Solving the Problem of "Many-to-One" Communicative Action: Re-designing the Webinar using Technology-Enabled Deliberative Theory

Kevin M. Esterling*
University of California Riverside

Ben Treves
University of California Riverside

Kelton Adey
University of California Riverside

Euchan Jang
The Ohio State University

March 12, 2025

*Corresondence to: Kevin M Esterling, University of California Riverside, 900 University Ave., Riverside, CA 92521 USA, kevin.esterling@ucr.edu

Abstract

The theory of communicative action is well-developed for discursive situations where only one person can speak at a time, but the theory has not been extended to the situation where communicative action should arise in a *many-to-one speaking situation*. A many-to-one situation occurs whenever a speaker wishes to understand the diverse perspectives of an audience in real time during a meeting. A solution to the problem of many-to-one communicative action requires the speaker and the audience to express their perspectives and understand each other simultaneously, even when the audience scales to a large number of participants. We develop a normative framework to re-imagine the webinar, showing that innovative webinar technology motivated by deliberative theory enables us to envision the counterfactual of communicative action in the many-to-one speaking situation, and to develop design solutions to enable deliberation in this context. We describe a case study of an innovative webinar platform, called Prytaneum, that was developed using this framework. Our normative framework for many-to-one communicative action provides design best practices for the webinar industry and enables the development of new communication theory relevant to this setting.

Keywords: Deliberation, Communicative Action, Large Language Models, Webinar Technology, Algorithmic Curation

1 Introduction

The concept of *communicative action* (Habermas, 1984) underwrites much of the modern normative theory of deliberation and of deliberative democracy (Gutmann and Thompson, 1996). Deliberation occurs under an ideal speaking situation where discussion participants freely exchange reasons for and against decision alternatives.[1] Communicative action within a discussion group materializes when the mutual scrutiny of such reasons creates an interpersonal understanding of the merits of the alternatives (Chambers, 2018). An interpersonal understanding over decision alternatives that results from deliberation can motivate collective action, enhance the legitimacy of decisions, or both (Cohen, 1989).

The theory of communicative action is well-developed for discursive situations where only one person can speak at a time, but the theory has not been extended to the situation where communicative action should arise in a *many-to-one speaking situation*. These many-to-one discursive situations are ubiquitous, and occur any time there is a speaker who wants to understand the perspectives among members of an audience during a synchronous meeting.[2] Such many-to-one situations occur, for example, in public town halls, scientific presentations, instruction in the classroom, corporate all-hands meetings, and so on. As we explain below, the core problems to solve, both in theory and in practice, include, how is it possible for an arbitrary number of audience members jointly to engage in deliberative reason-giving in real time during a meeting? How can the speaker accumulate, understand, and respond to the diverse perspectives and knowledge expressed among the audience members? And, how to best induce audience members to be active, accountable participants rather than passive listeners (Lerner and Tetlock, 1999)?

Communication theory has not yet addressed these problems because in-person discussion presents many practical constraints that disable effective or coherent communication in the many-to-one situation. The most obvious constraint is that it is not possible for all

[1] An ideal speaking situation in general is a counterfactual unrealized in practice.

[2] For simplicity and to limit our ambitions, in this paper we are not considering many-to-many communication, such as a committee hearing, where there are many speakers and many audience members.

members of the audience to speak at once and still be understood. Further, as we explain below, existing webinar technology worsens these constraints. The typical webinar platform is purposefully designed to enable only the speaker to communicate to the audience, and so by design makes it difficult for audience members to have any voice at all.[3]

Within these practical and technological constraints, it is not possible even to conceive of the problem of many-to-one discourse, much less to develop theory, design, and practice to enable communicative action in this situation. We show that innovative webinar technology motivated by democratic theory enables us to envision the counterfactual of communicative action in the many-to-one speaking situation. Here we explain the communicative action ideal in the many-to-one setting, then we describe the constraints for in-person communication in this setting and how those constraints are worsened with existing webinar technology. Then we develop a normative framework to motivate a solution to the problem of many-to-one communicative action. We then present a case study of a new online webinar platform, called Prytaneum, with a design that is fully motivated by this normative framework. Finally, we show how Prytaneum's innovative technology enables the new conception and theory of many-to-one communicative action that could not be theorized or addressed in the absence of the technology.

2 The Ideal of Many-to-One Communicative Action

The goal of many-to-one communicative action is to enable and promote collaborative, *deliberative perspective sharing* within a synchronous, speaker-led presentation and discussion – even when the audience scales to a large number of participants.

In its ideal form, deliberation prioritizes reason-based collaboration for collective actions and decisions (Cooper and Mackie, 1983; Elliot and Devine, 1994; Hinsz et al., 1997;

[3]As we discuss below, asynchronous deliberation platforms enable the scaling of deliberation among many participants, and each can submit comments simultaneously. Procedurally, however, asynchronous communication resides in the traditional design space where each person speaks in turn, and hence does not address or solve the many-to-one communication situation.

Matz and Wood, 2005). Deliberation envisions an ideal of reasoned, uncoerced, diverse, inclusive, and transparent argumentation for and against alternatives that must be decided collectively (Gutmann and Thompson, 2004; Habermas, 1996; Neblo, 2015). As Jürgen Habermas explains in his theory of communicative action, argumentation within diverse groups induces individuals to defend the *validity* of the reasons and justifications they offer for and against decision alternatives and to criticize those of others (Habermas, 1984). As participants examine each others' reasoning, they expose unsupported or invalid understandings (Chambers, 2018; Mercier and Landemore, 2012), enable participants to recognize and overcome framing effects (Druckman and Nelson, 2003), and allow the construction of superior alternatives (Esterling, 2004). Those claims that survive scrutiny enable interpersonal understanding regarding the relative merits of alternative decisions (Hinsz et al., 1997; Matz and Wood, 2005; Niemeyer et al., 2024).

Deliberation depends on the presence of ideal circumstances, which includes diversity, inclusion, equality, mutual respect and the absence of coercion (Gutmann and Thompson, 1996). In its ideal form, deliberation enhances social cognition and collaboration in two reinforcing ways. First, when discussion participants share diverse perspectives and accurate (perhaps falsifiable) information relevant to creating a collective understanding, they can enhance the group's knowledge and *epistemic capacity regarding the alternatives* (Bohman, 1998; Freeman, 2000; Gutmann and Thompson, 2004). Second, when audience members share reasons for the merits of alternative actions, the reasons they provide can underwrite the *coherence, legitimacy, and ethical foundations* of those actions (Cohen, 1989; Plemmons et al., 2020). Communicative action arises when an interpersonal understanding of the alternatives, informed by improvements in epistemic capacity and legitimacy, creates the basis for collective decision and action (Niemeyer et al., 2024).

