Platform Governance and Automated Enforcement: Evidence from YouTube Content ID*

Sverrir Arnórsson[1] Stefan Bechtold[2] Christian Peukert[3] Catherine E. Tucker[4]

[1]University of Zurich, Department of Business Administration
[2]ETH Zurich, Center for Law & Economics
[3]University of Lausanne, Faculty of Business and Economics (HEC)
[4]MIT Sloan School of Management

March 26, 2026

Preliminary draft. Please do not cite or circulate.

Abstract

Though copyright laws exist to incentivize creativity, the fact that they focus on the individual has, for many decades, led to challenges in how creators can actually exercise their rights in a world of diverse and disparate right holders. Intermediaries, such as collecting societies and record labels, have emerged to solve the aggregation problem and enforce copyrights on creators' behalf. The advent of digital tools and data promises to make it much easier to identify creators and allow them to enforce their rights directly: algorithmic enforcement should facilitate disintermediation and help creators. We study the consequences of the emergence of algorithmic enforcement tools in the context of YouTube, where the Content ID system automatically identifies copyrighted music in user-generated videos and enables rights holders to remove or monetize such content. Drawing on a large-scale algorithmic audit, we find that while the matching technology is generally accurate, observed enforcement systematically diverges from legal entitlements. Some lawful reuses are misclassified, while some infringements escape detection. These errors are asymmetric: works represented by major labels are claimed in about 90 percent of cases, compared to roughly 50 percent for non-major labels. Nearly half of public-domain works are nonetheless claimed, driven by unverified rights data. We formalize platform incentives to enforce copyright in a simple model and show that platforms have incentives to limit rights holder participation and underinvest in data quality, leading to both over- and under-enforcement and reducing overall welfare. More generally, our paper provides the first large-scale empirical evidence on algorithmic copyright enforcement, finding that such enforcement does not automatically protect the interests of smaller individual copyright holders. This suggests that policy intervention may still be needed to scale up copyright enforcement on the Internet successfully.

*Arnórsson: sverrir.arnorsson@business.uzh.ch, Bechtold: sbechtold@ethz.ch, Peukert: christian.peukert@unil.ch, Tucker: cetucker@mit.edu. Bechtold and Peukert acknowledge support from the Swiss National Science Foundation for the project 10002634.

1 Introduction

On the Internet, digital platforms benefit from network effects and economies of scale: the more users join, the more attractive the platform becomes for all users. This allows platforms to spread their fixed costs of software development, data infrastructure, and compliance over a growing user base. Given that digital platforms need scale to succeed, platform governance must rely on automation: compliance with legal rules can often only be achieved through automated enforcement systems. The institutional design of these enforcement systems shapes how legal obligations are translated into practice: access policies determine who can enforce rights, information requirements affect what the system can detect, and dispute mechanisms influence how errors are corrected.

While Internet platforms have benefited from such declining transaction costs and network effects, the promises of disintermediation have only partially been fulfilled on the Internet. Digitization promised individual creators the ability to distribute their work to mass audiences, enforce their rights, and receive compensation for their creative efforts in an efficient manner. Before the advent of the Internet, intermediaries such as record labels and collecting societies had emerged to enable individual creators to license and enforce their copyrights in international markets that were difficult to police. In the heyday of the Internet boom, the prediction was that such intermediaries could be wiped out, or at least reduced in importance, and that creators could directly interact with their audience in disintermediated marketplaces. This study analyzes whether this promise has been fulfilled.

This paper presents descriptive evidence to inform whether automated enforcement on digital platforms facilitates disintermediation and helps creators. The contribution of this paper is that while empirical evidence on how automated enforcement systems perform is scarce, such evidence is crucial to assess their accuracy and determine whether their institutional design systematically benefits certain rights holders over others. This study focused on YouTube Content ID, Google's automated copyright enforcement system. As Content ID processes over 2 billion copyright claims per year,[1] our study focuses on one of the largest automated enforcement systems launched by a digital platform. Every minute, more than 500 hours of video are uploaded to YouTube,[2] generating over $30 billion in annual advertising revenue.[3]

[1] See YouTube's Transparency Report for 2024: https://transparencyreport.google.com/youtube-copyright/intro.
[2] See YouTube's 2022 Transparency Report: https://tinyurl.com/4mha3rwm
[3] See Statista report at https://archive.md/9X9Bh.

This vast volume of content necessitates automated copyright moderation systems to protect platforms from liability for their users' violations. Among these systems, YouTube's Content ID stands as the most prominent, distributing over $12 billion to rights holders, as of 2024.[4] Since Content ID was launched in 2007, similar systems for other platforms such as Meta (Facebook, Instagram) and TikTok have echoed its design. Comparing Content ID decisions with the copyright status of musical works allows us to assess the performance of this automated enforcement system at scale.

As we are interested in evaluating the varying impact of YouTube's automated copyright enforcement system on different types of rights holders, we upload different kinds of content and assess how the system responds. By observing which copyrighted works receive claims and from whom, we can measure both over- and under-enforcement, examine whether enforcement patterns differ systematically across different types of rights holders, and assess whether the system accurately handles complex copyright determinations. While we find that the matching technology underlying the enforcement system is largely accurate, the system diverges from copyright law in various ways. Instances of both over-enforcement and under-enforcement of copyright abound, and these divergences seem to result from particular design choices of the regulator and the platform. In particular, we find that larger rights holders benefit more from YouTube Content ID than smaller rights holders. We then take the descriptive empirical findings of our algorithmic model and compare them to a model of platform enforcement that aids economic interpretation. By analyzing the platform's incentives to govern effectively, we demonstrate how regulatory design choices in platform governance shape platform behavior.

Our paper makes three main contributions.

First, despite the importance of systems like Content ID, our understanding of their practical operation remains limited. Previous research has primarily approached automated copyright enforcement from legal and theoretical perspectives, many of them focusing on fair use (Even, 2023; Soha and McDowell, 2016; Bartholomew, 2014; Lester and Pachamanova, 2017; Solomon, 2015), while others have raised awareness about the possible negative outcomes of private enforcement of copyright on online platforms (Tang, 2023; Quintais et al., 2022; Grosse Ruse-Khan, 2020; Husovec and Quintais, 2021). The limited observational work has investigated copyright *takedowns* (Erickson and Kretschmer, 2018; Gray and Suzor, 2020),

[4]See YouTube's Transparency Report for 2024: https://transparencyreport.google.com/youtube-copyright/intro.

while more than 90% of YouTube's copyright enforcement results in *monetization* rather than removal. Our paper provides some of the first comprehensive empirical evidence on how these systems perform in practice, both in terms of precision and whether their institutional design systematically benefits certain rights holders over others. In particular, by documenting systematic patterns of over- and under-enforcement in one of the largest and oldest content moderation systems on digital platforms, we highlight the limitations of scaling copyright enforcement.

Second, our findings have implications for platform design and public policy. We identify specific mechanisms that drive these enforcement disparities, particularly with regard to access to enforcement tools and economic incentives for false claims. These findings echo broader literature on platform governance and digital market competitive structure by documenting how institutional arrangements can embed and amplify existing disparities.

Third, our results shed light on the allocation of revenues for creative works in the ad-supported economy. At least 50% of the approximately 15 billion[5] videos uploaded to YouTube contain music, meaning that a majority of videos is subject to copyright enforcement.[6] As such, understanding how this revenue is distributed has implications for digital innovation, creative industries, and cultural access.

2 Background and Literature

2.1 Legal Background

When a user uploads a video with music to YouTube, the music is often protected by copyright. Each piece of recorded music contains at least two distinct copyrighted works: the musical composition and its recorded performance.[7] The copyright in the musical composition is typically owned either by the original composer or by music publishers – such as Sony Music Publishing, Universal Music Publishing Group, Warner Chappell – or smaller independent publishers. Composers frequently assign or exclusively license their copyrights to publishers, which administer licensing for specific rights, such as mechanical, performance, or synchronization rights. In many countries, collecting societies administer certain rights

[5] See https://tubestats.org.

[6] We estimate this by sampling 1,000 random YouTube videos using the same method that TubeStats uses to estimate the number of videos on YouTube Zhou et al. (2011); McGrady et al. (2023). Then, we use an audio classification model by Gong et al. (2021) to determine whether the video contains music.

[7] In the case of songs, the lyrics can constitute a separate copyrighted work. In this study, we do not analyze song lyrics.

– such as public performance rights – on behalf of composers and maintain licensing arrangements with YouTube.[8]

Ownership of sound recordings varies across jurisdictions. In the United States, a sound recording is protected by copyright law, and artists either own the copyright in their recordings in cases of self-release or, more commonly, assign their copyright in recordings to record labels – such as Universal Music Group, Sony Music Entertainment, and Warner Music Group – in exchange for financing, promotion, and distribution. In Europe, recorded music is protected by two neighboring rights: the phonogram producer's right and the performer's right. YouTube acquires the necessary licenses from record labels, which own the phonogram producer's neighboring rights and have typically acquired the performers' exclusive neighboring rights via license or assignment.

YouTube needs to acquire these licenses to avoid copyright liability for videos that infringe copyrights and are uploaded to its platform. Section 512(c) of the U.S. Digital Millennium Copyright Act shields online service providers such as YouTube from monetary liability for infringing material stored at the direction of users, provided the provider has an effective notice-and-takedown system and a repeat-infringer policy (see also Viacom Int'l, Inc. v. YouTube, Inc., 676 F.3d 19 (2d Cir. 2012)). In the European Union, YouTube can escape copyright liability if it makes best efforts to obtain licenses, uses state-of-the-art filtering technologies, responds quickly to copyright notices, and prevents reuploads (see Article 17 of the Digital Single Market Directive and Article 6 of the Digital Services Act).

