

Naughty Bits

An Empirical Study of What Consumers Would Mute and Excise from Hollywood Fare if Only They Could

Should parents have the freedom to block potentially offensive language, sexuality, and violence from the films their children watch at home? Should an adult with reservations about explicit material be allowed to experience the movie Titanic without that film's one notorious nude scene, or Schindler's List without its most uncomfortable audio and video moments? And are these freedoms rightly limited by the relevant decision-maker's ability to engage the fast-forward and mute buttons quickly enough, or should copyright law make room for more sophisticated solutions, even over the objections of a hostile copyright community? In this Article, we offer a unique contribution to this long-running debate: detailed data about what consumers would mute and excise from Hollywood films if only they could. Specifically, we report on the decisions made by roughly 300,000 viewers as they filtered and then watched nearly 4 million movie streams during calendar year 2016. The data, we argue, make a strong case in favor of a permissive copyright regime where viewers would have significant freedom to filter films according to their own religious, moral, and public policy convictions.

I. Introduction

In the late 1990s, the motion picture Titanic was a veritable blockbuster. It was the first film in history to reach more than one billion dollars in worldwide box office gross,¹ and it set records that still stand today for both the highest number of Oscar nominations² and the highest number of Oscar wins.³ When the movie was released on videotape in 1998, copies unsurprisingly flew off the shelves nationwide. Yet, in the heavily Mormon town of Pleasant Valley, Utah, almost no one was buying. "The problems were not only sexual," reported Don Biesinger, owner of the local video store, but also "violence, dead bodies and everything."⁴ So Biesinger took matters into his own hands: he purchased relatively crude video editing equipment, invited customers to bring to his store legitimately acquired videotaped copies of the film, and then, for five dollars, offered to physically cut the controversial scenes from each tape.

¹ Leonard Klady, *'Titanic' Sails to All-Time Box Office Record*, VARIETY (Mar 3, 1998).

² Joey Nolfi, *'La La Land' ties record for most Oscar nominations ever*, ENTERTAINMENT WEEKLY (Jan 24, 2017).

³ Mahita Gajanan, *Here's Who Has Won the Most Oscars Ever*, TIME (Feb 14, 2018).

⁴ Katharine Biele, *Videos: Who Makes the Final Cut?*, THE CHRISTIAN SCIENCE MONITOR (Oct 5, 1998) (quoting Biesinger, owner of Sunrise Family Video in Pleasant Valley, Utah).

Biesinger told reporters at the time that, in his view, his service was legal. As he explained, “We figure [our customers] own the movies and can do anything they want with them.”⁵ But copyright holders vehemently disagreed. Paramount Pictures publicly threatened to take “all necessary and appropriate” steps to protect the studio’s interest in the film.⁶ An industry-wide boycott of Biesinger’s mom-and-pop shop was briefly proposed,⁷ and there was even talk in Hollywood of adding a protective contractual provision (the Biesinger Clause) to the web of contracts under which movies wend their way from the backlot to the theater and ultimately to the video store.⁸ One prominent director well summarized the mood, complaining that Biesinger “is mutilating art. It’s like being a museum curator and taking a Picasso painting of a nude and painting a bathing suit on it and trying to pass it off as a Picasso.”⁹

No litigation was ever filed with respect to Biesinger’s do-it-yourself approach to film editing, and Biesinger himself passed away just a few months after the editing controversy exploded.¹⁰ But the genie was out of the bottle, and, by 2002, a who’s-who Hollywood litigation was framed: on one side, a handful of technology companies with names like Clean Flicks, Clear Play, Family Flix, and Play It Clean Video; on the other, movie studios including MGM, Time Warner, Sony, Disney, DreamWorks, Paramount, and Twentieth Century Fox, along with prominent movie directors, including Steven Soderbergh, Steven Spielberg, Robert Redford, and Sydney Pollack; and, in the middle, Judge Richard Matsch of the District of Colorado.¹¹

Two types of filtering technology were at issue. Firms that offered “Mechanical Editing” sold replacement versions of major Hollywood films, typically by purchasing legitimate videotapes or DVDs, copying the inoffensive portions from those originals onto new physical media, and then selling the new media while destroying each original. The idea was to preserve a one-to-one relationship between purchased and edited versions so as to stave off any allegation that these businesses were in fact engaged in piracy. Firms that offered “Electronic Editing,” by contrast, did not sell movies at all. Instead, these firms sold playback devices capable of recognizing whatever movie a customer inserted into the device, pulling data about the objectionable material present in that particular film, and then, during playback, automatically engaging mute and fast-forward accordingly. Here, too, the idea was to ensure

⁵ *Id.*

⁶ David M. Bresnahan, *Utah Video Business vs. Hollywood*, WND Independent News Network (Sept 3, 1998) (quoting a Paramount spokesman); Seeworthy, *PEOPLE MAGAZINE* (Oct 26, 1998) (same).

⁷ Jeff Call, *Video store may be winning battle with studio over edited films, but theater feels backlash from all in Hollywood*, DESERT NEWS (July 30, 1999).

⁸ *Id.*

⁹ Associated Press, *Film Director Furious at Am. F. video store*, DESERT NEWS (Oct 27, 1998) (quoting famed movie director John Frankenheimer).

¹⁰ Call, *supra* note 7 (noting Biesinger’s passing).

¹¹ See Docket, *Huntsman, et al v. Soderbergh, et al*, Civil Case 1:02-cv-01662 (filed Aug 29, 2002) (D Colo.).

that copyright holders would be compensated for at-home filtered performances. Electronic Editing only worked if the user first acquired a legitimate, unedited copy of the film.

For four years, copyright holders nevertheless argued that both of these approaches violate copyright law.¹² Mechanical Editing involved the creation of physical copies, thus implicating copyright holders' exclusive rights to reproduce their work. And, while Electronic Editing did not implicate reproduction in the traditional sense, Electronic Editing did involve the creation of a detailed database tracking the timing and extent of the language, violence, and intimacy in each film, and it likewise resulted in the creation of an at-home performance that was clearly based on, and largely redundant to, the performance associated with the unedited copyrighted original. The copyright allegations were thus sweeping. The Mechanical Editing Parties were accused of "directly infringing the Studios' copyrights in their films by creating, marketing, distributing, duplicating, renting, selling and/or offering for sale unauthorized edited versions of the Studios' films on VHS videocassettes and digital versatile discs (DVDs)," and the Electronic Editing Parties were accused of "creating, distributing, marketing, and selling the exact same infringing end-product" but "merely accomplishing these same ends utilizing different technologies."¹³

The technology companies responded by arguing public policy, both to the court¹⁴ and in parallel to Congress. They argued that any technical acts of infringement should be excused because their products and services protect children from indecent materials. They argued that filtering serves the public interest because it helps diverse communities come together behind shared cultural experiences while still respecting individual religious and moral constraints. And they argued that, because the technologies in dispute all preserved a one-to-one relationship between original and filtered films, filtering of the sort at issue in the case did not

¹² In addition to the copyright claims, Hollywood stakeholders also asserted claims under trademark and unfair competition law. See, e.g., *Plaintiffs' Brief in Support of Summary Judgment*, *Huntsman v. Soderbergh* (June 8, 2003) (D. Colo.) (summarizing the issues). Those theories strike us as uninteresting, however, because any confusion caused by filtering can be easily addressed through clear labeling of the various products, services, and films. Besides, the fact that unedited versions of every implicated film were also available in the marketplace seems likely to have further mitigated the risk that consumers would perceive an edit poorly and wrongly assign blame. Cf. Statement of Mary Beth Peters, Register of Copyrights, before the Subcommittee on Courts, The Internet and Intellectual Property of the House Committee on the Judiciary, 108th Congress (June 17, 2004) ("While the artistic integrity as well as the continuity of [an edited] motion picture may suffer, the person viewing the edited performance is fully aware that he or she is viewing a performance of less than the entire motion picture because that was his or her preference.").

¹³ *Motion Picture Studio Defendants' Statement Clarifying Claims*, *Huntsman v. Soderbergh* (March 11, 2003) (D. Colo.) at p. 3.

¹⁴ Copyright law's "fair use" doctrine invites accused infringers to explain, on policy grounds, why activity that might otherwise infringe copyright should nevertheless be excused in a particular case. The doctrine is codified at section 17 U.S.C. § 107, but it remains an explicitly flexible, equitable doctrine that welcomes a wide range of arguments from accused infringers. There is an enormous body of literature on the topic. For a sampling, see William W. Fisher III, *Reconstructing the Fair Use Doctrine*, 101 HARVARD L. REV. 1661 (1988); Pierre N. Leval, *Toward a Fair Use Standard*, 103 HARVARD L. REV. 1105 (1990).

harm copyright holders economically. Judge Matsch was not convinced. “What is protected [by copyright] are the creator’s rights to protect its creation in the form in which it was created,” wrote the judge in what was his final substantive ruling in the case.¹⁵ Congress, meanwhile, split the baby, enacting a limited exception to copyright law under which Electronic Editing, but oddly only Electronic Editing, was deemed not to infringe copyright.¹⁶

Fast forward to modern times, and the debate over film filtering is heating up anew. Last year, the Ninth Circuit issued a precedential opinion clarifying the degree to which the above-described Electronic Editing exception applies to companies that stream edited movies online.¹⁷ The short version: it’s complicated.¹⁸ Right now, that same case is back pending at the district court level, with the court receiving evidence on the question of whether other copyright law exceptions, including copyright’s fair use doctrine, might apply instead.¹⁹ In parallel, Congress is actively considering legislation that would protect from liability a broader range of at-home filtering technologies.²⁰ And technology companies continue to experiment with new approaches regardless of the law, including artificial intelligence systems that in the not-too-distant future might be able to engage in fully automated, context-sensitive, personalized movie filtering,²¹ at least if either copyright holders allow it, or the law is interpreted such that technology companies can pioneer these new products and services over Hollywood’s objections.

It is into this ever-changing technical and legal ecosystem that we introduce the one thing that has been long missing from the fight: data. Specifically, in this Article, we report and analyze previously unavailable, detailed information about what roughly 300,000 real people chose to mute and excise from Hollywood films when they were given the chance. This information is directly relevant to the legal and policy questions pending today before both

¹⁵ *Clean Flicks of Colo., LLC v. Soderbergh*, 433 F. Supp. 2d 1236, 1240 (D. Colo. 2006).

¹⁶ Perhaps Congress was worried about piracy. After all, under Electronic Editing, there is no way for a viewer to watch a filtered film without first independently acquiring access to the unedited original. Under Mechanical Editing, however, there is at least some risk that a dishonorable firm will create an edited version and then not actually destroy the original copy from which it came.

¹⁷ *Disney Enters., Inc. v. VidAngel, Inc.*, 869 F.3d 848 (9th Cir. 2017).

¹⁸ For discussion, see Thomas B. Norton, *Watch What You *Bleeping* Want: Interpretation of Statutes Dealing with Advancing Technology in Light of the Ninth Circuit Case of Disney Enterprises, Inc., v. VidAngel, Inc.*, 25 JOURNAL OF THE INTELLECTUAL PROPERTY LAW ASSOCIATION 287 (2018).

¹⁹ See Docket, *Disney Enterprises, Inc. v. VidAngel, Inc.*, Civil Case No. 2:16-cv-04109 (C.D. Cal. 2016) (pending).

²⁰ See *Family Movie Clarification Act of 2018*, H.R. 6816 (115th Congress) (sponsored by Utah Representatives Love, Biggs, Curtis, Bishop, and Stewart, and Illinois Representative Messer).

²¹ Facebook has already deployed systems of this sort, using artificial intelligence to find and filter nudity, graphic violence, hate speech, and the like. See Daniel Terdiman, *Here’s How Facebook Uses AI to Detect Many Kinds of Bad Content*, FAST COMPANY (May 2, 2018). Google, too. See Amelia Heathman, *Google is using AI to find child sexual abuse material online*, EVENING STANDARD (Sept 4, 2018).

courts and Congress. For example, if filtering technologies turn out to be primarily used to protect children from undesirable exposure to indecent materials, courts interpreting the relevant copyright exceptions and politicians debating potential changes to the law might both be more inclined to make room for film filtering. Protecting children from indecency, after all, is a concept with a long and distinguished legal pedigree, anchored in the Supreme Court's famous *Pacifica*²² decision and reaffirmed most recently in *Ashcroft v. ACLU*, a case that explicitly endorsed filtering as an effective response to salacious materials online.²³ By contrast, courts and lawmakers might be less forgiving if filtering turns out to simply be a way for adults to duck what would otherwise be socially beneficial opportunities to question their own religious, moral, and policy views on charged topics like race and sexuality.

