Asymmetric content moderation in search markets: The case of adult websites

Leonardo Madio (UNIPD) Matthew Mitchell (Toronto)
Martin Quinn (Rotterdam School of Management) Carlo Reggiani (JRC & Manchester)

Toulouse Digital
January 2025

Content Sharing and Moderation is Central to Many Platforms

- Platforms with user-uploaded content are some of the biggest platforms
  - Traditional Social Media (Facebook, Twitter, Instagram, TikTok etc)
  - Video (YouTube etc)
  - Audio (SoundCloud, Spotify etc.)
  - Answers and Bulletin Boards (Quora, Reddit)
  - Even some selling sites have similar features (Ebay, Amazon, Etsy)
- Big issue for these platforms (and regulators) is moderating content
  - To keep content quality high for consumers (Hidden ads)
  - To protect content creators (Fake posters)
  - *Social welfare concerns for third parties*

Content Moderation

- Platforms can moderate content with a combination of what they allow people to see (steering) and what messages they attach (certification).

Content Regulation is an Important Policy Issue

- Various legislation in an attempt to put some regulatory oversight
  - EU Digital Services Act
  - UK Online Safety Bill
  - Canada Online Harms Act
  - Disclosure Regulations in Many Places
- Policy concerns in many areas around third party harm in user generated content:
  - Hate speech
  - Misinformation
  - IP infringement
  - *Illegal/Non-consensual pornography*
- Regulatory methods for pornography include
  - Content moderation and age restrictions under DSA
  - HB 1181 (age restriction) in Texas and elsewhere in the US

Regulation of Content Moderation is Incomplete

- Can only regulate content on some platforms, e.g. VLOP under DSA
- *Does regulating content only on some websites just push the regulated content to less regulated websites?*
  - Critical feature: search across platforms; if moving/finding content is easy, these regulations might just push bad content to the shadows
- Secondary questions:
  - More traditional competition question: *how easily can content migrate to a different platform?*
  - Platform design question: *what are the costs of anonymity online?*

This Paper: Private Action Against a Major Platform

- Visa and Mastercard responded to allegations against non-consensual content on MindGeek (MG) sites — the largest group controlling several adult websites – by not processing their credit card transactions
- MindGeek responded by removing all "nonverified" content (80% of videos) *overnight*
- We find that:
  - MindGeek lost 37-51% of visits as a result; users has strong preferences for "nonverified" content
  - The rest of the industry slowly gained visits; probably all shifted within 6 months, much of which went to very unregulated spaces (often malicious and IP-infringing, and possibly gateway to other illegal activity)
- Direct evidence on search:
  - Origin of visits $\implies$ from direct access to search engines and aggregators

Literature

Content Moderation in online markets

- Theory Liu et al. (2022, MktScience), Madio and Quinn (2024 JEMS), Rendo (2024), Bar-Isaac et al. (2024)
  - key novelty: empirical application of asymmetric content moderation (as in Rendo's theoretical work) when a major platform competes against a fringe
- Empirics Andres & Slivko (2021), Andres, Rossi & Tremblay (2022), Agarwal et al. (2023), Jimenez-Duran (2023), Jimenez-Duran et al. (2023), Beknazar-Yuzbashev et al. (2024, AEA P&P), Rizzi (2023)
  - key novelty: exogenous shock involving a major platform + tracking traffic and engagement across major and fringe platforms

Literature

Liability for online platforms

- Non-formalized theory Buiten, De Strelle and Peitz (2021), Lefouili & Madio (2022 EJLE)
- Theory: Jeon, Lefouili & Madio (2021), Hua & Spier (2023), Hua & Spier (Forth. AEJ:Micro), De Chiara et al. (2021), De Chiara et al. (2024), Zennyo (2023)
  - key novelty: show that some of the effects highlighted by Lefouili and Madio are likely to take place with asymmetric policies if
    1. search frictions can be overcome and
    2. users have strong preferences for *bad* content.

