

Generative AI and Copyright: A Dynamic Perspective

S. Alex Yang

London Business School, sayang@london.edu

Angela Huyue Zhang

The University of Hong Kong, angelaz@hku.hk

The rapid advancement of generative AI is poised to disrupt the creative industry. Amidst the immense excitement for this new technology, its future development and applications in the creative industry hinge crucially upon two copyright issues: 1) the compensation to creators whose content has been used to train generative AI models (the fair use standard); and 2) the eligibility of AI-generated content for copyright protection (AI-copyrightability). While both issues have ignited heated debates among academics and practitioners, most analysis has focused on their challenges posed to existing copyright doctrines. In this paper, we aim to better understand the economic implications of these two regulatory issues and their interactions. By constructing a dynamic model with endogenous content creation and AI model development, we unravel the impacts of the fair use standard and AI-copyrightability on AI development, AI company profit, creators income, and consumer welfare, and how these impacts are influenced by various economic and operational factors. For example, while generous fair use (use data for AI training without compensating the creator) benefits all parties when abundant training data exists, it can hurt creators and consumers when such data is scarce. Similarly, stronger AI-copyrightability (AI content enjoys more copyright protection) could hinder AI development and reduce social welfare. Our analysis also highlights the complex interplay between these two copyright issues. For instance, when existing training data is scarce, generous fair use may be preferred only when AI-copyrightability is weak. Our findings underscore the need for policymakers to embrace a dynamic, context-specific approach in making regulatory decisions and provide insights for business leaders navigating the complexities of the global regulatory environment.

Key words: Generative AI, copyright, fair use standard, intellectual property, creative industry, creator economy, authorship

History: First draft: 4 February 2024; this draft: 19 March 2024.

1. Introduction

The rapid progress of generative artificial intelligence (AI) is propelling the next wave of productivity. Goldman Sachs has estimated that generative AI could drive a 7% (or almost \$7 trillion) increase in global GDP and lift productivity growth by 1.5 percentage points over a 10-year period (Goldman Sachs 2023). One sector poised to be transformed by generative AI is the creative industry (De Cremer et al. 2023). However, alongside the burgeoning excitement about this technology, it

has also sparked significant legal and economic conflicts, particularly between AI firms and content creators.

The first major controversy is whether the training of AI models falls under the exception of fair use, which would exempt AI firms from compensating copyright owners for using their content in AI training (“fair use standard”). This issue lies at the heart of numerous intellectual property (IP) infringement litigations filed by publishers, artists and authors against leading AI firms. For example, Getty Images sued Stability.AI for using their images in training AI models without authorization (Brittain 2023). Another high-profile case involved the New York Times suing Open AI and Microsoft, wherein the newspaper is seeking billions of dollars in damages for alleged copyright infringements (Grynbaum and Mac 2023). While it may take years before the US courts resolve this issue, leading AI firms are already facing pressures to negotiate content licensing agreements with media and music companies. For example, OpenAI has signed deals with publishers including Associated Press and Axel Springer to license their archived articles and new stories (O’Brien 2023, Thomas and Murgia 2023). And the ChatGPT-maker is reportedly offering between \$1 million and \$5 million annually to license news articles for training its AI models (David 2023). Meanwhile, the European Union, which is on the verge of introducing a sweeping AI Act, seems poised to take a quicker and tougher stance on the fair use issue. The latest draft of the AI Act mandates developers to disclose the copyrighted materials used for model training (European Parliament 2023). This transparency requirement could leave AI firms vulnerable to litigation while simultaneously bolstering the bargaining power of content creators in licensing negotiation.

The other major legal controversy revolves around the issue of whether content generated by AI should enjoy copyright protection (“AI-copyrightability”). In the United States, the Copyright Office and district courts have consistently denied protection to AI-generated content (Geiger 2023). This stands in contrast with a landmark ruling from the Beijing Internet Court in late 2023, in which the court held that an image created by feeding text prompts into Stable Diffusion can be copyrighted (Wininger 2023).

Both the issues of fair use and AI-copyrightability have posed tremendous challenges to the existing copyright doctrines, igniting fierce debate among scholars (Lemley 2023, Geiger 2023). Some scholars argue that only human-generated content should be eligible for copyright protection, while others believe that AI-created content could be copyrighted if human input plays a central role in its creation (Grimmelmann 2015, Burk 2020, Abbott and Rothman 2022). However, the extent of human input required remains a highly contentious issue. Similarly, scholars note that the transformative use of copyrighted materials is often considered fair use under US law, but this argument becomes less tenable when AI-generated content is substantially similar to the original copyrighted materials (Henderson et al. 2023). For example, in the New York Times lawsuit,

the newspaper cited several examples where the AI-generated content provided users with near-verbatim excerpts from its articles.

Beyond these legal complexities, these two issues also carry important business and policy ramifications. A primary concern is that restricting the use of copyrighted materials for training LLMs will significantly increase data acquisition costs, potentially hindering the development of the AI industry. At the same time, media experts and politicians have warned that the failure to adequately compensate content creators could lead to the decline of the creative industry (Henshall 2023). Similar considerations exist regarding the issue of AI-copyrightability. During a recent interview, the presiding judge in the aforementioned Chinese case argued that granting copyright protection to AI-generated content could incentivize the use of AI, thereby driving investment and fostering the industry's development.¹ Motivated by the above debates, this paper aims to provide an analytical investigation into how fair use and AI-copyrightability, in the context of generative AI, in particular, how these two copyright issues affect the incentives and actions of AI companies and content creators, and what the implications are for their profitability and consumer welfare?

In addition to their individual impacts, we are also interested in the interaction between the fair use standard and AI-copyrightability, a salient feature in the context of generative AI. Intuitively, compared to traditional content creation, the AI company plays a key role in the AI content creation process. For its perspective, the fair use standard affects the cost of their input data, and AI-copyrightability influences the value of the output of their service. Thus, in this paper, we seek to better understand what the interactions between the fair use standard and AI-copyrightability and their economic impacts are.

To answer the above questions, we construct a dynamic model consisting of a regulator, a generative AI company, and a group of content creators. The regulator makes two copyright-related decisions: 1) whether the fair use standard is *generous* (the AI company does not need to compensate copyright owners for using their content in model training) or *strict* (copyright owners need to be compensated); 2) the level of AI-copyrightability (relative to human content, the fraction of copyright protection AI-generated content enjoys). The content creators, who are endowed with heterogeneous skill levels, decide whether to create content themselves (human content), use an existing generative AI tool (provided by the AI company) to create content (AI content), or not create any content at all. In the model, while using AI tools reduces the cost of content creation, the quality of the content created is constrained by that of the AI model.² In turn, the AI company, who generates revenue from creators using its service for content generation, determines

¹ <http://tinyurl.com/5n84vyx6> (in Chinese)

² It has been demonstrated that a low-quality model can have a negative impact on the quality of the generated content. For example, CNET was forced to pause publishing AI-generated content after discovering the articles were "plagued with errors" (Thorbecke 2023).

the amount of human content to acquire for model training.³ This choice directly determines the quality of the AI tool available to creators for content generation in the next period. When the AI company makes this model improvement decision, we consider two scenarios that mimic different competition environments without explicitly modeling multiple AI companies: 1) in the *monopoly* setting, the AI company chooses the model quality that maximizes its profit (revenue minus cost of acquiring training data); 2) in the *competitive* setting, the AI company increases its model quality as long as its profit is non-negative.⁴

Using this model, we reveal that the content production decision of the creators follows a double-threshold structure: highly skilled creators continue creating contents themselves, those with intermediate skill levels use AI tools to produce content, while the least skilled creators choose not to produce anything. This decision is affected by both the fair use standard and AI-copyrightability, as well as AI model quality and other modeling parameters. For example, as the AI model improves, the range of creators opting for AI-assisted content creation expands, while human- and no-creation regions shrink. This dynamic ultimately shapes the demand for the AI model, which, in turn, affects the AI company's incentive to further improve the AI model.

By characterizing the dynamic interaction between the AI company and content creators under different copyright regulations, we reveal that the impact of copyright regulation depends crucially on the amount of existing data that could be used for AI model training. When there exists abundant content for model training (the “data abundant regime”), the relationship between the AI company and content creators is relatively simple: as the supplier, the AI company provides a service to content creators who opt for AI creation. In this case, we find that generous fair use, compared with strict fair use, always increases AI model quality, resulting in a higher profit for the AI company, higher aggregated incomes for content creators, and greater consumer surplus. The impact of AI-copyrightability, however, is more nuanced. In the short term, stronger AI-copyrightability always increases the profit of the AI company and the aggregated income for content creators, but may hurt consumer surplus. In the long run, however, stronger AI-copyrightability may diminish the AI company's incentive to further develop its AI model when the market competition is weak. Regarding the interaction between these two regulatory decisions, we find that AI-copyrightability and generous fair use are partial substitutes in terms of AI company profit (in the monopoly setting) and AI model quality and company revenue (in the competitive setting). However, their joint effect on other performance measures is more complicated.

³ Due to the consideration of model collapse (Rao 2023, Lutkevich 2023), we assume that only human content is usable for improving model quality. Our main findings are qualitatively robust as long as human content is more valuable than AI content for improving AI model quality.

⁴ As shown in Corollary 1, this objective is also equivalent revenue maximization. This setting corresponds to the case where the AI company, as the incumbent, sets its model quality so high that no other AI company with similar technical capability finds it profitable to enter the market.

When the existing training content is scarce and AI model training relies heavily on newly created content (the “data scarce regime”), the relationship between the AI company and content creators is more complicated: the AI company, in addition to its role as the service provider to some content creators, also becomes the customer of some creators as it acquires training data from them. As such, the impact of these two regulatory decisions become more intricate as the content production policy of the creators and the subsequent AI development decision are more tightly coupled. First, we show that generous fair use could lower model quality and social welfare relative to strict fair use. This is because generous fair use diminishes the incentive for creators to generate human content, thereby limiting the availability of training data. Similarly, while strong AI-copyrightability directly increases the demand for AI tools, it simultaneously reduces the supply of human generated content. Consequently, this limits further development of AI models, reducing consumer surplus and social welfare. Finally, we find that these two regulatory decisions interact in a more sophisticated fashion, and this interaction could in fact become a first-order consideration. For example, relative to strict fair use, generous fair use encourages AI development and enhances social welfare when AI-copyrightability is weak, but becomes a hindrance to both under strong AI-copyrightability.

As the first analytical research that jointly examines the two important copyright regulation issues in the context of generative AI, our paper uncovers how different regulatory decisions affect the dynamic interactions between AI developers and content creators and the welfare for different stakeholders in the creative industry, yielding significant business and policy implications. For policymakers, we caution against a one-size-fits-all approach to regulate AI related copyright issues. Instead, such regulations should be based on specific economic and operational factors such as training data availability, competition environment, and the current status of AI development, as well as a country’s regulatory priorities (e.g., whether to enhance AI development or to safeguard the welfare of content creators and consumers). Our findings also underscores the need for regulators, especially those constrained by existing legal jurisprudence, to carefully consider the interplay between fair use and AI-copyrightability. Finally, our paper urges business leaders to prepare for a diverse regulatory environment and strategize to adeptly navigate its complexities.

