

MSI WEBINAR, April 28, 2026

When Connections Depress Contributions: The Hidden Cost of Social Filtering

Ziwei Cong, Georgetown University

Jia Liu, HKUST

Shijie Han, HKUST

Recommender Systems Define How Users Discover Content

Percentage of views driven by recommendations (0 - 100)



YouTube



Instagram



TikTok

Source: SocialMediaToday

Social Filtering: Curate Feeds Through Online Social Networks

- **Ubiquitous adoption:** Facebook, Quora, LinkedIn, Yelp, and many others
- **Platform logic:** Social connections foster belonging, drive engagement and satisfaction
- **What we know:** Social filtering enhances content **consumption**
- **The blind spot:** We know almost nothing about its effect on content **production** (e.g., posts, comments)



Does Social Filtering Help or **Hurt** Content Production?

- Two opposite forces make the answer ambiguous

ARGUMENT FOR ↑

Social filtering creates a sense of community. Seeing your content reach familiar audiences may **motivate** content production.

ARGUMENT AGAINST ↓

Routing content within homogeneous social clusters narrows the audience and limits perspective diversity, **discouraging** content production.

- We empirically examine this using a large natural experiment on a knowledge-sharing platform

Zhihu: Largest Q&A Platform in China

- Launched in January 2011 (similar to Reddit, Quora)
- ~75.7 million MAU and ~130 million answers



ASK QUESTIONS



PROVIDE ANSWERS



EVALUATE ANSWERS (UPVOTE)



EVALUATE ANSWERS (DOWNVOTE)



SOCIAL NETWORKING



SUBSCRIBE QUESTIONS AND TOPIC TAGS



Tags indicate the topics of a question



A question can receive multiple answers



Voting activity reflect agreement or disagreement

NBA NBA Player Stephen Curry Super Star

NBA NBA 球员 斯蒂芬·库里 (Stephen Curry) 巨星

如果把库里放到上个世纪 90 年代，还能否有巨星的表现？

If Curry were placed in the 1990s, could he still perform like a superstar?

关注问题

写回答

邀请回答

好问题 4

3 条评论

分享 ...

查看全部 138 个回答



张佳玮
2021 新知答主

765 人赞同了该回答

巨星在哪个时代都会是巨星，只看不同教练治下，能巨成什么样了。

像马克·杰克逊带库里三年，库里也就是全明星、联盟二阵。科尔一用库里，库里立刻带队进王朝。

也不奇怪。达拉斯版本的纳什就是全明星，太阳版本的纳什就是MVP。太阳版本的基德就是个杰出后卫，网版本的基德就是准MVP。

看配阵，看教练。

怎么用很重要。

展开阅读全文

Upvote Downvote

赞同 765



添加评论

分享

收藏

喜欢

...