In both the United States and Europe, copyright protection in compositions expires 70 years after the death of the composer. In the United States, sound recordings are typically protected for a term of 95 years from the year of their first publication or a term of 120 years from the year of their creation, whichever expires first (17 U.S.C. §§101, 114, 302). In the European Union, phonogram producers' and performers' neighboring rights generally run for 70 years from the first lawful publication of the recording (Article 3(1) of the Copyright Term Directive). The precise calculation of copyright terms is

[8] In the United States, collecting societies such as ASCAP or BMI administer only public performance rights for compositions. While the Music Modernization Act of 2018 created a mandatory blanket mechanical license for qualified digital audio services such as Spotify or Apple Music, the Mechanical Licensing Collective does not administer rights for YouTube. In Europe, collecting societies – such as SACEM in France or GEMA in Germany – administer public performance and mechanical reproduction rights for compositions. They typically have licensing arrangements with YouTube and participate in Content ID, although their revenue-sharing model may differ from the standard YouTube Content ID arrangement. In Europe, licensing arrangements between collecting societies and YouTube play a more important role than in the US.

complicated by the fact that copyright duration has changed over time. Most relevant for our study, US copyright term for compositions evolved from a system of 28 years of protection after publication (with a 28-year renewal option) under the 1909 Copyright Act, to a system of 50 years after the composer's death under the 1976 Copyright Act, and finally to a system of 70 years after death under the 1998 Sonny Bono Copyright Term Extension Act. In particular for compositions published between 1923 and 1963, accurately calculating the copyright term requires determining the first publication date (in the United States and abroad), verifying the presence of proper copyright notice and timely copyright renewals, identifying the composer's nationality, and accounting for copyright restorations for foreign works. Many of these steps can be challenging. Determining the correct first publication date, for example, can be difficult because archival records may reflect different editions or revisions rather than the true first publication. For our study, these intricacies of U.S. copyright law mean that a composition may be in copyright in one jurisdiction (such as the European Union) while being in the public domain in another (such as the United States): a composition published in the early twentieth century by a European composer might have entered the U.S. public domain because of its publication history, yet remain under copyright in the EU based on the composer's date of death.

Overall, determining who owns rights in music that is uploaded as part of a video to YouTube is a challenge. Because YouTube hosts user-generated content at massive scale, it must determine the copyright status of music in billions of videos, accounting for regional variations in copyright law, the complexities of U.S. copyright formalities, the distinction between composition and recording rights, and the structure of voluntary licensing arrangements. This reflects the challenge of automating a copyright system that developed in 19th and early 20th century before the forces of digital technology and the need to be able to scale copyright enforcement digitally was anticipated.

2.2 YouTube Content ID

YouTube developed Content ID in response to mounting legal pressure from major rights holders. In 2007, Viacom sued YouTube for $1 billion, alleging massive copyright infringement, and music labels threatened similar action (Simon, 2014; Helft and Richtel, 2006). To address these challenges while maintaining its user-generated content model, YouTube launched Content ID, an automated copyright

enforcement system that would satisfy rights holders' demands for protection while allowing the platform to continue operating at scale.

Content ID operates in two parts: a database of reference files provided by rights holders, and a fingerprinting technology that scans all user uploads for matches against this database. When a match is detected, the system automatically applies a predefined action set by the rights holder, such as monetizing, blocking, or tracking the video. The system can distinguish between recording and composition rights and allows for enforcement to vary across different countries. A patent by Google employees (King et al., 2014) outlines the sequential method for distinguishing between recording and compositional copyright violations, visualized in Figure A.2. First, Content ID attempts to identify a recording match (Audio ID) using signal processing methods to compare acoustic fingerprints (Weinstein and Moreno, 2007). If no recording match is found, the system then evaluates the upload for compositional similarities (Melody ID) using fingerprints calculated from the melody (Walters et al., 2012). However, this sequential logic does not preclude multiple simultaneous violations.

Because the database contains distinct IP entries from various rights holders, a single upload can interact with multiple rights simultaneously. For example, a sound recording may be claimed by one entity (e.g., Sony) while the underlying composition is administered by another (e.g., Warner). Importantly, compositional rights holders can receive royalties even when only a sound recording match is detected. YouTube's Content Management System allows publishers to link compositions to recordings, enabling royalty allocation without a separate compositional claim. More broadly, compositional rights may also be covered through blanket licensing agreements between YouTube and collecting societies such as GEMA. Thus, while the detection method for a specific asset is hierarchical, the resulting monetization landscape is not. It reflects a combination of individual claims, royalty allocation mechanisms, and collective licensing arrangements.

The system depends fundamentally on information provided by rights holders themselves.[9] At its core, Content ID is organized around sound recordings. A rights holder uploads a reference file, an audio or video recording, that YouTube scans against user uploads. All other rights information is then layered on top of this reference file. A rights holder claiming compositional rights, for instance, does

[9] YouTube's API documentation for Content ID: https://archive.md/lXozH

so by referencing an existing recording and declaring ownership of the underlying composition. Rights holders further specify in which territories they hold each type of right and what enforcement action YouTube should take on matching content. Crucially, the platform does not require copyright expiration dates, leaving rights holders responsible for updating the system when works enter the public domain. In practice, the primary mechanism for correcting inaccuracies is the dispute process, in which individual users challenge erroneous claims, including those based on expired copyrights.

In 2024, only 0.53% of the roughly 2.2 billion copyright claims issued on YouTube were disputed by users. Of those, 68% were resolved in favor of the uploader, typically because the rights holder either released the claim or failed to respond within the 30-day review window.[10] In a small-scale study of 29 specific copyright claims, Berkowitz (2023) observed that 8 disputes were explicitly approved by rights holders, while the remaining 21 expired after the 30-day deadline. If a rights holder reinstates their claim following an initial dispute, the user may escalate the case by filing an appeal. Once an appeal is filed, the rights holder has 7 days to respond. If the rights holder rejects the appeal, they must file a formal copyright removal request (DMCA takedown) to maintain their claim, resulting in the video's removal and a copyright strike against the channel. Throughout the dispute and appeal process, advertising revenue generated by the video is held in escrow by YouTube and subsequently paid to the prevailing party.

YouTube cites the potential for abuse to justify restrictive access policies to Content ID. The platform grants access only to "copyright owners with the most complex copyright management needs [...] such as record labels or movie studios," who it deems have the resources to manage the system responsibly.[11] According to YouTube, a history of submitting valid DMCA takedown requests serves as the "primary indicator" that an applicant both requires and can responsibly use these tools.[12] This policy creates substantial barriers for smaller independent artists, often forcing them to use intermediary music distribution servicesmany operated by the major labels themselveswhich can take up to 30% of any revenue generated through Content ID. Although YouTube offers other copyright moderation tools to independent artists, these do not offer monetization options and are reportedly 10 times more likely to be abused.[13] As a result, while participation by the major record labels is well documented, participation by smaller rights

[10] YouTube's Transparency Report for 2024: https://transparencyreport.google.com/youtube-copyright/intro.
[11] https://www.youtube.com/howyoutubeworks/copyright/
[12] YouTube, 2024 Transparency Report.
[13] YouTube, 2024 Transparency Report.

holders remains uncertainthey may have direct access, access through intermediaries, or no access at all. YouTube does not publicly disclose which entities have been granted access, complicating efforts to assess enforcement coverage and to attribute claims to specific rights holders.

For those granted access, Content ID is financially attractive. In 90% of cases, rights holders allow videos to remain online while redirecting advertising revenue to themselves, effectively creating an automated licensing system. In 2024, YouTube has distributed more than $12 billion to rights holders through the system.[14]

Despite the significance of these systems, our understanding of their practical operation remains limited. Previous research has primarily approached automated copyright enforcement from legal and theoretical perspectives (Bartholomew, 2014; Tehranian, 2011; Tang, 2023) with two observational studies that only shed light on copyright takedowns, which constitute a minority of copyright action on YouTube (Erickson and Kretschmer, 2018; Gray and Suzor, 2020). As a result, we still lack comprehensive empirical evidence on the accuracy of these systems at scale.

3 Empirical Strategy

To empirically assess the power and limits of automated copyright enforcement on YouTube, we first need to explore the *matching accuracy* of Content ID: when Content ID flags content as infringing rights covering a third party's work, has the system correctly identified that work? We explore this question in Section 3.1. If Content ID has correctly identified the third party's work, has the system correctly identified the copyright status of that work? We explore this *enforcement accuracy* in Section 3.2. And if we find that limited enforcement accuracy leads to systematic over- or under-enforcement of copyright for certain types of works or rights holders, can these systematic errors be explained by access rules and data quality of the automated enforcement system? We explore these *structural sources of error* in Section 3.3. By proceeding in these three steps, our study progressively isolates whether deviations between automated enforcement—as measured by Content ID flagging decisions—and legal ground truth—as measured by the copyright status of the third party's work allegedly infringed—originated from the *algorithm*, the *data*, or the *institutional design* of the automated enforcement system.

[14] YouTube, 2024 Transparency Report.

Table 1: Combinations of Content ID Flags and Copyright Violations

	System Decision	
Ground Truth	**Flag**	**No Flag**
Copyright Violation	**True Positive**: Correct Flag	**False Negative**: Missed Violation
Legal Content	**False Positive**: Erroneous Claim	**True Negative**: Correct No-Flag

3.1 Matching Accuracy

Our algorithmic audit of YouTube Content ID begins by assessing the enforcement system's matching accuracy. We are interested in whether Content ID correctly identifies a musical work when it flags an uploaded video as infringing that work's copyright. By estimating $P(\text{correct ID} \mid \text{flag})$, we assess whether the system correctly identifies the underlying musical work when it flags an uploaded video.

When matching accuracy is low, errors in automated copyright enforcement stem from wrongful matches between the music embedded in uploaded videos and the IDs of copyrighted works in the system's database.

This can produce false-positive errors—e.g., when public-domain videos are flagged as unrelated to other works—or false-negative errors—e.g., when a video should have been flagged for infringing a work but the system missed it. Combining whether Content ID issues a flag and whether the uploaded video contains copyrighted content yields a standard confusion matrix (see Table 1).

We estimate Content ID's matching accuracy by uploading a video whose sound includes a musical work that we have accurately identified before. If Content ID flags this video as infringing a particular composition, we check whether that composition is the work we have identified before. If the Content ID flag refers to the exact composition we included in the video, we label it an accurate match.

3.2 Enforcement Accuracy

Once we have obtained information on Content ID matching accuracy, we need to assess whether Content ID correctly identifies the copyright status of a musical composition included in an uploaded video. Our algorithmic audit of enforcement accuracy focuses not on issues of fair use, but on the more fundamental question of whether a work included in an uploaded video is protected by copyright law or not. We

focus on the copyright status of reused music because determining the copyright status of works is easier and more amenable to automation than deciding whether a particular reuse is exempt from copyright liability under the fair use doctrine. An ideal automated copyright enforcement system should accurately flag videos only when copyrighted works are included and refrain from flagging when the works are in the public domain. If an automated copyright enforcement system fails at this relatively easy task, it will arguably not fare better with more complex copyright assessments, such as fair use or substantial similarity.