2. Related Literature

This paper is related to several strands of literature. First, there is a recent literature that studies the strategic interaction between different parties in supply chains in the presence of a market of information or data. Bimpikis et al. (2019) highlights an information provider should adjust the quality of information based on the strategic interaction between its customers when selling information to them. Mehta et al. (2021) examines the data provider’s pricing decisions. Also

focusing on the setting where a data provider sells data to a buyer who makes operational decisions, Drakopoulos and Makhdoumi (2023) establish conditions when it is optimal to provide a free sample to the potential buyer. Gurkan and de Véricourt (2022) study how a company could better leverage the AI flywheel effect (the virtuous cycle where more usage of an AI model will generate more data that could be used to further improve the model) by delegating certain tasks to an agent subject to moral hazard. Bimpikis et al. (2023) highlight how data generated in one market could be used to affect production market competition in another market by enabling price discrimination. By also focusing on a setting that features data selling (from content creators to AI company) and the consumption of an AI model/data by strategic agents, our research complements the above papers in two dimensions. First, by focusing on generative AI, our model captures a new form of data supply chain, where content creators choose to become the data provider for the AI company or the user of the model/service provided by the AI company. Second, while also capturing the strategic interactions between different players in a data supply chain, our focus is not on identifying optimal decisions for individual companies, but rather on the impact of copyright-related regulations.

Second, our paper is related to the emerging research stream focuses on the impact of generative AI on business, such as how the adoption of generative AI affects user productivity (e.g., Peng et al. 2023, Noy and Zhang 2023, Dell’Acqua et al. 2023) and creativity (e.g., Zhou and Lee 2023). Some of these papers focus on online labor market (Liu et al. 2023a) and the content-creator and knowledge economy (Burtch et al. 2023, Huang et al. 2023, Doshi and Hauser 2023). Complementing this literature, our paper focuses on how various regulatory approaches related to copyright may shape the action and performance measures of different stakeholders within the creative industry. More broadly, our paper is related to the strand of research on how managers/agents make decisions in the presence of AI-assistance. Both empirical and theoretical works have examined this question in settings such as forecasting (Ibrahim et al. 2021), retail (Kesavan et al. 2010, Sun et al. 2022, Caro and de Tejada Cuenca 2023), sales (Luo et al. 2021), project management (Beer et al. 2022), procurement (Cui et al. 2022), adjudication (Cohen et al. 2023), car-sharing (Cui et al. 2024), customer service (Xu et al. 2024), and across different professions and occupations (e.g., Brynjolfsson et al. 2018, Liu et al. 2023b). Similar to the above research, our paper is also related to the implication of AI tools on business decisions (the data acquisition and model development decisions of AI companies, and the content production policy of creators). Differently, our paper focuses on the tradeoffs imposed by different regulatory decisions by capturing the interaction between the AI provider (and AI service) and users (creators and consumers). In addition to generate managerial insights, we also provide guidelines to policymakers to better understand the implications of different policy decision on various stakeholders and AI development.

Our paper is also related to the stream of literature that focuses on the impact of copyright protection on creative content. Several papers have empirically documented that copyright encourages creativity by increasing payment to authors (Li et al. 2018, Giorcelli and Moser 2020). However, empirical research also finds that copyright protection reduces consumers access to information goods and impedes the reuse of copyrighted materials in the production of new information (Biasi and Moser 2021, Nagaraj 2018). Our paper complements these literature by highlighting that the interaction between two copyright issues, copyrightability and the fair use standard, a salient feature in the context of generative AI.

With the advent of generative AI, a burgeoning body of literature, mostly in law, addresses the complex legal challenges posed by this technology, in particular on the application of existing copyright doctrines to generative AI (Samuelson 2023, Epstein et al. 2023, Sag 2023). In a departure from these studies, through an analytical investigation into copyright issues in the context of generative AI, we unravel the potential ramifications of regulatory decisions by examining their impacts on the incentives and actions of various stakeholders. Finally, the work most closely related to ours is Gans (2024), who focuses on the issue of the fair use standard. Through an economic model featuring model size and negotiation cost, Gans (2024) reveals that content creators should be compensated (what we refer to as “strict fair use”) when the model is small, while the issue is more nuanced when the model is large. Our paper differs in two aspects: first, in addition to fair use, we also analyze the issue of AI-copyrightability and reveal the intricate link between these two important copyright issues. Second, instead of model size, our study unravels three other pivotal factors that significantly influence policy outcomes: the availability of human generated content as training data, the model quality and the competitive structure of the market.

3. Model

We consider a dynamic model with three types of strategic players: a regulator, a (generative) AI company and a continuum of content creators. The model consists of two periods: Period-1 and Period-2. Period-1 further comprises two stages: the regulation-setting stage, and the (first) content production stage. During the regulation-setting stage, the regulator makes two decisions related to copyright that shape the regulatory environment of generative AI.

1. The fair use standard: whether the fair use standard is *generous* (the AI company does not need to compensate copyright owners for using their content in AI model training) or *strict* (copyright owners need to be compensated).
2. AI-copyrightability: the level of AI-copyrightability, $\phi \in (0, 1]$ captures to what extent AI-generated content enjoys copyright protection relative to human content. A larger ϕ indicate that AI content enjoys more copyright protection. At the extreme, $\phi = 1$ implies AI content enjoys the same level of copyright protection as human content.

With the above regulation in place, a continuum of content creators, whose mass M_1 is normalized to 1, determines their content production policy. Each creator has a capacity of producing one piece of content in this period. To maximize their payoffs (income minus cost), each creator chooses from three options: 1) do not create any content, which leads to a payoff of zero; 2) create content in the traditional fashion without generative AI tools (human generation); and 3) create content using the generative AI tool developed by the AI company. The latter two options differ in terms of both cost and income.

On the cost side, a content creator incurs a cost c_H under human generation. Under AI generation, the cost, c_A , consists of two components: the human cost of generation (c_A^O), and the fee paid to the AI company for using the AI tool (p_m). Correspondingly, the AI company generates revenue by selling this AI service to these creators who adopt AI generation. Specifically, assuming that the number of AI creators during this period is N_A , then the AI company's profit is $p_m N_A$. To capture the cost advantage of AI generation, we assume $c_A = c_A^O + p_m$ is smaller than c_H .

On the income side, the amount that a creator earns from their content depends on the quality of the content and the related copyright regulation. Specifically, the content a creator produces could be either of high quality (with consumption valuation v_h) or of low quality (with consumption valuation v_l). For ease of exposition, we normalize $v_h = 1$ and $v_l = 0$. Through copyright protection, the creator of a piece of human-generated content (human content), which enjoys full copyright protection, could monetize β of the consumption value, resulting an income of $\beta v_h = \beta$. To avoid uninteresting cases, we assume that $\beta > c_H$. This ensures that without AI, it is profitable for some creators to produce content. On the other hand, the income earned from AI content depends on the degree of AI-copyrightability (ϕ). That is, a piece of human content generates income β for the creator, while a piece of AI content generates $\phi\beta$.

Regarding content quality, the probability that a piece of content is of high quality depends on the creator's skill level and the mode of creation. Specifically, creators are endowed with heterogeneous skills in generating high-quality content. Let the skill level $x \in [0, 1]$, which follows the cumulative distribution $F(\cdot)$. That is, among all content creators, the fraction of $F(x)$ of them have a skill level lower than or equal to $F(x)$. For tractability, we assume the skill levels of content creators are uniformly distributed, that is, $F(x) = x$. Under human generation, the probability that this piece of content is of high quality x_H equals x , the creator's skill level.

Under AI creation, on the other hand, a creator utilizes the generative AI tool developed by the AI company to create content (AI content). In this case, the probability that a piece of content is of high quality depends not only on the creator's skill level (e.g., initial ideas, effective prompting),

but also the AI model's quality. Let the model quality in this period be $X_1 \in [0, 1]$. Specifically, a creator with skill level x generates a piece of high-quality content with probability

$$x_A = \lambda X_1 + (1 - \lambda)x, \quad (1)$$

where $\lambda \in [0, 1]$ captures the degree of automation in content generation. The larger λ is, the less human intervention is needed in generating content. Such a specification reflects the empirical findings that generative AI tools tend to improve the (relative) performance of low-skilled individuals more than the high-skilled ones (Dell'Acqua et al. 2023).

Next, we enter Period-2, which also consists of two stages: the *model improvement* stage and the (second) content production stage. During the model improvement stage, the AI company chooses the amount of data to acquire to further improve the model. As AI generated content has been demonstrated to lead to a decline in model quality when used as training data (Martínez et al. 2023, Shumailov et al. 2023), we assume that only high-quality human content could be used to improve the AI model. Specifically, to improve the model quality from X_1 to $X_2 \in [X_1, 1]$, the amount of data needed is:

$$Q = q(X_2; X_1), \quad (2)$$

where the function $q()$ satisfies the following assumption.

ASSUMPTION 1. $q(X_2; X_1)$ convexly increases in X_2 , and decreases in X_1 ; $q(X_2; X_1) = 0$ for $X_2 \leq X_1$, and

$$\lim_{X_2 \rightarrow 1} q(X_2; X_1) = \infty. \quad (3)$$

Intuitively, this assumption states that the amount of data needed to improve a model increases in the target model quality (X_2) and decreases in the current one (X_1). In addition, for any further model improvement, the amount of data required increases, and it becomes prohibitively high to obtain a “perfect” model ($X_2 = 1$).

In terms of data source, the AI company can acquire data from an existing pool of data (let the total quantity of existing data that has not been used for model training be Q_0) and content creators in order to further improve the model. Depending on the amount of the data existing before the current period (let it be Q_0), we consider two regimes:

1. The data *abundant* regime: this regime corresponds to the scenario where the existing content for model training is abundant, that is, $Q_0 \rightarrow +\infty$. Therefore, the AI company does not depend on any content newly created in this period.

2. The data *scarce* regime: this regime corresponds to the scenario that there is no existing content for model training ($Q_0 = 0$). Thus, the AI company relies solely on content newly created in this period.

The cost that the AI company incurs when acquiring content depends on the fair use standard. Under generous fair use (G), the AI model can use all content for free. Under strict fair use, the AI company pays a fee p_t to the creator for each unit of high-quality human content used in model training. This in turn forms the stream of training income for the creators, potentially affecting their preferred creation mode.⁵ When the AI company determines the amount of data to acquire (and thus the improved model quality X_2), to capture the competition environment that the AI company operates in without explicitly modeling competition between multiple companies, we consider the following two settings.

1. In the *monopoly* setting (M), the AI company that chooses its improved model quality to maximize its Period-2 profit (Period-2 revenue minus the cost of acquiring training data). This corresponds to the scenario where the AI company has strong market power in this market.
2. In the *competitive* setting (C), the AI company maximizes the model quality X_2 subject to the constraint that its Period-2 profit is non-negative. This setting approximates the scenario where the AI company, as the incumbent, faces the threat of new entries with similar technical capability. As a result, it sets the AI model quality sufficiently high so that other companies find it unprofitable to enter this market by building a competing AI model.⁶

Finally, in the (second) content production stage, the setting is the same as in the content creation stage in Period-1, except for two differences. First, to capture possible market expansion, we assume the mass of content creators becomes $M \geq 1$. Second, the quality of AI-model used by creators becomes X_2 , as detailed above.

