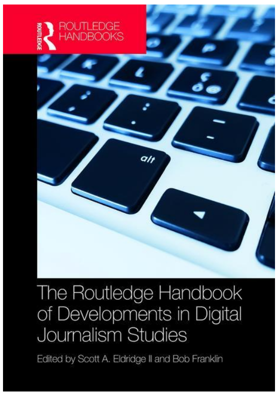

The Routledge Handbook of Developments in Digital Journalism Studies

Scott A. Eldridge, Bob Franklin

Testing the Myth of Enclaves

Publication details
https://www.routledgehandbooks.com/doi/10.4324/9781315270449-11
Jacob Ørmen
Published online on: 30 Aug 2018

10

TESTING THE MYTH OF ENCLAVES

A discussion of research designs for assessing algorithmic curation

Jacob Ørmen

More and more of the information we receive in the world is curated by algorithms. Every time people use digital intermediaries, such as social network sites or search engines, computer programs guide us to the information that serves our needs (supposedly). As these services play an increasingly important role as access points to news and information (Newman et al., 2016), users rely more and more on algorithms to guide them to what they encounter on the internet. News organizations also depend heavily on digital intermediaries to serve news to their readers. In this relationship, search engines and social network sites act as algorithmic gatekeepers (Bozdag, 2013) in the intersection between content producers and users. It remains a crucial task for journalism research and practitioners alike to understand how algorithmic curation affects the type of information users are exposed to and interact with.

To study this, this chapter takes a widespread myth about algorithmic curation as a starting point. The 'myth of enclaves' is the idea that digital intermediaries like search engines and social network sites drive people further into 'filter bubbles' (Pariser, 2012) or 'echo chambers', where they are only exposed to content and ideas in line with their own ideological beliefs. It is a myth not because it is (necessarily) fictional, but because it remains contested among scholars. Some research has found algorithmic curation to lead people into enclaves through personalization; that is the mechanisms whereby algorithms adapt the presentation of content to the individual user (Bakshy et al., 2015; Jacobson et al., 2016). In contrast, others have found algorithms to concentrate, rather than fragment, attention on the already well-established actors online (Dutton et al., 2017; Hindman, 2009). Moreover, a third group of scholars argue that digital intermediaries expose people to a greater variety of information than they seek out themselves (Flaxman et al., 2016; Vaccari et al., 2016). The 'myth of enclaves' lives on in part because the academic public is so mixed about what to think about it.

The aim of this chapter is neither to confirm nor debunk the 'myth of enclaves' but instead to discuss the various methods applied to investigate the role of algorithmic curation in fragmenting, concentrating, or polarizing the public's exposure to content on the internet. By critically assessing the methods and techniques used to test the myth of enclaves, we can achieve a firmer ground on which to evaluate the credibility of the claims in favor or against the myth. To do this, this chapter begins by outlining what 'algorithmic curation' covers. It then proceeds by analyzing the methodological strengths and weaknesses of four prototypical research designs for studying algorithmic curation. In the discussion, the chapter situates the research designs in context with a fifth approach that moves beyond mere exposure.

Algorithmic curation

Algorithmic curation describes a process whereby a computer program selects pieces of content to present to users based on a step of logical operations encoded by human programmers (Braun and Gillespie, 2011). A wide range of services online make use of algorithmic curation to provide users with the options that best match their presumed interests, be it TV and film (e.g., over-the-top distributors like Netflix and HBO), music (e.g., streaming services like Spotify or Apple Music), goods (e.g., marketplaces like Amazon and Taobao), information (e.g., search engines like Google or Baidu), or social activity (e.g., social network sites likes Facebook and WeChat).

There are at least three ways algorithmic curation matters for how users meet content on the internet. First, algorithms operate by *selecting* content to display to the users. The programs select content by searching through a library (e.g. all indexable documents on the web in the case of search engines or all posts by a person's network in a given time frame in the case of social network sites) from where a few pieces of content are chosen to be presented to the users (e.g., the search results list or the newsfeed). For the purpose of methodological discussions, we could call the library for a population and the content selected to be displayed to individual users for a sample. The key operation of curation algorithms is identifying the specific sample in the population that best matches the needs of the user in the particular moment of usage. The greater the size of the population and the smaller the sample needed, the larger is the role played by algorithmic curation. The growing number of websites, videos, songs, movies, goods, and social material online, combined with the limited time available for users, makes curation (in algorithmic or other forms) a necessary premise for finding content on the internet at all.