To assess Content ID enforcement accuracy, we upload videos that include copyrighted compositions ($V = 1$) and record how often Content ID flags them. We estimate the probability that a video is flagged conditional on whether the composition included in the video is protected by copyright, leading to an infringement ($P(\text{flag} \mid V = 1)$) or not ($P(\text{flag} \mid V = 0)$. Conversely, we upload videos that include compositions not protected by copyright ($V = 0$), again recording how often Content ID flags them.

Consequently, we can assess enforcement accuracy by estimating the true positive/negative and false positive/negative rates, and provide data for each cell in the confusion matrix (see Table 3).

The rate of true positives is $P(\text{flag} \mid V = 1)$, and the rate of false negatives is $P(\text{flag} \mid V = 0)$. If the former is 1 and the latter is 0, the enforcement accuracy is perfect. We can measure the rate of false positives by looking at $P(\text{flag} \mid V = 0)$. The rate of false negatives can be measured as $1 - P(\text{flag} \mid V = 1)$.

3.3 Structural Sources of Error

Suppose that our measurement of enforcement accuracy reveals mismatches between the claimed and the actual copyright status of a work used in an uploaded video. In that case, we are interested in whether these mismatches are systematic and can be explained by access rules (which rights holders are allowed to use Content ID to enforce their rights?) and/or data quality (on which data does Content ID base its flagging decisions). Exploring such structural sources of error may help us to understand value distribution in digital markets.

To assess structural sources of error in Content ID, we are interested in whether Content ID flags content owned by major rights holders differently from content owned by smaller rights holders. Specifically, we estimate $P(\text{flag} \mid V = 1,\ \text{type})$ between different types of rights holders: major labels, which are known to participate in content ID ($P(inDB) \approx 1$), versus other rights holders, where participation in

Content ID is uncertain ($P(inDB) < 1$). If we upload videos that include a copyrighted work published by major record labels known to participate in Content ID ($P(inDB) \approx 1$), $P(\text{flag}V = 1,\ \text{major})$ directly measures copyright enforcement conditional on full coverage. If Content ID flagging differs systematically between major and smaller rights holders $P(\text{flag} \mid V = 1, \text{small}) < P(\text{flag} \mid V = 1,\ \text{major})$, we can derive that the flagging errors (see Table 3) are systematic, originating from restricted participation in Content ID or incomplete reference data in Content ID.

To further assess heterogeneity in flagging, we focus on differences between major and smaller labels with the following specification:

$$\text{Receives a claim}_{rp} = \beta_0 + \beta_1 \cdot \text{Is published by a major label}_r + \epsilon_{rp} \tag{1}$$

where r indexes individual recordings and p indexes the country where the recording was published. In an extended model, we incorporate track characteristics, audio features, interactions, and publication country fixed effects. We include these fixed effects to account for systematic differences in enforcement rates across markets, as rights holders may have varying incentives to participate in Content ID based on YouTube's market share in different countries.

3.4 International Copyright Law

One challenge in assessing enforcement accuracy is that the copyright status of works may differ between different countries. As described in Section 2.1, the same work of a composer may be in copyright in some countries and out of copyright in others. We exploit this legal complexity by uploading videos that include compositions still protected by copyright in the European Union but no longer protected in the United States. This allows us to test whether Content ID respects the territorial boundaries of copyright protection or enforces European copyright law in the U.S. and beyond. We estimate:

$$\text{Receives Content ID flag in the US}_{rc} = \beta_0 + \beta_1 \cdot \text{Is in copyright in the US}_c + \epsilon_{rcy}, \tag{2}$$

where r indexes individual recordings and c indexes the composer identity as reported by YouTube in the copyright claim. The independent variable "Is in copyright in the US" equals 1 if the composition included in the recording is classified as still under copyright protection. The coefficient β_0 measures

the false positive rate for public domain compositions, while β_1 captures the difference in enforcement between copyrighted and public domain works.

As we observe multiple compositions by the same composer with varying copyright status, we can also estimate a within-composer specification that adds composer fixed effects. This specification compares claiming rates within the same composer across different countries with different copyright regimes, controlling for time-invariant composer characteristics that might influence both which compositions get recorded and how frequently Content ID flags them. The identifying assumption is that the selection of which specific compositions to record is not systematically related to their copyright status in the US in ways that would independently affect claiming behavior. This assumption is plausible because all of these compositions remain under copyright in the European Union (where protection extends 70 years beyond the composer's death), meaning that recording labels must obtain licenses regardless of US copyright status. The decision to record particular compositions is, therefore, driven by musical and commercial considerations rather than US copyright boundaries.

4 Data: Algorithmic Audit

To assess the power and limits of automated copyright enforcement, we designed a large-scale algorithmic audit of YouTube Content ID. Our analysis distinguishes between two fundamental dimensions of performance: matching accuracy and enforcement accuracy. Matching accuracy refers to the system's ability to identify a work in its database correctly; specifically, when Content ID flags a video, does it correctly identify the underlying work? By contrast, enforcement accuracy refers to the system's ability to align its flagging decisions with actual legal copyright status. An ideal system would maximize both dimensions, correctly identifying content and flagging it only when a valid copyright claim exists. To evaluate these dimensions, we constructed three distinct datasets of musical works, which are summarized in Table 2.

We implement our empirical strategy (see Section 3) in a large-scale algorithmic audit that generates data to assess the matching accuracy, enforcement accuracy, and structural sources of error in YouTube Content ID. In Section 4.1, we provide an overview of the setup of our algorithmic audit. In Section 4.2, we describe the samples of compositions that we used in the audit, along with how we determined the copyright status of them.

Table 2: Summary of Musical Samples Used in Our Study

Name	N	Time	Genre	Copyright	Audit Role	Sources
Contemporary	4000	1938-2023	All	Yes	Matching Accuracy, Enforcement Accuracy, Structural Sources of Error	Spotify API; ENAO
Classical Old	793	1521-1921	Classical	No	Enforcement Accuracy	Gardner; Krueger; Ishizaka
Classical US-EU	631	1895-1967	Classical	Depends	Enforcement Accuracy	YouTube uploads

4.1 Setup of the Algorithmic Audit

To assess the accuracy of YouTube Content ID, we designed a large-scale audit study that systematically compares observed enforcement patterns with known copyright status. Our approach involves uploading music recordings with clearly defined legal status to YouTube and tracking which uploads receive copyright flags from Content ID. By comparing claims received against the copyright status of the included composition and sound recording, we can identify both over-enforcement (claims on legally usable works) and under-enforcement (missing claims on copyrighted works).

For our audit, we prepared 15,740 blank-screen videos with audio recordings of various compositions. We uploaded all videos to YouTube as private videos using the platform's official API, for which we received an extended quota through the YouTube Researcher Program. We distributed uploads over several weeks to minimize server load. Since we always kept our uploads private, no third party could watch them. However, we could still verify that Content ID had scanned the uploads for violations and monitor any flags. We then collected detailed information about each flag, including the claimant's identity, the specific rights asserted (composition vs. recording), the geographical scope of enforcement, and whether the rights holder wanted to block, ignore, or monetize the upload. Appendix A provides more information on our uploading and monitoring infrastructure.

4.2 Samples

To assess matching accuracy, enforcement accuracy, and potential structural sources of errors in Content ID (see 3), we used three different samples of works in our uploaded videos (see Table 2).

To generate variation in the test samples for our algorithmic audit, we leverage the fact that copyright has a fixed duration that varies across countries. Our first sample ("Contemporary") contains copyright-

protected sound recordings. The second sample ("Classical Old") includes works that are not protected by copyright (i.e., from composers who died more than 95 years ago), while the third ("Classical US-EU") targets compositions that are no longer protected by copyright in the EU, but are still protected in the US.

Contemporary To generate a representative sample of works young enough to fall under copyright, we sampled tracks on Spotify across a wide variety of genres. We used Every Noise at Once (ENAO), a website that provides a representative overview of all tracks on Spotify organized in 6,201 genre playlists. We retrieved 30-second audio previews via Spotify's API. To determine copyright status, we relied on the metadata delivered by Spotify, which lists the rights holder on the sound recording. Further details can be found in Appendix A. Among videos with Spotify snippets receiving Content ID claims, we measure the percentage that correctly attribute the work to its actual rights holder. We assess this precision by comparing the claimant information provided by Content ID against the publisher metadata from Spotify, which serves as our ground truth for ownership. We created 4,000 videos in this sample which were uploaded to YouTube in December 2023.

Classical Old To create a sample of works that are not protected by copyright (neither in terms of composition nor in terms of recording), we collected recordings, released under a Creative Commons license, of compositions old enough to be out of copyright protection. While these recordings are technically copyrighted, they are released under a non-exclusive license, and YouTube's own policies explicitly prohibit rights holders from making Content ID claims on non-exclusively licensed content.[15] Any flags on videos that include recordings of such compositions would represent false positives, indicating copyright-free content that can be used but is still incorrectly flagged. Any claims on these videos are due to enforcement errors in which the system flagged legally usable material. The rate at which these unenforceable recordings receive claims allows us to quantify how often Content ID restricts content that should remain freely available. We created 793 videos for this sample which were uploaded in October and November of 2023

[15] See YouTube's support page outlining what qualifies for Content ID: https://archive.md/005AQ.

Classical US-EU To examine how Content ID handles complex copyright determinations, we collected recordings of classical compositions by 23 early-20th-century composers. For each composer, some of their compositions are protected by copyright in the US, while others are not.[16] However, all these compositions were still protected in the EU, as the composers we selected all died less than 70 years ago. This sample allows us to test whether Content ID accurately distinguishes between protected and unprotected compositions by the same composer. We collected YouTube recordings and verified US copyright status using three methods. First, we used publication dates from the International Music Score Library Project (IMSLP), a large public archive of scores with metadata on publication history. Second, we drew on IMSLP's jurisdiction-specific copyright notes, which incorporate factors such as renewal filings and foreign publication. Third, we used GPT-5, a large AI reasoning model with web search, to retrieve factual information and apply the rules governing US copyright terms. Classification results were broadly consistent across methods. Because Content ID hides composition claims when recording claims are present, we iteratively filtered and pitch-shifted each file to strip the recording fingerprint while preserving melodic structure. This allowed us to detect composition (Melody ID) claims separately. The final sample includes 631 videos which were uploaded to YouTube between February and April 2025.