Market and Data

The Market

- Mostly made by free(mium) streaming websites
  - Firms owning domains where users, studios, and creators upload videos in exchange for ads and subscriptions; similar to YouTube
- Main firms:
  - Aylo (former MindGeek), domains including `pornhub.com` and `youporn.com`
  - WGCZ, Czech Republic group owner of `xnxx.com` and `xvideos.com`
  - `xhamster.com`
  - Fringe of many small websites
    - ad-based, mostly IP-infringing, no rewards for creators, oftentimes malicious

Asymmetric Content Moderation

- In December 2020, the New York Times published an article titled "The Children of Pornhub", with allegations that MindGeek hosted non-consensual content
- In response, MasterCard and Visa stopped processing transactions the same month
- MindGeek's primary site, Pornhub, banned unverified content, reducing overnight videos *from 13.7 million to less than 3 million*
  - Some videos were ultimately restored.
  - Removed content included videos with more than 29 million views.
  - No change in the content supply of other mainstream rivals
    - approx. 8.7 mln videos for `xnxx.com` and 9.7 mln for `xvideos.com`

Content supply

Date	pornhub.com	xnxx.com	xvideos.com
03/07/20	12,421,848	8,502,428	9,287,882
03/10/20	13,235,519	8,603,945	9,454,872
09/12/20	13,770,758	8,806,304	9,751,026
16/12/20	2,901,218	8,829,659	9,743,823
21/02/21	3,046,924	8,696,321	9,728,227
20/06/21	3,284,569	8,666,671	9,753,277

Table: Website content supply on different dates on the top three adult websites by number of visits worldwide. Data on content supply was retrieved from `https://web.archive.org` using the URLs of the three websites for six different dates in the period under consideration

Data

- `Similarweb` daily data on traffic by domain (not page) for 48 adult streaming domains in 8 countries (e.g., US, UK, IT, ES, FR, AT, BE, DE):
  - Visits
  - Avg Visit Duration
  - Bounce Rate
  - Page Views
  - Pages/Visit
  - Source of traffic (direct, search, referral, etc)
- Similar data available for a sample of non-adult websites (e.g.,approx. 300 control domains in news & media, social media, movie and entertainment sites)

Summary Statistics

Panel A: Adult vs non-adult websites

	Adult		Non-adult	
Mean	Before	After	Before	After
visits	469,206.20	448,326.50	1,462,284	1,438,776
pages views	4,426,187	4,325,923	11,551,597	11,687,095
bounce rate	0.38	0.37	0.55	0.55
pages per visit	9.73	9.90	3.60	3.61
duration	367.68	371.51	271.67	272.24
total time	3,115,438,210	2,971,379,537	8,352,905,080	8,558,787,227
N. domains	48	48	249	249

Panel B: Adult websites only

	MindGeek		Non-MindGeek	
Mean	Before	After	Before	After
visits	2,131,211	1,607,478	306,831	334,146.20
pages views	16,691,923	12,840,403	3,227,844	3,487,218
bounce rate	0.39	0.33	0.38	0.38
pages per visit	6.32	7.15	10.06	10.18
duration	423.68	452.75	362.21	363.51
total time	10,363,028,893	7,378,330,360	2,407,359,835	2,537,280,146
N. domain	4	4	44	44

Main Empirical Goal

- How does traffic respond across sites?
  - Consumers might go to other sites, and we take this to mean they are seeking the removed content. Do they actually respond?
  - If so, where do they go? Back to MG? To mainstream competitors or fringe?
- We mostly measure traffic with visits but consider other measures in the paper

