

Better Feeds: Algorithms That Put People First

A How-To Guide for Platforms and US Policymakers

Every day, billions of people scroll through social media feeds, search results, and streaming recommendations that shape what they see, read, and watch. Algorithmic systems determine what to show each user, wielding enormous influence over our online experiences and, increasingly, our lives offline. While these algorithms have fueled some of the world’s most successful businesses, they have also sparked intense debate about amplifying unwanted content, fueling risks to kids, and sowing societal division. Chronological feeds and blanket bans on personalization are common go-to solutions, but they have important limitations and can reward spammer behavior. Better designs exist that put users’ interests front and center.

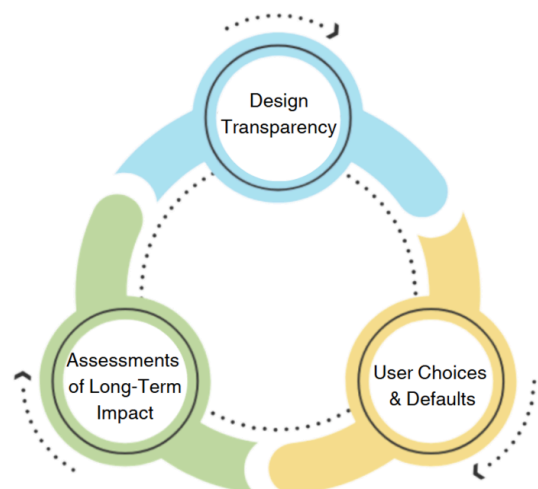
The Problem: Maximizing Short-Term Engagement

Some platforms optimize their recommender systems to maximize predicted “engagement” – the chances that users will click, like, share, or stream. This design aligns well with the business interests of tech platforms monetized through advertising. Product teams are rewarded for showing short-term gains in platform usage, and financial markets and investors reward companies that can deliver large audiences to advertisers. This design approach has been linked to a range of individual and societal harms, including the spread of low-quality or harmful information, reduced user satisfaction, problematic overuse, and increased polarization.

The Solution:

Focus on Long-Term Value to the User

Most everyone has had the experience of impulsively clicking, liking, or viewing content that they do not necessarily enjoy, or may even regret. Available evidence underscores the need for a shift towards designs that focus on **what people prefer or aspire to over the long term**. Finding out what users value over the long term and optimizing for it requires platforms to rely on data other than engagement, like user surveys and preferences.



Getting to Better Feeds

Platforms and policymakers can help to address the harms associated with recommender systems while preserving their potential to enhance user experiences and societal value. The Knight-Georgetown Institute [Expert Working Group on Recommender Systems](#) brought leading researchers working on recommender system design together with industry and policy experts to produce a [how-to guide](#) for achieving this through US policy and product design. Three components are key: **design transparency, user choices and defaults, and assessments of long-term impact.**

Design Transparency

Detailed disclosures about the design of recommender systems would allow outside experts, regulators, and the public to understand the tradeoffs being made in the design of these systems, and they would motivate platforms to optimize their products for long-term user value and satisfaction. The Expert Working Group proposes that platforms must publicly disclose:

- information about the specific input data and weights used in the design of their recommender systems;
- the metrics they use to measure long-term user value; and
- the metrics they use to evaluate product teams responsible for recommender system design.

User Choices and Defaults

User choices would allow individuals to tailor their platform experiences and switch to recommender systems that provide them with long-term value. The Expert Working Group proposes that:

- Platforms must offer users an easily accessible choice of different recommender systems. At least one of these choices must be optimized to support long-term value to users.
- Platforms must provide easily accessible ways for users to set their preferences about types of items to be recommended and to be blocked. Platforms must honor those preferences.
- By default, platforms must set minors' recommender systems to be optimized to support long-term value to these users. If platforms have insufficient information about long-term value to minors, they must default to non-personalized recommender systems.

Assessments of Long-Term Impact

Platforms can only deliver long-term value to users if they continuously test the impact of algorithmic changes over time. Many platforms maintain a “holdout” group – a group of users that are exempt from having design changes applied to their accounts, and who function as a control group for comparison with the rest of the user base. Auditing the results of these tests and publishing aggregate test results are keys to accountability. The Expert Working Group proposes:

- Platforms must run long-term (12-month or longer) holdout experiments on a continuous basis.
- Platforms must report the aggregate, anonymized results of the holdout experiments publicly.
- Holdout experiments must be subject to an audit by an independent third party.