DISCUSSION PAPER SERIES

DP14038

VIDEO KILLED THE RADIO STAR? ONLINE MUSIC VIDEOS AND RECORDED MUSIC SALES

Christian Peukert and Tobias Kretschmer

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

VIDEO KILLED THE RADIO STAR? ONLINE MUSIC VIDEOS AND RECORDED MUSIC SALES

Christian Peukert and Tobias Kretschmer

Discussion Paper DP14038
Published 06 October 2019
Submitted 25 September 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Christian Peukert and Tobias Kretschmer

VIDEO KILLED THE RADIO STAR? ONLINE MUSIC VIDEOS AND RECORDED MUSIC SALES

Abstract

We study the heterogeneous effects of online video platforms on the sales volume and sales distribution of recorded music. Identification comes from two natural experiments in Germany. In 2009, virtually all music videos were blocked from YouTube due to a legal dispute. In 2013, the dedicated platform VEVO entered the market, making videos of a large number of artists available overnight. Our estimates suggest that restricting (enabling) access to online videos decreases (increases) recorded music sales on average by about 5-10%. We show that the effect operates independently of the nature of video content, suggesting that user-generated content is as effective as official content. Moreover, we highlight heterogeneity in this effect: Online music video disproportionally benefits sales of new artists and sales of mainstream music.

JEL Classification: L15, L82

Keywords: Digital Distribution Platforms, User-generated content, Natural Experiment

Christian Peukert - christian.peukert@ucp.pt
Universita Catolica Portuguesa Lisbon

Tobias Kretschmer - t.kretschmer@lmu.de
University of Munich and CEPR

Acknowledgements
We would like to thank audiences at NBER Summer Institute, IIOC Chicago, Searle Internet Conference, VfS Industrieökonomischer Ausschuss, Toulouse Digital Economy Workshop, Florence Seminar on Media, INFORMS, Media Economics Workshop, ICT Conference Munich, ZEW ICT Conference, seminar participants at University of Zurich, LMU Munich, Toulouse School of Economics, ParisTech, USC Marshall, Rotterdam School of Management, Católica-Lisbon, Pompeu Fabre, Georgia Tech, TILEC Tilburg, IPTS Seville for helpful comments and feedback. We received very helpful comments from Joel Waldfogel, Lisa George, Christiaan Hogendorn, Julian Wright, Ulrich Kaiser, Miguel Godinho de Matos, Rahul Telang, Danielle Li, Johannes Koenen and Leon Zucchini. We thank Alexander Wolf (GEMA) and Tina Funk (VEVO) for insightful discussions, and financial support from FCT – Portuguese Foundation of Science and Technology for the projects UID/GES/00407/2013 and FCT-PTDC/EGE-OGE/2796/2017.

Video Killed the Radio Star? Online Music Videos and Recorded Music Sales

Tobias Kretschmer[1,2] and Christian Peukert[3,4]

[1]LMU Munich
[2]CEPR London
[3]UCP – Católica Lisbon School of Business and Economics
[4]ETH Zurich – Center for Law and Economics

revised version, July 13, 2019

Abstract

We study the heterogeneous effects of online video platforms on the sales volume and sales distribution of recorded music. Identification comes from two natural experiments in Germany. In 2009, virtually all music videos were blocked from YouTube due to a legal dispute. In 2013, the dedicated platform VEVO entered the market, making videos of a large number of artists available overnight. Our estimates suggest that restricting (enabling) access to online videos decreases (increases) recorded music sales on average by about 5-10%. We show that the effect operates independently of the nature of video content, suggesting that user-generated content is as effective as official content. Moreover, we highlight heterogeneity in this effect: Online music video disproportionally benefits sales of new artists and sales of mainstream music.

Keywords: Digital Distribution Platforms, User-Generated Content, Natural Experiment

1 Introduction

Digitization has brought important changes to the recorded music industry. Thanks to a substantial decline in the fixed cost of production, distribution and promotion, the number of new songs on the market has tripled (Aguiar and Waldfogel, 2016, 2017a). However, the precise way digital distribution and promotion affect artists' revenue streams (with implications for incentives to invest in new content in the long run) is not fully understood. We compare *open* and *closed* content distribution platforms in their effects on revenues in other distribution channels and highlight important heterogeneity in those effects.

Digital platforms distributing content differ in the amount of control rights holders have over available content and the corresponding revenue model. The spectrum ranges from *unlicensed platforms* like (the original) Napster or Megaupload, *licensed open platforms*, such as YouTube and Soundcloud that allow anyone to upload content, to *licensed closed platforms*, such as iTunes and Spotify, where only the rights holders themselves can contribute.

By definition, *unlicensed platforms* do not compensate rights holders. Indeed, the relatively large literature on digital piracy suggests that unlicensed consumption mostly substitutes for licensed consumption.[1] *Licensed platforms* typically pay some share of sales, advertising revenues, or per-use royalties. An emerging literature has started to look into *closed platform* streaming services and typically finds that services like Spotify substitute for piracy and digital ownership, but leave aggregate revenues largely unchanged (e.g. Thomes, 2013; Nguyen et al., 2014; Wlömert and Papies, 2016; Aguiar and Waldfogel, 2017b; Datta et al., 2017). Surprisingly little is known about *licensed open platforms*. They are particularly interesting because they create a market for derivative works. Rights holders are compensated for third-party use of their work without requiring individual licensing contracts. For example, on YouTube, where user-generated content (UGC) comprises the majority of available content (Liikkanen and Salovaara, 2015), rights holders can get a share of advertising revenues generated from third-party content.

We study how the open platform YouTube affects sales of recorded music and compare its effect to that of VEVO, a closed platform. We then investigate how availability on YouTube affects the type and variety of music consumers demand on other channels (Piolatto and Schuett, 2012; Datta et al., 2017). Like many

[1] See for example Hui and Png (2003); Rob and Waldfogel (2006); Zentner (2006); Rob and Waldfogel (2007); Oberholzer-Gee and Strumpf (2007) and the survey in Liebowitz (2016), although recent evidence in Peukert et al. (2017) suggests that there is significant heterogeneity across content types.

other digital platforms, YouTube offers tools that allow for search and social interaction, and provides up-to-date lists of the most popular videos and automated recommendations to help consumers comb through the vast amount of content available on the platform (Zhou et al., 2016), although their role in shaping the popularity distribution is not clear *ex-ante* (Fleder and Hosanagar, 2009; Tucker and Zhang, 2011; Oestreicher-Singer and Sundararajan, 2012b; Susarla et al., 2012; Godinho de Matos et al., 2016).

A unique setting in the German market helps us establish a causal link between YouTube availability and other consumption channels for recorded music. Because of a royalty dispute between YouTube and the de-facto monopolist royalty collection society (GEMA) that represents artists and publishers (not record labels), YouTube blocked access to almost all videos containing music in Germany on April 1, 2009.[2] For example, 85% of the 689 music videos in the list of the 1000 most viewed videos globally were blocked in Germany, while the same content remained accessible in a vast majority of other countries. For instance, 99% of those videos could be accessed in the US.[3] This standoff persisted until a consortium of record labels negotiated its own deal with GEMA and launched the dedicated platform VEVO on October, 1, 2013, which in most other countries is simply a channel on YouTube.

We use scanner data on the weekly number of units sold on a physical medium and as a digital download for the 1000 highest-grossing songs on the German market and match them with rich meta-information at the release- and artist-level. In 2013, we also observe the weekly number of (free and paid) streams per song. For each song in our data, we collect information on whether a corresponding video was available on YouTube (or VEVO). We estimate a difference-in-differences model to compare sales of songs *with* videos to sales of songs *without*, four weeks before and four weeks after the natural experiment(s) we observe.

Our results provide strong evidence that YouTube is complementary to other licensed consumption channels, at least in the short run. Across a variety of different specifications, our most conservative estimates suggest that removing access to music videos on YouTube *reduces* total weekly sales by about 6% on average. This result is robust to a number of falsification exercises, including placebo tests and data from Austria, a country that shares language and cultural history with Germany, but was not affected by the blocking on YouTube. We show that the size of our estimated effect does not vary across songs

[2] See New York Times, 'Royalty Dispute Stops Music Videos in Germany', April 2, 2009, http://www.nytimes.com/2009/04/03/technology/internet/03youtube.html. More than 7 years later, an agreement was reached, see http://www.bbc.com/news/technology-37839038.

[3] See http://apps.opendatacity.de/gema-vs-youtube/en.

that have a higher share of user-generated videos on YouTube, suggesting that UGC has a very similar complementary effect as official content. Strikingly, weekly sales and streams *increase* by a similar amount when VEVO, a closed platform without user-generated content, enters the market four years later. We do not find evidence that suggests that the complementary effect is moderated by overall popularity as measured in sales. We then study heterogeneity across songs and find that YouTube disproportionately benefits sales of new artists and mass-market artists (mainstream music).

By studying the (un)availability of a *licensed open platform* for music consumption and its heterogeneous effect on demand on other consumption channels, we make several contributions. Specifically, we emphasize the effect of digital distribution platforms on product discovery rather than just their substitution of paid channels. Demonstrating that mainstream artists and new artists benefit disproportionally from digital distribution platforms also adds to the literature on the effect of digitization on the sales distribution ("long tail", Brynjolfsson et al., 2010, 2011). Further, our results may be informative in the context of cultural trade policy (Hervas-Drane and Noam, 2017) and the debate on the reform of the compulsory licensing rules of interactive digital services (Lenard and White, 2015).

2 Related work and research framework

2.1 Channel competition: displacement or promotion?

A growing literature seems to have established that unlicensed consumption (digital piracy) harms sales of licensed products (e.g. Bhattacharjee et al., 2007; Adermon and Liang, 2014; Danaher et al., 2014), with substantial variation in estimated displacement rates (Hui and Png, 2003; Rob and Waldfogel, 2006; Zentner, 2006; Rob and Waldfogel, 2007; Oberholzer-Gee and Strumpf, 2007; Liebowitz, 2016). Recent empirical evidence on digital piracy (Peukert et al., 2017) supports theoretical work arguing that unpaid consumption can increase sales in licensed channels due to demand externalities (Takeyama, 1994) or because consumers can sample (vertical or horizontal) product quality at zero marginal cost (Peitz and Waelbroeck, 2006). The growing adoption of licensed services, technological advances and business model innovations have changed how and at which cost consumers access digital goods. Theoretically, (low-priced) licensed online offerings can combat piracy (Thomes, 2013), which is empirically supported in several settings (Danaher et al., 2010; Papies et al., 2011; Poort and Weda, 2015; Aguiar and Waldfogel, 2017b). For example, Zhang (2017) shows that making licensed content more usable (by removing digital rights management restrictions) increases sales of that content. Studies show that streaming services are associated with lower sales in conventional channels (Hiller, 2016; Wlömert and Papies, 2016; Aguiar and Waldfogel, 2017b; Datta et al., 2017), but lead to higher per-capita consumption (Datta et al., 2017), so that aggregate revenues of the music industry remain more or less unchanged (Wlömert and Papies, 2016; Aguiar and Waldfogel, 2017b). However, some studies also suggest complementary effects of streaming. For example, Nguyen et al. (2014) find that streaming services do not affect physical purchases, but increase attendance of live concerts. Online video in particular has been shown to indirectly increase sales of complementary products in the context of fashion retail (Kumar and Tan, 2015). For music however, there may be a more direct effect on sales. In October 2008, YouTube rolled out "click-to-buy" links to iTunes and Amazon next to music videos. Rights holders can choose to add such links to their own uploads and to user-generated videos for which the rights holder is identified through YouTube's Content ID system.[4] Given the mixed prior evidence and the specific institutional details therefore, the average effect of YouTube on sales of individual songs remains an empirical question.[5]

2.2 User-generated content

YouTube is an open licensed platform, meaning that anyone can provide content to the platform, and the platform has taken steps for rights holders to be compensated for the use of their content. Therefore, YouTube content linked to a particular song is usually a combination of official videos and user-generated content, which may affect its impact on recorded music sales.

YouTube is an immensely popular music platform. More than 80% of YouTube visitors use it for music (IFPI, 2016), and video-based streaming accounts for more than 50% of the 317.3 billion music streams in the US in 2015.[6] About 30% of YouTube videos (and 40% of total views) are music videos (Liikkanen and Salovaara, 2015). YouTube is not heavily invested in the creation of own music content, but provides financial incentives to third-parties to upload content (Tang et al., 2012). Consistent with evidence that ad revenue sharing and content commercialization shift incentives towards creating more mass-oriented

[4] See https://youtube.googleblog.com/2008/10/like-what-you-see-then-click-to-buy-on.html.

[5] See section 6.3.1 for a discussion of a related paper (Hiller, 2016) that studies the relationship of YouTube and album sales.

[6] See https://www.musicbusinessworldwide.com/youtubes-dominance-takes-a-beating-from-spotify-apple-music-and-co/. YouTube's influence on the music industry eventually grew so large that Billboard started incorporating YouTube data into their rankings in February 2013, see http://www.billboard.com/biz/articles/news/1549766/billboard-charts-add-youtube-views.

content (Sun and Zhu, 2013), the introduction of the so-called partner program in 2007 quickly led major record labels to make their music video libraries available on YouTube.[7] However, most content on YouTube is user-generated. Consumers upload videos of live performances, videos that embed the lyrics of the song or derivative works such as cover versions and parodies (Liikkanen and Salovaara, 2015). Through YouTube's Content ID technology, all videos are matched to a database of copyrighted material. When uploads are classified as infringing, rights holders can choose between blocking the infringing video or "monetizing" the content by sharing revenue from advertisements. In testing, YouTube reported that 90% of claims created through Content ID led to rights holders choosing monetization.[8] While this gives substantial control to rights holders compared to closed platforms, it is more difficult for artists or labels to entirely withdraw content from YouTube because new user-generated content is constantly uploaded.[9] Overall, little is known about the relationship between original innovation and adaptations and derivative works. The literature has focused on related, yet different questions. A stream of work studies how user-generated knowledge affects firm innovation (see Gambardella et al., 2016 for a recent example). Empirical evidence on the role of intellectual property rights shows that IP protection (or stronger enforcement) hinders follow-on innovation, measured as subsequent academic knowledge creation (Williams, 2013), developer activity in open source software projects (Wen et al., 2013), and the number of patent citations (Galasso and Schankerman, 2015). Hence, user-generated content (especially adaptations and derivative work) may act complementary to sales of official content or consumers may perceive such content as a substitute. Regarding online music videos and record sales, this implies that the fact that YouTube is an open platform can shape its impact on recorded music sales, and it is not clear in which direction.

2.3 Digital distribution platforms and effects on consumption patterns

Internet-enabled distribution platforms often offer a variety of tools to aid content discovery. The reduction in search cost has the potential to affect the sales distribution (Anderson, 2006). A first example are aggregate top lists of the best-selling products in a given product category, week and geographical area. Aggregate popularity information (observational learning) can drive concentration (Salganik et al., 2006;

[7]See `https://techcrunch.com/2007/05/04/youtube-launches-revenue-sharing-partners-program-but-no-pre-rolls/`.
[8]See `https://googleblog.blogspot.com/2008/08/making-money-on-youtube-with-content-id.html`.
[9]The case of Taylor Swift withdrawing her songs from Spotify, but not YouTube is a case in point. See `https://www.theguardian.com/music/2016/jul/13/taylor-swift-youtube-music-royalties-battle`.

Sorensen, 2007; Cai et al., 2009; Hinz et al., 2011), but can also benefit niche content when the same level of popularity implies higher quality of niche vs. broad-appeal products (Tucker and Zhang, 2011). More personalized recommendations can come directly from peers (often encouraged by platforms making it easy to share content) or from algorithms using social data. For example, evidence shows that social interactions determine which videos on YouTube become successful (Susarla et al., 2012). Theoretical work argues that consumers with less common preferences may benefit more from social recommendations, which triggers demand for niche products (Hervas-Drane, 2015). This is supported by evidence that the popularity distribution of e-commerce sales can shift towards the tail (Oestreicher-Singer and Sundararajan, 2012a). Although search and recommendation tools guide consumers to content they were previously unaware of, they can also increase the number of consumed units and the chance of consuming content that many others are also consuming (Hosanagar et al., 2014). In our empirical context, it is not clear how this trade-off plays out. As musical content is heterogeneous in its breadth of appeal and prior exposure, we also investigate whether the effects differ when we distinguish between niche content (realized appeal) and new content (ex-ante unknown appeal).

