

# Creative Ownership in the Age of AI

Annie Liang\*      Jay Lu†

March 11, 2026

## Abstract

Copyright law focuses on whether a new work is “substantially similar” to an existing one, but generative AI can closely imitate style without copying content, a capability now central to ongoing litigation. We argue that existing definitions of infringement are ill-suited to this setting and propose a new criterion: a generative AI output infringes on an existing work if it could not have been generated without that work in its training corpus. To operationalize this definition, we model generative systems as closure operators mapping a corpus of existing works to an output of new works. AI generated outputs are *permissible* if they do not infringe on any existing work according to our criterion. Our results characterize structural properties of permissible generation and reveal a sharp asymptotic dichotomy: when the process of organic creations is light-tailed, dependence on individual works eventually vanishes, so that regulation imposes no limits on AI generation; with heavy-tailed creations, regulation can be persistently constraining.

---

\*Department of Economics, Northwestern

†Department of Economics, UCLA

# 1 Introduction

Generative artificial intelligence systems are increasingly used to produce creative content including novels (Brodsky, 2024), comics (Andersen, 2022; Schor, 2024), and even entire digital persons (Heritage, 2025). Trained on large bodies of existing works, these systems can generate new creations that resemble the tone or structure of particular authors, or images that evoke specific artistic traditions, without reproducing any identifiable prior work. Figure 1, for instance, juxtaposes one of Sarah Andersen’s comics with an AI-generated image created using her name as a prompt. The latter is not a copy of any specific Andersen comic, yet it clearly reproduces recognizable features of her style (Andersen, 2022).

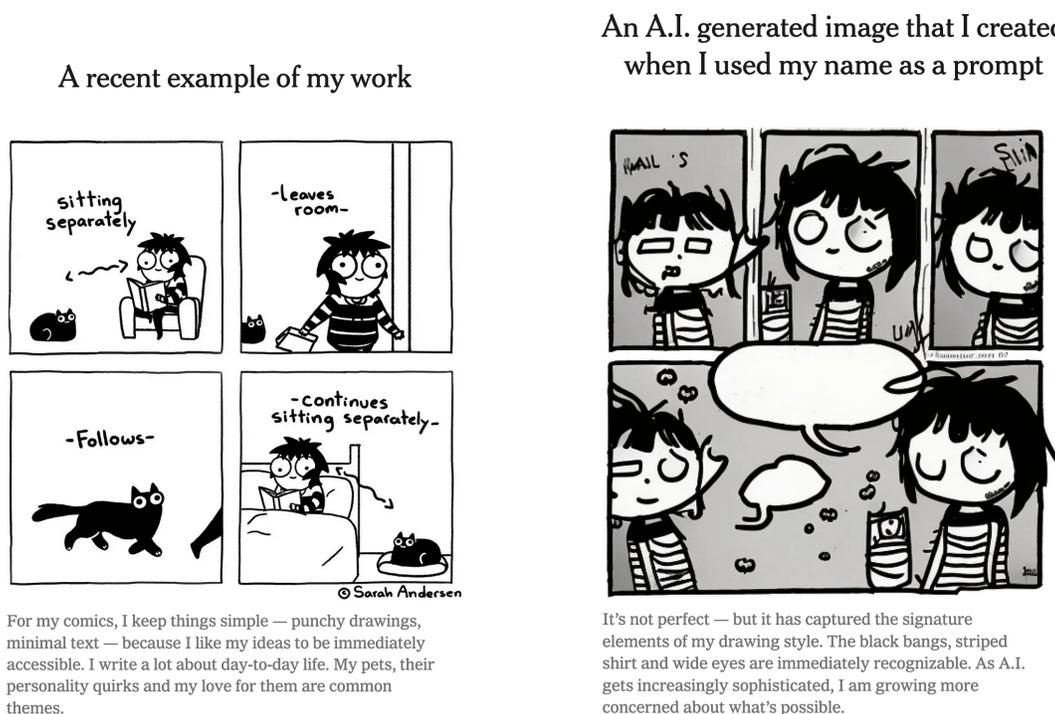


Figure 1: *Left*: an example of Andersen’s own work. *Right*: an A.I.-generated image created when her name was used as a prompt. Images taken from Andersen (2022).

Under standard copyright doctrine, which requires “substantial similarity,” such outputs need not constitute infringement. Yet the output may critically depend on the author’s work, in the sense that it could not have been generated without their

contributions to the training corpus. A wave of recent litigation—including *Andersen v. Stability AI* and *The New York Times v. OpenAI*—reflects growing discontent with existing notions of creative ownership.

This paper makes three primary contributions. First, we propose a new criterion for evaluating infringement on an existing work. Second, we characterize the structure of the permissible set of generations. Finally, we characterize the evolution of permissibility as the corpus of works grows large, thus studying the restrictiveness of regulation on generative AI in saturated markets.

In our framework (described in Section 2), a *creation* is modeled as a point in  $\mathbb{R}^d$ , which might be a novel, painting, or even artistic character. A *generator* maps a corpus of existing creations to a set of generable outputs. Formally, a generator is a closure operator: it satisfies preservation (existing creations remain generable), monotonicity (larger corpora expand the set of generable outputs), and idempotence (applying the generator to its own output produces nothing new). This abstraction encompasses various generative processes, from interpolation to feature recombination, while remaining agnostic to how the generator operates internally.

Our definition of infringement is counterfactual: a generated output *infringes on* an existing creation if and only if that output could not have been generated without this creation in the training corpus. The *violation set* consists of those outputs that infringe on at least one existing creation, and the *permissible set* is the complement.

Section 3 establishes structural properties of the permissible set. It expands monotonically as the corpus grows and is closed under further generation: combining permissible outputs cannot produce a violation (Proposition 1). We provide a sufficient condition for non-emptiness based on the *Radon number* from convex geometry (Corollary 1). We also show that comparative statics depend on the status of new creations: incorporating a previous violation strictly expands the permissible set, while incorporating an already-permissible creation leaves it unchanged (Proposition 3).

Creative corpora are not static: new works are constantly being produced and added to the training data of generative models. Section 4 turns to the evolution of the permissible set. Our main result characterizes how the permissible set evolves

as the corpus grows large. We show that when new creations are drawn from a distribution with light tails—so that extreme realizations are exponentially rare—the ratio of permissible to generable creations converges to one almost surely (Theorem 1). Intuitively, the corpus becomes sufficiently rich that multiple generative paths exist to any output, and no single creation in the corpus remains essential. By contrast, under heavy-tailed innovation, even very large corpora may continue to contain works that are genuinely distinct. As a result, a positive measure of outputs may remain violations indefinitely. Whether infringement risk vanishes or persists thus turns on whether creative production generates redundancy or is truly innovative.

These results relate to ongoing debates about generative AI and intellectual property. Standard infringement tests ask whether an output *resembles* a protected creation; our counterfactual criterion asks whether an output *relies* on that creation. The two can diverge: an output may be a violation without copying its source, and an output may closely resemble a creation yet remain permissible if alternative paths to generation exist.

## 1.1 Legal and Technological Context

### 1.1.1 Existing Copyright Law and Implications for Generative AI

Existing copyright law was not developed with generative artificial intelligence in mind, but its core rules nonetheless shape how generative 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 affect generative AI at two stages: model training and model output.

At the training stage, generative models are typically trained on large datasets that may include copyrighted works. Under traditional copyright principles, the “copying” of copyrighted works into data can constitute infringement unless it falls within an exception or limitation. In the United States, this question is largely governed by the fair use doctrine (United States Congress, 1976b). Courts applying fair use consider

factors such as whether the use is transformative, whether it is commercial, 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). 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).

While global policy responses to generative AI have largely focused on the input stage, there are arguments in favor of redirecting regulatory attention to outputs (Zhang, 2025). At this stage, infringement analysis follows more familiar rules. Copyright protects expression, not ideas, facts, or styles. An AI-generated output therefore infringes only if it is “substantially similar” to protected elements of a copyrighted work (United States Court of Appeals for the Second Circuit, 1946). Large models can sometimes memorize and reproduce portions of their training data (Carlini et al., 2021, 2023; Nasr et al., 2023), but most generative outputs do not constitute direct copying.