The traditional design for a speaker-led meeting can have underlying *organizational pathologies* that can disable the normative aspirations of communicative action even before the discussion begins. Deliberation cannot arise when the meeting excludes stakeholder

perspectives (Karpowitz et al., 2012), when participants lack mutual respect (Mendelberg and Oleske, 2000), and in the presence of coercion (Sanders, 1997). These pathologies are related to the organization of the meeting itself and can manifest irrespective of whether the meeting is in-person or online, or whether the communication is sequential or many-to-one. Solutions to these organizational pathologies must be in place before one can consider new methods for promoting communicative action of any kind.

There are at least four such pathologies. First, participants who self-select into attending a meeting generally are unrepresentative of the larger community that might have a stake in the meeting's outcomes. For example, the participants at public meetings often are homogeneously older, male, white and affluent (Einstein et al., 2019; Marble and Nall, 2020; Trounstine, 2020), and so tend to over-represent the powerful and exclude voices of underserved communities (Mansbridge, 1980; Sanders, 1997). In practice, online public meetings tend to reproduce the inequities of in-person meetings (Einstein et al., 2023). Second, meetings are often organized in an "open mic night" format, where participants can offer remarks that are off topic or misinformed (Mendelberg and Oleske, 2000) and where no one topic is explored in depth. Third, participants might not have access to factually accurate background reading material to adequately prepare for the discussion. And fourth, the standard format requires the audience to be passive listeners, which lessens the depth of cognitive processing compared to settings where the audience members are expected to share and defend their perspectives (Lerner and Tetlock, 1999). As a result the considerations that are distributed across audience members and that potentially can be shared in the meeting are less elaborate and informed compared to the actual capacity of the audience (Esterling et al., 2011).

Recent work in political science has identified contextual features that can overcome these organizational pathologies and create the circumstances conducive to deliberative democracy in public meetings, both online and in-person (Dryzek et al., 2019). This line of research identifies design components for hosting deliberative interactions that can ad-

dress the organizational pathologies that we list above, such as recruiting a representative cross section of the relevant community to ensure diversity of perspectives in the audience; providing factually-accurate reading material to ensure all participants can become informed and empowered to contribute to the discussion; and having a neutral moderator or facilitator (Neblo et al., 2018). None of this previous work tackles the problem of communicative action in the many-to-one situation. But understanding this previous work is essential in that, without solving the underlying organizational problems, deliberation of any kind cannot assert its normative justification. This previous work identifies necessary but not sufficient conditions for many-to-one communicative action.

The deliberative communicative ideal has been developed in the situation where only one person can speak at a time – that is, where each individual expresses their arguments sequentially. When applying the deliberative ideal to a speaker-led meeting, whether in-person or online, not only should the audience understand the speaker's perspective, but in addition the speaker also should understand the diverse perspectives of the audience regarding each of the statements and claims that the speaker makes, synchronously and in real time. The synchronous exchange of reasons among an arbitrarily-large number of participants is the core problem of many-to-one communicative action. Solving the problem of many-to-one communication is essential to enabling deliberation in the context of a speaker-led meeting. As James Fishkin notes (2019), in the absence of two-way communication there is a strong risk of the speaker holding coercive power to manipulate audience opinions and beliefs. The goal of many-to-one communication is to decenter the speaker and to empower the audience to speak and be heard.

3 Constraints within In-Person and Online Meetings

While speaker presentations can have immense value, the design of the traditional in-person meeting places fundamental constraints on the extent to which the interaction can meet the ideal of many-to-one communicative action.

Fundamentally, in any meeting there will be many diverse and relevant considerations that are distributed across audience members, but it is not possible for audience members to express those thoughts jointly in real time during the typical meeting. The audience members are generally reluctant to interrupt the speaker to share these considerations, and even if they do, only one audience member can speak at a time. The audience members might have an opportunity individually to share their thoughts at the conclusion of the presentation, but not at the moment when each comment is most relevant, and there is no assurance the sequence of comments will follow a coherent pattern or that the comments made represent the full diversity of perspectives in the audience. The speaker meanwhile has to multitask between understanding and answering questions and collecting and recording the ideas that the audience shares.

These limitations for audience participation and interaction are worsened when the meeting is hosted online using a standard webinar platform. Webinar platforms are designed for a speaker to communicate to a large audience. But webinar platforms are *purposefully designed to disempower the audience* and prevent the audience from having a voice during a meeting. A standard webinar lacks features that enable the audience members to share their thoughts effectively in real time during a meeting, often confining engagement to a meager chat stream that posts comments in a flat and rapidly-moving list. In addition, in online settings where the audience scales to a large number of people, the facilitator of a webinar can become overwhelmed with content posted to the chat stream and cannot be effective. Finally, webinar platforms often have a feature where the audience can respond to polls, but this captures only closed-ended responses and in addition, unless the audience is fully representative of the relevant stakeholder community, there is no reason to believe that the results of the poll have any meaning or legitimacy.

In sum, the standard webinar design for online presentations does nothing to overcome the underlying limitations of the traditional speaker-led talk, and only layers additional limitations. These constraints that disable the audience's voice are built into the webinar

technology itself. These constraints impede many-to-one communication but as we show next, it is possible to reimagine the webinar design to solve these constraints.

4 Designing Many-to-One Communicative Action

The degree to which an individual engages in a group cognitive task such as deliberation is not necessarily a fixed attribute of the individual. Instead, deliberation may be induced by the institutional and social context in which interactions take place (Hinsz et al., 1997; MacKuen et al., 2010). For example, in small group settings where the discussion participants are diverse and represent diverse perspectives, human facilitators can ensure that all different perspectives are given voice and that comments that are on-point, informed and constructive are prioritized, and can discourage misinformed and off-topic comments (Esterling et al., 2021). Within such a discussion, participants not only hear new perspectives, but they also must engage in cognitive effort in order to make a positive contribution that will be valued by the group, which in turn leads participants to more deeply encode information and have a more integratively complex understanding of the discussion topic (Lerner and Tetlock, 1999).

Creating a context to induce deliberation is manageable for a facilitator in a small group setting, but is more challenging at the scale of modern democracy where much of public discourse occurs through computer-mediated communication. When a town hall scales to hundreds, thousands or even millions of people, the sheer number of comments submitted can overwhelm a human facilitator. Understanding the diverse perspective of an audience at scale is the core problem of many-to-one communicative action.