Results

Event Study

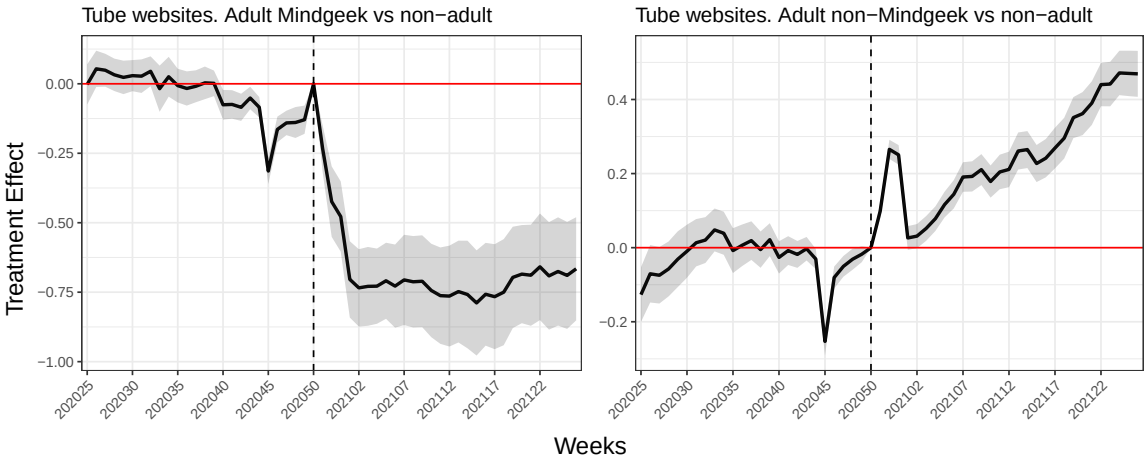

Figure: Event study of daily average weekly visits, log scale.

Main Specification

$$
\begin{aligned}
Y_{ict} = \alpha &+ \beta\, \mathit{After} + \gamma_1\, \mathit{AdultMindGeek}_{ict} + \gamma_2\, \mathit{AdultnonMindGeek}_{ict} \\
&+ \delta_1(\mathit{After} \times \mathit{AdultMindGeek}_{ict}) + \delta_2(\mathit{After} \times \mathit{AdultnonMindGeek}_{ict}) \\
&+ \delta_3(\mathit{After} \times \mathit{AdultMindGeek}_{ict} \times t) + \delta_4(\mathit{After} \times \mathit{AdultnonMindGeek}_{ict} \times t) \\
&+ \mathbf{\Psi}_{ict} + \varepsilon_{ict}
\end{aligned}
\tag{1}
$$

where

- Y_{ict}: Website activity metrics (e.g., visits, page views, bounce rate).
- *After*: 1 after the VISA/MasterCard ban (December 10, 2020).
- $\mathit{AdultMindGeek}_{ict}$ and $\mathit{AdultnonMindGeek}_{ict}$: MindGeek-owned vs. non-MindGeek adult websites.
- δ_1,δ_2,δ_3,δ_4: Treatment effects for MindGeek and non-MindGeek websites.
- Fixed effects: Country, website, and week fixed effects $\mathbf{\Psi}_{ict}$.

Main results

	Dependent variable:							
	log(Daily Visit per Domain-Country)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
After × Adult	0.134***	−0.021						
	(0.040)	(0.021)						
After × Adult MindGeek			−0.647***	−0.466***	−0.718***	−0.570***	−0.460***	−0.684***
			(0.065)	(0.049)	(0.052)	(0.046)	(0.053)	(0.058)
After × Adult non-MindGeek			0.277***	0.060***	−0.049	0.084***	0.088***	0.004
			(0.036)	(0.017)	(0.030)	(0.027)	(0.017)	(0.029)
After × Adult MindGeek × t						−0.005***	0.003	0.008***
						(0.002)	(0.002)	(0.002)
After × Adult non-MindGeek × t						0.014***	0.014***	0.013***
						(0.001)	(0.002)	(0.002)
Domain FE	✓	✓	✓	✓	✓	✓	✓	✓
Country FE	✓	✓	✓	✓	✓	✓	✓	✓
Domain-Country FE	✓	✓	✓	✓	✓	✓	✓	✓
Week FE	✓	✓	✓	✓	✓	✓	✓	✓
Domain Trends		✓		✓	✓		✓	✓
Holidays Included	✓	✓	✓	✓		✓	✓	
Observations	107,111	107,111	107,111	107,111	85,595	107,111	107,111	85,595

Interpretation

- MindGeek loses about half immediately
  - Change is sudden has substantial impact for MindGeek
  - It suggests strong user preferences for nonverified content
- Competitors gain 1.3-1.4% a week
  - Substitution is slow but persistent

Results robust to MC-NN, DID-PSM with website-domain FE, as well as analysis where all non-adult websites are in the control group

But which non-MindGeek sites gained the most? And how?