3 Institutional background

3.1 Variation from YouTube and the GEMA shock

YouTube is a unique setting to study the relationship between the availability of online music videos and recorded music sales. While YouTube has contracts with rights holders in most countries, the question of compensation was subject to a long-standing legal dispute between YouTube and GEMA in Germany. GEMA (Gesellschaft für musikalische Aufführungs- und mechanische Vervielfältigungsrechte, society for musical performing and mechanical reproduction rights) is the state-authorized (de-facto monopolist) collecting society and performance rights organization in Germany.[10] Collecting societies are bodies that ensure that royalties from any kind of reproduction (e.g. physical and digital reproduction, public performance, etc.) reach artists and publishers, making them important institutions for artists because royalties are a major part of income, independent of any private contracts with record labels (Kretschmer, 2005). A large international network of sister collection societies represents the rights of German artists/publishers

[10] Examples for international counterparts are BMI, ASCAP and SESAC in the United States of America, PRS in the United Kingdom, SACEM in France and SGAE in Spain.

in international markets, and GEMA does the same for international artists/publishers in the German market. That is, virtually every professional musician is either directly or indirectly a member of GEMA, which led to the so-called 'GEMA presumption', a case law presuming that rights of all musical works are managed by GEMA in Germany.[11]

The expiry of an initial agreement between YouTube and GEMA in 2009 triggered renewed negotiations about the appropriate level of compensation. In fear of high subsequent (and retrospective) payments, YouTube began blocking music videos on April 1st 2009.[12] The left-hand side panel of Figure A.1, showing Google Trends search volume for the term "gema" from April 2008 to April 2010 indicates a spike in the week when the blocking began, but not much systematic movement before and after. This suggests that the shock came unexpectedly to consumers and most artists, publishers and record labels.[13] The situation persisted until November 1st 2016, when YouTube and GEMA finally announced that they reached an agreement and the restrictions on music videos were lifted.[14]

However, this does not necessarily imply that German YouTube users did not have access to any music videos. Publishers/artists can negotiate independent contracts with any online and offline licensee, so publishers and artists may drop out of GEMA to reach individual agreements with YouTube in Germany.[15] However, this may not be optimal. First, royalty income from the collecting society comes from a variety of digital (e.g. download and streaming services) and physical (e.g. radio broadcasting and public performance) sources. According to GEMA's annual report, online rights accounted for less than 2% of the overall combined income from online, physical duplication and radio and TV in 2008.[16] Hence,

[11] See http://kluwercopyrightblog.com/2012/10/01/the-gema-presumption-and-the-burden-of-non-liquet-germany/. In 2009, GEMA had 64,534 members and distributed 713 million Euro in royalties.

[12] See New York Times, 'Royalty Dispute Stops Music Videos in Germany', April 2, 2009, http://www.nytimes.com/2009/04/03/technology/internet/03youtube.html?_r=1.

[13] There is no evidence that YouTube systematically warned content owners in Germany before blocking videos.

[14] Specific legal issues have made it complicated to reach an agreement between GEMA and YouTube. According to Rolf Budde, member of the GEMA advisory board, YouTube insists on a non-disclosure agreement (see https://www.youtube.com/watch?v=Hh3Ks4Kxvtk). However, GEMA is required by law to publish the exact royalty paying schemes. Reportedly, because of this deadlocked situation, the involved parties consulted the arbitration board of the German Patent and Trademark Office for mediation in January 2013, and an agreement was finally reached in November 2016. See http://www.bbc.com/news/technology-37839038.

[15] After careful research, we could only find anecdotal evidence of one band opting out of GEMA. Videos in the official YouTube channel of the successful German rock band 'Die Ärzte' were accessible in Germany. See http://www.spiegel.de/netzwelt/web/netzwelt-ticker-warum-das-neue-aerzte-album-komplett-auf-youtube-laeuft-a-828244.html. It is not clear whether the band opted out of GEMA. When we asked their management for a statement, they declined to comment on the issue.

[16] See https://drive.google.com/open?id=0Bxe11iVXrXgsM0NIQm1OV1QxSVk.

income from digital distribution may be too small an amount to forgo all other royalty income for most artists. Second, by joining a collecting society, individuals benefit from reduced contracting cost and increased bargaining power. This is even more beneficial for members of international collecting societies where it can be especially costly to negotiate with various potential licensees abroad across different legal systems. To avoid this potential endogeneity of selecting into (or out of) YouTube in our estimates, we focus on a very short time window of four weeks before and after the blocking began on April 1st 2009.[17]

3.2 Long-run supply-side reactions and the launch of VEVO

The royalty dispute triggered a controversial discussion in the German music industry. While some artists agree with the position of GEMA because they believe that YouTube's royalty rates are too low, others simply want their videos to be seen.[18] Representatives of Sony Music and Universal Music have publicly criticized GEMA for not working harder towards an agreement.[19]

Record labels are not members of GEMA (they do not create music) and therefore do not receive any royalty income. However, on top of a potential positive effect on record sales, they directly benefit from advertising revenues generated by YouTube. Not surprisingly therefore, record labels are heavily invested in online music video. Sony Music and Universal Music, with a joint market share of more than 46% in 2012, for example hold majority stakes in the music video service VEVO.[20] Since its launch in 2009, VEVO has been partnering with YouTube in most countries. Accordingly, 97% of its 51.6 million unique viewers accessed VEVO-content through YouTube in December 2012, making VEVO the most viewed channel on YouTube, accounting for a third of all unique viewers on YouTube.[21] As a workaround for the GEMA-YouTube deadlock, VEVO negotiated its own licensing deal with GEMA and launched the dedicated platform `vevo.com` – with content hosted outside of YouTube – in the German market on

[17] Further, there is no evidence that consumers switched to other video platforms in the short run.

[18] For example, Sven Regener, singer of Element of Crime, says (referring to YouTube): "A business model based on people who produce the content not getting any money is not a business model, it's crap.", see http://www.br.de/radio/bayern2/sendungen/zuendfunk/regener_interview100.html. The popular electro/hip-hop band Deichkind posted a raging comment on their Facebook page after finding out that their newly uploaded music video was being blocked, see http://www.spiegel.de/netzwelt/web/deichkind-zum-gema-streit-ihr-seid-evolutionsbremsen-a-820703.html.

[19] See billboard.com, 2011, http://www.billboard.com/biz/articles/news/publishing/1177342/gema-under-fire-for-royalties-dispute-with-youtube.

[20] Market share data according to Nielsen Soundscan for the US, see http://www.statista.com/statistics/317632/market-share-record-companies-label-ownership-usa/.

[21] See http://www.comscore.com/Insights/Press-Releases/2013/1/comScore-Releases-December-2012-U.S.-Online-Video-Rankings.

October 1st 2013.[22] Overnight, 75,000 music videos became available on the German internet. We use this event as an auxiliary test for our general findings at a different time (2009 vs. 2013), with a different scope (all musical content vs. content by two major labels), in a different direction (making videos available vs. unavailable), and a different type of platform (open vs. closed platform that does not allow for user-generated content).

4 Data

We construct a unique dataset by matching sales information to song- and artist-specific meta data and measures of online video availability from a variety of sources. Table A.1 gives a summary and definition of all variables used in the paper and descriptive statistics of the key variables are in Table A.2.

4.1 Sales data

Sales data comes from GfK Entertainment, a market research firm collecting weekly (scanner) data. The data includes information from 50 (online and offline) retail outlet chains and 27 digital retailers in Germany and Austria that collectively represent more than 90% of all retailers.

We have access to the weekly number of units sold for the subsample of all songs that were among the 1,000 highest grossing songs (based on cumulative sales across all distribution channels) at least once in weeks 10–23 of 2009 in the physical channel (in Germany and Austria) and the digital download channel (in Germany). We also have sales information for weeks 36–44 of 2013. In 2013 we observe physical sales (in Germany and Austria), digital downloads (in Germany and Austria), and we can distinguish between the number of weekly streams via free services and subscription-based services (in Germany). Note that YouTube and VEVO plays are not included in the latter variable. For most analyses we use a four-week window around the respective experiments.

The data cover sales of individual songs. In the physical sales channel, this is equivalent to the sales of "singles", which implies that the song has been released on a physical medium – and stocked in brick-and-mortar outlets – for us to observe positive sales figures.[23] In the digital channel, our data capture sales of songs independent of whether they are released as standalone products (digital single) or as part of an album. This is due to the fact that consumers can (almost) always buy any individual song on an album

[22] See https://www.ft.com/content/5bb2092e-117e-11e3-a14c-00144feabdc0.

[23] During the observation period, physical record stores drastically reduced shelf-space for singles and focused on physical albums. See https://www.telegraph.co.uk/news/uknews/2033246/CD-singles-off-the-shelves-at-Woolworths.html.

in digital record stores. This is reflected in the share of digital song sales in our data – on average, it is 79% in 2009 and 97% in 2013. Looking at aggregate sales figures covering the top 1000 best selling albums in the German market in the same period, the digital share of top 1000 album sales is 5% in 2009 and 18% in 2013. This is in line with industry reports where the digital share in overall sales figures (songs and albums, entire German market) is 11% in 2009 and 23% in 2013.[24] We prefer to run our analysis at the song-level because music videos are directly linked to individual songs, not entire albums. In fact, as we show in section 6.3.1, the choice of level of aggregation is an important aspect when comparing our results to that of the prior literature.

Comparing the head and tail of the top 1000 list shows that it covers a very large fraction of the market. In our data, songs at the bottom of the list never sell more than 12 units (physical and digital combined) in a given week in the 2009 sample and never more than 20 units/12 streams in the 2013 sample, while the top song sells on average 37,871 units in 2009, and 28,121 units and 543,653 streams in 2013. Hence, demand for songs that do not make it to the top 1000 list is small, both in relative and absolute terms. However, we do not only track songs while they are part of the top 1000 list. For each song that appears in the top 1000 list at some point in our observation period, our data lets us observe sales figures also in weeks where the song is not included in the top 1000 list but available for purchase or streaming.

4.2 Meta information

We match meta information such as the release date, information on other releases of the same artist, genre, geographic origin of the artist, and record label, using data from Musicbrainz, an online platform for music enthusiasts. Musicbrainz contains user-generated information on about 20.5 million songs. Crucial information for our analysis, especially exact release dates (not only month or year), are not available for every song, but in some cases we found additional information from other sources, such as iTunes, Wikipedia and Discogs. Data on prior international success of an artist comes from historical chart rankings in the US (2000w1-2013w36) from Billboard.

We note some technical details. We do not observe unique common identifiers across the various databases, which is why we rely on a comparison of text strings, i.e. artist names and song titles. Because the additional datasets are very large and we need to match multiple ones, going through each combination of

[24] See http://www.musikindustrie.de/umsatz/.

potential matches is close to impossible. Further, variations in artist names and song titles (e.g. "featuring", "feat.", "(Radio Version)", "(Club Mix)") make one-to-one matching too restrictive, potentially causing (too many) matching errors. Because manual inspection of match candidates is potentially error prone, and there is no structured way to quantify error rates, we develop a statistical matching algorithm.[25] We compare pairs of potential matches along a number of metrics (such as the Levenshtein distance or Soundex), manually code a random subsample (training data, n=7654) to estimate parameters of a logit model, and use those parameters to estimate match probabilities in the full sample. From the statistical properties of the underlying model and training data, we can estimate error rates. Our model performs at an estimated rate of 4.4% of type I errors and 4.6% type II errors. This is much more precise than the heuristic approaches that are commonly used in the related literature (for example a reduction of 16 percentage points in the type I error rate in the best performing method discussed in Raffo and Lhuillery, 2009, Figure 4). Our final dataset includes 1,542 songs from 999 artists.

4.3 Music video data

To build a song-level measure of video availability on YouTube, we would ideally observe which songs had corresponding videos on German YouTube just before the ban on April, 1st 2009. As such historical data is not available, we construct a proxy by gathering the first 20 search results (this reflects the first page of search results on YouTube) from a query of artist name and song title on the US version of YouTube (using YouTube's API).[26] In many cases, not all videos that YouTube returns for a song are directly related to that song. Sometimes we observe videos related to other songs of the same artist, songs from similar artists, etc. We treat videos as relevant if the video title includes the artist name and at least three words of the song title. Using the upload date of each thus defined video, we construct our measure of availability. We set the dummy variable $Video_i$ to 1 if at least one video corresponding to song i was uploaded before April 1st, 2009. Identification of the YouTube effect will thus come from differences between songs that had corresponding videos on YouTube and those that did not.

We use data from US YouTube because we can realistically assume that German YouTube would offer the same content as US YouTube had the GEMA shock not happened. A simple plausibility check for

[25] Python code can be found on the author's GitHub page: https://github.com/cpeukert/.

[26] This query was performed on April 15, 2015. In estimations not reported here, but available on request, we get very similar results if we use information obtained from the Austrian version of YouTube.

Table 1: YouTube in the United States, Germany and Austria

	Share of directly relevant videos	Share of total views	Official video share
United States			
Mean	0.7755	0.8250	0.0868
Standard Error	0.0014	0.0014	0.0025
Austria			
Mean	0.7726	0.8213	0.0893
Standard Error	0.0014	0.0014	0.0026
Germany			
Mean	0.7485	0.7483	0.0502
Standard Error	0.0016	0.0021	0.0020

Source: George and Peukert (2014).
Top 20 YouTube search results for 950 randomly selected songs released between 2006 and 2011 (based on data from Musicbrainz). The search was carried out on August 21st, 2014. Relevancy is defined as a YouTube video title containing the artist name and at least three words of the song title. Total views are calculated as the cumulative number of views of all 500 videos shown on the first 20 results pages. Official videos are identified by the word "official" in the title or uploader name.

the latter is to compare US search results to that from Austria, Germany's neighbor which shares the same language and similar culture, but was not affected by the GEMA shock. George and Peukert (2014) conduct such an exercise with a random selection of almost 1,000 songs released between 2006 and 2011, collecting search results on German, Austrian and US YouTube for each song. Table 1 shows that the share of directly relevant videos on the first results page (defined as above) is not significantly different in the US and Austria, but 2.7 percentage points lower in Germany. The share of total views (sum of views of all 500 videos on the first 20 results pages) on the first results page in the US and Austria is about 82% (no significant difference). This share is about 75% in Germany. Almost 9% of the relevant videos on the first results page are official videos in the US and Austria (no significant difference), in Germany this share is only 5%.[27] This suggests that the Austrian version of YouTube looks very much like the US version of YouTube, while top search results on German YouTube are clearly different.

In essence therefore, we have a measure of video availability just before the ban, based on US YouTube, not German YouTube, which, however, is likely to be essentially the same before April 2009. For 54% of the songs in our sample, there is at least one corresponding YouTube video that predates the GEMA shock.

[27] George and Peukert (2014) define official videos based on whether the video title includes the word "official", whereas our definition is more accurate in that it is based on the account that uploaded the video. See below for details.

The seemingly high share of songs without videos is plausible as is likely that a substantial number of artists and record labels believed that YouTube displaces record sales, and did therefore not upload official content and actively sent takedown requests for user generated content. Table A.3 lists the top 20 (based on their peak ranking in our observation period) songs with videos and without videos. Songs without videos tend to have lower peak ranks than songs with videos. Consistent with this, descriptive statistics in Table A.2 show that songs with videos have substantially higher average sales.[28] However, as long as this difference is constant, it will cancel out in the difference-in-differences model. In the econometric model, we control for unobserved factors that may be correlated to such differences via song-fixed effects and week-fixed effects and show that the identifying assumption of the difference-in-differences model holds.

Similarly, we would like to observe which songs had corresponding videos on the German VEVO website when it was launched on October 1st, 2013. VEVO does not provide such a list, but we can make use of the fact that VEVO is part of YouTube in many other countries, including the US. The underlying assumption is that the content on VEVO's standalone German platform is the same as VEVO's content on YouTube.[29] We take advantage of the fact that VEVO uses artist-specific usernames to upload videos to YouTube. For example, the corresponding username for official videos by Justin Bieber is `JustinBieberVEVO`. Accordingly, we define a song as having a VEVO-video if at least one song-specific video is uploaded by VEVO, which is the case for 37% of the songs in our sample.[30]

To identify official non-VEVO videos, we manually went through the official YouTube accounts of artists and record labels and flagged videos as official if they were uploaded by these accounts. The average share of official videos among the relevant YouTube results is 6%. Accordingly, the average share of user-generated videos is 94%. Casual inspection of those videos suggests that these are mostly live and lyrics videos, remixes and cover versions, and very few parodies (see Liikkanen and Salovaara, 2015).

[28] Further, in an analysis available on request, we show that songs with videos tend to be younger than songs without videos, songs with more recent videos tend to have higher sales and newer songs (that are as old or younger as YouTube itself, which started in April 2015) more or less immediately have a video, whereas older songs show no clear pattern.

[29] According to what Tina Funk, General Manager Germany at VEVO, told us, this mostly holds true, but VEVO Germany also has some exclusive content – especially in the first weeks of its launch.

[30] All 247 artists listed under VEVO's main YouTube account have "VEVO" in their YouTube username, see https://www.youtube.com/user/VEVO/channels?view=56. We also manually checked all artists in our sample to make sure we do not miss a VEVO account that does not follow this convention.

4.4 Heterogeneity

We distinguish between established and newcomer artists in several ways, using information on all historical releases of an artist. First, we define **Newcomer: 2 months** as a dummy variable indicating whether the first release of an artist did not appear before February 1st, 2009. We observe 59 songs in this category. Similarly, **Newcomer: 1st year** is defined as an indicator whether the first release of an artist did not appear before April 1st, 2008. This identifies 102 newcomer songs. Our third definition, **Newcomer: No album** indicates whether the artist has never released an album before the GEMA shock, i.e. April 1st, 2009. The number of newcomer songs for this definition is 68. Prominent examples of newcomers according to these definitions are Oceana, Steve Appleton, Katy Perry and Lady Gaga.

We then define three empirical measures of mainstream versus niche music. First, we use historical sales data from the US (Billboard Top 100 single and album charts) to define a measure of international success. Examples of German artists that appeared at least once on the Billboard charts include Cascada, Sarah Connor and Tokio Hotel. However, consistent with the evidence in Ferreira and Waldfogel (2013) showing that Germany almost exclusively exports music to Austria and Switzerland, these are exceptions. Accordingly, our first measure of niche artists consists of 638 artists that never appear in the US top charts, which we include in the category **Niche: Never US**. Second, based on data about their geographical origin, we classify 601 artists as **Niche: German**. This represents 60% of all artists in our data, which again is consistent with the estimate of the domestic share of music consumption in Ferreira and Waldfogel (2013). Third, we identify the best-selling genres by looking at cumulative sales in our data. Pop and Rock account for 74% of overall sales in the German market before the GEMA shock, such that we include all other genres in the category **Niche: Genre**.