Any communication platform that intends to advance deliberative interactions at the scale of modern democracy must rely on algorithms to curate content, just as the "big tech" social media platforms use algorithms to curate user content to advance the goals and purposes of their platforms (Lazer, 2015). But of course, a deliberation-reinforcing algorithm must be very different from the current social media algorithms that, as we

Table 1: Algorithmic Scaling of Deliberation

	Harms Deliberation	**Promotes Deliberation**
Small Group	**Cell A:** Homogeneous, Unfacilitated Discussion	**Cell C:** Diverse, Facilitated Discussion
At Scale	**Cell B:** "Big Tech" Social Media Algorithms	**Cell D:** Many-to-One Algorithmic Curation

describe below, undermine deliberation (Barrett et al., 2021; Rhodes, 2022; Sunstein, 2017). To date, no commercial webinar platform has designed an algorithmic environment that enables facilitators to incentivize and encourage deliberative discourse at scale.

In this section, we derive the design principles and technical requirements for enabling many-to-one communicative action in an online setting. Table 1 provides an overview and road map. The columns indicate circumstances that harm or promote deliberation, and the rows indicate interactions that occur in small groups or at scale. We first expand cell A to discuss how poorly-designed small group processes can lead to polarization, extremism and vulnerability to misinformation. We then expand cell B to show how "big tech" social media algorithms replicate these small group processes to undermine deliberative engagement at massive scale. Then we move to cell C to discuss the extensive literature on empirical deliberation, showing the cognitive benefits of deliberation and designs that have been implemented to overcome unproductive small group processes to induce deliberation in small group settings. Finally, in cell D we identify ways to scale up the lessons from the small group deliberation literature and describe the design principles and technical requirements of many-to-one communicative action.

Cell A: Small group polarization and extremism in homogeneous, unfacilitated discussion. A vast literature in small group research (Isenberg, 1986) demonstrates that when a group is composed of non-diverse or homogeneous participants, discussion within the group often leads to extremism in the direction of the group members' initial predis-

positions, partly but not entirely because group members are placed in an echo chamber and not exposed to countervailing arguments (Sunstein, 2002). Small groups that lack facilitation reduce incentives for participants to stay on-topic or to make well-informed contributions to the discussion. In these circumstances, participants often find themselves entrenching positions, resenting those who challenge them, increasing in partisan or ideological extremism, falling prey to misinformation, or withdrawing from the conversation altogether (Therriault et al., 2011).

Cell B. Social media curation algorithms undermine deliberation and democracy. "Big tech" social media algorithms create a computing environment that replicates and scales the harmful dynamics that occur in homogeneous, unfacilitated small group discussions.[4] Social media platforms such as X/Twitter and Facebook currently serve as the de facto town squares of modern democracy. Many observers share a concern, however, that democracy today is at such a scale that it does not and perhaps cannot foster deliberation in contemporary public life (Hills, 2019).

The algorithms driving these "big tech" social media platforms are designed to increase user engagement and persistence by exposing users to content they find agreeable, which in turn promotes filter bubbles, echo chambers, and the spread of misinformation (Barrett et al., 2021; Sunstein, 2017). At a fundamental level, social media algorithms exploit people's innate desire to seek out information and perspectives that confirm their own beliefs and predispositions (Lazer, 2015). These algorithms merely reinforce users' political predispositions, much like in homogeneous small groups, and this in turn leads to polarization, extremism, hostility to alternative perspectives, and hardened attitudes (Rhodes, 2022; Settle, 2018; Sunstein, 2002, 2009; Vosoughi et al., 2018). Because they organize communication on a massive scale, these platforms undermine society's ability

[4]Note that social media platforms are asynchronous, and as we discuss below asynchronous platforms by design do not offer a solution to the problem of many-to-one communicative action. That said, social media platforms serve as a useful example of how small group communication can scale through algorithmic content curation.

to understand and learn from different perspectives, and have the unfortunate effect of reducing democracy's ability to promote collaboration to solve important and complex public problems (Mann and Ornstein, 2016; Ognyanova et al., 2020; Prior, 2013). These algorithms serve as a foil as the kind of computing environment we wish to avoid.

Cell C. Designs to induce deliberation in diverse, facilitated small groups. While the prospects for deliberative democracy today seem bleak, at least at a scale that would matter for modern democracy, this is not a necessary state of affairs. The American political tradition has long been rooted in the idea that well-crafted institutional design can promote deliberation in public life (Bessette, 1994), and in that tradition political scientists examine designs of institutions that can overcome these harmful group dynamics (Karpowitz and Raphael, 2014) and that can create deliberative and collaborative environments for decentralized social problem solving (Noveck, 2105). We postulate that appropriately designed online platforms can induce deliberative, many-to-one communicative action, at the scale of modern communication, compared to what we currently observe in social media platforms, and even compared to traditional in-person or webinar-based meetings. In this subsection, we review the design elements that have proven effective in inducing deliberation in small group settings, and in the following subsection we will draw on these elements to propose methods to induce synchronous deliberation at scale.

An extensive literature on empirical deliberation identifies methods to change harmful small group dynamics into constructive, collaborative, deliberative dynamics (Dryzek et al., 2019). Certain personality traits might predispose individuals to deliberate, including openness to new experiences (Gerber et al., 2011) and a strong need for cognition (Cacioppo et al., 1984). But careful institutional design has the potential to induce all individuals, irrespective of their personality type, to increase their propensity to deliberate and to engage in constructive group reasoning (Hinsz et al., 1997; MacKuen et al., 2010).

At the most general design level, small group deliberative event organizers typically

articulate norms and admonitions to listen and participate constructively, to explore differences (Morrell, 1999; Vorauer et al., 2009), and to frame rigid disagreement as inappropriate (Conover et al., 2002; Vorauer and Sasaki, 2011). To create the contextual conditions conducive to deliberation, organizers typically will make efforts to recruit diverse participants, focus the discussion on a single substantive issue, provide ideologically-balanced factual information regarding the issue, and provide a facilitator to manage the discussion under a set of announced ground rules to govern debate (Neblo et al., 2018).

Each of these basic design features to promote deliberation in small groups is informed by cognitive theory and should lead to participants' deeper cognitive processing and integrative complexity. Consider first the role of diversity and the presence or absence of diverse perspectives in facilitated group discussion. Echo chambers lead to small group processes that rely on "System 1" cognition in which arguments that confirm one's prior beliefs are taken for granted as true (Kahneman, 2011). Removing participants from echo chambers can expose participants to unexpected arguments and so induce a deeper "System 2" cognitive processing (Cooper and Mackie, 1983; Elliot and Devine, 1994) that recognizes the limits of one's own understanding and the possible merits of alternatives.