Throughout the paper, we use double superscript to indicate the data availability regime (A for the data abundant regime, and S for the data scarce regime) and the fair use standard (G for generous fair use; under strict fair use, M for the monopoly setting and C for the competitive setting). Subscripts presents time period (1 or 2) and content type (H for human content, A for AI content, and N for no content).

⁵ In this model, to focus on the impact of regulations, p_t is assumed to be fixed. One interpretation is that p_t is the creators' perceived cost of sharing the data to the AI company (e.g., for privacy reasons). Thus, when the AI company does not have the power to commit to a certain data acquisition price ex ante, the firm will choose the price p_t under the strict fair use case, and 0 under the generous fair use case. Further, there is a practical concern among regulators and AI companies that some content creators may disguise their AI generated content as human content in order to earn this training income. In this paper, we assume there is a test with negligible cost that could distinguish AI content and human content with high accuracy. In this case, disguising AI contents as human ones becomes irrelevant.

⁶ As shown later in Corollary 1, this objective is equivalent to the AI company maximizing its revenue (market size).

4. Equilibrium Results under the Data Abundant Regime

This section focuses on the data abundant regime. In this case, the existing data stock Q_0 is sufficiently large, so the model training revenue for the new content created is negligible.

4.1. Period-2 Content Creation Policy

We solve the game in backward induction, starting from the content creators' production policy in the Period-2 content creation stage. Here, facing an AI model with quality X_2 , each creator faces three options: create human content, create AI content, or not produce any content. For the first option, the expected payoff for a creator with skill level x is:

$$u_H = \beta x - c_H. \quad (4)$$

If choosing to produce AI content, the creator's expected payoff is

$$u_A = \phi\beta[\lambda X_2 + (1 - \lambda)x] - c_A, \quad (5)$$

where X_2 is the current AI model quality, and ϕ represents the level of AI-copyrightability. Finally, if a creator who decides not to produce any content, their payoff is $u_N = 0$. By comparing their expected payoffs under these three options, the creator determines their content production policy.

In order to characterize the creators' Period-2 policy, we define the following quantities.

$$X_A = \frac{c_A - (1 - \lambda)\phi c_H}{\lambda\phi\beta}; \quad X_H = 1 + \frac{c_A + (1 - \phi)\beta - c_H}{\lambda\phi\beta}; \quad X_N = \frac{c_A}{\lambda\phi\beta}. \quad (6)$$

and

$$x_{HO} = \frac{c_H}{\beta}; \quad x_{HA} = \frac{\lambda\phi\beta X + (c_H - c_A)}{[1 - (1 - \lambda)\phi]\beta}; \quad x_{AO} = \frac{c_A - \lambda\phi\beta X}{(1 - \lambda)\phi\beta}. \quad (7)$$

We further define functions $\underline{x}(X)$ and $\bar{x}(X)$ in Table 1.

Table 1 **Thresholds for Second Period Production Policy**

Case I: $\phi \leq \frac{\beta - c_H}{\beta(1 - \lambda)}$				
X	$[0, X_A]$	$[X_A, X_N]$	$[X_N, X^{HN}]$	$[X^{HN}, 1]$
$(\underline{x}(X), \bar{x}(X))$	(x_{HO}, x_{HO})	(x_{AO}, x_{HA})	$(0, x_{HA})$	$(0, 1)$
Case II: $\phi > \frac{\beta - c_H}{\beta(1 - \lambda)}$				
X	$[0, X_A]$	$[X_A, X^{HN}]$	$[X^{HN}, X_N]$	$[X_N, 1]$
$(\underline{x}(X), \bar{x}(X))$	(x_{HO}, x_{HO})	(x_{AO}, x_{HA})	$(x_{AO}, 1)$	$(0, 1)$

PROPOSITION 1. *In Period-2, facing the AI tool with model quality X_2 , the creators' content production policy is:*

1. creators with skill level $x < \underline{x}(X_2)$ do not produce any content;
2. creators with $x \in [\underline{x}(X_2), \bar{x}(X_2))$ produce AI content; and
3. creators with $x \geq \bar{x}(X_2)$ produce human content.

The above result reveals that the creators follow an intuitive double threshold policy: those with the lowest skill level find it unprofitable to produce any content, and those with the highest skill level choose to produce content without the help of AI. The creators with medium skill levels will generate AI content. The size of these three ranges depends on the capability of the AI tool (X_2): as the AI tool becomes more capable, the range of AI creation expands, while the two ranges at the end shrink. That is, a stronger AI model incentivizes more AI creation by both replacing human creation with AI creation and by inducing new entries. AI-copyrightability (ϕ) has a similar effect. Further, the strength of AI-copyrightability determines the sequence that the human creation region and the no creation region disappears. With weak AI-copyrightability, as AI model quality improves, AI creation first squeezes out no creation, and after that it just “cannibalizes” human creation; with strong AI-copyrightability, as AI model quality increases, AI creation first squeezes out human creation, and further market expansion follows. Moreover, note that when $X \in [X_A, \min(X_N, X_H))$, $\underline{x} = x_{AO}$, $\bar{x} = x_{HA}$, a marginal increase in X leads a change in both $\underline{x}$ and $\bar{x}$. According to the partial derivatives,

$$\frac{\partial \underline{x}}{\partial X} = -\frac{\lambda}{1-\lambda}, \quad \frac{\partial \bar{x}}{\partial X} = \frac{\lambda\phi}{1-(1-\lambda)\phi}. \quad (8)$$

the impact of model quality on new AI entry is independent of copyrightability (ϕ). But the impact of model quality on replacing human with AI depends on copyrightability: the replacement effect is stronger as the degree of copyrightability increases (a greater ϕ). Finally, we note that when c_A is extremely small, $X_A < 0$. That is, regardless of the quality of the AI tool, its cost advantage alone could induce creators to enter the market.

COROLLARY 1. *In Period-2, the amounts of human content (Q_{2H}) and AI content (Q_{2A}) generated are:*

$$Q_{2H} = \frac{M(1-\bar{x}^2)}{2}, \quad Q_{2A} = M\lambda X_2(\bar{x} - \underline{x}) + \frac{M(1-\lambda)}{2}(\bar{x}^2 - \underline{x}^2). \quad (9)$$

The AI company's revenue (r_2), the content creators' aggregated income (u_2), and the consumer surplus (s_2) are:

$$r_2 = Mp_m(\bar{x} - \underline{x}); \quad u_2 = \beta(Q_{2H} + \phi Q_{2A}) - M[c_H(1 - \bar{x}) + c_A(\bar{x} - \underline{x})]; \quad (10)$$

$$s_2 = (1 - \beta)Q_{2H} + (1 - \phi\beta)Q_{2A}; \quad w_2 = Q_{2H} + Q_{2A} - M[c_H(1 - \bar{x}) + c_A^O(\bar{x} - \underline{x})]. \quad (11)$$

Further, Q_{2A} and u_2 increases in X_2 . r_2 increases concavely in X_2 when $X_2 \geq X_A$.

The above result reveals how the AI model quality affects the amount of content created, and the different parties' payoffs. As shown, when the AI model quality improves, more AI content is generated, and creators as a whole also benefit. Further, the improvement of the AI model also increases the revenue of the AI company by attracting more creators.

4.2. AI Model Improvement

Anticipating how the AI model quality affects its revenue in Period-2, the AI company chooses the amount of data used for model training, which in turn determines X_2 . Under generous fair use, it is clear that as X_2 always increases its revenue, without any cost or other constraints on acquiring data, the AI company will acquire as much data as possible to improve the AI model, leading to the following result.

PROPOSITION 2. *Under generous fair use standard, $X_2^{AG} = 1$.*

4.2.1. Model Improvement under strict fair use: the monopoly setting. Under strict fair use, by taking the data acquisition cost p_t into consideration, the AI company's Period-2 profit π_2 under improved model quality $X \geq X_1$ follows:

$$\pi_2(X) = Mp_m [\bar{x}(X) - \underline{x}(X)] - p_t q(X; X_1), \quad (12)$$

where the first term is the firm's revenue under model quality X , and the second term represents the cost incurred when acquiring training data based on the one-to-one correspondence between Q and X (Eq. 2). Analyzing the structure of the function leads to the following result.

LEMMA 1. 1. *For $X \in [X_1, X_A)$, $\pi_2(X)$ concavely decreases in X ;*
2. *For $X \geq \max(X_1, X_A)$, $\pi(X)$ is concave in X .*

Facing the above profit function, the AI company's model improvement decision depends on the competition environment. In the monopoly setting, the AI company chooses the improved model quality X_2 to maximize its profit (Eq 12), that is,

$$X_2^{AM} = \arg \max_{X \geq X_1} \pi_2(X), \quad (13)$$

Recall from Lemma 1 that for $X \geq \max(X_1, X_A)$, $\pi(X)$ is concave in X . Thus, it is easy to identify the single local maximizer X^I which satisfies:

$$X^I = \arg \max_{X_2 \geq \max(X_1, X_A)} \pi_2(X), \quad (14)$$

the AI company's data acquisition policy is characterized as follows.

PROPOSITION 3. *Under the data abundant regime,*

1. for $X_1 < X_A$, $X_2^{AM} = X_I$ if $\pi_2(X^I) \geq 0$ and $X_2^{AM} = X_1$ otherwise;
2. for $X_1 \geq X_A$, $X_2^{AM} = X_I$.

As shown, when the initial model quality (X_1) is so low such that it does not attract any AI creation ($X_1 \leq X_A$), improving the model to the optimal level that induces AI creation (X^I) becomes so costly that it may be more profitable for the AI company to not improve the model at all ($X_2^{AM} = X_1$).

4.2.2. Model improvement under strict fair use: the competitive setting. In the competitive setting, the firm chooses the improved model quality X_2 to maximize its Period-2 revenue (or equivalently, model quality) without losing money ($\pi_2 \geq 0$), that is.

$$X_2^{AC} = \max\{X : \pi_2(X) \geq 0\}. \quad (15)$$

The existence of X_2^{AC} is guaranteed as $\pi_2(X_1) \geq 0$ and $\pi_2(1) < 0$. The following proposition characterizes the monotonicity of X_2^{AC} with respect to model parameters.

PROPOSITION 4. X_2^{AC} increases in X_1 , ϕ , M and k , and decreases in p_t .

Intuitively, as a higher initial model quality lowers the cost of data acquisition, the AI company can afford to spend more on model training. Similarly, as stronger AI-copyrightability increases the firm's revenue potential (according to Corollary 1), it also increases its appetite for model improvement.

4.3. Period-1 Content Creation Policy

In the presence of sufficient training data, regardless of the fair use standard, creators' income from selling human content they generated in Period-1 is negligible. Thus, the creators' Period-1 content production policy is the same as that in Period-2 (Proposition 1), except that the thresholds $\bar{x}$ and $\underline{x}$ are based on the initial AI model quality is X_1 (instead of the improved one X_2). Put differently, creators choose their content production policy in a myopic fashion, without thinking about the fair use standard. Similarly, the different performance measures in Period-1 follow a similar form as in Corollary 1, except that the creator mass is 1 instead of M . For completeness, we summarize the policy and performance measures below.