The second way algorithmic curation plays out is by *ranking* the selected content. To do this, algorithms rely on relevance criteria encoded by programmers. In its early days, Google developed a search algorithm for assessing the relevance of individual web pages that takes both the hyperlink network between websites (structural level) and the content on web pages (semantic level) into account (Brin and Page, 1998). In this way, each website would get a popularity score – its structural importance in the link network, called PageRank – and a relevance score based on the search string. These criteria would be universal for all searchers using the search engine. Similarly, content providers often rank content on their websites based on general popularity (e.g., measured by clicks or time spent with content). Sometimes this information is fed back to users as "most popular" lists on news websites or trending stories on social network sites (Webster, 2014). Ranking plays an important role in guiding users to particularly relevant content in the sample selected.

The third way curation algorithms operate is by *personalizing* the sample to match each individual user. On top of the general relevance criteria discussed earlier, digital intermediaries integrate context-specific signals in the curation process. Over the years, the complexity of signals that feed the curation algorithms in Google Search has risen steeply and now includes factors such as geographical location, language-setting, and search history (Granka, 2010). Facebook originally used an algorithm (EdgeRank) to assess the relevancy of each post based on the affinity of user connections, previous interactions with the content, and how recently the content was posted. This has since been supplanted by a machine learning approach that takes "more than 100,000" factors into considering when ranking content in the Facebook newsfeed (McGee, 2013). Currently, personalization algorithms like those governing search results and newsfeeds are so complex in structure and execution that few – if any – know exactly how they function.

For research purposes, personalization has made it difficult to study curation empirically, since algorithms produce different samples of content for each user and each use situation. It makes a difference who interacts with the algorithm, when, and where.

This raises particular issues for how to research algorithms. For recommendation mechanisms such as those suggesting other goods to purchase (on Amazon), movies to watch (on Netflix), or music to listen to (on Spotify), personalization is driven largely by previous interactions *within* the service. However, social and search algorithms increasingly draw on information *outside* of the service. Facebook, for instance, tracks user activity across the internet and uses that information to sample content in the newsfeed (Debatin et al., 2009). Likewise, Google displays ads (through the AdSense network) based on search and click patterns around the web, both on and off Google services. This development has wide-ranging methodological consequences for studying algorithmic curation. It entails that there are no *universal* scores determining the output of the algorithm but rather a range of *particular* scores that depends in large part on signals picked up from the individual user. After the introduction of personalization factors in algorithmic curation, it is an inductive fallacy to generalize the specific results for one user (e.g., the search results list to a given query) to all other users across time and space.

This chapter focuses on the challenge personalization presents to empirical research on algorithmic curation. Therefore, by focusing on intermediaries like social network sites and search engines that rely heavily on personalization, it leaves out curation that only selects and rank but not personalizes, such as generic databases relying on keyword searches or popularity lists (e.g., most-read news stories on a news website) that are updated algorithmically but not adapted to the individual user. Understanding ways to study personalization makes a difference for how scholars can make sense of information curation in particular and the myth of enclaves in general.

Directed and undirected curation

Apart from the various ways algorithms work, there is also a fundamental difference between types of curation. When a user browses the feed on a social network site, the algorithm displays results without concrete input from the user. Curation is *undirected*. For this type of curation, the algorithms produce a ranked sample of content based solely on what it estimates to be in our greatest interest at that particular moment. The algorithm pushes information to the user. This is reminiscent of the way human editors have selected and ranked information (to be displayed in their publications) for centuries. In contrast, when a searcher enters keywords into a search engine, the algorithms deliver results in response. Here, curation is *directed*. The input provided by the searcher (a query or a search string) gives the algorithm further information to determine the relevance of each piece of content in the population. A user pulls information through the algorithm. This is similar to the work carried out by human librarians for millennia. In reality, many services integrate both types of curation, as is the case when users search on a social network site, or when search engines display information to individual users solely based on previous behavior. However, the distinction between undirected and directed curation makes an important difference for how we can analyze personalization mechanisms.

Directed curation is by far the easiest to track, since data are available from the services. To query Google, the searcher does not have to create an account and log on (the same goes for services like Twitter and YouTube, but not for Facebook). It is thus possible to access and retrieve data in simple ways, as each query forms a unique URL. For instance, a search on google.com with the string "barack obama" appears like this www.google.com/search?q=barack+obama (cleared for other type of metadata, such as browser information that Google integrates as well). In short, the search term is directly retrievable from the search URL. This makes it possible to study which terms people search for and the results they click on solely by knowing the URLs they have visited and in which sequence. Another benefit is that directed curation algorithms operate on one grand population (such as all indexable documents on the web). In theory, every searcher can receive exactly the same results to a query. Thereby, researchers can compare directly the sample of results (e.g., top 10 results) presented to each user as a response to the same search query. Accordingly, it is possible to study both the extend as well as the type of personalization occurring.