However, as Bail et al. (2018) demonstrate, exposure to diverse views is necessary but not sufficient for reducing polarization and extremism. For diversity to induce constructive collaboration, participants must attempt to offer *reasons* to support their positions that others can recognize as valid. Doing so requires a certain degree of cognitive integrative complexity, which is a measure of the ability of the individual to understand and communicate different aspects of a complex issue (Suedfeld et al., 1992). Participants with enhanced cognitive integrative complexity are better able to understand the often-competing aspects of policies, how those different aspects are connected, and the possible common ground for policy change. This leads participants to a more complete, empathetic and nuanced understanding of the discussion, and in particular to understand the perspectives and reasons among those with whom they disagree (Mutz, 2006).

As Tetlock et al. (1989) and Lerner and Tetlock (1999) show, individuals' cognitive integrative complexity is a function of accountability within institutional settings. When participants are held accountable for the merits of their contributions, they are motivated to learn and become well-informed about the issue, which in turn deepens their cognitive processing and learning. In traditional small group deliberation, facilitators create accountability by providing feedback to participants regarding the relevance of arguments and reasons they provide, or whether the contribution was off topic, disruptive or misinformed, and so make participants accountable for the merits of their contributions. Prioritizing contributions by their relevance reduces extremism and misinformation to the extent extreme or misinformed statements are excluded as off topic or disruptive. And effective facilitators will exclude toxic or disruptive contributions altogether.

There is an extensive empirical literature that demonstrates the capacity of individuals within small groups, both in-person and online, to deliberate given carefully designed institutions that promote diversity and accountability. A vast number of small group deliberative experiments have been evaluated. Among the many findings: Neblo et al. (2018) show that, under deliberative circumstances, diverse participants learn complex policy information and engage constructively with public officials. Broockman and Kalla (2016) find that conversations between voters and canvassers can be designed to evoke cognitively effortful or "System 2" (Kahneman, 2011) reasoning even on controversial topics. Esterling et al. (2021) and Grönlund et al. (2015) show that deliberation can reduce polarization and ideological extremism. Druckman (2004) shows that deliberation can counteract elite framing effects.

Cell D. Design requirements to scale up deliberation in the many-to-one setting. We draw on these insights from the empirical small-group deliberation literature to develop the normative framework for a webinar design that can induce deliberation at the scale of modern democracy. This framework will inform the computer-mediated

circumstances that can promote deliberation and informed reason-giving in an online, synchronous meeting, solving the problem of many-to-one communicative action even when the online event scales to a large number of participants. The webinar design should preserve deliberative interactions in a way that can become part of the everyday workflow in meetings in many contexts, including public town halls, science talks, classroom instruction, and employee all hands meetings.

Our framework consists of two levels. At the first level is the basic design task of enabling webinar participants to have robust opportunities to contribute content in real time during a meeting, even when the meeting scales to a large number of participants. As we mention above, the traditional webinar design does not have features that give participants the ability to express their views effectively, often confining participants to a meager chat stream or to closed-ended polling responses. Indeed, the traditional webinar's "solution" to the problem of many-to-one communication is to disempower the audience altogether. A re-imagined webinar should provide user interface features that enable participants to robustly share their thoughts in real time during the meeting, and to ask questions and make comments that are relevant to the presentation. This is just a straightforward task of improved user interface design.

At the second level of our normative framework is the more challenging design task of creating an environment that encourages participants to engage in mutual reasoning over decision alternatives constructively and on-topic, and that ensures the diverse views in the community are heard and understood, even when – and often especially when – those views are expressed among members of an under-represented stakeholder perspective. The facilitator performs these functions in a small group meeting. But when the event scales to a large number of participants, and when all participants have robust opportunities to share their views in real time during a meeting, a human facilitator can quickly become overwhelmed with the sheer volume of content (Hills, 2019). Tackling this second task requires developing carefully designed algorithmic content curation, relying on methods

for topic classification and text summarization to enable meeting facilitation at scale.

To create incentives for the discussion to remain informed, constructive and on-topic, even as the number of participants becomes very large, we propose developing methods to classify content by relevance. Here we consider meetings that are centered on a single issue, and that issue can be composed of a limited set of topics. The meeting organizer can preselect the relevant topics prior to the meeting, or they can allow topics to emerge endogenously during the meeting, or both. The webinar can use a classifier that prioritizes contributions that are on topic and so constructively add to the discussion, and that excludes contributions that are misinformed or disruptively off topic. Such classification can ensure that the discussion remains on-topic and well-informed, and that discussions are protected from malicious actors who seek to disrupt the meeting. Topic classification of content also enables the facilitator to ensure that the meeting is coherent in that each topic can be explored in depth before moving on to the next topic, a workflow that is preferred to a design that allows the meeting discussion to jump abruptly across topics.

To ensure that the full diversity of perspectives within a community of stakeholders are expressed and heard within an online town hall meeting, we propose developing methods to summarize content by extracting diverse reasoning for and against different decision alternatives. A summarization tool would allow a facilitator to curate content submitted by participants within a large-scale online town hall event to ensure that diverse reasoning is captured, and that that participants are exposed to the full and diverse range of perspectives present in their community. The summarization tool would use the topic classifier to filter out irrelevant or offensive content to ensure such content is excluded.

The priority here is to ensure that all perspectives are presented, irrespective of the relative popularity or unpopularity of that perspective. There are two key reasons why we encourage disregarding the popularity of a perspective when reporting summaries of the reasoning among audience members. First, as we discuss above, it is virtually never the case that the participants who show up at a meeting are representative of the full

set of stakeholders in the relevant community. Thus, the *distribution* of preferences and opinions among the participants at a meeting is generally not meaningful, and indeed simply reporting the views of the participants who happen to dominate a meeting might be objectionable. To ensure the full set of perspectives from a community are summarized in a meeting only requires that *at least one person holding that perspective* show up at the meeting, which is substantially easier than trying to ensure the full audience is representative of the community. And second, the merits of an argument are not necessarily correlated with the argument's popularity, and more importantly, we want to ensure that perspectives that might be underrepresented in the community are heard.