Taken together, existing copyright law does not treat generative AI as inherently unlawful. Instead, legal exposure depends on two main questions: how model training is evaluated under fair use or similar doctrines, and whether generated outputs cross the line from general inspiration into protected expression.

### **1.1.2 Criticisms and Ongoing Debates**

A central criticism of applying existing copyright law to generative AI concerns the distinction between style and expression. Copyright law has long drawn a firm line between protectable expression and unprotectable ideas, methods, systems, or artistic styles (17 U.S.C. §102(b); Supreme Court of the United States, 1879; Hand, 1930). As a result, copying an author’s or artist’s style—without copying protected expressive elements—does not amount to infringement, even when that style is distinctive or commercially valuable (United States Court of Appeals for the Ninth Circuit, 2003; Sag, 2024). This boundary reflects a longstanding policy choice to preserve creative

freedom, competition, and follow-on innovation (Leval, 1990).

Generative AI systems, however, are particularly well-suited to producing outputs that closely mimic stylistic features of human creators while avoiding literal copying. Large language models can produce text that resembles an author’s tone, pacing, or narrative structure, while image models can generate works that evoke an artist’s distinctive visual approach. Under current doctrine, such outputs generally fall outside copyright protection so long as they do not reproduce identifiable expressive elements from specific works (Sag, 2024; Mantegna, 2024).

Critics argue that this legal distinction fits poorly with the capabilities of generative AI. In many creative fields, style is often the primary source of value. A legal framework that protects expression but not style may allow AI systems to generate close substitutes that reduce demand for human creators, even when no infringement occurs in the legal sense. As a result, copyright law may permit significant displacement while offering limited remedies to affected creators (Pasquale and Sun, 2024; U.S. Copyright Office, 2025; de Rassenfosse et al., 2024).

Defenders of the existing framework respond that extending copyright protection to style would represent a sharp break from settled law. Granting exclusive rights over stylistic features risks blurring the distinction between ideas and expression and could give creators control over broad creative territories. Such an expansion would be difficult to define and enforce, and could inhibit follow-on creativity and innovation (Leval, 1990; Sag, 2024). From this view, the fact that many AI-generated outputs fall outside infringement doctrine reflects copyright’s intentionally limited scope rather than a flaw in the system.

This disagreement has fueled an ongoing debate over whether generative AI exposes a genuine gap in copyright protection or simply makes longstanding limits more visible. Some scholars argue that attempting to treat stylistic imitation as infringement would push copyright beyond its intended role (Mantegna, 2024). It is also highly subjective and whether a work meets the criterion of substantial similarity often on expert testimony. Others maintain that generative AI alters the scale, accuracy, and substitutability of stylistic imitation in ways that weaken the original

rationale for denying protection to style, even if the formal legal categories remain unchanged (Pasquale and Sun, 2024; de Rassenfosse et al., 2024).

## 1.2 Related Literature

Besides the legal literature discussed above, our paper contributes to a growing literature in economics on copyright and intellectual property in the context of generative AI (see Lutes (2025) and de Rassenfosse et al. (2024) for excellent surveys). Our paper in particular addresses the following problem, as described in de Rassenfosse et al. (2024):

“An essential issue in the copyright treatment of creative machines is the extent to which their outputs might be legally ‘derivative’ of the copyrighted works used to train the machine.”

Our paper presents a novel formalization of this concept, where a work is legally derivative of another if it could not have been generated without the original work in its training data.

Because large-scale generative models are recent, economic theory on their implications for intellectual property remains limited. Two early contributions are Gans (2024) and Yang and Zhang (2025), which study optimal copyright policy from different perspectives.

Gans (2024) emphasizes bargaining frictions and contracting feasibility. He demonstrates an important difference between “small” models trained on identifiable, contractible corpora and “large” models trained on web-scale data for which ex ante licensing is infeasible. In the former case, copyright protection strengthens creators’ incentives and improves welfare; in the latter, the welfare effects are ambiguous, motivating an ex post liability regime that permits training but allows harmed rights-holders to seek compensation.

Yang and Zhang (2025) adopt a dynamic perspective and study the interaction between two policy instruments: the fair-use standard governing training data and the copyrightability of AI-generated outputs. In their model, these policies jointly

shape creators’ incentives to supply new data and firms’ incentives to invest in model quality. A key insight is that generous fair use can be welfare-improving when training data are abundant but counterproductive when new human-created data are scarce.

Our focus is complementary but distinct. We study the more basic question of whether the previous notions of copyright are appropriate in a setting with generative AI, and how it might alternatively be defined. We view our proposed criterion as a building block that can be embedded into broader economic models of licensing, litigation, or regulation.

Finally, Theorem 1 connects to a broader literature on the nature of creative processes. Several papers have found that dividends to creative work—such as revenue from films, citations of academic papers, or attention to artworks—follow heavy-tailed distributions, with a small number of successes accounting for the bulk of observed outcomes (Rosen, 1981; Chung and Cox, 1994; De Vany and Walls, 1999). Like us, Drugov and Ryvkin (2020), considers the process of innovation itself, and shows that whether its distribution is light- or heavy-tailed has sharp implications for the optimal design of contests and incentives. More broadly, several theories explain why creative processes may exhibit heavy tails. Weitzman (1998) models idea generation as a process of recombinant expansion that can generate growth exceeding exponential rates, while Lutes (2025) argues that the use of AI as a tool in human creation may itself amplify variance in creative output.

## 2 Model

### 2.1 Setup

A *creation* is a vector  $x \in \mathbb{R}^d$ . A *corpus* is any (Borel measurable) subset of creations  $C \subset \mathbb{R}^d$ . 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 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 datasets like LAION-5B, which contains billions of images scraped from the web. 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—particularly of deceased or non-consenting actors—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 AI talent studio Xicoia’s actress Tilly Norwood (Heritage, 2025).

A *generator* is a technology for producing new creations from an existing corpus. Formally, it is a map from the set of all possible corpora  $\mathcal{C}$  into itself, which moreover satisfies three properties.

**Definition 1.** The map  $g : \mathcal{C} \rightarrow \mathcal{C}$  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:  $g(g(C)) = g(C)$

We interpret  $g(C)$  as the creations that are generable under  $g$  when  $g$  is trained on the works in  $C$ . A generator therefore (1) reproduces all input creations, (2) weakly expands the set of generable creations when the original corpus expands, and (3) is complete in that running it on its output produces no new creations. Such maps are also known as closure or consequence operators. Below, we present several examples.

**Example 4** (Convex Hull Generator). The AI firm generates creations by taking convex combinations of existing creations. We call the *convex hull* generator

$$g_{conv}(C) := \text{conv}(C).$$

which maps each  $C$  to its convex hull.

For example, GPT-5.2 suggests the following as a convex combination of Shakespeare’s *Sonnet 18* (“*Shall I compare thee to a summer’s day?*”) and Emily Dickinson’s *Because I could not stop for death*:

I weigh the hours against the light,  
They warm, then quietly recede;  
What time bestows, it also bends,  
Yet leaves a trace we read.

Of this creation, GPT-5.2 writes: “This poem softens Shakespeare’s rhetorical grandeur, relaxes Dickinson’s extreme compression, [and] blends temporal imagery with moderate regularity of rhythm. No single feature is taken wholesale from either poem; instead, tone, structure, and imagery lie between the two sources.”

**Example 5** (Splice Generator). The AI firm generates creations by recombining features from existing creations. We call this the *splice* generator

$$g_{splice}(C) := \{x \in \mathbb{R}^d : \forall k = 1, \dots, d, \text{ there is a } x' \in C \text{ s.t. } x_k = x'_k\}.$$

For example, GPT-5.2 suggests the following as a splice of the poems mentioned in Example 4:

The years—extend their measured grace—  
As Seasons—pass me by—  
So long as breath may utter praise—  
I walk—yet do not die.

Of this creation, GPT-5.2 writes: “[the] diction and thematic structure (“years,” “seasons,” endurance) are drawn from Sonnet 18, [while] meter, punctuation, stanza shape, and syntactic compression are drawn from Dickinson.”