Decomposition of Traffic Diversion

	Dependent variable:		
	log(Avg Daily Visits per Domain-Country)		
	(1)	(2)	(3)
After × Adult MindGeek	-0.570^{***}	-0.460^{***}	-0.684^{***}
	(0.046)	(0.053)	(0.058)
After × Adult non-MindGeek	-0.072^{***}	0.064^{***}	-0.119^{***}
	(0.013)	(0.011)	(0.018)
After × Adult non-MindGeek × fringe	0.181^{***}	0.028	0.142^{***}
	(0.032)	(0.021)	(0.036)
After × Adult MindGeek × t	-0.005^{***}	0.003	0.008^{***}
	(0.002)	(0.002)	(0.002)
After × Adult non-MindGeek × t	0.005^{***}	0.016^{***}	0.017^{***}
	(0.001)	(0.001)	(0.001)
After × Adult non-MindGeek × fringe × t	0.010^{***}	-0.002	-0.004^{*}
	(0.001)	(0.002)	(0.002)
Domain FE	✓	✓	✓
Country FE	✓	✓	✓

In Levels

- MindGeek sites lost 25-28 million visits per day within a month
- Over six months from the shock, the fringe sites in our sample gained approximately 6-7 million daily visits, while mainstream sites experienced a rise in traffic amounting to 27-43 million visits

Search and Aggregators as the Source of Substitution

Search through Aggregators

- Search might lead to a different pattern of traffic, i.e. more from search engines?
  - Including referrals from aggregators: a different kind of search method; more targeted
- Traffic indeed comes increasingly from these sources
  - Search engines benefit from the change, and become an important channel of competition
  - But referrals increased too! Big alternative to search for content.

