

Effects of Moderation and Opinion Heterogeneity on Attitude towards the Online Deliberation Experience

Simon T. Perrault*

Yale-NUS College,
Singapore University of Technology and Design (SUTD)
Singapore
Korea Advanced Institute of Science and Technology
(KAIST)
Korea
perrault.simon@gmail.com

Weiyu Zhang†

Department of Communications and New Media
National University of Singapore (NUS)
Singapore
weiyu.zhang@nus.edu.sg

ABSTRACT

Online deliberation offers a way for citizens to collectively discuss an issue and provide input for policy makers. The overall experience of online deliberation can be affected by multiple factors. We decided to investigate the effects of moderation and opinion heterogeneity on the perceived deliberation experience, by running the first online deliberation experiment in Singapore. Our study took place in three months with three phases. In phase 1, our 2,006 participants answered a survey, that we used to create groups of different opinion heterogeneity. During the second phase, 510 participants discussed about the population issue on the online platform we developed. We gathered data on their online deliberation experience during phase 3. We found out that higher levels of moderation negatively impact the experience of deliberation on perceived procedural fairness, validity claim and policy legitimacy; and that high opinion heterogeneity is important in order to get a fair assessment of the deliberation experience.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in collaborative and social computing**; *Collaborative and social computing systems and tools*; *Web-based interaction*.

*Both authors contributed equally to the work.

†Both authors contributed equally to the work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300247>

KEYWORDS

Civic engagement, deliberation, deliberative democracy, participation, public sphere

ACM Reference Format:

Simon T. Perrault and Weiyu Zhang. 2019. Effects of Moderation and Opinion Heterogeneity on Attitude towards the Online Deliberation Experience. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland UK*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3290605.3300247>

1 INTRODUCTION

Online deliberation is a young interdisciplinary field that studies deliberative processes with the use of information and communication technologies. The deliberation process may last for days, or even weeks, and involves exchanging ideas on deliberation platforms. The outcome of the discussion may be used to provide input for policy makers.

Compared to online discussion forums, online deliberation is more guided to have participants focus on specific and important topics, with the goal of achieving mutual understanding. One of the key metrics used in deliberation is attitude change [10], i.e. differences in the levels of agreement to specific topic statements before and after deliberation. However, few studies [4] have examined participants' attitudes towards the deliberation experience itself, especially the legitimacy of this relatively new political practice.

Previous work in the field generally showed positive effects of online deliberation for democracy [15, 24] and was conducted in mature democracies, e.g. the USA [10], or western European countries [24]. Many factors may impact the deliberation outcomes, such as the topics of deliberation [28], the rules governing the deliberation [39], and the participants of deliberation activities [38].

In this paper, we present the first study conducted in Singapore, dubbed as an "authoritarian democracy" [37], in which we focus on the effect of moderation and opinion heterogeneity on attitudes towards the deliberation experience.

Our study that lasted for three months was done in three steps: (1) pre-deliberation survey, (2) deliberation itself and (3) post-deliberation survey. Our pre-survey was done with 2,006 Singaporean citizens, representative of the country population. For the deliberation phase, we developed our own open-source platform that enables simple discussions and provides educational material before each round of discussion. A total of 510 participants followed phases 2 and 3.

Our results suggest that both moderation and opinion heterogeneity have an impact on attitudes towards the deliberation experience: (1) Participants who joined low moderation platforms perceived higher policy legitimacy, and thought of their discussion partners' claims as more valid, than those who joined high moderation platforms. However, moderation made a marginally significant difference in perceived procedural fairness; (2) Participants who were put in the group with those mainly agreed with the policy gave higher scores in perceived procedural fairness than both the group with those who mainly disagreed with the policy and the group with those who were mixed in the views. For validity claims and policy legitimacy, participants who were put in the group with those mainly disagreed with the policy gave lower scores than both the group with those who mainly agreed with the policy and the group with those who were mixed in their views.

To summarize, the contribution of this paper is three-fold:

- We present an open-source deliberation platform that we developed,
- We describe the first online deliberation experiment in Singapore,
- We investigate the effect of moderation and opinion heterogeneity on attitudes towards the online deliberation experience.

2 RELATED WORK

Our work specifically focuses on policy deliberation, which is different from other deliberative forum or more casual online discussions. The goal of the whole experiment is also to provide citizens a space to discuss social policies, deliberate and provide input for policy makers. We will first discuss the impact of moderation and opinion heterogeneity in online deliberation. Then we will discuss how existing platforms dealt with both moderation and heterogeneity, and how our platform builds from these existing efforts.

Moderation in Online Deliberation

Moderation plays an important role in online discussion forums, and has been extensively discussed. Some work, like Towne and Herbsleb [32], advocate for a minimal amount of moderation to avoid trolling and flame wars. Hanasono and

Yang [13] pointed that even on a support group for racially discriminated people, an undisclosed proportion of messages were inappropriate or insulting. Binns [1] described moderation as labour intensive and divided moderation to two different approaches. Pre-moderation, means that comments need to be validated by a moderator to be published, which prevents battle between users and moderators. With post-moderation, comments will be first published, then potentially moderated. They also discussed the effect of adding a "karma" system on the platform, with a karma threshold required to pass pre-moderation. Lazer et al. [22] discussed moderators' strategies during their moderation work, while Epstein and Leshed [8] highlighted the two main roles of moderators: (1) managing the stream of comments and (2) interacting with commenters [6, 23].

The impact of moderation on online discussion forums has been partially investigated. For example, Meyer and Carey [25] found that moderation would negatively affect the number of posts, as people are more likely to react to inappropriate posts, generating more activities and giving users a sense of virtual community. Other negative effects include higher level of suspicion towards the deliberation process, self-censorship from users [36] and even potential exclusion of traditionally underrepresented populations [33], which would then hurt the representativeness of online discussion.

Wise et al. [34] conducted the only study so far which included moderation as an independent variable in their design, with two levels: moderation vs. no moderation. However, their focus was on the moderation awareness, and how this awareness could lead to higher intent to participate in the process.

While these previous works guided decisions on how moderation should be done on our platform, e.g. to include underrepresented groups, none of them investigated the effect of moderation on participants' attitudes towards the deliberation experience.