One might be concerned that songs of newcomer or niche artists differ in their video availability. The descriptive statistics in A.2 do not confirm this. Applying Bayes' Theorem, we see that the probability that a song does not have a video if it is from a niche artist is roughly 50%, independent of how we measure *Niche*.[31] The probability that a song does not have a video if it is from a newcomer artist is

[31] Consider the example **Niche: Never US**. $P(B|A) = 0.821$ is the probability that a song is from niche artist if it does not have a video. $P(B) = 0.723$ is the probability that a song is from niche artist. $P(A) = 1 - 0.548 = 0.452$ is the probability that a song does not have a video. Using Bayes' Theorem, $P(A|B) = (0.821 * 0.452)/0.723 = 0.513$, gives the probability that a song does not have a video if it is from a niche artist.

Figure 1: Trends of treatment and control group, before and after GEMA shock

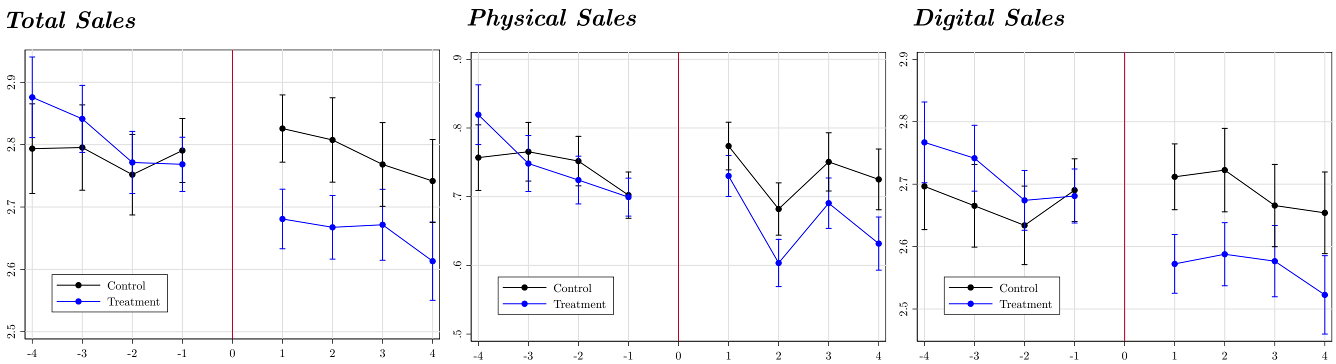

Vertical axis: *Average demeaned total/digital/physical sales, i.e. averaged residuals* $\overline{\hat{y}_{vt}} = \overline{\hat{y}_{it} - \hat{\mu}_i}$ *derived from the model* $y_{it} = log(Sales_{it} + 1) = \alpha + \sum_t \beta_0^t w_t + \sum_t \beta_1^t (w_t \times Video_i) + \mu_i + \varepsilon_{it}$ *for* $Video_i = 0$ *and* $Video_i = 1$.
Horizontal axis: *Weeks prior/after April 1st, 2009.*
Black (control group): *Average sales of songs without at least one video uploaded to U.S. YouTube before April 1st, 2009.*
Blue (treatment group): *Average sales of songs with at least one video uploaded to U.S. YouTube before April 1st, 2009.*
Bars indicate 90% confidence bands (standard error of the mean).

between 34% and 48%, depending on the definition of *Newcomer*.

5 Empirical specification and results

We first introduce and discuss the identification strategy and then report our baseline results on the average effect of online music videos on recorded music sales. We go on to test whether this effect is driven by differences in the song-specific amount of official and user-generated content on YouTube. Finally, we investigate heterogeneity and distinguish between new artists and niche artists. We report results of additional analyses (some of which are described in more detail in Appendix B) that largely support the robustness of our results and help us rule out alternative explanations.

5.1 Identification strategy

We identify the effect of music videos on sales of recorded music using exogenous variation from removing access to videos on YouTube in a difference-in-differences model. Essentially, we compare sales of songs *with* videos to sales of songs *without* videos before and after the natural experiment. Our baseline specification can be written as

$$log(Sales_{it}^k + 1) = \alpha + \sum_t \beta^t w_t + \delta\,(After_t \times Video_i) + \mu_i + \varepsilon_{it}, \tag{1}$$

Table 2: Group differences in the pre-period, GEMA shock

	(1) Total	(2) Physical	(3) Digital
t_{-4} × Video	0.113*	0.075	0.097
	(0.068)	(0.046)	(0.066)
t_{-3} × Video	0.091	-0.004	0.117*
	(0.062)	(0.042)	(0.060)
t_{-2} × Video	0.036	-0.018	0.055
	(0.058)	(0.036)	(0.056)
t_{-1} × Video	-0.013	-0.003	0.000
	(0.046)	(0.029)	(0.045)
Observations	7543	7543	7543
$\overline{R}^2$	0.922	0.920	0.924

Dependent variables: (Log+1) weekly sales in units.
Video indicates (at least one) song-specific video on U.S. YouTube uploaded prior to April 1st, 2009.
Only weeks before the respective experiments. Song and week fixed effects, constant not reported.
Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

where $Sales^k_{i,t}$ are unit sales of song i via channel k (physical, digital) in week t. $Video_i$ is our measure of video availability and defines the treatment group. $After_t$ indicates the time period after the GEMA shock. Our estimate of the causal effect of the experiment is δ. We further include week fixed effects β^t and song fixed effects μ_i to control for unobserved song-specific time-invariant and time-specific song-invariant heterogeneity. Because of these fixed effects, we implicitly control for, but cannot separately identify coefficients for $Video_i$ and $After_t$. Our preferred specification reports standard error estimates clustered at the song-level.

The identifying assumption in any difference-in-differences setting is that treatment and control group would have followed a similar trend in the dependent variable had the policy shock not happened. A necessary condition for this assumption to hold is that trends in the dependent variable of the treatment and control group are parallel before the experiment. To see whether this condition holds, we can plot a sales measure over time. In Figure 1, we partial out song-fixed effects and plot averaged residuals for each group and each week, regarding total sales, physical sales and digital sales. The plot shows that treatment and control group follow similar trends before the GEMA shock and start to diverge

Table 3: Baseline results: Average effect of the GEMA shock on song sales

	(1) Total (avg)	(2) Total	(3) Physical	(4) Digital	(5) Total (avg)	(6) Total	(7) Physical	(8) Digital
After × Video	-0.142**	-0.146***	-0.072*	-0.153***	-0.142**	-0.146***	-0.072*	-0.153***
	(0.068)	(0.050)	(0.037)	(0.049)	(0.070)	(0.052)	(0.037)	(0.051)
Observations	3084	13711	13711	13711	3084	13711	13711	13711
SE Cluster	Song	Song	Song	Song	Artist	Artist	Artist	Artist
$\overline{R}^2$	0.918	0.892	0.877	0.894	0.918	0.892	0.877	0.894

Dependent variables: Columns (1) and (5): Average Log(1+weekly total sales) of song i in the pre and post period. All other columns: Log(1+weekly total/physical/digital sales) of song i in week t.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded prior to April 1st, 2009.
After indicates weeks after April 1st, 2009.
All models include song fixed effects, and in columns (2)–(4) and (6)–(8) we additionally include week fixed effects.
Constant not reported. Standard errors in parentheses, clustered at the song/artist-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

substantially afterwards. Noting that the 90% confidence bands (calculated using the standard error of the mean) overlap in the pre-period, we can conclude that the parallel trends assumption is consistent with our data. More formally, we can test whether the difference in the dependent variable across points in time is zero in the pre-period, as is often done in the literature (see Autor, 2003 for one of the first applications of this idea). We estimate a model defined as

$$log(Sales^k_{it} + 1) = \alpha + \sum_t \beta^t_0 w_t + \sum_t \beta^t_1 \left(w_t \times Video_i\right) + \mu_i + \varepsilon_{it}, \quad (2)$$

in which we can test, week by week, if treatment and control groups differ in their sales dynamics ($H_0 : \beta^t_1 = 0$). Table 2 shows estimates of β^t_1 coefficients for total sales, physical sales and digital sales in the pre-period. Across almost all columns and pre-experiment weeks we cannot reject the hypothesis that the difference in sales in the treatment and control group is equal to zero. While these results are reassuring, we discuss an alternative identification strategy (using cross-country variation) and a falsification exercise (placebo country and timing) after the baseline results below.

5.2 Displacement or promotion?

5.2.1 Baseline results

Our baseline results are reported in Table 3.[32] In columns (1) and (5), we report the results of an aggregated model specification, where we look at total sales of a song as an average in the pre- and post-period. This specification lets us address the potential issue that serial correlation may lead to incorrect inference (Bertrand et al., 2004). We further report results of different methods of estimating standard errors (Clustered at the artist-level in column 1, clustered at the song-level in column 5). The estimated difference-in-differences coefficient is negative and significant at the 5% level. The point estimate is -0.142, which implies a decrease of 13%.[33]

Turning to the preferred disaggregated sample in columns (2) and (6), we get very similar results. The point estimate translates into a percentage reduction of 14% of total sales, with the 90% confidence interval between -21% and -6%. Because it is likely that the temporal correlation structure of sales of the same song is stronger than those of sales of different songs of the same artist, we continue to report results with standard errors clustered at the song-level in the rest of the paper.

In columns (3)–(4) and (7)–(8) we distinguish between physical and digital sales of the same song. Although the point estimate for physical sales is half the size of the point estimate for digital sales (-7% vs. -14%), the coefficients are not statistically different from each other in the sense that 90% confidence bands overlap. Hence, we conclude that YouTube's average effect on sales of songs is not statistically different regarding physical and digital sales.

5.2.2 Alternative identification strategy and falsification exercises

An alternative identification strategy that does not use song-specific information of video availability is to compare sales in Germany to sales in a different country not affected by the GEMA shock. Austria, Germany's neighbor which shares the same language and similar culture is a prime candidate. In this exercise, we assume that the GEMA shock affected all songs in the same way in the German market, i.e. that all songs have corresponding videos on YouTube which are all blocked. In Table 4 we report results of a model that compares sales of the same song in Austria to its sales in Germany, before and

[32] In results available on request, we show that the estimated effect of the GEMA shock is similar when we do not include fixed effects. Further, the estimated effect is consistent if we look at sales ranks, rather than units, as dependent variable.

[33] Here, and in what follows below, we calculate percentage effects as $(exp(-0.142) - 1) * 100 = -13.24$.

Table 4: Alternative identification strategy

	(1)
Germany	0.424***
	(0.046)
After × Germany	-0.068**
	(0.033)
Observations	27422
$\overline{R}^2$	0.462

Dependent variables: (Log+1) weekly physical sales in units (digital sales are not available for Austria).
Song and week fixed effects, constant not reported.
Standard errors in parentheses, clustered at the song-level. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

Figure 2: Falsification exercises: Timing of placebo experiments

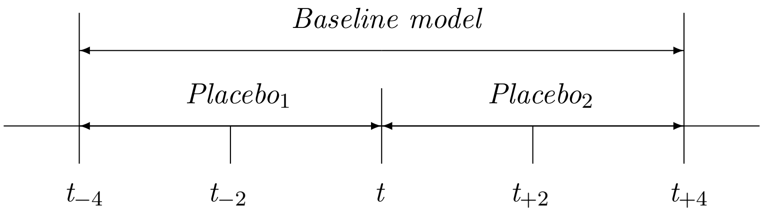

after the GEMA shock. Note that digital sales data does not exist for Austria, so that we can only compare physical sales. The number of observations is accordingly twice as large as in the disaggregated specifications (2)–(4) and (6)–(8) of Table 3. Because we have two observations per song, we can identify a country coefficient. Given that Germany is a larger market than Austria, this coefficient is positive and significant as expected. The difference-in-differences coefficient *After* × *Germany* is negative and significant. The effect size is −7% (90% CI [−12%,−1%]).

We also perform two falsification exercises. First, we estimate placebo-versions of our model pretending that the GEMA shock took place either two weeks before or two weeks after it actually did. As illustrated in Figure 2, we split the sample of our baseline model (running from t_{-4} to t_{+4}) in two parts and estimate models on the two subsamples running from t_{-4} to t and from t to t_{+4}, setting the respective dates of the placebo-experiments to t_{-2} and t_{+2}. That is, we run two types of placebos. For the first, if there is no general underlying trend, we should expect an effect that is close to zero. For the second, we expect an effect that is close to zero only if we assume that the effect of the real experiment is instantaneous

Table 5: Falsification exercises: Placebo experiments

	Timing		Country
	−2 weeks	+2 weeks	Austria
	(1)	(2)	(3)
After × Video	-0.034	-0.037	0.001
	(0.060)	(0.043)	(0.001)
Observations	4144	3852	13711
$\overline{R}^2$	0.803	0.734	0.755

Dependent variables: (Log+1) total weekly sales in units, in Germany (columns 1–2) and Austria (column 3).
Column (3) only includes physical sales because digital sales data are not available for Austria.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded prior to April 1st, 2009.
Song and week fixed effects, constant not reported. Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

and constant. Hence, we can take this exercise as a test of whether the effect of the true experiment is immediate and whether it changes in the observed time frame.[34]

In the results reported in columns (1) and (2) of Table 5,[35] we define the after period to include the week of the placebo experiment, but we get very similar estimates if we treat the placebo experiment week as part of the before period. The coefficients of *After* × *Video* in columns (1) and (2) are small and not significantly different from zero, suggesting that our results are not driven by a general trend that started before the GEMA shock, and that the effect of the shock persists. Table 5 only reports results concerning total sales as dependent variable, but we also do not find significant effects when we distinguish between physical and digital sales. Second, we run a country-based placebo exercise and estimate our model on data from Austria. If our results are driven by confounding temporal variation that coincides with the GEMA shock, we should see a similar effect in Austria. Alternatively, because music videos on YouTube remained available for Austrian consumers, we should not see any change in the Austrian sales of songs that are affected by the experiments in Germany. This is indeed what the data tells in column (3) of Table 5. We find that songs with videos on YouTube do not have significantly different sales in Austria

[34] If artists quickly adjust their digital strategy and drop out of GEMA or if consumers quickly discover technical measures to circumvent the blocking (e.g. VPN), we expect a positive coefficient of *After* × *Video* in the second placebo test.

[35] One may be concerned that the standard errors are large because the sample size is much smaller than in the baseline specifications of table 3. To make a fair comparison, we should therefore also estimate the "true" effect on a smaller sample. In results but available on request, we show that the coefficient of *After* × *Video* estimated on a sample covering a window of ± 2 weeks around the true date is −0.133 (s.e. 0.051), with a 90% confidence interval ranging from −0.233 to −0.033.

compared to songs without videos on YouTube before and after the GEMA shock in Germany. The coefficient is very close to zero.

While our results are robust to a number of specifications, some concerns regarding data structure and measurement error may remain. In appendix B, we speculate on the possible effects of measurement error, show that our results are broadly robust to different estimation windows, and conclude that the results are not likely to be driven by price changes.

5.3 The role of content type and platform type

One of the specific features of YouTube as a *licensed open platform* is that it can host both official and user-generated content (UGC). We provide two tests of whether this distinction can lead to different effects on sales. First, as described above, we separate the available song-specific videos by official and user-generated content. Second, we make use of the fact that VEVO is a closed platform where only rights holders themselves can contribute and estimate how its launch affected sales of recorded music across different distribution channels.

5.3.1 User-generated content

Table 6 reports results from a specification where we define dummy variables based on the share of official and user-generated content. We categorize songs based on the amount of official and user-generated content on YouTube. On average, the share of official videos in the directly related search results is 6%. The median is 0 and the 75th percentile is 10%. We define a song as having a *Small UCG Share* if the share of official videos is larger than 10%, which is roughly equivalent to two official videos on YouTube's first results page.[36] Songs without any official videos are categorized as *Only UGC*. Accordingly, songs with a share of official videos of between 0% and 10% are classified as *Official and UGC*. Similar to the baseline specification above, the omitted category comprises songs without videos. We report results for total sales (column 1), physical sales (column 2) and digital sales (column 3). The point estimates of $After\times$ *Small UGC Share*, $After\times$ *Official and UGC* and $After\times$ *Only UGC* are very similar. We interpret these results as evidence that the effect of music videos on record music sales operates in a similar fashion independent of whether there is more or less user-generated content on YouTube. In section B.2

[36] In results available on request, we show that the coefficients remain similar (qualitatively the same) if we use different thresholds, e.g. the sample mean of 6.4%.

Table 6: Official versus user-generated content

	(1) Total	(2) Physical	(3) Download
After × Small UGC Share	-0.183***	-0.085	-0.199***
	(0.069)	(0.062)	(0.066)
After × Official and UGC	-0.146*	-0.090	-0.155**
	(0.075)	(0.058)	(0.074)
After × Only UGC	-0.126**	-0.057	-0.126**
	(0.061)	(0.043)	(0.060)
Observations	13711	13711	13711
$\overline{R}^2$	0.892	0.877	0.894

Dependent variables: (Log+1) weekly sales in units, 2009.
Small UGC Share indicates that more than 10% of the song-specific videos
on U.S. YouTube, uploaded prior to April 1st, 2009, are uploaded from an official account.
Official and UGC indicates a share of official videos of between 0 and 10%.
Only UGC indicates that no video is uploaded from an official account.
The omitted category is "No Video". Song and week fixed effects, constant not reported.
Standard errors in parentheses, clustered at the song/artist-level. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

of the appendix, we show that the parallel trends assumption is supported by the data. The pre-GEMA shock sales dynamics of treatment (*Small UCG Share*, *Official and UGC*, *Only UGC*) and control songs do not systematically vary by the amount of user-generated video content on YouTube.

