. A more primitive generator may lack this property, and in that case we iterate it and define $g(C)$ as the set of all outputs reachable under repeated application.[7] Maps satisfying these three properties are known as *closure operators*.

The following examples illustrate the range of technologies that satisfy this definition.

Example 4 (Convex Hull Generator)**.** The AI system generates creations by taking convex combinations of existing works. The *convex hull* generator

$$g_{conv}(C) := \operatorname{conv}(C)$$

maps each C to its convex hull. See Figure 2 for an example.

Example 5 (Splice Generator)**.** The AI system generates creations by recombining features from existing works:

$$g_{splice}(C) := \left\{x \in \mathbb{R}^d : \forall k = 1, \ldots, d, \text{ there is a } y \in C \text{ s.t. } x_k = y_k\right\}.$$

We call this the *splice* generator.

Example 6 (Box Generator)**.** The AI system has access to both the convex hull and splice technologies. The *box* generator composes the two maps:

$$g_{box}(C) := g_{conv}(g_{splice}(C)).$$

For a corpus C, $g_{box}(C)$ is the smallest axis-aligned box containing C, hence the name.

[6] Monotonicity describes the capability frontier, not any particular training run: with a fixed architecture and imperfect optimization, adding data need not improve every realized model.

[7] If iteration never stabilizes, we take $g(C)$ to be the set of outputs reachable in finitely many steps. When this set is all of $\mathbb{R}^d$, the technology is effectively a universal generator (Example 7); our interest is in the case where the capability frontier is bounded.

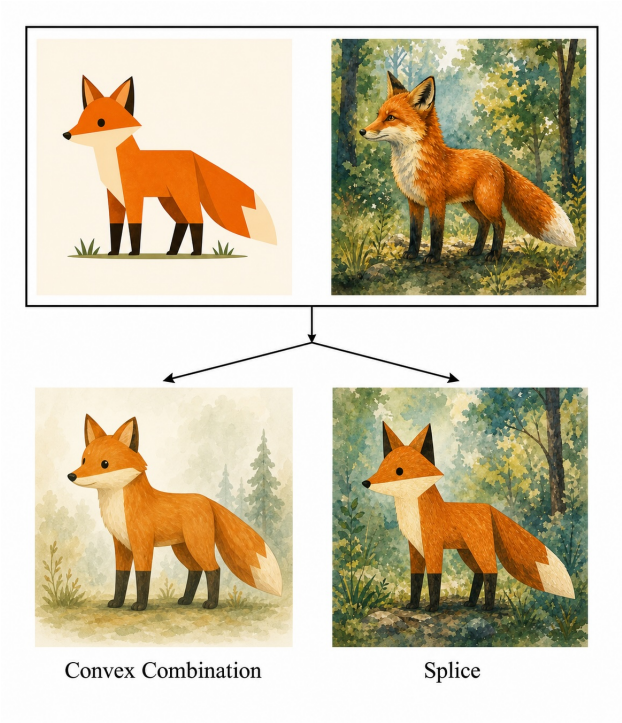

Figure 2: Illustration of the convex hull and splice generators. The images at the top represent two works in the corpus. The lower-left image is an example convex combination—its features lie in between those of the two source images—and the lower-right image is an example splice—some features are inherited from one source image and some from the other.

Example 7 (Universal Generator)**.** Fix a nonempty bounded corpus $\underline{C} \subseteq \mathbb{R}^d$. Suppose that once the AI system can produce this core corpus, its generative capacity becomes unlimited.[8] The *universal generator* is given by

$$g_{univ}(C) := \begin{cases} \mathbb{R}^d & \text{if } \underline{C} \subseteq \text{conv}(C), \\ \text{conv}(C) & \text{otherwise.} \end{cases}$$

The convex hull, box, and universal generators all output convex sets and are thus examples of what we will call convex-valued generators.

Definition 2. A generator g is *convex-valued* if $g(C)$ is convex for all $C \in \mathcal{C}$.

[8] One could interpret this as an example of the infinite monkey theorem or Artificial General Intelligence (AGI).

The following result characterizes convex-valued generators.

Proposition 1. *For any generator g, the following are equivalent:*

(1) *g is convex-valued*

(2) $g = g_{conv} \circ g$

(3) $g = g \circ g_{conv}$

Thus, assuming that g is convex-valued is equivalent to assuming that the AI system has access to convex interpolation, possibly alongside other generative technologies. This is a modeling abstraction rather than a literal description of AI systems. It is natural in applications where intermediate points in the chosen representation have meaningful interpretations, as in some forms of image morphing. It is not appropriate for domains such as proof generation, where interpolating between two valid outputs need not yield another valid output. We impose convex-valuedness for some, but not all, of our results, and discuss its role further in Section 5.4.

2.2 Discussion

Interpretation of the space. Our use of a feature vector to represent creative works is standard in modern machine learning, where text, images, audio, and video are routinely mapped to learned embeddings or latent spaces (Bengio et al., 2013; Mikolov et al., 2013; Radford et al., 2021; Rombach et al., 2022). While we might loosely think of these features as attributes such as subject matter, tone, or visual style, they are often not interpretable in practice. To apply our criterion (Section 3), the analyst does not need to know the feature space or the generator g, if they have access to the generative system itself. Our results in Section 5 do require the analyst to know that g satisfies certain qualitative properties in this space, such as convex-valuedness or continuity, which we will subsequently define.

Scope of the framework. A generator can be interpreted broadly as a description of any creative process, including human creation. For example, collaboration between two people could be represented as the composition of their individual generators, and human–AI collaboration as the composition of a human generator and an AI generator.[9] Since our goal is to define which AI-generated outputs are permissible, we interpret g as an AI system. This focus matters for the criterion in Section 3, which asks what remains generable when particular works are removed from the training corpus. Such counterfactual variation is natural for AI systems, whose training data can be manipulated, but has no direct analogue for human creators.

Probabilistic generation. Many generative AI systems are stochastic, with a corpus and a prompt inducing a distribution over possible outputs. One can interpret $g(C)$ as the set of outputs that the system produces with non-negligible probability across a class of reasonable prompts.[10] This binary formulation is convenient for our motivating legal question, since it allows us to classify outputs as either permissible or impermissible. Section 3.1 discusses various cardinal measures of influence that we can obtain by relaxing our definition of g.

3 Permissible Generation

We define when it is that a generated output *depends essentially* on an existing work. To formalize this, we consider the counterfactual question of whether the output would still be generable without that existing work in the corpus.

Definition 3. Fix a generator g and an existing work x in a corpus C. The *x-permissible set* is

$$p_g(x, C) := g(C \setminus \{x\})$$

[9] The composition of two commutative generators is also a generator.

[10] In principle the support of this distribution may be the entire space of works; we disregard outputs that arise only with negligible probability.

and the *x-violation set* is

$$v_g(x, C) := g(C) \setminus p_g(x, C).$$

An output y *depends essentially* on x in C if $y \in v_g(x, C)$.

Intuitively, the x-permissible set consists of all outputs that would remain generable if x were removed from the corpus. In contrast, outputs in the x-violation set are generable only because x is present: they could not have been produced had x never entered the corpus. These are interpreted as the outputs over which the creator of x has a claim.

Definition 4. Fix a generator g and corpus C. The *permissible set* is

$$p_g(C) := \bigcap_{x \in C} p_g(x, C)$$

and the *violation set* is

$$v_g(C) := \bigcup_{x \in C} v_g(x, C).$$

The permissible set consists of the outputs generable from C that would remain generable after the removal of any single existing work; we call such outputs *permissible*. All remaining generable outputs depend essentially on some existing work and constitute the violation set.

Depending on the application, the consequence of being in the violation set will be different. For example, if the creative works are academic papers, dependence might inform whether a work should be cited as a predecessor. In other applications, it could determine whether a creator can claim a share of the rents generated by the downstream creation. And in the strongest case, it may be prohibited for AI firms to generate works in the violation set.

Figure 3 illustrates the permissible set for an example corpus and two generators. Given existing works $x_1, x_2, \ldots, x_5$, any point in the convex hull of these works is generable under g_{conv}. Panel (a) splits this generable set into the outputs that depend

essentially on x_1, namely $v_g(x_1, C)$, and those that do not, namely $p_g(x_1, C)$. Panel (b) depicts the permissible set: the generable outputs that do not depend essentially on *any* existing work. In this example, the permissible set is a strictly interior subset of the generable set. This property does not hold in general: Panels (c) and (d) show that under the box generator, the extreme point x_4 is permissible given the same corpus.

In general, the structure of the permissible set depends both on how the generator aggregates the corpus and on which other existing works are present. Section 4 establishes basic properties of the permissible set, and Section 5 studies the permissible set for large corpora.

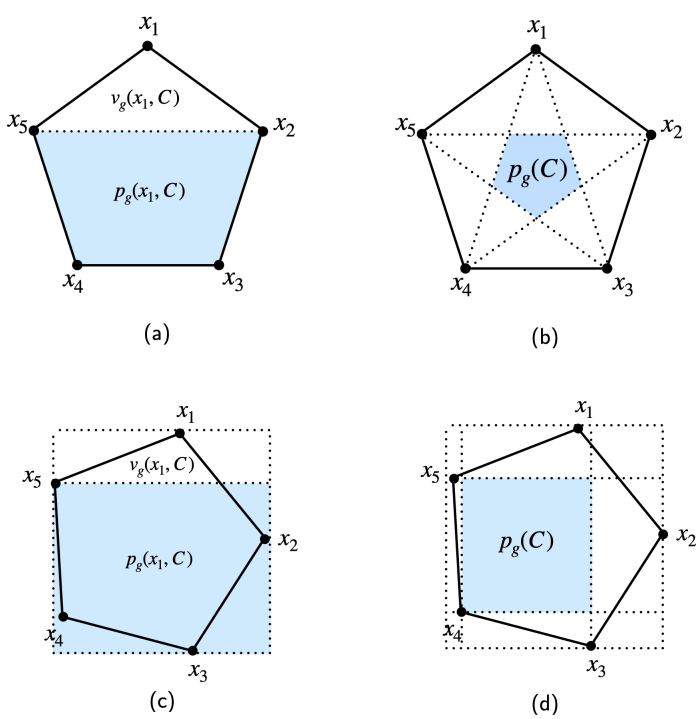

Figure 3: Let $C = \{x_1, x_2, \ldots, x_5\}$. Panels (a) and (b) consider the convex hull generator. Panel (a) shows the set of generable outputs that do and do not depend essentially on x_1. Panel (b) shows the permissible set $p_g(C) = \bigcap_{i=1}^{5} p_g(x_i, C)$. Panels (c) and (d) repeat the same construction for the box generator.