Advances in artificial intelligence (AI) and large language models (LLM) enables us to develop these classifiers and summarization tools so that they can curate content in real time during a meeting. To define the issue and topics, the LLMs can be fine tuned using factually accurate, politically-neutral reading material that can provide the substantive and epistemic foundations for the classification. To ensure that diverse perspectives are considered in the summarization tool, we propose using open ended polling questions in the meeting that require participants not only to choose the option they prefer in response to a prompt, but that also require participants to explain the reasoning to support their preferred option (or conversely, to explain their reasoning against options they dislike). In this way, the summarization tool can report not only the distribution of preferences among attendees, which is often not meaningful, but it can also report the reasoning that support and oppose each option, irrespective of the popularity of the option. Anchoring the classification and summarization tools to the epistemically correct reading material helps to address concerns about accuracy and fairness in LLMs.

Having this algorithmic curation in place can create new incentives for participants to engage the meeting constructively. When participants know that off-topic content will be excluded, and if they wish to make an impact, participants will need to strive to remain on-topic, avoid toxic content, and avoid trafficking in misinformation and conspiracy theories

that are incompatible with the epistemically correct reading material. And this curation ensures that all perspectives in the community are heard, including the voices that are normally excluded, not just the perspectives of those who happen to show up to dominate the meeting. In this way, the algorithmic curation will perform the same tasks as a human facilitator in a small group setting, except that the algorithmically enhanced curation will scale to an arbitrary number of audience members.

This normative framework creates a logic to support many-to-one communicative action. In this setting, participants hear perspectives from outside of their echo chambers, and each participant must expend cognitive effort to ensure their contributions have an impact on the discussion. Exposing webinar participants to diverse perspectives in this computing environment should induce participants to engage in reflective "System 2" cognition, and the incentives for participants to make sure their contributions are relevant, informed and on-topic should lead them to process information more deeply and have a more integratively complex understanding of the issue under discussion. As a result, such a webinar design would create an environment that promotes deliberative collaboration among audience members, even when there are many participants, resulting in epistemically better solutions for complex collective action problems. The webinar design cannot guarantee that discourse within a town hall will be constructive and deliberative, but it creates an environment that facilitates and incentivizes constructive input and collaborative learning, and so increases the prospects that deliberation will occur.

5 A Case Study of Prytaneum

In this section, we provide a case study of Prytaneum,[5] which is a novel webinar platform that is designed using our normative framework for many-to-one communicative

[5] https://prytaneum.io. Prytaneum is developed by the Laboratory for Technology, Communication and Democracy (TeCD-Lab) at the University of California Riverside, in collaboration with the Institute for Democratic Engagement and Accountability (IDEA) at the Ohio State University. The project has received generous support from the Democracy Fund, the Institute of Education Sciences, the UCR OASIS program, UCR Information and Technology Services, and IDEA at OSU.

action. Prytaneum has LLM-powered curation algorithms designed to enhance constructive collaboration and deliberative, social learning. The content curation algorithms assist facilitators who wish to preserve deliberative interactions even when an event scales to a large number of participants. Its artificial intelligence tools not only filter out toxicity, but also proactively assist event moderation in real time by curating the participant input, for example by grouping similar questions by topic and by clustering similar questions and comments together. Discussion facilitators can filter questions and comments by topic, ensuring only on-topic content is prioritized, and ensuring that all diverse perspectives are heard and each topic is fully considered before moving on to the next.

Most importantly, Prytaneum also has a "prompt-response" polling feature to prompt participants to provide open-ended responses or to choose among options on survey question prompts, where the response screen includes a text box soliciting their reasoning for their choices to ensure every participant is actively engaged and their perspectives are heard and understood. Prytaneum provides LLM-powered abstractive summaries of these responses to share the audience member's reasoning to both speaker and audience in real time during an event. The summarization is designed to ensure that underrepresented voices in an audience are heard, and to promote balanced and constructive user engagement. The prompt-response feature enables a redesign the traditional talk format, to make the talk more interactive for the audience and more informative for the speaker.

Prytaneum event workflow from the perspective of an audience member. Prytaneum is designed to make a synchronous event collaborative, and to give audience members robust opportunities to share their perspectives substantively and in real time during a meeting. The user engagement has a social media-like experience. Audience members can submit unlimited posts in the form of questions and comments addressed to the speaker at any time during the event. Speakers and facilitators can respond directly using the question and comment thread, and as we show below, facilitators have tools to

elevate relevant audience posts to the discussion. Audience members also can respond to open-ended survey prompts for their views and the reasoning that supports their views. This means that in events with a large number of participants, dozens or hundreds or even thousands of audience members can effectively contribute. Prytaneum's LLM-powered curation algorithms are designed around capturing and organizing this massive amount of information in real time, while ensuring that participation remains high-quality, balanced, and relevant to as many participants as possible.

Prytaneum event workflow from the perspective of a facilitator. Prytaneum provides the meeting host and facilitation team a fully collaborative user interface to curate content that is submitted by the audience, and to understand and make sense of that content in real time even when the event scales to a large number of participants. As we describe below, the curation algorithm tags content in real time by topics that are relevant to the discussion, enabling the facilitation team to quickly filter and organize audience submissions by topic. The host facilitator can choose to elevate questions and comments for discussion in the meeting, and with the topic filtering, can ensure that each topic is discussed thoroughly before moving onto the next topic. The facilitator also can deploy prompt-response real time polling. Finally, the facilitator can choose a toxicity threshold and can exclude abusive participants.

Curation algorithm. Now we describe the core elements of the curation algorithms in more detail. We highlight machine learning-based automation tools Prytaneum uses to automate certain parts of its content curation functionality. These tools emphasize human-in-the-loop feedback with the goal of ensuring only limited reliance on machine-learned prediction or classification.

Topic pre-selection. One of the main ways Prytaneum helps organize content submitted by the audience is by topic. Prytaneum events are most effective when they are

centered on an issue, and the organizer should plan for discussion of topics that compose the issue. When setting up the event, the host must choose their topics and give textual descriptions of each topic. If the topics are to be drawn from a written source (preferably, vetted, politically balanced and factually accurate), Prytaneum's LLM can extract suggested topics and descriptions from that written source. The user is then shown these suggested topic labels/descriptions and can manually amend them as they wish as a human in the loop. The automatic topic extraction tool scans the provided reading materials and uses the LLM to infer from them an initial set of relevant topics and topic definitions. The user is then prompted to make changes to the topics as they see fit. Possible changes include manually adding or removing topics, adjusting topic definitions, and regenerating a new set of topics via the LLM. Once satisfied, the user may confirm and exit the topic curation process, yielding a finalized set of topics and definitions.