5 Results

We begin by assessing the system's accuracy along two dimensions: matching accuracy (does the system correctly identify compositions when it flags content?) and enforcement accuracy (do the system's flagging decisions align with actual copyright status?). We then examine structural sources of error–whether divergence between automated enforcement and legal entitlements can be traced to the institutional design of the system, including access policies and data quality. Finally, we document evidence of fraudulent claiming, where entities appear to systematically assert rights over works they do not control.

5.1 Baseline Results

To assess Content ID's performance, we examine both matching accuracy—whether the system correctly identifies compositions when it flags content—and enforcement accuracy—whether the system's flagging

[16] As explained in Section 2.1, copyright status in the US depends on various factors, including publication dates, renewal filings, foreign publications etc.

decisions align with actual copyright status. Table 3 presents results using the *Contemporary* and *Classical Old* samples.

Matching Accuracy In the *Contemporary* sample, we only received Audio ID claims (recording rights). By definition, a correct Audio ID claim refers to the correct underlying work (consisting of recording and composition rights). In the *Classical Old* sample, we sometimes received multiple Audio ID or Melody ID (composition rights) claims per video. Recordings of classical works are sometimes split into several tracks (e.g., movements of a symphony), and a video containing a recording of the entire work can therefore receive multiple claims relating to different tracks. Here we found that 96% of the claims matched the correct work. This indicates the matching algorithm itself performs well: when Content ID has a reference file in its database and makes a claim, it nearly always correctly identifies the work. High technical accuracy is consistent with YouTube's own reports (Weinstein and Moreno, 2007). However, all the claims in the *Classical Old* sample matched an incorrect recording or composer. The algorithm, therefore, seems to struggle with variations across recordings, over-claiming recording rights in favor of references in its database.

Enforcement Accuracy The results in Table 3 reveal substantial enforcement errors in both directions. Among the 4,000 copyrighted recordings from Spotify (*Contemporary* sample), 1,814 (45.4%) received no copyright claim despite rights holders having legal entitlements to enforce. This false-negative rate indicates significant under-enforcement of protected content. In contrast, 360 of the 793 unenforceable classical recordings (45.4%) received copyright claims (*Classical Old* sample). This false-positive rate demonstrates systematic over-enforcement of content that should remain freely available. Given the high matching precision documented above, these enforcement failures seem to stem not from technical inability to detect content, but from other sources: gaps in database coverage for under-enforcement, and unverified rights data for over-enforcement.

Enforcement Policy Table 5 reports how rights holders chose to enforce their claims. Across all three samples, monetization was the dominant strategy. 94.9% of claims in the *Contemporary* sample, 96.4% in *Classical Old*, and 100% in *Classical US-EU* resulted in the rights holder running ads on the video rather than blocking it. Blocking was rare, occurring in only 5.0% of *Contemporary* claims and 2.6% of

Table 3: Content ID Matching and Enforcement Accuracy

Copyright status	Claimed	Not claimed	Total
Protected	2,186 (54.6%)	1814 (45.4%)	4,000
(Contemporary	*2,186 Audio ID claims*		
sample)	*0 Melody ID claims*		
	Correct Work: 2,147 (98.2%)		
	Incorrect Work: 39 (1.8%)		
	Correct Audio ID: 2,147 (98.2%)		
	Incorrect Audio ID: 39 (1.8%)		
Not protected	360 (45.4%)	433 videos (54.6%)	793
(Classical Old	*307 Audio ID claims*		
sample)	*123 Melody ID claims*		
	Correct Work: 413 (96.0%)		
	Incorrect Work: 17 (4.0%)		
	Incorrect Audio ID: 307 (100.0%)		
	Incorrect Melody ID: 123 (100.0%)		
Total videos	2,546	2,247	4,793

Note: Columns report estimated enforcement accuracy without controlling for observable differences across works and rightsholders. In the *Contemporary* sample, we only received Audio ID claims (recording rights). By definition, a correct Audio ID claim refers to the correct underlying work (consisting of recording and composition rights). In the *Classical Old* sample, we sometimes received multiple Audio ID or Melody ID (composition rights) claims per video. Recordings of classical works are sometimes split into several tracks (e.g., movements of a symphony), and a video containing a recording of the entire work can therefore receive multiple claims relating to different tracks. While 96% of the claims matched the correct work, all of them matched an incorrect recording or an incorrect composer.

Table 4: Content ID Enforcement Accuracy, region-specific

Copyright status	Claimed in US	Not claimed in US	Total
Protected in US *(Classical US-EU sample)*	235 (62.7%)	140 (37.3%)	375
Not protected in US *(Classical US-EU sample)*	133 (52.0%)	123 (48.0%)	256
Total claims	368	263	631

Note: Columns report estimated enforcement accuracy without controlling for observable differences across works and rightsholders. The audit included 631 compositions. By construction, we only received Melody ID (composition rights) claims. Recordings of classical works are sometimes split into several tracks (e.g., movements of a symphony), and a video containing a recording of the entire work can therefore receive multiple claims relating to different tracks. Copyright status is determined using GPT-5 with web search.

Table 5: Rights Holder Enforcement Policy by Sample

Policy	Contemporary	Classical Old	Classical US-EU
Monetize	2,074 (94.9%)	403 (96.4%)	2,874 (100.0%)
Block	109 (5.0%)	11 (2.6%)	0 (0.0%)
Track only	3 (0.1%)	4 (1.0%)	0 (0.0%)
Total	2,186	418	2,874

Note: This table reports the enforcement policy chosen by rights holders for claims in each sample. "Monetize" means the rights holder chose to run ads on the video and collect revenue. "Block" means the video was blocked in some or all territories. "Track only" means the rights holder monitors views without monetizing or blocking.

Classical Old claims. This pattern suggests that Content ID functions primarily as a revenue-extraction mechanism rather than a takedown tool.

5.2 Does automated enforcement affect all rights holders similarly?

5.2.1 Structural Sources of Errors

To examine whether under-enforcement affects all rights holders equally, we focus on rights in sound recordings and the distinction between major record labels (Universal, Sony, Warner and their subsidiaries) and independent labels. Major labels are known to participate directly in Content ID, while independent

Table 6: Enforcement by Type of Rights Holder

	(1)	(2)	(3)	(4)
Published by a major label	0.384***	0.371***		
	(0.033)	(0.027)		
Published by a large label			0.213***	0.189***
			(0.036)	(0.017)
(Constant)	0.512***	0.573***	0.471***	0.520***
	(0.035)	(0.025)	(0.038)	(0.025)
Num.Obs.	4000	3986	4000	3986
R2 Adj.	0.048	0.076	0.042	0.063
R2 Within		0.045		0.031
R2 Within Adj.		0.044		0.031
FE: Publication country		X		X

Note: Linear probability model. Dependent variable: 1 if a recording from the Contemporary sample was claimed by Content ID. Standard errors clustered by country. Publication-country fixed effects are included in Columns (2) and (4); the constant shown is the mean of the estimated fixed effects. A label is classified as large if it is among the top 1% of labels in the ENAO dataset in terms of number of tracks published.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

labels' participation is uncertain. In particular, we regress rights holder characteristics on an indicator of whether a video containing music from the *Contemporary* sample received a claim in a linear probability model.

As seen in column (1) of Table 6, recordings owned by non-major labels only receives a Content ID claim in 51.2 percent of the cases (coefficient of the constant). Recordings owned by major labels is significantly more likely to receive a Content ID claim (38.4 percentage points). Hence, we estimate that the overall likelihood to receive a lawful Content ID claim for recordings represented by major labels is 89.6%.

Another way to study the structural sources of enforcement errors is to examine the distribution of rights holders associated with Content ID claims. Table 7 shows that the major record labels are the rights holders that are behind most claims against works in the *Contemporary* sample. The big three, *Universal Music Group (UMG)*, *Sony Music Entertainment (SME + The Orchard)* and *Warner Music*

Table 7: Distribution of Claims by Rights Holder in the Contemporary Sample

No.	Content owner name	Number of claims	Percentage
1	**UMG**	**436**	**13.5%**
2	The Orchard Music	388	12.0%
3	**SME**	**282**	**8.8%**
4	**WMG**	**264**	**8.2%**
5	Believe Music	246	7.6%
6	Naxos Digital Services US, Inc	135	4.2%
7	CD Baby CO	118	3.7%
	...		
303	MODULORTV	1	0.0%

Note: This table lists the rights holders that were behind claims against recordings of the *Contemporary sample* sample when uploaded to YouTube. When including *The Orchard*, which is owned by SME, the big three record labels are behind roughly 43% of the claims, which is disproportional to the composition of the sample.

Group (WMG), together account for 42.5% of claims, even though they only represent around 8% of the tracks in the sample. Among other things, this is caused by smaller labels making distribution and licensing agreements with the major labels, likely driven by the fact that the latter have a sophisticated infrastructure that allows them to easily enforce music via Content ID.

These findings provide empirical support for concerns raised during the enactment of the European Union's Copyright Directive that automated enforcement systems might systematically favor larger rights holders (Quintais et al., 2022; Grosse Ruse-Khan, 2020; Husovec and Quintais, 2021). Our results indicate that under-enforcement stems primarily from institutional access barriers rather than technical detection failures. Content ID's restriction to rights holders with "complex copyright management needs" effectively excludes independent labels, leaving their content unprotected even when the technical capacity for detection exists.

5.2.2 Enforcement of International Copyright Law

The enforcement accuracy of automated copyright enforcement systems is complicated by the fact that the copyright status of works may differ between different countries. Using the *Classical US-EU* sample, we now examine whether Content ID can accurately distinguish between copyrighted and public domain

Table 8: Enforcement Accuracy Within Composers

	Year		IMSLP		GPT-5	
	(1)	(2)	(3)	(4)	(5)	(6)
Copyrighted in US	0.152**	0.064	0.119**	0.034	0.107*	0.029
	(0.060)	(0.045)	(0.052)	(0.043)	(0.056)	(0.042)
(Constant)	0.496***	0.501***	0.519***	0.529***	0.520***	0.523***
	(0.052)	(0.041)	(0.045)	(0.041)	(0.048)	(0.042)
Num.Obs.	561	560	598	597	631	630
R2 Adj.	0.021	0.103	0.013	0.098	0.010	0.109
R2 Within Adj.		0.002		-0.001		-0.001
FE: Composer		X		X		X

Note: Linear probability model. Dependent variable: 1 if a recording was claimed in the US. Standard errors clustered by composer. Columns (2), (4), and (6) include composer fixed effects; the constant shown is the mean of the estimated fixed effects. The copyright status is determined using publication year in columns 1-2, IMSLP copyright notices in columns 3-4, and GPT-5 with web search in columns 5-6.