The data that undergirds this paper comes from the technology startup VidAngel. During 2016, VidAngel offered a streaming service that empowered users to mute potentially offensive audio and cut potentially offensive video from Hollywood films. Copyright litigation forced VidAngel's service offline in December of that year. But, in the preceding eleven-and-a-half months, VidAngel managed to transmit roughly four million filtered streams and, for each, to record not only which filters each user applied to which film, but also how many minutes of the resulting film each user then watched. We use that data to explore many of the above-discussed issues. For example, we characterize the types of materials cut, discovering, surprisingly, that video filters are primarily used to filter scenes involving intimacy, rather than those related to violence. We also examine the extent to which viewers want to edit films in the first place, finding that even the most cautious viewers use filters as scalpels, not sledgehammers, muting and excising only a tiny fraction of a film's content. We ultimately use this information to argue in favor of an expanded copyright exception that would clearly protect modern filtering technologies from allegations of infringement.

The Article proceeds in six short sections. We begin in Section II with a discussion of the reasons Hollywood stakeholders are reluctant to offer viewers the option of watching filtered films. In Section III, we introduce and summarize VidAngel's service and the data it generated. Section IV reports descriptive information about the market for filtered films. Among other things, we report the number of filters triggered on average per stream, and the average amount of content cut. Section V examines the types of content being filtered. Here, for example, we report findings as to which words were most likely to be muted from a film's stream. Section VI addresses user satisfaction, reporting information about the extent to which users enjoyed the films they themselves chose to filter. Lastly, Section VII offers our view on the implications

²² *FCC v. Pacifica Foundation*, 438 U.S. 726 (1978) (upholding federal regulations designed to reduce the risk that children will be inadvertently exposed to indecent language during daytime hours on broadcast radio).

²³ *Ashcroft v. ACLU*, 542 U.S. 656 (2004) (striking down the Child Online Protection Act on the ground the user-defined filters were a sufficiently effective, less restrictive alternative).

of this work, focusing in particular on the question of whether copyright holders should have a veto right over technologies and services that would otherwise empower users to filter at home.

II. The Controversy

At first blush, one might think that Hollywood stakeholders would all largely support the idea of offering filtered films to interested viewers. Filtering, after all, looks like a clear win-win: viewers would be enticed to watch movies they might otherwise skip due to the presence of uncomfortable content, and the various Hollywood stakeholders would enjoy the artistic and financial benefits associated with reaching even larger audiences.

Yet, Hollywood stakeholders have historically engaged in only limited forms of self-filtering anyway. Nearly every major Hollywood studio works with television broadcasters to create versions of popular films that comply not only with regulations promulgated by the Federal Communications Commission but also with cultural norms about what words and images are appropriate for advertiser-supported broadcasts. Studios and directors similarly work with airline carriers to remove potentially offensive content, there not because of any legal rule but instead for fear of the public outcry that might follow if offensive materials were shown on an airline's big screen or offered privately on seat-back entertainment units. Studios and directors even occasionally remove potentially troubling content prior to a film's theatrical release, as when a film is originally rated NC-17 by the Motion Picture Association of America but then earns a less-restrictive, more audience-friendly rating simply by cutting a scene or editing some dialogue.²⁴ But no studio offers a comprehensive slate of filtered films ready for easy, at-home viewing. No studio offers a service on par with what VidAngel offered to its customers, namely tools through which viewers can choose which films, and what content, to quiet. And nearly every studio has at one time or another actively opposed the sale of third-party products and services like VidAngel's streaming service or the specialized DVD players described in the Introduction. So what explains this hesitant industry posture?

One possible explanation is that, for some Hollywood stakeholders, filtering is perceived as an affront to artistic integrity. Under this argument, removing profanity, cutting short an intimate scene, or otherwise limiting potentially uncomfortable content is problematic because, by definition, it distorts the original film. Schindler's List, from this perspective, is simply not the same movie if the Jewish prisoners are only shown fully clothed and the German guards speak only respectful words. Likewise, and contrary to what the Ninth Circuit Court of

²⁴ See Jason Bailey, *10 Films That Avoided the NC-17 Rating and Suffered for It*, THE ATLANTIC (Dec. 2, 2011) (offering examples where scenes and dialogue were cut for this reason from films including Eyes Wide Shut, Natural Born Killers, and Basic Instinct). Interestingly, studios sometimes add extraneous language or nudity in order to move a film to a higher rating category, again for marketing reasons.

Appeals recently stated in a published opinion on these issues,²⁵ Return of the Jedi is emphatically not the film George Lucas imagined if Princess Leia never dons that shimmering bikini. To at least some Hollywood stakeholders, then, any expansion in the availability of filtered films is problematic because filtering itself is problematic. More viewers might watch Straight Outta Compton if they could do so without hearing the film's vulgarity, but Dr. Dre might for his own reasons prefer not to make that option available. On this argument, that call is his to make.

Another possible explanation for the industry's awkward posture might be that Hollywood stakeholders are worried about the quality of the filtered film experience. Filters that remove objectionable audio or video might inadvertently also remove scenes or dialogue critical to the film's story. The conclusion of the Imitation Game, for example, would make little sense without the film's earlier segments focused on Alan Turing's sexuality. Jodi Foster's character in The Accused similarly cannot be understood without at least some of the film's discussion of the rape that launches the film's central story arc. Moreover, choppy video transitions and unscripted instances of silence can detract from the viewing experience even if the removed material is not otherwise consequential from a storytelling perspective. On this argument, then, filtering is undesirable because it might trigger a reputational backlash where viewers watch a filtered film, have a bad experience, and then think poorly of the implicated studio, film, writers, director, producer or actors.

A third reason Hollywood stakeholders might be reluctant to support an expanded market for filtered films is that filtering plausibly creates a mismatch risk where, because filtered films are available, a user might watch a film that is inappropriate for that user even in filtered form. For example, viewers who disapprove of sexual content might think they can enjoy Fifty Shades of Grey with the film's intimate scenes removed, but, in reality, viewers who dislike on-screen intimacy will probably still dislike that film given its overarching themes, story, and humor. (No one watches Fifty Shades for its witty, platonic banter.) Pulp Fiction is similarly unlikely to win rave reviews from viewers squeamish about on-screen violence even if every violent scene is excised. The problem in these examples is not that filtered films are necessarily less enjoyable than are unedited versions. Rather, the problem is that viewers might err, thinking that poorly-suited films can be sufficiently transformed by filters when, in fact, they cannot be.

A fourth explanation might be that Hollywood insiders doubt that the market for filtered films is large enough to warrant all the fuss. After all, how many viewers really want to

²⁵ *Disney Enterprises, Inc. v. VidAngel, Inc.* 869 F.3d 848, 861 (2017) ("*Star Wars* is still *Star Wars*, even without Princess Leia's bikini scene."). Note that the Court mixed up its *Star Wars* flicks, as Leia never dons a bikini in the original Star Wars film. The famous such scene comes from Return of the Jedi.

watch Deadpool without the film's iconic bathroom humor,²⁶ or experience Basic Instinct without Sharon Stone's infamous seduction scenes? And will those viewers in the aggregate pay enough to warrant not only the costs summarized above, but also the costs associated with creating filtered versions, educating the public as to what is and is not included, and weathering the inevitable controversies when (say) a scene involving a same-sex kiss is categorized as either offensive or innocuous, and some segment of the public vehemently disagrees?²⁷

The VidAngel data offers an unprecedented opportunity to evaluate these competing explanations. Specifically, on the question of market size, VidAngel's data show the extent to which approximately 300,000 subscribers filtered films over a twelve-month period, including details about how many audio and video segments were removed and the resulting impact on each film's audio and video presentation. On the question about movie quality, and likewise the possibility of movie/viewer mismatches, VidAngel's data quantify viewer satisfaction by showing how much of each movie each user actually watched. If viewers do not enjoy the films they filter, that displeasure should show up as a noticeable drop in movie viewing times. Lastly, on the issue of artistic integrity, the VidAngel data at a minimum help to quantify the extent of any potential affront. If viewer-selected changes are small in frequency and scope, for example, perhaps the threat to creative freedom is less than that which might be originally feared.

III. VidAngel and Its Data

As noted above, this Article is built on data gathered by VidAngel in the course of running its business during calendar year 2016. The data obviously began as confidential information owned and controlled by VidAngel. But, pursuant to an agreement negotiated in February 2017 and signed in March 2017, we were given unfettered access to the relevant databases and server logs, along with permission to use that data to produce academic papers on the topic of filtered films. VidAngel specifically agreed that we would be free to publish any information we derive from the data, no matter whether the implications are favorable or unfavorable to VidAngel's business, and we in exchange promised not to disclose the raw data publicly and not to disclose summary information that would plausibly help VidAngel's competitors reverse engineer the details of its system. Our work was approved by UCLA's Internal Review Board in March 2017 on the ground that, while VidAngel's data capture

²⁶ *But see*, or perhaps choose not to see, Once Upon a Deadpool, a PG-13 version of the R-rated film Deadpool 2. Chaim Gartenberg, *Deadpool 2 returns to theaters as PG-13 Christmas film Once Upon a Deadpool*, THE VERGE (Nov 5, 2018).

²⁷ Delta Airlines weathered such a controversy in 2016 when it chose to include as part of its inflight entertainment offerings an edited version of the film Carol, where the edits removed every intimate scene between the two female leads, including scenes that simply showed the two women kissing. Audiences objected, in part because of the perceived double standard being applied to homosexual, as opposed to heterosexual, intimacy. See Adam Lujan, *Carol: Controversy erupts after removal of same-sex kissing scenes on Delta Airlines flights*, ENTERTAINMENT WEEKLY (August 5, 2016).

information about real people, the data we received did not include any individually identifiable information, nor any information that would allow re-identification of individuals. We know VidAngel's customers only by the arbitrary numbers that VidAngel used internally to label each account.

We are confident that the data we were given is indeed VidAngel's genuine raw operating data, provided to us in good faith and in unedited form. Part of our confidence comes from our subjective judgment as to the credibility and intentions of VidAngel's founders and its general counsel. Part comes from the fact that most of the data came to us in the form of overwhelmingly large and complex JSON data files that were difficult to parse even using litigation-grade computational platforms and hence would have required Herculean effort for a startup like VidAngel to falsify convincingly. And part of our confidence comes from the reality that VidAngel is currently involved in litigation that directly pertains to exactly these topics and hence the company is vulnerable to court processes that could at some point force it to share this same data with its adversaries. Disingenuous interactions with us would therefore buy the company very little; any misdeeds will inevitably come to light, undermining whatever benefit VidAngel might have expected to obtain from our work, and harming the company's credibility with us, the court, and the public.

Turn now to the specifics. Whenever a user streamed a movie using the VidAngel system, that user was presented with an interface through which he or she could choose specific audio snippets to mute and specific video snippets to cut. A sample of the information presented to users is shown as Figure 1, where this particular sample screen shows some of the filtering options available for the movie 13 Hours: The Secret Soldiers of Benghazi.

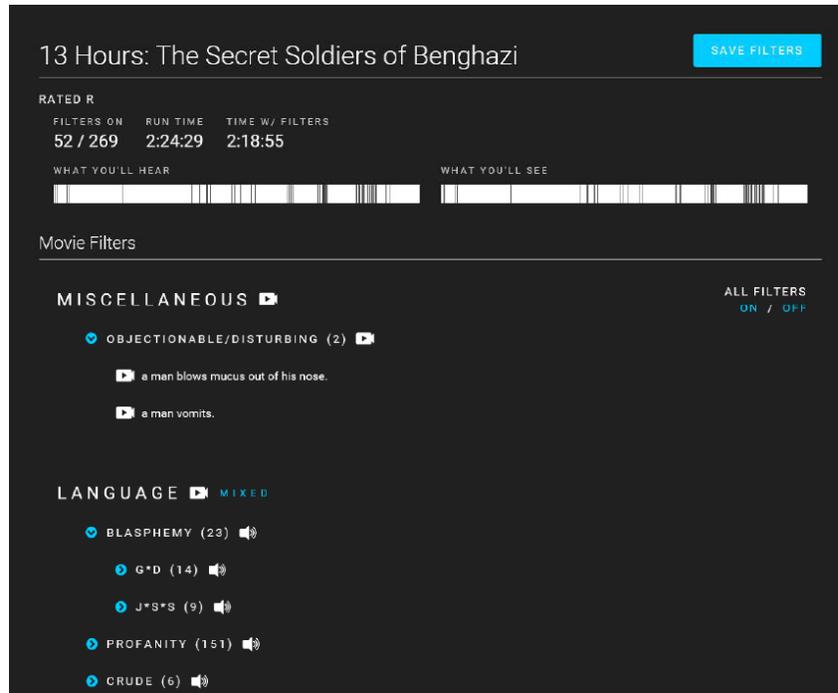


Figure 1: A sample of what a user would have seen when preparing to watch the movie 13 Hours: The Secret Soldiers of Benghazi on VidAngel.