Opinion Heterogeneity in Online Deliberation

Heterogeneity is a double-edge sword for online deliberation. On one hand, it may be perceived as a threat to the harmony of social relationships, and would overall make it harder to create political mobilization [27]. On the other hand, Price et al. [28] saw diverse opinions as benefiting democracy, because exposure to disagreement contributes to people's ability to generate reasons and improves argumentation. Brundidge [2] proposed the inadvertency thesis to explain how people get exposed to heterogeneous opinions. Their findings suggest that users tend not to seek out political difference, but rather end up finding it inadvertently. For example, Zhang and Chang [39] reported that everyday political talk could be a good opportunity for people to find political difference. Brundidge [2] also pointed out that the

frequency of online political discussion is positively related to the heterogeneity of one's political discussion network.

There are multiple levels of heterogeneity, depending on the level of agreement or disagreement to attitude items. High levels of disagreement or agreement would strongly affect participants' evaluations of deliberation processes [31]. This same work also suggests that one's satisfaction, potential reevaluation of opinion and expected future participation are not affected by low levels of disagreements, but high levels of disagreements could lower satisfaction and future engagement. Finally, Wojcieszak and Price [35] conducted a study on perceived vs. actual disagreement. Their results show that perceived disagreement has strong effects on satisfaction or future participation, whereas objective disagreement does not. Their results also confirm previous research on the negative effects of high disagreement on satisfaction [38] and future engagement [31].

Existing Deliberation Platforms

HCI works on online deliberation usually focus on platforms and unique features implemented by each of these. No prior work published in HCI venues has investigated the effect of moderation and opinion heterogeneity on the online deliberation experience.

Macintosh [24] defines three types of deliberation platforms: informative, consultative and participative. Informative platforms are platforms used by governments to produce and deliver information for use by citizens. Consultative platforms add a way for citizens to provide feedback on materials received. Participative platforms allow citizens to actively engage in defining the process and content of policy-making. Our work falls in this last category, which enables users to actively collaborate with each other, and discuss public issues. Our review does not cover aggregation platforms that aim at presenting different opinions and help users locate their own on the spectrum of opinions as well as reflect and refine their opinions, such as Opinion Space [9], ConsiderIt [19], Reflect [20], and Poli [30].

Our platform builds from existing collaborative platforms. MIT Deliberatorium [18] takes users' input in the form of a deliberation map, a tree-structured network of posts representing a single issue, idea, or pro/con argument. It is based on the Issue-Based Information System (IBIS) model of informal argumentation. Arrangement of posts is based on topic, as opposed to time as seen on regular forums. We used a similar hierarchical representation of discussion to make it easier to follow conversations.

Another early example of collaborative online deliberation was the Virtual Agora project [26]. Inspired by Deliberative Polls [11], the system allowed participants to share web-based information and discuss using audio both in real-time and asynchronously. The Virtual Agora [26] project was able

to enhance the salience of citizen identity, which in turn is correlated with higher political engagement. PICOLA was a two-phase experiment regarding online citizen deliberation [3]. Phase I of the study compared real-time online deliberation to face-to-face deliberation and the additional phase II explored the effects of asynchronous deliberations. PICOLA came with an education phase for participants (implemented through an online reading room), a discussion phase (using a real-time online interaction), and a reflection phase (where participants could continue discussions through asynchronous online forum and surveys). The results of the deliberation session of PICOLA were however not formally published. Finally, Deme enabled collaborative drafting, focused discussion and decision making [5]. Designed for small group interaction, it features meeting rooms where members could self-organize deliberative sessions.

Compared to aggregation platforms, our system offers an online space for citizens to discuss and deliberate. We followed a similar process to PICOLA with an education phase, to mitigate potential lack of knowledge of our participants on the discussion topics, and used a hierarchical structure of discussion roughly similar to MIT Deliberatorium [18].

Civic Tech in Asia

Civic Tech has been rising in Asia during the last few years. For example, Factful [16] is a platform that allows taxpayers to participate in the discussion of South Korea's government budget. The platform allows users to read about budget planning and proposes news article and other sources. This allows participants to understand the government choices, and be more critical while checking for facts. Similarly, BudgetMap [17] is an issue-driven visualization platform that allows taxpayers to explore and tag the links between budget and social issues and was deployed in Seoul. Finally, vTaiwan [14] is an open consultation process that brings citizens and government together in online and offline spaces, to deliberate and reach rough consensus on national issues, and to craft national digital legislation. Our online deliberation experiment is one of the many efforts ongoing in the Asian Civic Tech scene.

Hypotheses

From the previous work, we hypothesize that both moderation and opinion heterogeneity will have significant impacts on deliberation outcomes, such as participants' perception of their deliberation experience. However, due to the mixed findings from previous research, we are not able to make directional hypotheses.

H1: Different levels of moderation will lead to different levels of perceived procedural fairness, perceived validity claim, and perceived policy legitimacy.

H2: Different levels of opinion heterogeneity will lead to different levels of perceived procedural fairness, perceived validity claim, and perceived policy legitimacy.

3 PROCESS AND PLATFORM

We developed our own platform which allowed us to run our experiment. We will first discuss some of the challenges for running an online deliberation study and how the whole process was carried out, then we will present our platform and how the platform supports our process. Finally, we will explain our participants sampling strategy.

Challenges

When we started the project, we faced two issues, highlighted by previous literature: (a) lack of knowledge of the issue discussed and (b) representativeness of the sample chosen for the study. Lack of knowledge may lead to ill-informed opinions, which usually have low value for policy making [11]. Lack of representativeness could undermine the potential acceptability of decisions made through deliberation to the whole society.

Lack of Knowledge. We addressed the lack of knowledge by running an early education phase. During that phase, a team of policy researchers, in consultation with the government agency in charge of policies related to population (our focus for this study). The material produced explained each of the sub-issues, with numbers and sources, in a slideshow format, which was easy to read. These slides focused on numbers, and we were careful about trying not to express official opinions, in order to avoid potential biases. We pre-tested the sets of slides for each issue with small groups of participants. The material was then inserted in the platform. Participants had to review the material at least once and were free to review as many times as they wanted.

Representativeness. Another recurring challenge for online deliberation is the potential lack of representativeness of the participants pool [7]. Most of these projects sampled participants on a voluntary basis. This leads to participants being potentially more engaged, and more tech-savy, excluding populations such as elderly or people with a low technology proficiency.