3.1 Discussion

Restrictiveness of the criterion. Our definition of permissibility is neither uniformly more nor less restrictive than the two copyright standards reviewed in Section 1.1. Consider first the restriction on training on copyrighted works (Section 1.1.1). Our criterion is more permissive in that it allows copyrighted works to enter the corpus, so long as no output depends essentially on any one of them, and more demanding in that it has bite even for systems that were not trained on any copyrighted work. Relative to the restriction on outputs—that generated works may not be "substantially similar" to an existing copyrighted work (Section 1.1.2)—our criterion is more demanding in that it can classify outputs as violations even when they do not closely resemble any existing work, and more permissive in that close resemblance need not imply a violation when the output could have been generated without that work. We view our criterion as a complement to existing copyright doctrines rather than a substitute for them; effective regulation may require enforcing multiple standards simultaneously.

Implementation. Unlike "substantial similarity," our definition of permissibility is a formal criterion that can in principle be checked. The brute-force approach of retraining on $C \setminus \{x\}$ is computationally infeasible, but there are comparably demanding notions in computer science that are routinely approximated. For example, Shapley values require evaluating outcomes across exponentially many subsets, while counterfactual credit attribution requires stability uniformly over neighboring corpora. Our criterion could likewise be approximated using techniques from machine unlearning, which estimate the effect of removing individual data points from a trained model (Bourtoule et al., 2019; Xiong et al., 2024), or influence functions, which estimate a training example's marginal effect on model predictions without retraining (Koh and Liang, 2017).

An alternative approach is to treat permissible generation as an evidentiary ques-

tion and rely on qualitative proxies. For example, one could compare outputs across independently trained systems. If models trained without the plaintiff's work nonetheless produce outputs similar to a claimed violation, this suggests that the plaintiff's work was not in fact essential to generating that output. Conversely, if comparable outputs consistently fail to appear in systems trained without the plaintiff's work, this is evidence that the work played an important role.

Cardinal measures of influence. Our criterion is binary: an output either depends essentially on some existing work or it does not. There are various ways to extend it to a cardinal measure of how much an existing work influences the generability of an output.

One such approach is as follows: Suppose a corpus C induces a distribution μ_C over outputs. Given a neighborhood size $\varepsilon > 0$ and probability threshold $\delta > 0$, define the map

$$g_{\varepsilon,\delta}(C) := \{y : \mu_C(B_\varepsilon(y)) \geq \delta\}$$

that assigns to each corpus the set of outputs whose ε-neighborhood receives probability at least δ. The definitions of the permissible and violation sets then apply verbatim with $g_{\varepsilon,\delta}$ in place of g.[11] That is, y is x-permissible if its ε-neighborhood remains sufficiently likely to be generated even after x is removed from the corpus, i.e., $\mu_{C\setminus\{x\}}(B_\varepsilon(y)) \geq \delta$. Then, holding ε fixed, one can measure how much removing x lowers the probability of generating an output near y. And in the other direction, holding δ fixed, one can measure how much removing x increases the neighborhood size needed to generate y with probability at least δ. We leave the study of such influence measures to future work.

[11] The new map $g_{\varepsilon,\delta}$ need not satisfy our three generator properties, so the results that require all three simultaneously would not immediately extend.

4 Properties of Permissible Generation

4.1 Nonemptiness of the Permissible Set

We first consider the question of existence: under what conditions is the permissible set nonempty? In general, it may be that *all* generable outputs depend essentially on some existing work, as illustrated in the following example:

Example 8 (An Empty Permissible Set)**.** Let $d = 1$ so that each creation is a point on the real line. The corpus is $C = \{x_1, x_2\} = \{0, 1\}$ and the generator is the convex hull generator $g := g_{conv}$. Then $p_g(x_1, C) = \{1\}$ while $p_g(x_2, C) = \{0\}$, so the permissible set is empty: $p_g(C) = p_g(x_1, C) \cap p_g(x_2, C) = \varnothing$. With only two works, each is pivotal: removing either one collapses the generable set to the single remaining work.

We identify a sufficient condition for the permissible set to be nonempty, based on the Radon number of the generator, defined as follows.

Definition 5 (Radon Number)**.** The *Radon number* of a generator g, denoted $R(g)$, is the smallest integer k such that any corpus C of size at least k contains two disjoint and nonempty subsets $A, B \subseteq C$ satisfying

$$g(A) \cap g(B) \neq \varnothing.$$

If no such k exists, set $R(g) = \infty$.

Proposition 2. *Fix a generator g. If $|C| \geq R(g)$, then the permissible set is nonempty: $p_g(C) \neq \varnothing$.*

To see this, note that any output in $g(A) \cap g(B)$ can be generated in at least two ways—from the set A or B. Removing any single work from the corpus leaves at least one of these sets intact, so the output remains generable.

The Radon number is a standard complexity measure in convex geometry. For the convex hull generator in $\mathbb{R}^d$, Radon's theorem establishes that $R(g_{conv}) = d + 2$,

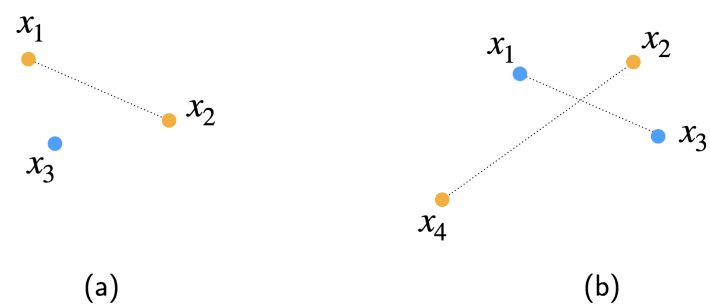

Figure 4: Let $d = 2$, so that creations are points in $\mathbb{R}^2$. *Panel (a):* No three points in general position can be separated into disjoint subsets whose convex hulls overlap, so the Radon number is at least four; *Panel (b):* Every set of four (or more) points can be separated in this way; for example, in the figure let $A = \{x_1, x_3\}$ and $B = \{x_2, x_4\}$.

as illustrated in Figure 4. Since every convex-valued generator g satisfies $g_{conv}(C) \subseteq g(C)$ (Corollary A.1), its Radon number is at most $d + 2$, which yields the following corollary.

Corollary 1. *Fix a convex-valued generator g. If $|C| \geq d + 2$, then the permissible set is nonempty: $p_g(C) \neq \varnothing$.*

In practice, we expect relevant corpora to contain more than $d+2$ works, and thus view the existence of permissible generation as typical in any application of interest. The substantive question is how large the permissible set is, which we take up next for a fixed corpus and, in Section 5, for a growing one.

4.2 Structure of the Permissible Set

We next establish two structural properties of the permissible set: it expands monotonically as the corpus grows, and is stable under further generation.

Proposition 3. *For every generator g and corpora C, D, the permissible set satisfies:*

(1) *Monotonicity: If $C \subseteq D$ then $p_g(C) \subseteq p_g(D)$*

(2) *Stability: $g(p_g(C)) = p_g(C)$*

Part (1) studies the effect of expanding the set of existing works. Because the generator is monotone, a larger corpus expands the set of generable outputs. At the same time, it introduces new works whose rights must be respected, potentially tightening the constraints on what can be produced. The result shows that the first effect dominates: expanding the corpus weakly enlarges the permissible set. Intuitively, if an output can be generated in a way that does not rely on any particular existing work, then adding further works cannot introduce such a dependence. For example, if an AI novel generator trained on 19th-century novels produces an output that does not depend essentially on any one of them, then adding 20th-century novels to the training corpus will not make this output depend essentially on any novel in the expanded corpus.

Part (2) says that the permissible set is closed under the generator: applying g to permissible outputs yields only permissible outputs. This follows because if every input to the generator remains generable without x, then so too must the output. Closure thus captures the natural principle that infringement cannot arise from combining non-infringing outputs.

The next result sharpens Part (1) by describing what determines whether the permissible set strictly expands.

Proposition 4. *Fix a generator g. For a new work $x \notin C$,*

(1) *If $x \in p_g(C)$, then the permissible set remains unchanged: $p_g(C) = p_g(C \cup \{x\})$.*

(2) *If $x \in v_g(C)$, then the permissible set strictly expands: $p_g(C) \subsetneq p_g(C \cup \{x\})$.*

(3) *If $x \notin g(C)$, then the effect on the permissible set is ambiguous: $p_g(C) \subseteq p_g(C \cup \{x\})$, with equality or strict inclusion both possible.*

Part (1) says that adding an already-permissible output to the corpus leaves the permissible set unchanged. Part (2) says that adding a *violation* strictly expands the permissible set. Finally, Part (3) says that the effect of adding a work that was not generable from the original corpus is ambiguous, with Example 9 illustrating both

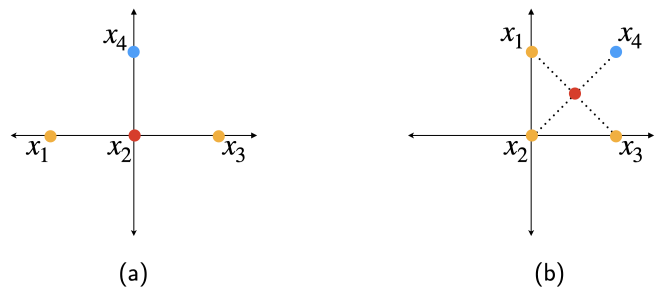

Figure 5: Two examples illustrating how adding a work can affect the permissible set under the convex hull generator. In Panel (a), adding x_4 leaves the permissible set unchanged: $p_g(\{x_1, x_2, x_3\}) = \{x_2\} = p_g(\{x_1, x_2, x_3, x_4\})$. In Panel (b), adding x_4 strictly expands the permissible set: $p_g(\{x_1, x_2, x_3\}) = \varnothing$ while $p_g(\{x_1, x_2, x_3, x_4\})$ is the singleton intersection point of the line segments $\overline{x_1x_4}$ and $\overline{x_2x_3}$, as indicated in red.

possibilities.

Example 9. Let $d = 2$ so that creations are points in $\mathbb{R}^2$.

(1) The corpus is $C = \{x_1, x_2, x_3\} = \{(-1, 0), (0, 0), (1, 0)\}$ and the generator is the convex hull generator. Consider the addition of the work $x_4 = (0, 1)$, as depicted in Panel (a) of Figure 5. Then

$$p_g(C) = p_g(C \cup \{x_4\}) = \{(0, 0)\}$$

so the inclusion of x_4 does not expand the permissible set.