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

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

For example, GPT-5.2 suggests the following as a box generation given the poems mentioned in Example 4:

Time offers warmth, then slips away,  
A season lent, not owned—  
Yet something spoken, lightly held,  
Outlasts the flesh and bone.

Of this creation, GPT-5.2 writes: “[this poem] inherits Shakespeare’s concern with endurance beyond time, [and] reflects Dickinson’s restraint and metaphysical inwardness, but smooths across multiple spliced variants to yield a coherent, unified style.”

The convex hull and box generators both have the property of outputting sets that are convex. They 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  $C \in \mathcal{C}$ .

Convex-valuedness ensures that if two creations can be generated, then so can any convex combination of them. This rules out gaps in the space of outputs and reflects the continuity of most real generative processes. For instance, AI systems that create digital actors or composite images can smoothly interpolate among existing examples. We characterize convex-valued generators in Appendix A, and show in particular that composing the convex-hull generator with any other generator yields a convex-valued generator. Thus one can view the assumption of a convex-valued generator as restricting to settings where the convex hull generator is available.

Our approach to modeling AI generation is intentionally broad and agnostic with respect to the specific algorithms employed in the process. In fact, our definition is not limited to generative AI but could equally apply to human creative processes as well. Moreover, the third property, idempotency, is to a certain degree without loss of generality. If the creative process  $g$  is not idempotent, we can define a new operator  $g' := g \circ g \circ \dots$  that repeatedly iterates the original process. With diminishing returns under repeated generative iterations,  $g'$  will approach idempotency. We can then treat  $g'$  as the effective generator in our analysis.

## 2.2 Permissible Generation

Our central definition of infringement specifies when a generated creation critically depends on the presence of some existing creation. To formalize this, we consider a counterfactual question: would the creation still be generable without that existing work in the corpus?

**Definition 3.** Fix a generator  $g$  and corpus  $C$ . For each existing creation  $x \in C$ :

- (1) The *x-permissible set* is  $p_g(x, C) := g(C \setminus \{x\})$ , i.e., all creations that are generable from  $C \setminus \{x\}$ .
- (2) The *x-violation set* is  $v_g(x, C) := g(C) \setminus p_g(x, C)$ , i.e., all creations that are generable from  $C$  but not from  $C \setminus \{x\}$ .

Intuitively, the *x-violation set* consists of creations whose generability relies essentially on the presence of the creation  $x$  in the corpus. If a generated creation is an *x-violation*, then the author of  $x$  has natural basis to claim infringement—the output would not exist without their contribution to the training data.

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

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

The *violation set* is

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

The permissible set consists of all creations that can be generated under  $g$  from  $C$ , and that would still be generable were any specific existing creation to be removed from the corpus. All remaining generable creations lie in the violation set and depend essentially on some existing creation.

Figure 2 illustrates these sets for the convex hull generator.<sup>1</sup> Given existing works  $x_1, x_2, \dots, x_5$ , any work in the convex hull of these points is generable under  $g_{conv}$ .

---

<sup>1</sup>For the convex hull generator, the permissible set corresponds to points with a Tukey depth of level 1.

But in Panel (a), the generability of the creations in  $V_{x_1}$  requires the presence of  $x_1$  as an input to the generator, while the creations in  $P_1$  do not. Panel (b) depicts the generable creations that are permissible, i.e., which do not depend on *any* given existing work. The remaining generable creations are violations of some existing work. For the convex hull generator, the violation set is loosely an outer shell of the space of generable creations. In particular, no extreme points of  $g(C)$  are permissible creations. This property is not however true for all generators; for example, Figure 3 shows that the extreme point  $x_4$  is permissible under the box generator.

In general, what the permissible set looks like depends on how the generator aggregates the underlying data and on what other existing works are present. Section 3 presents basic properties of the permissible set, and Section 4 studies the long-run evolution of the permissible set as new creations are added to the corpus.

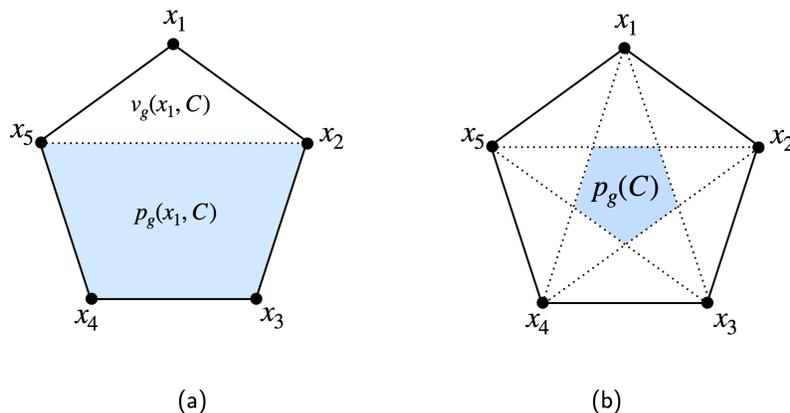


Figure 2: Consider the convex hull generator  $g_{conv}$  and the corpus  $C = \{x_1, x_2, \dots, x_5\}$ . *Panel (a)*:  $v_g(x_1, C)$  is the set of  $x_1$ -violations (which are only constructible using  $x_1$ ), and  $p_g(x_1, C)$  is its complement. *Right*:  $p_g(C) = \bigcap_{i=1}^5 p_g(x_i, C)$  is the set of points that are not  $x_i$ -violations for any  $i = 1, 2, \dots, 5$ .

### 2.3 Discussion

As reviewed in Section 1.1, existing copyright doctrine applies to generative AI at two distinct stages. First, training may violate copyright law if copyrighted works are included in the training corpus. Second, generated outputs may violate copyright law

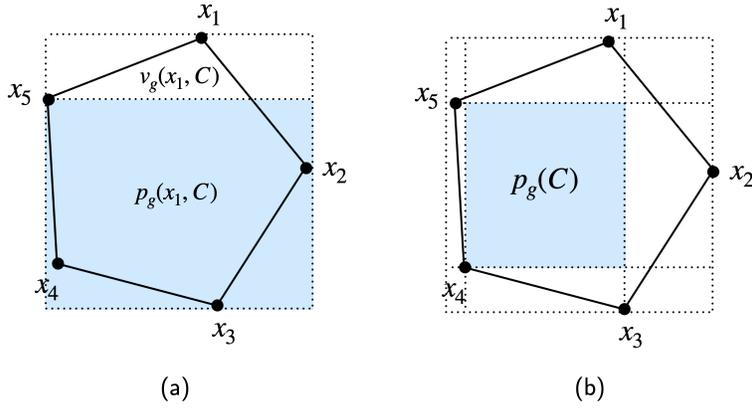


Figure 3: Replication of Figure 2 for the box generator  $g_{box}$ .

if they are *substantially similar* to an existing work. Our counterfactual definition departs from both of these standards. Below, we clarify how it relates to—and differs from—these prevailing legal approaches.

**Restrictiveness of the criterion.** Our definition can be either more or less restrictive than existing copyright standards. Relative to rules that prohibit training on copyrighted material altogether, our criterion is more restrictive in settings where AI systems are built by fine-tuning or composing existing models. For example, suppose a downstream generator can be written as  $g_2 = g' \circ g_1$ . Then even if the downstream firm does not itself train on copyrighted material, outputs of  $g_2$  can still be violations if they depend essentially on some copyrighted work used in the training corpus of  $g_1$ . At the same time, our criterion does not automatically treat the presence of copyrighted works in the training corpus as problematic: inclusion is permitted so long as no generated output depends essentially on any particular copyrighted work.

Compared to the *substantial similarity* standard applied to outputs, our definition again cuts in both directions. An output that is clearly not a copy of any existing work may nevertheless constitute a violation if its generation critically depends on a specific work in the corpus. Conversely, close resemblance to an existing work does not necessarily imply a violation if the output can be generated through multiple alternative paths that do not rely on that work.

Our criterion is intended as a complement rather than a substitute for existing copyright doctrines. Effective regulation may ultimately require the simultaneous enforcement of all three standards.