5.3.2 The entry of VEVO

The GEMA shock was partly reversed when Universal Music and Sony Music launched their own music video platform VEVO on October 1st, 2013. Although it is independently interesting to ask if the availability of music videos affects sales with a comparable magnitude (but opposite sign) as their unavailability, the entry of VEVO lets us speak to the role of user-generated content from another angle. VEVO only hosts official videos while YouTube hosts a large share of UGC videos. Table 7 replicates our baseline regressions with data covering four weeks before and after the VEVO launch. As described in section 4.3, we can measure whether song-specific music videos were available on the new VEVO platform because VEVO operates within the YouTube platform in almost all countries except Germany. In 2013 we can not only observe unit sales in the physical and digital purchase channel (columns 1–3), but also how often a particular song was streamed on free and premium *licensed closed streaming sites* (not including

Table 7: Average effect of VEVO's entry on record sales

	(1) Total	(2) Physical	(3) Digital	(4) Total	(5) Free	(6) Premium
After × Video	0.092**	0.021	0.087**	0.149*	0.133	0.143*
	(0.042)	(0.018)	(0.042)	(0.081)	(0.101)	(0.076)
Observations	14576	14576	14576	14576	14576	14576
$\overline{R}^2$	0.858	0.889	0.859	0.898	0.852	0.896

Dependent variables: (Log+1) weekly sales/streams in units, 2013.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded by VEVO.
Song and week fixed effects, constant not reported. Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

YouTube and VEVO, columns 4–6).

We find remarkably similar point estimates as in our analysis of the GEMA shock, yet with the (expected) opposite sign. Looking at the difference-in-differences estimate in columns (1) and (4) of Table 7, the entry of VEVO increased total record sales by about 10% (90% CI [2%,18%]) and the number of weekly streams by about 16% (90% CI [2%,33%]). Although it is important to note that these coefficients are less precisely estimated than the coefficients reported in Table 3 – perhaps because the entry of VEVO is a less clean experiment and because of the small number of physical copies sold in 2013 – our VEVO results are not statistically different from the GEMA results since the absolute confidence bands overlap for all common coefficients. In section B.3, we run a series of robustness checks that make us confident that the direction of our estimates reflects the causal effect of the VEVO entry. First, we show that the sales/streams of songs without VEVO videos and songs with VEVO videos follow similar trends before the entry. Second, we show that the results hold under an alternative identification strategy (treating all songs released by Universal and Sony as affected). Third, we show that there are no significant effects when we run falsification exercises based on placebo timing or placebo geography. Fourth, we can rule out that the effect is driven by price changes that temporally coincide with the entry of VEVO.

5.4 Effect heterogeneity

We now go beyond the average effect and look into effect heterogeneity. We allow for heterogeneity at the level of popularity in terms of sales, investigating the effect of the GEMA shock at various points of the sales distribution. We then allow for heterogeneity according to consumer awareness, and the breadth of an artist's appeal. We estimate triple difference models by adding the additional interaction terms $\delta_1(After_t \times X_i)$ and $\delta_2(After_t \times Video_i \times X_i)$ to equation (1), where $X_i \in \{Newcomer_i, Niche_i\}$. Under the null hypothesis $\delta_2 = 0$ we can directly test if the effect of the GEMA shock differs across observations where $X_i = 0$ and $X_i = 1$. The total effect for observations where $X_i = 1$ is $\hat{\delta} + \hat{\delta_2}$.

5.4.1 Overall popularity

The results in Table 8 suggest that the average baseline effect does not differ significantly by overall popularity. We test for differences in the effect of the GEMA shock on sales across the distribution of the sales variable in a number of ways. We begin by exploring variation in the amount of time songs stay in the top 1000 ranking to capture the notion that songs that stay in the top 1000 ranking for longer are more popular. In column (1), we test for differences between songs that stay in the top 1000 sales ranking throughout the observed period and those that drop out at some point. The effect of the GEMA shock on the latter, estimated as the coefficient $After \times Video$, is -10% and therefore about 4 percentage points smaller than the coefficient in our baseline results in column (2) of Table 3. The 90% confidence band ranges between -19% and -1%. We do not find a significant difference, as indicated in the coefficient of $After \times Video \times AlwaysInTop1000$. Adding the two, the total effect for songs that always stay in the top 1000 ranking is -5%, but its 90% confidence band overlaps substantially with that of the effect for songs that drop out of the top 1000 list ([-15%,5%]). Hence, we cannot interpret as evidence that more popular songs exhibit a less strong promotional effect from online music videos. In column (2), we go one step further and test whether the effect of the GEMA shock is different for even more popular songs. In the subset of songs that always remain in the top 1000 list, we distinguish between those that never fall below rank 200 over the observed period and those songs that do.[37] Again, we find no significant difference. The effect for songs that exit the top 200 is -10% (90% CI [-17%,-2%]), while the effect for those that stay within the top 200 is 2%, yet this estimate is very imprecise (90% CI [-14%,22%]). We continue to test whether songs at the top of the pre-experiment sales distribution respond differently to the GEMA shock in column (3). Comparing songs in 95th percentile to all others, we do not find a significant difference at the top of the pre-experiment sales distribution. The implied effect for songs outside the 95th percentile is -12% (90% CI [-19%,-4%]). The implied effect for songs in the 95th percentile is -6%, but very

[37] We choose the cut-off of 200 to be consistent with Hiller (2016). See the discussion in section 6.3.1.

imprecisely estimated (90% CI [−26%,19%]). Again, these results do not suggest that more popular songs exhibit a significantly weaker promotional effect from online music videos. In columns (4) and (5), we estimate quantile regressions for different percentiles of the dependent variable. Note that in these exercises, we look at pre- and post-experiment sales. The results suggest the effect of the GEMA shock is −16% at the 5th percentile of the total sales distribution (90% CI [−25%,−5%]), and −19% at the 95th percentile (90% CI [−37%,6%]). Finally, in column (6), we report the results of a weighted regression, where we give observations with higher average sales in the pre-experiment period a relative larger weight. The point estimate is −13% (90% CI [−24%,−1%]), still very similar compared to the baseline estimate in column (2) of Table 3.

These exercises do not support the notion that the average treatment effect in our setting affects different parts of the sales distribution differently. However, the above analysis did not consider that certain groups of observations may respond differently to the GEMA shock, which may affect sales of these groups, but is subtle enough to not affect the overall popularity distribution. From a managerial and policy perspective therefore, it is important to uncover this type of heterogeneity.

5.4.2 Consumer awareness

Our measures of consumer awareness (defined in section 4.4) are based on the idea that consumers are less aware of new artists than of established artists.[38]

Results in columns (1)–(3) of Table 9 show that newcomer artists are affected more strongly than established artists. The triple interaction term is negative in all columns and significant in columns (2)–(3), suggesting that total record sales of new artists decrease more when music videos are blocked on YouTube. The point estimate in column (1) translates into a −12% effect for established artists (90% CI [−19%,−5%]), and −50% for new artists (90% CI [−75%,2%]).[39] Column (2) suggests that total record sales decrease by 11% for established artists (90% CI [−18%,−3%]), and by 54% for new artists (90% CI [−75%,−14%]). The point estimate for established artists in column (3) is −11% (90% CI [−18%,−3%]), and −46% for new artists (90% CI [−50%,−41%]). These results are consistent with the notion that YouTube promotes artists that are less known to consumers. While only the specification in column (3) shows that the

[38] In section B.2 of the appendix, we show that the parallel trends assumption also holds for newcomer artists as a subset of the treatment group.

[39] $(exp(-0.133) - 1) * 100 = -12.45$; $(exp(-0.133 - 0.553) - 1) * 100 = -49.64$.

Table 8: Heterogeneity: Overall popularity

	(1) Time	(2) Pre-Rank	(3) Pre-Sales	(4) 5th Pctl.	(5) 95th Pctl.	(6) Weighted
After × Video	-0.110*	-0.102**	-0.125**	-0.171**	-0.207	-0.141*
	(0.060)	(0.052)	(0.052)	(0.074)	(0.159)	(0.080)
After × AlwaysInTop1000	-0.350***					
	(0.071)					
After × Video × AlwaysInTop1000	0.055					
	(0.088)					
After × AlwaysInTop200		-0.414***				
		(0.103)				
After × Video × AlwaysInTop200		0.124				
		(0.120)				
After × 95th Pctl.			-0.547***			
			(0.133)			
After × Video × 95th Pctl.			0.060			
			(0.154)			
Observations	13711	2799	13711	13711	13711	13711
$\overline{R}^2$	0.892	0.916	0.892			

Dependent variables: (Log+1) total weekly sales in units.
Song and week fixed effects in columns (1) and (5), columns (2)–(4) are quantile regressions with differenced-out song fixed effects and time dummies, (5) is a weighted regression using average pre-experiment sales as weights.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded prior to April 1st, 2009.
Constant not reported. Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

differences between established and new artists are significant, it is important to note that the statistical precision in Table 9 is impacted by the fact that the number of observations that qualify as newcomers (in either definition) is relatively small.[40]

5.4.3 Breadth of appeal

Our second set of results concerns heterogeneity across the inherent breadth of the artist's appeal. We do this by distinguishing between niche (narrow appeal) and mainstream (broad appeal) artists.[41] Because

[40] This is of course a reflection of the empirical distribution: the number of new artists in the music market is almost mechanically smaller than the number of established artists.

[41] In section B.2 of the appendix, we show that the parallel trends assumption also holds for niche artists as a subset of the treatment group.

Table 9: Heterogeneity: Consumer Awareness

	(1) 2 Months	(2) No Album	(3) 1st Year
After × Video	-0.133***	-0.114**	-0.114**
	(0.050)	(0.049)	(0.050)
After × Newcomer	0.668*	0.850***	0.637**
	(0.385)	(0.325)	(0.250)
After × Video × Newcomer	-0.553	-0.661*	-0.495*
	(0.430)	(0.384)	(0.287)
Observations	13711	13711	13711
$\overline{R^2}$	0.892	0.892	0.892

Dependent variables: (Log+1) total weekly sales in units.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded by prior to April 1st, 2009.
Newcomer defined as *2 Months* in column (1) and (3), as *No Album* in column (2) and (4), *1st Year* in column (3) and (5).
Song and week fixed effects, constant not reported. Standard errors in parentheses, clustered at the song/artist-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

any empirical definition may be arbitrary, we again report results for three distinct measures (defined in section 4.4). Results in Table 10 show that the triple interaction is positive and significant.

In column (1), the effect for non-niche artists is −33% (90% CI [−46%,−16%]), while the effect for niche artists (defined as never having appeared in the US charts) is −9% (90% CI [−17%,−0.04%]). In column (2), the effect for non-German artists is −23% (90% CI [−33%,−12%]), while for German-origin artists, we find a non-significant decrease of 8% (90% CI [−17%,3%]). The specification in column (3) distringuishes between Pop/Rock and other genres. The effect for Pop/Rock songs is −21% (90% CI [−30%,−12%]), for all other genres we find a non-significant increase of 1% (90% CI [−10%,12%]).

The 90% confidence bands overlap in the specifications in columns (1) and (2), and not in column (3). Hence, we can only confirm that total sales of artists in niche genres react differently to the GEMA shock than those of artists in mainstream genres for the genre-based definition of “niche”.

6 Discussion

In this section, we first discuss how to interpret our findings in the context of music discovery and speculate about the underlying mechanisms, before we assess the implications of the heterogeneous effects on aggregate sales. We then “zoom out” and provide a back-of-the-envelope calculation of the overall

Table 10: Heterogeneity: Breadth of Appeal

	(1) Billboard	(2) German	(3) Genre
After × Video	-0.394***	-0.265***	-0.241***
	(0.134)	(0.081)	(0.066)
After × Niche	-0.310**	-0.213**	-0.061
	(0.131)	(0.083)	(0.079)
After × Video × Niche	0.300**	0.187*	0.249**
	(0.146)	(0.105)	(0.102)
Observations	13711	13711	13711
$\overline{R^2}$	0.892	0.892	0.892

Dependent variables: (Log+1) total weekly sales in units.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded by prior to April 1st, 2009.
Niche defined as *Never US* in column (1), as *German* in column (2), *Genre* in column (3).
Song and week fixed effects, constant not reported. Standard errors in parentheses, clustered at the song/artist-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

economic significance and welfare effects of the blocking policy in the context of the GEMA shock. Finally, we discuss the external validity and general implications of our study in relation to prior literature.

We find that the effect of the removal of online music videos on song sales is negative. Point estimates vary, but -10% to -5% is a conservative estimate of the short-run effects of removing access to music videos on sales of recorded music. This suggests that despite while free interactive content could plausibly substitute for demand for paid content (as found in most of the piracy literature), free online music videos tend to complement sales of recorded music. The extent of user-generated content does not affect the positive effect of YouTube availability on music sales, a result established by directly measuring the share of UGC and further supported by the entry of VEVO. Our results further suggest that YouTube can play an important role in the discovery of music. Looking at artist subgroups, we find that new artists benefit more from content availability on YouTube (and consequently see a bigger drop in sales following the YouTube blackout). Conversely, niche artists do relatively better following the unavailability of YouTube content.

It is useful to think about these results in terms of the underlying process of music discovery. Since music is an experience good, consumers rely on sampling (Peitz and Waelbroeck, 2006), word-of-mouth recommendations (Susarla et al., 2012; Lee et al., 2015) or automated recommender systems (Liu et al., 2014; Zhou et al., 2016) for their choice of which new musical content to consume. Underlying all of them is the notion that music discovery incurs search costs, and each of the three cues for consumers reduces search cost for new music. Consider now the discovery process that turns music videos on YouTube into music purchases. Exposure to free content on YouTube leads to higher sales under three conditions: First, "new" music has to be new to the individual as otherwise the individual may have already purchased it in the past. Second, the "new" music has to meet the tastes of the listener. Finally, the sampled content must be a meaningful indicator of the quality of the song. Because the audio part of music videos tends to be either identical or a close adaptation of what can be purchased in a (digital) record store, this condition holds in the case of YouTube. Discovery on YouTube may directly transform into digital purchases via "click-to-buy" links to retailers like iTunes and Amazon (a mechanism that we cannot test in the absence of individual-level data). Other times, although there is a causal link between sampling on YouTube and music purchases, this may work more indirectly without leaving a trail in a consumer's clickstream.

What do our heterogeneous effects tell us about the way in which music discovery is affected by YouTube? First, the share of user-generated content plays no role in the effect of YouTube (un)availability. Conversion rates are not higher or lower if consumers hear the song accompanied by an official video compared to UCG, suggesting that the process of discovery does not need curated complements to the song. YouTube creates awareness of the original song (the function common to official and user-generated versions of the song) while the delivery channel is not considered a close substitute to actually purchasing the song.

We also find that new artists benefit relatively more from YouTube availability, as do mainstream artists. The former is in line with the awareness logic – listeners have not been exposed widely to these new artists and exposure on YouTube helps listeners discover them. However, the result that mainstream artists benefit more than niche artists from YouTube availability suggests another function of YouTube: holding everything else fixed, artists in mainstream genres gain more from YouTube exposure. This suggests that the likelihood of a purchase after watching a YouTube video is higher for a mainstream artist because the general appeal of mainstream artists on YouTube is broader.

Apart from demand-side awareness and mass appeal, what supply-side mechanisms may lead to the disproportional effects we observe in the data? First, music videos may be systematically different for new and mainstream artists. Music videos have become a widespread complement to audio recordings only in recent years. This implies that new artists are more likely to have both more and better music videos than established artists.[42] Similarly, because of mainstream artists' greater expected appeal, the incentives to invest in video content are bigger, which would lead to more and perhaps better videos for mainstream songs. Another mechanism consistent with our heterogeneous effects is based on active platform management. YouTube's objectives will likely be a mix of dynamic and static goals. Maximizing views at any one point in time increases YouTube's opportunities for monetization in the short run. This speaks for promoting already popular content with wide appeal to increase the chances of repeat consumption. However, consumers need to be exposed to new content systematically to avoid them losing interest in the platform. This gives YouTube an incentive to recommend new artists.[43] In sum, both artists and labels (through better/more engaging video content for mainstream and new artists) and the platform itself (through recommender systems) may reinforce the process of music discovery driven by awareness and mass appeal of music. These supply-side actions are, however, complementary to demand-side forces rather than substitutes.

6.1 Online music videos and changes in the distribution of music sales

Our findings are consistent with the idea that the process of music discovery is affected by the availability of YouTube videos by raising potential consumers' awareness of new music, which especially benefits artists about which consumers are initially relatively unaware and artists who have broad appeal and are therefore more likely to convert video views into record sales. This is in line with Lee and Hosanagar (forthcoming), who show that recommender systems in e-commerce tend to put weight on both extremes of the sales distribution, i.e. popular products become more popular, while long-tail products increase sales too. The fact that different types of music benefit to different extents from exposure on YouTube can have implications for the overall sales distribution. The literature has long speculated whether digital distribution platforms (and perhaps the recommender systems operating on them) help the most popular

[42] Note that established artists in our dataset include ones that started their musical careers before MTV and music videos were "invented" and widely adopted as marketing tools.