COROLLARY 2. In Period-1, the amounts of human content (Q_{1H}) and AI content (Q_{1A}) generated are:

$$Q_{1H} = \frac{1 - \bar{x}(X_1)^2}{2}, \quad Q_{1A} = \lambda X_2 [\bar{x}(X_1) - \underline{x}(X_1)] + \frac{(1 - \lambda)}{2} [\bar{x}(X_1)^2 - \underline{x}(X_1)^2]. \quad (16)$$

The AI company's revenue (r_1), the content creators' aggregated income (u_1), the consumer surplus (s_1), and social welfare (w_1) are:

$$\begin{aligned} r_1 &= p_m [\bar{x}(X_1) - \underline{x}(X_1)]; & u_1 &= \beta(Q_{1H} + \phi Q_{1A}) - [c_H(1 - \bar{x}(X_1)) + c_A(\bar{x}(X_1) - \underline{x}(X_1))]; \\ s_1 &= (1 - \beta)Q_{1A} + (1 - \phi\beta)Q_{1H}; & w_1 &= Q_{1H} + Q_{1A} - [c_H(1 - \bar{x}(X_1)) + c_A^O(\bar{x}(X_1) - \underline{x}(X_1))]. \end{aligned} \quad (17)$$

$$(18)$$

Further, Q_{1A} and u_1 increases in X_1 . r_1 increases concavely in X_1 when $X_1 \geq X_A$.

4.4. Impact of Copyright Policies


In this section, we uncover the impact of copyright policies on AI development (as measured by improved model quality X_2) and performance measures for different stakeholders (e.g., AI company's revenue and profit, aggregated creator income, consumer surplus and social welfare). We first examine the fair use standard and AI-copyrightability independently, and then the interaction between the two. Note that the interaction between the two regulatory decisions reflects a salient feature in generative AI due to the key role the AI company plays in content creation: the fair use standard directly affects the cost of inputs to the AI model, while AI-copyrightability influences the value of the output of the AI model (jointly with AI creators).

4.4.1. Impact of the fair use standard. As discussed above, because under the data abundant regime, the fair use standard does not affect the creators' first-period content production decisions. Thus, the following result focuses on its impact on the second period decisions, including AI model quality and content production decisions.

PROPOSITION 5. *Under the same level of AI-copyrightability,*

1. $X_2^{AG} > X_2^{Aj}$ for $j = M, C$.
2. $Q_{2A}^{AG} \geq Q_{2A}^{Aj}$, for $j = M, C$.
3. $r_2^{AG} \geq r_2^{Aj}$ for $j = M, C$, and $\pi_2^{AG} \geq \pi_2^{AM}$.
4. $u_2^{AG} \geq u_2^{Aj}$ for $j = M, C$.

This result reveals that in the data abundant regime, compared to strict fair use, generous fair use leads to higher model quality, more AI content, higher AI company profit (in the monopoly setting), and higher aggregated incomes for creators in Period-2. Thus, generous fair use is a beneficial policy for AI companies and content creators. Further, as shown in the following numerical results (Figure 1, where the y-axis represents the relative quantity under strict fair use as a percentage of the same quantity under generous fair use), in a monopoly market, generous fair use is also preferred by consumers, and thus the society as a whole. This benefit, however, decreases in the initial AI model quality. Similar results hold in a competitive market.

Figure 1 The Impact of Fair Use Standard under Monopoly Market: Period-2 Payoff

Notes. **Parameters:** $c_H = 0.3$, $c_A^0 = 0.05$, $p_m = 0.1$, $p_t = 0.2$, $\lambda = 0.6$, $\beta = 0.6$, $\phi = 1$, $M = 2$, $k = 2$.

4.4.2. Impact of AI-copyrightability. Unlike the fair use standard, AI-copyrightability affects decisions in both periods.


PROPOSITION 6. *In Period-1, the level of AI-copyrightability increases the amount of AI content, AI company revenue, and aggregated creator income, but reduces the amount of human content.*

The above result suggests that stronger AI-copyrightability leads to more AI creation, and consequently more AI company revenue. It also increases the aggregated creator income by making the AI creation option more attractive. However, as shown in Figure 2, stronger AI-copyrightability is not always favorable for consumers. This is due to two competing forces: on the one hand, strong AI-copyrightability benefits consumers by attracting more creators to enter the market and encouraging existing AI creators to produce more high-quality AI content, thus benefiting consumers. On the other, strong AI-copyrightability allocates a larger portion of the consumption value of AI content to the creators, leaving less to consumers. For the same reason, we find that when it comes to social welfare (the combined payoffs for the AI companies, creators, and consumers), full AI-copyrightability is not necessarily optimal.

Finally, we note that the same results hold in Period-2 under generous fair use (regardless of market structure), where the improved model quality X_2 is independent of AI-copyrightability.

Regarding the impact of AI-copyrightability on model improvement under strict fair use, the results depend on the competitive landscape. Recall from Proposition 4 that X_2^{AC} increases in ϕ . Thus, in the competitive setting, stronger AI-copyrightability also incentivizes AI development.

In contrast, when the firm aims to maximize profit in the monopoly setting, the impact of AI-copyrightability on X_2 is more nuanced. Specifically, according to Proposition 1, the firm's Period-2 revenue (r_2) follows a piece-wise linear structure. That is, within each segment (e.g.,


Figure 2 Impact of AI-Copyrightability on Period-1 Performance

Notes. **Parameters:** $c_H = 0.3$, $c_a^0 = 0.05$, $p_m = 0.1$, $\lambda = 0.6$, $\beta = 0.6$, $X_1 = 0.5$. All quantities are normalized by their value at $\phi = 1$.


$X_2 \in [X_A, \min(X_H, X_N)]$, r_2 is a linear function on X_2 . Within the segment, AI-copyrightability increases the marginal return of higher model quality, that is, $\frac{\partial}{\partial \phi} \left(\frac{\partial r}{\partial X} \right) > 0$. In this case, stronger AI-copyrightability leads to a higher model quality. However, we note that with an increase in ϕ , the AI company's revenue function could shift to another linear segment, where the marginal return of model quality improvement ($\frac{\partial}{\partial X_2} r(X_2)$) is smaller. Consequently, stronger AI-copyrightability could be a disincentive for AI model development. Figure 3 numerically illustrate the impact of AI-copyrightability on AI model improvement. Here, we specify the function $q(X_2; X_1) = \frac{1}{k} \log \left(\frac{1-X_1}{1-X_2} \right)$. We could easily verify that $q(X_2; X_1)$ satisfies all conditions in Assumption 1. This functional specification is used in all subsequent numerical results. As shown, consistent with Proposition 6, when market competition is intense, stronger AI-copyrightability incentivizes a higher investment in model improvement. In contrast, under weak market competition, the impact of AI-copyrightability on AI model improvement is non-monotone, as elaborated above.

4.4.3. Interaction between the fair use standard and AI-copyrightability. As the fair use standard only affects parties' Period-2 policies and payoffs, the interaction between these two copyright policies is also reflected in Period-2. In the competitive setting, by Proposition 4, we have $\frac{\partial}{\partial \phi} X_2^{AC} > 0$. Further, as the model quality under generous fair use $X_2^{AG} = 1$, independent of ϕ , we have $\frac{\partial}{\partial \phi} (X_2^{AG} - X_2^{AC}) < 0$. That is, in the competitive setting generous fair use and AI-copyrightability act as partial substitutes in the development of the AI model.

Figure (4a) shows that the other performance measures, such as AI company revenue, creator incomes and consumer surplus, largely follow the same trend. In contrast, in the monopoly setting, as the improved model quality (X_2^{AM}) is not monotone under ϕ , the interactive impact of the

Figure 3 Impact of AI-Copyrightability on AI Model Improvement (X_2)

Notes. **Parameters:** $c_H = 0.3$, $c_a^0 = 0.05$, $p_m = 0.1$, $\lambda = 0.6$, $\beta = 0.6$, $M = 2$, $p_t = 0.2$, $X_1 = 0.5$. The amount of data required to improve the model quality from X_1 to X_2 is: $q(X_2; X_1) = \frac{1}{k} \log \left(\frac{1-X_1}{1-X_2} \right)$, where $k = 2$.

Figure 4 Interaction of fair use standard and AI-copyrightability

Notes. All quantities are normalized by their counterparts under generous fair use. **Parameters:** $c_H = 0.3$, $c_a^0 = 0.05$, $p_m = 0.1$, $\lambda = 0.6$, $\beta = 0.6$, $M = 2$, $p_t = 0.2$, $X_1 = 0.5$. $q(X_2; X_1) = \frac{1}{k} \log \left(\frac{1-X_1}{1-X_2} \right)$, where $k = 2$.

two regulatory policies on other quantities is more complicated. As shown in Figure 4b, the AI company's profit under strict fair use approaches that under generous fair use as the degree of AI-copyrightability increases. Intuitively, this is because the AI company could adjust the AI model quality according to AI-copyrightability. However, the benefit of generous fair use on creator income may be more pronounced when AI-copyrightability is weak. Finally, regarding consumer surplus and social welfare, we note that the relationship is also non-monotone, and largely tracks the shape of AI model quality. Collectively, these results suggest that the interaction between these two regulatory policies is dependent on market structures.

5. Equilibrium Results under the Data Scarce Regime

This section focuses on the equilibrium under the data scarce regime ($Q_0 = 0$). That is, the training data required for the AI model improvement in Period-2 comes solely from the human content produced in Period 1. As shown in this section, this feature not only creates inter-temporal interaction between the two periods, but also changes how the two policies interact with each other.

5.1. AI Model Improvement

Under the data scarce regime, the creators' Period-2 content production policy is the same as those under the data abundant regime (Proposition 1). Thus, we focus on the AI model improvement policy in this section, and the creators' Period-1 content production policy in the next.

PROPOSITION 7. *In the data scarce regime, under generous fair use, the AI company utilizes all human content generated in Period-1 (Q_{1H}^{SG}), and the following AI model quality is*

$$X_2^{SG} = q^{-1}(Q_{1H}^{SG}; X_1). \quad (19)$$

On the other hand, when the fair use standard is strict, due to the cost associated with data acquisition, the AI company may not utilize all human content generated in Period-1 (Q_{1H}^{Sj} , $j = M, C$). For example, recall that from the data abundant regime, the AI company chooses to only improve their model quality up to X_2^{Aj} ($j = M, C$). Let Q^{Aj} be the amount of data used to achieve this target, that is,

$$Q^{Aj} = q(X_2^{Aj}; X_1), \quad j = M, C. \quad (20)$$

Intuitively, if $Q^{Aj} < Q_{1H}^{Sj}$, the AI company would not want to acquire all human content produced in Period-2. In another scenario, recall from Lemma 1 that when $X_1 < X^A$, the AI company's profit decreases in $X \in [X_1, X^A]$. Thus, if the amount of available data (Q_{1H}^{Sj}) is too small to improve the model beyond the level of X^A , it may not be worth improving the model at all.

By considering the above scenarios, we characterize the AI company's model improvement policy in Proposition 8 with the following definition. Let

$$\bar{X}_2^{Sj} = \sqrt{1 - 2Q_{1H}^{Sj}}, \quad j = M, C. \quad (21)$$

PROPOSITION 8. *In the data scarce regime, under strict fair use standard, the improved model quality X_2^{Sj} ($j = M, C$) follows:*

1. for $Q^{Aj} \leq Q_{1H}^{Sj}$, $X_2^{Sj} = X_2^{Aj}$;
2. for $Q^{Aj} > Q_{1H}^{Sj}$, $X_2^{Sj} = \bar{X}_2^{Sj}$ if $\pi_2(\bar{X}_2^{Sj}) \geq 0$, and $X_2^{Sj} = X_1$ otherwise.