In contrast, undirected curation presents a number of unique challenges to research. First, to get access to undirected curation on social network sites a researcher has to create a profile or log in through existing user accounts. It used to be possible to access the newsfeed through the API (Application Programming Interface), but that service has been discontinued as of October 2015.[1] Currently, scraping newsfeeds directly through the user interface is prohibited by Facebook's Terms of Service.[2] Instead, it is necessary to access the newsfeed directly through each individual user account. Second, the algorithms rely solely on the content shared by connections in each user's (as well as sponsored) content to compose the sample. It is difficult to compare newsfeeds between users directly, since the samples of content are not drawn from the same population. Instead, it is possible to study the relationship between sample and population *for each* user – that is, which content the algorithms deem more or less relevant to the specific user.

Accordingly, there are substantial differences in how researchers can access and study algorithmic curation. In the next section, I go through widespread strategies for studying both indirect and direct curation. To make the analysis more concrete, the search engine Google and the social network site Facebook serve as cases of directed and undirected curation, respectively.

Four prototypical designs for assessing algorithmic curation

Now I look into ways research has assessed the effect of algorithmic curation on the type of information users are exposed to and interact with. The studies I discuss here all rely on behavioral data collected with/on users as they interact with algorithms. This is not to discount the importance of theoretical or ethical discussions of algorithms – which remain a pertinent and prolific activity in communication research (for discussion see Nahon, 2016; Dörr, this volume, Chapter 23) – but merely to zoom in on the part of the literature that focus on methodological issues of researching curation mechanisms.

To guide the discussion, I have constructed a matrix that encapsulates what can be considered from four prototypical research designs for studying algorithmic curation (see Table 10.1). In the columns, the model distinguishes between whether the research designs operate with artificial or real users. Artificial user accounts are those constructed for the purpose of research. Typically, this is done by setting up specific profiles on a social network site or a search engine. It also includes research profiles that are modeled on real user accounts but inscribed with artificial agency for the purpose of research. In the rows, the matrix distinguishes between artificial and real settings. The idea is to make a delineation between studies that operate in a preexisting and live environment (real setting) and those that construct a specific environment for the purpose of research (artificial setting). When the rows and columns are crossed, four prototypical designs emerge: *simulation* (artificial setting and artificial user), *experimentation* (artificial setting and real

Table 10.1 Matrix of prototypical algorithmic curation designs

	Artificial user	*Real user*
Artificial setting	Simulation	Experimentation
Real setting	Manipulation	Observation

user), *manipulation* (real setting and artificial user), and *observation* (real setting and real user). These designs configure the extent to which the researcher has access and control over the user experience as well as the research environment.

The four designs are prototypical in that each design is a rough characterization of typical features shared by studies in general but not applicable to any study in particular. Many of the projects working on algorithmic curation draw on elements from different designs, and for good reason, as will be returned to in the discussion. Nonetheless, the designs are helpful as sketches to think about the different methods and techniques researchers can apply to conduct empirical studies.

Simulation

With simulation, researchers have the benefit of controlling both the setting and user input. This approach is strictly not empirical but computational. The researcher does not study real users or real settings but instead constructs a research environment where the interaction between artificial user profiles and platform characteristics are simulated. This approach draws on game-theoretical assumptions to model human behavior in socio-technical systems, termed agent-based simulation.[3] In recent years, it has been a popular approach to test and design information systems (Ryczko et al., 2017). The goal is often to improve ways of displaying information, for instance through personalization steps, based on assumptions about which outcomes serve the individual user better or worse (Micarelli et al., 2007). Thus, a core element of simulation studies is to set up parameters, such as assumptions about agents' rationality, as well as methods for evaluation outcomes. Often the models draw on data from real-world users, but they are treated as computational agents inscribed with certain preferences and behavioral tendencies.

Simulation studies have dealt with a range of issues in algorithmic curation. The spread of information such as news stories through social network sites is estimated with simulations of millions of user accounts (Del Vicario et al., 2016; Goel et al., 2016). Here, the goal is to predict information cascades (or virality) of various types of content. Other studies have sought to find optimal levels of personalization for individual users by simulating interactions among artificial profiles (Guo et al., 2013; Chung et al., 2016). Together, simulation studies are helpful in showing general tendencies across many or all users of a search engine or social network site. This design provides an important insight into how signals affect algorithms, such as how likely it is for content to show up in individual newsfeeds given previous patterns of information diffusion.