[43] This is highly consistent with the actual goals of YouTube's video recommendation system as practiced during our time period, as the following quote (from a paper by researchers at Google that describes YouTube's recommendation algorithm) shows: "We want recommendations to be reasonably recent and fresh, as well as diverse and relevant to the user's recent actions." (Davidson et al., 2010, p. 294).

products become even more popular (Celma and Cano, 2008; Fleder and Hosanagar, 2009; Zhou et al., 2010; Oestreicher-Singer and Sundararajan, 2012b) or if they help niche products with relatively low sales (the "long tail") reach a wider audience (Tucker and Zhang, 2011; Oestreicher-Singer and Sundararajan, 2012a). Related work finds evidence that streaming can also change the distribution of sales (Hiller, 2016; Datta et al., 2017). Using individual-level consumption data from 900 consumers, Datta et al. (2017) show that users switching from owning (iTunes) to streaming (Spotify) listen to music from a more diverse set of artists and discover artists that are new to them. Similarly, the displacement effect of streaming in Hiller (2016) is weaker for less known albums.[44]

In the specific context of our short-run analysis of the removal of music videos due to the licensing dispute between GEMA and YouTube, a back-of-the-envelope calculation shows that newcomer artists suffer more than niche artists benefit relative to mainstream artists. A conservative reading of our results implies that sales of newcomer artists change by about $\hat{\delta}_{new} = -41\%$, and sales of niche-genre artists change by about $\hat{\delta}_{niche} = -10\%$ because of the GEMA shock. To understand the aggregate effect of this, we need to compare realized market shares to the counterfactual market shares that would have realized had the GEMA shock not happened. We can estimate the counterfactual sales of newcomers and mainstream artists by dividing the observed total sales in the post period, i.e. $\hat{T}^*_j = T_j/(1+\hat{\delta}_j)$, and then calculate their counterfactual market shares as $\hat{m}^*_j = \hat{T}^*_j/(\hat{T}^*_j + \hat{T}^*_{-j})$. Realized market shares can be calculated accordingly, i.e. $\hat{m}_j = \hat{T}_j/(\hat{T}_j + \hat{T}_{-j})$. Finally, we can express counterfactual market shares as percentages of realized market shares such that $\hat{\delta}^m_j = \hat{m}_j/\hat{m}^*_j - 1$. Average total weekly sales in the post period for newcomers (defined as in the first year of their career) and niche artists with music videos are $T_{new} = 24301$ units and $T_{niche} = 22384$ units, respectively. With this, we can calculate that the market share of newcomer artists decreases by 13%, whereas the market share of niche artists increases by 20% as a result of the GEMA shock. Using more conservative estimates, i.e. from the respective confidence bands of other measures of "newcomer" and "niche", leads to the same qualitative conclusions. On average, these calculations suggest that YouTube helps increase the sales of newcomer artists about two thirds more than it helps increase the sales of mainstream artists.

After establishing the distributional changes across types of artists within the group of songs that have

[44]See section 6.3.1 for a detailed discussion of Hiller (2016) in relation to our findings.

Figure 3: Aggregate effects of online music video platforms

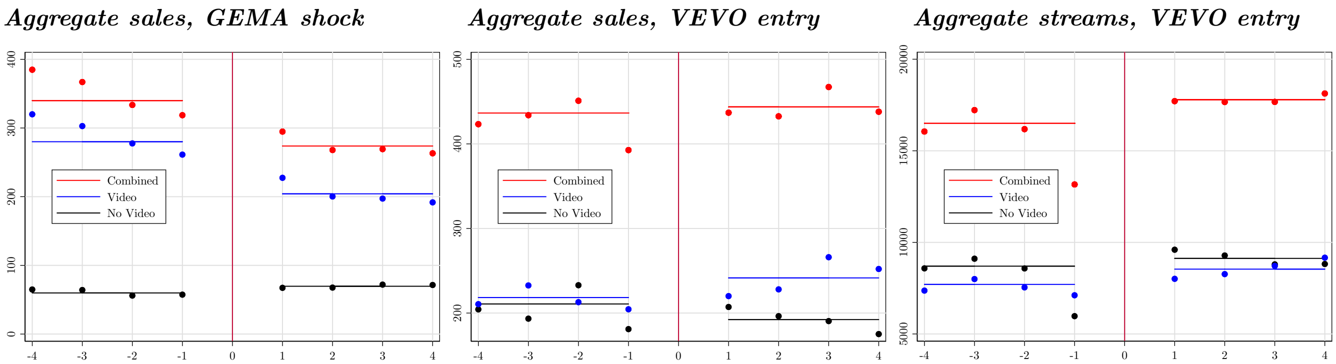

__Vertical axis__: Aggregate of physical and digital sales in thousand units.
__Horizontal axis__: Weeks prior/after GEMA shock/VEVO entry.
Dots indicate weekly aggregates, lines indicate averages in the before and after period.
__Black:__ Aggregate sales of songs without music videos. __Blue:__ Aggregate sales of songs with music videos.
__Red:__ Aggregate sales of all songs, with and without music videos.

music videos, we now change the level of aggregation and investigate distributional changes across songs with and without music videos. So far, a valid concern would be that the availability of online music videos is a zero-sum game, i.e. leading to changes in the composition of music sales without changing aggregate music sales. Further, as in Liebowitz (2007), a fallacy of composition could lead to effects in the aggregate that are the opposite of the relationship on the individual song level.

We first analyze this issue at the song-level. From visual inspection in Figures 1 and B.2, we see no indication that the trend in total sales/streams of songs without videos (our control group) changes after either the GEMA shock or the entry of VEVO. Accordingly, our results do not suggest that online music video platforms have externalities on the sales/streams of songs not available on these platforms. This is already indicative that the results we document are not merely driven by a transfer of demand from songs without videos to songs with videos. By looking at aggregate sales numbers, we can investigate this further. The plot of aggregate sales around the GEMA shock in the first panel of Figure 3 shows that also the aggregate sales trend of songs without videos does not change from the pre- to the post-period. Aggregate sales of songs with videos, however, decline substantially. Accordingly, we also see substantial decrease in total sales (songs with and without videos). Turning to aggregate sales in the time frame of the VEVO entry in the second panel of Figure 3, we see that songs without videos continue their negative sales trend, whereas aggregate sales of songs with videos increase. As a result, total record sales do not change by much. Finally, in the third panel of Figure 3, there is again no indication that the trend of aggregate streams of songs without videos changes with the entry of VEVO – especially if week t-1 is considered as an outlier. We do see, however, that aggregate streams of songs with videos increase. This leads to an increase in total streams (songs with and without videos).

Combining the insights from the disaggregated and the aggregated analysis, we conclude that the experiments we study in this paper do not only affect the composition of sales/streams, but also aggregated sales/streams. However, note that the analysis of aggregated sales is of course less rigorous than the analysis using disaggregated data as the plots in Figure 3 are ultimately only based on 8 data points.

6.2 Economic significance and welfare

Because it is difficult to attach a monetary value on the utility of free music consumption, it is difficult to derive an estimate of consumer surplus and therefore draw conclusions about overall welfare effects. However, we can calculate the average economic size of our estimates and derive estimates of total industry and artist surplus.

According to our price data, the average price for a song on a physical medium was 4.4 EUR and 1.1 EUR for digital songs after the GEMA shock. Multiplied by average total weekly sales units of songs with videos in the post period, this implies average total weekly revenues of about 360,000 EUR. With the most conservative estimate of a 6% reduction, this is 94% of the counterfactual revenues that would have been realized had the GEMA shock not happened. The monetary equivalent of 6% less sales than in the counterfactual world is thus about 23,000 EUR. Using the information that artists earn about 10 cent from a downloaded song and 13% of the physical revenue, we arrive at an estimate of a weekly decrease of about 2,500 EUR in total artist income from record sales.[45]

It is difficult to say much about the loss in royalty income to artists because we do not have access to song-level data on the number of streams on YouTube prior to the GEMA shock. However, we do have data on the number of streams on services like Spotify and Deezer for 2013. Hence, we can roughly calculate how the entry of VEVO has affected the surplus of the recorded music industry and the surplus

[45] See `http://www.informationisbeautiful.net/2010/how-much-do-music-artists-earn-online` and `http://www.bbc.com/news/magazine-23840744`

of artists. Before launching in October 2013, VEVO signed a licensing deal with GEMA.[46] Hence, artists benefit from an increase in record sales and receive royalties collected by GEMA from VEVO.

Using our data on the average weekly sales and prices in Germany after the VEVO entry, the most conservative estimate of a 2% increase in total record sales translates to an increase in total income of about 4,500 EUR, and an increase in total artist income from record sales of about 450 EUR. According to GEMA's official royalty rates schedule for ad-funded streaming offerings (VR-OD-9), the licensing fee for a highly interactive service such as VEVO is 0.00375 EUR per stream. We do not have exact data on the number of streams on VEVO, but data from the US can be helpful as an approximation. In 2013 there were 49.5 billion music streams, 22.4 billion (45%) from audio services, 27.1 billion (55%) from video services.[47] Our data on the average number of weekly (free & paid) audio streams in Germany does not include video streaming. Assuming the ratio in the US is the same in Germany, we can then calculate the hypothetical total royalties from video streaming as $(0.55/0.45) * audio * 0.00375$. The total number of audio streams of songs with VEVO videos in the post period is 7,894,000. This implies an increase in video streaming royalties worth roughly 730 EUR. We do not have data on VEVO's market share in the music video streaming market in Germany, but industry reports suggest that VEVO has had a market share of 2.8% in the overall online video market in 2016, which we can use as a conservative estimate.[48] We thus arrive at a lower bound of the average weekly increase in video streaming royalties of about 300 EUR. The average weekly increase in audio streaming royalties is about 245 Euro. Our estimate of the total increase in artist surplus is therefore some 700 EUR per week. However, this includes only artists that are represented by VEVO (about 70% of the market).

Now we can calculate the total lost artist surplus from the unavailability of online music video – the sum of lower record sales and foregone royalty income from video streaming. Assuming that the number of video streams is stable over time, foregone royalties of all artists amount to about 400 EUR per week. This implies that the promotional externalities of online music video (2,500 EUR total artist surplus) can offset forgone royalty income by a factor of six. In the 235 weeks between the GEMA blocking and the VEVO entry, our most conservative estimate is a total loss in artist welfare of about 0.6 million EUR,

[46] See for example https://www.ft.com/content/5bb2092e-117e-11e3-a14c-00144feabdc0)

[47] See https://www.musicbusinessworldwide.com/youtube-is-the-no-1-music-streaming-platform-and-getting-bigger/

[48] See https://de.statista.com/statistik/daten/studie/209329/umfrage/fuehrende-videoportale-in-deutschland-nach-nutzerant

and a loss in industry revenues of 5.4 million EUR. Note that both numbers do not include potential advertising revenues from YouTube.

6.3 External validity and contribution

6.3.1 Prior work on the effect of YouTube on recorded music sales

At face value, our baseline result seems in direct contrast to Hiller (2016), who concludes that a royalty dispute that led to the takedown of Warner Music content on YouTube is related to an increase in sales of albums released by Warner. However, a careful analysis of the differences and commonalities in the two studies reveals that some of Hiller's results are line with our findings.

We first discuss key differences between Hiller's and our approach and test some of the resulting empirical implications. We then try to replicate Hiller's approach as closely as possible in our setting and discuss alternative interpretations of Hiller's results and how they resonate with our findings.

Hiller's sample comprises the 200 bestselling albums in the US market. He does not link sales and corresponding video availability at the individual song level, but compares album sales of Warner artists to album sales of non-Warner artists. The distinction between songs and albums can have important implications. Since albums are bundles consisting of individual songs, YouTube availability, by allowing consumers to sample, may lead consumers to substitute purchases of the entire album in favor of purchases of an individual song which is promoted by a music video. In the absence of the sampling mechanism on YouTube, sales of other songs that are not promoted by a music video (or the entire album bundle) could therefore increase. We can least partly test this hypothesis.

In column (1) of Table 11 we investigate how the removal of video(s) for a specific song affects sales of other songs (that do not have videos) on the same release. Releases with multiple songs are typically extended singles, mini albums (EPs), and in the digital channel also entire albums. We find a positive and significant coefficient.[49] The point estimate is 22% (90% CI [7%,38%]). Hence, this suggests that Hiller's results can be driven by unbundling effects and substitutive effects within releases.

In columns (2)–(4) we test whether we can replicate Hiller's results when we aggregate sales at the release level. While this is not the same as actual album sales – especially not in the physical channel –, it can serve as a proxy of album sales. The first identification strategy is to compare releases that have at

[49]Note that this does not affect the identification strategy in the main analysis, see tests in section 5.1.

Table 11: Aggregate (release-level) effects of the GEMA shock

	(1) OtherSongs	(2) Release	(3) ReleaseTop	(4) Release	(5) ReleaseTop
After × OtherVideo	0.197**				
	(0.077)				
After × AtLeastOneVideo		-0.173**	-0.126		
		(0.078)	(0.223)		
Germany				0.435***	2.156***
				(0.043)	(0.187)
After × Germany				-0.024	0.272
				(0.021)	(0.166)
Observations	9919	6841	1180	25102	2259
$\overline{R^2}$	0.902	0.894	0.939	0.031	0.282

Dependent variables: (Log+1) weekly total sales in units, 2009.
Column (1) uses physical and digital sales of individual songs.
Column (2) uses physical and digital sales sales of the release bundle.
Column (3) uses physical and digital sales sales of the release bundle, only when songs are part of the top 200 list.
Column (4) uses physical sales of the release bundle.
Column (5) uses physical sales of the release bundle, only when songs are part of the top 200 list.
OtherVideo and *AtLeastOneVideo* indicate at least one (other) song on the same release has a video on U.S. YouTube, uploaded prior to April 1st, 2009.
Song and week fixed effects in column (1), release and week fixed effects in columns (2)–(4), constant not reported.
Standard errors in parentheses, clustered on the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

least one song with a video on YouTube to those that do not. The effect in column (2) is -16% (90% CI [-26%,-4%]), which is similar to our baseline effect. This implies that, in our setting, the song-level substitution effects which we find in column (1) are not strong enough to outweigh the release-level promotional effects of music videos.[50]. In column (3) we follow Hiller and restrict our sample to the sum of sales within a release while songs are in the top 200 list. This decreases precision of the estimates, leading to an effect -12% that is not significant (90% CI [-39%,24%]).

Finally, we adjust the identification strategy to better align with Hiller (2016). The equivalent to Hiller's approach of comparing sales of Warner artists to sales of non-Warner artists in our setting is to compare sales in Germany to sales in a comparable country, say Austria. The caveat is that we can only run this

[50] Note that this reconfirms our findings in section 6.1 regarding the fallacy of composition hypothesis.

analysis for sales in the physical channel, because sales data for the digital channel are not available for Austria. In column (4), we find a negative, yet imprecisely estimated coefficient for $After \times Germany$, with a point estimate of -2% (90% CI [-6%,1%]). Much like in section 5.2.2, the less precise identification strategy leads to smaller point estimates, but the sign of the effect remains. Finally, our closest replication of Hiller's approach in column (5) only considers the sum of sales within a release while songs are in the top 200 list. We find a positive coefficient for $After \times Germany$ which implies an effect of 31% (90% CI [-0.2%,73%]). This is very similar to the effect reported in column (1) of Table 4 in Hiller (2016). The stark disparity between the results in column (4) and (5) suggests that the empirical approach of Hiller, a combination of sample restrictions and measurement error in the definition of treatment and control group, may be responsible for the displacement effect that he finds. Further, it seems likely that Hiller's finding is mainly driven by the most popular albums.

Indeed, this can be seen directly by looking more closely at Hiller's results. He shows that albums that have a very successful debut face more displacement from YouTube videos, while the effect on lower debuting albums may be moderated by a promotional effect. In fact, the estimated interaction effect is so strong that the relationship is reversed and the promotional effect dominates for many albums in the sample. The results in Table 5 in Hiller (2016) suggest that *Debutrank* moderates the average effect such that the sign changes with *Debutrank* greater than values between 15.64 and 44.41 (depending on the specification, $x = -\beta_{Warnereffect}/\beta_{Warnerdebutrank}$), which is much lower than the average *Debutrank* of Warner albums in his sample (79.4). A valid alternative interpretation of Hiller's results therefore is that most albums' sales benefit from the added awareness generated through YouTube and only the most popular albums (those with very small debut ranks) experience net *displacement* rather than *promotion*. As we have shown in section 5.4, we do not find heterogeneity across the overall sales distribution (as Hiller does), but we do find that the promotional effect is driven by new artists and by artists in non-niche genres (which tend to rank slightly better, i.e. are more popular, when they enter the top 1000).

All this leads us to conclude that the differences between our results and Hiller's findings are largely due to the of level of aggregation, sample restrictions, and the choice of identification strategy. Hiller provides a partial view on the phenomenon at hand, while our analysis covers a broader range of the popular appeal of music content through the deeper analysis afforded by our broader and more detailed data, and our more fine-grained identification strategy. Our findings show important heterogeneity suggesting that the average effect hides sizable differences in the strength of one of the counteracting forces, the role of YouTube availability on music discovery. Thus, we offer a more detailed unpacking of the process of music discovery than Hiller (2016). This lets us draw more comprehensive policy implications below.

