

Reviewer 1:

We want to thank you for your very detailed and helpful feedback. You greatly helped us improve the paper. Below, we list your comments and our reactions in bold. Please note that we respond only briefly to your general comments, as you kindly detailed them in your “further comments”. These comments are addressed in detail.

Summary of paper in a few sentences:

This paper discusses the effect of personalisation technologies on polarisation.

Reasons to accept:

- 1.- The paper addresses a relevant topic, that of the effect of personalisation technologies on polarisation
- 2.- The highlights important points of the debate
- 3.- The paper proposes an abstraction of the problem, identifying three levels or dimensions: individual, local and global

Reasons to reject:

- 1.- The paper addresses the problem mainly from a social science perspective, sidelining the technological and data science perspectives. In this sense, the paper may not fit well the topics of the journal

Thank you for your comment. The aim of the paper is indeed to advocate a perspective that is, in our opinion, underrepresented in the data-science literature. The personalization-polarization-hypothesis proposes that technology has an effect on a social dynamic. We argue, in a nutshell, that an empirical test of the hypothesis and the design of technology that prevents undesired effects requires a thorough theoretical understanding of this social dynamic. The field of complexity research, which originates from physics and mathematics, provides a framework for such a theoretical analysis. In our responses to your detailed comments, we show how we tried to better convey this message. That is, we tried to more explicitly demonstrate what the perspective of our paper adds to the technological and data-science perspective. In addition, we also tried to make more explicit that theoretical modeling can contribute to the development of new personalization technology that prevents undesired effects on opinion dynamics.

- 2.- The paper presents a wide range of concepts in an ambiguous manner. For example, recommendation algorithms, and personalisation algorithms are not the same, and acknowledging this distinction is important for the problem being discussed. Similarly with the effect of personalisation algorithms vs. the effect of communication via social networks.

We agree with both points. To improve, we added a paragraph defining the concept of personalization and acknowledging the diversity of different algorithms and purposes. We also make explicit that the personalization-polarization-hypothesis and our work is focused on one core aspect of these algorithms: their tendency to expose users to political and cultural views that are in line with their own opinions. What is more, we tried to use the terms “personalization algorithms” and “recommend systems” more consistently throughout the paper. That is, we now refer only to “personalization algorithms”.

Concerning your second point, we also agree and thank you for pointing us to the confusion. There are indeed two distinct hypotheses. First, Sunstein [1] argued that online communication intensifies polarization. Second, Pariser and, in particular, modelers inspired by his work added that personalization has an effect on top of the effect proposed by Sunstein [2–4]. The two hypotheses are closely related, as they are based on very similar theoretical arguments. What is more, the effects

proposed by the two hypotheses may reinforce each other. Nevertheless, they need to be clearly distinguished, as you point out. Below, we describe how we adjusted the manuscript accordingly.

3.- The paper seems to be missing important literature from computer science and from computational social science research.

Thanks. We describe below how we tried to tackle this point, listing further references.

Further comments:

The paper addresses an important topic and discusses important points within the debate. It is also well written and structured. However, multiple important issues are not being considered in this discussion.

Thank you for your nice words and your detailed and helpful comments.

First, please note that there are important distinctions between personalisation, recommendation, and the effect of social networks. Communication via social networks have enabled us to connect with individuals that are physically far from us, to be exposed to high amounts of information from a variety of sources, to hide on the anonymity of user accounts, etc. These are aspects of the communication medium that may incentivise polarisation, but these aspects are different than the algorithmic aspects of personalisation or recommendation methods.