We addressed this issue by working with an online panel provider during the recruitment phase. This provider had access to a panel of over 20,000 Singaporean citizens. We set our demographic quotas to match the latest census data and monitored the sampling closely. We manually corrected under-representation of given groups by sending e-mails to individuals of said groups. As the deliberation phase lasted for three-weeks, we expected some drop-out. Drop-out in itself is not an issue, unless only specific groups keep dropping out, by lack of time or resources.

To mitigate this issue, we provided incentives to sustain participation, especially for lower-income brackets. Participants were thus compensated 50 SGD (37 USD) for taking part in the online deliberation phase. They could also earn up to 30 additional SGD (22 USD) depending on their participation (for posting, reading, reacting). While our original sample of 2,006 participants taking part in the pre-deliberation survey was overall representative of Singapore population on gender, ethnicity, housing, age, education and income criterion, we could not achieve a perfect representativeness for our sample of 510 participants taking part in the deliberation and reflection phase.

Deliberation Phases

The study ran in three different phases: (1) recruitment/pre-deliberation phase, (2) deliberation phase, (3) reflection/post-deliberation phase.

Pre-deliberation Phase. During this phase, we recruited 2,006 participants (see Representativeness for more information). Participants were invited to fill our pre-deliberation survey, which contained 40 close-ended questions about current policies regarding population. At the end of the survey, participants were asked if they wanted to join the deliberation. We sent an invitation to all participants from the respondents who were willing to join the deliberation. This phase lasted for one month.

We used the answers from the pre-deliberation survey to group our participants into the three levels of opinion heterogeneity as explained in the Independent Variable subsection of the next section.

Deliberation Phase. During this three-weeks phase, participants were invited to join our platform. Upon logging for the first time, participants were shown a video detailing the features of the platform and how to use it. Note that participants were identified with a unique numerical ID on the platform, and no identifying information were published for privacy reasons.

Each week, citizens were invited to discuss one of the three sub-issues (fertility, foreign workforce, integration of new citizens). Before proceeding to the discussion itself, participants would have to go through the education material for each sub-issue. Discussion would start with the title, and a brief, carefully crafted opening message from the admins. Administrators and moderators would not post any additional message throughout the process. At the end of each discussion, an online poll was set up based on moderators' reading of opinions/suggestions that emerged during the discussion. Participants were reminded to answer the poll before proceeding to the next sub-issue.

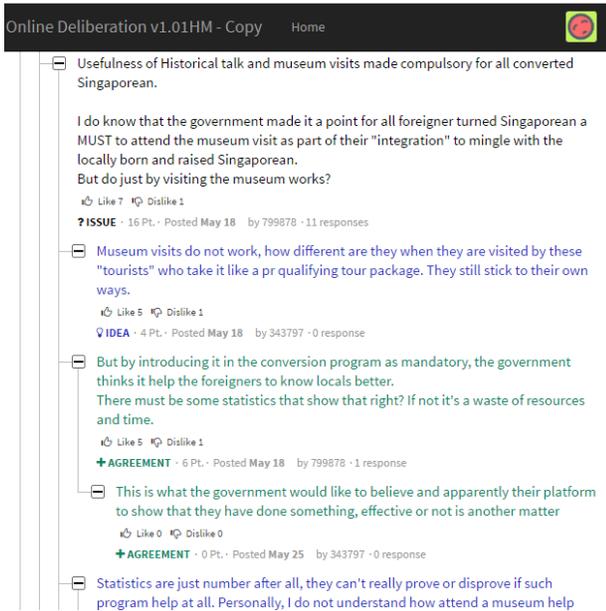


Figure 1: Basic discussion interface. The discussion is structured as a tree, users can either expand or shrink nodes by clicking on the "+/-" button.

Post-deliberation Phase. During this one-month phase, participants were asked to fill a post-deliberation survey. This survey essentially contained the same questions related to existing policies in Singapore, plus additional questions on their perception of their online deliberation experience. Additionally, we asked some questions related to their experience with the platform.

We specifically used the post-deliberation survey to measure our three dependent variables: (1) perceived procedural fairness, (2) perceived validity claim and (3) perceived policy legitimacy. The dependent variables are described in the subsection Dependent Variables of the next section.

Platform

Implementation. The platform was developed based on Vanilla Forums¹, an open source forum platform. The core of the platform is written in PHP, and data were stored in a MySQL database. Since Vanilla Forums supports multiple plugins, we did tweak or develop new plugins to fit our needs.

More advanced features were developed using JavaScript libraries. We used Underscore.js, Vis.js and JQuery for the video display and graph creation/visualization tools. The whole source code of our platform is available online at the following address: <http://onlinedeliberation.org/about>. An example of a discussion can be seen on Figure 1.

¹<https://vanillaforums.org/>

Part 1 Background

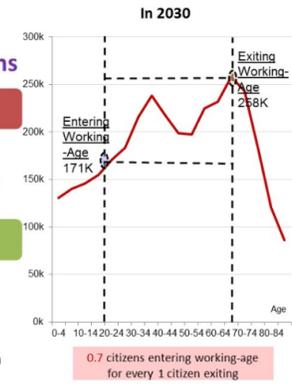
The Need for New Citizens

Shrinking Population

- Citizens' life expectancy increased to 83 years in 2013
- Median age to rise to 45 years old in 2025

Ageing Population

- Fewer younger citizens
- More citizens retiring than entering workforce
- The number of citizens aged 65 and above will triple to 900,000 in 2030



Source: A Sustainable Population for a Dynamic Singapore: Population White Paper, 2013.

Figure 2: Example of educational slide used for the sub-issue of Social Integration and New Citizens. This slide highlights the aging population issues, with less citizens entering the workforce and more exiting it by 2030. Each sub-issue came with a set of slides with data presented in a simple way.

Educational Slides & Quiz. Inspired by PICOLA [3], the platform featured an educational phase at the beginning of each sub-issue discussion. This phase consisted of a set of interactive PowerPoint slides (Figure 2) where each user would read through multiple educational slides and some self-paced interactive quiz that prompt citizen users to furnish their opinions about the issue. The user was given the option of re-watching the slides and re-taking the quiz as many times as they wanted, to encourage deeper engagement with the learning material.