The obvious downside of such a design is the loss of ecological validity. The simulated environment can only approximate a real setting but never replicate the complexity of a Facebook or Google ecology. Predicting future patterns of user behavior from logs of past activities might function well for very general phenomena such as the spread of information cascade – although scholars do debate whether simulation models can even do this reliably (Cheng et al., 2014) – but it is not well-suited to show how algorithmic curation affects the individual user. A related issue is the behavioral assumptions programmed into the simulations. Without an input from real-world users, researchers have to rely on general heuristics such as rational choice theory to simulate user agency. These assumptions will naturally be rather crude and in some cases misleading. Lastly, the need to operate with a ground-truth to evaluate the results of simulations is problematic in itself. Scholars should always be wary of attributing fixed preferences to individuals from the outset (such as person x always preferring content y to content z). Nonetheless, social research is well advised to draw inspiration from simulation studies to understand personalization mechanisms at play.

Manipulation

The most widespread approach to study directed curation has been to set up artificial research profiles and let them loose in real-world settings (notably on Google Search). This approach makes it possible to control and manipulate some personalization signals – such as language settings, geo-location of IP-address, anonymous browsing – while holding others constant. In this way, the researcher can seek to mitigate (or work with) the effects of the algorithm on the outcome presented to users (for an elaborated discussion see Ørmen, 2015). Whereas some have advocated a strategy where the researcher keeps a clean research profile with as few personalization signals as possible (Rogers, 2013), others have actively manipulated one or more signs such as language settings or geographical location to study the causal effect on personalization (Ørmen, 2015; Kliman-Silver et al., 2015). A third line of research has sought to construct "extreme cases" (Kuzel, 1999) to compare profiles that are as different as possible (Dutton et al., 2017; Feuz et al., 2011). In the latter case, research teams have trained profiles to display ideological preferences by feeding each profile with a predefined set of search strings belonging to either the left or right side of the political spectrum. Subsequently, a set of common keywords have been queried through each profile, and the differences in results have then been attributed largely to personalization processes. By constructing profiles bottom-up, from a blank slate to an ideological extremist, these studies attempt to set up a quasi-experimental design in semi-controlled conditions. In this way, researchers can attempt to 'game' the algorithm to maximize (or minimize) personalization of results.

However, operating with artificial research profiles also creates issues of ecological validity. It is difficult to mimic real human behavior on the internet solely by setting up a fake Google profile. First, algorithms on search engines and social network sites take a plethora of signals into account, including general browsing behavior (including beyond the boundaries of the website or app). For instance, if one is an avid user of popular services online, these services are likely to show up in the Google search results as well, in particular if they are Google related, such as YouTube (Edelman and Lai, 2016). Second, real users interact with content in ways that can be hard to replicate with artificial accounts (such as clicking patterns). Third, digital intermediaries experiment with and randomize the samples produced, which introduce substantial – and unknown – noise into analysis. There are simply too many unknown factors at play for artificial research profiles to truly mimic real-world user behavior.

Experimentation

The opposite approach has been to study real-world users in artificial settings. A number of studies have designed applications mimicking well-known social network sites to experiment with real-world users. This approach follows the controlled experiment approach known from laboratory research, where participants are recruited and randomly assigned to either a treatment or a control group. One study extracted newsfeed data through the now-unavailable Facebook API to test various personalization designs with participants in a lab setting (Eslami et al., 2016). In a similar manner, studies have shown the influence of social endorsements on social network sites by randomly assigning types of social interactions (kinship, tie-strength) to news stories on an experimental platform designed for the purpose of research (Kulkarni and Chi, 2013; Messing and Westwood, 2014).

The greatest advantage of an experimental approach is the ability to assess causality. By controlling which algorithmic signals influence which users, it is possible to test hypotheses directly. The downside is that external validity might be harmed. If the lab settings (artificial platforms designed to look like well-known services) fail to reproduce the complexity of behavior online, then the conclusions drawn do not hold. For instance, the importance of social cues for interactions on social network sites might not be as clear or consistent when users are exposed to content in the messy real-world settings where there are many more variables at play. A further issue is that the experimental setting itself might interfere with the validity of the results. A lab setting, where a person is asked to focus on one or more specific tasks (and possibly rewarded to do so), does not replicate a use situation in daily life, where people interact with content in the midst of all the other tasks, thoughts, and duties that preoccupy their minds in daily life. Accordingly, the controlled experiment has merits but suffers greatly from the artificial situation that lab settings create, both offline and online.