As the screenshot makes clear, users were given a great deal of information about which filters were available for each film. First, each filter was labelled with a word or phrase that described the underlying content and indicated whether the content was audio that could be muted or video that could be cut. In the sample screen, for instance, there are audio snippets labeled “G*D” and “J*S*S,” and there are video snippets labelled with the phrases “a man blows mucus out of his nose” and “a man vomits.” Second, each filter was presented as part of a hierarchy of categories and subcategories, and users could select filters by clicking on those groupings as a whole or on specific filters one by one. For 13 Hours, available groupings ranged from broad combinations such as “all filters” or all “language,” to narrower groupings such as “all language involving blasphemy” or any audio of the word “G*D.” Third and finally, there was always a graphic representation at the top of the screen showing how much audio or video would be trimmed by the selected filters. Audio and video were represented in separate graphics, with vertical black bars showing the size of each snippet as well as its approximate location in the film.

VidAngel offered a comparable, flexible hierarchy for nearly every movie in its library. Figure 2, for example, shows sample screenshots for two other popular films, Back to the Future and The Big Short. For these movies, too, users could choose to eliminate extremely broad

categories of material like “sex/nudity/immodesty” or smaller groupings like “profanity,” “kissing,” or “sexual reference/innuendo.”

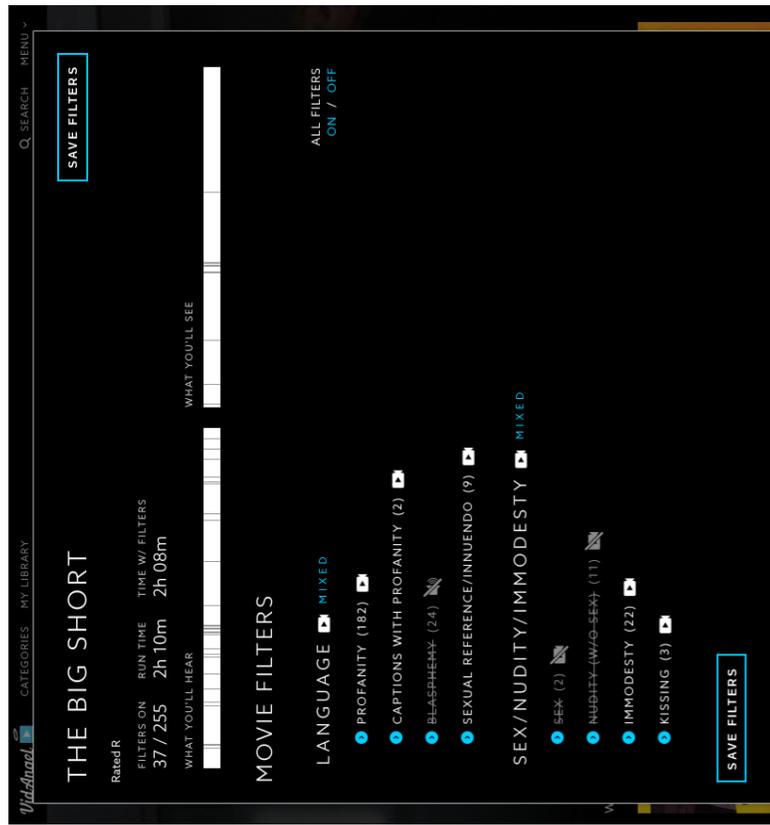
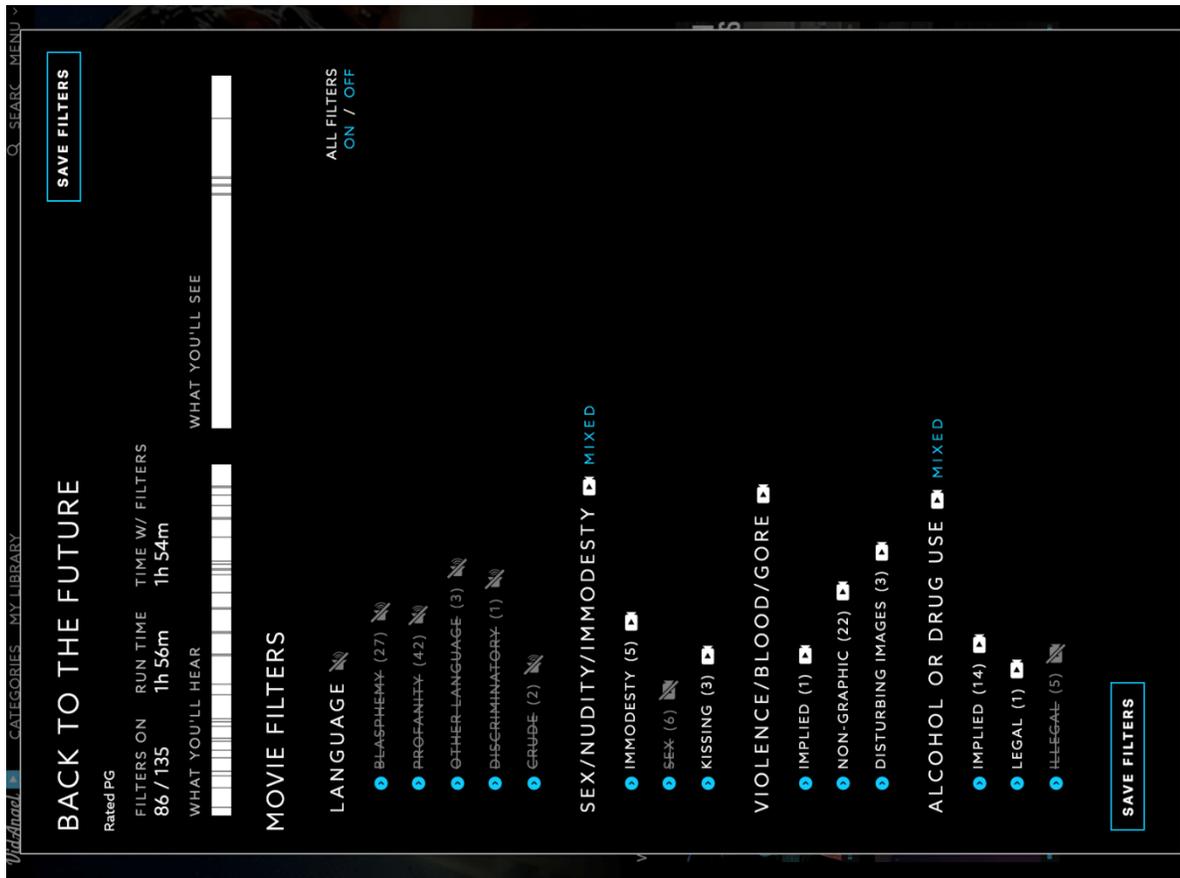


Figure 2: Sample screens showing what users might have seen when preparing to watch the films Back to the Future or The Big Short, respectively.

VidAngel provided us with nearly all the data used to generate these screen displays.²⁸ Specifically, for each filter, we were given a short description of the filter, such as “immodesty_male” or “racial_slurs”; identification of the movie to which the filter applied; information about whether the filter impacted audio, video, or both; and timing information sufficient to calculate how much audio or video was at stake with respect to that filter. The resulting dataset allowed us to interpret user requests. For example, in any instance where we knew the movie being streamed and the filters selected, we could use this data to understand exactly what content was targeted and how much the film’s resulting audio and video presentation would change.

Figure 3 offers a snapshot summary of this filter data. Along the horizontal axis, we track the number of individual audio and video filters available for each movie; on the vertical axis, we count the number of movies in the dataset for which that number of filters was available. Thus, for example, because users watching Back to the Future could choose from among 135 filters, that film is one of the roughly 400 films counted in the fourth bar from the left. Users watching The Big Short were able to choose from among 255 filters, making that movie one of the roughly 230 films counted in the ninth bar from the left. In the aggregate, for the 2,767 movies that VidAngel offered in 2016, users could choose from among 477,736 individual filters.

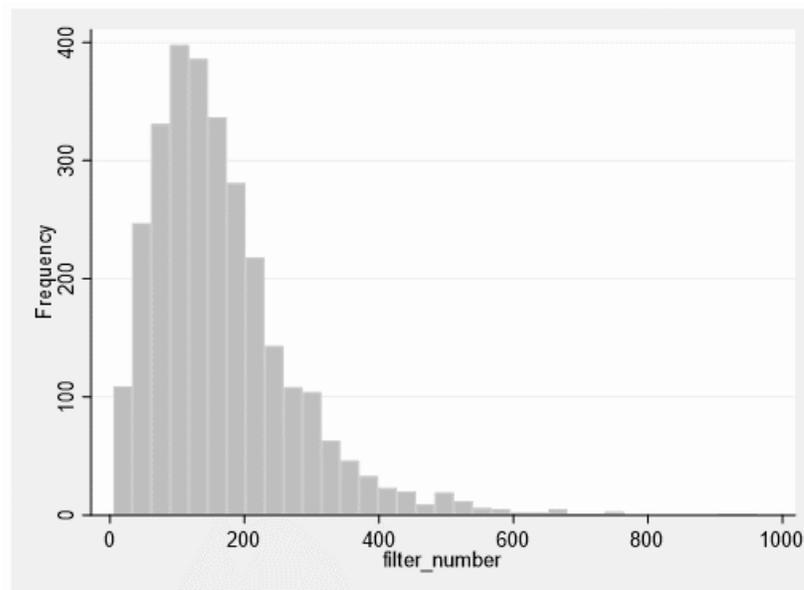


Figure 3: A histogram showing how many filters were available per movie. The horizontal axis counts the number of filters available for each specific movie, and the vertical axis counts the number of movies that offered that number of filters.

²⁸ The only data we lacked was information sufficient to know exactly where within each movie each filter applied. So, while we knew which filters were applicable to which films, and how many seconds of content were at stake each time, we could not ourselves fully recreate the graphics showing “what you’ll hear” and “what you’ll see.” See *infra* note 37 for further discussion of this missing data and its implications.

The above discussion focuses on VidAngel’s data about movies. But VidAngel also provided us with a database that logged every streaming request the company received in 2016. Users in this data were represented by numbers, not names, and thus each entry in the database included a customer number, a movie identifier, a list of the movie-specific filters selected by that user prior to the start of the relevant stream, and a number indicating the total expected runtime of the resulting filtered film. VidAngel also included a “lineup id” that could be used to match the above information to specific data packets sent or received by VidAngel’s servers. A sample data entry along these lines is shown in Figure 4. In this particular lineup, a viewer was watching Batman v. Superman and chose to filter one instance of a man kissing another man and one use of the word “f*ck.” The resulting film had an expected runtime of 10,953 seconds, which is roughly 30 seconds shorter than the published runtime associated with the unedited film.

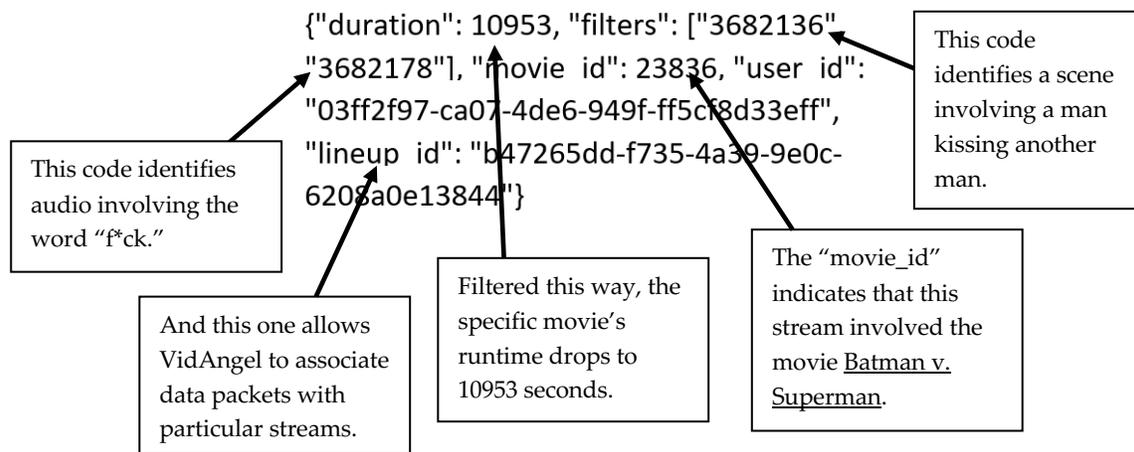


Figure 4: An example of the data VidAngel provided about each requested movie stream.