(2) The corpus is $C = \{x_1, x_2, x_3\} = \{(0, 0), (0, 1), (1, 0)\}$ and the generator is the convex hull generator. Consider the addition of the work $x_4 = (1, 1)$, as depicted in Panel (b) of Figure 5. Then

$$p_g(C) = \varnothing \subsetneq \{(1/2, 1/2)\} = p_g(C \cup \{x_4\})$$

so the inclusion of x_4 strictly expands the permissible set.

5 Permissible Generation for Large Corpora

The previous section characterized permissibility for a fixed corpus. We now ask how permissibility changes as creative works accumulate. If an AI system is trained on a sufficiently large corpus, will individual creators retain significant ownership rights under our criterion, or will all generable outputs eventually become permissible?

The answer depends on the process of creative innovation, which we model in a reduced-form way in Section 5.1. We distinguish between two regimes. Section 5.2 studies *light-tailed* innovation, in which new works advance the frontier only gradually. In this regime, we show that every individual work eventually contributes only marginally to the corpus, and the permissible set eventually approximates the entire generable set. Section 5.3 shows that the conclusion is starkly different under *heavy-tailed* innovation. In that regime, the fraction of generable outputs that are permissible not only fails to converge to one, but falls arbitrarily close to zero infinitely often.

5.1 Creative Process for Corpus Formation

As before, the ambient creative space is $\mathbb{R}^d$. There is a finite set of $m \geq d$ unit vectors $U := \{u_1, \ldots, u_m\} \subseteq \mathbb{R}^d$ representing creative directions, which we assume spans $\mathbb{R}^d$.[12] Associated with each direction k is a distribution F_k over positive innovation magnitudes. At date n, a direction $I_n \in \{1, \ldots, m\}$ is drawn, possibly as a function of the entire history up to that date. This allows creators to choose which direction to pursue based on the existing body of work. Conditional on $I_n = k$, a magnitude $R_n > 0$ is independently drawn from F_k, and the resulting work is $X_n = R_n u_k$. Thus, I_n determines the kind of creative advance, while R_n determines how far the new work advances in that direction. We assume that each F_k is atomless and has unbounded support, so that works almost surely do not coincide and there is no exogenous upper

[12] That is, the linear span of U is equal to $\mathbb{R}^d$.

bound on the magnitude of innovation in any direction.

There is a null creation corresponding to the origin and the corpus at date n is $C_n = \{0, X_1, \ldots, X_n\}$. We assume the null creation is available to the generator but not owned by any creator, e.g. it is part of the public domain.[13]

Given the creative process X_n, a direction is *persistent* if it is almost surely sampled infinitely often, and *transient* if it is a.s. sampled only finitely often. For simplicity, we assume every direction is either persistent or transient. Let U_p denote the set of persistent directions, and assume that every direction is a positive combination of persistent directions: $U \subseteq \operatorname{cone}(U_p)$.[14]

We measure the restrictiveness of our criterion for a given corpus as the share of generable outputs that are permissible. Let $\operatorname{Vol}(A)$ denote the d-dimensional Lebesgue measure of A.[15]

Definition 6 (Permissible Ratio)**.** For any corpus C, the *permissible ratio* is

$$r_g(C) := \frac{\operatorname{Vol}(p_g(C))}{\operatorname{Vol}(g(C))}$$

whenever the denominator is strictly positive.[16]

A permissible ratio near one means that the criterion leaves the AI system largely unconstrained, while a ratio near zero means that most generable outputs are violations.

5.2 Light-Tailed Creative Processes

Our first result pertains to light-tailed creative processes, defined as follows. Given a cumulative distribution function (cdf) F, let $\overline{F} := 1 - F$ denote the survival function.

[13] Because the origin is unprotected, throughout this section the permissible set is defined with respect to removal of the realized works only: $p_g(C_n) := \bigcap_{j \le n} g(C_n \setminus \{X_j\})$. This is a special case of the general framework of Section 6.2, in which the collection of protected works consists of all singletons other than the origin.

[14] The conic hull $\operatorname{cone}(U')$ of a set U' is the set of all positive combinations $\sum_i a_i u_i$ for $u_i \in U'$ and $a_i \ge 0$.

[15] If A is not measurable, then we take the outer Lebesgue measure.

[16] If $\operatorname{Vol}(p_g(C)) = \infty$, we set the ratio $r_g(C) = 1$ since $p_g(C) \subseteq g(C)$.

Definition 7. A random variable on $\mathbb{R}_+$ with cdf F has *light tails* if for all $c \geq 1$ and all $t \geq 0$,

$$\overline{F}(ct) \leq \overline{F}(t)^c .$$

This condition is satisfied with equality by exponential distributions, since

$$\overline{F}(ct) = e^{-\lambda ct} = \overline{F}(t)^c .$$

Thus, a light-tailed random variable has a tail no heavier than that of an exponential distribution.[17] Examples include the Weibull and gamma distributions with shape parameter at least one, as well as the truncated normal distribution.

We also impose a continuity property on the generator.

Definition 8. g is *uniform lower semicontinuous (ULS)* if for every $\varepsilon \in (0,1)$, there is a $\delta \in (0,1)$ such that for all $C \in \mathcal{C}$ with $0 \in C$,

$$(1-\varepsilon)\, g(C) \subseteq g((1-\delta)\, C) .$$

ULS says that if we shrink inputs by a small proportion, then the generator can still produce almost all of the old outputs. The condition rules out technologies in which an arbitrarily small loss at the frontier of the input destroys a large interior region of the output space. One simple sufficient condition—which is satisfied by the convex hull, splice, and box generators—is homogeneity of degree one, i.e., $g(\lambda C) = \lambda g(C)$ for $\lambda > 0$.[18]

Theorem 1 shows that under light-tailed innovation, the permissible ratio converges to one: with enough accumulated works, essentially every generable output is permissible.

Theorem 1. *Fix a convex-valued and ULS generator g. If F_k has light tails for every*

[17] A sufficient condition is that the distribution has an increasing hazard rate.
[18] To see this, set $\delta = \varepsilon$ so $g((1-\varepsilon)\, C) = (1-\varepsilon)\, g(C)$.

persistent direction u_k, then a.s.

$$\lim_{n\to\infty} r_g\left(C_n\right) = 1.$$

Thus in mature creative fields with gradual innovation, our criterion becomes asymptotically permissive: the share of generable outputs that it rules out vanishes. For intuition, consider the convex hull generator. Light-tailed draws can continue to generate breakthrough works, but the proportional expansion of the frontier caused by each new work eventually becomes arbitrarily small. It follows that deleting any one work can at most slightly contract the frontier of the existing corpus. Under the convex hull generator, this barely changes the generable set. The ULS property allows us to transfer this argument from the convex hull to the original generator and show that the permissible ratio converges to one.

5.3 Heavy-Tailed Creative Processes

Our next result pertains to a heavy-tailed creative process.

Definition 9. A random variable on $\mathbb{R}_+$ with cdf F has *heavy tails* if there is some $\alpha > 0$ such that for all $c \geq 1$ and all $t \geq 0$,

$$\overline{F}\left(ct\right) \geq c^{-\alpha}\overline{F}\left(t\right). \tag{1}$$

The condition (1) is satisfied with equality (for t above the scale parameter b) by Pareto distributions with any shape parameter $\alpha > 0$, since

$$\overline{F}\left(ct\right) = \left(\frac{b}{ct}\right)^{\alpha} = c^{-\alpha}\overline{F}\left(t\right).$$

Thus, a heavy-tailed random variable has a tail at least as heavy as that of a Pareto distribution, with examples including the Lomax, Burr, and Fréchet distributions.[19]

[19] The light- and heavy-tailed classes are not exhaustive: for example, the lognormal distribution satisfies neither definition.

If the generator is unrestricted, with $g(C) = \mathbb{R}^d$ for every corpus C, then the permissible ratio is trivially 1 regardless of the distribution of innovation, since by definition the training corpus does not restrict what the generator can produce. The interesting case is therefore one in which the generator's capabilities are constrained by its training corpus. We impose the following condition.

Definition 10. g is *convex-bounded* if there are $a \geq 0$ and bounded $B \subseteq \mathbb{R}^d$ such that for all $C \in \mathcal{C}$ with $0 \in C$,

$$g(C) \subseteq a \ \operatorname{conv}(C) + B.$$

Convex-boundedness ensures that the technology is not so powerful that it can generate unlimited outputs from limited inputs. More precisely, every generable output must lie within a scaled convex hull of the corpus, up to a uniformly bounded, corpus-independent residual.

Theorem 2 shows that the conclusion of Theorem 1 reverses under heavy-tailed innovation: the permissible ratio falls arbitrarily close to zero infinitely often, no matter how large the corpus grows.

Theorem 2. *Fix a convex-valued and convex-bounded generator g. Suppose $|U_p| = d$. If F_k has heavy tails for some persistent direction u_k, then a.s.*

$$\liminf_{n\to\infty} r_g(C_n) = 0.$$

That is, under heavy-tailed innovation, even arbitrarily large corpora repeatedly contain works so pivotal that their inclusion is responsible for (close to) the full volume of the generable set. Consequently, constraints on AI generation do not wash out with scale, and the permissibility criterion continues to bind even for very large corpora.

To show this result, we prove that infinitely often, a new work is a large proportional advance over all previous works in its direction. Because there are exactly d spanning persistent directions, the convex hull of the corpus is a simplex whose

volume is proportional to the product of the running maxima across directions. (The assumption of exactly d persistent directions is stronger than what the result requires, but simplifies our volume calculations in this way.) Removing such a breakthrough work therefore shrinks this volume by (nearly) the same proportional factor, and no earlier work in its direction can replace it. Convex-boundedness extends the argument to any generator whose outputs remain on the same volume scale as the convex hull.

5.4 Discussion

Role of volume. The permissible ratio, which compares volumes, is not intended to reflect economic demand or welfare. For example, if demand is not uniform over the creative space, then a small violation set can still carry significant economic weight. Conversely, a large violation set may have little value if it lies in regions of the creative space where consumer demand is insignificant. Nevertheless, since our results pertain to the extreme cases where the limit ratio is zero or one, the qualitative conclusions that we prove for volume extend to any alternative measure that is uniformly comparable to volume on the relevant generable sets.[20]

Role of convex-valuedness. Convex-valuedness allows us to work with the geometry of convex hulls, but it plays a distinct role in each theorem. In Theorem 1, we show that deleting a work only slightly shrinks the convex hull of the corpus. Since $g(C) = g(\mathrm{conv}(C))$ for convex-valued generators (Proposition 1), the ULS property translates this statement about hulls into a statement about generable sets. In Theorem 2, convex-valuedness combines with Preservation to give the lower bound $\mathrm{conv}(C) \subseteq g(C)$, which together with the upper bound from convex-boundedness places the generable set on the same "volume scale" as the hull, so that a breakthrough that expands the hull expands generable volume comparably.