Live Content Curation. In order to help the facilitator manage a large volume of incoming audience content, Prytaneum's machine learning algorithms automatically tag and organize the content by topic (where the "topics" are those preselected above). The facilitator can easily filter content by topic to quickly see the diverse perspectives and variety of content relevant to that topic, in order to ensure each topic is considered thoroughly before moving on to the next topic. This filtering also removes irrelevant and/or offensive comments, and mitigates redundancy. The content sorting of user comments functions is a pipeline of independent modules. The first step in the pipeline verifies that the content is substantive to the overall topic, preventing the content from being classified if it is deemed unsubstantive. The next step is detecting offensive language in the comment, where the facilitator can select whether or not they wish to classify content deemed offensive. As the default behavior, Prytaneum does not classify offensive comments. As the final step in the pipeline, the content is tagged with the relevant topic label(s) derived from the previous topic extraction feature via an LLM.

Clustering of similar content. When an event scales to a large number of participants, it is common that many audience members will submit semantically-similar content but using different wordings, grammar or syntax. Prytaneum uses word embeddings and sophisticated k-means clustering algorithms to group similar audience posts together so the facilitator can quickly see which comment version is best for advancing into the discussion for further consideration. After the comments are sorted into topics, Prytaneum utilizes content clustering to merge similar questions together. Specifically, every comment within every topic is internally converted to word embeddings vector, yielding a 384-dimensional vector per comment. Next, the bisecting k-means algorithm is used to cluster together the vectors with every K from 2 to 10. Finally, the silhouette score method is used to select the optimal K, yielding K clusters of comments with semantically similar meanings.

Feedback-response live polling. This is perhaps Prytaneum's most powerful feature to promote engagement and interactivity during a synchronous event. With this feature, the speaker can deploy either open-ended or closed-ended survey questions, but in either case, the respondent cannot submit a response until they provide the *reasoning* to support their choice or perspective. The LLM summarizes the perspectives and reasoning among the audience for each option presented, and the host can share the summarization of the reasoning and perspectives with the audience. This helps to ensure that diverse perspectives are heard and that all audience members have a voice and see each others' perspectives even when the event scales to a large number of participants.

Prompt summarization begins by aggregating all user responses and utilizing the LLM to filter out off-topic and offensive content and then to extract the different viewpoints. For each prompt, the facilitator is given the option to summarize viewpoints in general, or to focus on perspectives on a particular topic. Depending on the prompt type, the responses can be summarized in one of two granularities: overall viewpoints and per-

stance (agree, disagree, neutral) viewpoints. In the future, Prytaneum will implement a new feature that enables audience members to respond to the summaries themselves, to indicate areas of continuing disagreement, and these responses to the summaries will themselves be summarized and shared with the audience and speaker. In this way, each perspective is voiced in the meeting, and each audience member has the opportunity to offer reasons for and against each option that the speaker proposes in each response-feedback prompt.

Real-time language translation. Prytaneum currently allows users to select English or Spanish language. In the future, Prytaneum will be language agnostic, where participants can choose any languge as their preferred language.

Typical Prytaneum meeting workflow. Prytaneum's feedback-response feature enables a re-envisioning of the traditional format for speaker-led meetings. In the traditional format, audience members are able to ask questions and make comments one at a time – and typically at the end of the presentation. When the meeting scales to a large number of people, the Q & A format ensures that the speaker can address only a small number of participants' concerns or suggestions during the synchronous part of the meeting. As a result, most of the participants will experience dissatisfaction in that their perspectives were not addressed in the meeting. In addition, the Q & A format does not allow participants to participate dynamically or interactively beyond contributing questions and comments to the list, and so might lead audience members to become passive listeners.

The feedback-response feature allows the meeting facilitator to organize the meeting in manner that is much more interactive and dynamic. In planning a meeting, the organizer creates a set of topics that will be the focus of the discussion. In a typical workflow, the facilitator can work with the speaker to create a set of prompts relevant to each of the topics. During the meeting, for each topic under consideration, the facilitator can deploy the relevant prompt to the audience, and then can ask the speaker to share some

thoughts about the prompt; for example, at a public meeting, the speaker can share how their agency or office is currently addressing the topic of the prompt. As the speaker shares these thoughts, the audience is able to process the prompt and the speaker's insights to deeply process the information and develop a considered opinion regarding the prompt.

Once the speaker is ready, the facilitator will ask the audience to then complete the response to the prompt, where they select their preferred option and share their reasons for their choice.[6] Once the facilitator closes the feedback-response, they will use the LLM to summarize the perspectives among the submissions, and then they can share the summaries back to the audience. In this way, the audience is able to read and understand the reasoning for why different stakeholders support the different options, exposing community members to the diverse reasoning for collective decisions among community members. In this redesign of the traditional meeting, each audience member has a voice on each topic, and participants are exposed to the full diversity of perspectives in the community, not just the loudest or most persistent voices. Furthermore, each audience member has an opportunity to be actively engaged throughout the meeting and so will experience the meeting as being more interactive, dynamic and deliberative.

6 Comparison to Asynchronous Deliberative Platforms

There are number of asynchronous deliberation platforms that enable the scaling of deliberation among many participants.[7] Asynchronous communication is the core mode of communication in online settings and has immense value. However, asynchronous and synchronous communication are very different in the many-to-one context. In particular, even though many participants on an asynchronous platform can submit comments simultaneously, in practice asynchronous communication resides in the design space where

[6] Likewise, the facilitator can pair this with a related prompt asking each member to choose their least preferred option and their reasons for opposing that option. By pairing prompts in this way, the participants can understand both arguments for and arguments against each option.

[7] Perhaps the best known asynchronous deliberation platform is Polis.

each participant speaks in turn and so does not solve the many-to-one communication situation. In particular, because asynchronous communication unfolds over time, participants experience the flow of the conversation as sequential, and there can be dependence in participants' contributions due to the sequence. Any summary of the platform content in turn will depend on that sequence. In addition, it is likely that social cognition (Chambers, 2018) in synchronous and asynchronous contexts differ, and deliberative theory has mostly been developed with the synchronous case in mind.

7 Technology-Enabled Deliberative Theory

As an environment designed specifically for promoting many-to-one communication, Prytaneum enables us to consider the counterfactual of how communication can be improved in this setting. In the absence of a platform designed to promote many-to-one communication, it is hard even to conceive of many-to-one communication, much less to consider theory, design and practice to promote it. In virtually all settings, in the absence of this new technology, it is simply impractical for all audience members to speak at once during a meeting. As a result, communication theory has been developed for the case where only one person can speak at a time, and there currently is no normative theory to support the case of many-to-one communication. However, the situations that call for many-to-one communication are ubiquitous, and occur whenever there is a synchronous, speaker-led meeting that has more than a handful of participants in the audience.