Our work required yet one more type of data, namely VidAngel’s complete server logs for the year 2016. Whenever VidAngel streamed a movie, its server would request and then endeavor to log minute-by-minute confirmation that specific movie snippets were successfully received by specific users. Each entry included a variety of information, but we focused on six items. Three were simply labels: the movie identification number for the movie being streamed, a user identification number for the customer account at issue, and the lineup ID that allowed us to link log entries to the above-described filter data. The fourth item was a timestamp that showed, down to the fraction of a second, the time and date each log entry was made. The fifth item was a “location” variable that reported exactly where in the movie a given transmission fit. Location data was measured in seconds, such that the tenth minute of a film would be reported as position 600, whereas the thirteenth minute would be reported as position 780. Lastly, the sixth item was a “runtime” variable that showed how long the filtered film would run if the user were to watch the entire edited film. A sample server log excerpt is shown in Figure 5.

| lineup | work_id | location | user_id | created | runtime |
|--------------------------------------|---------|----------|--------------------------------------|----------------------------|---------|
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1066 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:42:19.649664 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1126 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:43:19.561599 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1186 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:44:19.441274 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1246 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:45:19.805650 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1306 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:46:19.519907 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1366 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:47:19.802936 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1426 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:48:19.471114 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1486 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:49:19.597822 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1546 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:50:19.439551 | 8669 |
| 0027d4a0-f72c-4928-afab-4750ab1cd16f | 20802 | 1606 | 1ecca336-92ea-4f18-bee7-99cd53fa28b9 | 2016-12-26T21:51:19.514523 | 8669 |

Figure 5: A portion of a VidAngel server log. The columns, in order from left to right, show information about which filters were triggered, which movie is being streamed, where in the movie the viewer was at that moment, which user account is at issue, the time the log entry was created, and the total expected runtime for the edited film.

The server log was, by far, the most challenging and imperfect data VidAngel provided. The idea behind the log was straightforward: whenever VidAngel streamed a movie to a user, the user’s computer was supposed to send minute-by-minute confirmations back to VidAngel’s servers, acknowledging that a specific part of a specific movie was successfully received. The reality of the log, however, was that a significant percentage of the expected minute-by-minute confirmations turned out to be missing. Missing confirmations in theory could mean that significant portions of many movies did not in fact reach the relevant users. That said, because users presumably would have abandoned VidAngel had its systems in fact failed to deliver movie minutes that frequently, we believe that a better explanation is some combination of (a) user equipment failing to generate confirmations at the desired pace; (b) confirmations being lost while traveling from user equipment to VidAngel; and (c) confirmations arriving at VidAngel’s servers but then not being recorded because VidAngel’s servers were at the relevant time overwhelmed with other, higher-priority processes and activities.

The server logs were important to our work because they alone told us how much of any given stream was watched. Had the server logs been complete, that would have been easy information to extract. Because they were imperfect, however, we instead estimated viewing time by implementing three distinct measurement strategies and then, for each stream, adopting the shortest of the three results. Our first strategy estimated the amount of time each user spent watching his or her personalized stream by subtracting the timestamp associated with the first logged transmission from the timestamp associated with the final logged transmission. Our second strategy subtracted the minimum movie “location” listed for that stream from the maximum “location” listed for that stream. Our third estimation strategy counted the number of log entries associated with each specific stream and treated each as representing one minute of viewership.

Different types of errors and behaviors impacted these various measures differently. A user who fast-forwarded through some portion of a film, for example, inadvertently padded our estimate based on the difference between the smallest and largest “location” transmitted, but did not corrupt our measurement based on timestamps and barely affected our measurement based on the number of log entries received. A user who paused a film to grab a snack or use the restroom, by contrast, inadvertently increased our estimate based on

timestamps and inadvertently increased our estimate based on the number of entries in the log, but did not in that way affect our estimate based on location. Thus, by choosing the shortest of these three measures, we adopted what we believe to be a conservative and yet relatively accurate estimate of the time each user spent watching each stream. Moreover, because we applied this same strategy to every stream under review, our work comparing users who watched largely unedited films to users who watched heavily edited films was further insulated from this unavoidable data challenge.²⁹

Figure 6 shows the viewing length estimates generated by this process. There are two charts. For the chart on the left, the horizontal axis reports our estimate of the minutes a viewer spent watching each stream, and the vertical axis shows the number of streams for which viewers invested that amount of time. For the chart on the right, we report the same information but, instead of counting minutes, we report time as a percentage of the relevant stream’s full runtime. Note that the length of an edited film might be slightly smaller than the full length of the corresponding original because, obviously, video cuts make films shorter.

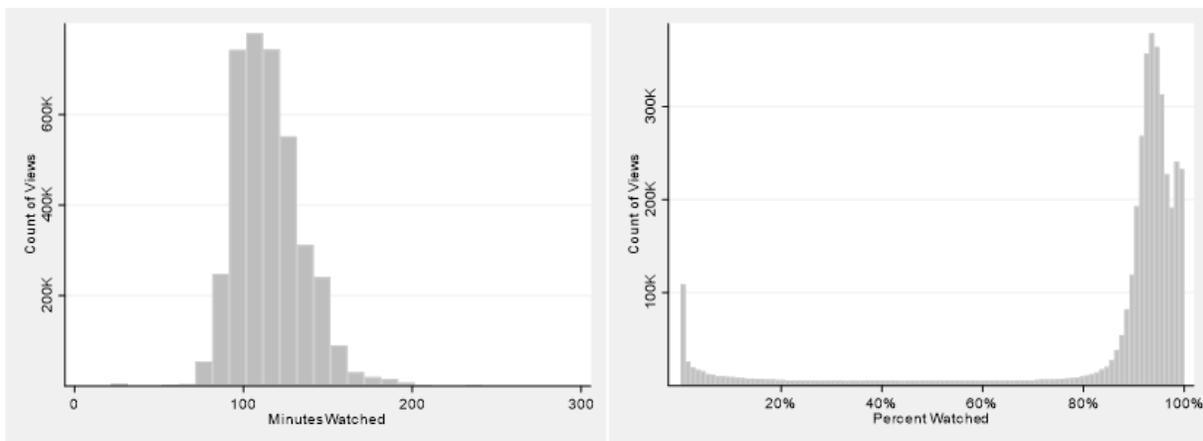


Figure 6: Our estimates as to the amount of time each viewer spent watching each stream, reported on the left in terms of the total number of minutes watched and on the right as a percentage of total runtime.

In the analysis that follows, we exclude streams that lasted ten minutes or less. Those streams appear on the far left side of the first chart shown in Figure 6. We did this because our review of the data suggested that many short streams were either (a) system tests in which a user seemed to have been playing with the system but not actually trying to watch a coherent

²⁹ A small number of our estimates struck us as implausible even after we implemented the above strategy, specifically because they suggested viewing times greater than the length of the relevant edited films. For those, if we could determine the source of the problem, we made the correction. If not, we excluded the data. We ultimately corrected obvious errors in the data associated with roughly 1% of all streams, and we discarded approximately 2,500 streams that we could not reliably interpret. An example of a typical problem: VidAngel’s computers sometimes put an error code in the “location” field, and thus an error during transmission could inadvertently pad our timestamp data (by delaying the film), our count of log entries (because errors were logged), and our location measure (because the error code was sometimes a very large number).

sequence from a film or (b) system failures where the log primarily consists of error codes relevant to failed transmissions rather than data about actual movie viewing. We used all streams that lasted more than ten minutes, however, because some users presumably requested a filtered movie, started to watch, disliked the film, and stopped watching after ten minutes. The dataset we used for our analysis thus included just over 3.8 million streams. Had we used every available stream, the number instead would have been 4.1 million.

IV. Demand for Filtered Films

The VidAngel data offer insight into a variety of questions, but the conversation most naturally starts with the question of whether there actually exists a significant market for filtered films. That is, as evidenced by the VidAngel data, do a substantial number of consumers truly prefer to watch filtered versions of Hollywood films rather than unedited originals? If so, an interesting puzzle is framed, because Hollywood stakeholders have for the last twenty years been leaving that money on the table by neither themselves serving those customers nor allowing others to do so.

VidAngel's data help answer this question because VidAngel's customers were asked to pay the same fee regardless of whether they chose to filter a little, a lot, or possibly not at all. That is, although VidAngel encouraged its customers to meaningfully filter the films they streamed³⁰ and also required users to activate at least one filter in order to stream a film,³¹ as a practical matter users were not obligated to filter films aggressively. Indeed, this is presumably one of the reasons that several studios sued VidAngel. Because VidAngel's service was not licensed by the studios, VidAngel was able to offer a streaming service that was in some respects better than the streaming services offered by licensed firms like Netflix and Amazon, even from the perspective of customers who had no interest in filtering. For example, because VidAngel was not paying per-stream fees to the studios, VidAngel was able to offer standard-

³⁰ Throughout 2016, any time a user tried to stream a movie without selecting any filters, VidAngel would notify the user that filtering was not optional and suggest that, if the user wanted to watch unfiltered fare, the user should instead use another streaming service. VidAngel would then provide links to traditional streaming services, including those offered by Amazon and Apple.

³¹ From January through May 2016, VidAngel required users to select at least one filter before they would be allowed to stream a film. VidAngel deemed this requirement satisfied by any filter, however, including a handful of filters that arguably were not focused on objectionable content. For example, a user could satisfy this requirement by selecting a filter that blocked the opening credits in a movie, or could satisfy this requirement by selecting a filter that blocked the relevant studio's logo. On June 9, 2016, VidAngel adopted a more stringent policy, from that time forward requiring users to select at least one filter targeting specific objectionable content. Users could then filter other things as well, but at a minimum they had to filter at least one substantive element. As we explain *infra* footnote 35 and accompanying text, in our work we took an even more conservative approach, excluding from consideration any filter that did not obviously and exclusively target objectionable content.

definition streams at a below-market rate of \$1 per stream.³² Moreover, because VidAngel would stream any film as long as that film was available for purchase at retail in DVD form, VidAngel streamed at least some films that were not at the relevant time available to rent via licensed streaming platforms.³³ Again, these distinctions understandably drew studio ire; VidAngel justified its service by pointing to its high-minded social purpose but then implemented its technology in a way that gave it several advantages that had nothing to do with filtering. But, from the perspective of our study, these are features, not bugs. If a substantial number of VidAngel users chose to meaningfully filter films even though there were obvious reasons to instead use the service to watch largely unedited films, that would be compelling evidence that there really does exist a meaningful market for filtered Hollywood fare.

To explore this issue, we first amassed descriptive statistics about user choices. In one cut, we calculated for every movie the median number of filters selected by users who streamed it. In another, we looked at the number of streams where the relevant user selected zero filters, one filter, two filters, and so on, up through streams where the relevant user selected more than 200 filters. Both approaches were designed to show the degree to which VidAngel's users took advantage of the filtering options that VidAngel offered. Users who chose the service simply to stream a movie inexpensively, or simply to access a movie that was not conveniently available elsewhere, would presumably have chosen as few filters as possible. Yet, as we show in greater detail below, in 76.5% of the streams, users chose more than just the one filter that VidAngel required. And, in roughly half the streams, users chose not one filter, but more than ten.

Figures 7(a) through 7(c) tell the story more fully. In Figure 7(a), we report the median number of filters applied to each of the 30 most streamed movies in our dataset. There is substantial variation, consistent with the idea that users were intentionally filtering those movies. Figure 7(b) shows comparable information for all 2,767 movies VidAngel streamed in 2016. The horizontal axis counts the number of filters selected per stream, and the vertical axis counts the number of streams to which that count applied. Again, there is substantial variation. The final panel, Figure 7(c), summarizes that same data but in tabular form. Each row is phrased in a "more than" formulation in order to highlight that the filtering shown in the table is filtering above and beyond VidAngel's minimum filtering requirements. That is, because VidAngel required users to select one filter as a prerequisite to streaming any film, we framed Figure 7(c) to focus on the number of filters chosen in addition to that first filter.

³² See Complaint for Copyright Infringement and Violation of Digital Millennium Copyright Act, *Disney Enterprises, Inc. et al. v. VidAngel, Inc.*, No. 2:16-cv-04109 (C.D. Ca. June 9, 2016) at ¶ 4 ("by cutting out payments to copyright owners, VidAngel is able to offer prices that undercut licensed services and charge only \$1 for daily access to movies in standard definition format").