[20] Its density with respect to Lebesgue measure must be bounded above and bounded away from zero, with bounds that remain stable along the sequence of corpora. The weaker condition of absolute continuity alone preserves zero-volume conclusions, but not ratio comparisons.

Neither argument requires convex-valuedness per se: Theorem 1 needs small contractions of the hull to imply small contractions of the generable set, and Theorem 2 needs $\lambda \operatorname{conv}(C) \subseteq g(C) \subseteq a \operatorname{conv}(C) + B$ for some $\lambda > 0$. These weaker conditions extend beyond the class of convex-valued generators, at a cost in interpretability.

Incentives and regime selection. We study a reduced-form arrival process of creative works. A richer equilibrium model would derive the tail behavior of innovation from underlying incentives, such as consumer demand for familiar versus novel work, the cost of moving into unexplored regions, reputational rewards for originality, and rents created by legal protection. If these forces generate light-tailed innovation, our results say that individual works eventually become replaceable. If instead they generate heavy-tailed innovation, then some works remain pivotal.

Different incentive environments may select different regimes. In certain commercial genres, audience demand may penalize works that move too far from familiar conventions, pushing creative arrivals toward light-tailed innovation. Alternatively, rewards for originality and differentiation—or the rents from controlling the violation set, as we have defined it—may push creators toward larger frontier jumps. A model of this feedback between legal protection and creative choice would endogenize the regimes studied here, and we leave this for future work.

6 Extensions

6.1 New Creative Dimensions

Section 5 treated breakthroughs as unusually large advances along directions already represented in the corpus. Some breakthroughs are instead different in kind: they introduce a feature that no earlier work possessed, as when a new genre, instrument, or visual technique enters a field. In this section, we extend the model to allow breakthroughs that open new creative dimensions. Because the number of dimensions

that may open over time is unbounded, the ambient creative space is now $\mathbb{R}^{\mathbb{N}}$. There is an infinite set of basis vectors $U := \{u_1, u_2, \dots\} \subseteq \mathbb{R}^{\mathbb{N}}$ where u_k is the basis vector for dimension k. As before, at date n, a direction $I_n \in \mathbb{N}$ is drawn; conditional on $I_n = k$, a magnitude $R_n > 0$ is independently drawn from F_k. Given $I_n = k$ and $R_n = r$, the new work is $X_n = ru_k$ and the corpus is $C_n = \{0, X_1, \dots, X_n\}$, with the origin unprotected and the permissible set defined with respect to removal of the realized works $X_1, \dots, X_n$, as in Section 5.1.

The *dimension* of a corpus $C \subseteq \mathbb{R}^{\mathbb{N}}$, denoted by $\dim(C)$, is the number of non-zero coordinates across all works in C (note this is finite for all C_n). Let $d_n = \dim(C_n)$, and let Vol_n denote the d_n-dimensional Lebesgue measure on $\mathbb{R}^{d_n}$. With this definition of volume, we define the *n-period permissible ratio*

$$r_g(C_n) := \frac{\mathrm{Vol}_n(p_g(C_n))}{\mathrm{Vol}_n(g(C_n))}$$

whenever the denominator is strictly positive. This measures the share of generable outputs that are permissible within the finite-dimensional creative space of the corpus at date n.

We assume that the AI technology is incapable of opening new dimensions.

Definition 11. g is *dimension-bounded* if $\dim(g(C)) = \dim(C)$ for every corpus C.

A dimension-bounded generator can interpolate and recombine features already represented in the corpus but cannot invent an entirely missing dimension. As before, this rules out universally powerful generators such as $g(\cdot) = \mathbb{R}^{\mathbb{N}}$. The convex hull, splice and box generators are all dimension-bounded.

Proposition 5 shows that when the creative process can open new dimensions, the n-period permissible ratio repeatedly collapses to zero.

Proposition 5. *Let g be convex-valued and dimension-bounded. If a new direction is drawn with probability at least $p > 0$ at every period, then a.s. infinitely often*

$$r_g(C_n) = 0.$$

Thus, in creative fields where innovations can open new dimensions, creators retain significant ownership rights even in large corpora. The mechanism is straightforward: deleting the work that opened a new dimension reduces the corpus to a lower dimension, and since a dimension-bounded generator cannot restore the missing dimension, the permissible set has zero volume relative to the corpus.

The conclusion is stronger than in the heavy-tailed case in two respects. First, the permissible ratio is exactly zero infinitely often, not merely arbitrarily close to zero. Second, no condition on the magnitude distributions is required: even if innovation along every existing dimension is light-tailed, the works that open new dimensions remain pivotal.

6.2 General Violations

Apart from the unprotected origin in Sections 5 and 6.1, the analysis so far has treated every work in the corpus as protected, with violations defined at the level of individual works. In practice, however, not all works are protected—some belong to the public domain—and copyright disputes often arise at the level of collections rather than single works. For example, a novelist may sue over their entire catalog, or a class action may aggregate works by many creators. This subsection considers a more general framework that allows for unprotected works and groupwise violations.

A *collection* of protected works is a nonempty family of sets $\mathcal{A} \subseteq \mathcal{C}$. The case studied in Sections 3–4 corresponds to $\mathcal{A}$ being the family of all singleton works. Works that do not belong to any set in $\mathcal{A}$ are unprotected, and interpreted as belonging to the public domain.

Definition 12. Fix a generator g, a collection of protected works $\mathcal{A}$, and a corpus C. The *permissible set* is

$$p_g^{\mathcal{A}}(C) := \bigcap_{A \in \mathcal{A}} g(C \setminus A)$$

and the *violation set* is

$$v_g^{\mathcal{A}}(C) := g(C) \setminus p_g^{\mathcal{A}}(C).$$

For each $A \in \mathcal{A}$, the set $g(C \setminus A)$ consists of all outputs that can be generated without using any work in A, i.e., that are permissible with respect to A.[21] The permissible set $p_g^{\mathcal{A}}(C)$ thus consists of all outputs permissible with respect to every set in the collection $\mathcal{A}$.

Proposition 3 extends fully.

Proposition 6. *For every generator g, collection $\mathcal{A}$, and corpora C, D, the permissible set satisfies*

(1) *If $C \subseteq D$, then $p_g^{\mathcal{A}}(C) \subseteq p_g^{\mathcal{A}}(D)$*

(2) $g\left(p_g^{\mathcal{A}}(C)\right) = p_g^{\mathcal{A}}(C)$

Part (1) is monotonicity in the corpus: expanding the corpus weakly expands the permissible set, just as for individual works. Part (2) is stability: generating from permissible outputs yields only permissible outputs.

We now consider what happens to the permissible set under different protected collections.

Definition 13. A collection $\mathcal{A}$ is *richer* than another collection $\mathcal{B}$ if for every $B \in \mathcal{B}$, there exists some $A \in \mathcal{A}$ such that $B \subseteq A$.

A collection can have "more" protected works in two senses: it can contain additional protected sets, or its sets can be enlarged, as when novelists pool their catalogs to sue as a coalition. Richness captures both—adding a set to a collection makes it richer, and so does enlarging any of its sets.

The following result justifies Definition 13 by showing that moving to a richer collection weakly shrinks the permissible set.

[21] Note that if $A \cap C = \varnothing$, then $g(C \setminus A) = g(C)$.

Proposition 7. *Fix a generator g and corpus C. If a collection $\mathcal{A}$ is richer than a collection $\mathcal{B}$, then*

$$p_g^{\mathcal{A}}(C) \subseteq p_g^{\mathcal{B}}(C).$$

An immediate corollary is that violations are *superadditive*. Since $v_g^{\{A,B\}}(C) = v_g^{\{A\}}(C) \cup v_g^{\{B\}}(C)$, the result says that the violation set of a coalition contains—and can strictly exceed—the union of the individual violation sets.

Corollary 2 (Superadditivity)**.** *Fix a generator g and corpus C. For any two protected sets $A, B \in \mathcal{C}$,*

$$v_g^{\{A,B\}}(C) \subseteq v_g^{\{A \cup B\}}(C)$$

with strict inclusion possible.

The economic content of superadditivity is that when creators join a class action, their combined claim covers outputs that none of them could claim alone. A natural question for future work is which coalitions of creators would form in equilibrium.

7 Conclusion and Open Questions

This paper proposes a new definition of creative ownership developed with generative AI in mind. In our framework, an output depends essentially on an existing work if it could not have been generated without that work in the training corpus. An output is permissible if it does not depend essentially on any existing work. We give conditions under which the permissible set is nonempty, show that it is monotone and stable, and study its long-run size: under light-tailed innovation our criterion eventually permits essentially all generable outputs, while under heavy-tailed innovation—or innovation that opens new creative dimensions—some works remain pivotal and the criterion continues to bind.

There are many possible directions for extending these ideas. One natural next step is to move from our binary notion of dependence to a cardinal measure of influ-

ence. Section 3.1 suggested one route, measuring how much the removal of a work reduces the probability of generating a given output. A complementary route would quantify how much a work expands the system's generative capacity as a whole: given a measure ν over the creative space, the influence of a work x could be measured by comparing $\nu(g(C))$ with $\nu(g(C \setminus \{x\}))$.

Another interesting direction would be to endogenize both the corpus and the generator. Creators who anticipate the use of our criterion to determine creative ownership may strategically choose creative directions and effort, while AI firms may invest in generative capabilities that respond to those endogenous choices. A richer model of this form would connect the legal definition of dependence to the economic forces that determine which works become influential in the first place.

Appendix

A Proof of Proposition 1

The proposition asserts that (1)–(3) are equivalent; we show that each of (2) and (3) is equivalent to (1). We first show that g is convex-valued if and only if $g = g_{conv} \circ g$. In one direction, if g is convex-valued then $g(C)$ is convex, implying $\text{conv}(g(C)) = g(C)$ as desired. In the other direction, if $g(C) = \text{conv}(g(C))$ then $g(C)$ is clearly convex.