6.3.2 Policy implications

While our study is clearly limited in its scope, being a short-run measurement exercise that makes use of natural experiments that happened in specific contexts, it can still be useful to think about the policy implications that arise from our results.

First, the policy of some countries to mandate a share of broadcast music to be local (or in the local language) may be affected by the role of interactive and global platforms like YouTube. Second, US copyright law currently stipulates licensing fees for interactive digital platforms, but none for radio broadcasting, based on assumptions about the degree of substitution between (free) digital and paid content (Liebowitz, 2007; Lenard and White, 2015). Finally, from the perspective of the German right holders' association, the differential effects on domestic versus international artists may imply that it was acting to maximize their direct stakeholders' revenues. We elaborate on each of these domains below.

Several countries (e.g. France and Portugal) consider local music a cultural good in danger of being overrun and eventually replaced by international, often English-language content. This has led to quasi-protectionist policies stipulating a minimum percentage (often 40 or 50%) of local content on radio and TV broadcasts (Hervas-Drane and Noam, 2017). Our results suggest that while YouTube does not displace music sales (which would provide a direct, unregulated channel for international music to enter the local market), music discovery via YouTube favors mainstream (non-German in our case) artists substantially more than local ones.[51] This implies that policies to protect local (national) creative music industries are less effective in the presence of digital platforms for music discovery.

In the US, the rules of compulsory licensing imply that the music industry receives licensing fees from interactive digital platforms while radio broadcast is traditionally exempt from licensing. It is interesting to consider the historical reasoning behind the exemption: radio is a unidirectional medium on which

[51] The results discussed in section 5.4.3 imply that removing videos of non-German artists leads to a significantly lower sales (point estimate: −23%). The sales effect of removing videos of German artists is about three times smaller and not significant (point estimate: −8%).

users cannot choose particular songs and it was heavily influenced by the music industry, who offered payment or other inducements to radio stations to play particular songs (Coase, 1979). The perception was that radio acted as a promotional channel for actual purchases, relieving broadcasters of the need to compensate artists for playing their music because the compensation would come in the form of higher sales.[52] Conversely, digital platforms are on-demand and users can (at least to some extent) choose the songs they want to hear, which in principle allows for more substitution. Our findings suggest that digital platforms also fulfill an important promotional function, at least for some types of music. If replicated in other settings, such evidence might justify extending the license fee exemption to plays on digital platforms, and for firms, this result offers a case for differential online royalties, with new music especially prone to benefitting from exposure on digital platforms.

Finally, it is also worth revisiting the outcome for one of the instigators of the initial dispute, *GEMA*. GEMA's members are (virtually all) German artists, with international artists only associated by virtue of their membership to other national rights holders' associations. Our findings show that German artists suffered relatively less from the YouTube blackout and gained market share as a result.[53] Therefore, the removal of videos on YouTube may have (deliberately or not) resulted in the main constituency of GEMA increasing their market position vis-à-vis their international competitors.

7 Conclusions

In this paper, we exploit two natural experiments in the German market for online music videos to identify the effect of free sampling on sales of recorded music. The first experiment lets us identify the effect of removing access to online music videos on YouTube (in April 2009), while in the second we identify the effect of making official music videos available on the proprietary platform VEVO (in October 2013). Our analysis is based on a rich dataset that combines sales data that cover a large fraction of all music sales with song-level information on music video availability.

We find robust evidence that online videos are complementary to record sales. We believe that our findings

[52] The results in Liebowitz (2007) suggest that the opposite is true. He shows that radio is more of a substitute for the purchase of sound recordings than it is a complement.

[53] We can estimate market share changes in a similar fashion as described in section 6.1. Our results imply that the market share of non-German artists with videos decreased by 6% (average total realized sales are 130898 units), whereas the market share of German artists with videos increased by 13% (average total realized sales are 72864 units) due to the removal of music videos on YouTube.

carry some important general implications. Three results especially suggest to us that the promotional effect we robustly identify is not driven by the fact that YouTube is an open platform, but rather by the fact that YouTube offers differentiated content – music *video* rather than music. First, we show that YouTube's effect on music sales is not only driven by the availability of official music videos, but also by user-generated content on YouTube. Second, our results are very similar when we look at the entry of the closed platform VEVO. Third, we show that the promotional effect prevails when we look at the number of plays on audio streaming platforms, which should be a very close substitute to music video streaming. While we do not find the effect to differ based on overall popularity in terms of sales, we show that sales dynamics of newcomer artists are affected much more by the (non-)availability of online music videos than those of established artists. We also find that YouTube disproportionally favors music genres with greater mass appeal. We discuss several possible mechanisms consistent with the observed dynamics on the demand-side and on the supply-side.

In reference to a song by *The Buggles* – which happens to be the first music video shown on MTV in 1981 – we conclude that our study does not provide much evidence that "video killed the radio star".[54] If anything, we find the opposite. While it is straightforward to conclude that free consumption increases consumer surplus, conclusions about overall welfare are more difficult to reach. We calculate that positive externalities of access to an open platform like YouTube can offset forgone royalties (if YouTube did not pay any royalty fees to artists) income by a factor of six. Further, we show that the distributional effects are substantial. We calculate that YouTube helps to increase sales of newcomer artists by about two thirds more than it helps to increase sales of mainstream artists.

[54] According to music scholar Timothy Warner, the song's lyrics are "concerned with the adverse effect of technological change" (Warner, 2003, p. 47) on artists that used to be commercially successful (in the "Golden Age of Radio", Warner, 2003, p. 44). This relates to our study, which is is essentially an exercise of measuring the substitutionary or complementary effects of different distribution technologies.

References

Adermon, A., and Liang, C.-Y. (2014). "Piracy and Music Sales: The Effects of an Anti-Piracy Law." *Journal of Economic Behavior & Organization*, *105*, 90–106.

Aguiar, L., and Waldfogel, J. (2016). "Even The Loser Get Lucky Sometimes: New Products and The Evolution of Music Quality Since Napster." *Information Economics and Policy*, *34*, 1–15.

Aguiar, L., and Waldfogel, J. (2017a). "Quality Predictability and the Welfare Benefits from New Products: Evidence from the Digitization of Recorded Music." *Journal of Political Economy*, *forthcoming*.

Aguiar, L., and Waldfogel, J. (2017b). "Streaming reaches flood stage: Does spotify stimulate or depress music sales?" *International Journal of Industrial Organization*, *forthcoming*.

Anderson, C. (2006). *The long tail: How endless choice is creating unlimited demand*. Random House.

Autor, D. H. (2003). "Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing." *Journal of Labor Economics*, *21*(1), 1–42.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). "How Much Should We Trust Differences-In-Differences Estimates?" *The Quarterly Journal of Economics*, *119*(1), 249–275.

Bhattacharjee, S., Gopal, R., Lertwachara, K., Marsden, J., and Telang, R. (2007). "The Effect of Digital Sharing Technologies on Music Markets: A Survival Analysis of Albums on Ranking Charts." *Management Science*, *53*(9), 1359–1374.

Brynjolfsson, E., Hu, Y. J., and Simester, D. (2011). "Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales." *Management Science*, *57*(8), 1373–1386.

Brynjolfsson, E., Hu, Y. J., and Smith, M. D. (2010). "Long Tails vs. Superstars: The Effect of Information Technology on Product Variety and Sales Concentration Patterns." *Information Systems Research*, *21*(4), 736–747.

Cai, H., Chen, Y., and Fang, H. (2009). "Observational learning: Evidence from a randomized natural field experiment." *The American Economic Review*, *99*(3), 864.

Celma, Ò., and Cano, P. (2008). "From hits to niches?: or how popular artists can bias music recommendation and discovery." In *Proceedings of the 2nd KDD Workshop on Large-Scale Recommender Systems and the Netflix Prize Competition*, 5, ACM.

Coase, R. H. (1979). "Payola in radio and television broadcasting." *The Journal of Law and Economics*, *22*(2), 269–328.

Danaher, B., Dhanasobhon, S., Smith, M. D., and Telang, R. (2010). "Converting Pirates Without Cannibalizing Purchasers: The Impact of Digital Distribution on Physical Sales and Internet Piracy." *Marketing Science*, *29*(6), 1138–1151.

Danaher, B., Smith, M. D., Telang, R., and Chen, S. (2014). "The effect of graduated response anti-piracy laws on music sales: evidence from an event study in france." *Journal of Industrial Economics*, *62*(3), 541–553.

Datta, H., Knox, G., and Bronnenberg, B. (2017). "Changing their tune: How consumers' adoption of online streaming affects music consumption and discovery." *Marketing Science*, *forthcoming*.

Davidson, J., Liebald, B., Liu, J., Nandy, P., Van Vleet, T., Gargi, U., Gupta, S., He, Y., Lambert, M., Livingston, B., et al. (2010). "The youtube video recommendation system." In *Proceedings of the fourth ACM conference on Recommender systems*, 293–296, ACM.

Ferreira, F., and Waldfogel, J. (2013). "Pop internationalism: has half a century of world music trade displaced local culture?" *The Economic Journal*, *123*(569), 634–664.

Fleder, D., and Hosanagar, K. (2009). "Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity." *Management science*, *55*(5), 697–712.

Galasso, A., and Schankerman, M. (2015). "Patents and cumulative innovation: Causal evidence from the courts." *The Quarterly Journal of Economics*, *130*(1), 317–369.

Gambardella, A., Raasch, C., and von Hippel, E. (2016). "The user innovation paradigm: impacts on markets and welfare." *Management Science*, *forthcoming*.

George, L. M., and Peukert, C. (2014). "YouTube Decade: Cultural Convergence in Recorded Music." *NET Institute Working Paper 14-11*, SSRN–ID 2506357.

Godinho de Matos, M., Ferreira, P., Smith, M. D., and Telang, R. (2016). "Culling the Herd: Using Real-World Randomized Experiments to Measure Social Bias with Known Costly Goods." *Management Science*, *62*(9), 2563–2580.

Hervas-Drane, A. (2015). "Recommended for you: The effect of word of mouth on sales concentration." *International Journal of Research in Marketing*, *32*(2), 207–218.

Hervas-Drane, A., and Noam, E. (2017). "Peer-to-peer file sharing and cultural trade protectionism." *Information Economics and Policy*, *41*(12), 15–27.

Hiller, R. S. (2016). "Sales Displacement and Streaming Music: Evidence from YouTube." *Information Economics and Policy*, *34*, 16–26.

Hinz, O., Eckert, J., and Skiera, B. (2011). "Drivers of the long tail phenomenon: an empirical analysis." *Journal of management information systems*, *27*(4), 43–70.

Hosanagar, K., Fleder, D., Lee, D., and Buja, A. (2014). "Will the global village fracture into tribes? recommender systems and their effects on consumer fragmentation." *Management Science*, *60*(4), 805–823.

Hui, K.-L., and Png, I. (2003). "Piracy and the legitimate demand for recorded music." *The B.E. Journal of Economic Analysis & Policy*, (1), 11.

IFPI (2016). "Music consumer insight report." *Industry Report*, https://drive.google.com/open?id=0Bxe11iVXrXgsYXcxNUxUMFBOWHc.

Kretschmer, M. (2005). "Artists' Earnings and Copyright: A Review of British and German Music Industry Data in the Context of Digital Technologies." *First Monday*, *Special Issue 1*.

Kumar, A., and Tan, Y. (2015). "The demand effects of joint product advertising in online videos." *Management Science*, *61*(8), 1921–1937.

Lee, D., and Hosanagar, K. (forthcoming). "How do recommender systems affect sales diversity? a cross-category investigation via randomized field experiment." *Information Systems Research*.

Lee, Y.-J., Hosanagar, K., and Tan, Y. (2015). "Do i follow my friends or the crowd? information cascades in online movie ratings." *Management Science*, *61*(9), 2241–2258.

Lenard, T. M., and White, L. J. (2015). "Moving music licensing into the digital era: More competition and less regulation." *Working Paper*.

Liebowitz, S. J. (2007). "Don't Play it Again Sam: Radio Play, Record Sales, and Property Rights." *Working Paper*, SSRN–ID 956527.

Liebowitz, S. J. (2016). "How much of the decline in sound recording sales is due to file-sharing?" *Journal of Cultural Economics*, *40*(1), 13–28.

Liikkanen, L. A., and Salovaara, A. (2015). "Music on YouTube: User engagement with traditional, user-appropriated and derivative videos." *Computers in Human Behavior*, *50*, 108–124.

Liu, J.-H., Zhou, T., Zhang, Z.-K., Yang, Z., Liu, C., and Li, W.-M. (2014). "Promoting cold-start items in recommender systems." *PloS one*, *9*(12), e113457.

Nguyen, G. D., Dejean, S., and Moreau, F. (2014). "On the Complementarity between Online and Offline Music Consumption: The Case of Free Streaming." *Journal of Cultural Economics*, *38*(4), 315–330.

Oberholzer-Gee, F., and Strumpf, K. (2007). "The Effect of File Sharing on Record Sales: An Empirical Analysis." *Journal of Political Economy*, *115*(1), 1–42.

Oestreicher-Singer, G., and Sundararajan, A. (2012a). "Recommendation networks and the long tail of electronic commerce." *MIS Quarterly*, *36*(1).

Oestreicher-Singer, G., and Sundararajan, A. (2012b). "The Visible Hand? Demand Effects of Recommendation Networks in Electronic Markets." *Management Science*, *58*(11), 1963–1981.

Papies, D., Eggers, F., and Wlömert, N. (2011). "Music for free? how free ad-funded downloads affect consumer choice." *Journal of the Academy of Marketing Science*, *39*(5), 777–794.

Peitz, M., and Waelbroeck, P. (2006). "Why the Music Industry May Gain From Free Downloading – The Role of Sampling." *International Journal of Industrial Organization*, *24*(5), 907–913.

Peukert, C., Claussen, J., and Kretschmer, T. (2017). "Piracy and box office movie revenues: Evidence from megaupload." *International Journal of Industrial Organization*, *52*, 188–215.

Piolatto, A., and Schuett, F. (2012). "Music piracy: A case of 'the rich get richer and the poor get poorer'." *Information Economics and Policy*, *24*, 30–39.

Poort, J., and Weda, J. (2015). "Elvis is returning to the building: Understanding a decline in unauthorized file sharing." *Journal of Media Economics*, *28*(2), 63–83.

Raffo, J., and Lhuillery, S. (2009). "How to play the names game: Patent retrieval comparing different heuristics." *Research Policy*, *38*(10), 1617–1627.

Rob, R., and Waldfogel, J. (2006). "Piracy on the High C's: Music Downloading, Sales Displacement, and Social Welfare in a Sample of College Students." *Journal of Law and Economics*, *49*(4), 29–62.

Rob, R., and Waldfogel, J. (2007). "Piracy on the Silver Screen." *Journal of Industrial Economics*, *55*(3), 379–395.

Salganik, M. J., Dodds, P. S., and Watts, D. J. (2006). "Experimental study of inequality and unpredictability in an artificial cultural market." *Science*, *311*(5762), 854–856.

Sorensen, A. T. (2007). "Bestseller lists and product variety." *Journal of Industrial Economics*, *55*(4), 715–738.

Sun, M., and Zhu, F. (2013). "Ad revenue and content commercialization: Evidence from blogs." *Management Science*, *59*(10), 2314–2331.

Susarla, A., Oh, J.-H., and Tan, Y. (2012). "Social Networks and the Diffusion of User-Generated Content: Evidence from YouTube." *Information Systems Research*, *23*(1), 23–41.

Takeyama, L. N. (1994). "The Welfare Implications of Unauthorized Reproduction of Intellectual Property in the Presence of Demand Network Externalities." *Journal of Industrial Economics*, *42*(2), 155–166.

Tang, Q., Gu, B., and Whinston, A. B. (2012). "Content contribution for revenue sharing and reputation in social media: A dynamic structural model." *Journal of Management Information Systems*, *29*(2), 41–76.

Thomes, T. P. (2013). "An Economic Analysis of Online Streaming Music Services." *Information Economics and Policy*, *25*(2), 81–91.

Tucker, C. E., and Zhang, J. (2011). "How Does Popularity Information Affect Choices? A Field Experiment." *Management Science*, *57*(5), 828–842.

Warner, T. (2003). *Pop Music – Technology and Creativity*. Ashgate.

Wen, W., Forman, C., and Graham, S. J. (2013). "The impact of intellectual property rights enforcement on open source software project success." *Information Systems Research*, *24*(4), 1131–1146.

Williams, H. L. (2013). "Intellectual property rights and innovation: Evidence from the human genome." *Journal of Political Economy*, *121*(1), 1–27.

Wlömert, N., and Papies, D. (2016). "On-Demand Streaming Services and Music Industry Revenues – Insights from Spotify's Market Entry." *International Journal of Research in Marketing*, *33*, 314–327.

Zentner, A. (2006). "Measuring the effect of file sharing on music purchases." *The Journal of Law and Economics*, *49*(1), 63–90.

Zhang, L. (2017). "Intellectual Property Strategy and the Long Tail: Evidence from the Recorded Music Industry." *Management Science*, forthcoming.

Zhou, R., Khemmarat, S., and Gao, L. (2010). "The impact of youtube recommendation system on video views." In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 404–410, ACM.

Zhou, R., Khemmarat, S., Gao, L., Wan, J., and Zhang, J. (2016). "How youtube videos are discovered and its impact on video views." *Multimedia Tools and Applications*, *75*(10), 6035–6058.