This comment details your second general comment from above. We agree that the effects online communication and personalization need to be distinguished. To avoid confusion, we made more explicit that there are indeed two separate arguments. The following paragraphs were added on Page 4:

Homophily is a strong force in human interaction also in the absence of personalization [19,33]. There is a rich empirical literature documenting that humans tend to interact with others who hold similar demographic attributes, have similar social status, and hold similar opinions [34–36]. In addition, it has been proposed that the Internet makes it especially easy to find and contact like-minded individuals, allowing in particular users with extreme opinions to form online enclaves that would be very difficult to establish and maintain offline [37]. Sunstein argued that this high degree of homophily is potentially harmful, as it intensifies processes of opinion polarization, the development of antagonistic groups, where opinion differences between groups intensify and positions between the two extremes of an opinion spectrum are increasingly sparsely occupied [38]. Informed by social-psychological research [39,40], he argued that strong homophily intensifies users' opinions, as they are mainly exposed to online content containing persuasive information that reinforces their initial opinions. As opinions of users from the left end of the political spectrum grow more leftist and users identifying with rightist political views also grow more extreme, opinion differences between the political camps increase and the opinion distribution polarizes.

Scholars and experts have noted that personalization technology is yet another source of homophily, a hypothesis that found empirical support with research on Facebook [41]. Personalization algorithms might have further intensified opinion polarization and may even be responsible for the growing opinion polarization observed in many western countries [42,43]. Here, we refer to this conjecture as the personalization-polarization hypothesis.

Next, your comments showed us that our illustration of local aspects did not focus on personalization effects but on the effects of communication in online social networks. The central point that we did not convey clearly is that personalization can intensify the effects of local aspects of online communication on polarization. As this insight has actually never been formally demonstrated, we conducted a new simulation experiment and added a figure driving home the insight that the one-to-many communication regime of many online social networks fosters polarization more when

personalization is strong. We rewrote large parts of Section 3.2 and repeat below the main paragraphs:

“The increased tendency to generate opposing clusters under the one-to-many regime is relevant for the personalization-polarization debate, because the difference between the two communication regimes is greater when homophily is increased. To demonstrate this, we conducted a simulation experiment with Axelrod’s model of cultural dissemination, extending the analyses of [105].

[...]

Figure 3 reports the results of our simulation experiment. In this experiment, we studied seven different values of personalization strength h and conducted 200 independent simulation runs per condition with one-to-one and with one-to-many communication.

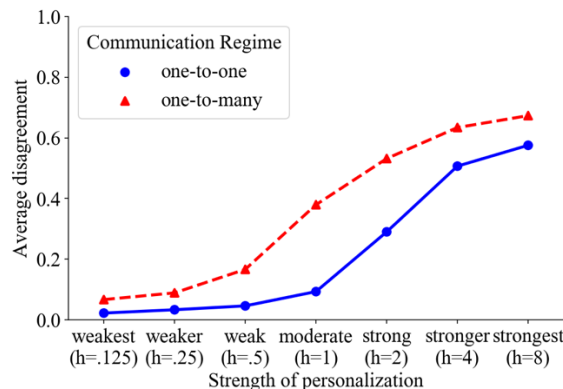


Figure 3. Average disagreement when increasing homophily in the Axelrod model under two different communication regimes (averaged over 200 replications per homophily condition and regime).

[...]

In sum, the example of one-to-many communication effects illustrates that local-level aspects can impact opinion dynamics in social networks. The example also shows that the effect of personalization technology on opinion polarization can depend on local-level aspects. We conclude that a thorough analysis of the personalization-polarization hypothesis needs to consider relevant local-level aspects. Unfortunately, researchers are only starting to explore local-level aspects, which suggests that it is too early to draw conclusions about the truth of the personalization-polarization hypothesis.”

At no point in the paper the algorithmic aspects are being discussed. What are the particularities of collaborative filtering, matrix factorisation, content-based, or other recommendation methods that may increase polarisation? How can different user profiles, item profiles or matching methods may influence polarisation? How can the data that is being personalised influence polarisation?