Note that the amounts of data generated in the first period under generous fair use (Q_{1H}^{SG}) and that under strict fair use (Q_{1H}^{Sj} , $j = M, C$) could be different. Intuitively, with strict fair use, creators have more incentive to generate human data, and thus the data availability increases, that is, $Q_{1H}^{Sj} \geq Q_{1H}^{SG}$. Therefore, even though the amount of data used for training is bounded by an extra threshold Q^{Aj} , strict fair use may still result in a higher model quality, a result we will later establish formally.

5.2. Period-1 Content Production Policy

Under generous fair use, as content creators only earn income from the content's consumption value (e.g., β for human content), their content production policy in Period-1 is the same as that under the data abundant regime.

On the other hand, when the fair use standard is strict, the creator's Period-1 policy is more involved: when producing human content, creators can earn income not only through the content's consumption value, but also (potentially) through its training value (p_t). For example, by considering Proposition 8 (Scenario 1), when the data demanded by the AI company under the strict fair use case is less than the total amount of human content produced ($Q^{Aj} < Q_{1H}^{Sj}$), the AI company would only acquire an amount of Q^{Aj} . Thus, each piece of human content will be chosen for training the new model with probability $\frac{Q^{Aj}}{Q_{1H}^{Sj}}$, and a creator's payoff of producing human content is:

$$u_H^j = \left[\beta + \left(\frac{Q^{Aj}}{Q_{1H}^{Sj}} \right) p_t \right] x - c_H, \quad j = M, C. \quad (22)$$

Compared to Eq. (4), the newly added term, $\left(\frac{Q^{Aj}}{Q_{1H}^{Sj}} \right) p_t x$, captures the expected income a human content creator receives by selling content to the AI company. As shown, the creator's payoff depends not only on Q^{Aj} , which reflects the AI company's unconstrained demand for data, but also on Q_{1H}^{Sj} , capturing the collective supply of human content from creators. This relationship helps moderating the supply of human content: as the AI company's demand of human content increases due to its revenue or profit projection in Period-2, the creators, anticipating a higher training income, move to create more human content. Similarly, as more creators move to create human content, they should expect that an increase in Q_{1H}^{Sj} lowers the training income they will receive.

On the other hand, when $Q_{1H}^{Sj} < Q^{Aj}$, that is, the amount of human content produced is less than what the AI company would want under the unconstrained case, by Proposition 8 (Scenario 1), the AI company either acquires all generated data at unit-price p_t , or does not acquire any data, leaving the payoff of the human creators as Eq. (4). By considering these scenarios jointly, the following proposition characterizes the content creators' Period-1 production policy, with thresholds $(\underline{x}^j, \bar{x}^j)$ for $j = M, C$ defined in the Appendix A for brevity of exposition.

PROPOSITION 9. *In Period-1, under the data scarce regime and strict fair use standard ($j = M, C$), the creators' content production policy is:*

1. *creators with skill level $x < \underline{x}^j$ do not produce any content;*
2. *creators with skill level $x \in (\underline{x}^j, \bar{x}^j]$ produce AI content; and*
3. *creators with $x > \bar{x}^j$ produce human content.*

Under this policy, the amount of human generated content is:

$$Q_{1H}^{Sj} = \frac{1 - (\bar{x}^j)^2}{2}. \quad (23)$$

As shown, despite the new dynamics affecting creators' payoff, their production policy remains structurally the same as in Proposition 1. That said, the thresholds are affected by additional factors, including the training data price p_t , as well as the Period-2 market size M and the AI company's training efficiency k . Together with Proposition 8, this proposition uncovers inter-temporal linkage that is not present in the data abundant regime. That is, in the data scarce regime, the creators' Period-1 content production policy and the AI company's model improvement policy are tightly coupled. As shown later, such dynamics introduce new channels through which copyright regulations will affect stakeholder actions and performance measures.

5.3. Impact of Copyright Policies


Similar to the data abundant regime, we first focus on the impact of the fair use standard and AI-copyrightability separately, and then their interactions.

5.3.1. Impact of the fair use standard. Different from the data abundant regime, where generous fair use always leads to a higher AI model quality, the following proposition reveals that this is not the case in the data scarce regime.

PROPOSITION 10. *When $X_1 > X_H$, the AI model quality in Period-2 under the strict fair use standard (X_2^{Sj} , for $j = M, C$) is (weakly) higher than that under generous fair use (X_2^{SG}).*

The above result states that in the extreme case where the initial model quality X_1 is so high such that creators do not produce human content under generous fair use, there is no data available for the AI company to make any improvement. In contrast, under strict fair use, the additional training income may induce some human creation, allowing more human content for model training and thus resulting in a higher AI model quality. While intuitive, this result highlights the drawback of generous fair use in the data scarce regime – limiting the supply of training data. In this aspect, it acts as a counterforce to its benefit of lowering data acquisition cost, which incentivizes AI development from the demand side.

A further implication of Proposition 10 is by encouraging AI development, strict fair use may also lead to higher social welfare in Period-2, as the training cost (p_t) is a transfer payment between the

Figure 5 Impact of the fair use standard under the data scarce regime

Notes. **Parameters:** $c_H = 0.3$, $c_A^0 = 0.05$, $p_m = 0.1$, $\lambda = 0.6$, $\beta = 0.6$, $M = 2$, $k = 5$, $p_t = 0.4$, $\phi = 1$.

AI company and creators. This is indeed what we observe in Figure 5b. That is, the Period-2 social welfare tracks the AI model development. Interestingly, as illustrated in Figure 5c, the fair use standard also affects Period-1 payoff, a phenomenon not observed under the data abundant regime. The reason is as follows. Relative to generous fair use, strict fair use (especially in the competitive setting) encourages more human content creation from the creators. This forward-looking behavior thus hurts short-term social welfare. However, by combining these two periods together, we observe that strict fair use (in the competitive setting) could improve long-term social welfare.


5.3.2. Impact of AI-copyrightability. We first focus on the scenario of generous fair use, and then turn to the strict fair use case.

PROPOSITION 11. *In the data scarce regime, under generous fair use, stronger AI-copyrightability increases the amount of AI content and AI company revenue in Period-1, but reduces the model quality in Period-2.*

Under generous fair use, the impact of AI-copyrightability on Period-1 content generation and profit is the same as those in the abundant data regime. However, strong AI-copyrightability lowers the amount of human content generated in Period-1, which is all used for AI model improvement in Period-2. Consequently, it results in a lower AI-model quality in Period-2.


In terms of Period-2 content production and performance measures, the impact of AI-copyrightability is more complicated. Directly, stronger AI-copyrightability increases creators' incentive to produce AI content; Indirectly, it lowers the quality of the AI model, and thus possibly making AI generation less attractive.

Figure 6 Impact of AI-copyrightability in the data scarce regime: Generous Fair use



Notes. All Period-2 performance measures are normalized by their counterparts at $\phi = 1$. **Parameters:** $c_H = 0.3$, $c_A^0 = 0.05$, $p_m = 0.1$, $\lambda = 0.6$, $\beta = 0.6$, $M = 2$, $k = 5$, $p_t = 0.2$, $X_1 = 0.5$.

When the fair use standard is strict, in addition to the aforementioned channel through which stronger AI-copyrightability lowers model quality by reducing training data availability, there exists one more economic force that links AI-copyrightability and AI development: as shown earlier, by encouraging AI content creation in Period-2, stronger AI-copyrightability affect the AI company's demand for training data (Q^{Aj}), with the directions depending on the competitive landscape and other factors. For example, according to Proposition 4, stronger AI-copyrightability leads to a higher demand for training data in the competitive setting. By Eq. (22), an increase in Q^{Aj} boosts the per-unit training income received by the creator, thus induces more human content in Period-1, lessening the data availability constraint. This indirect effect, which is only present under strict fair use standard, acts as a counterforce to the negative impact of AI-copyrightability on AI-development. Consequently, as shown in Figure (7a), the model quality under strict fair use (competitive) is less sensitive on the change on AI-copyrightability than that under generous fair use.

Figure 7 Interaction of fair use standard and AI-copyrightability under the data scarce regime

Notes. **Parameters:** $c_H = 0.3$, $c_A^0 = 0.05$, $p_m = 0.1$, $\lambda = 0.6$, $\beta = 0.6$, $M = 2$, $k = 5$, $p_t = 0.2$, $X_1 = 0.5$.

5.3.3. Interaction between the fair use standard and AI-copyrightability. Due to the mechanisms explained above, the interaction between these two copyright regulations becomes stronger. For example, recall from Eq. (6) that X_H decreases in AI-copyrightability ϕ . Thus, one implication of Proposition 10 is that even when generous fair use may enhance AI development under weak AI-copyrightability (high X_N), it could become an obstacle when AI-copyrightability is strong (low X_N). This is consistent with the numerical results presented in Figure (7a) (the competitive setting). We note that this is in sharp contrast with the result in the data abundant regime. There, even though stronger AI-copyrightability could act as a partial substitute to the generous fair use standard in enhancing AI development, it would never make strict fair use the preferred option (over generous fair use). Relatedly, Figures (7b) and (7d) reveal that similar trends are present for both Period-2 social welfare and the aggregated one.

Together, our analytical and numerical results have shown that under the data scarce regime, the impact of the fair use standard and AI-copyrightability, as well as their interactions, has become more complicated than that under the data abundant regime. This is because, through the two-way influence between Period-1 content creation and Period-2 model improvement, these two regulation issues affect decisions made in both periods. For instance, while the fair use standard has no impact on Period-1 content creation in the data abundant regime, it affects the same decisions when training data is scarce. As a result, although generous fair use is in general beneficial in the data abundant regime, it could backfire when AI-copyrightability is stronger.

6. Policy and Business Implications

By analyzing how different regulatory regimes affect the dynamic interactions between AI companies and content creators, as well as various related performance measures, our analysis yields findings with significant policy and managerial implications.

First, our findings underscore that there is no one-size-fits-all doctrine for addressing these AI-related copyright issues. Instead, a flexible and agile governance approach, attuned to the unique institutional contexts of various industry sectors and regions, is preferred. More specifically, our results provide guiding principles on how such regulations should be specific to various economic and operational factors, such as training data availability, competition environment, and AI model quality. For example, we show that while generous fair use is beneficial to all parties when there exists abundant training data, it could backfire when training data is scarce. Similarly, although strong AI-copyrightability stimulates AI development in highly competitive markets, it could become a hindrance for less competitive markets.

Our paper also offers insights for policymakers on how to tailor optimal policy decisions to align with their specific policy goals. Different countries, with their unique institutional environments, will prioritize different objectives when making regulatory decisions. For example, China and the United States, the two AI superpowers, may lean towards prioritizing AI growth, whereas the European Union is more inclined to focus on the welfare of their consumers and content creators. Based on our results (e.g., strong AI-copyrightability benefits AI companies at the cost of consumers. See Figure 2 and the related discussion), these different priorities are likely to shape regulatory decisions in these major jurisdictions.