A Figures and Tables

Figure A.1: Google search volume for GEMA and VEVO

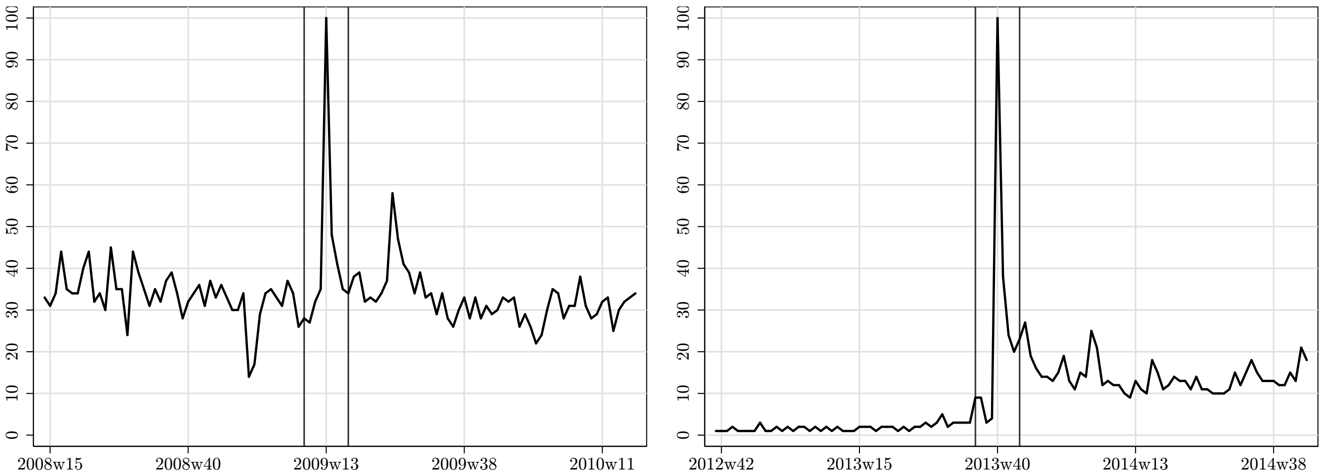

Relative Google search volume for "gema"/"vevo" in Germany, April 2008 – April 2010 and October 2012 – October 2014. Vertical lines indicate the sample period for the econometric analysis.
Source: *Google Trends.*

Table A.1: Variable definition

Variable	Definition
Dependent Variables	
Total (ln+1)	Weekly number of sold units (physical + digital)
Physical (ln+1)	Weekly number of physically sold units
Digital (ln+1)[a]	Weekly number of digitally sold units
Streaming: Total (ln+1)[b]	Weekly number of streams (free + premium)
Streaming: Free (ln+1)[b]	Weekly number of streams on free streaming services (not including YouTube and VEVO)
Streaming: Premium (ln+1)[b]	Weekly number of streams on subscription-based streaming services
Independent Variables	
After (0/1)	Weeks after 2009w14/2013w40 (GEMA shock/entry of VEVO)
Video (0/1)	(At least one) song-specific video on U.S. YouTube uploaded prior to April 1st, 2009/by VEVO
Germany (0/1)	Data from Germany
VEVO Label (0/1)	Universal Music or Sony Music (and their subsidiaries)
UGC: Small Share (0/1)	Share of official song-specific videos is bigger or equal to the average song in the sample
UGC: Official and UGC (0/1)	Share of official song-specific videos is smaller than the average song in the sample
UGC: Only UCG (0/1)	Only user-generated song-specific videos
Newcomer: 2 Months (0/1)	Earliest release of artist not older than 2 months before GEMA shock
Newcomer: No album (0/1)	Artist has never released an album before the GEMA shock
Newcomer: 1st Year (0/1)	Earliest release of artist not older than 1 year before GEMA shock
Niche: Never US (0/1)	Artist never appeared on the US charts (album and single), 2000w1–2009w14/2000w1–2013w36
Niche: German (0/1)	Artist has German origin
Niche: Genre (0/1)	Song is not in a main stream genre (i.e. not Pop or Rock)

[a] Not available for Austria in 2009.
[b] Only available for Germany in 2013.

Table A.2: Descriptive statistics

	No Video		Video		Total	
	Mean	SD	Mean	SD	Mean	SD
Sample of the GEMA shock, 2009w10–2009w18						
Log(Total+1)	1.783	2.277	3.575	2.168	2.764	2.390
Log(Physical+1)	0.507	1.306	0.910	1.675	0.728	1.532
Log(Download+1)	1.673	2.229	3.482	2.158	2.664	2.369
After	0.451	0.498	0.449	0.497	0.450	0.497
Video	0	0	1	0	0.548	0.498
UGC: Small Share	0	0	0.278	0.448	0.152	0.359
UGC: Official and UGC	0	0	0.221	0.415	0.121	0.326
UGC: Only UCG	0	0	0.501	0.500	0.275	0.446
Newcomer: 2 Months	0.028	0.165	0.045	0.206	0.037	0.189
Newcomer: No Album	0.046	0.210	0.040	0.195	0.043	0.202
Newcomer: 1st Year	0.065	0.247	0.065	0.246	0.065	0.247
Niche: Never US	0.821	0.383	0.642	0.480	0.723	0.448
Niche: German	0.616	0.486	0.470	0.499	0.536	0.499
Niche: Genre	0.469	0.499	0.355	0.479	0.407	0.491
Number of songs	700		842		1542	
Number of artists	535		549		999	
Observations	6202		7509		13711	
Sample of the VEVO entry, 2013w36–2013w44						
Log(Total+1)	3.333	2.058	3.869	2.030	3.576	2.062
Log(Physical+1)	0.199	0.796	0.207	0.835	0.203	0.814
Log(Download+1)	3.315	2.065	3.857	2.031	3.561	2.068
Log(Total Streams+1)	4.776	4.079	5.583	3.908	5.142	4.023
Log(Free Streams+1)	3.150	4.118	3.726	4.240	3.411	4.184
Log(Premium Streams+1)	4.547	3.832	5.323	3.669	4.899	3.779
After	0.588	0.492	0.574	0.495	0.582	0.493
Video	0	0	1	0	0.454	0.498
VEVO Label	0.016	0.127	0.021	0.144	0.019	0.135
Number of songs	937		759		1696	
Number of artists	631		485		1027	
Observations	7962		6614		14576	

Table A.3: Top 20 songs, with and without videos

Video		No video	
Peak Rank	Song	Peak Rank	Song
1	Lady Gaga - Poker face	25	James Morrison - Broken strings
6	Razorlight - Wire to wire	29	Lady Gaga - Just dance
6	Silbermond - Irgendwas bleibt	46	Eisblume - Eisblumen
21	Peter Fox - Haus am See	106	Jason Mraz - Lucky
26	Kings of Leon - Use somebody	109	Sido - Beweg dein Arsch
31	Polarkreis 18 - Allein allein	118	Big Ali - Hit the floor
38	Tim Toupet - So ein schöner Tag	118	Kid Cudi - Day 'n' nite
45	Peter Fox - Schwarz zu blau	121	Curse - Bis zum Schluss
52	Jeanette - Undress to the beat	139	Uschi Blum - Sklavin der Liebe
54	Lily Allen - The fear	158	Sido - Nein
58	Katy Perry - Thinking of you	167	Bushido - Für immer jung
60	Amy Macdonald - This is the life	175	Queensberry - I can't stop feeling
63	Reamonn - Million miles	177	T.I. - Live your life
64	Beyoncé - Halo	180	Michael Wendler - I don't know
64	Rihanna - Rehab	186	The Pussycat Dolls - Bottle pop
67	Die Toten Hosen - Alles was war	191	Axel Fischer - Amsterdam
69	Kevin Rudolf - Let it rock	204	Jem - It's amazing
73	Maria Mena - All this time	222	Xavier Naidoo - Wann
81	MGMT - Kids	230	Jason Mraz - I'm yours
84	Coldplay - Viva la vida	231	Mark Ronson - Valerie

Note: Video is defined as at least one video that matches artist name and up to three words of the song title appearing in the top 20 (according to "relevance") search results for "song - title" on U.S. YouTube, with an upload date prior to April 1st, 2009. Data collection was carried out on April 15, 2015 through the YouTube API v2.

B Technical appendix

B.1 Measurement error

B.1.1 Estimation window

One might be worried that the results are sensitive to our choice of the estimation window (four weeks after the GEMA shock). In Figure B.1, we report the point estimate and 90% confidence band of *After* $\times$ *Video* in samples of 1–9 weeks after the GEMA shock. The estimates are statistically indistinguishable from the baseline results reported in Table 3 (depicted at week 4 in Figure B.1). Note that the definition of *Video* is slightly different for $>$ 4 weeks after the GEMA shock, because the data from YouTube was collected at a different point in time (Oct, 4, 2016 vs. Apr, 15, 2015). This may explain why the point estimates in the samples of 5, 6 and 7 weeks after the experiment are smaller than the point estimate we obtain in the sample of 4 weeks after the experiment. The upper end of the 90% confidence bands varies between -8% and -4%, with an average of -6%, which is equivalent to the the estimate of -6% from our baseline specification which we use for the welfare calculations in section 6.2.

Figure B.1: Shorter and longer estimation windows, GEMA shock

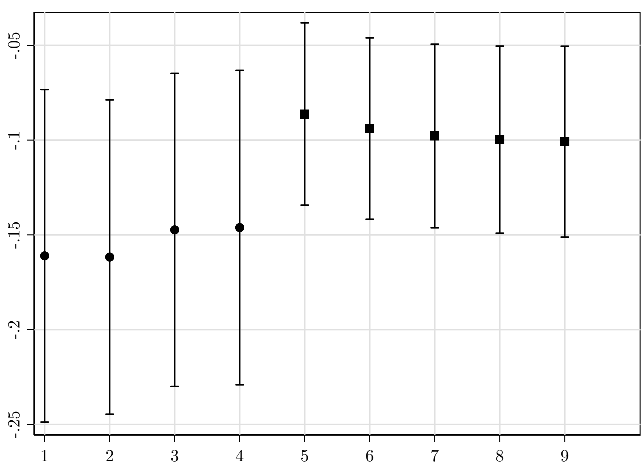

Note: *Point estimates and 90% confidence bands of* After $\times$ Video *estimated according to equation (1), with log(1+total sales) as dependent variable, on varying samples. The before period is defined as 4 weeks prior to the GEMA shock, the vertical axis indicates the number of weeks in the after period. Note that the definition of* Video *is slightly different for* $>$ *4 weeks after the GEMA shock, because the data from YouTube was collected at a different point in time (Oct, 4, 2016 vs. Apr, 15, 2015).*

B.1.2 Measurement error in the treatment group

Our measure of video availability is not perfect because we do not have exact historical data, but instead have to rely on what YouTube returns as search results some time after the respective experiments happened. This may lead to measurement error issues, because some music videos might have been removed, and/or uploaded with a new version. However, we argue that neither type I or type II errors in the definition of our treatment group can change the sign of the effect estimate, nor overestimate the effect in a difference-in-differences model.

Let V be the true set of songs with online music videos prior to the GEMA shock. Assume that we cannot observe V, but V^*, an imprecise measure of the treatment group. Type I errors (false positives) occur if song $i \in V^*$ but $i \notin V$, type II errors (false negatives) occur if $i \in V$ but $i \notin V^*$. An example of a type I error is a scenario where not all (newly uploaded) music videos are immediately blocked after the GEMA shock. An example of a type II error is a scenario where music videos have been removed before we collect our data (because they were blocked in a specific country and the uploader decided to remove them).

Let us first consider the case where the process that makes $V \neq V^*$ is *independent of the outcome variable.* Assume the true effect is positive. If the measurement error is of type I, the estimate of the treatment group average in the post-period is too low (biased towards the mean of the true control group), because we incorrectly take untreated observations into account, that are not affected by definition. The estimate of the control group average remains unbiased. As a result, the difference between treatment and control group averages (difference-in-differences estimate) is too low (biased towards zero). If the measurement error is of type II, the estimate of the control group average in the post-period is too high (biased towards the average of the true treatment group), because we incorrectly take treated observations into account. As a result, the difference between treatment and control group averages is too low (biased towards zero). The bias is symmetric when we assume that the true effect is negative.

Now consider the perhaps more realistic case where measurement error is *correlated with pre-experiment values of the outcome variable.* Assume that the true effect is positive and there is a negative correlation between measurement error and the pre-experiment outcome. The bias is symmetric if we consider the case where the true effect is negative. If type I or type II errors are correlated with high pre-period values of the outcome variable, the bias is stronger, and the diff-in-diff estimate is closer to zero. If the measurement error is of type I, the estimate of the treatment group average in the post-period is too low (biased towards the mean of the true low pre-period values of the control group), because we incorrectly take untreated observations into account. The estimate of the control group average is biased upwards, because we only take higher values into account. As a result, the difference between treatment and control group averages is biased towards zero. If the measurement error is type II, the estimate of the control group average in the post-period is too low (biased towards the average of the true treatment group), because we incorrectly take treated observations (yet only those with low pre-period values of the outcome variable) into account. Again, this means that the diff-in-diff estimate is biased towards zero.

In results available on request, we can show that the estimated effect of the GEMA shock is very similar when we use information from Austrian YouTube that was collected more recently (October 4th, 2016) than the information from US YouTube that we collected on April 15th, 2015. The point estimate of the effect of the GEMA shock on total sales in this specification is -13% (90% CI $[-19\%, -6\%]$). That is, the estimate is not different from that in our preferred specification in Table 3 in the sense that 90% confidence bands overlap.

B.2 Parallel trends assumption: subgroups

B.2.1 User generated content

For extraneous reasons, some songs may attract more user-generated content on YouTube than others. If this unobserved heterogeneity is time-invariant, it should be captured by the song-fixed effects in our baseline model. However, if consumers for example create more or less cover versions as a song becomes more or less popular over time, our estimates in Table 6 can be biased. Similar to the test of the parallel trends assumption described in section 5.1, we can estimate whether sales dynamics of treatment and control songs vary by the amount of user-generated content on YouTube before the GEMA shock. Results from a regression model similar to equation (2) are reported in Table B.1. With the exception of a few coefficients, this exercise suggests that there are no pre-existing trends or time-variant differences across songs with more or less user-generated content on YouTube.

Table B.1: Group differences in the pre-period, UGC

	(1) Total		(2) Physical		(3) Digital	
t_{-4} × Small UGC Share	0.168*	(0.089)	0.062	(0.062)	0.151*	(0.089)
t_{-3} × Small UGC Share	0.133*	(0.077)	-0.003	(0.055)	0.160**	(0.077)
t_{-2} × Small UGC Share	0.099	(0.065)	-0.029	(0.051)	0.113*	(0.065)
t_{-1} × Small UGC Share	0.014	(0.059)	-0.055	(0.047)	0.029	(0.059)
t_{-4} × Official and UGC	0.072	(0.100)	0.046	(0.088)	0.042	(0.104)
t_{-3} × Official and UGC	0.032	(0.096)	-0.023	(0.079)	0.055	(0.095)
t_{-2} × Official and UGC	-0.018	(0.085)	0.031	(0.051)	-0.002	(0.084)
t_{-1} × Official and UGC	0.004	(0.081)	0.077**	(0.039)	0.004	(0.081)
t_{-4} × Only UGC	0.102	(0.089)	0.096*	(0.054)	0.091	(0.087)
t_{-3} × Only UGC	0.094	(0.078)	0.003	(0.052)	0.120	(0.076)
t_{-2} × Only UGC	0.025	(0.075)	-0.033	(0.048)	0.049	(0.072)
t_{-1} × Only UGC	-0.034	(0.057)	-0.010	(0.035)	-0.018	(0.057)
Observations	7543		7543		7543	

Dependent variables: (Log+1) weekly total sales in units.
Song and week fixed effects, constant not reported. Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

B.2.2 Newcomer and niche artists

Unobserved time-variant heterogeneity may be correlated to newcomer and niche status such that the results reported in section 5.4 are biased. To test this, we estimate a regression model similar to equation (2), described in section 5.1. Table B.2 shows that newcomer artists follow similar, in fact statistically mostly indistinguishable trends, in the four weeks before the GEMA shock. We carry out a similar exercise regarding our measures of niche artists/music in Table B.3, with very similar results. Looking jointly at the $t_{-\tau}\times$ Video × Niche coefficients, we do not find that sales of niche artists/music follow a different time trend than sales of mainstream artists/music before the GEMA shock. With just a few exceptions, the coefficients are not significantly different from zero. In total, the tests in Tables B.2 and B.3 strongly suggest that the potential endogeneity of our empirical measures of newcomer status, consumer awareness, and niche status are unlikely to cause substantial estimation bias.