Indeed, these aspects are not discussed in the paper, and we feel that they are not in the scope of this paper. We, therefore, chose to acknowledge the diversity of personalization techniques and to make explicit why we abstract from it. We also made clear that we focus on the central aspect that personalization technology has been designed and criticized for: their tendency to expose users to content that is similar to their interests. In addition, we added references to articles describing

different personalization approaches. [19–24] We added the following paragraph to Section 2 of the paper:

“Personalization algorithms have been developed for various online services including online social networks, search engines, and online markets [23–28]. What is more, for each of these services there is a vast number of different technical approaches to personalization. What these approaches share, however, is that they seek to infer individual users’ interests from information they provided, from their earlier behavior, and from the behavior of other individuals who share relevant attributes with the respective user. For instance, if a YouTube user regularly watches a certain car show on the platform, its algorithms will recommend other car-related content and content that other users who watched the same car show have selected in the past. As a consequence, users are exposed to content that is in one way or the other similar to the content they chose to consume earlier. This tendency to provide users to similar content and to limit their exposure to content that deviates from users’ interests and opinions is central to the debate about undesired social effects of the technology. Accordingly, we also focus on this central aspect, abstracting from the large variety of technical implementations of personalization.”

It is also important to consider that 'social connection' does not mean influence. Users may receive information that does not influence them in any way. Note that some users are more prone to polarisation than others. There is important research on the field of misinformation related to the topics discussed in this article, including the effects of confirmation biases, or the effect of human values and personality types on the spread of misinformation and polarisation. There are also relevant works on social media 'influence' and engagement. Note that not all users have the same influence than others (authorities, celebrities, etc.) have a higher degree of influence. They also tend to have more followers and hence higher chances to spread their messages.

We agree that there are many other aspects that could affect whether or not personalization affects polarization. The aim of our paper is to illustrate that factors on the individual, local, and global level might matter critically, highlighting three examples that are underrepresented in the debate about personalization. To better acknowledge that there are also further potentially important aspects, we mention possible aspects at the end of the three sections focusing on the three levels of analysis. For instance, we added at the end of Section 3.1 that individuals differ in the degree to which they are open to influence and their persuasiveness, as you describe. We also added references to the modeling literature on this aspect, mentioning that this is an important research . We write at the end of Section 3.1: *“In conclusion, alternative theories of the individual-level processes in communication network make opposing predictions about whether the personalization-polarization hypothesis is true or false. In reality, online communication may be best described by a hybrid of assumptions from rejection and reinforcement models, but without empirical information about which theory is true under what conditions, there are too many possible ways to combine assumptions of the competing theories into a single model. Furthermore, there are further individual-level factors that may have critical effects on whether polarization emerges or not. For instance, the described models abstract from possible heterogeneity between individuals. Some individuals may be more open to influence from online contacts than others, and some may exert stronger influence than others. To our knowledge, such heterogeneity has not been studied in the context of the two models, but for alternative opinion-dynamics models researchers documented critical effects [101–103].”*

Likewise, we added further aspects that we could not cover in Section 3.3.:

“The presented analysis of the effects of network clustering illustrates, in a nutshell, that the structure of the communication network can affect the degree to which personalization technology affects the outcomes of opinion-dynamics processes. Network clustering was taken here as an example and is just one of many potentially important global aspects. Other potentially relevant global aspects that have been shown to influence-model dynamics are demographic diversity [88,115,116], network segregation [108,117], the number of bridges connecting otherwise disconnected network clusters [115], and the existence of agents with many connections [107,118]. Considerable empirical and theoretical research is needed to understand whether and under what conditions global aspects affect how personalization technology affects opinion polarization. Without this research, however, it is not possible to evaluate whether or not online communication systems are a setting where personalization could affect polarization or whether global aspects prevent any desired or undesired effects.”

These features, such as the number of followers, as well as network structures and information cascades, have been studied in the context of multiple election campaigns, on the spread of misinformation and in the context of radicalisation studies. However, the authors mentioned that local-level aspects of communication and their effects on polarisation have not been considered. I would recommend the authors to take a look to the work of Filippo Menczer, Emilio Ferrara or Claudia Wagner.

