Copyright Policy Options for Generative AI

AEA Annual Meeting

Joshua Gans, January 2025

Andreessen Horowitz is warning that billions of dollars in AI investments could be worth a lot less if companies developing the technology are forced to pay for the copyrighted data that makes it work.

The VC firm said AI investments are so huge that any new rules around the content used to train models “will significantly disrupt” the investment community's plans and expectations around the technology, according to comments submitted to the US Copyright Office.

“The bottom line is this,” the firm, known as a16z, wrote. “Imposing the cost of actual or potential copyright liability on the creators of AI models will either kill or significantly hamper their development.”

a16z, 2023

… a story about venture capitalists building a runaway artificial intelligence that will likely enslave or destroy humankind, only to be thwarted by a minor poet suing them for copyright violations for scraping her poems.

Terminator: Fair Use Doctrine. What if fastidious enforcement of intellectual property rights is all that stands between us and annihilation by robots?

Matt Levine (2023)

"Hilarious, provocative, and supersmart, *Year Zero* is a brilliant novel to be enjoyed in perpetuity in the known universe and in all unknown universes yet to be discovered."

—**JOHN HODGMAN**, RESIDENT EXPERT, *THE DAILY SHOW WITH JON STEWART*

YEAR ZERO

NEW YORK TIMES BESTSELLER

ROB REID

a novel