Here we use the case study of Prytaneum to develop our normative framework for many-to-one communicative action. This framework will help to lay the groundwork for a technology-enabled normative theory to support synchronous deliberation at scale. And while many-to-one communicative action can improve meetings in virtually any context, we particularly highlight the urgency of developing new methods for the democratic use case of public meetings. Long-term declines in trust in government, combined with the

rise of polarization and misinformation, strongly suggest the need to develop new methods to connect the people with their democratic government; to enable individuals to have a voice in policy making; for government officials to understand the diverse views of their communities, to promote constructive discussions across differences; and to ensure that diverse and under-represented perspectives in the community are heard and empowered (Neblo et al., 2018).

These democratic aspirations apply to virtually any meeting context where a speaker wishes to not only communicate their perspective to an audience, but also to understand the audience's diverse perspectives as a way to inform their own work and decisions. The goal is to design computing environments that empower participants to have effective voice, and to ensure the participants' contributions remain constructively on topic; that participants and the speaker are exposed to and understand diverse perspectives in their community, especially for perspectives that give voice to the concerns of underrepresented stakeholders; and that encourage participants to be actively engaged and accountable through out the meeting. In modeling a means to enable many-to-one communicative action, Prytaneum can usher in a new generation of computing environments designed to enhance democratic deliberation and group-based collaboration at the scale of modern democracy.

References

Bail, C. A., L. P. Argyle, T. W. Brown, J. P. Bumpus, H. Chen, M. B. F. Hunzaker, J. Lee, M. Mann, F. Merhout, and A. Volfovsky (2018, 9). Exposure to opposing views on social media can increase political polarization. *Proc. Natl. Acad. Sci. U. S. A. 115*(37), 9216–9221.

Barrett, P., J. Hendrix, and G. Sims (2021, 9). How tech platforms fuel political polarization and what government can do about it.

https://www.brookings.edu/blog/techtank/2021/09/27/how-tech-platforms-fuel-u-s-political-polarization-and-what-government-can-do-about-it/.

Bessette, J. M. (1994). *The Mild Voice of Reason: Deliberative Democracy and American National Government*. Chicago, Ill.: University of Chicago Press.

Bohman, J. (1998). Survey Article: The Coming of Age of Deliberative Democracy. *Journal of Political Philosophy 6*(4), 400–425.

Broockman, D. and J. Kalla (2016). Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science 352*(6282), 220–224.

Cacioppo, J. T., R. E. Petty, and C. Feng˜Kao (1984). The Efficient Assessment of Need for Cognition. *Journal of Personality Assessment 48*(May), 306–307.

Chambers, S. (2018). Human Life is Group Life: Deliberative Democracy for Realists. *Critical Review 30*(1-2), 36–48.

Cohen, J. (1989). Deliberation and Democratic Legitimacy. In A. Hamlin and P. Pettit (Eds.), *The Good Polity: Normative Analysis of the State*, pp. 17–34. New York, N.Y.: Basil Blackwell.

Conover, P. J., D. D. Searing, and I. M. Crewe (2002). The Deliberative Potential of Political Discussion. *British Journal of Political Science 32*, 21–62.

Cooper, J. and D. Mackie (1983). Cognitive Dissonance in an Intergroup Context. *Journal of Personality and Social Psychology 44*(3), 536–544.

Druckman, J. N. (2004). Political Preference Formation: Competition, Deliberation, and the (Ir)relevance of Framing Effects. *American Political Science Review 98*(Nov.), 671–686.

Druckman, J. N. and K. R. Nelson (2003). Framing and Deliberation: How Citizens' Conversations Limit Elite Influence. *American Journal of Political Science 47*(Oct.), 729–745.

Dryzek, J. S., A. Bächtiger, S. Chambers, J. Cohen, J. N. Druckman, A. Felicetti, J. S. Fishkin, D. M. Farrell, A. Fung, A. Gutmann, H. Landemore, J. Mansbridge, S. Marien, M. A. Neblo, S. Niemeyer, M. Setälä, R. Slothuus, J. Suiter, D. Thompson, and M. E. Warren (2019, 3). The crisis of democracy and the science of deliberation. *Science 363*(6432), 1144–1146.

Einstein, K. L., D. Glick, L. G. Puig, and M. Palmer (2023). Still Muted: The Limited Participatory Democracy of Zoom Public Meetings. *Urban Affairs Review 59*(4), 1279–1291.

Einstein, K. L., M. Palmer, and D. M. Glick (2019, 3). Who Participates in Local Government? Evidence from Meeting Minutes. *Perspectives on Politics 17*(1), 28–46.

Elliot, A. J. and P. G. Devine (1994). On the Motivational Nature of Cognitive Dissonance: Dissonance as Psychological Discomfort. *Journal of Personality and Social Psychology 67*(3), 382–394.

Esterling, K. M. (2004). *The Political Economy of Expertise: Information and Efficiency in American National Politics*. Ann Arbor, Mich.: University of Michigan Press.

Esterling, K. M., A. Fung, and T. Lee (2021). When Deliberation Produces Persuasion Rather than Polarization: Measuring and Modeling Small Group Dynamics in a Field Experiment. *British Journal of Political Science 51*(2), 666–684.

Esterling, K. M., M. A. Neblo, and D. M. Lazer (2011). Means, Motive, and Opportunity in Becoming Informed about Politics: A Deliberative Field Experiment Involving Members of Congress and their Constituents. *Public Opinion Quarterly 75*(Fall), 483–503.

Fishkin, J. (2019). Response to Michael A. Neblo, Kevin M. Esterling, and David M. J. Lazer's review of Democracy When the People Are Thinking: Revitalizing Our Politics Through Public Deliberation. *Perspectives on Politics 17*(02), 531–531.

Freeman, S. (2000). Deliberative Democracy: A Sympathetic Comment. *Philosophy and Public Affairs 29*(4), 371–418.

Gerber, A. S., G. A. Huber, D. Doherty, and C. M. Dowling (2011). The Big Five Personality Traits in the Political Arena. *Annual Review of Political Science 14*, 265–287.

Grönlund, K., K. Herne, and M. Setälä (2015). Does Enclave Deliberation Polarize Opinions? *Political Behavior 37*, 995–1020.

Gutmann, A. and D. Thompson (1996). *Democracy and Disagreement: Why Moral Conflict cannot be Avoided in Politics, and What Should be Done about It*. Princeton, N.J.: Princeton University Press.

Gutmann, A. and D. Thompson (2004). *Why Deliberative Democracy?* Princeton, N.J.: Princeton University Press.