We now show that g is convex-valued if and only if $g = g \circ g_{conv}$. First suppose g is convex-valued. Then since $C \subseteq g_{conv}(C)$ we have

$$g(C) \subseteq g(g_{conv}(C)) \tag{2}$$

by monotonicity of g. Moreover,

$$\begin{aligned} g_{conv}(C) &\subseteq g_{conv}(g(C)) && \text{since } C \subseteq g(C) \text{ by Preservation} \\ &= g(C) && \text{since } g = g_{conv} \circ g \end{aligned}$$

Applying the mapping g to both sides of this set inclusion yields

$$g(g_{conv}(C)) \subseteq g(g(C)) = g(C) \tag{3}$$

where the last equality follows from idempotence. Combining (2) and (3) yields $g(g_{conv}(C)) = g(C)$ as desired.

For the converse, suppose $g = g \circ g_{conv}$. Then

$$\begin{aligned} g_{conv}(g(C)) &\subseteq g(g_{conv}(g(C))) && \text{by Preservation} \\ &= g(g(C)) && \text{by assumption that } g = g \circ g_{conv} \\ &= g(C) && \text{by idempotence} \\ &\subseteq g_{conv}(g(C)) && \text{by Preservation} \end{aligned}$$

so $g = g_{conv} \circ g$, which directly implies that g is convex-valued.

An immediate corollary is the following.

Corollary A.1. *The convex hull generator is minimal among all convex-valued generators, i.e.,*

$$g_{conv}(C) \subseteq g(C)$$

for every corpus $C \in \mathcal{C}$ and every convex-valued generator g.

Proof. For any convex-valued generator g,

$$g_{conv}(C) \subseteq g(g_{conv}(C)) = g(C)$$

where the set inclusion follows from Preservation and the equality follows from Proposition 1. □

B Proofs for Section 4

B.1 Proof of Proposition 3

Part (1): Let $C \subseteq D$ and without loss write $D = C \cup A$ where C and A are disjoint sets. Then

$$\begin{aligned} p_g(C \cup A) &= \bigcap_{x \in C \cup A} p_g(x, C \cup A) \\ &= \bigcap_{x \in C \cup A} g((C \cup A) \setminus \{x\}) \\ &= \bigcap_{x \in C} g((C \cup A) \setminus \{x\}) \cap \bigcap_{y \in A} g((C \cup A) \setminus \{y\}) \end{aligned}$$

while

$$p_g(C) = \bigcap_{x \in C} g(C \setminus \{x\})$$

Fix an arbitrary $x \in C$. Then $C \setminus \{x\} \subseteq (C \cup A) \setminus \{x\}$, so

$$g(C \setminus \{x\}) \subseteq g((C \cup A) \setminus \{x\})$$

by monotonicity of g. Moreover, $C \setminus \{x\} \subseteq C \subseteq (C \cup A) \setminus \{y\}$ for every $y \in A$, so also

$$g\left(C \setminus \{x\}\right) \subseteq g\left((C \cup A) \setminus \{y\}\right).$$

Thus $p_g\left(C\right) \subseteq p_g(C \cup A) = p_g(D)$ as desired.

Part (2): In one direction, $p_g(C) \subseteq g(p_g(C))$ by Preservation. In the other direction, define $E_x := g(C \setminus \{x\})$ for each $x \in C$, so that $p_g(C) = \bigcap_{x \in C} E_x$. Fix $y \in C$. Since $\bigcap_{x \in C} E_x \subseteq E_y$, monotonicity of g implies

$$g\left(\bigcap_{x \in C} E_x\right) \subseteq g\left(E_y\right)$$

for each $y \in C$. So

$$g\left(\bigcap_{x \in C} E_x\right) \subseteq \bigcap_{y \in C} g\left(E_y\right) \tag{4}$$

But for each $y \in C$,

$$g(E_y) = g(g(C \setminus \{y\})) = g(C \setminus \{y\})$$

where the second equality follows from idempotence. So

$$\bigcap_{y \in C} g\left(E_y\right) = \bigcap_{y \in C} g(C \setminus \{y\}) = p_g(C).$$

Combining this with (4) yields

$$g(p_g(C)) = g\left(\bigcap_{x \in C} E_x\right) \subseteq p_g(C)$$

completing the proof.

B.2 Proof of Proposition 2

Since $|C| \geq R(g)$, the definition of the Radon number yields disjoint nonempty subsets $A, B \subseteq C$ such that $g(A) \cap g(B) \neq \varnothing$. Let z be an element in this intersection. We show that $z \in p_g(C)$.

Recall that $p_g(C) = \bigcap_{y \in C} g(C \setminus \{y\})$. Fix an arbitrary existing work $y \in C$. Since A and B are disjoint, y cannot belong to both sets. This leads to two cases:

(1) If $y \notin A$, then $A \subseteq C \setminus \{y\}$. By monotonicity, $g(A) \subseteq g(C \setminus \{y\})$. Since $z \in g(A)$, it follows that $z \in g(C \setminus \{y\})$.

(2) If $y \notin B$, then $B \subseteq C \setminus \{y\}$. By monotonicity, $g(B) \subseteq g(C \setminus \{y\})$. Since $z \in g(B)$, it follows that $z \in g(C \setminus \{y\})$.

In either case, z remains generable after the removal of y. Since this holds for all $y \in C$, we conclude that $z \in p_g(C)$.

B.3 Proof of Corollary 1

Assume $|C| \geq d + 2$. By Radon's theorem, there exist disjoint subsets $A, B \subseteq C$ whose convex hulls intersect, i.e., $g_{conv}(A) \cap g_{conv}(B) \neq \varnothing$. By Proposition 2, the permissible set for the convex hull generator is nonempty: $p_{g_{conv}}(C) \neq \varnothing$.

Recall from Corollary A.1 that g_{conv} is minimal among convex-valued generators, meaning $g_{conv}(C) \subseteq g(C)$ for any corpus C. This inclusion extends to the permissible sets:

$$p_{g_{conv}}(C) = \bigcap_{x \in C} g_{conv}(C \setminus \{x\}) \subseteq \bigcap_{x \in C} g(C \setminus \{x\}) = p_g(C).$$

Since $p_{g_{conv}}(C)$ is nonempty, it follows that $p_g(C)$ is nonempty as well.

B.4 Proof of Proposition 4

Part (1): Choose some $y \in C$. Since by assumption x is permissible,

$$x \in p_g(C) \subseteq g(C \setminus \{y\}) \tag{5}$$

i.e., x is generable from C without y. Moreover,

$$\begin{aligned}
(C \cup \{x\}) \setminus \{y\} &= (C \setminus \{y\}) \cup \{x\} \\
&\subseteq (C \setminus \{y\}) \cup g\,(C \setminus \{y\}) && \text{by (5)} \\
&\subseteq g\,(C \setminus \{y\}) \cup g\,(C \setminus \{y\}) && \text{by Preservation} \\
&= g\,(C \setminus \{y\})\,.
\end{aligned}$$

Applying g to both sides yields

$$g\,((C \cup \{x\}) \setminus \{y\}) \subseteq g\,(g\,(C \setminus \{y\})) = g\,(C \setminus \{y\}) \tag{6}$$

where the set inclusion follows from monotonicity of g, and the final equality follows from idempotence.

Thus,

$$\begin{aligned}
p_g\,(C \cup \{x\}) &= g(C) \cap \bigcap_{y \in C} g\,((C \cup \{x\}) \setminus \{y\}) \\
&\subseteq g(C) \cap \bigcap_{y \in C} g\,(C \setminus \{y\}) && \text{by (6)} \\
&= \bigcap_{y \in C} g\,(C \setminus \{y\}) = p_g(C).
\end{aligned}$$

The reverse inclusion $p_g\,(C) \subseteq p_g\,(C \cup \{x\})$ follows from Part (1) of Proposition 3, so the two sets are equal, as desired.

Part (2): Let $x \in v_g\,(C)$ so $x \notin p_g\,(C)$. Since g is monotone, $p_g\,(C) \subseteq p_g\,(C \cup \{x\})$ by Part (1) of Proposition 3. We will show that $x \in p_g\,(C \cup \{x\})$, which then implies that $p_g\,(C \cup \{x\})$ is strictly larger.

Suppose otherwise. Then there must be some $y \in C \cup \{x\}$ such that

$$x \in v_g\,(y, C \cup \{x\}) = g\,(C \cup \{x\}) \setminus g\,((C \cup \{x\}) \setminus \{y\}) \tag{7}$$

First suppose $y \in C$. Then

$$x \in (C \cup \{x\}) \setminus \{y\} \subseteq g\left((C \cup \{x\}) \setminus \{y\}\right)$$

where the set inclusion follows from Preservation. But this contradicts (7).

Suppose instead $y = x$. Then

$$x \in v_g\left(x, C \cup \{x\}\right) = g\left(C \cup \{x\}\right) \setminus g(C)$$

so $x \notin g\left(C\right)$, yielding a contradiction with our initial assumption that $x \in v_g(C)$.

Part (3): The set inclusion $p_g(C) \subseteq p_g(C \cup \{x\})$ is a direct consequence of Part (1) of Proposition 3, and the possibility of both equality and strict inclusion is shown in Example 9.

C Proofs for Section 5

C.1 Redundancy of Transient Works

We first prove the redundancy of transient works, which will be used in the proofs of both theorems below. Let $C_n^* \subseteq C_n$ consist of the works in C_n that are not transient:

$$C_n^* := C_n \setminus \left\{X_i \in C_n : X_i \text{ is transient}\right\}.$$

Thus, a work in a transient direction belongs to C_n but not to C_n^*.

Lemma C.1. Almost surely, for sufficiently large n and any $i \leq n$, $\operatorname{conv}\left(C_n^*\right) = \operatorname{conv}\left(C_n\right)$ and $\operatorname{conv}\left(C_n^* \setminus \{X_i\}\right) = \operatorname{conv}\left(C_n \setminus \{X_i\}\right)$.

Proof. Since $C_n^* \subseteq C_n$, $\operatorname{conv}\left(C_n^*\right) \subseteq \operatorname{conv}\left(C_n\right)$ and $\operatorname{conv}\left(C_n^* \setminus \{X_i\}\right) \subseteq \operatorname{conv}\left(C_n \setminus \{X_i\}\right)$. Thus, all we need to prove is the reverse inclusion.