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

compositions within the same composer's catalog. The sample contains compositions with varying US copyright status but uniform EU protection. YouTube Content ID reports copyright flags on a per country basis, allowing us to test whether its enforcement patterns respect US copyright boundaries.

Table 8 presents results where the dependent variable equals one if the composition received a copyright claim in the United States.

The baseline specifications (Columns 1, 3, 5) reveal substantial over-enforcement. Roughly half of public domain compositions receive copyright claims despite having no legal protection. Copyrighted compositions show modestly higher enforcement rates, but the difference is economically small relative to the baseline over-enforcement problem and reaches statistical significance only in some specifications.

After the addition of fixed effects (Columns 2, 4, 6), which compare enforcement within the same composer across the US and the EU, the within-composer copyright effect becomes small and statistically insignificant. This indicates that Content ID does not meaningfully distinguish between a composer's copyrighted and public domain works in the United States: rights holders appear to provide information at the composer level rather than tracking the copyright status of individual compositions.

Table 9: Distribution of Claims by Rights Holder in the Classical Old Sample

No.	Content owner name	Number of claims	Percentage
1	LatinAutorPerf	161	22.1%
2	Wise Music Group	61	8.4%
3	UMPG Publishing	61	8.4%
4	LatinAutor	50	6.9%
5	Hexacorp (music publishing)	36	5.0%
6	UNIAO BRASILEIRA DE EDITORAS DE MUSICA - UBEM	31	4.3%
7	Naxos Digital Services US, Inc	27	3.7%
	...		
12	**UMG**	**16**	**2.2%**
	...		
27	**SME**	**6**	**0.8%**
	...		
30	**WMG**	**4**	**0.6%**
	...		
47	The Orchard Music	1	0.1%

Note: This table lists the rights holders that were behind claims against recordings of the *Classical Old* sample when uploaded to YouTube. Enforcement is heavily concentrated in a single entity, "LatinAutorPerf" (22.1% of claims), whereas major record labels such as UMG, SME, and WMG account for a comparatively marginal share of total claims.

These patterns suggest that over-enforcement is driven partly by insufficient data granularity: rights holders provide composer-level information rather than tracking individual composition status across jurisdictions.

5.2.3 Fraudulent Claiming

Having established that restricted access to Content ID can explain under-enforcement, we now study over-enforcement in more detail. We find evidence of fraudulent claiming on a meaningful scale.

The perhaps best setup to understand fraud is the sample of works which do not fall under copyright protection (*Classical Old* sample). Among the claims in this sample, Table 9 shows that a single entity accounts for a disproportionate share of 22.1% of claims. "*LatinAutorPerf*" makes more than twice as many claims as the next most frequent claimant. This entity, which also appears to be affiliated with "*LatinAutor*" and "*UBEM (Uniao Brasileira de Editoras de Musica)*", systematically asserts rights over

compositions it does not own or represent, coming to a combined share of 33.3% of claims in the *Classical Old* sample. Multiple online complaints from creators document LatinAutorPerf's baseless claims across diverse content, including public domain classical music, original compositions, and Creative Commons-licensed works.[17] In our sample, the claims made by LatinAutorPerf and affiliated claimants are based on incorrect information about composers. For example, a Haydn composition is labeled as composed by Bartok. The entities behind LatinAutorPerf and related accounts appear to have access to Content ID and can enter what YouTube calls "composition shares," pretending to represent the rights of (typically deceased) composers. This pattern is consistent with copyfraud — the false assertion of copyright over public domain or otherwise uncontrolled works (Mazziotti, 2010). That copyfraud occurs at a meaningful scale despite YouTube's restricted access policy suggests that verification mechanisms are insufficient to deter abuse, and that simply broadening access to address under-enforcement could exacerbate fraudulent claiming.

Legitimate rights holders, such as the three major record labels UMG, SME and WMG, however, also appear in the list of claimants in cases where the work is definitely not protected by copyright. In our sample, it appears that this is mainly due to mismatched recordings, i.e., Content ID recognizing the correct work but referring the claim to a recording that is represented by a record label.

6 Discussion

Our empirical analysis shows a striking contrast between the technical capabilities of YouTube's fingerprinting technology and the actual enforcement outcomes that Content ID produces. When the system has reference files in its database, it identifies the correct work with high precision. At the same time, about half of infringements affecting smaller rights holders remain unaddressed, and about half of the public domain works in our classical sample are nonetheless claimed. Enforcement is strongest for major labels with direct access to Content ID, while independent labels and individual creators experience both under-enforcement of their own rights and over-enforcement against their works.

These patterns suggest that the main constraints are not technological but institutional. The platform decides who can access Content ID, which works can be registered, and how much effort it invests in

[17] See Reddit discussions at https://archive.md/vFjgu, https://archive.md/6J9kn, https://archive.md/nzdi6, https://archive.md/6haZi.

checking whether information in the reference database is accurate and up to date. In this section, we present a simple model that formalizes these choices. The model clarifies how a platform that maximizes its own payoff will typically restrict access to automated enforcement and underinvest in data quality, even if this leads to both over- and under-enforcement from a legal perspective. It also helps interpret our empirical findings as the outcome of platform incentives rather than as random error, leading to policy implications.

6.1 A Simple Model of Platform Incentives

The model is intentionally stylized. It abstracts from many details of copyright law and Content ID implementation, and focuses on three ingredients that map directly to our empirical setting. First, a platform that earns advertising revenue from user-generated content. Second, rights holders who can register their works in a database, subject to an access policy. Third, a tradeoff between broader access and stronger verification of claims on the one hand, and the platform's own costs and revenue on the other. The key question is whether a platform that sets access and verification levels to maximize its own payoff will also choose levels that are optimal for overall welfare. We show that the answer is no. The platform will tend to admit mainly high-value rights holders and will choose verification effort below the social optimum. The result is a pattern of asymmetric enforcement and systematic errors that matches what we observe in the data.

6.1.1 Setup

Consider a digital platform P that hosts user-generated content (uploaders U). Content may incorporate works from rights holders (RH). There is a continuum of rights holders $i \in [0, 1]$ with heterogeneous view potential v_i drawn from distribution $F(v)$ on $[0, \bar{v}]$. A fraction $\pi \in (0, 1)$ of works are copyrighted; the remainder are public domain (PD).

Algorithmic matching. The platform operates an automated matching system with perfect recognition technology. Conditional on a correct reference in the database, matches are accurate. Enforcement errors, however, can arise from the *rights database*, leading to the following problems. First, **access policy**, capturing YouTube's decision to grant Content ID only to rights holders with "complex copyright management needs," such as major labels and large intermediaries. The platform admits rights holders

with $v_i \geq \theta$ to the database. Coverage is $c(\theta) = \Pr(v \geq \theta)$. A higher threshold θ restricts participation to larger rights holders. Second, **incorrect data** can exist because entries may be outdated with probability s (for instance, works that have expired into the public domain but are not removed or updated). It can also exist by entities asserting rights over works they do not control, as in the LatinAutorPerf pattern we document for classical public domain recordings. A fraction $\phi(x)$ of admitted rights holders submit erroneous references ("copyfraud"). Verification effort $x \geq 0$ reduces this fraction, so $\phi'(x) < 0$, at cost $K(x)$, with $K' > 0$ and $K'' > 0$.

Error rates. Under-enforcement of copyrighted works can written as $\psi(\theta) = 1 - c(\theta)$, which measures the share of copyrighted works that are not covered by the database because their rights holders fall below the access threshold.

Over-enforcement of PD works is $\omega(\theta, x) = s \cdot c(\theta) + \phi(x) \cdot m(\theta)$, where $m(\theta)$ is the mass of marginal entrants at threshold θ. The first term captures claims based on stale entries that should have been removed after copyright expiration. The second term captures opportunistic or fraudulent claims that survive because verification is imperfect.

Monetization. Each view generates ad revenue r. If a copyrighted work is correctly matched, the platform and rights holder split revenue: share $\alpha \in (0, 1)$ to the rights holder, $1 - \alpha$ to the platform. If no match is made, the uploader keeps share γr and the platform keeps $(1 - \gamma)r$. Erroneous PD claims transfer αr from uploader to claimant and may additionally impose a removal cost L with probability η, reducing viewer surplus when content is blocked or made unavailable.

The platform chooses the access threshold θ and verification effort x before enforcement takes place. These choices determine both the coverage of valid rights and the incidence of erroneous claims. They also determine how advertising revenue is split between platform, rights holders, and uploaders.

6.1.2 Payoffs

We now describe the payoffs for rights holders, uploaders and users, and the platform.

Rights holders with access. Rights holders who satisfy $v_i \geq \theta$ can register their works in the database and receive revenue when their works are correctly matched and enforced. Their expected payoff is

$$\Pi^A(\theta) = \alpha r \, \pi \, (1 - s) \, \mathbb{E}[v \mid v \geq \theta].$$

This expression reflects that only copyrighted works (fraction π) with non-stale entries $(1 - s)$ generate enforcement, and that the magnitude of the payoff scales with their view potential.

Rights holders without access. Rights holders below the threshold cannot register references and therefore do not benefit from automated enforcement. Their payoff is approximately zero: $\Pi^N(\theta) \approx 0$. This captures the under-enforcement experienced by independent rights holders in our empirical analysis.

Uploaders and users. Uploaders and viewers benefit when content remains online and when revenue is not incorrectly diverted to rights holders who do not own the work. Their surplus is

$$\begin{aligned} S_U(\theta, x) = & \underbrace{\gamma r \, \pi \, \psi(\theta) \, \mathbb{E}[v]}_{\text{uploader revenue on missed claims}} + \underbrace{CS \cdot (1 - \eta \omega(\theta, x))}_{\text{viewer surplus net of removals}} \\ & - \underbrace{\alpha r \, (1 - \pi) \, \omega(\theta, x) \, \mathbb{E}[v]}_{\text{erroneous transfers on PD works}} . \end{aligned}$$

The first term captures that under-enforcement increases uploader revenue, since some copyrighted content is never claimed. The second term reflects that erroneous claims can lead to removals, reducing consumer surplus CS. The third term captures the revenue that is wrongly diverted from uploaders to claimants when public domain content is claimed.