Thanks for the comment as it helped us clarify our paper. First, it is not our intention to claim that there is no research on individual, local, or global-level variables, but we argue that there is too little of it and that the consequences that individual, local, and global-level variables have on personalization effects are often complex and poorly understood. Throughout the text, we corrected formulations that suggested the opposite, in order to make clearer that the research we highlight serves as illustrations of urgently needed work. For instance, we now write on Page 16:

“In sum, the example of one-to-many communication effects illustrates that local-level aspects can impact opinion dynamics in social networks. The example also shows that the effect of personalization technology on opinion polarization can depend on local-level aspects. We conclude that a thorough analysis of the personalization-polarization hypothesis needs to consider relevant local-level aspects. Unfortunately, researchers are only starting to explore local-level aspects, which suggests that it is too early to draw conclusions about the truth of the personalization-polarization hypothesis.”

Second, you pointed to research on the diffusion of fake and true content on the web. This work is indeed very relevant and needs to be mentioned. However, models of diffusion are concerned with the spreading of, for instance, a rumor in a network and, thus, typically do not model opinions, which makes it difficult to derive conclusions about opinion polarization from these models. To acknowledge this literature, to distinguish diffusion models from the models of social influence, and to explain the reader why we focus on social-influence models, we added the following paragraph to the introduction of the paper:

“Opinion dynamics models differ in important ways from models of diffusion in networks, a class of models that has been used to model, for instance the spreading of rumors and (mis)information in online social networks [16,17]. In both model classes, populations are represented as a set of nodes integrated in a network of social relationships. These connections allow nodes to pass information around or to exert influence on each other. In diffusion models,

it is assumed that nodes receive content from their neighbors and subsequently share it with their network neighbors. Models of social influence, in contrast, often do not attempt to model this spreading of information explicitly, but focus on opinion values to describe nodes. When two connected nodes interact, they exert social influence on each other, influencing the opinion value of their network neighbor. The central difference between the two classes of models is that diffusion models assume that for instance a piece of fake news can be passed on from one node to another. However, a node that, for instance, has never been exposed to the piece cannot pass its unawareness on to its network neighbors. In social-influence models, in contrast, influence can be bi-directional in that nodes can push and pull each other's opinions in all possible directions independent of their current state. We focus here on opinion dynamics models rather than diffusion models, as opinion dynamics models have a direct representation of opinions and can, therefore, be used to study the conditions of opinion polarization.”

We added references to the work you mentioned:

- Nematzadeh A, Ferrara E, Flammini A, Ahn Y-Y. Optimal Network Modularity for Information Diffusion. *Phys Rev Lett.* 113(8):088701.
- Weng L, Flammini A, Vespignani A, Menczer F. Competition among memes in a world with limited attention. *Sci Rep.* 2(335):1–9.

In summary, while this paper targets an important problem, and brings important aspects into discussion, the article seems to miss the technical angle, as well as a wider review of technical works from the computer science and computational social science fields. Moreover, the article does not seem to provide a novel and potentially disruptive view of the topic (as required for position papers <https://datasciencehub.net/content/guidelines-reviewers>) but more of an overview and an in depth discussion of the different aspects of the topic.

Whether a paper is disruptive or not is not a fruitful discussion, mainly because disruptiveness is not a scientific criterion. What you comment makes clear, however, is that we failed to clearly state implications of our insights for future research and technology development. In a nutshell, we feel that our paper has two general messages for data-scientists interested in opinion dynamics on the Internet. First, online communication platforms are complex systems and, therefore, require a rigorous analysis of the system's complexity. The examples provided in our manuscript illustrate that formal modeling helps identify aspects that have been overlooked before even in a field that attracts as much scholarly attention as the personalization-polarization-hypothesis. Second, formal models that have been carefully calibrated can serve as a tool to test and improve personalization technology, and to explore the consequences of technology on societal dynamics before being implemented. In order to better communicate these two insights, we rewrote large parts of the introduction and the conclusion section. In the introduction, we included the following paragraph.