For each persistent direction u_k, let $\hat{M}_n^k$ denote the second-largest magnitude realized in direction u_k by date n (which is defined once the direction has been sampled twice). Since $U \subseteq \operatorname{cone}\left(U_p\right)$, we have for any $u_i \in U$, $u_i = \sum_k a_k u_k$ for $a_k \geq 0$ and

$u_k \in U_p$. Note that for any realized X_i,

$$X_i = R_i u_i = R_i \sum_k a_k u_k = R_i \sum_k \frac{a_k}{\hat{M}_n^k} \hat{M}_n^k u_k$$

If we let $\lambda_k := \frac{a_k/\hat{M}_n^k}{\sum_k a_k/\hat{M}_n^k}$, then

$$X_i = \left(\sum_k R_i \frac{a_k}{\hat{M}_n^k} \right) \sum_k \lambda_k \hat{M}_n^k u_k$$

First, suppose X_i is in a transient direction. Since each transient direction is a.s. sampled only finitely often, R_i is bounded a.s. Since every persistent direction is sampled infinitely often and F_k has unbounded support, $\hat{M}_n^k \to \infty$ a.s.. This implies that for sufficiently large n, $\sum_k R_i \frac{a_k}{\hat{M}_n^k} \leq 1$ so X_i is a convex combination of 0 and $\sum_k \lambda_k \hat{M}_n^k u_k$. This establishes that $X_i \in \operatorname{conv}(C_n^* \setminus \{X_i\})$. The case for when X_i is in a persistent direction is straightforward, so the result obtains. □

C.2 Proof of Theorem 1

We next prove two lemmas. The first shows that with light tails, the running maximum of an iid sequence is eventually never far ahead of the second-best draw; the second converts this into a statement about the convex hull of the corpus.

Lemma C.2. Suppose Z_n are iid random variables on $\mathbb{R}_+$ with light tails and unbounded support. For any $c > 1$, almost surely there exists $N(\omega)$ such that

$$\max_{i \leq n} Z_i \leq c \max_{i \leq n, i \neq j} Z_i$$

for all $n > N(\omega)$ and all $j \leq n$.

Proof. Let F denote the cdf of Z and $\overline{F} := 1 - F$. It suffices to prove the lemma for $c \in (1, 2)$, since the conclusion for a given c implies it for every larger c; fix such a c.

Define $M_n := \max_{i \le n} Z_i$ (with $M_0 := 0$) and the event

$$E_n := \{Z_n > cM_{n-1}\}.$$

Then, by the light-tailed assumption,

$$\begin{aligned}
\mathbb{P}(E_n) &= \mathbb{E}\left[\mathbb{P}\left(Z_n > cM_{n-1} \mid Z_1, \ldots, Z_{n-1}\right)\right] \\
&= \mathbb{E}\left[\overline{F}\left(cM_{n-1}\right)\right] \\
&\le \mathbb{E}\left[\left(\overline{F}\left(M_{n-1}\right)\right)^c\right] \\
&= \mathbb{E}\left[\left(1 - \max_{i \le n-1} F\left(Z_i\right)\right)^c\right]
\end{aligned}$$

Note that $\mathbb{P}\{F(Z_i) \le t\} \le t$ for every $t \in [0,1]$, so $F(Z_i)$ first-order stochastically dominates a uniform random variable on $[0,1]$. Thus, $\max_{i \le n-1} F(Z_i)$ stochastically dominates the maximum of $n-1$ iid uniform random variables, which has cdf t^{n-1}, so

$$\begin{aligned}
\mathbb{P}(E_n) &\le \int_0^1 (1-t)^c (n-1) t^{n-2} dt = (n-1) B(n-1, 1+c) \\
&= \frac{1}{n} \Gamma(1+c) \frac{\Gamma(n+1)}{\Gamma(n+c)}
\end{aligned}$$

where B and Γ are the beta and gamma functions respectively. From Gautschi's inequality, we have

$$\begin{aligned}
\frac{\Gamma(n+1)}{\Gamma(n+c)} &= \frac{\Gamma((n+c-1)+(2-c))}{\Gamma(n+c)} < (n+c-1)^{-(1-(2-c))} \\
&= \frac{1}{(n+c-1)^{c-1}} \le \frac{1}{n^{c-1}}
\end{aligned}$$

Thus,

$$\sum_n \mathbb{P}(E_n) \le \Gamma(1+c) \sum_n \frac{1}{n^c} < \infty.$$

By the Borel-Cantelli lemma, we have a.s.

$$\sum_n 1_{E_n}(\omega) < \infty,$$

so the last record time $N(\omega) := \max\{n : \omega \in E_n\}$ is finite (with $N(\omega) := 0$ if no E_n occurs). In other words, there is an almost sure event E such that $Z_n > cM_{n-1}$ finitely often.

Since Z has unbounded support, there is an almost sure event H such that $\max_{i \le n} Z_i$ converges to infinity. Note that $\mathbb{P}(H \cap E) = 1$ and consider $\omega \in H \cap E$. We will show that for sufficiently large n and for all $j \le n$

$$\max_{i \le n} Z_i \le c \max_{i \le n, i \ne j} Z_i$$

Suppose otherwise, so we can find an infinite subsequence (s_ℓ) and $j_\ell \le s_\ell$ such that

$$\max_{i \le s_\ell} Z_i > c \max_{i \le s_\ell, i \ne j_\ell} Z_i$$

This means that $Z_{j_\ell} > \max_{i \le s_\ell, i \ne j_\ell} Z_i$ so $\max_{i \le s_\ell} Z_i = Z_{j_\ell}$. Thus,

$$Z_{j_\ell} > c \max_{i \le j_\ell - 1} Z_i$$

Since $\omega \in E$, this means $j_\ell \le N(\omega)$ so

$$M_{s_\ell} = Z_{j_\ell} \le M_{N(\omega)}$$

However, since $\omega \in H$, the left side is unbounded, yielding a contradiction. □

Lemma C.3. Suppose F_k has light tails for every persistent direction u_k. For any $\delta \in (0,1)$, almost surely there exists $N(\omega)$ such that

$$(1-\delta)\operatorname{conv}(C_n) \subseteq \operatorname{conv}(C_n \setminus \{X_j\})$$

for all $n > N(\omega)$ and all $j \le n$.

Proof. Recall $U_p \subseteq U$ is the set of persistent directions. For each persistent direction k, let

$$T_1^k < T_2^k < \cdots$$

denote the dates at which direction u_k is sampled, and let R_ℓ^k be the corresponding

magnitude of the ℓ-th such work, so

$$X_{T_\ell^k} = u_k R_\ell^k.$$

Now fix a persistent direction $u_k \in U_p$. Since $Z_\ell := R_\ell^k$ has light tails, by Lemma C.2 for any $\delta \in (0,1)$, there is an a.s. event E_k on which, for some $L_k(\omega)$,

$$(1-\delta) \max_{i \le m} R_i^k \le \max_{i \le m, i \ne \ell} R_i^k$$

for all $m > L_k(\omega)$ and all $\ell \le m$. Since u_k is persistent, the sequence $\left(T_m^k\right)_m$ is almost surely infinite, so $T_{L_k(\omega)+1}^k < \infty$.

Recall that $C_n^* \subseteq C_n$ consists of the origin together with the works in persistent directions, and that by Lemma C.1, a.s. for all sufficiently large n, $\operatorname{conv}(C_n^*) = \operatorname{conv}(C_n)$ and $\operatorname{conv}(C_n^* \setminus \{X_j\}) = \operatorname{conv}(C_n \setminus \{X_j\})$ for all $j \le n$. If X_j is in a transient direction, then $X_j \notin C_n^*$ and hence, for every $n \ge j$,

$$C_n^* \subseteq C_n \setminus \{X_j\} \subseteq C_n.$$

Taking convex hulls gives

$$\operatorname{conv}(C_n^*) \subseteq \operatorname{conv}(C_n \setminus \{X_j\}) \subseteq \operatorname{conv}(C_n).$$

Therefore, on the event where $\operatorname{conv}(C_n^*) = \operatorname{conv}(C_n)$, we obtain

$$\operatorname{conv}(C_n \setminus \{X_j\}) = \operatorname{conv}(C_n).$$

So for all sufficiently large such n,

$$(1-\delta) \operatorname{conv}(C_n) \subseteq \operatorname{conv}(C_n) = \operatorname{conv}(C_n \setminus \{X_j\}).$$

Thus, all we need to check are those X_j that are in persistent directions. Consider the a.s. event $E := \bigcap_k E_k$ and choose

$$N(\omega) \ge \max_{u_k \in U_p} T_{L_k(\omega)+1}^k$$

large enough that, for all $n > N(\omega)$, the two hull equalities above hold. Now fix $n > N(\omega)$ and suppose X_j lies in a persistent direction k. Let

$$T_1^k, \ldots, T_m^k$$

be the dates up to n at which direction k is sampled. Then $j = T_\ell^k$ for some $\ell \leq m$, and $m > L_k(\omega)$. Therefore the event E_k gives

$$(1-\delta) \max_{i \leq m} R_i^k \leq \max_{i \leq m, i \neq \ell} R_i^k.$$

Let $M_n^{k'}$ denote the record magnitude realized in each persistent direction k' by date n. Since $0 \in C_n^*$ and every work in direction k' lies on the segment from 0 to $M_n^{k'} u_{k'}$,

$$\operatorname{conv}(C_n) = \operatorname{conv}(C_n^*) = \operatorname{conv}\left(\{0\} \cup \left\{M_n^{k'} u_{k'} : u_{k'} \in U_p\right\}\right),$$

and hence

$$(1-\delta) \operatorname{conv}(C_n) = \operatorname{conv}\left(\{0\} \cup \left\{(1-\delta) M_n^{k'} u_{k'} : u_{k'} \in U_p\right\}\right).$$

It therefore suffices to show that each generating point of the right-hand side belongs to $\operatorname{conv}(C_n \setminus \{X_j\})$. First, $0 \in C_n \setminus \{X_j\}$, since the origin is never removed. For $k' \neq k$, the record work $M_n^{k'} u_{k'}$ remains in $C_n \setminus \{X_j\}$, so $(1-\delta) M_n^{k'} u_{k'}$ lies on the segment from 0 to it. For $k' = k$, the display above gives $(1-\delta) M_n^k \leq \max_{i \leq m, i \neq \ell} R_i^k$, and the work attaining this maximum is a.s. a distinct point from X_j (since F_k is atomless), so $(1-\delta) M_n^k u_k$ lies on the segment from 0 to a work in direction k that remains in $C_n \setminus \{X_j\}$. Therefore

$$(1-\delta) \operatorname{conv}(C_n) \subseteq \operatorname{conv}(C_n \setminus \{X_j\})$$

as desired. □

Proof of Theorem 1. Consider some $\varepsilon > 0$. Since g is ULS, there is some $\delta > 0$ such that $(1-\varepsilon) g(C) \subseteq g((1-\delta) C)$ for every corpus C with $0 \in C$; without loss of

generality $\delta \in (0,1)$, since for a convex corpus containing 0, shrinking δ only enlarges $(1-\delta)\,C$ and hence, by monotonicity, $g\left((1-\delta)\,C\right)$. Since $0 \in C_n \subseteq \operatorname{conv}(C_n)$, applying the ULS property to $\operatorname{conv}(C_n)$ gives

$$(1-\varepsilon)\,g\left(\operatorname{conv}(C_n)\right) \subseteq g\left((1-\delta)\operatorname{conv}(C_n)\right).$$