Answer 1



演无三点
最难的是避免思维惯性和跳出舒适圈

12 人赞同了该回答

能。

上世纪90年代，非法防守还没被取消，只要球队不是傻缺，库里大概率能得到大量的一对一机会。而且就凭不能包夹弱侧无球人这一点，以库里的无球跑位穿插能力，对手得头疼死

至于handcheck，对喜欢原地持球三威胁的人有用，对库里真没多大用

发布于 2021-04-29 12:41

赞同 12



16 条评论

分享

收藏


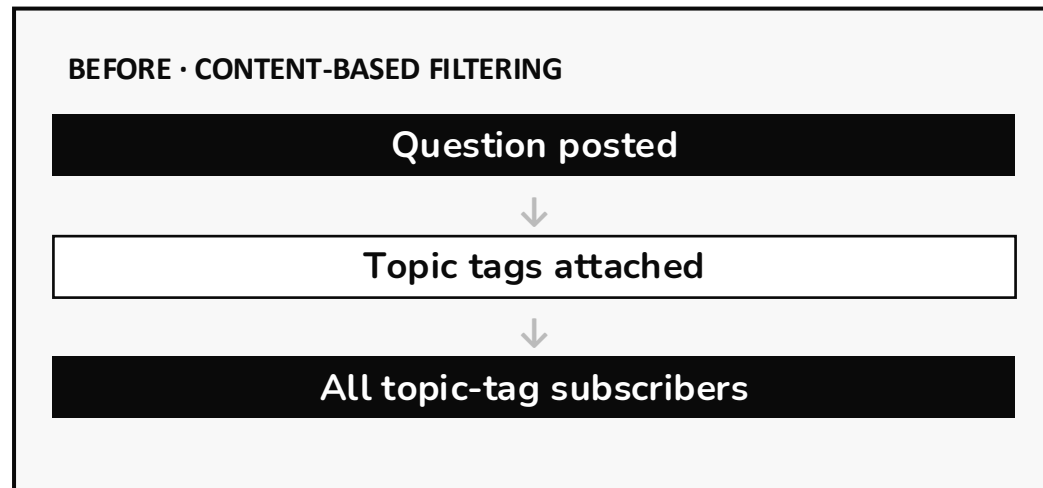
喜欢

...

Answer 2


Intervention on Zhihu User Feed

- Zhihu uses algorithm to recommend Q&As to users
- On August 16, 2012, the algorithm changed from content-based to social filtering



Intervention on Zhihu User Feed

- Simple, transparent designs
 - Isolate the impact of different types of algorithms
- Abrupt, platform-level change
 - Eliminate spillovers in social networks



Data


- Observation window: 90 days before and after the intervention
- 116,714 new questions, 284,419 new answers
- An average question receives 2.19 answers
 - 84% happened within the **first day** of question creation
- An average answer receives 2.5 votes
 - 95% happened within the **first day** of answer creation

Regression Discontinuity *in Time* (RDiT)

We compare answer creation in the days just **before and after** the invention, controlling for seasonality and smooth time trends

INTUITION

A sharp jump at $t = 0$, smooth on either side.



Regression Discontinuity *in Time* (RDiT)

We compare answer creation in the days just **before and after** the invention, controlling for seasonality and smooth time trends

 Y_1

Answer quantity

Num of answers per Q, answer odds, first answer response time

 Y_2

Answer quality

Avg and max quality across answers per Q


 Y_3

Viewer feedback (voting)


Num of upvotes and downvotes, upvote ratio

Effect on Content Quantity


- Social filtering caused a dramatic decrease in answer quantity



-45% ($p < 0.01$)



-33% ($p < 0.01$)



+555% ($p < 0.01$)

Effect on Content Quality

- LLM-based quality evaluation via prompt engineering (Liu et al 2023; Zheng et al 2023)

DIMENSION	WEIGHT	CRONBACH's α
Answer Length	0.10	0.748
Logic & Structure	0.20	0.882
Grammar & Diction	0.20	0.756
Relevance	0.20	0.864
Accuracy / Insight	0.30	0.867
Total Score	—	0.847

- **Advantage of LLM vs. human labeling:** scalable, can handle subjective languages and wide topical coverage
- **6 frontier models cross-validated:** GPT-4o, GPT-5, Claude 3.7 Sonnet, Gemini 2.5 Pro, DeepSeek V3.1, Qwen3-235B
- **Strong inter-model consistency:** Cronbach's $\alpha > 0.74$ on all dimensions

You are an evaluator for Zhihu, a Q&A platform in China. Your task is to assess an answer to a specific question along several independent dimensions. You will first read the question and the answer, then evaluate the answer along five dimensions using a 7-point scale. Importantly, evaluate each dimension independently.

The 5 Evaluation Dimensions

1. Answer Length

This dimension evaluates information richness and depth. Rather than assessing absolute word count, you should consider whether the response aligns with the expected depth given the question's complexity. A concise answer to a simple question can score just as highly as a detailed answer to a complex one. Penalize when the length is clearly mismatched with what the question demands — either unnecessarily bloated or superficially thin.

2. Logic & Structure

This dimension evaluates the internal coherence of the response, including the presence of clear argumentation, transitions, and hierarchical organization. Assess whether the answer exhibits logical flow and structured argumentation.

3. Grammar & Diction

This dimension evaluates the fluency and clarity of expression, including syntactic correctness and lexical choice. Penalize noticeable grammatical or wording errors that affect readability.

4. Relevance

This dimension measures the extent to which an answer addresses the posed question. Penalize content that is off-topic or only loosely related to the question.

5. Accuracy/Insight

For Fact-based Questions: Evaluate the correctness and precision of the information.

For Open-ended Questions: Evaluate the depth, originality, and practical value of the perspective.