“In a nutshell, we demonstrate that models' predictions about the effects of personalization on polarization hinge on assumptions about (i) individual behavior, (ii) individuals' local information environment and local communication structure, and (iii) global characteristics of the whole communication network. We conclude that these aspects need to be studied both theoretically and empirically before one should draw conclusions about the effects of personalization and we criticize that important contributions to the debate have so far failed to do so. While we echo the warning that personalization might have serious effects on societal processes, we demonstrate that experts, politicians, and scientists leap to conclusions when

they propose that personalization is responsible for increased polarization. Unlike other recent contributions to the debate [18,19], however, we do not conclude that personalization is an innocent technology, but point to gaps in the empirical literature that need to be filled before one can draw conclusions. Accordingly, we call for more research on communication in online environments, pointing to the potential of approaches that combine rigorous theoretical modeling with the emerging fields of data science and computational social science. While these fields provide innovative sources of data and powerful methods of data analysis, we argue that their potential may not be exploited if it is not combined with rigorous theoretical modeling of the complex dynamics emerging in online communication systems, an approach that is increasingly popular [20–22]. We also discuss how carefully calibrated formal models can inform the development of personalization technology that prevents undesired effects on opinion dynamics on the web.”

Also the concluding section has been rewritten. The two final paragraphs address the two general take-home messages of the paper:

“From our perspective, the most promising approach to deriving predictions about the future effects of personalization on opinion polarization is to develop empirically calibrated models, an endeavor that requires empirical and theoretical research from various disciplines [13]. Theoretical research is needed to identify those theoretical assumptions that have a critical impact on model predictions, as these assumptions need to be put to the test by empirical research. Our review has covered several aspects that require empirical investigation, but this list is not conclusive. To identify the most important mechanisms, modelers should invest more into comparing the predictions of alternative models [52,53,125–127]. Unfortunately, a recent review of the literature concluded that many contributors fail to highlight the similarities and differences between the model underlying their work and existing models [13], hampering the field’s ability to accumulate knowledge and move forward. To improve, modelers should invest more into identifying these critical model assumptions, understanding why their model generates outcomes that other models do not. Furthermore, theoretical work should not only derive predictions about when a given model generates certain outcomes, but should find conditions under which different models provide different predictions. These insights will point empirical researchers to the empirical settings where competing models can be tested against each other, which in turn will help modelers develop validated models.

The emerging fields of data science and computational social science provide novel computational tools, sources of data, and methods of analysis to study opinion dynamics in online environments. Without proper theoretical foundations, however, attempts to empirically quantify the amount of online polarization or network segregation will remain underutilized [128–130]. Informing research on the individual level, many online services offer application programming interfaces (APIs) that provide researchers with information about the content that users share online. In tandem with novel methods of sentiment analysis and topic modeling, this may allow testing assumptions about who is communicating what content to whom on the Internet [95,131]. In addition, controlled online experiments shed light on how users adjust their opinions as a result of online communication [68,96,132–134]. On the local level, models need to be enriched with empirical information on how often users are exposed to online content on different online platforms and when they decide to contribute to online debates. Finally, there have been advances in gathering, storing, and analyzing detailed information about global-level factors [120–122]. In particular, there is considerable research

on the structure of online communication networks, which make it possible to directly implement or regrow realistic communication networks in models of opinion dynamics [135–137]. When this empirical information is fed into a formal model of opinion dynamics, it will be possible to predict the collective dynamics arising from social influence and to study whether and to which degree personalization technology affects model dynamics. Empirically validated models of social influence dynamics will not only make it possible to predict the consequences of web personalization, but they can also serve as a powerful tool to virtually experiment with alternative personalization algorithms and to develop technology that prevents undesired effects on public debate and opinion dynamics. Theoretically informed and empirically grounded computational models allow programmers to experiment with alternative specifications of personalization algorithms and analyze when they outperform each other on dimensions such as accuracy, scalability, user experience, and computational efficiency. In addition, validated models will make it possible to predict undesired effects of personalization technology on societal processes such as public debate, opinion polarization, and political decision-making. These predictions will yield new tools to understand and design new algorithms that generate the best browsing experience for individual users without harming societal dynamics. As communication technology is critical to deliberative democracy, tools that help us understand its consequences are urgently needed.”