Deliberation Graph. Our platform also incorporated a deliberation graph (see Figure 3), a mind map of salient posts (posted within a specific discussion) organized as an abstract topic model. Inspired in part by the deliberation map from MIT deliberatorium [18], the Deliberation Graph acts as a visual summary of an ongoing discussion. But unlike the deliberatorium, the Deliberation Graph does not allow users to structure or populate the graph, and is instead operated by the moderator (for high moderation instances) who can use it to organize ongoing discussion along topic-hierarchies that they deem fit. For low-moderation instances, the graph contains all the posts with their hierarchical relations and is automatically generated. This flexibility with re-organization of arguments helps to further counteract the common knowledge problem, namely the tendency of groups to focus on the knowledge held by the majority and discount information held by the minority [18]. The deliberation graph allowed for drawing shapes and directed connectors (each with changeable color and size attributes). Once the moderator modifies the graph on a specific instance, changes are auto-saved and

rendered instantly on the forum page for all users on the same instance. Figure 3 shows an example of such a graph. As a normal user, the graph allows zooming in and out on the graph and navigating from a node to a post that it refers to. This mapping from node to post allows the moderator to frame posts in interesting and thought provoking ways without having to quote them verbatim.

IBIS based color coding of posts. With the platform, we were also interested in providing argumentation support for users. However, we did not want to make the experience of posting arguments counter-intuitive as is the case with MIT deliberatorium [18]. We wanted to retain the normal time-wise sequential interface of traditional forums, but include interface element that could be used to organize posts based on the IBIS model of informal argumentation. In terms of implementation, this meant the user was given a choice of categorizing each of their posts into either an issue, idea, agreement (pro argument) or disagreement (con argument). A post belonged to each of these categories would have its own color code when displayed on the forum interface. To further situate the IBIS structure, we used a nested style of indentation for posts, so that a reply to a post would appear at a level nested under the original post. This way a user who has raised an issue or idea can have other users contribute pros and cons arguments under their post, offering the ability to have more focused discussions.

Gamification. We implemented both point systems and rewards in the platform. As discussed in the design process section, on top of the baseline reimbursement, an additional discretionary reward was made available based on the level of user’s participation during the deliberation. The user can accrue participation points, by both submitting posts and by engaging with others’ posts: 5 points gained by poster for

every instance of their post getting ‘liked’, 5 points lost by poster for every instance of their post getting ‘disliked’, 5 points gained by Replying to posts, 2 points gained through liking/disliking another’s post, 1 point gained for every instance of reading or viewing another’s post. The point system was tied to 3 milestone levels of participation, with each level unlocking additional monetary reward. These milestones were capped at 1000, 2000 and 3000, with a commensurate reimbursement of 10SGD, 20SGD and 30SGD respectively.

We did not implement additional feature or content unlocking with these milestone levels, because this can lead to conflict with one of the core principles of deliberation namely accessible reasons [12] (p.144). This principle requires that the reasons and content used while deliberating should be accessible to all citizens. Additionally, the use of social currency based incentives can lead to online hierarchy which impedes equal participation [29]. Keeping this in mind, we decided to not implement a leader board, although the individual participants’ particulars were still made publically visible on the user’s profile page.

4 ONLINE DELIBERATION EXPERIMENT

In this experiment, we consider the effect of moderation (2 levels) and opinion heterogeneity (3 levels) on the perceived deliberation experience. We rely on a between-subject design, and therefore split our participants into six different instances of our platform.

Participants and Procedure

We sent out an invitation to respondents from our pre-survey who consented to join the platform. Prospective participants would receive an email containing a unique link to one of the six instances of the platform they were assigned to. Among the 1200+ invitations we sent out, 510 participants logged into our platform at least once. Among these 510 citizen users, 56% were male; 86% were ethnic Chinese; 55% of people aged between 21 to 39 years old; 58% had a university or high degree: and 57% had a household monthly income below 10,000 SGD. Compared to the general population, our participants tended to be male, Chinese, younger, better educated and having higher income. A comparison between our participants and the singaporean population is shown in Table 1.

Following the pre-deliberation survey, during which participants provided their opinion on current policies, participants were assigned to one of our six instances (one per heterogeneity \times moderation combination). The deliberation phase was described in the Process and Platform section. Each week, participants were invited to discuss one of the sub-issues, in that order:

- (1) Singapore’s Low Birthrate

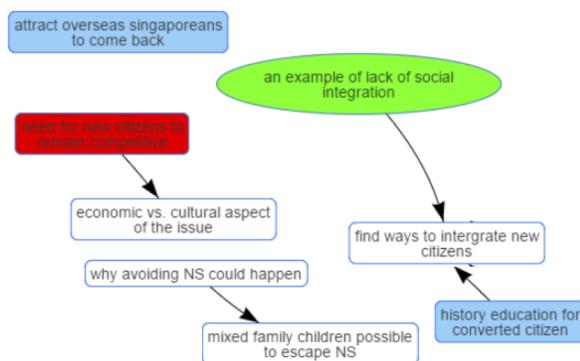


Figure 3: Details of a deliberation graph generated to facilitate the discussion. Suggestions appear in boxes with a blue background, negative points in red and positive ones in green.

Table 1: Population comparison between our participants and the Singaporean population. Data computed and aggregated from https://www.singstat.gov.sg/-/media/files/publications/reference/yearbook_2018/yos2018.pdf.

Criterion	Participants	SG Population
Male	56%	49%
Chinese	86%	74.4%
Age 20-39	55%	28.4%
University Degree	58%	30.7%
Income <S\$ 10,000	57%	63.4%

- (2) Foreign Domestic Workforce
- (3) Integration of New Citizens

After the deliberation, participants would answer a post-deliberation survey, in which we assessed the procedural fairness, validity claim and policy legitimacy as described in the Dependent Variables subsection.

Independent Variables

We considered two independent variables: moderation and opinion heterogeneity.

Moderation. We differentiated two levels of moderation, namely, **low** and **high**. There are significant differences in terms of amount of human moderator's work. In the **low** moderation conditions, moderators were instructed to keep the basic order of the discussion, such as posting the initial posts, removing duplicates and hiding offensive posts. In the **high** moderation conditions, moderators have to do much more work: firstly, moderators greeted every participant with a welcome message that links to the help page which specifies discussion rules, when the participant posted the first post. Secondly, moderators for the **high** moderation condition were in charge of generating and updating a deliberation graph that highlights good contributions and put them into a thought map at a daily frequency. An automated deliberation map was also available for the **low** moderation conditions, which means all posts will automatically enter a network map without any moderator's interference. Last but not least, moderators in the **high** moderation conditions were also managing inappropriate posts. In fact, the only two instances in which we had to temporarily ban users both happened in high moderation conditions.