**Implementation.** Unlike “substantial similarity,” counterfactual generability has a clear meaning in principle: an output depends on a given work if it would no longer be producible once that work is removed from the training data. In the most direct implementation, this would involve retraining the model on  $C \setminus \{x\}$  and checking whether the output persists. While this brute-force method is usually computationally infeasible, a growing literature on *machine unlearning* develops practical techniques for approximating the effect of removing individual datapoints from trained models (Bourtole et al., 2019; Xiong et al., 2024). Closely related are influence-function methods, which estimate the marginal contribution of a single training example to a model’s predictions without full retraining (Koh and Liang, 2017).

An alternative approach is to treat counterfactual generability as an evidentiary question 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 when the plaintiff’s work is absent from the training data, that absence provides evidence that the work played a necessary role.

Finally, although our analysis is motivated by copyright law, the criterion we propose extends beyond strictly legal settings and beyond a binary determination of infringement. More generally, dependence (as we have defined it) can be understood as a measure of whether a prior work meaningfully contributed to a given output. In academic research, for example, such dependence might inform whether a work should be cited as a predecessor. In other contexts, it could determine whether a creator has a legitimate claim to recognition or to a share of the rents generated by downstream use.

Our framework also applies to settings in which the underlying objects are not literally copyrighted. Consider the generation of digital actors or persons: while an individual’s identity is not protected by copyright, it may nonetheless be morally or socially relevant to ask whether a generated character appropriates an existing person’s voice, personality, or likeness. In this sense, the framework provides a general approach to assessing attribution and appropriation, even outside the formal boundaries of intellectual property law.

### 3 Properties of Permissible Creation

This section clarifies the structure of the permissible set  $p_g(C)$ . Section 3.1 shows that permissibility behaves like a “core” of generative capability, in the sense that the permissible set is monotone and self-stable. Section 3.2 provides a sufficient condition for nonemptiness of the permissible set. Section 3.3 considers comparative statics in  $p_g(C)$  as the corpus  $C$  is augmented with an additional creation. All proofs can be found in the Appendix.

#### 3.1 Structure of the Permissible Set

**Proposition 1.** *For every generator  $g$ , 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 creations. 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 set of permissible creations. Intuitively, this is because if a creation can be generated in a way that does not rely on any particular existing work, then adding additional source works cannot introduce such a dependence. For example, if an AI novel generator is trained

on novels from the 19th century and produces an output that does not essentially rely on any 19th century novel, then adding 20th century novels to the AI’s training set will not make this novel dependent on some novel in the expanded corpus.

Part (2) says that the permissible set is closed under the generator: applying  $g$  to permissible creations yields only permissible creations. 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 works.

### 3.2 Non-Emptiness of the Permissible Set

We now turn to the question of existence: under what conditions is the permissible set non-empty? In general, it may be that *all* generable creations are violations of some existing work, as illustrated in the following simple example:

**Example 7.** 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:  $p_g(C) = p_g(x_1, C) \cap p_g(x_2, C) = \emptyset$ .

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  $A, B \subseteq C$  ( $A \cap B = \emptyset$ ) satisfying:

$$g(A) \cap g(B) \neq \emptyset$$

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

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  $R(g_{conv}) = d + 2$ , as illustrated in Figure 4.

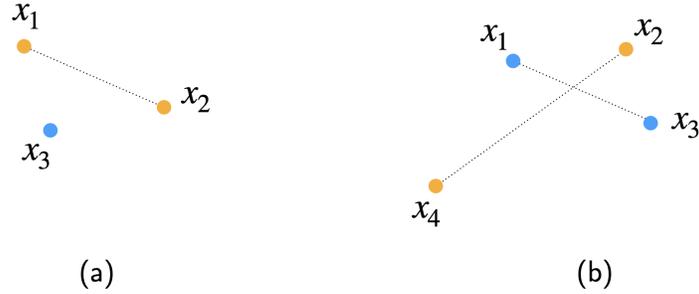


Figure 4: Let  $d = 2$ , so that creations are points in  $\mathbb{R}^2$ . *Panel (a)*: No three points in general position can be separated into disjoint subsets whose convex hulls overlap, so the Radon number is at least three; *Panel (b)*: Every set of four (or more) points can be separated in this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

**Proposition 2.** *Fix any generator  $g$ . If  $|C| \geq R(g)$ , then the permissible set is non-empty:  $p_g(C) \neq \emptyset$ .*

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

### 3.3 Comparative Statics

The next result tells us how permissibility changes when new creations enter the world and are included in the set of existing works whose intellectual property must be respected. By Part (a) of Proposition 1 (monotonicity of the permissible set), we know that  $p_g(C) \subseteq p_g(C \cup \{x\})$  for any addition of  $x$ , so the set of permissible creations weakly expands. The proposition below provides a more precise account depending on  $x$ .

**Proposition 3.** *Fix a generator  $g$  and corpus  $C$ . Let  $x$  be a new creation not already in  $C$ . Then exactly one of the following cases applies:*

- (1) *If  $x$  is permissible, the permissible set remains unchanged:*

$$x \in p_g(C) \implies p_g(C) = p_g(C \cup \{x\})$$

(2) If  $x$  is a violation, the permissible set strictly expands:

$$x \in v_g(C) \implies p_g(C) \subsetneq p_g(C \cup \{x\})$$

(3) If  $x$  is not generable, the effect on the permissible set is ambiguous:

$$x \notin g(C) \implies p_g(C) \subseteq p_g(C \cup \{x\})$$

with equality or strict inclusion both possible.

Part (1) says that including a new permissible generation in the set of accepted existing works does not expand the set of generations that are permissible. On the other hand, Part (2) says that adding a *violation* to the set of existing works strictly expands the permissible set. Intuitively, when a creation that was previously a violation is added to the corpus, it becomes a primary source in its own right. This breaks the counterfactual dependence on whatever original work it had relied upon: the generator now has a second path to that part of the creative space.

Finally, Part (3) considers the introduction of a genuine innovation which was not generable from the original corpus. In this case the effect on the permissible set is ambiguous and we provide an example of each case below.

**Example 8.** 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 creation  $x_4 = (0, 1)$ , 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 creation  $x_4 = (1, 1)$ , as depicted in Panel (b) of Figure 5. Then

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

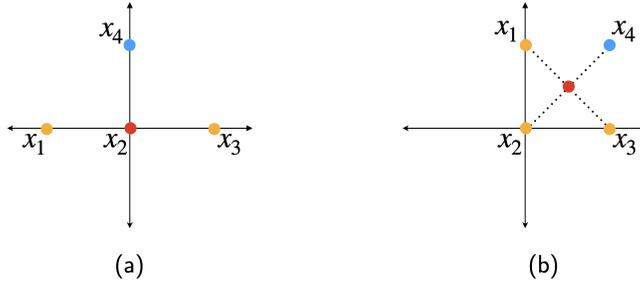


Figure 5: Two examples illustrating how adding a work can affect the permissible set under the convex hull generator. In Panel (a), adding  $x_4$  leaves the permissible set unchanged:  $p_g(\{x_1, x_2, x_3\}) = \{x_2\} = p_g(\{x_1, x_2, x_3, x_4\})$ . In Panel (b), adding  $x_4$  strictly expands the permissible set:  $p_g(\{x_1, x_2, x_3\}) = \emptyset$  while  $p_g(\{x_1, x_2, x_3, x_4\})$  is the singleton intersection point of the line segments  $\overline{x_1x_3}$  and  $\overline{x_2x_4}$ , as indicated in red.

so the inclusion of  $x_4$  strictly expands the permissible set.

## 4 Permissible Creations in the Long Run

The previous section characterized the permissible set for given corpora. But creative corpora are not static: new works are constantly being produced and added to the training data of generative models. This raises a natural question: when the corpus is large, does infringement become a vanishing concern—with the permissible set eventually encompassing nearly all generable creations—or does it remain a persistent feature of generative technology?

This section shows that the answer depends on the nature of creative innovation. When new works push the frontier only “gradually” (in a suitably defined sense), the permissible ratio converges to one: almost everything generable becomes permissible. But when breakthrough innovations regularly occur, a positive fraction of generable creations can remain violations even as the corpus grows without bound.