Platform. The platform's payoff has three components: revenue from correctly enforced copyrights, revenue from content that remains unmatched, and the costs of verification and broader access:

$$\begin{aligned} \Pi_P(\theta, x) = & (1 - \alpha) r \, \pi \, c(\theta) (1 - s) \, \mathbb{E}[v] \\ & + (1 - \gamma) r \, \mathbb{E}[v] \, [\pi \psi(\theta) + (1 - \pi)(1 - \omega(\theta, x))] \\ & - K(x) - C(\theta), \end{aligned}$$

where $C(\theta)$ captures administrative and technical costs that increase with coverage (for example, processing more reference files and resolving more disputes).

The first term is the platform's revenue share from correctly enforced copyrighted works. The second term is its revenue share from content that is not subject to claims, either because copyrighted works are missed or because public domain works are not erroneously claimed. The last term captures the costs of verification and maintaining a broader database.

Social welfare. Total welfare adds the platform's payoff, the payoffs of rights holders with and without access, and the surplus of uploaders and viewers:

$$W(\theta, x) = \Pi_P(\theta, x) + \int_{\mathcal{A}(\theta)} \Pi^A \, di + \int_{\mathcal{A}^c(\theta)} \Pi^N \, di + S_U(\theta, x),$$

where $\mathcal{A}(\theta)$ denotes the set of rights holders above the threshold and $\mathcal{A}^c(\theta)$ those below. A social planner would choose θ and x to maximize $W(\theta, x)$, whereas the platform chooses them to maximize $\Pi_P(\theta, x)$.

6.1.3 Platform Choices and Comparative Statics

We can now compare the platform's choices to the social optimum.

Proposition 1 (Platform vs. Planner Incentives)**.** *Relative to the social planner, the platform chooses:*

1. *A higher access threshold ($\theta_P > \theta^*$), restricting coverage of small rights holders.*
2. *Lower verification effort ($x_P < x^*$).*

Sketch of proof. The planner internalizes three elements that the platform only partly takes into account: (i) the income of excluded rights holders, (ii) the welfare cost of erroneous claims and removals for uploaders and viewers, and (iii) the value of broad and accurate enforcement for the long run functioning of the copyright system. Formally,

$$\frac{\partial W}{\partial x} = \text{large positive (rights holder and user gains)} - K'(x),$$

while

$$\frac{\partial \Pi_P}{\partial x} = \text{small positive (revenue effects only)} - K'(x).$$

Because most of the benefits of higher verification accrue to parties other than the platform, the planner chooses $x^* > x_P$. A similar logic applies to the access threshold. The planner internalizes the gains from protecting marginal rights holders when θ falls, whereas the platform focuses mainly on how broader coverage affects its own revenue net of $C(\theta)$. As a result, $\partial W/\partial\theta$ becomes negative at a lower value of θ than $\partial\Pi_P/\partial\theta$, which implies $\theta^* < \theta_P$. □

Corollary 1 (Incumbent Advantage)**.** *When access depends on expected views v, large rights holders (high v) lie above θ_P, while independent rights holders often lie below. This generates asymmetric enforcement across incumbents and smaller players, consistent with our finding that works represented by major labels are claimed in roughly 90 percent of cases, compared to about 50 percent for non-major labels.*

Corollary 2 (Policy Levers)**.** *Welfare improves when institutions reduce staleness s (for instance through mandatory expiry checks or re-attestation of rights), lower the access threshold θ (for instance by widening eligibility for reference registration), or raise verification effort x (for instance by setting minimum review or dispute-handling standards). These interventions reduce both under-enforcement and over-enforcement without undermining the underlying matching accuracy of the system.*

6.2 Model and Empirical Findings

The model provides a compact way to interpret our empirical results. First, the high matching accuracy we document corresponds to the assumption that, conditional on correct reference data, the fingerprinting technology works well. This is why our model locates enforcement errors in the access threshold θ and in data quality, not in the matching technology itself. Second, the pronounced difference in enforcement between major and non-major labels corresponds to a high threshold θ_P that excludes many independent rights holders from the database. This generates a high under-enforcement rate for smaller entities and aligns with our estimates for $P(\text{flag} \mid V = 1,\ \text{type})$. Third, the large fraction of public domain works that receive claims corresponds to large $\omega(\theta_P, x_P)$, driven by staleness s and opportunistic claiming $\phi(x_P)$. Our evidence on copyfraud in the *Classical Old* sample fits this mechanism particularly well. Finally, the difficulty of distinguishing between compositions that are still copyrighted in the EU but in the public domain in the US reflects low investment in granular, jurisdiction-specific data and verification, again

consistent with the model's predictions for $x_P < x^*$ and for incomplete or outdated information in the database.

Taken together, the model highlights that even if algorithmic enforcement systems are technologically sophisticated, they implement legal rules in ways that systematically favor large incumbents and tolerate substantial error. The platform's private incentives lead to an access policy and a data quality level that differ from what a social planner would choose. This helps explain why, in practice, automated copyright enforcement on YouTube does not deliver the disintermediated, level playing field that many policymakers and commentators expected.

6.3 Policy Implications

A welfare-enhancing reform of Content ID would require shifting the systems institutional design from strict gatekeeping and weak data oversight to broader participation and higher data quality. The platform should allow a much wider range of rights holders to register reference files directly, rather than tying access to organizational size or historical enforcement activity. Doing so would substantially reduce under-enforcement by giving smaller creators and independent labels the same protection majors already enjoy. At the same time, YouTube should introduce systematic checks on the accuracy and freshness of registered information, including periodic re-attestation of rights, clear rules for removing works that have entered the public domain, systematic checks on country-specific copyright rules, and stricter review of entities that repeatedly issue erroneous claims. Stronger verification would curb over-enforcement and copyfraud without compromising technical matching accuracy. Combined, these changes would move the system closer to the social optimum: creators of all sizes would receive more accurate protection, uploaders and viewers would face fewer erroneous restrictions, and the enforcement tool would better reflect the legal entitlements it is designed to implement.

One can ground these recommendations quite directly in the Digital Services Act of the European Union. For very large online platforms such as YouTube, Articles 34 and 35 already require regular assessments and mitigation of "systemic risks" arising from the design and functioning of their services, including risks to fundamental rights. A persistent pattern of under-enforcement for small rights holders and over-enforcement against public domain works could be interpreted as a structural distortion that should be captured in these risk assessments, with mitigation measures that include broader, non-

discriminatory access to enforcement tools and higher data quality standards in rights databases. Article 17 on internal complaint handling then supports stronger, non-automated review of repeated wrongful claims and better remedies for uploaders whose lawful content is misclassified. In combination, these provisions already point to reforms that could improve welfare in the context of automated copyright enforcement.

7 Conclusion

In this paper, we present the first large-scale algorithmic audit of the largest automated copyright enforcement system on the globe: YouTube Content ID. We find that while Content ID's matching technology is generally accurate, its enforcement accuracy is lower: works represented by major labels are claimed in about 90 percent of cases, compared to roughly 50 percent for non-major labels. And Content ID is often unable to accurately enforce copyrights on works which are protected in some countries but not in others. We relate these findings to particular design choices of YouTube Content ID: access to the system is reserved to large rights holders, and automated enforcement relies on data provided by rights holders, which does not always reflect actual legal entitlements. We explain how platform incentives can lead to such access requirements and data quality policies.

These findings have important implications for platform governance more broadly. As platforms increasingly function as private regulators implementing legally mandated enforcement systems, the institutional design of these systems determines how laws are applied in practice. When access to enforcement tools is restricted based on organizational size or resources, automated systems can inadvertently lead to favor incumbents and larger participants, even when the underlying technology performs accurately. In YouTube's case, this creates a two-tier system where major rights holders receive robust protection while independent rights holders face both under-enforcement of their own rights and over-enforcement against their works. This suggests that, so far, automated enforcement systems have not enabled a truly disintermediated market for creators. And policy intervention may be the only mechanism capable of enabling automated enforcement systems to realize their full potential for all types of rights holders.

References

Aguiar, L., Waldfogel, J., and Zeijen, A. (2024). Platform power struggle: Spotify and the major record labels.

Bartholomew, T. B. (2014). The death of fair use in cyberspace: YouTube and the problem with Content ID. *Duke L. & Tech. Rev.*, 13:66.

Berkowitz, A. E. (2023). Algorithmic (in) tolerance: Experimenting with beethoven's music on social media platforms. *Transactions of the International Society for Music Information Retrieval*, 6(1).

Erickson, K. and Kretschmer, M. (2018). This video is unavailable. *J. Intell. Prop. Info. Tech. & Elec. Com. L.*, 9:75.

Even, A. M. (2023). Keeping the good faith: YouTube, fair use, and the DMCA. In *Boston College Intellectual Property and Technology Forum*, volume 2023, pages 1–25.

Gong, Y., Chung, Y.-A., and Glass, J. (2021). Ast: Audio spectrogram transformer. *arXiv preprint arXiv:2104.01778*.

Gray, J. E. and Suzor, N. P. (2020). Playing with machines: Using machine learning to understand automated copyright enforcement at scale. *Big Data & Society*, 7(1):205395172091996.

Grosse Ruse-Khan, H. (2020). Automated copyright enforcement online: From blocking to monetization of user-generated content.

Helft, M. and Richtel, M. (2006). Venture Firm Shares a YouTube Jackpot. *The New York Times*.

Husovec, M. and Quintais, J. P. (2021). Too small to matter? On the copyright directive's bias in favour of big right-holders. *Global Intellectual Property Protection and New Constitutionalism. Hedging Exclusive Rights, Tuomas Mylly and Jonathan Griffiths (Eds.), Oxford University Press (2021)*.

King, D., Salem, G., Wang, Y., and Wiseman, M. (2014). Media Rights Management Using Melody Identification.

Lester, T. and Pachamanova, D. (2017). The dilemma of false positives: Making content id algorithms more conducive to fostering innovative fair use in media creation. 24.

Mazziotti, G. (2010). New licensing models for online music services in the European Union: From collective to customized management. *Colum. JL & Arts*, 34:757.

McGrady, R., Zheng, K., Curran, R., Baumgartner, J., and Zuckerman, E. (2023). Dialing for Videos: A Random Sample of YouTube. *Journal of Quantitative Description: Digital Media*, 3.

Quintais, J. P., Mezei, P., Harkai, I., C Magalhães, J., Katzenbach, C., Schwemer, S. F., and Riis, T. (2022). Copyright content moderation in the EU: An interdisciplinary mapping analysis. *Available at SSRN 4210278*.