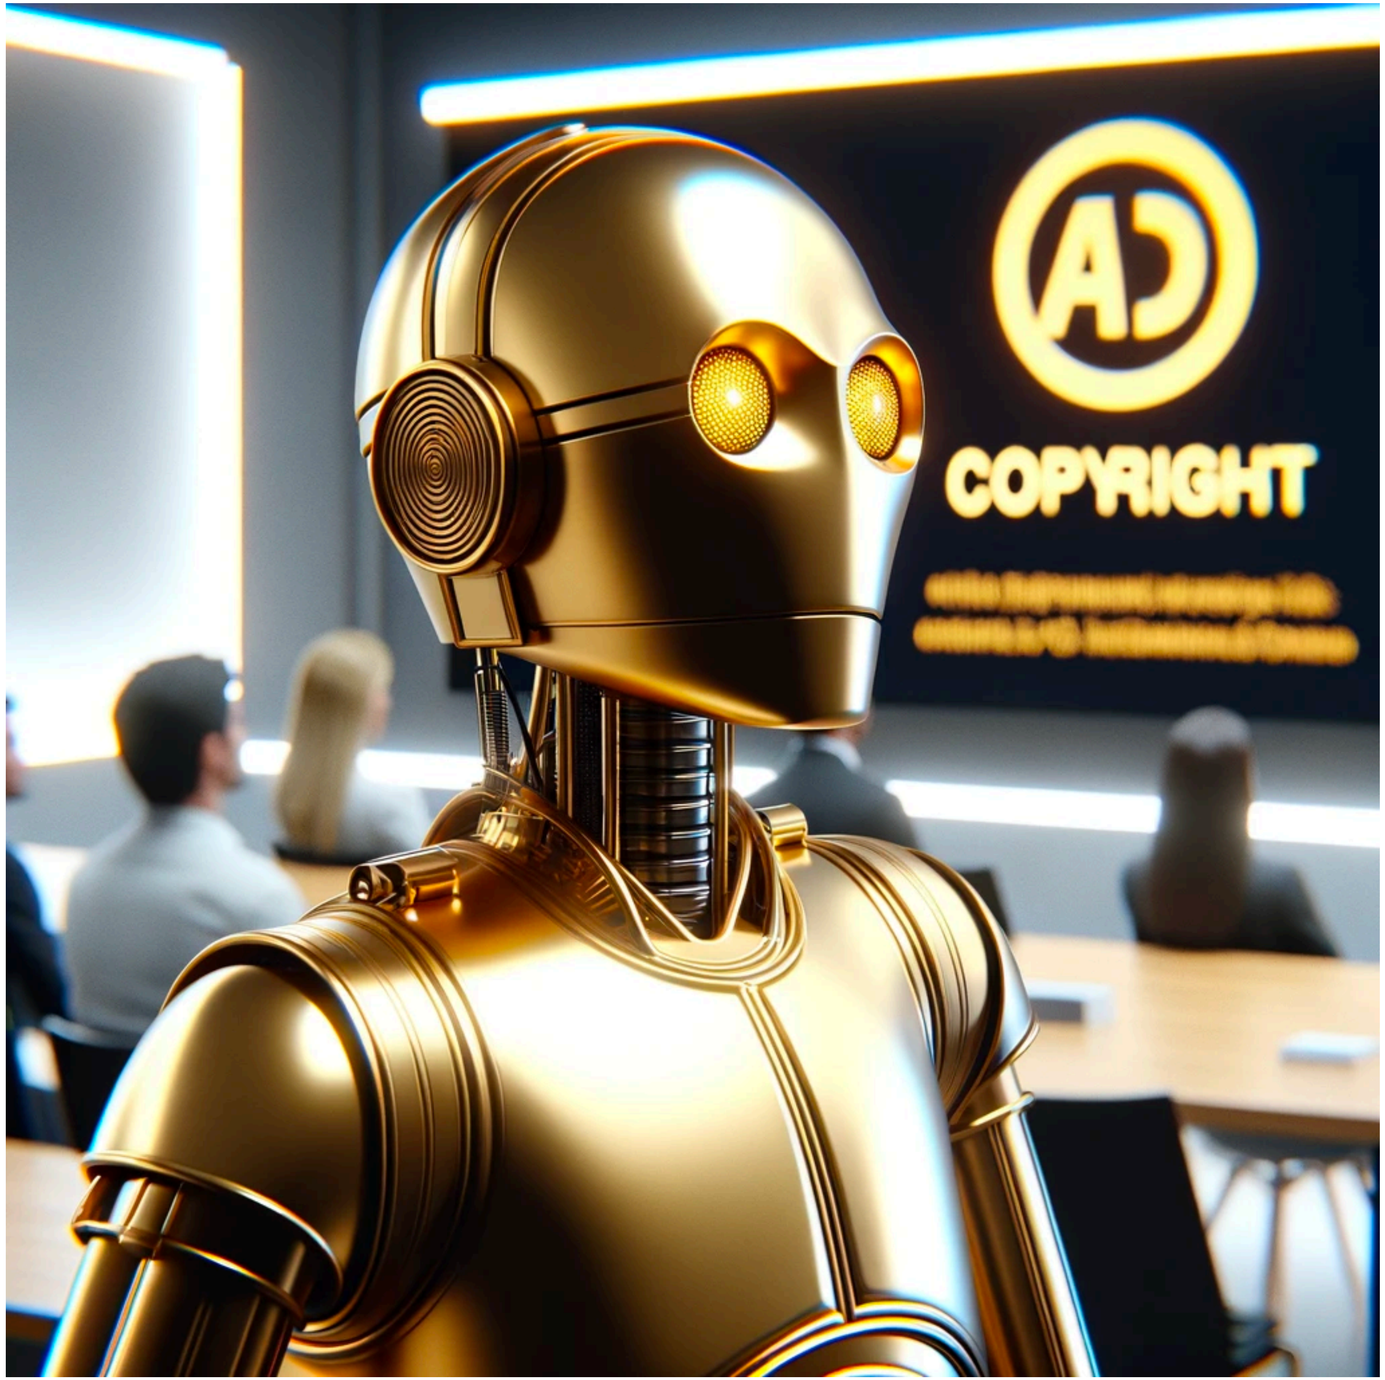

DALL-E Prompt: Show me a droid looking interested during a presentation about copyright. Make it a golden droid like in the movies.

Output from GPT-4:

exempted it from regulations, subsidized its operations and promoted its practices, records and interviews showed.

Their actions turned one of the best-known symbols of New York — its yellow cabs — into a financial trap for thousands of immigrant drivers. More than 950 have filed for bankruptcy, according to a Times analysis of court records, and many more struggle to stay afloat.

"Nobody wanted to upset the industry," said David Klahr, who from 2007 to 2016 held several management posts at the Taxi and Limousine Commission, the city agency that oversees medallions. "Nobody wanted to kill the golden goose."

New York City in particular failed the taxi industry, The Times found. Two former mayors, Rudolph W. Giuliani and Michael R. Bloomberg, placed political allies inside the Taxi and Limousine Commission and directed it to sell medallions to help them balance budgets and fund key initiatives.

During that period, much like in the mortgage lending crisis, a group of industry leaders enriched themselves by artificially inflating medallion prices. They encouraged medallion buyers to borrow as much as possible and ensnared them in interest-only loans and other one-sided deals that often required borrowers to pay hefty fees, forfeit their legal rights and give up most of their monthly incomes.