To understand the extent of permissible generation, we will consider the fraction of generable creations that are permissible. Formally, let  $\text{Vol}(A)$  denote the Lebesgue

measure (or *volume*) of a set  $A \subset \mathbb{R}^d$ . The *permissible ratio* is defined as

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

i.e., the proportion by volume of permissible creations that lie in the generable set.<sup>2</sup>

We study the following creative process. Let  $X_i$  be an i.i.d. random variable on  $\mathbb{R}^d$  with some strictly positive density. For each  $n \in \mathbb{N}$ , define the random corpus  $C_n = \{X_1, \dots, X_n\}$ .

Recall that a random variable (on  $\mathbb{R}$ ) has *light tails* if its hazard rate is increasing.<sup>3</sup> In higher dimensions, tail behavior must account for how mass decays in different directions. Let  $U := \{u \in \mathbb{R}^d : \|u\| = 1\}$  denote the unit sphere in  $\mathbb{R}^d$ . We express  $X \in \mathbb{R}^d$  in terms of its polar coordinates. The radial coordinate is

$$r = \|X\| \in \mathbb{R}_+,$$

i.e., the Euclidean norm of  $X$ , and the angular coordinate is

$$\theta = \frac{X}{\|X\|} \in U,$$

the unit vector pointing in the direction of  $X$ . Every  $X \neq 0$  admits the decomposition  $X = r\theta$ .

**Definition 6.**  $X$  has *uniform light tails* if

- (1)  $\|X\|$  has light tails
- (2) there exists a  $c > 0$  such that for all  $r > 0$ , the conditional density satisfies

$$f(\theta : \|X\| \geq r) \geq c$$

The uniform light tails condition captures domains where expansion of the creative frontier is gradual. Formally, condition (1) ensures that extreme realizations become increasingly rare, while condition (2) rules out “pockets” of transient innovation that

---

<sup>2</sup>When  $g$  is convex-valued, then generically  $\text{Vol}(g(C)) > 0$  and so this ratio is well-defined.

<sup>3</sup>Formally, if we let  $F$  denote the cdf of the random variable and  $h(t) = \frac{f(t)}{1-F(t)}$ , then  $F$  has light tails if  $h(\cdot)$  is increasing.

concentrate in particular directions. Together, these conditions mean that as the corpus grows, each new work is unlikely to lie far outside the convex hull of existing works. The standard multivariate Normal distribution satisfies uniform light tails, as does any standard elliptical distribution from a light-tailed family.

We also require a technical continuity condition on the generator.

**Definition 7.**  $g$  is *uniform lower semicontinuous (ULS)* if for every  $\varepsilon > 0$ , there is a  $\delta > 0$  such that for all  $C \in \mathcal{C}$  where  $0 \in C$ ,

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

Intuitively, ULS says that contracting the corpus slightly causes only a slight contraction in the generable set. This ensures that the generative frontier does not depend too sensitively on the precise location of any single training point. Positive homogeneity—i.e.,  $g(\alpha C) = \alpha g(C)$  for all  $\alpha > 0$ —is a sufficient condition, and is satisfied by the convex hull, splice, and box generators.

We now state our main asymptotic result.

**Theorem 1.** *Fix a convex-valued and ULS generator  $g$ . If  $X_i$  has uniform light-tails, then a.s.*

$$\lim_{n \rightarrow \infty} r_g(C_n) = 1$$

Theorem 1 says that in creative domains with gradual innovation, copyright infringement becomes a vanishing concern as the corpus grows. The violation set—creations whose generability depends essentially on some specific existing work—shrinks to measure zero relative to the full generable set. Intuitively, when new works arrive gradually, the corpus eventually becomes rich enough that every generable creation has multiple “generative paths” leading to it; removing any single work from the training data leaves the creation still reachable.

This result implies that in creative markets where innovation is incremental, regulation (in the sense of our proposed criterion) only has bite in the early stages of a creative domain. If it is imposed sufficiently late, or if sufficiently many generations

have already flooded the market, then very few creators will be able to defend critical dependence on their work.

The theorem also clarifies which types of creative works are most vulnerable to losing effective intellectual property protection within our framework. Works located in the interior of the creative space—surrounded by many similar works—quickly become dispensable as training data accumulate. By contrast, works at the frontier, in sparsely populated regions of the space, may remain essential for generating nearby outputs even as the corpus grows large.

Our result extends beyond creative processes that satisfy uniform light tails. For example, generic multivariate random variables can be normalized to identity covariance via a linear transformation (whitening). Since Theorem 1 is preserved under linear transformations, any elliptical distribution from a light-tailed family (e.g., Normal, exponential) satisfies the premise. What the theorem rules out are heavy-tailed distributions, where breakthrough innovations regularly occur. The following example demonstrates that in such domains, a positive fraction of generable creations remain violations even asymptotically.

**Example 9** (Heavy tails). Suppose  $X_i$  are Pareto distributed on  $\mathbb{R}_+$  and let  $g = g_{\text{conv}}$  be the convex hull generator. The Pareto distribution is the canonical model of heavy-tailed phenomena, arising in contexts where “superstar” outcomes regularly occur—bestselling novels, viral content, or paradigm-shifting artistic innovations. Note that

$$r_g(C_n) = \frac{\text{Vol}(p_g(C_n))}{\text{Vol}(g(C_n))} \leq \frac{\max\{X_1, \dots, X_{n-1}\}}{\max\{X_1, \dots, X_n\}}$$

For Pareto distributions, the ratio of successive maxima does not converge to one almost surely. The  $n$ -th observation has a non-negligible probability of dramatically exceeding all previous observations, no matter how large  $n$  becomes. Consequently,  $r_g(C_n)$  remains bounded below one: the most recent breakthrough always defines a substantial portion of the violation set.

This contrast between light-tailed and heavy-tailed domains implies that, in creative fields characterized by occasional breakthroughs—such as avant-garde art and literary fiction—protection for individual works may retain its importance even as

corpora grow large. Creators of paradigm-shifting works can continue to have legitimate infringement claims because their contributions remain essential for generating outputs near the creative frontier. By contrast, in mature creative markets where innovation is largely incremental—possibly including genre fiction and commercial movies—the marginal contribution of any single work diminishes rapidly as the corpus expands, making individual infringement claims correspondingly weaker.

## 5 Extension: General Violations

The analysis so far has focused on a setting in which the entire training corpus consists of protected works and infringement is 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. A novelist may sue over their entire catalog; a photographer may claim infringement of their portfolio; 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 family of subsets  $\mathcal{A} \subseteq \mathcal{C}$ . In the main case we have considered thus far,  $\mathcal{A}$  is the collection of all singleton works. We allow the collection of protected works to be different from the collection of all singleton works and extend the notions of permissibility and violation correspondingly.

**Definition 8.** 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 set of protected works  $A \in \mathcal{A}$  in the collection,  $g(C \setminus A)$  consists of all outputs that can be generated without using any input from  $A$ , i.e. permissible with respect to  $A$ . The permissible set  $p_g(\mathcal{A}, C)$  thus consists of all works permissible with

respect to every set in the collection  $\mathcal{A}$ .

Proposition 1 extends fully.

**Proposition 4.** *For every generator  $g$  and collection  $\mathcal{A}$ , the permissible set satisfies*

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

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 works yields only permissible works.

We now consider what happens to the permissible set under different collections. We say a collection or protected works is *richer* than another if for every set of protected works in the former, there is a larger set of protected works in the latter.

**Definition 9.** 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  $B \subseteq A$ .

There are two notions of what it means for a collection to have “more” protected works. First, the collection could simply consist of a greater number of protected works. Second, the collection could consist of larger groups of protected works, e.g., novelists coming together to sue as a coalition. The first notion corresponds to greater volume while the second to greater aggregation. Our definition of richness captures both.

The following result shows that the permissible set is smaller in richer collections.

**Proposition 5.** *Fix a generator  $g$  and corpus  $C$ . If a collection  $\mathcal{A}$  is richer than a collection  $\mathcal{B}$  then*

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

To illustrate the implications of this result, consider what happens when two human creators combine their individual creations to form a coalition. An immediate corollary is that violations are *superadditive*: the violation of a union can exceed the union of their individual violations.