To summarize, the role of the moderator in the **high** moderation condition covers the two roles highlighted by Epstein and Leshd [8] in addition to updating the deliberation graph daily. In the **low** moderation condition, moderators would strictly stick to basic order keeping without directly interacting with the users. Users were aware of the presence of the

Table 2: Group Composition for the Heterogeneity independent variable.

Heterogeneity Level	Group Composition
Low-Agree	$\frac{2}{3}$ agree, $\frac{1}{3}$ neutral
Low-Disagree	$\frac{1}{3}$ neutral, $\frac{2}{3}$ disagree
High-Mixed	$\frac{1}{3}$ agree, $\frac{1}{3}$ neutral, $\frac{1}{3}$ disagree

moderators by seeing moderators' posts such as the initial posts, and seeing messages such as "this post was hidden by a moderator" for hidden posts.

Four moderators were involved in moderating the discussions. A 7-page moderation guide (available as supplementary material) was developed to provide standard instructions, including technical details, strategies to take when certain situations emerge, and templates to use for greeting, warning, banning users. All moderators had to go through multiple rounds of training to ensure that they understand the moderation instructions fully. All four moderators are communication scholars with at least a master degree in the major. Three moderators who moderated the high moderation conditions rotated every week during the three-weeks discussion, to avoid conditions being confounded with moderators. Moderators in the high moderation conditions spent at least two hours a day to complete the moderation tasks. Thanks to the low workload, one moderator was tasked to moderate all low moderation conditions, to ensure the consistency among conditions.

Opinion Heterogeneity. To create our three groups of different opinion heterogeneity, we proceeded in two steps. First, we generated a tripartite grouping based on the heterogeneity in group members' opinions towards policy positions, measured in the pre-deliberation survey. The participants were first ranked on their level of agreement with policy positions (averaged across all three deliberation issues) and separated into 3 equal parts. The first tertile of this partition was assigned a classification of Disagree (agreement score $\leq \frac{2.67}{5}$), second tertile a classification of Neutral (agreement score of $\frac{3}{5}$ or $\frac{3.33}{5}$) and the third quintile a classification of Agree (agreement score $\geq \frac{3.67}{5}$).

Second, we randomly selected the $\frac{2}{3}$ of the "Agree" users and $\frac{1}{3}$ of the "Neutral" users to create the Low Heterogeneity-Agree condition. We repeated this process for the Low Heterogeneity-Disagree condition. The remaining users ($\frac{1}{3}$ of each original pool) was grouped as High Heterogeneity-Mixed. Table 2 summarizes how the groups were created. We will refer to each group according to their heterogeneity and opinion level. As such, our three levels for heterogeneity are: **Low-Mainly Agree (L-A)**, **Low-Mainly Disagree (L-D)**, **High-Mixed (H-M)**.

Dependent Variables

We measured three dependent variables, all related to their overall experience of deliberation.

Perceived policy legitimacy. This variable assessed participants' overall approval of the incorporation of deliberation in policymaking regarding population policies. It was measured by four items: (a) I feel it is right for citizens to deliberate on taking action to control Singapore's population; (b) I feel citizen deliberation make just decisions to control Singapore's population; (c) I feel citizen deliberation can generate population policies that are worth of my support; and (d) I feel citizen deliberation can generate population policies that are likely to work. A 5-point Likert scale was used (1 = strongly disagree, 5 = strongly agree) for all the questions. Results of principal component analysis indicated that the four items loaded on one factor (Eigenvalue = 2.75, factor loadings between .78 and .89) and therefore, the four items were averaged to form one overall measure (Cronbach's $\alpha = .85$, $M = 3.73$, $SD = .66$).

Perceived procedural fairness. We used a four-item measure in the post-deliberation survey to measure perceived procedural fairness. Building on prior studies on perceived speech conditions [4, 38, 39], we asked participants how much they agree about these statements: (a) I had full capacity to freely raise questions about the population policies during the discussions; (b) I feel all of us had equal opportunities to express our opinions on the population issues during the discussions; (c) There was good balance in whose opinion about the population issues is being heard during the discussions; (d) I feel the people I discussed with gave a fair consideration to what I thought about the population issues. The same 5-point Likert scales were used. Results of principal component analysis indicated that the four items loaded on one factor (Eigenvalue = 2.90, factor loadings between .83 and .86) and therefore, the four items were averaged to form one overall measure (Cronbach's $\alpha = .87$, $M = 2.31$, $SD = .99$).

Perceived validity claim. Four validity items were used to measure this concept [4]: (a) I think I understand fellow discussants' points about the population issues; (b) I think fellow discussants' discussions about the population issues were based on accurate facts; (c) I think fellow discussants communicated their arguments on the population issues in an appropriate way; and (d) I think fellow discussants expressed their sincere intentions to communicate with me on the population issues. The same 5-point Likert scales were used. All items loaded on one factor (Eigenvalue=2.71, factor loadings between .79 and .86) and hence they were combined into one variable (Cronbach's $\alpha = .84$, $M = 3.68$, $SD = .64$).

Table 3: Summary of the online deliberation phase in terms of Posts and Users. A user was considered active if they logged in at least once on the platform. L-A stands for Low Heterogeneity-Agree, H-M for High-Mixed and L-D for Low-Disagree.

Moderation	Heterogeneity	Posts	Active Users
High	L-A	453	74
High	H-M	893	90
High	L-D	1324	90
Low	L-A	529	71
Low	H-M	2162	94
Low	L-D	1147	91
Total	-	6508	510

Data Analysis

We ran a 2×3 between-subject study to empirically test the effects of *Moderation* {Low, High} and *Opinion Heterogeneity* { Low-Mainly Agree, Low-Mainly Disagree, High-Mixed} on three dependent variables that measure participants' perceive experience with online deliberation. We then performed a series of two-way ANCOVAs to analyze the effects of moderation and opinion heterogeneity.

Covariates

We ran a series of two-way ANOVAs with moderation and opinion heterogeneity as the two factors and with age, gender, ethnicity, education, housing type, and household income as dependent variables. Only gender, ethnicity, and education showed some significant differences across heterogeneity conditions. We decided to include these three variables as covariates in our main analyses.