When the market collapsed, the government largely abandoned the drivers who bore the brunt of the crisis. Officials did not bail out borrowers or persuade banks to soften loan

Actual text from NYTimes:

exempted it from regulations, subsidized its operations and promoted its practices, records and interviews showed.

Their actions turned one of the best-known symbols of New York — its signature yellow cabs — into a financial trap for thousands of immigrant drivers. More than 950 have filed for bankruptcy, according to a Times analysis of court records, and many more struggle to stay afloat.

"Nobody wanted to upset the industry," said David Klahr, who from 2007 to 2016 held several management posts at the Taxi and Limousine Commission, the city agency that oversees cabs. "Nobody wanted to kill the golden goose."

New York City in particular failed the taxi industry, The Times found. Two former mayors, Rudolph W. Giuliani and Michael R. Bloomberg, placed political allies inside the Taxi and Limousine Commission and directed it to sell medallions to help them balance budgets and fund priorities. Mayor Bill de Blasio continued the policies.

Under Mr. Bloomberg and Mr. de Blasio, the city made more than $855 million by selling taxi medallions and collecting taxes on private sales, according to the city.

But during that period, much like in the mortgage lending crisis, a group of industry leaders enriched themselves by artificially inflating medallion prices. They encouraged medallion buyers to borrow as much as possible and ensnared them in interest-only loans and other one-sided deals that often required them to pay hefty fees, forfeit their legal rights and give up most of their monthly incomes.

---

[26] For original article, *see* Brian M. Rosenthal, *As Thousands of Taxi Drivers Were Trapped in Loans, Top Officials Counted the Money*, N.Y. TIMES (May 19, 2019), https://www.nytimes.com/2019/05/19/nyregion/taxi-medallions.html.

For a distinguished example of investigative reporting, using any available journalistic tool, Fifteen thousand dollars ($15,000).

Brian M. Rosenthal of The New York Times

For an exposé of New York City's taxi industry that showed how lenders profited from predatory loans that shattered the lives of vulnerable drivers, reporting that ultimately led to state and federal investigations and sweeping reforms.

Brian M. Rosenthal accepts the 2020 Pulitzer Prize for Investigative Reporting from Columbia University President Lee Bollinger. (Jose Lopez/The Pulitzer Prizes)

WINNING WORK

2. Permanently enjoining Defendants from the unlawful, unfair, and infringing conduct alleged herein;

OpenAI

3. “Regurgitation” is a rare bug that we are working to drive to zero

Our models were designed and trained to learn concepts in order to apply them to new problems.

Memorization is a rare failure of the learning process that we are continually making progress on, but it’s more common when particular content appears more than once in training data, like if pieces of it appear on lots of different public websites. So we have measures in place to limit inadvertent memorization and prevent regurgitation in model outputs. We also expect our users to act responsibly; intentionally manipulating our models to regurgitate is not an appropriate use of our technology and is against our terms of use.

Just as humans obtain a broad education to learn how to solve new problems, we want our AI models to observe the range of the world’s information, including from every language, culture, and industry. Because models learn from the enormous aggregate of human knowledge, any one sector—including news—is a tiny slice of overall training data, and any single data source—including The New York Times—is not significant for the model’s intended learning.

https://openai.com/blog/openai-and-journalism

Difficult Challenges

- Without being ‘wilful,’ AI generated output can potentially create commercial harm to content creators
  - Cause consumers to substitute from purchasing from original content provider
  - Lead to unauthorised use of icons and assets
- These issues were always present with humans …

Scenario 1: Spoilers

*A **developer** creates a chatbot that allows users to ask questions about specific television series, including plot details, characters and key quotes from the series. The chatbot is trained on transcripts and other data from the television series.*

*A **person** answers questions about a specific television series over social media, including plot details, characters and key quotes from the series. The person has watched the television series multiple times.*

Scenario 2: TL;DR

*A **developer** builds and sells access to a website that provides summaries of business books. The summaries were generated by AI and trained on scanned text from books the developer purchased.*

*A **person** sells summaries of business books. These summaries were written by the person after the person had purchased and read a copy of each book.*

Scenario 3: Fan Art

*A **developer** builds an image generator that can create comic frames for Marvel superheroes.*

*A **person** sells hand drawn comics based on Marvel superheroes.*

If the human version is not infringing …
why is the AI version infringing?