³³ *Id.* at ¶ 5 ("because VidAngel absolves itself of having to abide by contractual restrictions, VidAngel offers content that is not available on licensed VOD services . . . [and then] flaunts this unfair competitive advantage by expressly promoting a selection of titles that are available on VidAngel but 'Not Available on Netflix'").

| Most Viewed Movies | Viewings | Median Filters |
|-------------------------------------|----------|----------------|
| Zootopia | 46,628 | 4 |
| Captain America: Civil War | 41,690 | 24 |
| Batman v. Superman: Dawn of Justice | 38,758 | 17 |
| Finding Dory | 36,849 | 2 |
| Now You See Me 2 | 36,849 | 23 |
| The Jungle Book | 36,572 | 1 |
| The Legend of Tarzan | 35,405 | 9 |
| The Secret Life of Pets | 34,111 | 6 |
| Star Trek Beyond | 33,483 | 11 |
| Central Intelligence | 32,414 | 77 |
| Jason Bourne | 32,391 | 13 |
| Ghostbusters | 30,806 | 28 |
| X-Men: Apocalypse | 29,624 | 10 |
| The Revenant | 29,526 | 83 |
| Star Wars: The Force Awakens | 28,842 | 2 |
| Independence Day: Resurgence | 27,736 | 21 |
| Me Before You | 27,173 | 20 |
| Deadpool | 26,819 | 151 |
| Allegiant | 26,404 | 8 |
| Bridge of Spies | 25,649 | 19 |
| The Good Dinosaur | 23,982 | 1 |
| The Intern | 23,190 | 33 |
| The Angry Birds Movie | 23,162 | 10 |
| Pete's Dragon | 22,884 | 5 |
| Ice Age: Collision Course | 22,601 | 2 |
| Free State of Jones | 21,938 | 11 |
| Miracles from Heaven | 21,344 | 5 |
| The Martian | 21,136 | 30 |
| The BFG | 21,019 | 2 |
| The Huntsman: Winter's War | 19,465 | 9 |

Figure 7(a)

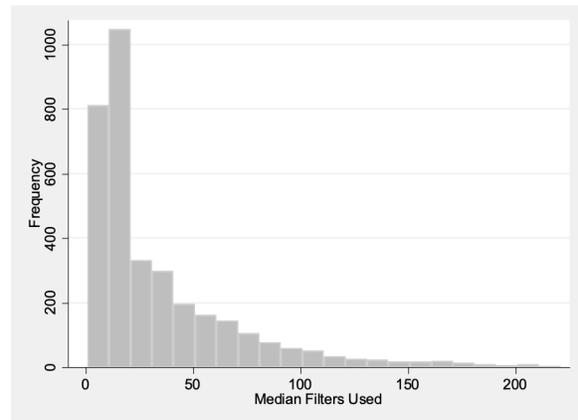


Figure 7(b)

| Number of Filters | Views | Percent |
|-------------------|-----------|---------|
| More than 1 | 2,961,523 | 76.5% |
| More than 2 | 2,667,056 | 69.9% |
| More than 5 | 2,337,429 | 61.3% |
| More than 10 | 2,031,072 | 53.2% |
| More than 15 | 1,804,308 | 47.3% |
| More than 20 | 1,614,846 | 42.3% |
| More than 30 | 1,305,536 | 34.2% |
| More than 40 | 1,076,731 | 28.2% |
| More than 50 | 889,995 | 23.3% |
| More than 100 | 380,341 | 10.0% |
| More than 200 | 105,262 | 2.8% |

Figure 7(c)

Figure 7: A snapshot of the variation in user behavior while using VidAngel to stream films.

Variation was also apparent when we organized the above data not by movie but instead by user account. For example, out of the roughly 196,000 accounts that were used to stream five or more films in 2016, just over 96,400 triggered at least 100 filters for at least one viewing but fewer than 10 filters for at least one other viewing. That is, roughly half of VidAngel’s moderately active accounts exemplified a Jekyll-and-Hyde pattern when it came to their filtering decisions, sometimes aggressively cutting and muting the films they streamed and sometimes leaving the films largely intact. Some of this variation is likely explained by the above-documented variation between movies. That is, a specific VidAngel user likely filtered

differently depending on whether he or she was watching Me Before You or Pete's Dragon. And some is likely attributable to the fact that user accounts were presumably shared to some degree, for example with Mom and Dad watching some movies alone but watching other movies alongside Junior and maybe Grandmother.³⁴

Note that, in all of these counts, we have excluded a handful of filters that VidAngel made available to its customers, and customers in fact triggered, but we thought inappropriate for our project. Some of these were filters that eliminated the character "Jar Jar Binks" from a relevant Star Wars film. VidAngel users might well have valued those filters (weesa woulda) but we excluded them from our data because personalization of that sort struck us as importantly different from personalization driven by concerns about potentially offensive language, sex, violence, and the like. The rest of the excluded filters blocked a film's opening credits, closing credits, or the video shorts that might otherwise have run before or after the main movie. As VidAngel's executive team reminded us, there are substantive reasons to excise these portions of a film, in that credits and shorts do sometimes contain objectionable material. We excluded these filters anyway, however, because we could not tell whether particular users triggered them for substantive reasons or, instead, simply to skip content that viewers in general might similarly prefer not to watch. Excluded filters were not only excluded from the above-reported measures but also from all the substantive analysis that follows, except where explicitly noted.³⁵

The preceding discussion is focused on filter counts, but another way to evaluate user behavior is to look at filter impact as measured in minutes and seconds. To do this, we needed to first establish a baseline against which to compare specific streams. Should we expect that a user truly interested in filtering offensive audio would mute 1 percent of the average PG-13 film? More? Less? And what about video? Should we expect video filters to have a greater impact on a stream because controversial audio tends to be of short duration whereas a violent movie scene can plausibly last several minutes? Or should we expect video filters to have less impact, for instance because users might be reluctant to significantly interrupt a movie's visual presentation or might self-select away from movies they would otherwise need to cut substantially?

³⁴ Interestingly, although VidAngel allowed users to stream films as long as at least one filter was triggered, out of the 196,400 accounts that were used five or more times in 2016, only 2.2% consistently opted to filter at only that required, bare-bones level. The remaining 97.8% of the accounts voluntarily used more than one filter at least once during the course of the year.

³⁵ One implication is that Figure 7(c) understates the degree to which users in fact chose to filter. Consider, for example, a user who in February 2016 streamed a film using two filters: a substantive filter that blocked the word "shit" and a filter that cut the film's opening credits. In reality, this user chose to filter at a level greater than the then-relevant minimum because VidAngel would at that time have deemed it sufficient to simply cut the movie's credits. Yet this stream is not included in any of the tallies in Figure 7(c) because, in light of our exclusion rule, we would count this as a stream with only one filter triggered.

Our approach to these questions was to assume that, for each movie streamed in 2016, there was at least one user who used VidAngel to aggressively filter the film's audio and at least one (possibly different) user who used VidAngel to aggressively filter the film's video. For each film, we thus separately identified the streams with the greatest amount of muted audio and greatest amount of cut video, and used them as baselines for how much audio and video would be cut by a particularly cautious viewer of that specific film. We then developed two additional baselines by repeating those calculations but focusing not on the streams with the highest rates of exclusion but on the streams at the 95th and 90th percentiles.³⁶

For example, Ghostbusters was streamed 30,303 times in 2016. The stream with the highest amount of muted audio quieted 2.2 percent of the sound originally present in the film. The stream at the 95th percentile muted 2.1 percent, and the stream at the 90th percentile muted 1.6 percent. With respect to video, meanwhile, the Ghostbusters stream with the highest amount of excised video cut 10.8 percent of the original film, and the streams at the 95th and 90th percentiles cut 4.5 percent and 3.7 percent, respectively. Figure 8 captures these numbers in a table and also shows comparable measurements for three other popular movies, The Revenant, X-Men: Apocalypse, and Central Intelligence. Note that, while some of these percentages might seem small, small percentages are still meaningful. Even an R-rated movie with a lot of potentially offensive language in fact dedicates only a small percentage of its total runtime to troubling words. Similarly, a PG-13 movie with potentially problematic physical intimacy is nevertheless typically overwhelmingly comprised of uncontroversial visuals.

³⁶ An alternative approach would have been to use as a baseline the maximum amount of content that VidAngel allowed a user to filter for each specific movie. We shied away from this approach because it depends so heavily on choices made by VidAngel and its team. For example, if VidAngel chose to list as potentially offensive a large number of audio segments that in fact turned out to trouble none of its customers, our actual baselines would not be affected but this alternative baseline would have been. Note that VidAngel's choices do nevertheless impact our baselines in part, because a cautious viewer could filter at most only what VidAngel thought to offer. That said, because VidAngel offered so many filters for every movie, and because VidAngel was trying to entice cautious viewers to use its service, we doubt that this was much of a constraint.

| | Audio | | | Video | | |
|----------------------|-------|-----------------|-----------------|--------|-----------------|-----------------|
| | Max | 95th Percentile | 90th Percentile | Max | 95th Percentile | 90th Percentile |
| Central Intelligence | 6.59% | 6.46% | 5.42% | 15.20% | 11.30% | 10.03% |
| Ghostbusters | 2.22% | 2.11% | 1.58% | 10.84% | 4.52% | 3.70% |
| The Revenant | 1.21% | 1.21% | 1.17% | 16.39% | 10.55% | 6.87% |
| X-Men: Apocalypse | 0.77% | 0.77% | 0.41% | 24.16% | 8.43% | 5.20% |

Figure 8: Baselines for how much audio and video might be cut by a particularly cautious viewer.

Using these baselines, we can now characterize the demand for filtered films. Consider an example using only Ghostbusters and, at that, focusing first only on audio. The above chart shows that the most cautious user who used VidAngel to watch Ghostbusters cut 2.22% of its audio. One way to represent demand, then, is to ask how many other users who watched Ghostbusters filtered (say) 90% as much audio as this extremely cautious user, 75% as much audio as this extremely cautious user, and 50% as much audio as this extremely cautious user. We can then ask that same question again but, instead of using the most extreme user as the baseline, we can use the users at the 95th or 90th percentiles. Then we can repeat all nine calculations but using the percentages applicable to video instead of audio. Figure 9 captures those results, again focused only on Ghostbusters. The table reports, for example, that 8.78 percent of all users filtered at least 75 percent as much audio as did the user whose audio filtering decisions put that user at the 95th percentile overall.

| Comparative Levels of Filtering: Audio | | | | Comparative Levels of Filtering: Video | | | |
|--|--------|--------|--------|--|--------|--------|--------|
| | 90% | 75% | 50% | | 90% | 75% | 50% |
| Maximum | 5.35% | 6.76% | 18.68% | Maximum | 0.44% | 0.50% | 2.69% |
| 95th Percentile | 5.56% | 8.78% | 26.50% | 95th Percentile | 7.22% | 13.08% | 26.50% |
| 90th Percentile | 12.17% | 16.77% | 35.39% | 90th Percentile | 13.51% | 19.60% | 35.10% |

Figure 9: The demand for filtering audio and video in Ghostbusters.

Now we can expand this analysis to the full data set. In Figure 10, we report the same type of information as shown in Figure 9, except that the new figure summarizes this information for every movie in our dataset, not just for Ghostbusters. The data shown in Figure 10 are therefore averages, and specifically they are weighted averages calculated by giving each movie a weight based on the number of times it was streamed in 2016. Thus, for example, because Ghostbusters was streamed 30,303 times in 2016, its max audio percentage was given correspondingly more weight than was the max audio percentage associated with The Huntsman: Winter's War, which was viewed just 19,465 times. (The same basic trends and

numbers result from even a straight average, but in our view a weighted average is a more appropriate representation of the level of filtering experienced by viewers.)

| Comparative Levels of Filtering: Audio | | | | Comparative Levels of Filtering: Video | | | |
|---|--------|--------|--------|---|--------|--------|--------|
| | 90% | 75% | 50% | | 90% | 75% | 50% |
| Maximum | 17.36% | 24.30% | 34.95% | Maximum | 1.52% | 3.35% | 9.55% |
| 95th Percentile | 18.65% | 25.68% | 37.15% | 95th Percentile | 7.99% | 12.44% | 22.88% |
| 90th Percentile | 21.58% | 29.00% | 41.34% | 90th Percentile | 13.93% | 19.39% | 30.99% |

Figure 10: The demand for filtered audio and video as defined by various baselines.

These percentages reflect what we interpret to be a substantial demand for both audio and video filtering across a broad range of Hollywood films. Focusing, for example, on the third row and third column in each chart, we see that in roughly 40% of the streams users voluntarily filtered at least half as much audio content as did the more cautious viewer whose filtering decisions put that user at the 90th percentile for audio, and in roughly 30% of the streams users voluntarily filtered half as much video content as did the more cautious viewer whose filtering decisions put that user at the 90th percentile for video. Depending on the baseline chosen, the various percentages move up or down. And different counting decisions would similarly cause modest fluctuations in nearly every number.³⁷ But, in our view, these numbers powerfully reject the hypothesis that VidAngel customers were using the system solely because the movies were cheap and accessible, and powerfully support the contention that there exists a substantial market for filtered films.

³⁷ For example, as noted *supra* note 28, the data to which we had access did not include information about the exact location within each movie where each filter applied. VidAngel understandably viewed that information as proprietary because, with that data, anyone could fully reverse-engineer VidAngel’s movie database. So, while we would know that a given filter muted 1.5 seconds of audio in a specific film, we would not know exactly which 1.5 seconds were quieted. Because of that limitation, we could not account in our estimates for the possibility that two filters might both remove the same portion of a film and that, hence, counting both filters would inadvertently double-count their impact. Were we able to account for this issue, nearly every number in Figure 10 would change, although likely only modestly. Specifically, the baseline estimates would drop because the likelihood of overlap increases as the number of chosen filters increases, and those high-end baselines are therefore the most vulnerable to accidental overstatement. The reported percentages would as a result then also increase, however, because, while all of our estimates would drop, the drops would be more significant for the high-filtering baselines than for the lower-filtering data. The entire effect would likely be small, however, for two reasons. First, we report results for audio and video separately, thereby eliminating what is likely the most common type of overlap. Second, overlap seems unlikely for most streams in our dataset, given the relatively modest number of filters triggered in each stream.