Online Activity Data

We recorded a variety of online activities of our participants. On average, each citizen user logged into our platform 11 times (range = 138; $SD = 17$), liked 42 posts (range = 1087; $SD = 125$), and disliked 2 posts (range = 28; $SD = 15$). Each of them posted an average of 12 posts (range = 305; $SD = 36$) and viewed 84 posts (range = 1846; $SD = 218$). They clicked, on average, at least once the node in the deliberation graph (including both single clicks and double clicks; range = 41; $SD = 3$). Citizen users viewed our educational slides for an average of 1.43 times (range = 10; $SD = 1$). They accumulated an average of 590 points (range = 12505; $SD = 1517$) through their platform activities. Table 3 shows a summary of the posting activity by instances and conditions.

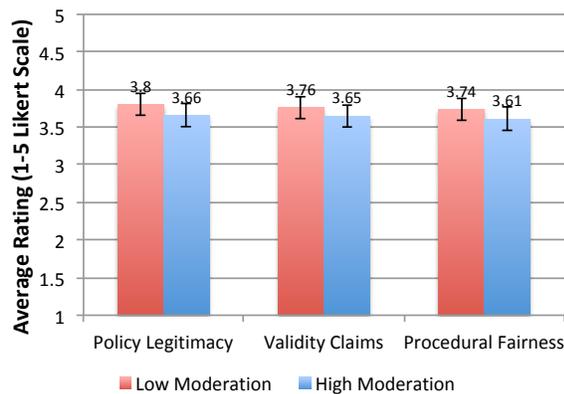


Figure 4: Perceived policy legitimacy, validity claim and procedural fairness across moderation conditions. Error bars represent .95 confidence intervals.

Results

Moderation. For the three dependent variables that gauge participants' perceptions about the online deliberation experience, both moderation and opinion heterogeneity made significant differences. Moderation had significant impacts on perception of validity claim ($F_{1,437} = 4.96, p < .05$, partial $\eta^2 = .011$), perceived legitimacy of deliberation as a way of making population policies ($F_{1,437} = 5.51, p < .05$, partial $\eta^2 = .010$), and a marginally significant effect perception of procedural fairness ($F_{1,437} = 3.08, p < .1$, partial $\eta^2 = .007$). Specifically, a consistent finding emerged: low moderation groups consistently had higher perception ratings than high moderation groups, as illustrated in Figure 4. **H1 is supported.**

Opinion Heterogeneity. Opinion heterogeneity also had significant impacts on perception of procedural fairness ($F_{2,437} = 6.97, p < .01$, partial $\eta^2 = .031$), perception of validity claim ($F_{2,437} = 7.19, p < .01$, partial $\eta^2 = .032$), and perception of policy legitimacy ($F_{2,437} = 7.67, p < .01$, partial $\eta^2 = .034$).

Specifically, a consistent finding emerged with respect to perceptions of policy legitimacy and validity claim: the L-D group members perceived the lowest level of procedural fairness compared (policy legitimacy: $M = \frac{3.54}{5}$ and validity claim: $M = \frac{3.5}{5}$) to the L-A (policy legitimacy: $M = \frac{3.84}{5}, p < .001$; validity claim: $M = \frac{3.79}{5}, p < .001$) and H-M (policy legitimacy: $M = \frac{3.77}{5}, p < .01$; validity claim: $M = \frac{3.7}{5}, p < .01$) group members. No significant differences were seen between the H-M and L-A group on these two measures.

The pattern is slightly different in perceptions of procedural fairness: participants in the L-A group perceived the highest level of procedural fairness ($M = \frac{3.85}{5}$) compared to those in

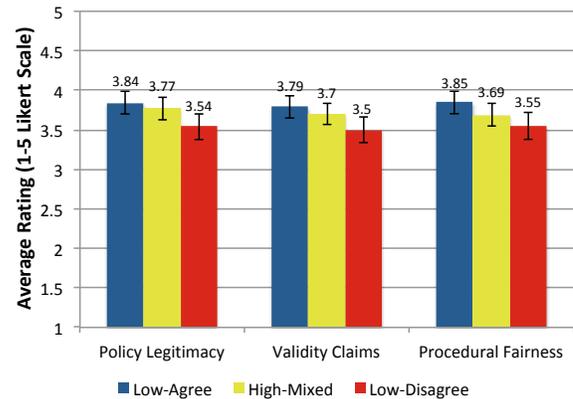


Figure 5: Perceived policy legitimacy, validity claim and procedural fairness across opinion heterogeneity conditions. Error bars represent .95 confidence intervals.

the H-M ($M = \frac{3.69}{5}, p < .05$) and L-D ($M = \frac{3.55}{5}, p < .001$) groups. We did not observe any significant difference between the L-D and H-M groups on this measure. The patterns can be seen in Figure 5. **H2 is supported.**

Interactions. We did not find any interaction effects between moderation and opinion heterogeneity (policy legitimacy: $p = .228$, validity claim: $p = .582$, procedural fairness: $p = .221$).

5 DISCUSSION AND GUIDELINES

We discuss the implications of our results for online deliberation, and provide guidelines to conduct future studies in the field.

Challenges for Online Deliberation

Use Educational Material. Regarding the lack of knowledge problem, we designed educational slides for both using the platform and each sub-issue. Most participants did consult the slides more than once, and a majority of our participants (72%) also found that "the educational slides were helpful". Our participants' interest in educational slides suggest that future online deliberation works could benefit from applying a similar process. It would also be interesting to quantify to which extent participants did learn from the slides, and how this might have impacted their opinions on issues.

Oversample Underrepresented Populations. Our second challenge was representativeness. While we did start with a participant pool representative of Singapore's population in the pre-deliberation survey, we could not keep that representativeness. We did still achieve a better representativeness compared to previous works based on voluntary participation. Achieving representativeness in online deliberation

requires a lot of work in view of the technological barriers to some groups of the population (e.g., the elderly). To compensate for the attrition of less tech-savvy and older participants after the pre-deliberation survey, researchers may want to oversample participants of these underrepresented groups and match their numbers to the population quota after the initial survey.

Keep Moderation to Low Levels

Moderation had an impact on participants' perception of the deliberation experiment. Lower levels of moderation leads to better perception of validity claim, perceived legitimacy of deliberation for policy making, and suggests that it may also affect procedural fairness. Previous work also showed that lower levels of moderation could increase participation, in terms of number of posts [25], but higher levels of moderation lead to self-censorship [36] and exclude underrepresented populations from the deliberation process [33]. Given all these negative effects on discussion and perception of the deliberation process, we would argue to keep moderation to a strict minimum, in which moderators are simply required to prevent participants from resorting to insult or inappropriate behavior.