Observation

The most desirable approach is to study real-world users in real time as they are exposed to and interact with content in real settings. This is often done by tracking users with cookies and through user accounts as they move across the internet. It is an observational approach in the sense that it is reminiscent to a "fly on the wall" approach in ethnography, where researchers seek to blend in with the environment and observe unnoticed. The major advantage of tracking users online is that researchers can document behavioral patterns of individuals in an unobtrusive manner. In this way, researchers can collect detailed levels of data on a large number of users. Unfortunately, it is typically only industry researchers that have access to this kind of data. Most of their research is for optimizing business and thus considered proprietary information not made available to researchers.

Therefore, research on algorithmic curation using observational data primarily comes from the digital intermediaries themselves. In the case of directed curation, researchers at Google have assessed the impact of various personalization measures on the click-through patterns of Google News users in real time (Liu et al., 2010; Das et al., 2007). For nondirected curation, Facebook's research team has carried out studies on personalization using observation data on millions of users (Bakshy et al., 2015). Working with the raw data collected by Facebook, the researchers are able to compare the total population of content available to be displayed to each user (all the posts created or shared by their connections) and the sample of content presented to each user (the personalized newsfeed) as well as the select few pieces of content that users interact with (by clicking, liking, sharing, etc.). These studies provide crucial insight into algorithmic curation, since they can cover the whole user population in real time as well as past time (historical data). This allows for comparisons on an unprecedented scale and detail.

If researchers gain access to archived or live observational data, there is the possibility to move beyond mere observation and conduct natural and field experiments. For a natural experiment, researchers compare two or more preexisting groups that vary on key characteristics and then investigate differences between groups in time. Thus, the researcher cannot manipulate the variables as in a controlled experiment but will have to attribute causality based on the differences observed. For instance, one natural experiment showed how users tend to click on Google products more often when the algorithm ranks such products higher in search results lists (Edelman and Lai, 2016). The major downside of a natural experiment is, naturally, the lack of possible interference by the researcher. Causality can be assumed from observations but not tested in practice. In contrast, field experiments offer the ability to work with real-world users in their real settings. Field experiments in online settings can be done unobtrusively and are thus preferred to controlled experiments, discussed earlier. A Facebook study showed how a get-out-to-vote button significantly increased the propensity of users to actual get out and vote on election day (Jones et al., 2017). The downside here is that field experiments require full access to live user data, and it is challenging to do while abiding by academic ethics codes. In short, rich observational data allow for studies that attribute causality by comparison (natural experiments) or interference (field experiments) in live environments.

There are challenges to the observation design as well. First of all, few researchers have access to the scale and depth of data required to assess personalization mechanisms. Forming partnerships with or buying data directly from industry actors is typically the only way to gain access to large-scale observational data. For field experiments – where one needs access not only to the user base but also to back-end programming – this is even more difficult. The most promising way to produce observational data from the outside has been to recruit human participants (e.g., through Mechanical Turk) and then instruct them to query particular keywords on Google (Dillahunt et al., 2015; Hannak et al., 2013) or record their Facebook newsfeed (Bhargava et al., 2015). However, it is not ethically unproblematic to hire people to participate in research in this way, in particular if the project relies on precarious labor platforms like Mechanical Turk. Accordingly, in academic practice large-scale observational studies will often be infeasible or unethical – or both.

Moving forward: an alternative design

So far, the analysis of methods has dealt solely with the type of content algorithms tend to expose users to. Such a focus overlooks the part played by users in decoding curation processes. As media research has shown in the past 50 years, users are not mere passive consumers but reflexive and critical beings. We need to take user agency seriously in assessing how algorithmic curation affects their practices. Some studies have attempted to do so through interviews or surveys with users (Bucher, 2017; Rader and Gray, 2015; Powers, 2017). Getting user discourses serves an important purpose in outlining general attitudes of and affections toward algorithmic curation. However, it tells us little about how users relate to the specific acts of curation, for instance the particular constellation of search results or stories in the newsfeed. It remains an open question to which extend users navigate, peruse, and contest the results displayed by the algorithm. To do this, we need to integrate the focus on exposure as outlined in the four designs with user interpretations.