Further, by uncovering the interplay between the fair use standard and AI-copyrightability, this paper highlights that regulators, especially those constrained by existing legal jurisprudence, should consider these two decisions jointly. For instance, leading US scholars have argued that AI model training likely constitutes fair use under existing US copyright laws in most circumstances (Henderson et al. 2023). In that case, it might not make sense to grant AI-generated content copyright protection in the United States, particularly in contexts where training data is scarce.

For business leaders, our findings underscore the importance of grasping the intricate impact of AI-related regulatory decisions on their operational landscape. For instance, policies that grant extensive copyright protection to AI-generated content may initially bolster the profits of AI enterprises but could also risk diminishing the reservoir of human-generated content. This suggests that a policy seemingly beneficial to AI entities might, in fact, precipitate adverse long-term impacts. Moreover, our analysis predicts that nations will likely tailor their regulatory approaches to align with their unique policy objectives. Therefore, business leaders must prepare for a diverse regulatory environment and strategize to adeptly navigate its complexities. Finally, as the major technology players could shape the policy debate regarding AI regulation through various lobbying efforts (Mackrael 2023, Tracy and Seetharaman 2023, Zhang 2024), our findings recommend AI companies to advocate for a regulatory framework that supports agile and responsive governance of generative AI technologies.

7. Concluding Remarks

In this paper, we provide an equilibrium analysis of how generative AI-related copyright policies, namely the fair use standard and AI-copyrightability, affect the development of AI and various stakeholders in the creative industry. By endogenizing the content production policy of creators and the model development decisions of AI companies, our findings underscore the short-term and long-term impact of these copyright policies, as well as the intricate interplay between them. As such, our results yield important policy and business implications.

As one of the first analytical investigations of copyright regulation in the setting of generative AI, this paper could be extended along multiple directions. For example, our study has made several deliberate modeling choices, such as endogenizing key high-level decisions like choice of creation model and AI model improvement, while treating certain variables—namely, the pricing of AI services (p_m) and the costs associated with data acquisition (p_t)—as exogenous. Future research could explore the optimization of these decisions from the perspectives of individual firms or entire supply chains under a given regulatory regime. Additionally, while our analysis concentrates on the regulatory actions of a single agency, expanding the scope to include multiple regulatory bodies could illuminate complex dynamics such as innovation arbitrage (Samuelson 2023) and international collaboration.

References

- Abbott R, Rothman E (2022) Disrupting creativity: Copyright law in the age of generative artificial intelligence. *Florida Law Review*.–August .
- Beer R, Qi A, Ríos I (2022) Behavioral externalities of process automation. *Available at SSRN 4295527* .

- Biasi B, Moser P (2021) Effects of copyrights on science: Evidence from the wwii book republication program. *American Economic Journal: Microeconomics* 13(4):218–260.
- Bimpikis K, Crapis D, Tahbaz-Salehi A (2019) Information sale and competition. *Management Science* 65(6):2646–2664.
- Bimpikis K, Morgenstern I, Saban D (2023) Data tracking under competition. *Operations Research* .
- Brittain B (2023) Getty Images lawsuit says Stability AI misused photos to train AI. *Reuters.com* .
- Brynjolfsson E, Mitchell T, Rock D (2018) What can machines learn and what does it mean for occupations and the economy? *AEA papers and proceedings*, volume 108, 43–47 (American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203).
- Burk DL (2020) Thirty-six views of copyright authorship, by Jackson Pollock. *Hous. L. Rev.* 58:263.
- Burch G, Lee D, Chen Z (2023) The consequences of Generative AI for UGC and online community engagement. *Available at SSRN 4521754* .
- Caro F, de Tejada Cuenca AS (2023) Believing in analytics: Managers' adherence to price recommendations from a DSS. *Manufacturing & Service Operations Management* 25(2):524–542.
- Cohen MC, Dahan S, Khern-Am-Nuai W, Shima H, Touboul J (2023) The use of AI in legal systems: determining independent contractor vs. employee status. *Artificial intelligence and law* 1–30.
- Cui R, Li M, Zhang S (2022) AI and Procurement. *Manufacturing & Service Operations Management* 24(2):691–706.
- Cui Y, Jiang ZZ, Li Q (2024) Unlocking the value of real-time ai assistance: Who benefits, and why? *SSRN (4497014)* .
- David E (2023) OpenAI's news publisher deals reportedly top out at \$5 million a year. *Verge.com* URL <http://tinyurl.com/rdbmrkf3>.
- De Cremer D, Bianzino NM, Falk B (2023) How generative AI could disrupt creative work. *Harvard Business Review* 13.
- Dell'Acqua F, McFowland E, Mollick ER, Lifshitz-Assaf H, Kellogg K, Rajendran S, Kraymer L, Candelon F, Lakhani KR (2023) Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (24-013).
- Doshi AR, Hauser O (2023) Generative artificial intelligence enhances creativity. *Available at SSRN* .
- Drakopoulos K, Makhdoumi A (2023) Providing data samples for free. *Management Science* 69(6):3536–3560.
- Epstein Z, Hertzmann A, Herman L, Mahari R, Frank MR, Groh M, Schroeder H, Smith A, Akten M, Fjeld J, et al. (2023) Art and the science of generative ai: A deeper dive. *arXiv preprint arXiv:2306.04141* .
- European Parliament (2023) Artificial Intelligence Act: deal on comprehensive rules for trustworthy AI. *European Parliament Press Release* URL <http://tinyurl.com/2du7j4sa>.

- Gans J (2024) Copyright Policy Options for Generative Artificial Intelligence. *Available at SSRN* URL <https://dx.doi.org/10.2139/ssrn.4707911>.
- Geiger C (2023) Elaborating a human rights friendly copyright framework for generative AI. *Available at SSRN* .
- Giorelli M, Moser P (2020) Copyrights and creativity: Evidence from italian opera in the napoleonic age. *Journal of Political Economy* 128(11):4163–4210.
- Goldman Sachs (2023) Generative AI could raise global GDP by 7%. URL <http://tinyurl.com/24mw97hk>.
- Grimmelmann J (2015) There’s no such thing as a computer-authored work-and it’s a good thing, too. *Colum. JL & Arts* 39:403.
- Grynbaum M, Mac R (2023) The times sues openai and microsoft over a.i. use of copyrighted work. *New York Times* .
- Gurkan H, de Véricourt F (2022) Contracting, pricing, and data collection under the ai flywheel effect. *Management Science* 68(12):8791–8808.
- Henderson P, Li X, Jurafsky D, Hashimoto T, Lemley MA, Liang P (2023) Foundation models and fair use. *arXiv preprint arXiv:2303.15715* .
- Henshall W (2023) Experts Warn Congress of Dangers AI Poses to Journalism. *Time.com* URL <http://tinyurl.com/3uvh982m>.
- Huang H, Fu R, Ghose A (2023) Generative AI and content-creator economy: Evidence from online content creation platforms. *Available at SSRN 4670714* .
- Ibrahim R, Kim SH, Tong J (2021) Eliciting human judgment for prediction algorithms. *Management Science* 67(4):2314–2325.
- Kesavan S, Gaur V, Raman A (2010) Do inventory and gross margin data improve sales forecasts for us public retailers? *Management Science* 56(9):1519–1533.
- Lemley MA (2023) How Generative AI turns copyright law on its head. *Available at SSRN 4517702* .
- Li X, MacGarvie M, Moser P (2018) Dead poets’ property – how does copyright influence price? *The RAND Journal of Economics* 49(1):181–205.
- Liu J, Xu X, Li Y, Tan Y (2023a) “Generate” the future of work through AI: Empirical evidence from online labor markets. *arXiv preprint arXiv:2308.05201* .
- Liu Y, Liu C, Huang T (2023b) Does AI Reduce Inequality? A Study with a New Occupational AI Exposure Measure. *available at SSRN: 4515629* .
- Luo X, Qin MS, Fang Z, Qu Z (2021) Artificial Intelligence Coaches for Sales Agents: Caveats and Solutions. *Journal of Marketing* 85(2):14–32.
- Lutkevich B (2023) Model collapse explained: How synthetic training data breaks AI. *TechTarget* URL <http://tinyurl.com/mu58f8ke>.

- Mackrael K (2023) AI Law Draws Pushback From Big Brands in Europe. *Wall Street Journal* URL <http://tinyurl.com/3bw45y4b>.
- Martínez G, Watson L, Reviriego P, Hernández JA, Juárez M, Sarkar R (2023) Towards Understanding the Interplay of Generative Artificial Intelligence and the Internet. *arXiv preprint arXiv:2306.06130* .
- Mehta S, Dawande M, Janakiraman G, Mookerjee V (2021) How to sell a data set? Pricing policies for data monetization. *Information Systems Research* 32(4):1281–1297.
- Nagaraj A (2018) Does copyright affect reuse? evidence from google books and wikipedia. *Management Science* 64(7):3091–3107.
- Noy S, Zhang W (2023) Experimental evidence on the productivity effects of generative artificial intelligence. Available at SSRN 4375283 .
- O'Brien M (2023) Chatgpt-maker openai signs deal with ap to license news stories. *APNews.com* URL <http://tinyurl.com/yc52axuz>.
- Peng S, Kalliamvakou E, Cihon P, Demirer M (2023) The impact of AI on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590* .
- Rao R (2023) AI-Generated Data Can Poison Future AI Models. *Scientific American* URL <http://tinyurl.com/mwzjanbu>.
- Sag M (2023) Fairness and fair use in generative ai. *Fordham Law Review, Forthcoming* .
- Samuelson P (2023) Generative ai meets copyright. *Science* 381(6654):158–161.
- Shumailov I, Shumaylov Z, Zhao Y, Gal Y, Papernot N, Anderson R (2023) The curse of recursion: Training on generated data makes models forget. *arXiv preprint arXiv:2305.17493* .
- Sun J, Zhang DJ, Hu H, Van Mieghem JA (2022) Predicting human discretion to adjust algorithmic prescription: A large-scale field experiment in warehouse operations. *Management Science* 68(2):846–865.
- Thomas D, Murgia M (2023) Axel springer strikes landmark deal with openai over access to news titles. *Financial Times* URL <http://tinyurl.com/yc7c7ku4>.
- Thorbecke C (2023) Plagued with errors: A news outlet's decision to write stories with AI backfires. *CNN* URL <http://tinyurl.com/28yu9wve>.
- Tracy R, Seetharaman D (2023) Tech Titans Zuckerberg, Musk, Altman Bring AI Debate to Washington. *Wall Street Journal* URL <http://tinyurl.com/2p7vc62z>.
- Wininger A (2023) Beijing Internet Court Releases Translation of Li vs. Liu Recognizing Copyright in Generative AI. *China IP Law Update* URL <http://tinyurl.com/533dr947>.
- Xu Y, Dai H, Yan W (2024) Identity disclosure and anthropomorphism in voice chatbot design: A field experiment. *Management Science* .
- Zhang AH (2024) *High Wire: How China Regulates Big Tech and Governs Its Economy* (London: Oxford University Press).
- Zhou E, Lee D (2023) Generative AI, human creativity, and art. Available at SSRN 4594824 .

Appendix A: Definition of $(\underline{x}^j, \bar{x}^j)$ for $j = M, C$ in Proposition 9.