Habermas, J. (1984). *The Theory of Communicative Action, Volume One: Reason and the Rationalization of Society* (McCarthy ed.). Boston, Mass.: Beacon Press.

Habermas, J. (1996). *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy* (Rheg ed.). Cambridge, Mass.: MIT Press.

Hills, T. T. (2019). The Dark Side of Information Proliferation. *Perspectives on Psychological Science 14*(3), 323–330.

Hinsz, V. B., R. S. Tindale, and D. A. Vollrath (1997). The Emerging Conceptualization of Groups as Information Processors. *Psychological Bulletin 121*(1), 43–64.

Isenberg, D. J. (1986). Group Polarization: A Critical Review and Meta-Analysis. *Journal of Personality and Social Psychology 50*(6), 1141–1151.

Kahneman, D. (2011). *Thinking, Fast and Slow*. New York, N.Y.: Farrar, Straus and Giroux.

Karpowitz, C. F., T. Mendelberg, and L. Shaker (2012, 8). Gender Inequality in Deliberative Participation. *Am. Polit. Sci. Rev. 106*(03), 533–547.

Karpowitz, C. F. and C. Raphael (2014). *Deliberation, Democracy, and Civic Forums: Improving Equality and Publicity*. New York, NY: Cambridge University Press.

Lazer, D. M. (2015). The Rise of the Social Algorithm. *Science 348*(6239), 1090–1091.

Lerner, J. S. and P. E. Tetlock (1999). Accounting for the Effects of Accountability. *Psychological Bulletin 125*(2), 255–275.

MacKuen, M., J. Wolak, L. Keele, and G. E. Marcus (2010). Civic Engagements: Resolute Partisanship or Reflective Deliberation. *American Journal of Political Science 54*(April), 440–458.

Mann, T. E. and N. J. Ornstein (2016, 4). *It's Even Worse Than It Looks: How the American Constitutional System Collided with the New Politics of Extremism*. Basic Books.

Mansbridge, J. J. (1980). *Beyond Adversary Democracy*. New York, N.Y.: Basic Books.

Marble, W. and C. Nall (2020). Where interest trumps ideology: Homeownership and liberal opposition to local housing development. *Journal of Politics 83*(4), 1747–1763.

Matz, D. C. and W. Wood (2005). Cognitive Dissonance in Groups: The Consequences of Disagreement. *Journal of Personality and Social Psychology 88*(1), 22–37.

Mendelberg, T. and J. Oleske (2000). Race and Public Deliberation. *Political Communication 17*, 169–191.

Mercier, H. and H. Landemore (2012, 4). Reasoning Is for Arguing: Understanding the Successes and Failures of Deliberation. *Political Psychology 33*(2), 243–258.

Morrell, M. E. (1999). Citizens' Evaluations of Participatory Democratic Procedures: Normative Theory Meets Empirical Science. *Political Research Quarterly 52*(June), 293–322.

Mutz, D. C. (2006). *Hearing the Other Side: Deliberative versus Participatory Democracy*. New York, N.Y.: Cambridge University Press.

Neblo, M. A. (2015). *Deliberative Democracy between Theory and Practice*. Cambridge: Cambridge University Press.

Neblo, M. A., K. M. Esterling, and D. M. Lazer (2018). *Politics with the People: Building a Directly Representative Democracy*. New York, N.Y.: Cambridge University Press.

Niemeyer, S., F. Veri, J. S. Dryzek, and A. Bächtiger (2024, 2). How Deliberation Happens: Enabling Deliberative Reason. *American Political Science Review 118*(1), 345–362.

Noveck, B. S. (2105). *Smart citizens, smarter state: The technologies of expertise and the future of governing*. Cambridge, MA: Harvard University Press.

Ognyanova, K., D. Lazer, R. E. Robertson, and C. Wilson (2020, 6). Misinformation in action: Fake news exposure is linked to lower trust in media, higher trust in government when your side is in power.

Plemmons, D. K., E. N. Baranski, K. Harp, D. D. Lo, C. K. Soderberg, T. M. Errington, B. A. Nosek, and K. M. Esterling (2020, 1). A randomized trial of a lab-embedded

discourse intervention to improve research ethics. *Proceedings of the National Academy of Sciences 117*(3), 201917848.

Prior, M. (2013). Media and Political Polarization. *Annual Review of Political Science 16*(1), 101–127.

Rhodes, S. C. (2022, 1). Filter Bubbles, Echo Chambers, and Fake News: How Social Media Conditions Individuals to Be Less Critical of Political Misinformation. *Political Communication 39*(1), 1–22.

Sanders, L. M. (1997). Against Deliberation. *Political Theory 25*(3), 347–376.

Settle, J. M. (2018). *Frenemies: How Social Media Polarizes America*. New York, N.Y.: Cambridge University Press.

Suedfeld, P., P. E. Tetlock, and S. Streufert (1992). Conceptual/integrative complexity.

Sunstein, C. R. (2002). The Law of Group Polarization. *Journal of Political Philosophy 10*(2), 175–195.

Sunstein, C. R. (2009, 5). *Going to Extremes: How Like Minds Unite and Divide*. Oxford University Press.

Sunstein, C. R. (2017). *#Republic: Divided Democracy in the Age of Social Media*. Princeton, NJ: Princeton University Press.

Tetlock, P. E., L. Skitka, and R. Boettger (1989). Social and Cognitive Strategies for Coping with Accountability: Conformity, Complexity and Bolstering. *Journal of Personality and Social Psychology 57*(4), 632–640.

Therriault, A., J. A. Tucker, and T. Brader (2011). Cross-Pressures and Political Participation. *OpenSIUC Paper 23*, \url{http://opensiuc.lib.siu.edu/pnconfs_2011/23}.

Trounstine, J. (2020, 5). The Geography of Inequality: How Land Use Regulation Produces Segregation. *Am. Polit. Sci. Rev. 114*(2), 443–455.

Vorauer, J. D., A. Gagnon, and S. J. Sasaki (2009). Salient Intergroup Ideology and Intergroup Interaction. *Psychological Science 20*(7), 838–845.

Vorauer, J. D. and S. J. Sasaki (2011). In the Worst Rather than the Best of Times: Effects of Salient Intergroup Ideology in Threatening Intergroup Interactions. *Journal of Personality and Social Psychology Advance On*(doi: 10.1037/a0023152), 1–14.

Vosoughi, S., D. Roy, and S. Aral (2018, 3). The spread of true and false news online. *Science 359*(6380), 1146–1151.