Why is the human version is not infringing?

International Journal of Industrial Organization 41 (2015) 76–83

Contents lists available at ScienceDirect

International Journal of Industrial Organization

journal homepage: www.elsevier.com/locate/ijio

Remix rights and negotiations over the use of copy-protected works☆

Joshua S. Gans

University of Toronto, Toronto, ON Canada
NBER, Cambridge, MA, USA

ARTICLE INFO

Article history:
Received 23 December 2014
Received in revised form 8 June 2015
Accepted 20 June 2015
Available online 2 July 2015

JEL Classification:
O34

Keywords:
Copyright
Incentives
Remixing
Compensation

ABSTRACT

This paper examines an environment where original content can be remixed by follow-on creators. The modelling innovation is to assume that original content creators and remixers can negotiate over the 'amount' of original content that is used by the follow-on creator in the shadow of various rights regimes. The following results are demonstrated. First, traditional copyright protection where the original content creators can block any use of their content provides more incentives for content creators and also more remixing than no copyright protection. This is because that regime incentivises original content creators to consider the value of remixing and permit it in negotiations. Second, fair use can improve on traditional copyright protection in some instances by mitigating potential hold-up of follow-on creators by original content providers. Finally, remix rights can significantly avoid the need for any negotiations over use by granting those rights to follow-on innovators in return for a set compensation regime. However, while these rights are sometimes optimal when the returns to remixing are relatively low, standard copyright protection can afford more opportunities to engage in remixing when remixing returns are relatively high.

1. Introduction

Remixing is a term used to describe taking content (sound, music, photos and words) and altering them in some manner to create new content. While the notion of derivative works has existed in copyright law for some time, digital technology has allowed a wider range of content types to be combined to produce remixed creative works and for that work to reach a wider audience. For example, users now attract millions of views on YouTube with their own video representations of popular songs (e.g., the Harlem Shake phenomenon) or re-working of television shows and movies (e.g., the synchronisation of video of George W. Bush and Tony Blair to the duet, Endless Love). For the purposes of this paper, *remixing* occurs whenever someone takes copy-protected content and repurposes it in some manner. Thus, while it does not include direct copying, parts, or maybe all, of the copy-protected material are used in the derivative work.

Legal and economic scholars have been challenged in considering how such remixing should be treated. In law, there is a position that it may fall under 'fair use.'[1] Fair use is an exemption to copyright protection, that exists in some jurisdictions, related to works that use copy-protected materials for reviews, discussion and parody. This may, in some circumstances, encompass remixed content. However, it should be noted that a fair use exemption does not exist everywhere; including, notably, Europe and Australia. Moreover, while fair use can potentially be readily applied to non-commercial remixing and use of copy-protected material, when it is uploaded or distributed using commercial platforms (such as YouTube or SoundCloud), the precise nature of the content becomes ambiguous.

In this paper, we explore remix rights in the context of considering the optimality of copyright law from a property rights perspective. That is, the law sets default rights on various parties. Usually, in fact, no copy-protected material can be used in any form by others unless express permission is given by the copyright owner. Thus, potential user/creators need to obtain permission to use such materials. As Lessig (2008) notes, this has occurred on open platforms such as YouTube where amateurs, for instance, uploaded videos with songs playing in the background only to be cited with takedown notices because appropriate permissions had not been sought. These takedown notices invariably occurred after remix effort had been expended. In other cases cited by Lessig, users who sought permissions found transaction costs prohibitive. Google and other platforms have since opted for solutions that encourage remixed content and lower transaction costs for obtaining permissions.

Lessig (2008) argues that these measures still stifle creativity. He claimed that technology has made copying so easy that it is hard to base the structure of the law on the presumption that copying can be prevented. He argues that remix rights should be applicable where the user has a demonstrably amateur characteristic. For instance, a song in the background of a family video is not copy-protected but if it is

☆ I thank Ariel Katz, Shane Greenstein, Avi Goldfarb, an anonymous referee, and seminar participants at the University of Toronto Law School for helpful comments. The latest version of this paper is available at research.joshuagans.com.

E-mail address: joshua.gans@gmail.com.

[1] For a review see Katz (2013).

http://dx.doi.org/10.1016/j.ijindorg.2015.06.004
0167-7187/

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: 26 February 2024.