The 7-Point Evaluation Scale

For each dimension, choose a single score from **1 to 7**. Apply the description to the specific dimension being evaluated.

1=Very Weak

2=Weak

3=Somewhat Weak

4=Moderate

5=Good


6=Very Good

7=Excellent

Effect on Content Quality


- Social filtering decreased answer quality, especially for the “best” answers

Average Quality (Across Answers) per Question



-2% ($p < 0.01$)


Max Quality (Across Answers) per Question




-7% ($p < 0.01$)

Effect on Viewer Feedback (Voting)


- Social filtering significantly improved consensus in voting



+5% ($p < 0.01$)




NS ($p > 0.1$)



-19% ($p < 0.01$)


The Core Paradox

- **Worse Content. Happier Voters. At the Same Time.**



Relying solely on voting metrics, a platform would erroneously conclude social filtering has positive effects.


Mechanisms




Mechanism: Reduced Distribution Scale

Distribution scale determines the probability of willing, qualified answerers encounter the question

Distribution scale under social filtering:
Asker's follower size




Distribution scale under content filtering:
subscribers of related tags




Mechanism: Reduced Distribution Scale

- 353,140 users and 63,666 tags before intervention




- 116,741 questions posted in the study period



Does Distribution Scale Influence Answer Contribution?

- If yes, the decrease in answer quantity and quality should be stronger for askers with fewer followers
- We split askers by the median pre-intervention follower size




When questions can't reach enough people, both quantity and quality suffer


Mechanism: Amplified Homophily in Distribution Composition (Shi et al. 2019; McPherson et al., 2001)

Route content within homophilic clusters inflates voting consensus, but depresses discussion diversity


Homophily in social networks



Homophily among subscribers



Measure Homophily




1

User = average of her questions

Represent a user by the questions she answered or subscribed to before intervention

Measure Homophily



1

User = average of her questions

Represent a user by the questions she answered or subscribed to before intervention

2


Question = average of its tags

3


How to get tag embedding?

Measure Homophily

- Zhihu Topic Tree + Graph Neural Network + Pre-trained Zhihu tag embeddings (Li et al 2018)




Zhihu Topic Tree



GNN Update Process

Measure Homophily

- Embedding users and content on the same space
- Quantity interest homophily using Cosine Similarity




Measure Homophily

FOR EACH FOCAL **USER** u

SocialHomo_u


Average cosine similarity across her followers (before the intervention)



FOR EACH FOCAL **TAG** j


SubscriberHomo_j

Average cosine similarity across its subscribers (before the intervention)



Measure Homophily

- On average, $SocialHomo_u$ is **higher** than $SubscriberHomo_j$, across all users and tags before the intervention
- Social ties on Zhihu link more similar users than topic subscriptions do



Does Homophily Influence Answer Contribution?

- If yes, the reduction in quantity and quality should be stronger for askers with higher social network homophily
- We split askers by median pre-intervention social network homophily

ANSWER QUANTITY

High-Homophily

-31.6%

Low-Homophily

-23.2%

~36%
larger drop

BEST ANSWER QUALITY

High-Homophily

-6.1%

Low-Homophily


-3.1%

~97%
larger drop


When a question circulates within homophilic social clusters, both quantity and quality suffer
Homophily matters even controlling for distribution scale

Mechanism: Amplified Homophily in Distribution Composition


Does exposure homophily translate into homophily among engaged users?




(1) Asker-Answerer Similarity



(2) Asker-Voter Similarity




(3) Answerer-Answerer Similarity




(4) Answerer-Voter Similarity

Mechanism: Amplified Homophily in Distribution Composition

Find evidence that exposure homophily translates into homophily in actual engagement



What do We Find?



Rethinking Algorithm Design

On content platforms, the recommender algorithm is **distribution** — the "place" of the 4Ps

①

Platform Designers

Balance homophily and diversity.
Many algorithms (not just social filtering) are homophily-driven.

②

Algorithm Evaluators

Multi-metric assessment is essential.
An algorithm's impact on platform ecosystem is multi-dimensional.

③

Policymakers

Filter bubbles aren't just a consumption problem; they also depress knowledge exchange and production.

Thank you!

Jia Liu: jialiu@ust.hk

Ziwei Cong: zc260@georgetown.edu