V. What Users Choose to Filter

The prior section characterized the market for filtered films by looking at the degree to which VidAngel users activated filters while streaming movies in 2016. This section uses the VidAngel data for a slightly different purpose: to understand the types of audio and video content that users chose to mute and excise.

At a high level, VidAngel users apparently filtered the same sorts of materials that the MPAA looks at in deciding whether to rate movies as G, PG, PG-13 or R. The MPAA does not make public the exact list of considerations upon which it relies, and the criteria are notoriously of the “I know it when I see it” variety regardless.³⁸ But, when we looked for correlations between each film’s MPAA rating and the extent to which users filtered that film, we found that the intuitive patterns held. Roughly speaking, the higher the MPAA rating, the more content VidAngel’s users filtered.

Consider in this light Figure 11. To generate the figure, we took the data about user streams, sorted by movie, and then calculated two averages: the average percentage of audio muted in that movie, and the average percentage of video cut. We then grouped movies based on their MPAA rating and calculated averages within each rating category. As with earlier calculations, for this step we again used a weighted average, giving more-watched movies appropriately greater weight in the aggregate statistics. The result was an unsurprising alignment between increased MPAA ratings and increased user-driven filtering.

| | Audio | Video |
|-------|-------|-------|
| G | 0.09% | 0.65% |
| PG | 0.19% | 0.88% |
| PG-13 | 0.38% | 1.33% |
| R | 1.12% | 2.81% |

Figure 11: Generally speaking, the higher the MPAA rating, the more audio and video filtered.

The trend here does call into question an interesting counter-hypothesis: namely, that the amount of material filtered would not increase with increases in MPAA ratings because users who watch movies with higher ratings are also more willing to experience language and/or watch intimate and violent content. Put differently, an increase in the MPAA rating associated with a given stream has two plausible implications. First, the higher the rating, the more potentially objectionable content likely present in the stream originally. This would lead to greater amounts of filtering, all else held equal. Second, viewers who choose to watch

³⁸ See [This Film Is Not Yet Rated](#) (a 2006 documentary about the MPAA rating system); *This editorial is not yet rated*, LOS ANGELES TIMES (October 14, 2006) (inspired by that documentary, describing and criticizing the MPAA rating system); Kirby Dick, *Rated R for ridiculous*, LOS ANGELES TIMES (Jan. 24, 2007) (same); Steven Zeitchik & John Horn, *More MPAA film ratings appealed*, LOS ANGELES TIMES (March 23, 2012) (same).

movies that have higher ratings might have a higher tolerance for potentially uncomfortable content. That would lead to lower amounts of filtering, all else held equal. Figure 10 suggests that the first consideration dominates, as it would if (for example) filtering is primarily used to remove certain types of content that a given viewer finds objectionable regardless of the movie in which it appears.³⁹

Another way to understand what VidAngel users filtered is to count the number of times each filter was selected for use. To do that, we first took the list of every stream triggered by a user in 2016 and again eliminated those that the requesting user watched for ten minutes or less. As explained previously, we eliminated these short streams because our review of the data suggested that many of these streams represented situations where a user was testing the system, rather than actually watching a movie, and hence that user's selections might not accurately reflect his or her content preferences. Using the roughly 3.4 million streams that remained, we next counted the number of times each filter was selected by a user, focusing on VidAngel's internal label for the category (e.g., "audio, damn"; "video, blowjob") rather than the longer textual descriptions (such as "man touches woman inappropriately"). We focused on the labels because the textual descriptions are often movie-specific and hence not amenable to summation, whereas category labels tended to repeat across movies.

Totals for all 85 available labels are shown in Figure 12, with audio and video filters listed separately and the results sorted by magnitude. Note that this list includes filters, such as the "Jar Jar Binks" filters and the various credits filters, that we did not otherwise use in our analysis. We report them here for completeness, but we excluded them elsewhere because they do not necessarily target violence, language, sex or other classically objectionable content. Typographical mistakes in this list are intentional. In VidAngel's internal data, for example, the word "nigger" was intentionally misspelled as "niger," apparently in an effort to be sensitive to how that word might be perceived by even VidAngel's own internal team. Finally, in addition to reporting the number of times each filter was triggered, we also report a percentage that

³⁹ To disentangle these competing considerations, we pulled statistics for the potentially offensive word "shit" across all our data. For PG streams, the word "shit" turned out to be used on average 1.35 times per stream; and, when users were offered the chance to mute those utterances, they chose to do so 72.7 percent of the time. The word "shit" was used much more often in PG-13 streams, specifically showing up on average 5.91 times per stream. To hold constant the number of times "shit" was muted, then, viewers of PG-13 streams would have needed to be significantly more comfortable with that word, muting only 16.6 percent of the larger number. We found, however, that viewers in fact muted 51.1 percent of those utterances. The same pattern repeated as we moved from PG-13 to R. Streams of R movies on average include 18.54 utterances of "shit" per stream. To hold constant the number of mutings as between PG-13 and R films, viewers of R films would have needed to be so tolerant that they would have filtered only 16.3 percent of the larger number. In fact, viewers of R films chose to mute the word "shit" 42.2 percent of the time. Thus, while viewers are more tolerant of the word as we move from PG (72.7%) to PG-13 (51.1%) to R (42.6%), that increased tolerance is drowned out by the significantly more pronounced increase in the use of the word, which rose from 1.35 (PG), to 5.91 (PG-13), to 18.54 (R). Similar trends are seen for other content, including the words "God," "Jesus," and "f*ck."

represents how that number compares to the maximum number of times the filter could have been triggered had users selected it every time it was offered to them. For example, the audio filter “shit” was offered to viewers as a possible filter 27,745,832 times in 2016 and was selected 13,781,850 times; we thus report for “shit” a selection percentage of 45 percent.

| <u>Filter Type (Audio)</u> | <u>Total Times</u> | | <u>Filter Category (Video)</u> | <u>Total Times</u> | |
|-------------------------------|--------------------|----------------|--------------------------------|--------------------|----------------|
| | <u>Filtered</u> | <u>Percent</u> | | <u>Filtered</u> | <u>Percent</u> |
| fuck | 26,425,075 | 60% | immodesty_female | 7,194,069 | 12% |
| god | 20,318,645 | 49% | non-graphic | 5,158,855 | 4% |
| shit | 13,781,850 | 45% | graphic | 4,866,412 | 10% |
| damn | 6,827,097 | 45% | nudity_female | 2,115,812 | 32% |
| hell | 6,549,235 | 37% | violence_images | 2,055,007 | 8% |
| sexual_reference | 5,757,094 | 24% | sexually_suggestive | 1,706,272 | 25% |
| Jesus | 5,537,462 | 53% | negative_elements | 1,494,518 | 10% |
| language_crude | 5,495,167 | 12% | implied_nudity | 1,339,624 | 25% |
| ass | 5,068,875 | 41% | implied_not_shown | 1,274,004 | 34% |
| bitch | 3,036,410 | 46% | nudity_male | 1,267,503 | 35% |
| bastard | 755,080 | 45% | immodesty_male | 1,076,799 | 16% |
| racial_slurs | 732,138 | 14% | drugs_legal | 967,252 | 5% |
| other_language | 727,217 | 24% | kissing_normal | 502,874 | 7% |
| dick | 714,723 | 32% | language_captions | 489,561 | 28% |
| language_sexual_other | 654,680 | 30% | drugs_implied | 460,315 | 6% |
| language_vulgar | 577,182 | 26% | credits | 443,366 | 10% |
| piss | 450,693 | 28% | immodesty_both | 437,400 | 19% |
| british_profanity | 363,884 | 29% | kissing_passion | 418,955 | 15% |
| puss | 219,392 | 38% | drugs_illegal | 375,404 | 11% |
| balls | 208,714 | 32% | shown_w_o_nudity | 308,396 | 44% |
| cock | 205,119 | 35% | studio_logos | 301,066 | 7% |
| prick | 178,437 | 41% | nudity_both | 286,574 | 33% |
| crap | 161,572 | 22% | implied | 272,695 | 4% |
| niger | 123,556 | 26% | shown_w_nudity | 266,033 | 57% |
| cunt | 96,045 | 43% | jar_jar_binks | 185,796 | 8% |
| language_crude_other | 75,183 | 22% | nonconsensual | 142,669 | 44% |
| sex | 40,941 | 29% | NA | 83,504 | 9% |
| penis | 36,395 | 35% | kissing_homo_normal | 50,718 | 23% |
| horny | 23,127 | 37% | kissing_homo_passion | 15,064 | 28% |
| twat | 21,963 | 46% | nudity_w_o_sex | 1,798 | 14% |
| fag | 20,404 | 40% | climax_scene | 1,122 | 52% |
| pecker | 19,928 | 48% | immodesty | 616 | 5% |
| fart | 14,600 | 20% | short_film | 72 | 12% |
| porn | 12,659 | 36% | lyrical_song | 46 | 14% |
| language_discriminatory_other | 11,464 | 30% | kissing | 35 | 4% |
| jerk | 10,530 | 38% | | | |
| dildo | 6,998 | 41% | | | |
| blowjob | 6,321 | 40% | | | |
| cum | 4,212 | 37% | | | |
| douche | 3,811 | 45% | | | |
| testicles | 3,775 | 38% | | | |
| clit | 3,742 | 32% | | | |
| christ | 3,516 | 69% | | | |
| language_blasphemy_other | 3,032 | 41% | | | |
| jiz | 1,938 | 35% | | | |
| orgasm | 924 | 57% | | | |
| queers | 797 | 24% | | | |
| ejaculate | 635 | 27% | | | |
| dink | 377 | 68% | | | |
| gangbang | 173 | 30% | | | |

Figure 12: Filter counts using VidAngel's category labels, separated by audio and video, reported in terms of both absolutes and percentages.

As the figure shows, audio filters were triggered much more often than video filters, with each of the three most popular audio filters being triggered more than twice as often as was the most popular video filter. The audio filter targeting the word “f*ck” tops the audio list in terms of the absolute number of times it was chosen (roughly 26 million), but, interestingly, users were most likely to mute the words “Christ” and “dink” if they were uttered in a film. Specifically, “Christ” was muted 69 percent of the times it was uttered, and “dink” was muted 68 percent of the time. Of the more common words used in films, however, “f*ck” was the most often filtered (60 percent), with “Jesus” (53 percent) and “God” (43 percent) close behind.

On the video side, meanwhile, users most often cut scenes involving female immodesty (roughly 7 million cuts) but viewers most consistently cut scenes involving nudity (57 percent), sexual climax (52 percent), and nonconsensual sex (44 percent). Filters related to violence, by contrast, were rarely selected. The filter for “violent images” was selected only 8 percent of the times it was offered, for example, and the filters for “graphic” violence and “non-graphic” violence similarly clocked in at just 10 and 4 percent, respectively. When scenes were specifically labeled by gender, viewers reacted to male and female nudity at comparable rates, with the men being cut 35 percent of the time and the women 32 percent. Interestingly, scenes labelled as involving homosexual activity were in the aggregate slightly less controversial than male and female nudity, with “kissing_homo_normal” being cut 23 percent of the time and “kissing_homo_passion” being filtered at a rate of 28 percent. Obviously, however, these numbers do not account for the possibility that viewers uncomfortable with homosexual activity simply avoided the relevant films entirely.⁴⁰

As for the filters that we typically excluded from our work, those turned out to be relatively unpopular anyway. Filters for a film’s opening credits, closing credits, and studio logos were all chosen at rates of under 10 percent. And, somewhat surprisingly, only 8 percent of the scenes involving Jar Jar Binks were cut by VidAngel’s users. Somewhere out there, George Lucas is smiling.

Counts like these are informative but they suffer an important constraint: the numbers and percentages turn heavily on VidAngel’s decisions about the level of abstraction at which to define categories. For Deadpool, for instance, the lowest level of VidAngel’s hierarchy separately labelled audio use of the words “cock” and “dick,” and also had a separate category for “language_crude” that primarily included yet other examples of (how shall we say?)