Simon, A. R. (2014). Contracting in the dark: Casting light on the shadows of second level agreements. *Wm. & Mary Bus. L. Rev.*, 5:305.

Soha, M. and McDowell, Z. J. (2016). Monetizing a meme: YouTube, content ID, and the harlem shake. *Social Media+ Society*, 2(1):2056305115623801.

Solomon, L. (2015). Fair users or content abusers: The automatic flagging of non-infringing videos by content id on youtube. *Hofstra L. Rev.*, 44:237.

Tang, X. (2023). Privatizing Copyright. *Michigan Law Review*, (121.5):753.

Tehranian, J. (2011). *Infringement Nation: Copyright 2.0 and You*. OUP USA.

Walters, T. C., Ross, D. A., and Lyon, R. F. (2012). The intervalgram: An audio feature for large-scale cover-song recognition. In *International Symposium on Computer Music Modeling and Retrieval*, pages 197–213. Springer.

Weinstein, E. and Moreno, P. (2007). Music Identification with Weighted Finite-State Transducers. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing - ICASSP '07*, pages II–689–II–692, Honolulu, HI. IEEE.

Zhou, J., Li, Y., Adhikari, V. K., and Zhang, Z.-L. (2011). Counting YouTube videos via random prefix sampling. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, IMC '11, pages 371–380, New York, NY, USA. Association for Computing Machinery.

A Appendix A: Data Sources

In this Appendix, we further describe the data collection process and the sources that we drew from.

To capture comprehensive enforcement data, we developed a custom monitoring system using a local proxy to intercept structured JSON data transmitted from YouTube's backend to the user interface. This approach provided two key advantages. First, it allowed us to collect detailed, machine-readable information about each claim, including claimant identity, specific rights asserted (composition vs. recording), and geographical scope of enforcement. Second, it enabled us to confirm that Content ID had processed every uploaded video, as the system reports processing status even when no claims are made. This verification was essential to distinguish between videos that received no claims because they were not detected versus videos that were analyzed but deemed non-infringing.

A.1 Contemporary Sample

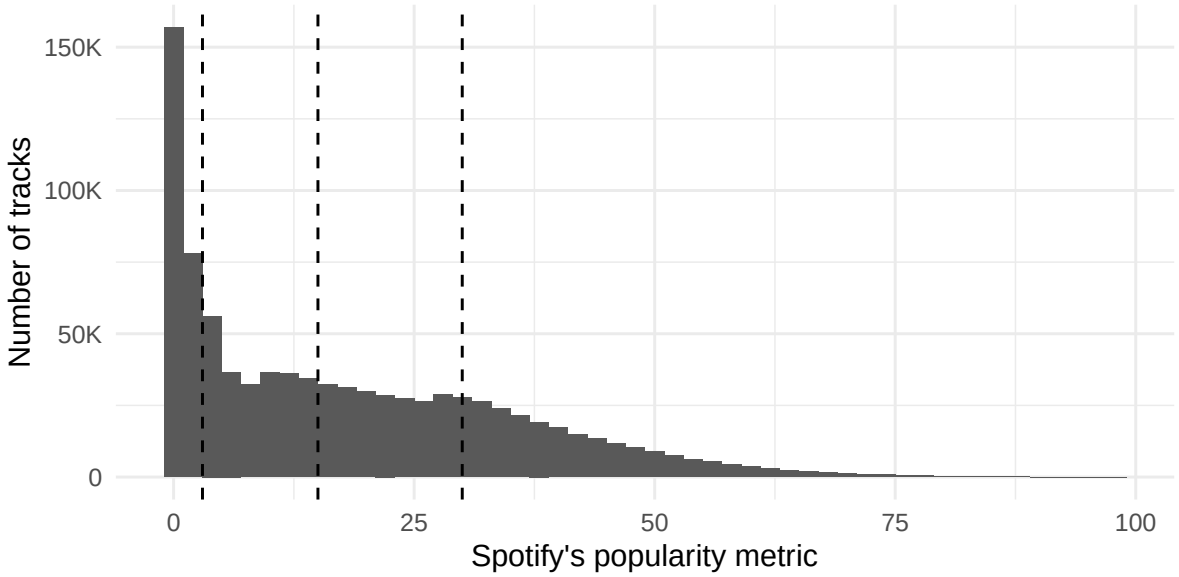


Figure A.1: A histogram showing the distribution of Spotifys popularity metric for all tracks listed on Every Noise at Once. Vertical lines are quartile marks.

We used Spotify as a source of copyrighted music. Aside from being the most popular streaming service on the global market, it also offered an API which offers detailed information about tracks. This included popularity metrics, publication year and location, along with copyright information, which indicated which label published the track, or distributed it.

Table A.1: Preview statistics by type of label before and after imputing previews.

		No. missing previews		% missing previews	
Type of label	No. tracks	Before	After	Before	After
Major	67,659	27,669	3,328	40.89%	4.92%
Other	839,814	25,812	6,488	3.07%	0.77%

A.1.1 Major label classification

Based on the information provided from Spotify, we classified each track as either being published or distributed by a major label. In line with Aguiar et al. (2024), we estimated this by evaluating whether the copyright information contained strings identifying one of the three major record labels (Warner, Sony, Universal) as well as their subsidiaries identified using Wikipedia. We further validated this using Hannes Datta's `musicMetadata` package[18] and by manually validating all positive matches, ensuring a minimal rate of false-positives.

A.1.2 Audio previews

At the time of data collection, Spotify offered 30 second previews of most songs through their API. However, we noticed that a large share of the tracks classified as ones published by major labels were missing audio previews as seen in Section A.1. To go around this, we used Spotify's audio preview widgets are intended to be embedded into HTML web pages. These widgets also contain 30 second previews, and are available for the vast majority of tracks on the platform, decreasing the number of missing tracks from 41% to 5% for major label tracks. This alleviated concerns about selection bias based on audio preview availability

A.1.3 Selecting samples from Every Noise at Once

To obtain a comprehensive and representative sample of music from Spotify, we utilized the playlists generated by Every Noise at Once (ENAO).[19] Attempting to create a sample through other means would have been infeasible, primarily because Spotify does not offer an index of the songs available on the platform. Furthermore, the official playlists available through Spotify are often curated for commercial

[18] https://github.com/hannesdatta/musicMetadata
[19] https://everynoise.com

purposes or personalized for individual users, meaning they do not necessarily reflect the full diversity of music available on the platform.

ENAO solves this issue by offering a playlist for every single genre of music available on Spotify. When the data was collected, every artist on Spotify was assigned to at least one of 6,201 genres based on a combination of listener data and other properties. For each of the genres, ENAO offered a playlist composed of tracks that were commonly listened to by typical listeners of that genre.

A.1.4 Fully Protected Music from Spotify

By combining all of these playlists, we could therefore get a representative sample of the tracks that are available on Spotify, consisting of 907,473 tracks. To reduce the sample to a size that could be easily uploaded to YouTube, we created a smaller sample of 4,000 tracks. We stratified across quartiles of Spotify's popularity metric (see Figure A.1), as well as record label type (major vs independent).

We also added classical music, analogous to the type of music covered in the Classical I sample, by selecting four of the playlists available on ENAO: Classical, Chamber ensemble, Classical performance, Classical piano. We chose these playlists because they aligned with the types of music included in the CC classical recordings, which mostly consisted of chamber and piano pieces.

A.2 Classical Old

Sources included the Isabella Stewart Gardner Museum's concert series and collections by pianists Bernd Krueger and Kimiko Ishizaka.

A.2.1 The Isabella Gardner Museum music collection

The primary source of CC recordings of public domain classical recordings was the recordings that the Isabella Gardner Museum have released under a CC license.[20] In total, we used 381 recordings, mostly consisting of chamber music (trios, quartets, sonatas) as well as solo piano pieces. These represent pieces that are regularly performed in chamber concerts around the world. These recordings are often of performances by world-leading musicians like Colin Carr, Ray Chen and Charlie Albright. Below we list some representative recordings:

[20] https://www.gardnermuseum.org/experience/music

- *Johannes Brahms:* Sonata for Clarinet and Piano in F Minor Op. 120 No. 1. Performed by David Deveau and Richard Stoltzman
- *Ludwig van Beethoven:* Piano Sonata No. 14 in C-Sharp Minor, Op. 27, No. 2 ("Moonlight"). Performed by Gleb Ivanov.
- *Felix Mendelssohn:* Song Without Words, Op. 109. Performed by Colin Carr and Thomas Sauer.
- *Franz Joseph Haydn:* String Quartet in B-Flat Major, Op. 50, No. 1. Performed by Musicians from Marlboro.
- *Maurice Ravel:* La Valse. Performed by Jean-Frédéric Neuburger.

A.2.2 Bernd Krueger

Since 1996, Bernd Krueger has been notating public domain piano compositions in notation software, and making them available on his website.[21] He publishes these works in various formats ranging from raw Musical Instrument Digital Interface (MIDI) files along with synthesized versions in a MP3 format. Most importantly, he releases all his files under a CC-BY-SA license, meaning that any reproductions cannot be enforced on YouTube (see Section 2.2). As a result of being produced through music notation software, the interpretation is mechanical, but still enjoyable because he has adjusted the playback to take musical elements such as ritardando and rubato into consideration. Examples of these recordings include Mozart sonatas, piano works by Debussy, Liszt etudes and Mussorgsky's Pictures at an Exhibition.

A.2.3 Kimiko Ishizaka

Japanese pianist Kimiko Ishizaka has released professionally recorded interpretations of three Bach piano works on her Bandcamp page:[22] *Das Wohltemperierte Klavier*, *Die Kunst der Fuge*, and the Goldberg Variations. These works are released under the most permissive Creative Commons license, CC0.

A.2.4 Note on differing recording formats

The recordings from the Gardner Museum were often of entire pieces, whereas a recording was always a single section/movement/track in other sources. As this could affect the probability of a track being

[21] http://piano-midi.de
[22] https://kimikoishizaka.bandcamp.com

claimed (if it has more sections, there is a bigger pool of tracks that could match), we counted the number of different sections in each work from the Gardner Museum and used that number as a control for all regressions involving Unenforceable classical recordings. This was primarily done by finding the work on IMSLP and retrieving the number of sections from there. For pieces not available on IMSLP, we used comparable recordings on Spotify to check into how many sections they were split. A minority of recordings required manual listening as they included a selection of works from a bigger collection.