Now, from Lemma C.3, there is an a.s. event $E(\varepsilon)$ such that for large n and all $j \leq n$,

$$\begin{aligned}
(1-\delta)\operatorname{conv}(C_n) &\subseteq \operatorname{conv}(C_n \setminus \{X_j\}) \\
g\left((1-\delta)\operatorname{conv}(C_n)\right) &\subseteq g\left(\operatorname{conv}(C_n \setminus \{X_j\})\right) \\
(1-\varepsilon)\,g\left(\operatorname{conv}(C_n)\right) &\subseteq g\left(\operatorname{conv}(C_n \setminus \{X_j\})\right) \\
(1-\varepsilon)\,g(C_n) &\subseteq g(C_n \setminus \{X_j\})
\end{aligned}$$

where the second line follows from monotonicity of g, the third from the choice of δ, and the last from Proposition 1. Intersecting over all j yields

$$(1-\varepsilon)\,g(C_n) \subseteq \bigcap_{j=1}^{n} g(C_n \setminus \{X_j\}) = p_g(C_n).$$

This implies that

$$(1-\varepsilon)^d \operatorname{Vol}\left(g(C_n)\right) = \operatorname{Vol}\left((1-\varepsilon)\,g(C_n)\right) \leq \operatorname{Vol}\left(p_g(C_n)\right).$$

Note that if $\operatorname{Vol}(g(C_n)) = \infty$ for some n, then $\operatorname{Vol}(p_g(C_n)) = \infty$ so $r_g(C_n) = 1$. Moreover, $\operatorname{Vol}(g(C_n)) > 0$ for all sufficiently large n: once every persistent direction has contributed a work, $\operatorname{conv}(C_n) \subseteq g(C_n)$ contains the simplex spanned by 0 and works in d linearly independent persistent directions, which exist since the persistent directions span $\mathbb{R}^d$ (as U does and $U \subseteq \operatorname{cone}(U_p) \subseteq \operatorname{span}(U_p)$). Otherwise,

$$(1-\varepsilon)^{\,d} \leq \frac{\operatorname{Vol}\left(p_g(C_n)\right)}{\operatorname{Vol}\left(g(C_n)\right)} \leq 1.$$

Taking $\varepsilon_k \searrow 0$ and considering the a.s. event $\bigcap_k E(\varepsilon_k)$ yields the result. $\square$

C.3 Proof of Theorem 2

We first prove a lemma: with heavy tails, the running maximum of an iid sequence is beaten by a large proportional margin infinitely often.

Lemma C.4. Suppose Z_n are iid random variables on $\mathbb{R}_+$ with an atomless, heavy-tailed distribution. For any $c > 1$, there is an a.s. event such that

$$\max_{i \leq n} Z_i > c \max_{i \leq n-1} Z_i$$

infinitely often.

Proof. Let F denote the cdf of Z and $\overline{F} := 1 - F$. Define $M_n := \max_{i \leq n} Z_i$ and the event

$$A_n := \{Z_n > M_{n-1}\}$$

By Rényi (1962), the events A_n are independent and $\mathbb{P}(A_n) = \frac{1}{n}$. Note that

$$\sum_n \mathbb{P}(A_n) = \infty$$

so by the Borel-Cantelli lemma, a.s.

$$\sum_n 1_{A_n}(\omega) = \infty.$$

By the conditional Borel-Cantelli lemma, a.s.

$$\infty = \sum_n \mathbb{P}(A_n \mid M_{n-1}) = \sum_n \overline{F}(M_{n-1}).$$

Now, fix some $c > 1$ and define the event

$$E_n := \{Z_n > cM_{n-1}\}$$

and note that a.s.

$$\sum_n \mathbb{P}\left(E_n \mid M_{n-1}\right) = \sum_n \overline{F}\left(cM_{n-1}\right)$$
$$\geq c^{-\alpha} \sum_n \overline{F}\left(M_{n-1}\right) = \infty.$$

Applying the conditional Borel-Cantelli lemma one more time, we have a.s.

$$\sum_n 1_{E_n}(\omega) = \infty.$$

In other words, there is an a.s. event such that $Z_n > cM_{n-1}$ infinitely often. □

Proof of Theorem 2. Recall $U_p \subseteq U$ is the set of persistent directions. For each persistent direction k, let

$$T_1^k < T_2^k < \cdots$$

denote the dates at which direction u_k is sampled, and let R_ℓ^k be the corresponding magnitude of the ℓth such work, so

$$X_{T_\ell^k} = u_k R_\ell^k.$$

Without loss of generality, let u_1 be a persistent direction whose magnitudes are drawn from a heavy-tailed distribution F_1. By Lemma C.4, for any $c > 1$, almost surely

$$\max_{i \leq \ell} R_i^1 > c \max_{i \leq \ell - 1} R_i^1 \tag{8}$$

for infinitely many ℓ.

Recall that $C_n^* \subseteq C_n$ consists of the origin together with the works in persistent directions, and that by Lemma C.1, a.s. for all sufficiently large n, $\operatorname{conv}(C_n^*) = \operatorname{conv}(C_n)$ and $\operatorname{conv}(C_n^* \setminus \{X_j\}) = \operatorname{conv}(C_n \setminus \{X_j\})$ for all $j \leq n$. Since every persistent direction is sampled infinitely often, for all sufficiently large n each persistent direction has appeared at least once. For such n, let

$$M_n^k := \max\{R_i^k : T_i^k \leq n\}$$

denote the record magnitude realized in direction k by date n. Since $0 \in C_n^*$ and every work in a persistent direction lies on the segment from 0 to its direction's record, $\operatorname{conv}(C_n^*)$ is the simplex with vertices $0, M_n^1 u_1, \ldots, M_n^d u_d$ (the persistent directions are linearly independent, since U spans $\mathbb{R}^d$ and $U \subseteq \operatorname{cone}(U_p) \subseteq \operatorname{span}(U_p)$), and we have

$$\operatorname{Vol}(\operatorname{conv}(C_n^*)) = \frac{1}{d!} M_n^1 \cdots M_n^d \left|\det A^*\right|$$

where A^* is the matrix whose columns are the persistent directions.

Almost surely, there are infinitely many ℓ for which (8) holds. Fix such an $\ell \geq 2$ and let $n := T_\ell^1$. By (8), the newest work $X_n = u_1 R_\ell^1$ is the record in direction u_1: $M_n^1 = R_\ell^1 > c \max_{i \leq \ell - 1} R_i^1$. Let $L_n := \operatorname{conv}(C_n^* \setminus \{X_n\})$. Removing X_n leaves the previous record $\max_{i \leq \ell - 1} R_i^1$ in direction u_1 and does not affect the other directions, so L_n is the simplex with vertices 0, $(\max_{i \leq \ell - 1} R_i^1) u_1$, and $M_n^k u_k$ for $k \geq 2$. Therefore

$$\frac{\operatorname{Vol}(L_n)}{\operatorname{Vol}(\operatorname{conv}(C_n^*))} = \frac{\max_{i \leq \ell - 1} R_i^1}{\max_{i \leq \ell} R_i^1} < \frac{1}{c}. \tag{9}$$

Since $0 \in L_n$ and g is convex-bounded, there is some $a \geq 0$ and bounded B such that

$$g(L_n) \subseteq a L_n + B.$$

Since U_p spans $\mathbb{R}^d$, there is some ball $B_\varepsilon(z)$ centered at some point z such that

$$B_\varepsilon(z) \subseteq \operatorname{conv}(U_p \cup \{0\})$$
$$B_\varepsilon(0) = B_\varepsilon(z) - z \subseteq \operatorname{conv}(U_p \cup \{0\}) - z.$$

Let r_n denote the minimum over persistent directions u_k of the largest magnitude realized in direction u_k within $C_n^* \setminus \{X_n\}$. Then $r_n u_k \in L_n$ for every persistent direction u_k, and $0 \in L_n$, so

$$r_n B_\varepsilon(0) \subseteq r_n (\operatorname{conv}(U_p \cup \{0\}) - z) \subseteq L_n - r_n z.$$

Each persistent direction is sampled infinitely often and its magnitude distribution has

unbounded support, so even the second-largest magnitude realized in each persistent direction diverges (as in the proof of Lemma C.1), and hence $r_n \to \infty$ a.s. Since B is bounded, we have a.s., for all sufficiently large n, $B \subseteq r_n B_\varepsilon(0) \subseteq L_n - r_n z$. Thus,

$$a(L_n - r_n z) + B \subseteq (a+1)(L_n - r_n z).$$

Since $\mathrm{Vol}(\cdot)$ is translation invariant,

$$\begin{aligned}\mathrm{Vol}(aL_n + B) &= \mathrm{Vol}(a(L_n - r_n z) + B) \\ &\leq \mathrm{Vol}((a+1)(L_n - r_n z)) \\ &= (a+1)^d \mathrm{Vol}(L_n).\end{aligned}$$

Thus a.s.

$$\mathrm{Vol}(g(L_n)) \leq (a+1)^d \mathrm{Vol}(L_n).$$

Next, note that a.s., by Proposition 1 and Lemma C.1,

$$p_g(C_n) \subseteq g(C_n \setminus \{X_n\}) = g(\mathrm{conv}(C_n \setminus \{X_n\})) = g(L_n)$$

so

$$\mathrm{Vol}(p_g(C_n)) \leq (a+1)^d \mathrm{Vol}(L_n). \tag{10}$$

On the other hand, since $\mathrm{conv}(C_n) \subseteq g(\mathrm{conv}(C_n))$,

$$\mathrm{Vol}(g(C_n)) = \mathrm{Vol}(g(\mathrm{conv}(C_n))) \geq \mathrm{Vol}(\mathrm{conv}(C_n)) = \mathrm{Vol}(\mathrm{conv}(C_n^*)). \tag{11}$$

Combining (9), (10) and (11), we have a.s.

$$r_g(C_n) = \frac{\mathrm{Vol}(p_g(C_n))}{\mathrm{Vol}(g(C_n))} \leq \frac{(a+1)^d \mathrm{Vol}(L_n)}{\mathrm{Vol}(\mathrm{conv}(C_n^*))} < \frac{(a+1)^d}{c}$$

infinitely often. Setting $c \to \infty$ yields the result. □

D Proofs for Section 6

D.1 Proof of Proposition 5

By assumption, a previously undrawn direction is drawn with probability at least $p > 0$ in every period, regardless of history, so by the conditional Borel–Cantelli lemma, a.s. infinitely often

$$\dim(C_n) > \dim(C_n \setminus \{X_n\}).$$

Since g is dimension-bounded,

$$\dim(g(C_n)) > \dim(g(C_n \setminus \{X_n\})).$$

Since $p_g(C_n) \subseteq g(C_n \setminus \{X_n\})$ by definition, and the latter set has dimension less than d_n,

$$\mathrm{Vol}_n(p_g(C_n)) \leq \mathrm{Vol}_n(g(C_n \setminus \{X_n\})) = 0$$

and $r_g(C_n) = 0$ infinitely often.

