When Should Platforms Break Echo Chambers?

Abstract

Recent calls for regulation of social media platforms argue that they serve as conduits of extremism. Several platforms have responded by banning communities that peddle extreme or misleading ideas. These communities are usually echo chambers, consisting of users with similar ideologies repeating the same information to each other. This amplifies harmful beliefs and makes them more likely to metamorphose into dangerous offline actions. We develop a novel community formation model to show that this traditional view of echo chambers is incomplete. We then use this model to study several interventions that are implemented in practice and investigate their efficacy in preventing costly offline actions from originating on the platform.