**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.*

Superadditivity implies that when multiple creators join a class action, their combined violation set can exceed the sum of their individual claims. This strengthens the collective bargaining position of creator coalitions and suggests that group licensing may command a premium over the sum of individual licenses.

Another implication of Proposition 5 is that removing protected works from the collection naturally expands the permissible set. Combining this with Theorem 1, we have the following asymptotic result.

**Corollary 3.** *Fix a convex-valued and ULS  $g$ , and suppose the collection  $\mathcal{A}$  consists of some subset of all singleton works. If  $X_i$  has uniform light tails, then a.s.*

$$\lim_{n \rightarrow \infty} \frac{\text{Vol}(p_g^{\mathcal{A}}(C_n))}{\text{Vol}(g(C_n))} = 1$$

In other words, Theorem 1 is conservative: its implication that AI generation becomes unrestricted in the long run under light-tailed innovation assumes that every work in the corpus is protected. If there are works in the corpus that are unprotected (e.g., public domain), then this renders AI generation more permissible and strengthens the asymptotic result.

A natural question is to ask what happens if we permit group violations. In this case, we conjecture that as long as the number of individuals in each group is finitely bounded, then Theorem 1 would still hold. This is because as  $n$  grows large, any finite number of individuals banding together as a group would still be negligible and the same asymptotic forces will prevail. On the other hand, if the group size grows unboundedly with the corpus—for instance, if it includes all works by authors who have ever contributed—the limiting behavior depends on how it grows relative to  $n$ . Characterizing this relationship is an open question.

## 6 Conclusion and Open Questions

Our paper characterizes the structure of permissible generation but takes as given both the generator and the distribution of creative works. A natural extension is to endogenize these objects. If creators anticipate how generative systems will use existing works, they may adjust what they produce. Analyzing equilibrium in this environment requires modeling the strategic interaction between creators, who choose where in the creative space to locate new works, and AI firms, who choose which generative capabilities to develop. One plausible conjecture is that creators concentrate effort near the frontier of the creative space rather than in its interior. Such strategic positioning would thicken the tails of the creation distribution and could sustain a larger violation set than would arise under non-strategic creation.

A second open question concerns social optimality—specifically, what level of permissiveness would we ideally sustain? At one extreme, if the permissible ratio converges to one then nearly all generative outputs are non-infringing, and this may weaken creators' *ex ante* incentives to produce new works. At the other extreme, a large and persistent violation set preserves creator rents but limits the social gains from generative technologies. Our framework and proposed definition provide a foundation for formalizing and evaluating these tradeoffs.

# Appendix

## A Preliminary Results

**Proposition 6.** *The following are equivalent:*

- (1)  $g$  is convex-valued
- (2)  $g = g_{conv} \circ g$
- (3)  $g = g \circ g_{conv}$

*Proof.* We will first show that  $g$  is convex-valued if and only if  $g = g_{conv} \circ g$ . In one direction, when  $g$  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

$$g(C) \subseteq g(g_{conv}(C)) \tag{1}$$

by monotonicity of  $g$ . Moreover,

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

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

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

where the last equality follows from idempotency. Combining (1) and (2) yields  $g(g_{conv}(C)) = g(C)$  as desired.

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

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

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

**Corollary 4.** *The convex hull generator is minimal among all convex-valued generators, i.e.,*

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

for all convex-valued generators  $g$ .

*Proof.* For any convex-valued generator  $g$ ,

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

where the set inclusion follows by containment and the inequality follows from Proposition 6.  $\square$

## B Proofs of Results in Section 3

### B.1 Proof of Proposition 1

*Part (a):* 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

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

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

*Part (b):* In one direction,  $p_g(C) \subseteq g(p_g(C))$  by containment. 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

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

for each  $y \in C$ . So

$$g\left(\bigcap_{x \in C} E_x\right) \subseteq \bigcap_{y \in C} g(E_y) \tag{3}$$

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

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

Combining this with (3) yields

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

completing the proof.

## C Proof of Proposition 2

Since  $|C| \geq R(g)$ , by definition there exist disjoint subsets  $A, B \subseteq C$  such that  $g(A) \cap g(B) \neq \emptyset$ . Let  $x$  be an element in this intersection. We show that  $x \in p_g(C)$ .

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

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

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

## D 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.,  $g_{conv}(A) \cap g_{conv}(B) \neq \emptyset$ . By Proposition 2, the permissible set for the convex hull generator is non-empty,  $p_{g_{conv}}(C) \neq \emptyset$ .

Recall from Corollary 4 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 non-empty, it follows that  $p_g(C)$  is non-empty as well.

## E Proof of Proposition 3

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

$$x \in p_g(C) \subseteq g(C \setminus y) \tag{4}$$

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

$$\begin{aligned}
(C \cup \{x\}) \setminus y &= (C \setminus y) \cup \{x\} \\
&\subseteq (C \setminus y) \cup g(C \setminus y) && \text{by (4)} \\
&\subseteq g(C \setminus y) \cup g(C \setminus y) && \text{by containment} \\
&= g(C \setminus y)
\end{aligned}$$

Applying  $g$  to both sides yields

$$g((C \cup \{x\}) \setminus y) \subseteq g(g(C \setminus y)) = g(C \setminus y) \quad (5)$$

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

Thus,

$$\begin{aligned}
p_g(C \cup \{x\}) &= g(C) \cap \bigcap_{y \in C} g((C \cup \{x\}) \setminus y) \\
&\subseteq g(C) \cap \bigcap_{y \in C} g(C \setminus y) && \text{by (5)} \\
&= \bigcap_{y \in C} g(C \setminus y) = p_g(C)
\end{aligned}$$

as desired.

*Part (b):* 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 (a) of Proposition 1. 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  $y \in C$  such that

$$x \in v_g(y, C \cup \{x\}) = g(C \cup \{x\}) \setminus g((C \cup \{x\}) \setminus \{y\}) \quad (6)$$

First suppose  $x \neq y$ . Then

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

where the set inclusion follows from containment. But this contradicts (6).

Suppose instead  $x = y$ . Then

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

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

*Part (c):* The set inclusion  $p_g(C) \subseteq p_g(C \cup \{x\})$  is a direct consequence of Part (a) of Proposition 1, and the possibility of both equality and strict inclusion is shown in Example 8.

## F Proof of Theorem 1

The proof proceeds in the following steps. We first state two key intermediate results. Lemma 1 shows that almost surely, for all sufficiently large  $n$ , each new point  $X_n$  lies inside a slight radial expansion of the convex hull of the previous points  $X_1, \dots, X_{n-1}$ . Lemma 2 shows that if removing the *last* prior observation shrinks the convex hull by at most a multiplicative factor  $1 + \varepsilon$ , then the same is true for removing *any* prior observation.

Throughout this proof, let  $U$  denote the unit sphere in  $\mathbb{R}^d$  and let  $\lambda$  denote the surface area measure on  $U$ . Recall that  $C_n = \{X_1, \dots, X_n\}$  is the size- $n$  corpus, and define  $K_n = \text{conv}(C_n)$  to be its convex hull.

**Lemma 1.** For any  $\varepsilon > 0$ , almost surely there exists  $N < \infty$  such that for all  $n \geq N$ ,

$$X_n \in (1 + \varepsilon)K_{n-1}.$$

**Lemma 2.** Suppose that  $(Z_i)_{i \geq 1}$  is a sequence of positive random variables such that:

- (1) For any  $\varepsilon > 0$ , almost surely there exists  $N(\varepsilon)$  such that  $Z_n \leq (1 + \varepsilon) \max_{i \leq n-1} Z_i$  for all  $n \geq N(\varepsilon)$ .
- (2)  $\max_{i \leq n} Z_i \rightarrow \infty$  almost surely as  $n \rightarrow \infty$ .