D.2 Proof of Proposition 6

Part (1). Suppose $C \subseteq D$. For any $A \in \mathcal{A}$,

$$C \setminus A \subseteq D \setminus A.$$

By monotonicity of the generator g, for every $A \in \mathcal{A}$,

$$g(C \setminus A) \subseteq g(D \setminus A).$$

Taking intersections over all A yields

$$\bigcap_{A \in \mathcal{A}} g(C \setminus A) \subseteq \bigcap_{A \in \mathcal{A}} g(D \setminus A),$$

which is exactly $p_g^{\mathcal{A}}(C) \subseteq p_g^{\mathcal{A}}(D)$.

Part (2). Let $P := p_g^{\mathcal{A}}(C)$. By Preservation, $P \subseteq g(P)$. To show the reverse inclusion, fix any $A \in \mathcal{A}$. Since

$$P \subseteq g(C \setminus A),$$

we have

$$g(P) \subseteq g(g(C \setminus A)) = g(C \setminus A)$$

where the inclusion follows from monotonicity of g, and the equality follows from idempotence of g. Intersecting over all A yields

$$g(P) \subseteq \bigcap_{A \in \mathcal{A}} g(C \setminus A) = P$$

as desired.

D.3 Proof of Proposition 7

Since $\mathcal{A}$ is richer than $\mathcal{B}$, for each $B \in \mathcal{B}$ we can choose some $A_B \in \mathcal{A}$ with $B \subseteq A_B$. Since $C \setminus A_B \subseteq C \setminus B$, monotonicity of g yields

$$g\left(C \setminus A_B\right) \subseteq g\left(C \setminus B\right).$$

Thus,

$$p_g^{\mathcal{A}}\left(C\right) = \bigcap_{A \in \mathcal{A}} g\left(C \setminus A\right) \subseteq \bigcap_{B \in \mathcal{B}} g\left(C \setminus A_B\right) \subseteq \bigcap_{B \in \mathcal{B}} g\left(C \setminus B\right) = p_g^{\mathcal{B}}\left(C\right)$$

as desired.

D.4 Proof of Corollary 2

Let $\mathcal{U} = \{A \cup B\}$ and $\mathcal{S} = \{A, B\}$. The collection $\mathcal{U}$ is richer than $\mathcal{S}$, so Proposition 7 gives

$$p_g^{\mathcal{U}}(C) \subseteq p_g^{\mathcal{S}}(C).$$

Taking complements within $g(C)$ yields

$$v_g^{\{A,B\}}(C) \subseteq v_g^{\{A\cup B\}}(C).$$

Strict inclusion can occur. Let $g = g_{conv}$ and let

$$C = \{(0,0),(1,0),(1,1),(0,1)\} \subset \mathbb{R}^2, \quad A = \{(0,0)\}, \quad B = \{(1,0)\}.$$

The point $z = (1/2, 3/4)$ lies in both $g(C \setminus A)$ and $g(C \setminus B)$, so $z \in p_g^{\{A,B\}}(C)$. But $C \setminus (A \cup B) = \{(1,1),(0,1)\}$, whose convex hull is the top edge of the square, so $z \notin p_g^{\{A\cup B\}}(C)$. Hence $z \in v_g^{\{A\cup B\}}(C) \setminus v_g^{\{A,B\}}(C)$.

References

Andersen, S. (2022): "The Alt-Right Manipulated My Comic. Then A.I. Claimed It." *The New York Times*.

Bengio, Y., A. Courville, and P. Vincent (2013): "Representation Learning: A Review and New Perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35, 1798–1828.

Bourtoule, L., V. Chandrasekaran, C. A. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie, and N. Papernot (2019): "Machine Unlearning," *CoRR*, abs/1912.03817.

Brodsky, S. (2024): "The AI-Generated Books Trend is Getting Worse," Accessed: 2026-01-29.

Chung, K. H. and R. A. K. Cox (1994): "A stochastic model of superstardom: An application of the Yule distribution," *The Review of Economics and Statistics*, 76, 771–775.

Cohen, A. and C. Zhang (2026): "Barriers to Counterfactual Credit Attribution for Autoregressive Models," ICML 2026.

Cullen, Z., D. Li, and S. Li (2026): "Labor as Capital: AI and the Ownership of Expertise," Preliminary and incomplete manuscript.

Dai, Y., M. Farboodi, N. Golrezaei, and S. Shahshahani (2026): "Market Design for AI: Beyond the Copyright Binary," Working paper.

de Rassenfosse, G., A. Jaffe, and J. Waldfogel (2024): "Intellectual Property and Creative Machines," *NBER Working Paper*.

De Vany, A. and W. D. Walls (1999): "Uncertainty in the movie industry: Does star power reduce the terror of the box office?" *Journal of Cultural Economics*, 23, 285–318.

Drugov, M. and D. Ryvkin (2020): "Tournament rewards and heavy tails," *Journal of Economic Theory*, 188, 105092.

Dwork, C., F. McSherry, K. Nissim, and A. Smith (2006): "Calibrating Noise

to Sensitivity in Private Data Analysis," in *Theory of Cryptography*, Springer, vol. 3876 of *Lecture Notes in Computer Science*, 265–284.

Gans, J. S. (2024): "Copyright Policy Options for Generative Artificial Intelligence," *NBER Working Paper*.

Garicano, L. and L. Rayo (2025): "Training in the Age of AI: A Theory of Apprenticeship Viability," Working paper, Kellogg School of Management, Northwestern University.

Ghorbani, A. and J. Zou (2019): "Data Shapley: Equitable Valuation of Data for Machine Learning," in *Proceedings of the 36th International Conference on Machine Learning*, PMLR, vol. 97 of *Proceedings of Machine Learning Research*, 2242–2251.

Hand, L. (1930): "Nichols v. Universal Pictures Corp." 45 F.2d 119, u.S. Court of Appeals for the Second Circuit.

Heritage, S. (2025): "Tilly Norwood: how scared should we be of the viral AI 'actor'?" Accessed: 2026-01-29.

Koh, P. W. and P. Liang (2017): "Understanding Black-box Predictions via Influence Functions," in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, Sydney, Australia: PMLR, 1885–1894.

Leval, P. N. (1990): "Toward a Fair Use Standard," *Harvard Law Review*, 103, 1105–1136.

Livni, R., S. Moran, K. Nissim, and C. Pabbaraju (2024): "Credit Attribution and Stable Compression," in *Advances in Neural Information Processing Systems*, vol. 37.

Lutes, B. A., ed. (2025): *Identifying the Economic Implications of Artificial Intelligence for Copyright Policy: Context and Direction for Economic Research*, U.S. Copyright Office.

Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean (2013): "Distributed Representations of Words and Phrases and their Compositionality,"

in *Advances in Neural Information Processing Systems*, vol. 26.

Radford, A., J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever (2021): "Learning Transferable Visual Models From Natural Language Supervision," in *Proceedings of the 38th International Conference on Machine Learning*, PMLR, vol. 139 of *Proceedings of Machine Learning Research*, 8748–8763.

Rényi, A. (1962): "Théorie des éléments saillants d'une suite d'observations," in *Proceedings of the Colloquium on Combinatorial Methods in Probability Theory*, Aarhus, Denmark: Matematisk Institut, Aarhus Universitet, 104–115.

Rombach, R., A. Blattmann, D. Lorenz, P. Esser, and B. Ommer (2022): "High-Resolution Image Synthesis With Latent Diffusion Models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.

Rosen, S. (1981): "The economics of superstars," *The American Economic Review*, 71, 845–858.

Sag, M. B. (2024): "Copyright Safety for Generative AI," *Houston Law Review*, 61.

SAG-AFTRA (2024): "Keeping the Workforce Human in the Age of A.I." Accessed: 2026-01-29.

Schor, Z. (2024): "Andersen v. Stability AI: The Landmark Case Unpacking the Copyright Risks of AI Image Generators," .

Shapley, L. S. (1953): "A Value for n-Person Games," in *Contributions to the Theory of Games II*, ed. by H. W. Kuhn and A. W. Tucker, Princeton: Princeton University Press, 307–317.

Supreme Court of the United States (1879): "Baker v. Selden," 101 U.S. 99, u.S. Supreme Court decision.

——— (1994): "Campbell v. Acuff-Rose Music, Inc." 510 U.S. 569.

United States Congress (1976a): "17 U.S.C. §106: Exclusive Rights in Copy-

righted Works," United States Code.

——— (1976b): "17 U.S.C. §107: Limitations on Exclusive Rights: Fair Use," United States Code.

United States Court of Appeals for the Ninth Circuit (2003): "Satava v. Lowry," 323 F.3d 805.

United States Court of Appeals for the Second Circuit (1946): "Arnstein v. Porter," 154 F.2d 464.

——— (2015): "Authors Guild v. Google, Inc." 804 F.3d 202.

Weitzman, M. L. (1998): "Recombinant growth," *The Quarterly Journal of Economics*, 113, 331–360.

Xiong, P., H. Liu, T. Zhu, and P. S. Yu (2024): "A Survey on Machine Unlearning: Techniques and New Emerged Privacy Risks," *arXiv preprint arXiv:2406.06186*.

Yang, S. A. and A. H. Zhang (2025): "Generative AI and Copyright: A Dynamic Perspective," Working Paper.

Zhang, J. (2025): "Input out, output in: towards positive-sum solutions to AI-copyright tensions," *Journal of Intellectual Property Law & Practice*, 20, 594–604.