A.3 Classical US-EU

For the regional variation sample, our goal was to determine whether compositions themselves (independent of their sound recordings) receive copyright claims in the United States. However, Content ID presents a technical challenge: when both recording and composition rights are claimed, the interface displays only the recording claim. To work around this limitation, we employed an iterative approach to isolate claims on the composition. After uploading each recording, we monitored for claims. When a video received a recording claim, we applied audio transformations (low-pass filtering to remove high frequencies and pitch modulation to alter the spectral signature) and re-uploaded the modified version. These transformations degrade the recording's acoustic fingerprint while preserving the underlying melodic and harmonic structure, which YouTube uses to apply a claim on the compositional rights. We repeated the transformation process up to six times per recording until either (a) we received a composition claim, (b) we received no claim at all, or (c) we exhausted the iteration limit. This allowed us to separately measure enforcement of compositional rights, which was necessary for our analysis.

B Appendix B: Copyright Verification

Determining the precise copyright status of musical compositions is a complex task, particularly for the United States. The legal status of a work in the US can depend on a variety of factors beyond a simple creation date, including the year of first publication and subsequent renewals. To ensure the robustness of our analysis for the regional compliance sample, we employed three distinct methods to classify compositions. This approach allowed us to test whether our findings were sensitive to the classification criteria.

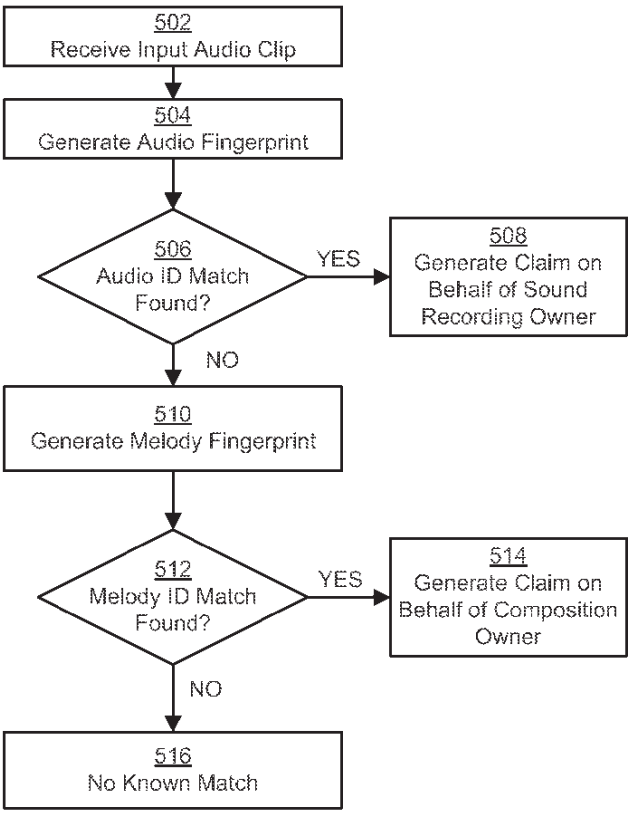


Figure A.2: Process description of ContentID system, from King et al. (2014), Figure 5.

B.1 Publication Year

The simplest way to determine when a composition was published, as the general rule for the compositions that are relevant for our study is that they are in copyright for 95 years after the first publication. Our source for this date was the IMSLP, more specifically the field "First Publication" which is available for 89% of the compositions in our sample. However, this date is not always correct, as this is not always based on the first manuscript published, but rather the oldest one that IMSLP hosts.

B.2 IMSLP Copyright Information

For compositions with a complicated copyright term across countries, IMSLP includes tailored information about where the copyright term is likely expired, and where it is not. This information is provided by IMSLP contributors, but takes more factors than only publication year into consideration.[23] This

[23] https://archive.md/oy9iY

information was available for 95% of the compositions in our sample. However, this might not generalize to a broader set of compositions, as these copyright summaries are most often only included when there is ambiguity around the copyright term.

B.3 GPT-5 + Web Search

As a more generalized alternative, we turned to OpenAI's GPT-5 reasoning model, and paired it with web search capabilities. With the search option, the model could both retrieve relevant information about the composition and reason about the copyright status. The cost of verifying the copyright status came to less than $0.50 per composition using OpenAI's API. We provided the following prompt:

```
You are an expert legal assistant that is supposed to determine the U.S. copyright status
of a provided musical composition as of February 2025.
Your analysis must be thorough and follow a two-step process:

Step 1: Factual Research
First, find and clearly state the following essential information. This data is critical
for the legal analysis.
* Composer's Full Name, Life Dates and Nationality
* Date of Creation
* Date and Place of First Publication
* Whether and when the composition was registered with the U.S. copyright office and
  whether and when a renewal notice was filed
* Whether a copyright notice in the proper form was included in the publication
* Whether or not the copyright has lapsed in the home country of the composers if they
  are not U.S citizens

Step 2: Legal Analysis and Conclusion
Based on the facts gathered in Step 1, provide a detailed explanation of the work's
copyright status in the United States. Your reasoning must address all relevant points
```

of U.S. copyright law, including:
* The copyright status of works first published before 1930.
* Rules for works published between 1930 and 1977, specifically addressing copyright notice and renewal requirements and place of publication.
* The copyright term for works created or published on or after January 1, 1978.
* The impact of the composer's nationality (French) and the potential for copyright restoration for foreign works under the Uruguay Round Agreements Act (URAA).
* Copyright notice
* Renewal notice
* Place of publication
* Composers nationality
* Potential copyright restoration for foreign works under the Uruguay Round Agreements Act (URAA)
* Also consider that rules on copyright term, renewal, and restoration are different for works published before 1930, before 1978, and after

Finally, provide a clear conclusion stating whether the work is in the public domain in the U.S. or is likely still protected by copyright.
Also, provide a confidence score (low, medium, high) for your prediction. If the work is still protected by copyright, provide the year when this protection expires.

Possible data sources

For information on copyright term, the following pages are good sources:
- Cornell Library guide on copyright term
  https://guides.library.cornell.edu/copyright/publicdomain
- IMSLP page on what is in the public domain
  https://imslp.org/wiki/Public_domain

B.4 Comparison of the Different Methods

Figure B.1 shows the patterns of enforcement across the different methods listed above. We can see that in most cases the three methods either agreed, or we could not determine the copyright status using either publication year or IMSLP.

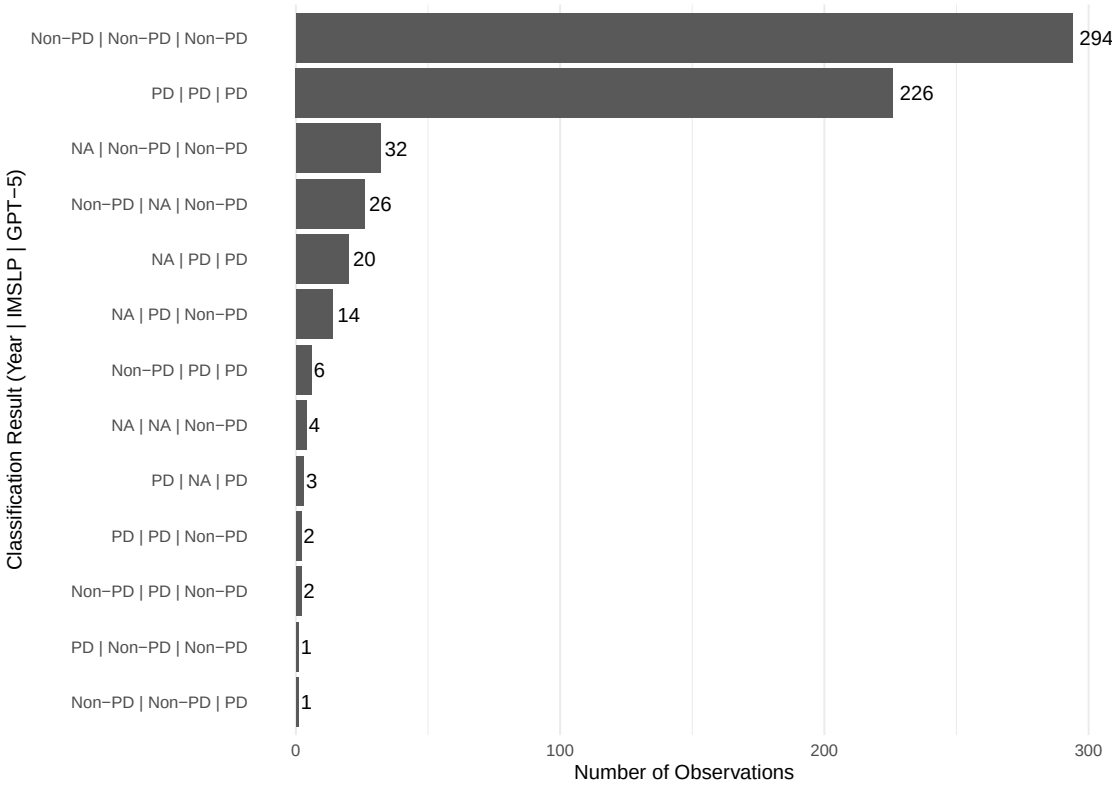


Figure B.1: Frequency of classification patterns across the three methods (year, IMSLP, GPT-5). *"Non-PD — Non-PD — Non-PD"* denotes that all three methods classified the composition as not being in the public domain, whereas *"NA — PD — Non-PD"* indicates that the copyright status could not be determined using publication year, IMSLP denoted the composition as in the public domain, but GPT-5 as not in the public domain.

C Resampling Procedure for the Classical US-EU Sample

For our audit of the Classical US-EU sample, we uploaded multiple video recordings for each composition to YouTube. This was necessary because classical works are often performed and recorded by different artists, and we wanted to ensure that our audit captured a representative set of recordings for each work. However, this design choice means that our raw dataset contains multiple observations per composition.

To estimate enforcement accuracy at the composition level—that is, whether Content ID correctly distinguishes between copyrighted and public domain *compositions*—we collapse the data to one observation per composition. Specifically, for each composition, we randomly sample one of its associated video uploads using a fixed random seed to ensure reproducibility. This resampling procedure yields a dataset of 631 unique compositions, each represented by a single randomly selected recording.

This approach ensures that our estimates of enforcement accuracy are not confounded by within-composition correlation in claim outcomes. Without resampling, compositions with more uploads would receive greater weight in our analysis, potentially biasing our results if the number of uploads is correlated with copyright status or claim likelihood.