Low Opinion Heterogeneity is Good?

While statistical analyses suggests an effect of different heterogeneity levels on all three dependent variables, the trend that seems to emerge is that the instances of the platform where participants tended to agree (L-A) with current policies had a better perception of the deliberation process. But does it mean that low opinion heterogeneity is good? We would argue for the opposite: high opinion heterogeneity is not that bad. If citizens already all agree/disagree to a policy, there is no need for deliberation. Our findings show that at least H-M is better than L-D. In fact, H-M showed no differences compared to L-A in two out of three outcomes, policy legitimacy and validity claim. We would thus advocate to create groups of high heterogeneity with all kinds of opinions mixed, which should not only justify the need for deliberation but also reduce any evaluation of deliberation biased by one's agreement or disagreement to the policies.

Implications for Online Discussion in General

Our results are in line with past works on moderation (see above) which generally suggest to keep moderation to the minimum. For instance, Towne and Herbsle [32] pointed out the importance of keeping at least some moderation to avoid trolling and flame wars. In order to keep the basic order of discussions, moderation should also be quick. But having a moderator working all the time on the platform is not always possible. A solution is to use crowdsourced moderation [21], randomly assigning users the task to remove/hide offensive

posts, which may act as a complement or replacement for centralized moderation. However, this raises new challenges, such as small organized groups being able to moderate posts going against their ideas.

Online Deliberation in Singapore

This work is the first of its kind done in Singapore. Because of the unique political context, the high moderation level could also trigger some mistrust or defiance towards the online deliberation process. In a few isolated cases, some participants reacted very negatively to the presence of a moderator, but apart from these cases, we did not observe many negative reactions and/or signs or mistrust.

As shown by the results, the participants gave positive scores to all three of our dependent variables ($\frac{3.3}{5}$ and above). In other questions of the post-deliberation survey, participants also reported "enjoying the discussions" (66.4%) and found that the "discussions were interesting" (69.8%). The overall positive feedback on the experience suggests that there is space to conduct more of these studies in Singapore, giving citizens the opportunities to discuss online and provide input for local policy-makers. As one of the participant reported: *"I enjoyed very much my participation in this online deliberation and hope there will be more such activity/events organized. Perhaps, on other topics that concern Singapore as well, not just on Singapore's population issues"*.

Limitations

In our study, we used monetary incentives, as a way to address the representativeness issue [7]. Not providing incentives for a three-week long experiment would have meant excluding elderly, lower income and less tech-savvy people, as they tend to be the ones who are short of resources to support their participation in online deliberation. Another possible effect of not providing financial incentives would be that people with strong opinions regarding the discussion topic would join voluntarily, while those with mild opinions would not care enough to participate. This possible effect would render our experimental design futile, as we would not be able to attract enough participants with neutral opinions. We acknowledge this limitation of our work, as voluntary participants are different from those who are motivated by extrinsic factors such as financial incentives.

6 CONCLUSION

We introduced our own open-source platform, which we built based on design recommendations from previous works. During the deliberation phase, we evaluated how moderation and opinion heterogeneity may impact users' overall experience of deliberation. We found that higher levels of moderation may degrade the overall experience. We also found that opinion heterogeneity seems to be strongly linked to the

perceptions of the participants, and therefore suggest to use high heterogeneity with mixed opinion in order to avoid any perception bias, positive or negative, based on one's existing opinions on the policies. In terms of the Singapore context, our results suggest that despite of the unique political setting, participants were overall satisfied with their experience, and would be likely to support future deliberation practices like such.

7 ACKNOWLEDGMENTS

This research is supported by Singapore's Ministry of Education (MOE2013-T2-1-105). Views expressed are those of the authors alone and do not necessarily reflect opinions of the sponsoring agencies. The authors want to thank Carol Soon, Leanne Chang, Shengdong Zhao, Tian Yang, Sanju Menon, Yang Wang, Jingjing Liang, Joseph Goh Seow Meng, Jeffrey Effendy, Atima Tharatipyakul and a group of undergraduate and graduate research assistants for their collaboration and assistance that made this research possible.