Referral Traffic Important

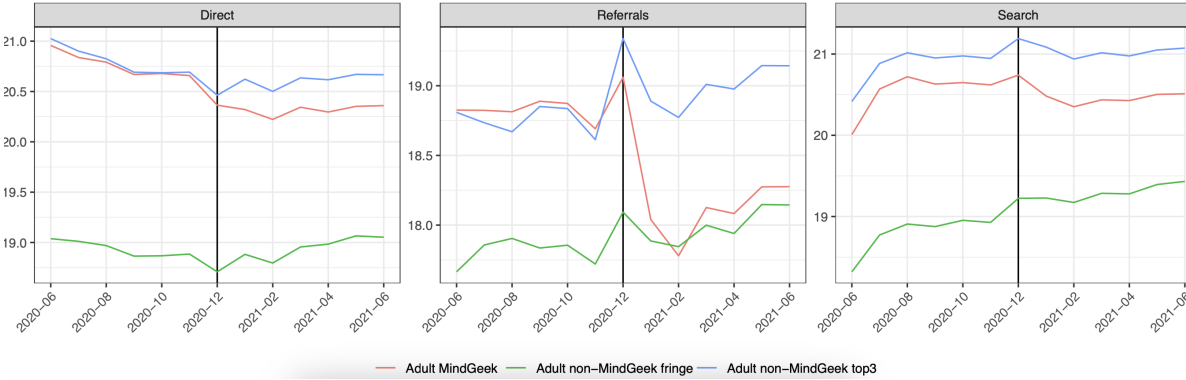

Source of Referrals: Aggregators

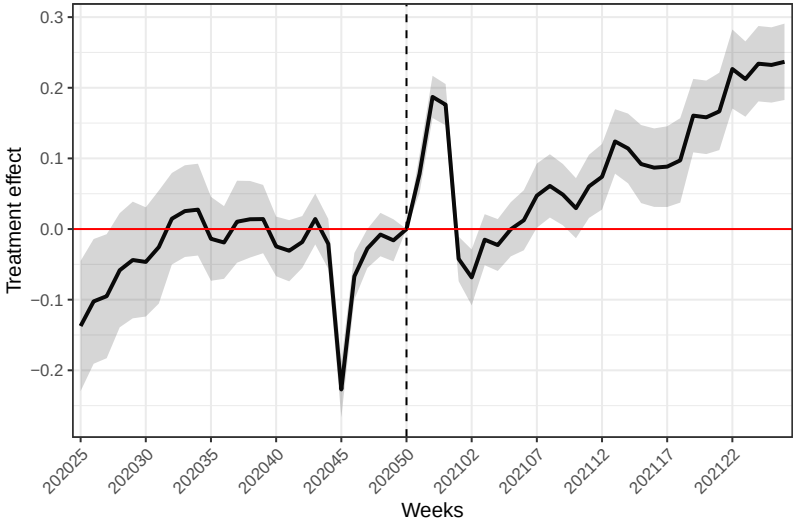

Conclusions

Conclusions

- Asymmetric content moderation leads to migration to other platforms
  - while MindGeek lost as much as half of their traffic, shady rivals gained as much or more over six months
  - Unintended consequence if
    - some platforms are unregulated,
    - search frictions are overcome (due to search engines and aggregators)
    - users show strong preferences for nonverified content
  - Asymmetric policies such as age restrictions on VLOPs are likely to generate similar effects
- Evidence that external force for content moderation was warranted: the drop in visits for MindGeek's sites and diversion towards rivals consistent with rivals having content with third party harm, but without incentives for content moderation

(Hopeful) Conclusions

- MindGeek did seem to change the nature of content by eliminating anonymity
  - May be more effective in other contexts
- While some substitution was bad from the standpoint of harm reduction, the substitution did change competition in the industry
  - Maybe useful for thinking about competition between platforms
- `xnxx.com` was declared VLOP in June 2024 as a result of the market it gained from this event
  - Unintended benefit? Asymmetric content regulation impacted competition and, probably, regulation more generally
- Timing also coincides with increased market share for creator-led sites

"Cam" Sites

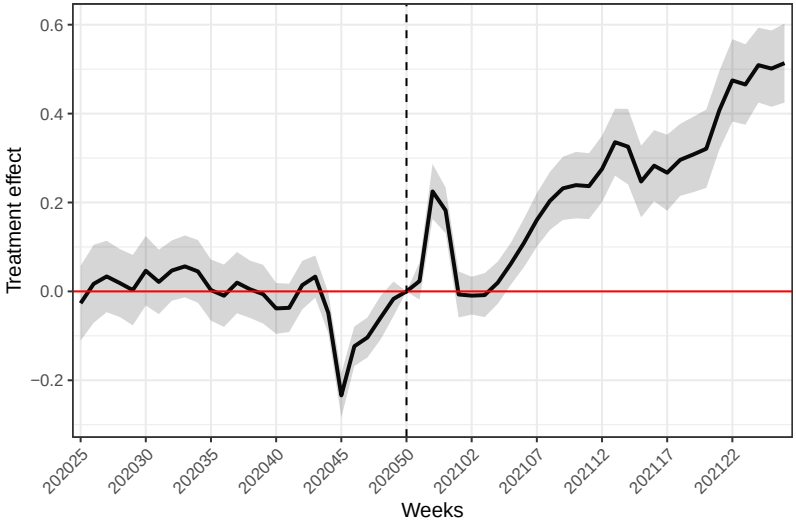
Treatment effect
0.6
0.4
0.2
0.0
−0.2
202025
202030
202035
202040
202045
202050
202102
202107
202112
202117
202122
Weeks