Model setup

Original content creator (OC): quality x at cost $c_{OC}(x)$

AI provider (AI): quality y at cost $c_{AI}(y, sx)$ where $s \in [0,1]$ is a sample of original content

Consumers: mass $\theta \sim U[0,1]$ with WTP for content of $x\theta$ and WTP of AI of $u(y)$

Prices: OC at p_{OC} and AI at p_{AI}

Leakage: use of s can allow AI consumers to access original content with probability $\rho(s)$

Consumer choices

Purchase both OC and AI: $\theta x - p_{OC} + u(y) - p_{AI}$

Purchase AI only: $u(y) + \rho(s)\theta x - p_{AI}$

Incremental WTP for OC is: $(1 - \rho(s))\theta x$

Leakage reduces demand for original content from OC

First Best

$$\max_{x,y,s} \int_0^1 \theta x \, d\theta - c_{OC}(x) + u(y) - c_{AI}(y, sx)$$

Set $s^* = 1$

x^* takes into account benefit to AI training

Timing

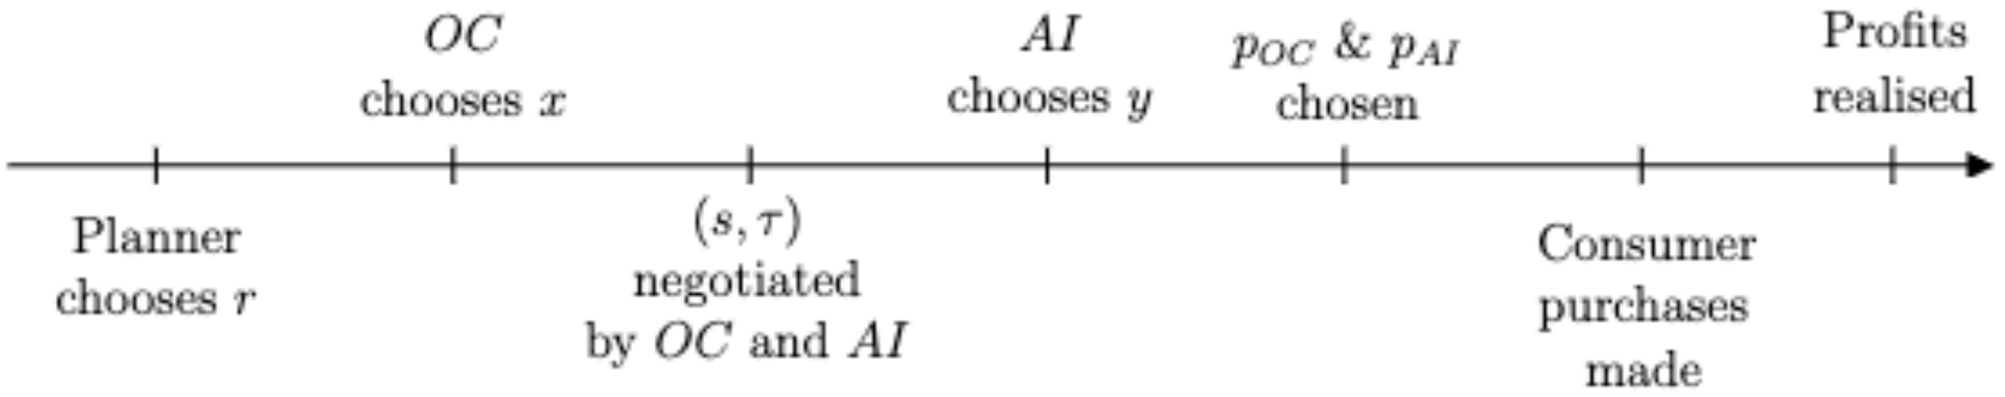

Figure 1: **Model Timeline (with Negotiations)**

Timing

Can this happen?