Define x_Q^j ($j = M, C$) be

$$x_Q^j = \sqrt{1 - \min(1, 2Q^{Aj})}, \quad (24)$$

where Q^{Aj} is the amount of human content wanted by the AI company, as characterized in Proposition 8. Next, define

$$x_{HO}^F := \frac{c_H}{\beta + p_t}; \quad x_{HA}^F := \frac{c_H - c_A + \lambda\phi\beta X_1}{[1 - (1 - \lambda)\phi]\beta + p_t}, \quad (25)$$

and x_{HO}^{Pj} and x_{HA}^{Pj} through the following implicit functions:

$$\left[\beta + \frac{1 - (x_Q^j)^2}{1 - (x_{HO}^{Pj})^2} p_t \right] x_{HO}^{Pj} - c_H = 0, \quad (26)$$

and

$$\left[\beta - (1 - \lambda)\phi\beta + \frac{1 - (x_Q^j)^2}{1 - (x_{HA}^{Pj})^2} p_t \right] x_{HA}^{Pj} - \lambda\phi\beta X_1 - (c_H - c_A) = 0. \quad (27)$$

Further, define

$$X_{HO}^F := 1 - (1 - X_1)e^{-\frac{k}{2}[1 - (x_{HO}^F)^2]}; \quad X_A^F := \frac{c_A - (1 - \lambda)\phi\frac{\beta c_H}{\beta + p_t}}{\lambda\phi\beta}, \quad (28)$$

and

$$X_A^{Pj} := \frac{\left[\beta - (1 - \lambda)\phi\beta + \frac{1 - (x_Q^j)^2}{1 - (x_{HO}^{Pj})^2} p_t \right] x_{HO}^{Pj} - (c_H - c_A)}{\lambda\phi\beta}. \quad (29)$$

With the above quantities and x_{AO} defined in Eq. (7), the thresholds $\bar{x}^j$ and $\underline{x}^j$ are defined as follows:

1. When $x_Q^j \leq \max(x_{HO}^F, x_{HA}^F)$,
 - (a) for $X_1 \leq X_A^F$, if $\pi_2(X_{HO}^F) < 0$, $\underline{x}^j = \bar{x}^j = x_{HO}^F$; otherwise, $\underline{x}^j = \bar{x}^j = x_{HO}^F$.
 - (b) for $X_1 > X_A^F$, $\underline{x}^j = \max(0, x_{AO})$ and $\bar{x}^j = \min(1, x_{HA}^F)$.
2. When $x_Q^j > \max(x_{HO}^F, x_{HA}^F)$,
 - (a) for $X_1 \leq X_A^{Pj}$, $\underline{x}^j = \bar{x}^j = x_{HO}^{Pj}$;
 - (b) for $X_1 > X_A^{Pj}$, $\underline{x}^j = \max(0, x_{AO})$ and $\bar{x}^j = \min(1, x_{HA}^{Pj})$.

Appendix B: Proofs

Proof of Proposition 1. For a creator with skill level x , the (net) income of creating human content is

$$u_H = \beta x - c_H. \quad (30)$$

Next, facing AI model with quality X , the utility of creating AI content is:

$$u_A = \phi\beta[\lambda X + (1 - \lambda)x] - c_A. \quad (31)$$

Finally, $u_N = 0$ for creators not producing anything.

To prove the double-threshold policy, we want to show that for $x_1 < x_2$, if x_1 chooses human creation, then x_2 will also chooses human creation. To see that, note that x_1 chooses human creation if and only if

$$u_H(x_1) > \max(u_A(x_1), 0), \quad (32)$$

then we have: $u_H(x_2) > u_H(x_1) > 0$, and

$$u_H(x_2) \geq u_H(x_1) + \beta(x_2 - x_1) \geq u_A(x_1) + (1 - \lambda)\phi\beta(x_2 - x_1) = u_A(x_2) \quad (33)$$

Thus, we have $u_H(x_2) > \max(u_A(x_2), 0)$, that is, creator with x_2 also chooses human creation. Symmetrically, we could show that for $x_1 < x_2$, if creator with x_2 chooses to produce no content, then creator with x_1 also choose not to produce anything. Combining these two results, we obtain the double-threshold policy.

Next, we establish the structure of the thresholds $\underline{x}$ and $\bar{x}$ as characterized in Table 1. Here, as mentioned above, a creator chooses to produce human content if and only if $u_H \geq \max(u_A, u_N)$. That is,

$$\beta x - c_H \geq \max(\phi\beta[\lambda X + (1 - \lambda)x] - c_A, 0). \quad (34)$$

Or equivalently,

$$x \geq \max(x_{HO}, x_{HA}), \quad (35)$$

where

$$x_{HO} = \frac{c_H}{\beta}; \quad x_{HA} = \frac{\lambda\phi\beta X + (c_H - c_A)}{[1 - (1 - \lambda)\phi]\beta}. \quad (36)$$

Similarly, a creator chooses to generate AI-content if and only if

$$x \in (x_{AO}, x_{HA}). \quad (37)$$

where x_{HA} is defined above, and

$$x_{AO} = \frac{c_A - \lambda\phi\beta X}{(1 - \lambda)\phi\beta}. \quad (38)$$

Finally, a creator will choose not to enter the market if and only if

$$x \leq \min(x_{HO}, x_{AO}). \quad (39)$$

Let

$$X_A := \frac{c_a - (1 - \lambda)\phi c_h}{\lambda\phi\beta}, \quad (40)$$

we have:

1. when $X < X_A$, it could be verified that $x_{HA} \leq x_{HO} \leq x_{AO}$, and the corresponding policy is:

- (a) for $x \in [0, x_{HO})$, creators choose not to produce any content.
- (b) for $x \in (x_{HO}, 1]$, creators choose to produce human content.

Note that as $c_H \in (0, \beta)$, we have $x_{HO} \in (0, 1)$, so both regions are non-empty.

2. when $X \geq X_A$, it could be verified that $x_{AO} \leq x_{HO} \leq x_{HA}$, and the corresponding content production policy is

- (a) for $x \in [0, x_{AO})$, creators choose not to produce any content.
- (b) for $x \in [x_{AO}, x_{HA}]$, creators choose to produce AI content
- (c) for $x \in (x_{HA}, 1]$, creators choose to produce human content.

To further refine the three regions in the second scenario ($X \geq X_A$), we could show that when

$$X \geq X_H := 1 - \frac{(c_H - c_A) - (1 - \phi)\beta}{\lambda\phi\beta}. \quad (41)$$

we have $x_{HA} > 1$, thus there is no human creation. Similarly, when

$$X \geq X_N := \frac{c_A}{\lambda\phi\beta}. \quad (42)$$

we have $x_{AO} < 0$, thus every creator produces some content. Further, by comparing X_A , X_N , and X^{HN} , we have that $X_A < \min(X_N, X_H)$ and $X_H > X_N$ if and only if $c_H < [1 - (1 - \lambda)\phi]\beta$. Reorganizing the above scenarios lead to the thresholds and cases in Table 1. $\square$

Proof of Corollary 1. Under the above policy, the number of (high-quality) human content is

$$Q_{2H} = M \int_{\bar{x}}^1 x dF(x) = \frac{M(1 - \bar{x}^2)}{2} \quad (43)$$

The amount of AI content is:

$$Q_{2A} = M \int_{\underline{x}}^{\bar{x}} [\lambda X + (1 - \lambda)x] dF(x) = M \left[\lambda X(\bar{x} - \underline{x}) + \frac{1 - \lambda}{2} (\bar{x}^2 - \underline{x}^2) \right]. \quad (44)$$

The AI company's revenue is:

$$r_2 = M \int_{\underline{x}}^{\bar{x}} p_m dF(x) = M p_m (\bar{x} - \underline{x}). \quad (45)$$

The human content creators' income is:

$$u_{2H} = \beta Q_H - M c_H (1 - \bar{x}) = M \left[\frac{\beta}{2} (1 - \bar{x}^2) - c_H (1 - \bar{x}) \right], \quad (46)$$

The total AI-content creators' profit is:

$$u_{2A} = \phi\beta Q_A - M c_A (\bar{x} - \underline{x}) = M \left(\phi\beta \left[\lambda X(\bar{x} - \underline{x}) + \frac{1 - \lambda}{2} (\bar{x}^2 - \underline{x}^2) \right] - c_A (\bar{x} - \underline{x}) \right). \quad (47)$$

And the content creators' aggregated income total creators' profit is:

$$u_2 = u_{2H} + u_{2A} = \beta(Q_{2H} + \phi Q_{2A}) - M [c_H(1 - \bar{x}) + c_A(\bar{x} - \underline{x})]. \quad (48)$$

The total consumer surplus is:

$$s_2 = (1 - \beta)Q_{2H} + (1 - \phi\beta)Q_{2A}. \quad (49)$$

Finally, for social welfare:

$$w_2 = u_2 + s_2 + \pi_2 = Q_{2H} + Q_{2A} - M [c_H(1 - \bar{x}) + c_A^O(\bar{x} - \underline{x})] \quad (50)$$

Regarding monotonicity, it is clear that as $\bar{x}$ increases concavely in X_2 and $\underline{x}$ decreases convexly in X_2 , we have that Q_{2A} , which increases in $\bar{x}$ and decreases in $\underline{x}$, increases in X_2 . Similarly, r_2 , which increases linearly in $\bar{x} - \underline{x}$, increases concavely in X_2 .

For the monotonicity of u_2 , note that we could re-write a single creator's optimal content production policy as $\max(u_H, u_A, 0)$, which (weakly) increases in X_2 through u_A . By writing u_2 as

$$u_2 = M \int_0^1 \max(x - c_H, \lambda X_2 + (1 - \lambda)x - c_A, 0) dx, \quad (51)$$

we could show that u_2 also increases in X_2 . $\square$

Proof of Proposition 2. According to Corollary 1, the AI company's revenue r_2 increases in X_2 . On the other hand, the AI company could access abundant existing content with zero cost. Thus, whether the AI company is revenue-maximizing (competitive setting) or profit-maximizing (monopoly setting), they always prefer to improve the model quality to the highest possible level, that is, $X_2^{AG} = 1$. $\square$

Proof of Lemma 1. We note that by the definition of $\pi_2(X)$, we could examine its shape by two segments:

1. For $X \in [X_1, X_A)$, note that this interval is non-empty if and only if $X_1 < X_A$. By Proposition 1, we have that $r_2(X) = 0$ for $X < X_A$. Thus, $\pi_2(X) = -p_t q(X_2; X_1)$, which is concavely decreasing in X according to Assumption 1.
2. For $X \geq \max(X_1, X_A)$, by examining the AI company's profit, we observe that the revenue part is piece-wise linear and concave, and the cost part is convex, so the difference is concave. $\square$

Proof of Proposition 3. Consider two scenarios.

1. When $X_1 < X_A$, by Lemma 1, there are two possible candidates for X_2^{AM} : 1). X_1 , which is the optimal skill level for $X \in [X_1, X_A)$. In this case, $\pi_2(X_1) = 0$; 2) X^I , which is the optimal skill level for $X \geq \max(X_1, X_A)$. Thus, $X_2^{AM} = X^I$ if and only if $\pi_2(X^I) > 0$, and $X_2^{AM} = X_1$ otherwise.
2. When $X_1 \geq X_A$, the only possible candidate is X^I . $\square$

Proof of Proposition 4. By the structure of π_2 , we have the following structure for X_2^{AC} .