Arguably, the best way to test the 'myth of enclaves' would be to document not what people get exposed to but what they think about what they are exposed to. One way to do this would be to recruit a small-*n* sample of participants, for instance following a maximum variation sampling strategy (Kuzel, 1999), to explore differences across individuals. Researcher could install logging software on their primary devices and interview them about their use experiences. Currently, several applications exist that can log usage across devices (e.g. laptop, smartphone, and tablets) and some that can collect detailed descriptions of web behavior (e.g. www.webhistorian.org). Following the ideas of experience sampling methodology (Larson and Csikszentmihalyi, 1983), users could be prompted with small questions right after their visits to specific services, such as Facebook and Google, to get immediate reactions to and interpretations of the curation process. Likewise, the logged behavioral data (or screenshots of a search results list or newsfeed) could be used as prompts in follow-up interviews to trigger reflections on routinized and habitual behavior (Ørmen and Thorhauge, 2015). In this way, researchers and participants can discuss algorithmic curation on a detailed level with data on the participant's own behavior as the basis. Such an approach, an *interaction* design, would also enable researchers to situate ordinary people's practices with algorithmic curation in their broader practices of media use. We know from other studies that most people keep a varied media diary, moving far beyond social media (for an overview, see Helles et al., 2015). Including this contextual information about individuals in a study of algorithmic curation would be a strong test of the 'myth of enclaves' and a supplement to existing research designs.

Conclusion

In the introduction, we saw how scholars have come up with widely different empirical results when they have investigated the 'myth of enclaves'. Now that the chapter has discussed and evaluated methods for testing this myth, part of an explanation for these inconsistencies emerges. The fundamental differences in how scholars study algorithmic curation and personalization mechanisms naturally affect the type of conclusions that can be drawn. The empirical material used in simulation, manipulation, experimentation, and observation studies is so different that it renders direct comparisons between design types untenable. The findings from one study do not translate well into results from a study using a fundamentally different design. Unfortunately, the research design that tends to produce the most convincing empirical results (observational studies of real users in real-world settings) requires special access to proprietary data owned by digital intermediaries. One way to remedy this is for academic researchers to seek partnerships with industry researchers to get access to real-world observational data and to conduct field experiments with real-world users in a live environment. More of this research would surely strengthen the empirical basis in the debate on 'enclavization'. To move the discussion forward, I suggested an alternative route that combines observational data (logged by research applications) with people's own discourses about algorithmic curation. Such a design would shed further light on poorly covered aspects of polarization and fragmentation – namely, the experiences of ordinary people in the broader context of their media-saturated lives. The need to move beyond single services or platforms in the study of personalization effects is the key takeaway from this chapter. Given all these challenges to empirical research, it is likely that the 'myth of enclaves' will neither be confirmed nor debunked in the near future. Instead, the myth will live on and, hopefully, inspire researchers to design new and innovative ways to study algorithms at work in our daily lives.

Further reading

The most comprehensive overview of the relationship between algorithms and users is *The Marketplace of Attention* (2014) by James G. Webster. This book also offers a sobering view on the polarization and fragmentation debate underlying the 'myth of enclaves'. Matthew Hindman's 2009 book *The Myth of Digital Democracy* continues to be a relevant reminder of the power of algorithms to concentrate rather than disperse public attention. Solomon Messing and Sean Westwood's paper from 2014, "Selective Exposure in the Age of Social Media", provides a good overview of empirical studies on polarization and critically evaluates the experimental approach in online settings.

Notes

1 https://developers.facebook.com/docs/graph-api/reference/v2.9/user/home, accessed May 4, 2017.
2 www.facebook.com/terms, accessed May 4, 2017. It is possible to apply for a scraping permit from Facebook, but Facebook explicitly prohibits scraping for the purpose of "academic consumption" www.facebook.com/apps/site_scraping_tos.php, accessed May 4, 2017.
3 Thanks to Andreas Gregersen for making me aware of this connection to game theory.

References

Bakshy, E., Messing, S. and Adamic, L. (2015) "Exposure to ideologically diverse news and opinion on Facebook." *Science*, 348(6239), 1130–1132. doi:10.1126/science.aaa1160

Bhargava, P., Brdiczka, O. and Roberts, M. (2015) "Unsupervised modeling of users' interests from their Facebook profiles and activities." *Proceedings of the 20th International Conference on Intelligent User Interfaces*, ACM, New York, NY (pp. 191–201).

Bozdag, E. (2013) "Bias in algorithmic filtering and personalization." *Ethics and Information Technology*, 15, 209–227.

Braun, J. and Gillespie, T. (2011) "Hosting the public discourse, hosting the public: When online news and social media converge." *Journalism Practice*, 5(4), 383–398.