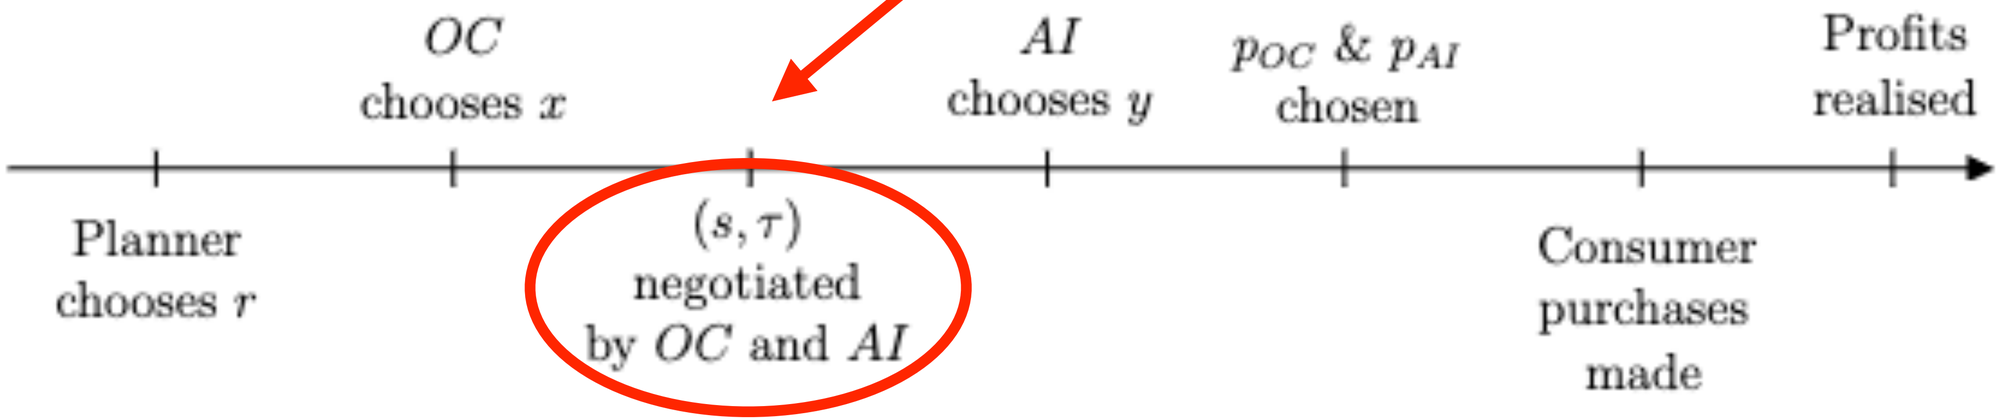

Figure 1: **Model Timeline (with Negotiations)**

Outcomes

Proposition: Comparing (CP), with $s_r = 0$ and (NC) with $s_r = 1$:

- $\hat{x}_{CP} > \hat{x}_{OC}$
- $\hat{y}_{CP} > \hat{y}_{OC}$
- Expected social welfare higher under (CP) than (NC)

“Small” AI Model

Use specific content to train the AI

- Negotiations over the use of content are possible *prior* to AI training
- Negotiations take place *after* original content created (hold-up)
- Copyright regime impacts on what happens if negotiations break down
- **Copyright protection is superior to no copyright because it creates the maximal incentives for the original content provider.**
- But … leads to too little use of the original content either in AI training or in the consumption of that original content.