Reviewer 2.

Thank you very much for your careful review and the helpful comments. You greatly helped us improve the paper, pointing to missing information and unclear arguments. We have rewritten large parts of the manuscript and added new analyses, which we detail in the following. Below you find your comments and our reactions in bold.

Summary of paper in a few sentences:

The authors argue that to address the personalization- polarization hypothesis, a complex approach is needed. In three levels of analysis (individual, local, and, global) the authors set out their take on how to approach the personalization- polarization hypothesis, encouraging other researchers in the field to build upon this approach.

Reasons to accept:

The authors address an important topic, and an innovative way to study this topic. Their discussion in section 3 is very useful to go beyond the state of the art.

Reasons to reject:

Despite the manuscript being quite lengthy, there is no engagement with the key literature from political science/ political communication on this topic.

We added references to contributions to this literature reporting findings on ideological segregation of online and offline media and processes of opinion polarization:

- **Morris JS. The Fox News factor. Harvard Int J Press. 10(3):56–79.**
- **Stroud NJ. Media use and political predispositions: Revisiting the concept of selective exposure. Polit Behav. 30(3):341–66.**

- Barberá P. How Social Media Reduces Mass Political Polarization. Evidence from Germany, Spain, and the U.S. 2015.
- Abramowitz AI, Saunders KL. Is polarization a myth? J Polit. 70(2):542–55.
- Iyengar S, Hahn KS. Red Media, Blue Media: Evidence of Ideological Selectivity in Media Use. J Commun. 59(1):19-U6.
- Peterson E, Goel S, Iyengar S. Partisan selective exposure in online news consumption: evidence from the 2016 presidential campaign. Polit Sci Res Methods. :1–17.

Moreover, the authors do not provide clear definition of either of the key concepts in the manuscript.

To better define the concept of personalization, we adjusted the first two paragraphs of Section 2. The first paragraph provides an intuitive definition and a set of examples. The second responds to a comment by Reviewer 1, acknowledging that there is a vast amounts of alternative approaches to personalization and identifying the central aspects of personalization that the debate about personalization and polarization is concerned with.

“Personalization is ubiquitous on the Internet. Providers of Internet services seek to tailor their products to the needs and interests of individual users. Search engines, for instance, rank the results of users’ search queries according to the interests of the individual user. When the authors of the present article google the term “polarization”, for example, websites discussing political polarization should be ranked higher than websites of manufacturers selling “polarized” sunglasses, even though both websites contain the search term. Likewise, online markets recommend products based on the purchases of other customers who bought similar products in the past and online social networks sort incoming messages according to the similarity between the user and the source of the message. Personalization has tremendously improved online companies’ services, making it easier for users to navigate the immense and rapidly growing amount of online content. Personalization has also turned into a multibillion-dollar business area, increasing engagement on online platforms using this technology, and allowing advertisers to directly target potential customers.

Personalization algorithms have been developed for various online services including online social networks, search engines, and online markets [23–28]. What is more, for each of these services there is a vast number of different technical approaches to personalization. What these approaches share, however, is that they seek to infer individual users’ interests from information they provided, from their earlier behavior, and from the behavior of other individuals who share relevant attributes with the respective user. For instance, if a YouTube user regularly watches a certain car show on the platform, its algorithms will recommend other car-related content and content that other users who watched the same car show have selected in the past. As a consequence, users are exposed to content that is in one way or the other similar to the content they chose to consume earlier. This tendency to provide users to similar content and to limit their exposure to content that deviates from users’ interests and opinions is central to the debate about undesired social effects of the technology. Accordingly, we also focus on this central aspect, abstracting from the large variety of technical implementations of personalization.”