1. When $X_1 < X_A$ and $\pi_2(X^I) < 0$, $X_2^{AC} = X_1$, which increases in X_1 , and is independent of ϕ , M , k and p_t .
2. When $X_1 \geq X_A$, or $\pi_2(X^I) \geq 0$, then by the concavity of π_2 for $X \geq \max(X_1, X^I)$, there exists exact one skill level, let it be X^C , such that $\pi_2(X^C) = 0$. Therefore, $X_2^{AC} = X^C$. In this region, as $\pi_2()$ increases in X_1 , ϕ , M , and k , and decreases in p_t . Thus X_2^{AC} also increases in X_1 , ϕ , M , and k , and decreases in p_t .

Finally, we note that as X_1 increases, the choice of X_2^{AC} also could only shifts from the first candidate to the second, but not the other way around. Thus combining these two scenarios, we also have X_2^{AC} increases in X_1 . For the same argument, we could prove the monotonicity with respect to the other parameters. $\square$

Proof of Corollary 2. The proof is similar to that in Corollary 1 and the detail is omitted for expositional brevity. $\square$

Proof of Proposition 5. First, we have $X_2^{AG} = 1 > X_2^{Aj}$ for $j = M, C$. This proves the first statement. By Corollary 1, we have $Q_{2A}^{AG} \geq Q_{2A}^{Aj}$, $r_2^{AG} \geq r_2^{Aj}$ and $u_2^{AG} \geq u_2^{Aj}$. Finally, as $\pi_2^{AG} = r_2^{AG}$ and $\pi_2^{AM} \leq r_2^{AM}$, we have $\pi_2^{AG} \geq \pi_2^{AM}$. $\square$

Proof of Proposition 6. Note from the expressions in Proposition 1, we have that $\bar{x}$ increases in ϕ and $\underline{x}$ decreases in ϕ . Next, by the expressions in Corollary 1, Q_{2A} and r_2 increase in ϕ , and Q_{2H} decreases in ϕ . Regarding u_2 , by writing it as

$$u_2 = \int_0^1 \max(u_H, u_A, 0) dx \quad (52)$$

we observe that as u_A increases in ϕ and u_H is independent of ϕ , we have u_2 increases in ϕ . $\square$

Proof of Proposition 7. Under generous fair use, as the cost of data acquisition is zero, the AI company acquires as much data as available, which is Q_{1H}^{SH} in this case. The improved model quality X_2^{SG} follows directly from the definition. $\square$

Proof of Proposition 8. The proof for $j = M$ is similar to that for $j = C$. Thus, we focus on the case with $j = C$ and omit the $j = M$ case. Consider the two scenarios that correspond to the two cases in the proposition.

1. when $Q^{AC} \leq Q_{1H}^{SC}$, we have that for any $X_2 \geq X_2^{AC}$, $\pi_2(X_2) < 0$. Thus, the AI company only acquires the amount of Q^{AC} , resulting in improved model quality $X_2^{SC} = X_2^{AC}$.
2. When $Q^{AC} > Q_{1H}^{SC}$, the feasible region for X_2 is $[X_1, \bar{X}_2^{SC}]$. To identify X_2^{SC} , we consider the following scenarios.
 - (a) When $X_1 < X_A$, $\pi_2(X_1) = 0$. For $X \in [X_A, X_2^{AC})$, there is at most one X with $\pi_2(X) = 0$. if $\pi_2(\bar{X}_2^{SC}) < 0$, then for $X \in [X_1, \bar{X}_2^{SC}]$, we have $\pi_2(X) < 0$. Thus, $X_2^{SC} = X_1$.
 - (b) When $X_1 \geq X_A$, for $X \in [X_1, \bar{X}_2^{SC}]$, $\pi_2(X) \geq 0$, thus, X_2^{SC} should take the largest value, that is, $\bar{X}_2^{SC}$. $\square$

Proof of Proposition 9. The proof that the content production policy follows a double-threshold structure is similar to that in Proposition 1 and is omitted for brevity. Next, we focus on characterizing the thresholds $\bar{x}^S$ and $\underline{x}^S$. To do so, we first define the following auxiliary utility functions for human content creation. Let $u_H^t()$ be a creator's utility of producing human content and receives the full training price p_t , that is,

$$u_H^t(x) = (\beta + p_t)x - c_H; \quad (53)$$

Similarly, let $u_H^j()$ ($j = M, C$) be a creator's utility of producing human content and receives only partial training price (with total wanted amount Q^{Aj} and total produced human content Q_{1H}^{Sj}), that is,

$$u_H^j = \left(\beta + \frac{Q^{Aj}}{Q_{1H}^{Sj}} p_t \right) x - c_H; \quad (54)$$

Intuitively, u_H^t is the creators' relevant utility function when the total amount of content produced, Q_{1H}^{Sj} is no greater than Q^{Aj} , and u_H^j is the relevant one when $Q_{1H}^{Sj} > Q^{Aj}$. Based on the above definitions, we have the following auxiliary results.

Auxiliary Claim 1. When assuming the utility function of human creation is $u_H^t()$, by comparing that with u_A and u_O (note that the training income does not affect the utility function under AI-generation and no generation), the creators' content production policy is:

1. when $X_1 \leq X_A^F$, we have: $x_{HO}^F \geq x_{HA}^F$. Thus, a creator produces human content when $x \geq x_{HO}^F$ and does not create any content for $x < x_{HO}^F$.
2. when $X_1 > X_A^F$, we have $x_{HO}^F < x_{HA}^F$. Thus, a creator produces human content when $x \geq x_{HA}^F$, AI content when $x \in (x_{AO}, x_{HA}^F)$, and does not create any content for $x < x_{AO}$.

With this auxiliary claim, we characterize $\underline{x}^j$ and $\bar{x}^j$ by the following two scenarios.

Scenario I: $x_Q^j \leq \max(x_{HO}^F, x_{HA}^F)$. In this case, if the creators follow the production policy under Auxiliary Claim 1, the amount of human content created in Period-1, $\frac{1 - (\max(x_{HO}^F, x_{HA}^F))^2}{2}$, is smaller than the amount that the AI company wanted under the data abundant regime, Q^{Aj} . To obtain the creators' actual production policy, which forms an equilibrium, we consider the following two scenarios.

1. When $X_1 < X_A^F$, we have $x_{HO}^F > x_{HA}^F$, thus, the amount of human content created based on Auxiliary Claim 1 is X_{HO}^F . We then have two sub-cases.
 - (a) If $\pi_2(X_{HO}^F) < 0$, by the proof of Proposition 8, for all $X \in (X_1, X_{HO}^F)$, we have $\pi_2(X) < 0$. Thus, the AI company will not acquire any human content. Anticipating that, the creators' content production policy is the same as under Proposition 1. Further, note that $\pi_2(X_{HO}^F) < 0$ implies $X_1 < X_A$. Thus, we have $\underline{x}^j = \bar{x}^j = x_{HO}$.
 - (b) If $\pi_2(X_{HO}^F) \geq 0$. Thus, assuming the amount of human content is created, the AI company will acquire all human content to improve the model quality to X_{HO}^F . Anticipating that, the creators' content production policy is the same as in Auxiliary Claim 1, that is, $\underline{x}^j = \bar{x}^j = x_{HO}^F$.
2. When $X_1 \geq X_A^F$, we have $x_{HO}^F \leq x_{HA}^F$. Further, $X_1 \geq X_A^F$ implies that $\pi_2(X_1) > 0$. Thus, the first sub-case under the previous scenario (the AI company does not acquire any data) is irrelevant. Anticipating that the AI company will acquire all human content at training price p_t , the creators' content production policy is the same as in Auxiliary Claim 1, that is, $\underline{x}^j = \max(0, x_{AO})$, and $\bar{x}^j = \min(1, x_{HA}^F)$.

Combining the above scenarios, we have obtained the first set definition of $\underline{x}^j$ and $\bar{x}^j$ in Appendix A.

Scenario II: $x_Q^j > \max(x_{HO}^F, x_{HA}^F)$. Here, we could verify that the amount of data wanted by the AI company in Period-2 is less than the amount of human content crated in Period-1 assuming the creators follow the production policy in Auxiliary Claim 1. Anticipating that, creators should decide their production policy by assuming the utility function of creating human content follows $u_t^j()$ defined above. In particular, by the double-threshold policy, the total human content created $Q_{1H}^{Sj} = \frac{1 - (\bar{x}^j)^2}{2}$. In this case, we define thresholds x_{HO}^{Pj} and x_{HA}^{Pj} according to Equations. (26) and (27) in Appendix A. Intuitively, x_{HA}^{Pj} is the skill level of the

marginal creator who is indifferent between human and AI creation, and x_{HO}^{Pj} is that for the marginal creator who is indifferent between human creation and no creation.

Further, we note that a unique $x_{HA}^{Pj} \in [0, x_Q^j)$ is guaranteed because of the following three properties: 1) Equation (27) is less than 0 at $x_{HA}^{Pj} = 0$; 2) Equation (27) strictly increases in x_{HA}^{Pj} , and 3) by the condition that $(\beta + p_t)x_Q^j - c_H > \max(u_A(x_Q^j), 0)$, when $x_{HA}^P = x_Q^j$, Equation (27) is greater than 0. Similarly, a unique $x_{HO}^P \in [0, x_Q^j)$ is also guaranteed. Based on this definition, we have that a creator prefers human creation over AI creation if and only if $x > x_{HA}^{Pj}$, and prefers human creation over no creation if and only if $x > x_{HO}^{Pj}$. Finally, by the definition of X_A^{Pj} in Equation (29) in the Appendix, we have $x_{HO}^{Pj} \geq x_{HA}^{Pj}$ if and only if $X_1 \leq X_A^{Pj}$. Thus, we could characterize the creators' production policy as follows:

1. When $X < X_A^{Pj}$, we have $x_{HA}^{Pj} < x_{HO}^{Pj} < x_{AO}$. Thus,
 - (a) for $x \in [0, x_{HO}^{Pj})$, creators choose not to produce any content;
 - (b) for $x \in [x_{HO}^{Pj}, 1]$, creators choose to produce human content.
2. When $X \geq X_A^{Pj}$, we have $x_{AO} < x_{HO}^{Pj} < x_{HA}^{Pj}$, thus,
 - (a) for $x \in [0, x_{AO})$, creators choose not to produce any content;
 - (b) for $x \in [x_{AO}, x_{HA}^{Pj}]$, creators choose to produce AI content;
 - (c) for $x \in (x_{HA}^{Pj}, 1]$, creators choose to produce human content.

Reorganizing these terms leads to the second set condition in Appendix A ($x_Q^j > \max(x_{HO}^F, x_{HA}^F)$). □

Proof of Proposition 10. When $X_1 > X_H$, no human content is produced in Period-1 under generous fair use. Thus, $X_2^{SG} = X_1$. On the other hand, under strict fair use, $X_2^{Sj} \geq X_1$ for $j = M, C$. Thus, $X_2^{Sj} \geq X_2^{SG}$, as desired. □

Proof of Proposition 11. According to the proof of Corollary 1, Q_{1A}^{SG} and r_1^{SG} both increase in ϕ , and Q_{1H}^{SG} decreases in ϕ . Further, by Proposition 7, X_2^{SG} increases in Q_{1H}^{SG} , and thus decreases in ϕ . □