“Large” AI Model

Cannot negotiate over terms

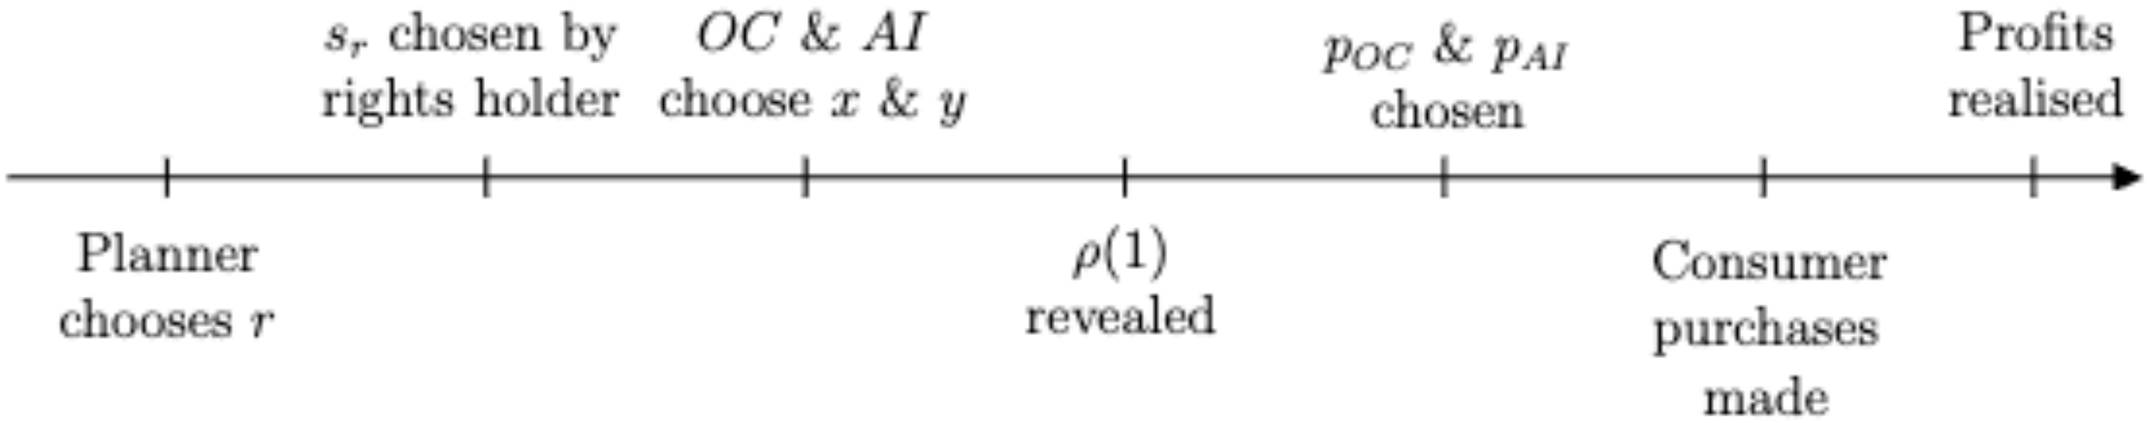

Figure 2: **Large AI Model Timeline**

Outcomes

Comparing traditional rights regimes

Let γ be the marginal value of $s \int x_i \, di$ in reducing AI training costs

Proposition: Expected social welfare under (NC) exceeds that under (CP) if and only if:

$$\gamma > \frac{\mathbb{E}[\rho_i(1)]^2}{2(1 - \mathbb{E}[\rho_i(1)])}$$

“Large” AI Model

Use large amount of general content to train the AI

- Negotiations over the use of content are not possible (too many content providers)
- Copyright regime impacts incentives to invest in content & AI training
- **Copyright protection is only preferred to no copyright if the commercial harm to original content providers is greater than the general value of content in lowering AI training costs.**
- Humans using content is OK because it is thought not to create commercial harm to content providers. Same should apply to AI.

Ex post ‘Fair Use’ Regime (FU)

An alternative mechanism

1. (Ownership) OC holds copyright
2. (Fair Use) Cannot prevent use in AI training
3. (Ex post) OC can (potentially) make a claim for damages
   - If $\rho_i(1) < \Gamma$, AI not liable for damages
   - If $\rho_i(1) \geq \Gamma$, AI must pay full compensation of $\frac{1}{4}\rho_i(1)\hat{x}_{FU}$

Ex post ‘Fair Use’ Regime (FU)

Outcomes …

As $\Gamma \to 0$, $\hat{x}_{FU} \to \hat{x}_{CP}$ and $\hat{y}_{FU} \to y^*(\hat{x}_{CP})$

Set liability threshold to ensure AI profitable in expectation

Ex post `fair use' like mechanism

Infringement damages only payable if above a threshold

- All content providers hold copyright but cannot prevent use in AI training.
- Ex post, after commercial harm realised, content providers can obtain damages from AI provider if harm greater than some threshold.
- This mechanism leads to …
  - more original content than no copyright but less the full copyright
  - superior AI training than either no copyright or full copyright
- As the threshold gets smaller, this mechanism dominates other rights regimes