REFERENCES

- [1] Amy Binns. 2012. DON'T FEED THE TROLLS! Managing troublemakers in magazines' online communities. *Journalism Practice* 6, 4 (2012), 547–562.
- [2] Jennifer Brundidge. 2010. Encountering “difference” in the contemporary public sphere: The contribution of the Internet to the heterogeneity of political discussion networks. *Journal of Communication* 60, 4 (2010), 680–700.
- [3] Robert Cavalier, Miso Kim, and Zachary Sam Zaiss. 2009. Deliberative democracy, online discussion, and project PICOLA (Public informed citizen online assembly). *Online deliberation: Design, research, and practice/Eds. Davies T., Gangadharan SP Stanford, CA: Center for the study of language and information* (2009), 71.
- [4] Leanne Chang and Thomas Jacobson. 2010. Measuring Participation as Communicative Action: A Case Study of Citizen Involvement in and Assessment of a City's Smoking Cessation Policy-Making Process. *Journal of Communication* 60, 4 (2010), 660–679. <http://dx.doi.org/10.1111/j.1460-2466.2010.01508.x>
- [5] Todd Davies, Brendan O'Connor, Jonathan J. Cochran, Alex Angiolillo and Effrat, Andrew Parker, Benjamin Newman, and Aaron Tam. 2009. An online environment for democratic deliberation: motivations, principles, and design. (2009), 275–292.
- [6] Arthur R Edwards. 2002. The moderator as an emerging democratic intermediary: The role of the moderator in Internet discussions about public issues. *Information Polity* 7, 1 (2002), 3–20.
- [7] Stephen Elstub and Peter McLaverty. 2014. *Deliberative democracy: Issues and cases*. Edinburgh University Press.
- [8] Dmitry Epstein and Gilly Leshed. 2016. The magic sauce: Practices of facilitation in online policy deliberation. *Journal of Public Deliberation* 12, 1 (2016), 4.
- [9] Siamak Faridani, Ephrat Bitton, Kimiko Ryokai, and Ken Goldberg. 2010. Opinion Space: A Scalable Tool for Browsing Online Comments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1175–1184. <https://doi.org/10.1145/1753326.1753502>
- [10] James S Fishkin. 2011. *When the people speak: Deliberative democracy and public consultation*. Oxford University Press.
- [11] James S Fishkin and Robert C Luskin. 2005. Experimenting with a democratic ideal: Deliberative polling and public opinion. *Acta politica* 40, 3 (2005), 284–298.
- [12] Amy Gutmann and Dennis F Thompson. 1998. *Democracy and disagreement*. Harvard University Press.
- [13] Lisa K Hanasono and Fan Yang. 2016. Computer-mediated coping: Exploring the quality of supportive communication in an online discussion forum for individuals who are coping with racial discrimination. *Communication Quarterly* 64, 4 (2016), 369–389.
- [14] Yu-Tang Hsiao, Shu-Yang Lin, Audrey Tang, Darshana Narayanan, and Claudina Sarahe. 2018. vTaiwan: An Empirical Study of Open Consultation Process in Taiwan. (2018).
- [15] Davy Janssen and Raphaël Kies. 2005. Online forums and deliberative democracy. *Acta politica* 40, 3 (2005), 317–335.
- [16] Juho Kim, Eun-Young Ko, Jonghyuk Jung, Chang Won Lee, Nam Wook Kim, and Jihee Kim. 2015. Factful: Engaging Taxpayers in the Public Discussion of a Government Budget. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 2843–2852. <https://doi.org/10.1145/2702123.2702352>
- [17] Nam Wook Kim, Jonghyuk Jung, Eun-Young Ko, Songyi Han, Chang Won Lee, Juho Kim, and Jihee Kim. 2016. BudgetMap: Engaging Taxpayers in the Issue-Driven Classification of a Government Budget. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. ACM, New York, NY, USA, 1028–1039. <https://doi.org/10.1145/2818048.2820004>
- [18] M. Klein. 2011. The MIT deliberatorium: Enabling large-scale deliberation about complex systemic problems. In *2011 International Conference on Collaboration Technologies and Systems (CTS)*. 161–161. <https://doi.org/10.1109/CTS.2011.5928678>
- [19] Travis Kriplean, Jonathan T. Morgan, Deen Freelon, Alan Borning, and Lance Bennett. 2011. ConsiderIt: Improving Structured Public Deliberation. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems (CHI EA '11)*. ACM, New York, NY, USA, 1831–1836. <https://doi.org/10.1145/1979742.1979869>
- [20] Travis Kriplean, Michael Toomim, Jonathan Morgan, Alan Borning, and Andrew Ko. 2012. Is This What You Meant?: Promoting Listening on the Web with Reflect. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1559–1568. <https://doi.org/10.1145/2207676.2208621>
- [21] Cliff Lampe, Paul Zube, Jusil Lee, Chul Hyun Park, and Erik Johnston. 2014. Crowdsourcing civility: A natural experiment examining the effects of distributed moderation in online forums. *Government Information Quarterly* 31, 2 (2014), 317–326.
- [22] David M Lazer, Anand E Sokhey, Michael A Neblo, Kevin M Esterling, and Ryan Kennedy. 2015. Expanding the conversation: Multiplier effects from a deliberative field experiment. *Political Communication* 32, 4 (2015), 552–573.
- [23] David López García. 2017. Mediation Styles and Participants' Perception of Success in Consultative Councils: The case of Guadalajara, Mexico. *Journal of Public Deliberation* 13, 2 (2017), 10.
- [24] Ann Macintosh. 2004. Characterizing e-participation in policy-making. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*. IEEE, 10–pp.
- [25] Hans K Meyer and Michael Clay Carey. 2014. In Moderation: Examining how journalists' attitudes toward online comments affect the creation of community. *Journalism Practice* 8, 2 (2014), 213–228.
- [26] Peter Muhlberger. 2005. The Virtual Agora Project: A research design for studying democratic deliberation. *Journal of Public Deliberation* 1, 1 (2005), 5.
- [27] Diana C Mutz. 2002. The consequences of cross-cutting networks for political participation. *American Journal of Political Science* (2002), 838–855.

- [28] Vincent Price, Joseph N Cappella, and Lilach Nir. 2002. Does disagreement contribute to more deliberative opinion? *Political communication* 19, 1 (2002), 95–112.
- [29] Shawn Rosenberg. 2005. The Empirical Study of Deliberative Democracy: Setting a Research Agenda. *Acta Politica* 40, 2 (01 Jul 2005), 212–224. <https://doi.org/10.1057/palgrave.ap.5500105>
- [30] Bryan Semaan, Heather Faucett, Scott P. Robertson, Misa Maruyama, and Sara Douglas. 2015. Designing Political Deliberation Environments to Support Interactions in the Public Sphere. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3167–3176. <https://doi.org/10.1145/2702123.2702403>
- [31] Jennifer Stromer-Galley and Peter Muhlberger. 2009. Agreement and disagreement in group deliberation: Effects on deliberation satisfaction, future engagement, and decision legitimacy. *Political communication* 26, 2 (2009), 173–192.
- [32] W Ben Towne and James D Herbsleb. 2012. Design considerations for online deliberation systems. *Journal of Information Technology & Politics* 9, 1 (2012), 97–115.
- [33] Matthias Trénel. 2009. Facilitation and inclusive deliberation. *Online deliberation: Design, research, and practice* (2009), 253–257.
- [34] Kevin Wise, Brian Hamman, and Kjerstin Thorson. 2006. Moderation, response rate, and message interactivity: Features of online communities and their effects on intent to participate. *Journal of Computer-Mediated Communication* 12, 1 (2006), 24–41.
- [35] Magdalena E Wojcieszak and Vincent Price. 2012. Perceived versus actual disagreement: Which influences deliberative experiences? *Journal of communication* 62, 3 (2012), 418–436.
- [36] Scott Wright. 2009. The role of the moderator: Problems and possibilities for government-run online discussion forums. *Online deliberation: Design, research, and practice* (2009), 233–242.
- [37] Weiyu Zhang. 2012. The effects of political news use, political discussion and authoritarian orientation on political participation: Evidences from Singapore and Taiwan. *Asian Journal of Communication* 22, 5 (2012), 474–492.
- [38] Weiyu Zhang. 2015. Perceived procedural fairness in deliberation: Predictors and effects. *Communication Research* 42, 3 (2015), 345–364.
- [39] Weiyu Zhang and Leanne Chang. 2014. Perceived speech conditions and disagreement of everyday talk: A proceduralist perspective of citizen deliberation. *Communication Theory* 24, 2 (2014), 124–145.