Then almost surely there exists  $\tilde{N}(\varepsilon)$  such that for all  $n \geq \tilde{N}(\varepsilon)$  and all  $j \in \{1, \dots, n\}$ ,

$$\max_{i \leq n} Z_i \leq (1 + \varepsilon) \max_{i \in \{1, \dots, n\} \setminus \{j\}} Z_i.$$

By Lemma 1, for any  $\frac{\delta}{1-\delta} > 0$  there is an a.s. event such that for all sufficiently large  $n$

$$X_n \in \left(1 + \frac{\delta}{1-\delta}\right) K_{n-1}$$

Equivalently, for all  $u \in U$

$$u \cdot X_n \leq \left(1 + \frac{\delta}{1-\delta}\right) \max_{i \leq n-1} u \cdot X_i$$

Let  $\sigma(\cdot, u)$  denote the support function in direction  $u$ . Then almost surely for all sufficiently large  $n$ , all  $u \in U$ , and all  $j \in \{1, \dots, n\}$ ,

$$\begin{aligned} \sigma((1-\delta)\text{conv}(C_n), u) &= (1-\delta)\sigma(C_n, u) && \text{by homogeneity of } \sigma \\ &= (1-\delta) \max_{i \leq n} u \cdot X_i \\ &\leq \max_{i \in \{1, \dots, n\} \setminus \{j\}} u \cdot X_i && \text{by Lemma 2} \\ &= \sigma(C_n \setminus \{X_j\}, u) \\ &= \sigma(\text{conv}(C_n \setminus \{X_j\}), u) && \text{since } \sigma(\text{conv}(\cdot), u) = \sigma(\cdot, u). \end{aligned}$$

Equivalently,  $(1-\delta)\text{conv}(C_n) \subseteq \text{conv}(C_n \setminus X_j)$ , which implies

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

by monotonicity of  $g$ , and finally that

$$g((1-\delta)C_n) \subseteq g((C_n \setminus X_j))$$

by convex-valuedness of  $g$  (see Proposition 6). Intersecting over  $j$  yields

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

Moreover,  $p_g(C_n) \subseteq g(C_n)$  by definition, and therefore,

$$g((1-\delta)C_n) \subseteq p_g(C_n) \subseteq g(C_n).$$

This implies that

$$\text{vol}(g((1-\delta)C_n)) \leq \text{vol}(p_g(C_n)) \leq \text{vol}(g(C_n))$$

Dividing through by  $\text{vol}(g(C_n))$  (and noting  $\text{vol}(g(C_n)) > 0$  for large  $n$ ), we also have

$$\frac{\text{vol}(g((1-\delta)C_n))}{\text{vol}(g(C_n))} \leq \frac{\text{vol}(p_g(C_n))}{\text{vol}(g(C_n))} \leq 1$$

Consider some  $\varepsilon > 0$ . By assumption,  $g$  is ULS and moreover  $0 \in \text{conv}(C_n)$  for large  $n$  without loss. We can then find some  $\delta$  such that

$$(1-\varepsilon)g(\text{conv}(C_n)) \subseteq g((1-\delta)\text{conv}(C_n))$$

Since  $g$  is convex-valued, we also have

$$(1-\varepsilon)g(C_n) \subseteq g((1-\delta)C_n)$$

by Proposition 6. This implies that

$$(1-\varepsilon)^d \text{vol}(g(C_n)) = \text{vol}((1-\varepsilon)g(C_n)) \leq \text{vol}(g((1-\delta)C_n))$$

so

$$(1-\varepsilon)^d \leq \frac{\text{vol}(g((1-\delta)C_n))}{\text{vol}(g(C_n))} \leq \frac{\text{vol}(p_g(C_n))}{\text{vol}(g(C_n))}$$

Taking  $\varepsilon \rightarrow 0$  yields the result.

## F.1 Proof of Lemma 1

Since  $K_{n-1} = \text{Conv}(X_1, \dots, X_{n-1})$  is (by definition) a convex set, it suffices to control the support function in every direction. That is, we will prove that for every  $u \in U$ ,

$$u \cdot X_n \leq (1+\varepsilon) \max_{i \leq n-1} u \cdot X_i.$$

We separate this argument into two steps. First, we control the magnitude of the new draw  $X_n$ , showing that its norm is eventually bounded by a radius  $r_n$  determined by the tail behavior of  $\|X\|$ . Second, we control the geometry of the existing convex hull

$K_{n-1}$  at this same scale, showing that the hull already extends a distance comparable to  $r_n$  in every direction. Combining these two facts implies that no new point of magnitude comparable to  $r_n$  can exceed the previous maximum in any direction, yielding  $X_n \in (1 + \varepsilon)K_{n-1}$  for all sufficiently large  $n$ .

### F.1.1 Controlling the magnitude of the new draw $X_n$

Fix an arbitrary constant  $K > 0$  (to be chosen later). Let  $F$  be the cdf of  $|X|$  and  $\bar{F} := 1 - F$ . For all  $n$  such that  $K \frac{\log n}{n} \in (0, 1]$ , define  $r_n \in \mathbb{R}_+$  by

$$\bar{F}(r_n) = K \frac{\log n}{n}.$$

Such an  $r_n$  exists by continuity of  $\bar{F}$ .

**Lemma 3.** Almost surely for large  $n$ ,

$$|X_n| \leq (1 + \delta) r_n$$

*Proof.* Fix  $\delta > 0$  and write  $\bar{F}(t) = \mathbb{P}\{|X_n| > t\}$ . By the light-tail assumption,  $\phi(t) := \log \bar{F}(t)$  is concave on  $[0, \infty)$ .<sup>4</sup> Since  $\phi$  is concave and  $\phi(0) = \log \bar{F}(0) = 0$ , we have  $\phi((1 + \delta)t) \leq (1 + \delta)\phi(t)$  for all  $t \geq 0$ . Therefore,

$$\begin{aligned} \mathbb{P}\{|X_n| > (1 + \delta)r_n\} &= \bar{F}((1 + \delta)r_n) = \exp(\phi((1 + \delta)r_n)) \\ &\leq \exp((1 + \delta)\phi(r_n)) = (\bar{F}(r_n))^{1+\delta} \leq K^{1+\delta} \left(\frac{\log n}{n}\right)^{1+\delta}. \end{aligned}$$

Thus,

$$\sum_n \mathbb{P}\{|X_n| > (1 + \delta)r_n\} \leq K^{1+\delta} \sum_n \left(\frac{\log n}{n}\right)^{1+\delta} < \infty$$

as  $\delta > 0$ . Since this converges, we have the desired result by the first Borel-Cantelli lemma.  $\square$

---

<sup>4</sup>This is equivalent to a nondecreasing hazard rate  $h(t) = f(t)/\bar{F}(t)$ , since  $\phi'(t) = -h(t)$ .

### F.1.2 Directional coverage at the frontier

Let  $V \subseteq U$  be a finite  $\theta$ -cover of the sphere, i.e., a finite set of directions such that every  $u \in U$  lies within angle  $\theta$  of some  $v \in V$ :

$$u \cdot v \geq \cos(\theta) \quad \forall u \in U.$$

Fix  $v \in U$ . We say that an observation  $X_i$  is *v-good at level n* if

$$|X_i| \geq r_n \quad \text{and} \quad \frac{X_i}{|X_i|} \cdot v \geq \cos(\theta).$$

Thus, a *v-good* observation is both large in magnitude and aligned with the direction  $v$  up to angle  $\theta$ .