⁴⁰ Although we do not discuss it in the text, it is jarring to see how much potentially inappropriate content is incorporated into the average Hollywood film. For example, in the 2,914 movies in the VidAngel movie database, the word “shit” is uttered 24,783 times, which means that there are 8.5 such utterances per film on average, a number that breaks down to 18.54 per R-rated film, 5.91 per PG-13 film, 1.35 per PG film, and 0.05 per G-rated film. Scenes that incorporate graphic violence number 31,238 overall in the data. That’s 10.72 per film in general, or 2.69 on average for G films, 5.26 for PG films, 8.27 for PG-13 films, and 19.32 for R-rated films. We focus in the text on the question of what exactly users choose to filter out of Hollywood films, but our data raise a separate set of questions about what Hollywood chooses to put into its films in the first place.

anatomical slang. Streams of that movie thus pushed up the counts for “cock” and “dick”; had VidAngel instead chosen to lump both of those words under the more general label “language_crude,” by contrast, streams of Deadpool would have moved “language_crude” up the chart at the expense of the “cock” and “dick” counts. Counting the number of times each label was selected thus gives some information about the types of material VidAngel users targeted, but the results also unavoidably reflect VidAngel’s judgment about how best to categorize content.

To mitigate this concern, we reorganized VidAngel’s separate filter categories into six larger buckets: audio filters related to sex, to violence, and to language; and video filters related to sex, to violence, and to language. Some labels were included in more than one of the six buckets. For instance, we characterized VidAngel’s label “audio, blowjob” as a filter relevant to both audio sex and audio language, and we characterized VidAngel’s “video, nonconsensual [sex]” as relevant to both video violence and video sex. By the same token, some of VidAngel’s labels fit none of the categories. For example, the label “video, drugs_implied” did not seem to be relevant to sex, violence, or language, even though it is easy to imagine a viewer wanting to filter out that content. When interpreting labels, meanwhile, we read VidAngel’s textual descriptions in order to make sure we were construing ambiguous labels accurately. Thus, for instance, we looked at specific movies in the database to determine that “video, implied” in fact meant “video, implied *violence*” and not “video, implied *immodesty*.”

Figure 13 summarizes the results. In the first column, we report statistics across all streams in our dataset, showing (a) the absolute number of filters triggered in each audio and video category; (b), in parentheses, per-stream averages for that category; and (c) a percentage measure, equivalent to the one shown in Figure 12, that shows how those numbers compare to the maximum number of times the filter could have been triggered had users selected it every time it was offered. For example, we report that users in 2016 triggered over 37 million audio filters related to sex, which works out to a per-stream average of roughly 9.7 and approximately 47 percent of the number of audio filters related to sex that users could have triggered. Similarly, we report that users triggered over 20 million video filters related to sex, which corresponds to a 5.4 per-stream average and a selection percentage of almost 18 percent. The next four columns report comparable information but according to each stream’s MPAA rating. Comparisons between MPAA ratings are best made using the per-stream numbers and the percentage scores, but not the absolute numbers, because the absolute numbers are influenced both by filter selection and by the number of streams associated with that rating. The per-stream and percentage measures, by contrast, are scaled; the per-stream measure reports the number of filters chosen divided by the number of streams in that category, and the percentage measure reports the number of filters chosen divided by the number of filters available across all the relevant streams.

| Audio | | All Movies | | G Movies | | PG Movies | | PG-13 Movies | | R Movies | | | | | |
|--------------|------------|-------------------|-------|-----------------|-------|------------------|-----------|---------------------|-------|-----------------|--------|-------|------------|--------|-------|
| Filters | | | | | | | | | | | | | | | |
| Sex | 37,189,746 | (9.7) | 47.4% | 67,126 | (0.4) | 19.8% | 984,059 | (1.0) | 31.0% | 6,434,371 | (3.7) | 37.2% | 29,704,193 | (34.7) | 51.6% |
| Violence | 1775 | (0.0) | 12.8% | 151 | (0.0) | 3.1% | 347 | (0.0) | 21.1% | 1,031 | (0.0) | 15.4% | 246 | (0.0) | 35.3% |
| Language | 91,310,037 | (23.9) | 41.1% | 509,036 | (2.8) | 15.6% | 7,666,670 | (7.5) | 25.6% | 32,735,959 | (18.6) | 40.1% | 50,398,382 | (58.8) | 46.9% |
| Other | 5,533,209 | (1.5) | 51.8% | 39,271 | (0.2) | 38.7% | 414,704 | (0.4) | 44.5% | 2,359,272 | (1.3) | 55.1% | 2,719,962 | (3.2) | 50.7% |
| Video | | All Movies | | G Movies | | PG Movies | | PG-13 Movies | | R Movies | | | | | |
| Filters | | | | | | | | | | | | | | | |
| Sex | 20,731,125 | (5.4) | 17.7% | 333,626 | (1.8) | 12.4% | 3,270,455 | (3.2) | 17.6% | 9,219,578 | (5.2) | 15.5% | 7,907,466 | (9.2) | 21.6% |
| Violence | 12,184,430 | (3.2) | 5.8% | 341,318 | (1.9) | 5.7% | 2,046,678 | (2.0) | 3.8% | 4,827,479 | (2.7) | 4.8% | 4,968,954 | (5.8) | 10.5% |
| Language | 30,81,104 | (0.8) | 27.4% | 36,289 | (0.2) | 12.6% | 328,163 | (0.3) | 12.8% | 1,019,702 | (0.6) | 26.5% | 1,696,950 | (2.0) | 37.1% |
| Other | 4,242,663 | (1.1) | 8.3% | 185,782 | (1.0) | 7.3% | 1,390,242 | (1.4) | 7.6% | 1,704,930 | (1.0) | 7.4% | 961,708 | (1.1) | 13.1% |

Figure 13: Filter use, by category and MPAA rating.

Start with the absolute numbers. Across all movie ratings, users primarily selected audio filters that target language. For PG-13 movies, for example, users selected over 32 million audio filters related to language, but fewer than 7 million audio filters related to sex and barely 1 thousand related to violence. Across all movie ratings except for films rated G, users primarily selected video filters that target content related to sex. For PG-13 movies, the counts this time put video sex above 9 million, video violence just below 5 million, and everything else below 3 million in total. G movies broke the pattern in that video filters related to violence slightly outnumbered video filters related to sex, but that result is likely driven by the fact that both sex and violence are rare in G-rated movies. Indeed, the shocking thing about the data related to G films is that there *is* data in the first place. Yet, on closer inspection, G films do regularly include violence (Mufasa dies a horrible on-screen death in The Lion King) and at least hint at sexuality (when streaming Aladdin, VidAngel quite accurately warns viewers that, in many scenes, “scantly clad women are seen”). Older films especially contain surprising content, like, for example, scenes in both Pinocchio and One Hundred and One Dalmatians where cartoon protagonists smoke cigars or pipes.

Turning to the per-stream averages, we find that these averages generally increase with increases in the MPAA rating. So, for example, the average G movie was streamed with 2.8 audio language filters selected, and that number grew to 7.5 for PG, 18.6 for PG-13 and ultimately 58.8 for R. This trend captures the interaction of two factors we identified earlier in the Article: movies with high MPAA ratings likely have more potentially objectionable content, which explains the upward trajectory; but they are likely watched by viewers who have more tolerance for uncomfortable content, which would at least slow and, in theory, could have reversed the rise. Looking at the numbers themselves, meanwhile, we see that per-stream edits tend to be small in number across all ratings and categories except for audio language filters applicable to R movies. There, viewers are watching, but they are muting an average of nearly 60 audio snippets per stream.

Lastly, as for selection percentages, perhaps the most shocking numbers in the chart are the selection percentages reported for video violence. Users, it seems, were interested in filtering only a tiny percentage of the violent scenes that VidAngel flagged as potentially inappropriate, with viewers opting to filter only 5.7 percent of the violent scenes in G movies, 3.8 percent in PG movies, 4.8 percent in PG-13 movies, and 10.5 percent in R movies. These percentages are noticeably lower than the comparable percentages for video filters related to sex, and even more noticeably lower than the comparable percentages for audio filters related to language.

VI. How Users Perceive Filtered Films

In the preceding two sections, we told two basic stories. We first used the VidAngel data to show that there does in fact exist a market for filtered Hollywood films, even though the

studios have historically resisted serving that market. We then offered some information about what exactly users seem to want to filter, finding (for example) that users focus more on visuals related to sex than they do on visuals related to violence; and that, even in the hands of the most aggressive users, filters are used modestly and with discretion, targeting specific types of content and in the aggregate excluding only small amounts of material. Here, we turn to one final question: do users who filter films enjoy those films to a similar extent as do viewers who opt instead to watch the corresponding unfiltered originals?

As we explained in Section II, this is a critically important question because Hollywood stakeholders might reasonably worry that viewers of filtered films will have a bad experience. Filters might remove scenes or dialogue critical to a film's story. Filters might mar the movie-watching experience by introducing unexpected silence and choppy video transitions. And filters might contribute to what we have described as a mismatch effect, where viewers make bad choices in essence because they overestimate the degree to which filters can transform a film. These concerns are particularly plausible as applied to the filtering facilitated by VidAngel, because VidAngel gave users great discretion to choose what and how much to filter. Users thus could have easily caused all of these problems for themselves, inadvertently trying to enjoy Saved! without any of its religious vocabulary, or Se7en with all of its violent scenes abandoned to the cutting room floor.

To quantify these potential harms, we embraced what we think is a reasonable hypothesis: viewers are likely to watch more of a film they are enjoying, and to cut short their time with a film they dislike. If that is true, then by looking at the extent to which users watch the films they filter, we can evaluate how filtered films are perceived. Specifically, if viewers tend to turn off filtered versions of a film noticeably sooner than viewers turn off the unfiltered version of that same film, that would be evidence consistent with a theory of reputational harm. By contrast, if viewers tend to watch filtered films to a similar or greater degree as compared to viewers who watch unfiltered originals, that evidence would suggest that the above-described harms are theoretically plausible but not in fact experienced by viewers.

Figure 14 answers the question. The first row reports that, on average, for the 1.689 million streams for which users triggered between 1 and 10 filters, VidAngel's viewers watched roughly 88.9 percent of each filtered movie. The second row shows that, on average, for the 416,000 streams for which users triggered between 11 and 20 filters, viewer similarly watched roughly 88.8 percent of the resulting filtered films. Later rows report the results for larger numbers of filters, but the viewership percentage remains nearly constant. Streams involving more than 100 filters were on average viewed to roughly the same degree as were streams involving fewer than 10 filters, and, across the entire dataset, roughly 90% of all streams fell in one of the buckets where the viewership percentage was between 87.2% and 88.9%. Viewers, it seems, enjoyed their filtered films.

| Number of Filters | Percent Viewed | Thousands of Streams |
|------------------------------|---------------------------|---------------------------------|
| 1-10 | 88.9% | 1,689 |
| 11-20 | 88.8% | 416 |
| 21-30 | 88.5% | 309 |
| 31-40 | 88.4% | 229 |
| 41-50 | 88.1% | 187 |
| 51-100 | 87.2% | 510 |
| 101+ | 84.8% | 380 |

Figure 14: Viewers seem to watch roughly the same percentage of a film regardless of how much material they filter from it.

The aggregate numbers reported in Figure 14 incorporate two substantial simplifications that more sophisticated analysis can disentangle. First, the numbers in Figure 14 treat all filters the same, even though some filters mute or excise larger portions of a film. Second, the numbers ignore the fact that viewers sometimes turn off a movie for reasons other than their displeasure with filtering. For instance, some movies are terribly long (Gone with the Wind); some are simply unpleasant in even their original form (The Emoji Movie); and some movies (Finding Dory) cater to younger audiences whose members may not always have sufficient patience to watch an entire film from start to finish regardless of its content.

We can account for these differences by using statistical models that include “fixed effects” for each individual movie. As readers with statistical backgrounds will know, fixed effects in this context are in essence movie-specific baselines against which filtered streams of each particular movie can then be compared. Simplifying a bit, if viewers in general tend to watch only half of Gone with the Wind regardless of the number of filters engaged, fixed effects adjust for that reality by ignoring the half-a-movie baseline and focusing instead on variation above and beyond that default level.

We report key findings in Figure 15. In this regression model, we include movie-specific fixed effects, along with measures of how much of the stream was filtered across each of eight categories: audio and visual filtering for sex, violence, language and “other.” To allow for the possibility of non-linear and non-monotonic filtering effects, our model includes not just these raw measures of filter impact, but also each measure squared and cubed. Ultimately, for five of the eight categories, moving from zero filtering to the 90th percentile of filtering gave a predicted marginal difference of less than 1% in terms of how much of the stream would be watched. The other three categories, however, showed slightly more variation, as shown below. In each graph, the x-axis captures the percentage of material filtered as compared to the total runtime of the film, and the y-axis captures the percentage of runtime actually watched by the relevant user.