Brin, S. and Page, L. (1998) "The anatomy of a large-scale hypertextual web search engine." *WWW7: Proceedings of the Seventh International Conference on World Wide Web 7*, Elsevier, Amsterdam (pp. 107–117).

Bucher, T. (2017) "The algorithmic imaginary: Exploring the ordinary affects of Facebook algorithms." *Information, Communication & Society*, 20(1), 30–44. doi:10.1080/1369118X.2016.1154086

Cheng, J., Adamic, L., Dow, P. A., Kleinberg, J. M. and Leskovec, J. (2014) "Can cascades be predicted?" *WWW '14: Proceedings of the 23rd International Conference on World Wide Web*, ACM, New York, NY (pp. 925–936).

Chung, T. S., Wedel, M. and Rust, R. T. (2016) "Adaptive personalization using social networks." *Journal of the Academy of Marketing Science*, 44(1), 66–87. doi:10.1007/s11747-015-0441-x

Das, A. S., Datar, M., Garg, A. and Rajaram, S. (2007) "Google news personalization: Scalable online collaborative filtering." *WWW '07: Proceedings of the 16th International Conference on World Wide Web*, ACM, New York, NY (pp. 271–280).

Debatin, B., Lovejoy, J. P., Horn, A.-K. and Hughes, B. N. (2009) "Facebook and online privacy: Attitudes, behaviors, and unintended consequences." *Journal of Computer-Mediated Communication*, 15(1), 83–108. doi:10.1111/j.1083-6101.2009.01494.x

Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H. E. and Quattrociocchi, W. (2016) "The spreading of misinformation online." *Proceedings of the National Academy of Sciences*, 113(3), 554–559. doi:10.1073/pnas.1517441113

Dillahunt, T. R., Brooks, C. A. and Gulati, S. (2015) "Detecting and visualizing filter bubbles in Google and Bing." *CHI EA '15: Proceedings of the 33rd Annual ACM Conference: Extended Abstracts on Human Factors in Computing Systems*, ACM, New York, NY (pp. 1851–1856).

Dutton, W. H., Reisdorf, B. C., Dubois, E. and Blank, G. (2017, January 5) *Search and Politics: The Uses and Impacts of Search in Britain, France, Germany, Italy, Poland, Spain, and the United States*. Quello Center Working Paper. Retrieved from https://ssrn.com/abstract=2960697

Edelman, B. and Lai, Z. (2016) "Design of search engine services: Channel interdependence in search engine results." *Journal of Marketing Research*, 53(6), 881–900. doi:10.1509/jmr.14.0528

Eslami, M., Karahalios, K., Sandvig, C., Vaccaro, K., Rickman, A., Hamilton, K. and Kirlik, A. (2016) "First I 'like' it, then I hide it: Folk theories of social feeds." *CHI '16: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, ACM, New York, NY (pp. 2371–2382).

Feuz, M., Fuller, M. and Stalder, F. (2011) "Personal web searching in the age of semantic capitalism: Diagnosing the mechanisms of personalisation." *First Monday*, 16(2).

Flaxman, S., Goel, S. and Rao, J. M. (2016) "Filter bubbles, echo chambers, and online news consumption." *Public Opinion Quarterly*, 80(S1), 298–320. doi:10.1093/poq/nfw006

Goel, S., Anderson, A., Hofman, J. and Watts, D. J. (2016) "The structural virality of online diffusion." *Management Science*, 62(1), 180–196. doi:10.1287/mnsc.2015.2158

Granka, L. A. (2010) "The politics of search: A decade retrospective." *The Information Society*, 26(5), 364–374. doi:10.1080/01972243.2010.511560

Guo, J., Zhang, P., Zhou, C., Cao, Y. and Guo, L. (2013) "Personalized influence maximization on social networks." *CIKM '13 Proceedings of the 22nd ACM International Conference on Information & Knowledge Management*, ACM, New York, NY (pp. 199–208).

Hannak, A., Sapiezynski, P., Molavi Kakhki, A., Krishnamurthy, B., Lazer, D., Mislove, A. and Wilson, C. (2013, May 13–17) "Measuring personalization of web search." *WWW '13 Proceedings of the 22nd International Conference on World Wide Web*, ACM, New York, NY (pp. 527–538).

Helles, R., Ørmen, J., Radil, C. and Jensen, K. B. (2015) "The media landscapes of European audiences." *International Journal of Communication*, 9, 299–320.

Hindman, M. (2009) *The Myth of Digital Democracy*. Princeton, NJ: Princeton University Press.