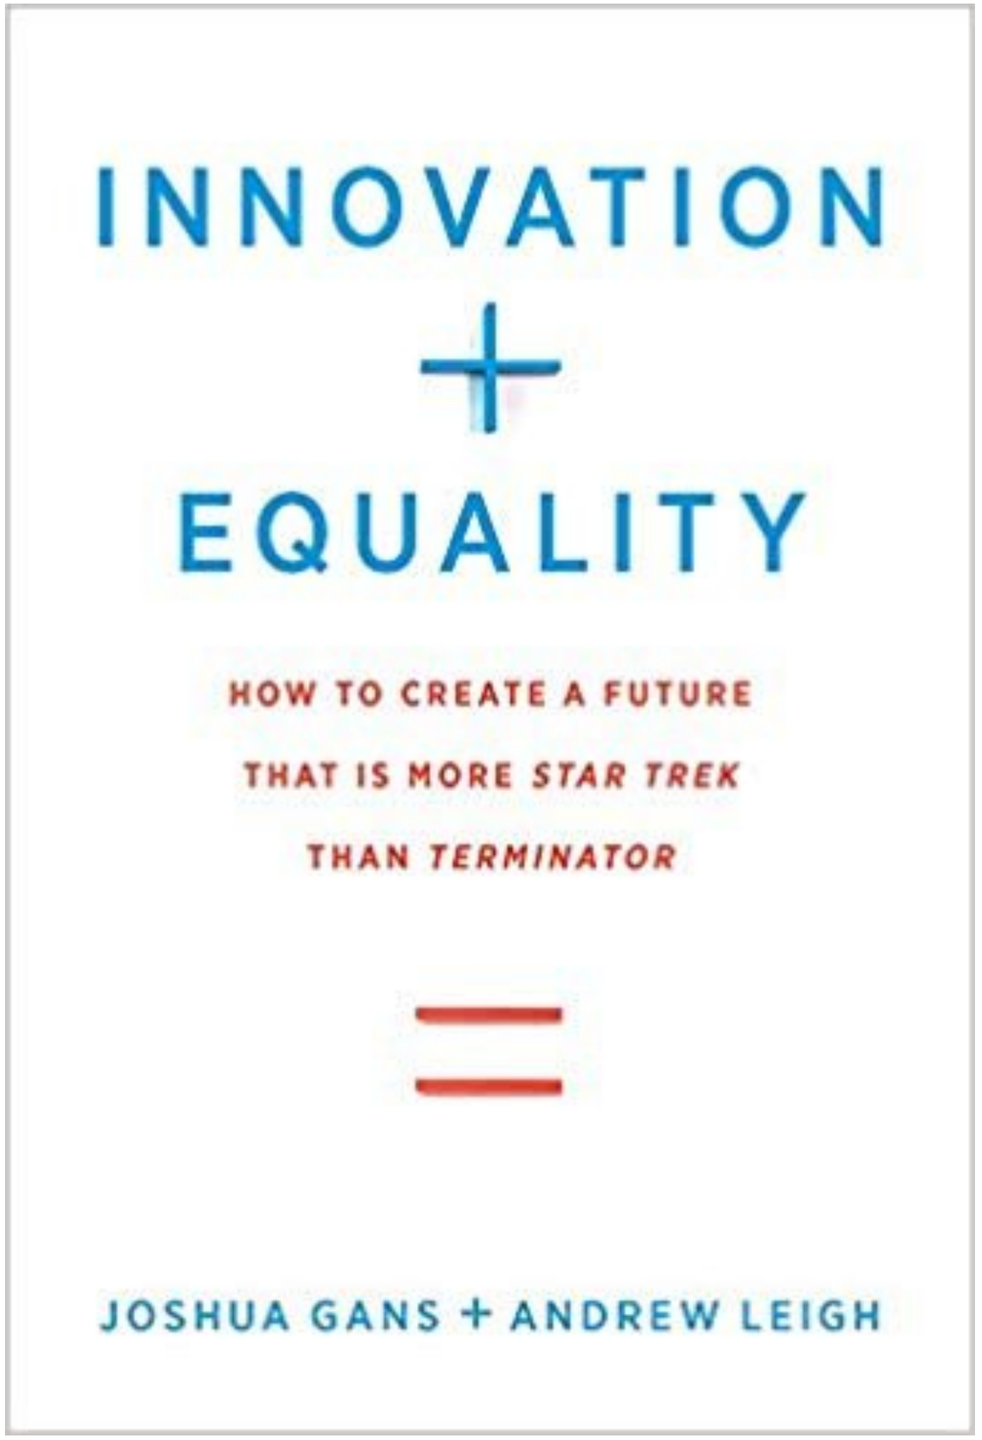
INNOVATION
+
EQUALITY
HOW TO CREATE A FUTURE
THAT IS MORE *STAR TREK*
THAN *TERMINATOR*
=
JOSHUA GANS + ANDREW LEIGH