In the new version of the paper, the concept of polarization is defined already in the first paragraph as “a dynamic where a population falls apart into subgroups with increasingly opposing opinions.” The definition is mentioned a second time on Page 4: “As opinions of users from the left

end of the political spectrum grow more leftist and users identifying with rightist political views also grow more extreme, opinion differences between the political camps increase and the opinion distribution polarizes.”

For a great recent lit review on media diversity in online environment (and potential risks), see <https://www.tandfonline.com/doi/full/10.1080/21670811.2020.1764374>

Thanks for pointing us to the brand-new paper. We now refer to it on Page 4.

Moreover, the authors do not give any recommendations on how to (not) use the method for a broad variety of scholars interested in this hypothesis.

We feel that an introduction to the complexity approach is out of the scope of this paper, as complexity research is a vast field comprised of researchers from diverse disciplines who, in addition, also use a variety of empirical and theoretical methods. However, to respond to your comment, we tried to make clearer recommendations about future research in the conclusion section. Rather than advocating the complexity approach in general, we now argue that the integration of empirical and theoretical research is required and point to gaps in the empirical and the theoretical literature.

“We advocate here an approach that combines formal theoretical modeling with empirical research. On the one hand, a purely empirical approach to testing the personalization-polarization hypothesis can lead to false conclusions. Assume, for instance, that an empirical study quantified the degree of personalization-induced homophily in various settings and found no correlation with opinion polarization in these settings. This finding certainly challenges the personalization-polarization hypothesis. However, in complex systems effects can take very long to unfold and can then be very abrupt and strong. In Panel A of Figure 1, for instance, polarization remained low for a long time, until it grew rapidly [52]. In addition, personalization algorithms are still being elaborated. The fact that they have not contributed to opinion polarization so far, does not imply that further advances in personalization will also remain without negative effects [86]. This suggests that the empirical observation that personalization so far appears to be relatively mild and its effects on opinions modest [18,41], should not lead one to conclude that personalization will remain an innocent technology in the future. On the other hand, also a purely theoretical approach will fail to generate reliable predictions about personalization effects, even when analytical and computational tools are used to derive predictions. Our review of the opinion-dynamics literature provided several examples of modeling decisions that can have big impact on the model’s predictions. As a consequence, models relying on assumptions that have not been backed up by rigorous empirical research in the context of online social networks may fail to make true predictions and, in addition, will not be considered reliable tools for anticipating future opinion dynamics. From our perspective, the most promising approach to deriving predictions about the future effects of personalization on opinion polarization is to develop empirically calibrated models, an endeavor that requires empirical and theoretical research from various disciplines [13]. Theoretical research is needed to identify those theoretical assumptions that have a critical impact on model predictions, as these assumptions need to be put to the test by empirical research. Our review has covered several aspects that require empirical investigation, but this list is not conclusive. To identify the most important mechanisms, modelers should invest more into comparing the predictions of alternative models [52,53,126–128]. Unfortunately, a recent

review of the literature concluded that many contributors fail to highlight the similarities and differences between the model underlying their work and existing models [13], hampering the field's ability to accumulate knowledge and move forward. To improve, modelers should invest more into identifying these critical model assumptions, understanding why their model generates outcomes that other models do not. Furthermore, theoretical work should not only derive predictions about when a given model generates certain outcomes, but should find conditions under which different models provide different predictions. These insights will point empirical researchers to the empirical settings where competing models can be tested against each other, which in turn will help modelers develop validated models. ”

To be sure, we do not argue that the complexity approach is superior to contributions from other fields. In contrast, the three levels of analysis that our manuscript is concerned with show that an interdisciplinary effort is needed.