Jacobson, S., Myung, E. and Johnson, S. L. (2016) "Open media or echo chamber: The use of links in audience discussions on the Facebook pages of partisan news organizations." *Information, Communication & Society*, 19(7), 875–891. doi:10.1080/1369118X.2015.1064461

Jones, J. J., Bond, R. M., Bakshy, E., Eckles, D. and Fowler, J. H. (2017) "Social influence and political mobilization: Further evidence from a randomized experiment in the 2012 U.S. presidential election." *PLoS ONE*, 12(4), e0173851. doi:10.1371/journal.pone.0173851

Kliman-Silver, C., Hannak, A., Lazer, D., Wilson, C. and Mislove, A. (2015) "Location, location, location: The impact of geolocation on web search personalization." *IMC '15: Proceedings of the 2015 Internet Measurement Conference*, ACM, New York, NY (pp. 121–127).

Kulkarni, C. and Chi, E. (2013) "All the news that's fit to read: A study of social annotations for news reading." *Chi'13: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, New York, NY (pp. 2407–2416).

Kuzel, A. J. (1999) "Sampling in qualitative inquiry." B. F. Crabtree and W. L. Miller (eds.), *Doing Qualitative Research*. Thousand Oaks, CA: Sage Publications (pp. 31–44).

Larson, R. and Csikszentmihalyi, M. (1983) "The experience sampling method." *New Directions for Methodology of Social Behavioral Science*, 15, 41–56.

Liu, J., Dolan, P. and Pedersen, E. R. (2010) "Personalized news recommendation based on click behavior." *IUI '10: Proceedings of the 15th International Conference on Intelligent User Interfaces*, ACM, New York, NY (pp. 31–40).

McGee, M. (2013) "EdgeRank is dead: Facebook's news feed algorithm now has close to 100K weight factors." *Markering Land*. Retrieved from http://marketingland.com/edgerank-is-dead-facebooks-news-feed-algorithm-now-has-close-to-100k-weight-factors-55908

Messing, S. and Westwood, S. J. (2014) "Selective exposure in the age of social media." *Communication Research*, 41(8), 1042–1063. doi:10.1177/0093650212466406

Micarelli, A., Gasparetti, F., Sciarrone, F. and Gauch, S. (2007) "Personalized search on the world wide web." In B. Peter, K. Alfred and N. Wolfgang (eds.), *The Adaptive Web*. Berlin: Springer-Verlag (pp. 195–230).

Nahon, K. (2016) "Where there is social media there is politics." In A. Bruns, G. Enli, E. Skogerbo, A. O. Larsson and C. Christensen (eds.), *The Routledge Companion to Social Media and Politics*. New York, NY and London: Routledge.

Newman, N., Fletcher, R., Levy, D. A. L. and Nielsen, R. K. (2016) *Digital News Report 2016*. Reuters Institute for the Study of Journalism, University of Oxford. Retrieved from http://reutersinstitute.politics.ox.ac.uk/sites/default/files/Digital-News-Report-2016.pdf?

Ørmen, J. (2015) "Googling the news: Opportunities and challenges in studying news events through Google search." *Digital Journalism*, 4(1), 107–124.

Ørmen, J. and Thorhauge, A. M. (2015) "Smartphone log data in a qualitative perspective." *Mobile Media & Communication*, 3(3), 335–350. doi:10.1177/2050157914565845

Pariser, E. (2012) *The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think*. New York, NY: Penguin Books.

Powers, E. (2017) "My news feed is filtered?: Awareness of news personalization among college students." *Digital Journalism*, 1–21. doi:10.1080/21670811.2017.1286943

Rader, E. and Gray, R. (2015) "Understanding user beliefs about algorithmic curation in the Facebook news feed." *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, New York, NY (pp. 173–182).

Rogers, R. (2013) *Digital Methods*. Cambridge, MA: The MIT Press.

Ryczko, K., Domurad, A., Buhagiar, N. and Tamblyn, I. (2017) "Hashkat: Large-scale simulations of online social networks." *Social Network Analysis and Mining*, 7(1). doi:10.1007/s13278-017-0424-7

Vaccari, C., Valeriani, A., Barberá, P., Jost, J. T., Nagler, J. and Tucker, J. A. (2016) "Of echo chambers and contrarian clubs: Exposure to political disagreement among German and Italian users of Twitter." *Social Media + Society*, 2(3). doi:10.1177/2056305116664221

Webster, J. G. (2014) *The Marketplace of Attention: How Audiences Take Shape in a Digital Age*. Cambridge, MA and London: The MIT Press.