Table B.2: Group differences in the pre-period, Newcomer

	(1) 2 Months	(2) No Album	(3) 1st Year
t_{-4} × Video	0.116*	0.114*	0.114*
	(0.068)	(0.068)	(0.069)
t_{-3} × Video	0.073	0.070	0.072
	(0.054)	(0.054)	(0.055)
t_{-2} × Video	-0.005	0.000	-0.003
	(0.044)	(0.044)	(0.045)
t_{-1} × Video	-0.011	-0.012	-0.013
	(0.035)	(0.035)	(0.036)
t_{-4} × Video × Newcomer	0.236	0.303	0.179
	(0.304)	(0.339)	(0.234)
t_{-3} × Video × Newcomer	0.082	0.164	0.067
	(0.276)	(0.315)	(0.217)
t_{-2} × Video × Newcomer	0.145	0.020	0.069
	(0.299)	(0.273)	(0.172)
t_{-1} × Video × Newcomer	0.143	0.182*	0.136*
	(0.093)	(0.106)	(0.073)
Observations	7543	7543	7543

Dependent variables: (Log+1) weekly total sales in units.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded by prior to April 1st, 2009.
Not Hit includes artists that never appeared on the German top 100 single charts before the GEMA shock. Song and week fixed effects, constant not reported. Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

B.3 Robustness checks regarding VEVO entry

B.3.1 Parallel trends assumption

Very much like in section 5.1, we can test whether the parallel trends assumption is consistent with the data also regarding the VEVO entry. First, we can plot the dependent variable over time. In Figure B.2, we partial out song-fixed effects and plot averaged residuals for each group and each week, regarding total sales, physical sales, digital sales, total streaming, free streaming, and premium streaming. The plot shows that treatment and control group follow similar trends before the entry of VEVO and start to diverge substantially afterwards. The 90% confidence bands (calculated using the standard error of the

Table B.3: Group differences in the pre-period, Niche

	(1) Billboard	(2) German	(3) Genre
t_{-4} × Video	0.061	0.044	0.042
	(0.046)	(0.041)	(0.036)
t_{-3} × Video	0.064	0.031	0.052
	(0.040)	(0.036)	(0.033)
t_{-2} × Video	0.027	0.014	0.049
	(0.034)	(0.028)	(0.031)
t_{-1} × Video	-0.020	-0.002	0.035
	(0.028)	(0.023)	(0.027)
t_{-4} × Video × Niche	-0.024	0.004	-0.048
	(0.049)	(0.048)	(0.115)
t_{-3} × Video × Niche	-0.070	-0.023	-0.086
	(0.044)	(0.041)	(0.096)
t_{-2} × Video × Niche	-0.043	-0.031	-0.164*
	(0.040)	(0.039)	(0.096)
t_{-1} × Video × Niche	0.023	-0.008	-0.189***
	(0.034)	(0.032)	(0.067)
Observations	15086	15086	15086

Dependent variables: (Log+1) weekly total sales in units.
Video indicates (at least one) song-specific video on U.S. YouTube, uploaded by prior to April 1st, 2009.
Song and week fixed effects, constant not reported. Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

mean) overlap in the majority of the pre-period leads us to conclude that the parallel trends assumption is consistent with our data. More formally, we can test whether the difference in the dependent variable across points in time is zero in the pre-period. Table B.4 shows estimates of β_1^t coefficients (defined as in equation 2) for all dependent variables in the pre-period. Across all columns and most pre-experiment weeks, we cannot reject the hypothesis that the difference in sales in the treatment and control group is equal to zero.

B.3.2 Alternative identification strategy

The alternative identification strategy that we introduced in section 5.2.2 can also be utilized in the context of the VEVO entry. We can treat all songs in Germany as having videos on the VEVO platform,

Figure B.2: Trends of treatment and control group, before and after VEVO entry

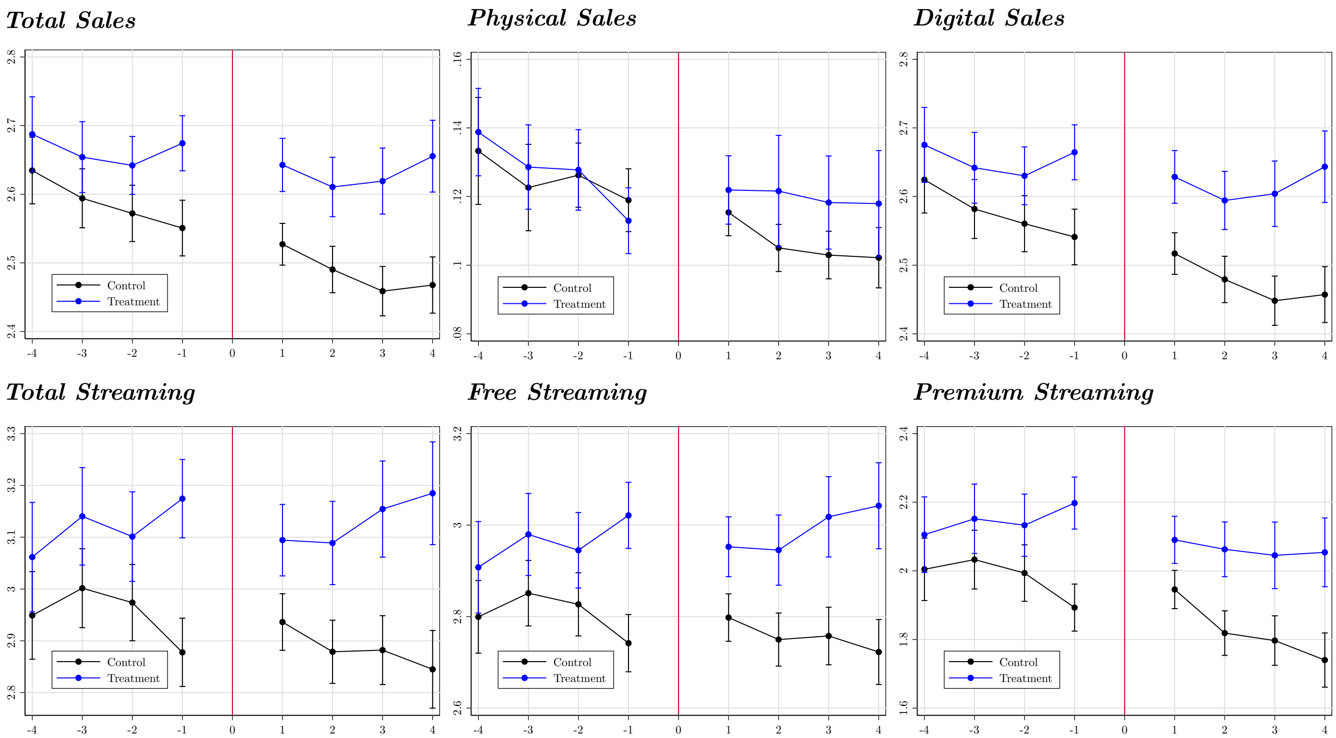

Vertical axis: *Average demeaned total/digital/physical sales or total/free/premium streaming, i.e. averaged residuals* $\overline{\hat{y}_{vt}} = \overline{\hat{y}_{it} - \hat{\mu}_i}$ *derived from the model* $y_{it} = log(Sales_{it}+1) = \alpha + \sum_t \beta_0^t w_t + \sum_t \beta_1^t (w_t \times Video_i) + \mu_i + \varepsilon_{it}$ *for* $Video_i = 0$ *and* $Video_i = 1$.
Horizontal axis: *Weeks prior/after October 1st, 2013.*
Black (control group): *Average sales of songs without at least one video uploaded to U.S. YouTube by VEVO.*
Blue (treatment group): *Average sales of songs with at least one video uploaded to U.S. YouTube by VEVO.*
Bars indicate 90% confidence bands (standard error of the mean).

and contrast their sales in Germany to their sales in Austria. Because the entry of VEVO is by definition is a much smaller experiment than GEMA shock in that it only affected some record labels in the market, treating all sales in Germany equally may introduce too much noise. This is what column (1) of Table B.5 suggests. We do not find evidence that the average song had significantly different sales in Germany compared to Austria before and after the entry of VEVO. We therefore take a different route and define the treatment group to include artists that are contracted by Universal Music or Sony Music (and their subsidiaries). Column (2) shows that we get similar (not significantly different) results concerning the effect of VEVO on total sales as in our baseline specification in column (1) of Table 7. We can also qualitatively replicate our baseline results regarding total streams in column (3).

Table B.4: Group differences in the pre-period, VEVO entry

	(1) Total	(2) Physical	(3) Digital	(4) Total	(5) Free	(6) Premium
t_{-4} × Video	-0.052	0.010	-0.055	0.018	0.027	0.021
	(0.072)	(0.026)	(0.072)	(0.115)	(0.143)	(0.109)
t_{-3} × Video	-0.028	0.004	-0.024	0.056	0.065	0.049
	(0.067)	(0.023)	(0.067)	(0.103)	(0.140)	(0.098)
t_{-2} × Video	0.017	-0.001	0.018	0.049	0.116	0.044
	(0.058)	(0.019)	(0.058)	(0.095)	(0.125)	(0.090)
t_{-1} × Video	0.079	-0.000	0.075	0.142*	0.174*	0.138*
	(0.052)	(0.016)	(0.052)	(0.079)	(0.099)	(0.075)
Observations	7792	7792	7792	7792	7792	7792
$\overline{R}^2$	0.865	0.903	0.865	0.920	0.878	0.917

Dependent variables: (Log+1) weekly sales/streams in units.
Video indicates (at least one) song-specific video on U.S. YouTube uploaded by VEVO.
Only weeks before the respective experiments. Song and week fixed effects, constant not reported.
Standard errors in parentheses, clustered at the song-level.
* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

Table B.5: Alternative identification strategy, VEVO entry

	Sales		Streams
	AT+DE (1)	DE (2)	DE (3)
Germany	1.966***		
	(0.059)		
After × Germany	0.018		
	(0.021)		
After × VEVO Label		0.073*	0.215***
		(0.042)	(0.080)
Observations	29152	14576	14576
$\overline{R}^2$	0.560	0.858	0.898

Dependent variables: (Log+1) weekly sales/streams in units (streams are not available for Austria).
Song and week fixed effects, constant not reported.
Standard errors in parentheses, clustered at the song-level. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

Falsification exercises To further ascertain the robustness of our results, we perform two different falsification exercises. First, we estimate placebo-versions of our model pretending that the experiments

Table B.6: Falsification exercises: Placebo experiments, VEVO entry

	Timing				Country
	Sales		Streaming		Sales
	−2 weeks	+2 weeks	−2 weeks	+2 weeks	Austria
	(1)	(2)	(3)	(4)	(5)
After × Video	0.042	0.101*	0.128	0.242**	-0.011
	(0.056)	(0.060)	(0.099)	(0.098)	(0.034)
Observations	4917	3844	4917	3844	14576
$\overline{R}^2$	0.902	0.894	0.932	0.943	0.754

Dependent variables: (Log+1) total weekly sales/streams in units,
in Germany (columns 1–4) and Austria (column 5). Data on streaming are not available for Austria.
Video indicates (at least one) song-specific video on U.S. YouTube uploaded by VEVO.
Song and week fixed effects, constant not reported.
Standard errors in parentheses, clustered at the song-level. * $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

took place either two weeks before or two weeks after they actually did. As illustrated in Figure 2, we split the sample of our baseline model (running from t_{-4} to t_{+4}) in two parts and estimate models on the two subsamples running from t_{-4} to t and from t to t_{+4}, setting the respective dates of the placebo-experiments to t_{-2} and t_{+2}. As described in section 5.1, we can think of the t_{-2} placebo experiment as a test of pre-existing trends, whereas we can interpret the t_{+2} placebo experiment as a test of the dynamics of the true effect. In the results reported in Table B.6, we define the after period to include the week of placebo-experiment, but we get very similar estimates if we treat the placebo-experiment week as part of the before period. The diff-in-diff coefficients across columns (1) and (3) are not significantly different from zero, suggesting that there is no general trend that can explain our main results. However, the coefficient in column (3) is relatively large, suggesting that there might be a trend that started before the entry of VEVO. The coefficients in columns (2) and (4) are larger and significant. This implies that some weak evidence that the effect is growing over time (in the short run), which seems to be somewhat consistent with Figure B.2.

Finally, we run a counterfactual experiment in which we pretend that VEVO had entered in Austria and compare sales (only in Austria) of songs that have a video available on the VEVO platform (in Germany) to songs that do not. As expected, we find a small and non-significant coefficient in column (5) of Table B.6, which further supports our main results.

Note that Table B.6 only reports results concerning total sales/streams as dependent variable, but we also do not find significant effects when we distinguish between physical and digital sales, and free and premium streaming.

B.4 Price changes

Figure B.3: Average song prices over time

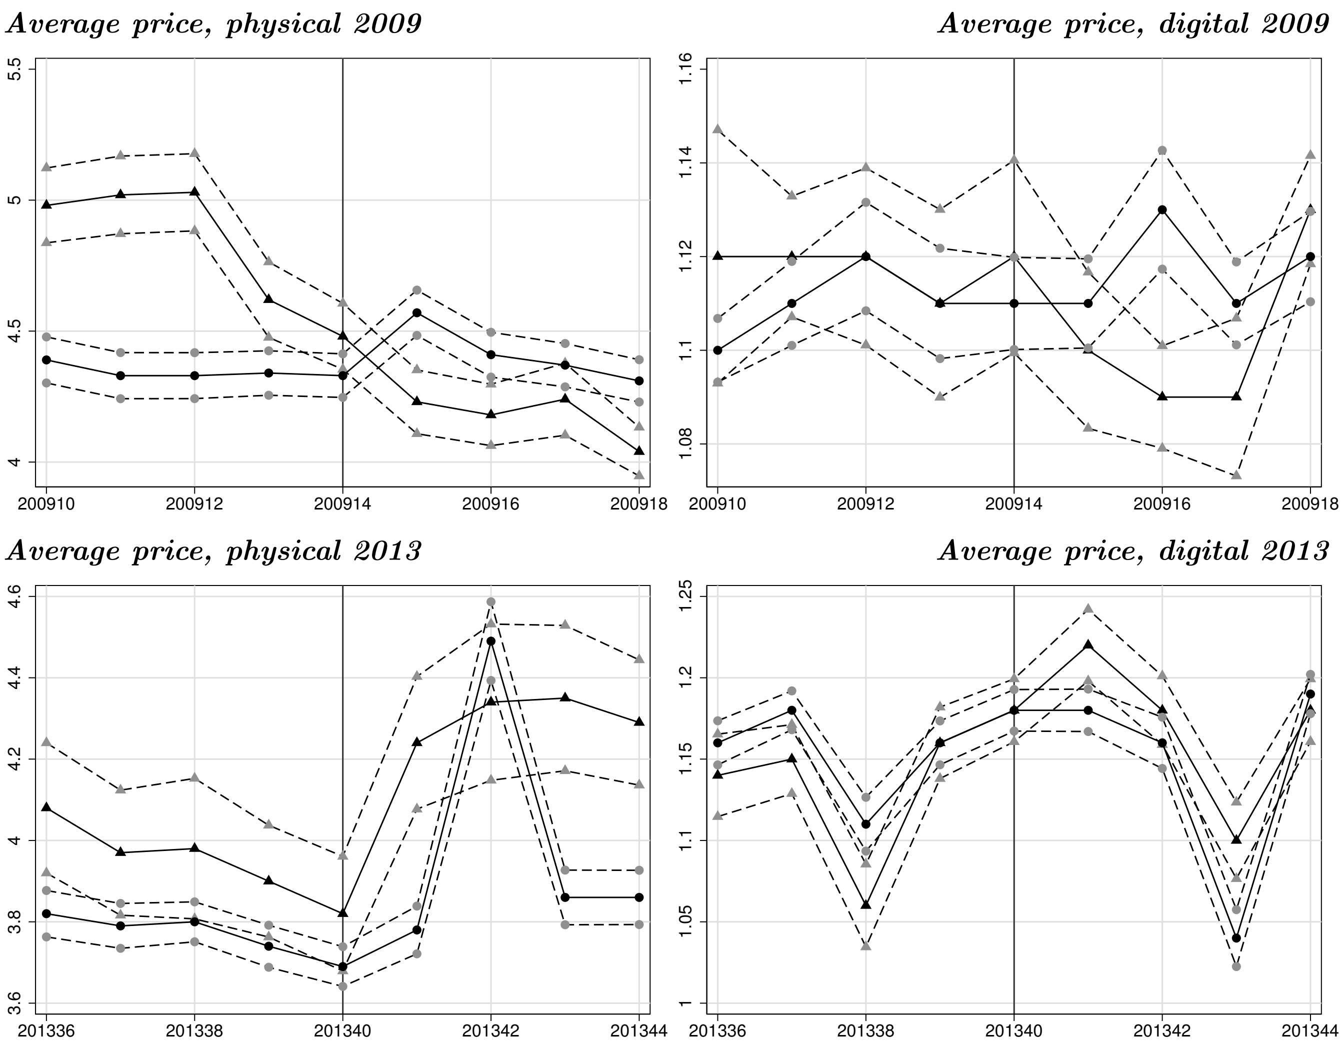

—●— Songs with videos, —▲— Songs without videos, dashed lines represent 90% confidence bands.
Data source: *GfK Entertainment*

If record labels or retailers reacted to the GEMA shock by increasing the prices of songs with blocked videos (perhaps in the erroneous belief that videos heavily substitute for sales and the disappearance of a close substitute relieves downward pressure on prices), our diff-in-diff estimates would not be due to a lack of promotion through videos, but due to the change in relative prices. Similarly, price changes might explain the positive effect of VEVO's entry. For privacy reasons, GFK Entertainment does not provide weekly price data at the individual song level. However, we were able to obtain average weekly prices, along with standard errors of the mean, for our treatment and control groups. In that way we can compute confidence bands and test whether group means differ before and after the GEMA shock and VEVO entry (see Figure B.3). Although the data suggests that prices vary over time (to a much lesser degree in the digital than in the physical channel), there is no clear pattern of change that coincides with the respective experiments (especially in the digital channel). Most importantly, changes in prices of songs with videos are overall not significantly different from changes in prices of songs without videos in the sense that the respective 90% confidence bands overlap.