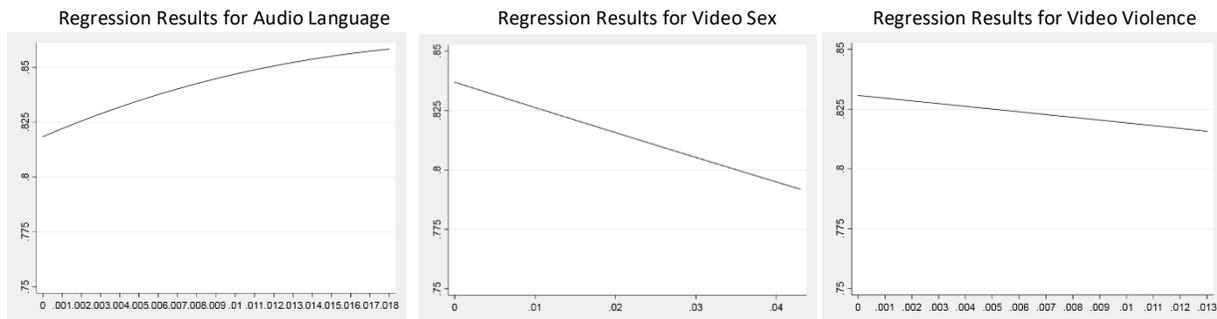


Figure 15: Viewership trends as a function of the amount of audio language, video sex, and video violence filtered.

Two details bear mention. First, as is clear from comparisons across the three charts, viewing time is most at risk when video filters excise content related to sex. There are many possible explanations for why this type of filter might have the most pronounced impact. One possibility is that the absence of visual content related to sex particularly interferes with storytelling, disrupting the communication of key plot points and interfering with character development to a greater degree than does muting audio language or cutting video violence. Another possibility is that visual content related to sex is a particularly enjoyable aspect of the films in which it appears (is there really anything else redeeming about *Fifty Shades of Grey*?) and hence removing these scenes has a comparatively larger impact on viewer satisfaction. A third possibility is that movies with significant content related to sex are particularly vulnerable to the mismatch effect. That is, viewers might think that their objections apply only to scenes where intimacy is readily apparent, but when watching films that originally included intimate visuals, viewers might realize that their objections in fact extend to the adult themes and mature dialog that likely pervade those films. Language and violence, by contrast, might be easier to isolate from other aspects of a film.

Second, viewing time seems to increase as the amount of language filtered increases. This is puzzling in that it implies that filtering actually increases enjoyment, but the trend might be explained by three observations. First, audio filters that target language might not have much impact on a film's story. A viewer probably does not need to know exactly which profanity John Goodman is yelling in order to appreciate the scene in *The Big Lebowski* where he trashes his adversary's car, nor does a viewer necessarily need to hear any of the vulgar words in *The Wolf of Wall Street* in order to understand the various characters' emotions, motivations, and personalities. Second, audio filters that target language might not significantly interfere with the presentation of a movie either. Visuals are not at all interrupted by audio filters. And audio disruptions tend to be quick, with the average length of an audio filter in our dataset clocking in at barely one second in length. Filters applicable to visual violence and visual sex, by contrast, can be much more disruptive, breaking the video's visual

continuity and lasting an average of 7 and 9 seconds, respectively. Third, and the factor that might be driving viewing time up: viewers who take the time to choose specific filters might also have a particularly strong interest in seeing the associated film. For visual violence and visual sex, this factor might be outweighed by the disruptions caused to the film's storytelling and presentation. But, if audio language does not much impact either storytelling or presentation, this correlation between intensity of filtering and intensity of interest might explain the unique upward slope of this one line.

VII. Implications

We began this Article by telling the stories of a few early filtering technologies, and we used those stories to situate our empirical work in the context of what has been a decades-long debate about who should decide whether and how films can be filtered at home. We intentionally held in abeyance, however, the most recent and most telling vignette. In June 2017, Sony Pictures announced that the studio itself was ready to sell to the public edited versions of some of its most popular films.⁴¹ The details of the program were all going to be controlled by Sony, a company whose direct financial ties to viewers, writers, directors, producers, and actors presumably gave it a strong incentive to balance every competing interest when deciding which films to filter and how. Yet, Hollywood responded to even this program with breathtaking hostility.

The Directors Guild of America spoke up within hours. "Throughout the years," the industry group intoned in a harshly-worded press release, "the DGA has achieved hard-fought creative rights gains protecting our members from such practices. As creators of their films, directors often dedicate years of hard work to realize their full vision, and they rightfully have a vested interest in protecting that work. We are committed to vigorously defending against the unauthorized alteration of films."⁴² Director Judd Apatow was more pointed, writing on Twitter, "This is absolute bullshit and @sony and @SonyPictures is gonna get hell for FUCKING with our movies. Shove the clean versions up your asses!"⁴³ The actor/director Seth Rogen added simply, "Holy shit please don't do this to our movies."⁴⁴

Sony understandably relented. Within a week, division president Man Jit Singh sheepishly affirmed that directors "are of paramount importance to us" and promised that "if

⁴¹ Pamela McClintock, *Sony's Sanitized Movie Initiative Faces Growing Opposition*, THE HOLLYWOOD REPORTER (June 13, 2017) (quoting a Sony executive describing the program as "a pilot program . . . that offers viewers the option of watching an airline or TV version of certain movies when they purchase the original version").

⁴² Josh Rottenberg, *Judd Apatow and the Directors Guild blast Sony plan to release 'clean' versions of films*, THE LOS ANGELES TIMES (June 14, 2017) (quoting the press release).

⁴³ Dave McNary, *Judd Apatow to Sony: 'Shove the Clean Versions of Movies Up Your A—!*', VARIETY (June 13, 2017) (quoting Judd Apatow).

⁴⁴ McClintock, *supra* note 41 (quoting Seth Rogen).

any of them are unhappy” with the new program, Sony will “discontinue it for their films.”⁴⁵ And, since then, not a single film has been released as part of the “Clean Version” program, and its once-live URL (www.cleanversionmovies.com) today points nowhere.

Why is this an important vignette from our perspective?

But for this real-world evidence, we might have been tempted to end this Article by arguing that, when it comes to filtering, neither courts nor Congress need force Hollywood’s hand. Viewers, our data suggest, want to filter Hollywood fare. Moreover, viewers then happily watch the resulting films, enjoying edited versions to roughly the same extent as do viewers of the corresponding unedited originals. According to our data, then, filtering really can be a win-win where viewers are able to participate in shared cultural experiences despite their individual religious, moral, and personal constraints, and Hollywood stakeholders can enjoy the artistic and financial benefits associated with having larger audiences for their movies. The law need only stay out of the way, it would seem. The demand for filtered films should be willingly met without court or legislative intervention.

And yet, clearly not. Why? Because directors, writers, actors, and other Hollywood stakeholders perceive filtering as an affront to artistic integrity.

The example we used in Section II about Schindler’s List, for instance, was drawn from real events: in 1994, Steven Spielberg refused to allow a movie theater to remove violence and nudity from the film even though he knew that patrons at that theater preferred to watch an edited version and, by denying them that chance, Spielberg was likely reducing the audience for his own film.⁴⁶ And, in 1985, when ABC proposed cutting nine minutes from the broadcast version of Warren Beatty’s three-and-a-half hour Oscar-winning film Reds, Beatty not only scuttled the deal, but also sued the affiliated movie studio and launched a public relations campaign that Martin Scorsese would later describe as an attempt to “lay the foundation for a philosophical change where the emphasis is off economics and onto art.”⁴⁷ The tweets we quote from Judd Apatow and Seth Rogen lack Scorsese’s eloquence, but they strike a similar chord. Hollywood is not objecting because of problems about film quality, distribution, or price. This is a fight about perceived artistic insult.

The implication for us is clear: copyright law should be interpreted so as to overrule Hollywood’s high-class objections and open the market for filtered films. We have two reasons. First, our data shows that any potential harm to artistic integrity is actually quite small. Viewers are not aiming to substantially change the films they otherwise want to watch. Quite the opposite, even viewers who want to filter are looking to trim on average less than 1% of the

⁴⁵ Rottenberg, *supra* note 42 (quoting Man Jit Singh).

⁴⁶ Rebecca Buckman, *A Cottage Industry in Utah Cleans Up Hollywood’s Act: Video Stores, Software Firms Cater to Mormon Tastes*, THE WALL STREET JOURNAL (Sept. 19, 2002) at A1.

⁴⁷ David T. Friendly, *Movie Directors Versus TV Editing*, The Los Angeles Times (April 29, 1985).

audio and less than 2% of the video from the average film. Spider-Man would still save the day at high personal cost in the typical edited version of Sam Raimi's Spider-Man. Adam Sandler and Drew Barrymore would still touchingly and repetitively fall in love in almost every filtered version of 50 First Dates. And, yes, viewers who opt to watch filtered versions of Superbad will turn out to enjoy almost exactly the comedic experience that Seth Rogan originally intended for them. Deep breath, everyone. Schindler's List is an outlier. The artistic stakes are low when the movies at issue are Step Brothers and Talladega Nights: The Ballad of Ricky Bobby, the two Apatow films that Sony had planned to filter as part of its now-abandoned family-friendly offering.⁴⁸

Second, even where the harm is larger, the right to prevent unauthorized alternations is simply not a right that copyright law was meant to recognize. European copyright systems have long enforced "moral rights" of this sort, including not only an explicit right to prevent unauthorized alteration, but also related rights, like the right to demand attribution, and the right to withdraw a work from public view.⁴⁹ But, for 100 years, the United States went so far as to refuse to join the international Berne Convention on Copyright because that agreement required its signatories to provide protection for moral rights, and American negotiators saw those rights as inconsistent with the American approach.⁵⁰ Even today, the closest federal law comes to recognizing a right to prevent alteration is the Visual Artists Rights Act of 1990; and, while that statute does give authors the right "to prevent any intentional distortion, mutilation, or other modification . . . which would be prejudicial to his or her honor or reputation," the protection applies only to works that exist "in a limited edition of 200 copies or fewer" and, at that, are signed by the relevant author and numbered.⁵¹ Hollywood films do not remotely fit the bill.

All that said, we are not unsympathetic to Hollywood's intuitive concern. Spielberg, Rogan and Apatow probably should publicly beseech their fans to watch Hollywood films in pristine form. They likewise probably should petition high schools across the country to use movies like Schindler's List and Brokeback Mountain to facilitate the learning that takes place when people are exposed to words, images, and ideas with which they are initially

⁴⁸ See Matt Goldberg, *Sony to Provide 'Clean Versions' of Its Movies on Home Entertainment*, COLLIDER (June 6, 2017) (listing the films initially planned for edited release).

⁴⁹ See Henry Hansmann & Marina Santilli, *Authors' and Artists' Moral Rights: A Comparative Legal and Economic Analysis*, 26 J. LEGAL STUD. 95 (1997) (discussing the European approach); Roberta Kwall, *Copyright and the Moral Right: Is an American Marriage Possible?*, 38 VAND. L.REV. 1 (1985) (same).

⁵⁰ Hansmann, *supra* note 50, at 96-97 ("For more than 100 years the United States refused to sign the Berne Convention, in part because of objections to the moral rights clause."). See also Kwall, *supra* note 50 (highlighting inconsistencies between the American and European approaches).

⁵¹ 17 U.S.C. § 101 (defining a "work of visual art"); § 106A (creating the above-quoted rights).

uncomfortable.⁵² Spielberg, Rogen and Apatow perhaps should even petition Congress to expand copyright protection to include moral rights. But, in our view, Hollywood stakeholders should be forced to wage and win those fights on the merits, rather than using copyright law to take away tools and technologies that would otherwise empower viewers to honor their own religious, moral, and policy convictions.

⁵² Society admittedly has a strong interest in finding ways to ensure that each of us is exposed to a wide variety of conflicting perspectives, an interest that has grown only more urgent in these, society's Facebook years. Indeed, one of us wrote in 2005 that government should in certain situations promote the exchange of ideas specifically by *not* helping listeners filter out uncomfortable messages. See Douglas Lichtman, *How the Law Responds to Self-Help*, 1 J. LAW, ECON & PUBLIC POLICY 215 (2005) at 255-259. We depart from that principle here, however, because we doubt that this high-minded interest is meaningfully served through the medium of Hollywood films, shown in the home, in front of children. Put simply, it is one thing to forcibly expose people in a courthouse to a jacket bearing the inscription "Fuck the Draft" (*Cohen v. California*, 403 U.S. 15 (1971)), and quite another to pressure a person in his living room, possibly with kids within earshot, to listen to John Goodman yell about "fucking people up the ass." (*The Big Lebowski* (1998)). The home need not be sacrosanct, and children are more resilient than they might at first seem, but the combination of factors at issue here seems to weigh against a policy of take-it-or-leave-it film distribution. Compare *FCC v. Pacifica Foundation*, 438 U.S. 726 (1978) (upholding restrictions on the use of profanity in broadcast radio in part because radio invades the privacy of the home, and in part because radio is uniquely accessible to children); *Ashcroft v. ACLU*, 542 U.S. 656 (2004) (upholding a preliminary injunction against a federal statute regulating Internet speech in part on the ground that "filter software is an alternative that is less restrictive" than was the proposed regulation and "likely more effective as a means of restricting children's access to materials harmful to them").