By the uniform light-tails assumption, there exists a constant  $c > 0$  such that for all  $r > 0$  and all  $u \in U$ ,

$$\mathbb{P}\left(\frac{X}{\|X\|} \cdot u \geq \cos(\theta) \mid \|X\| \geq r\right) \geq c s_u.$$

In particular, this bound holds for all  $v \in V$ . Moreover, the surface area of the spherical cap of directions within angle  $\theta$  of  $v$  is strictly positive:

$$s_v := \lambda\{u \in U : u \cdot v \geq \cos(\theta)\} > 0 \quad \forall v \in V.$$

Since  $V$  is finite, we may choose a constant  $K > 0$  such that

$$c s_v \geq \frac{4}{K} \quad \forall v \in V.$$

For each  $v \in V$ , let  $A_{n-1}(v)$  denote the event that there exists at least one *v-good* observation among  $\{X_1, \dots, X_{n-1}\}$ . By independence of  $X_i$ ,

$$\mathbb{P}(A_{n-1}(v)^c) = (1 - p_n)^{n-1},$$

where

$$p_n := \mathbb{P}\left(\|X\| \geq r_n, \frac{X}{\|X\|} \cdot v \geq \cos(\theta)\right), \quad X := X_1.$$

Writing  $f$  for the marginal density of  $\|X\|$  and  $f_r$  for the conditional density of

$X/\|X\|$  given  $\|X\| = r$ , we have

$$\begin{aligned} p_n &= \int_{r_n}^{\infty} \left( \int_{u \cdot v \geq \cos(\theta)} f_r(u) du \right) f(r) dr \\ &\geq cs_v \mathbb{P}\{\|X\| \geq r_n\} \geq cs_v K \frac{\log n}{n} \geq 4 \frac{\log n}{n}. \end{aligned}$$

It follows that

$$\mathbb{P}(A_{n-1}(v)^c) \leq \exp(-(n-1)p_n) \leq \exp\left(-4 \frac{n-1}{n} \log n\right).$$

Since  $(n-1)/n \geq 1/2$  for large  $n$ ,

$$\mathbb{P}(A_{n-1}(v)^c) \leq n^{-2}.$$

Thus,

$$\sum_n \mathbb{P}(A_{n-1}(v)^c) < \infty,$$

and the first Borel–Cantelli lemma implies that, almost surely, there exists a finite (random) index  $N_v$  such that  $A_{n-1}(v)$  holds for all  $n \geq N_v$ .

Since  $V$  is finite, we may intersect over all  $v \in V$ . Hence, with probability one, there exists

$$N := \max_{v \in V} N_v < \infty$$

such that for all  $n \geq N$  and all  $v \in V$ , there exists a  $v$ -good observation among  $\{X_1, \dots, X_{n-1}\}$ .

### F.1.3 Uniform control of the support function

We now combine directional coverage with the magnitude bound to obtain a uniform estimate over all directions  $u \in U$ . We work on the intersection of the almost-sure events established above.

Fix  $u \in U$ . Choose  $v \in V$  such that  $u \cdot v \geq \cos(\theta)$ . On the almost-sure event above, there exists a  $v$ -good observation  $X_i$  with  $i \leq n-1$ . Since both  $u$  and  $X_i/|X_i|$

lie within angle  $\theta$  of  $v$ , the angle between  $u$  and  $X_i/|X_i|$  is at most  $2\theta$ , and therefore

$$u \cdot \frac{X_i}{|X_i|} \geq \cos(2\theta).$$

Since  $|X_i| \geq r_n$  for a  $v$ -good observation, it follows that

$$u \cdot X_i \geq r_n \cos(2\theta). \tag{7}$$

Now fix  $\varepsilon > 0$ . Choose  $\delta > 0$  and  $\theta \in (0, \pi/4)$  sufficiently small so that

$$\frac{1 + \delta}{\cos(2\theta)} < 1 + \varepsilon. \tag{8}$$

Then almost surely for all sufficiently large  $n$ ,

$$\begin{aligned} u \cdot X_n &\leq |X_n| && \text{by Cauchy–Schwarz} \\ &\leq (1 + \delta)r_n && \text{by Lemma 3} \\ &\leq \frac{1 + \delta}{\cos(2\theta)} u \cdot X_i && \text{by (7)} \\ &\leq (1 + \varepsilon) u \cdot X_i && \text{by (8)} \\ &\leq (1 + \varepsilon) \max_{j \leq n-1} u \cdot X_j. \end{aligned}$$

Since this bound holds uniformly for all  $u \in U$ , we conclude that

$$X_n \in (1 + \varepsilon)K_{n-1},$$

as desired.

## F.2 Proof of Lemma 2

By assumption (1), since  $Z_i > 0$  a.s., we have a.s.

$$\limsup_n \frac{Z_n}{\max_{i \leq n-1} Z_i} \leq 1 + \varepsilon$$

We will show that in this event,

$$\limsup_n \left( \max_{j \in \{1, \dots, n\}} \frac{\max_{i \in \{1, \dots, n\}} Z_i}{\max_{i \in \{1, \dots, n\} \setminus j} Z_i} \right) \leq 1 + \varepsilon$$

Suppose otherwise, so we can find an infinite sequence  $k$  and  $j_k \leq k$  such that

$$\begin{aligned} \max_{i \in \{1, \dots, k\}} Z_i &> (1 + \varepsilon) \max_{i \in \{1, \dots, k\} \setminus j_k} Z_i \\ Z_{j_k} &> (1 + \varepsilon) \max_{i \in \{1, \dots, k\} \setminus j_k} Z_i \geq (1 + \varepsilon) \max_{i \leq j_k - 1} Z_i \end{aligned}$$

This implies that  $j_k \leq N$ . But this implies for all  $k > N$ ,

$$\max_{i \leq N} Z_i \geq Z_{j_k} > (1 + \varepsilon) \max_{i \in \{N, \dots, k\}} Z_i$$

But by assumption (2),  $\max_{i \in \{N, \dots, k\}} Z_i \rightarrow \infty$  as  $k \rightarrow \infty$ , yielding a contradiction.

## G Proofs for Section 5

### G.1 Proof of Proposition 4

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

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

By monotonicity of the generator  $g$ ,

$$g(C \setminus A) \subseteq g(D \setminus A) \quad \text{for all } A \in \mathcal{A}.$$

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  $p_g^A(C) \subseteq p_g^A(D)$ .

*Part (2).* Let  $P := p_g^A(C)$ . By definition,  $P \subseteq g(P)$ . 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.

## G.2 Proof of Proposition 5

Since  $\mathcal{A}$  is richer than  $\mathcal{B}$ , let  $A_B \in \mathcal{A}$  denote the set in  $\mathcal{A}$  such that  $A_B \supseteq B$  for every  $B \in \mathcal{B}$ . By the monotonicity of  $g$ , we have

$$g(C \setminus A_B) \subseteq g(C \setminus B)$$

Thus,

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

as desired.

## References

- ANDERSEN, S. (2022): “The Alt-Right Manipulated My Comic. Then A.I. Claimed It.” *The New York Times*.
- BOURTOULE, L., V. CHANDRASEKARAN, C. A. CHOQUETTE-CHOO, H. JIA, A. TRAVERS, B. ZHANG, D. LIE, AND N. PAPERNOT (2019): “Machine Un-learning,” *CoRR*, abs/1912.03817.
- BRODSKY, S. (2024): “The AI-Generated Books Trend is Getting Worse,” Accessed: 2026-01-29.
- CARLINI, N., D. IPPOLITO, M. JAGIELSKI, K. LEE, F. TRAMER, AND C. ZHANG (2023): “Quantifying Memorization Across Neural Language Models,” in *International Conference on Learning Representations*.
- CARLINI, N., F. TRAMER, E. WALLACE, M. JAGIELSKI, A. HERBERT-VOSS, K. LEE, A. ROBERTS, T. BROWN, D. SONG, U. ERLINGSSON, A. OPREA, AND C. RAFFEL (2021): “Extracting Training Data from Large Language Models,” in *USENIX Security Symposium*.
- 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.
- 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.
- GANS, J. S. (2024): “Copyright Policy Options for Generative Artificial Intelligence,” *NBER Working Paper*.
- 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.
- 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.
- MANTEGNA, M. (2024): "ARTificial: Why Copyright Is Not the Right Policy Tool for Generative AI," *Yale Law Journal Forum*.
- NASR, M., S. ZHANG, M. JAGIELSKI, N. CARLINI, B. BALLE, D. IPPOLITO, AND F. TRAMER (2023): "Scalable Extraction of Training Data from (Production) Language Models," *arXiv preprint arXiv:2311.17035*.
- PASQUALE, F. AND H. SUN (2024): "Consent and Compensation: Resolving Generative AI's Copyright Crisis," *Virginia Law Review Online*.
- 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," .
- 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 Copyrighted 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.
- U.S. COPYRIGHT OFFICE (2025): “Identifying the Economic Implications of Artificial Intelligence for Copyright Policy,” Pre-publication report.
- 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.