Collaborative e-Rulemaking, Democratic Bots, and the Future of Digital Democracy

OREN PEREZ, Bar-Ilan University, Dean, Faculty of Law

This article focuses on “deliberative e-rulemaking”: digital consultation processes that seek to facilitate public deliberation over policy or regulatory proposals [1, 2]. The main challenge of e-rulemaking platforms is to support an “intelligent” deliberative process that enables decision makers to identify a wide range of options, weigh the relevant considerations, and develop epistemically responsible solutions. This article discusses and critiques two approaches to this challenge: The Cornell RegulationRoom project and model of computationally assisted regulatory participation by Livermore et al. It then proceeds to explore two alternative approaches to e-rulemaking: One is based on the implementation of collaborative, wiki-styled tools. This article discusses the findings of an experiment, which was conducted at Bar-Ilan University and explored various aspects of a wiki-based collaborative e-rulemaking system. The second approach follows a more futuristic approach, focusing on the potential development of autonomous, artificial democratic agents. This article critically discusses this alternative, also in view of the recent debate regarding the idea of “augmented democracy.”

CCS Concepts: • **Social and professional topics** → **Governmental regulations**; • **Applied computing** → **E-government**;

Additional Key Words and Phrases: Collaborative e-rulemaking, digital democracy, public good, e-democracy artificial agents, social bots, augmented democracy

ACM Reference format:

Oren Perez. 2020. Collaborative e-Rulemaking, Democratic Bots, and the Future of Digital Democracy. *Digit. Gov.: Res. Pract.* 1, 1, Article 8 (January 2020), 13 pages.
<https://doi.org/10.1145/3352463>

1 INTRODUCTION

In the past decade, the landscape of “e-democracy” has expanded significantly. Governments and international organizations have experimented with various initiatives, ranging from digital consultation platforms and e-petition mechanisms to various crowd-sourcing platforms. In this article, I focus on a particular type of e-democracy scheme that I term “deliberative e-rulemaking”: digital consultation platforms that aim to facilitate public deliberation over policy or regulatory proposals [3].¹ Such platforms can be used on different scales,

¹I do not examine in this article e-rulemaking systems designed to produce binding decisions, which are rarely used in regulatory contexts.

Author’s address: O. Perez, Bar-Ilan University, Faculty of Law, Ramat Gan, 52900, Israel; email: oren.perez@biu.ac.il.

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.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2639-0175/2020/01-ART8 \$15.00

<https://doi.org/10.1145/3352463>

ranging from the international and national levels to the local level (e.g., in the context of urban and university communities) [4–6].² Despite the increasing popularity of the idea of e-democracy, digital participatory schemes still occupy a peripheral place in regulatory praxis. This reflects a certain disillusionment with the capacity of e-democracy to generate processes that serve concurrently the goals of democratization and of “good governance” [1, 2]. The literature flags two problems in this respect. The first problem concerns the “intelligence” of the deliberation process: “whether it enables decision makers to identify a wide range of options, to weigh relevant considerations, and to act on accurate or at least defensible beliefs” [8]. Various writers have noted that the epistemic quality of online deliberative processes can be disappointing [9, 10]. Many of the comments received by agencies amount to short statements of approval or disapproval, which echo the opinions of other dominant players. Such comments do not make a significant epistemic contribution to the deliberation [11].

The second problem concerns the dynamics of e-rulemaking processes. Many such processes fail to attract a sufficient number of participants [3, 12]. For example, participation levels in Regulations.gov, the central hub for citizens’ comments on U.S. federal regulations, seems to follow a power law distribution: most of the proposals published on Regulations.gov get little or no attention, and a small number of proposals receive a significant number of comments [2, 13].³ Low participation levels can undermine both the legitimacy of e-rulemaking initiatives and their capacity to generate novel ideas.

There appears to be a dissonance between the optimistic vision that has underpinned the project of e-democracy in its first days, which was driven by a combination of pro-democratic sentiment and a belief in the transformative power of technology, and the relatively negligible effect of e-democracy platforms on the regulatory decision-making process in the U.S. and elsewhere [3]. This dissonance also characterized President Obama’s “Open Government Directive” (OG Directive) and its rulemaking program [17].⁴

The Cornell eRulemaking Initiative (CeRI) represents one prominent approach for dealing with this dissonance. CeRI was based on the insight that simply providing citizens with the opportunity to comment online (the approach of the Regulations.gov portal) is not sufficient to facilitate widespread and epistemically valuable deliberation. The Cornell team developed a web-based consultation platform (RegulationRoom) based on purposeful design of digital tools and intensive human support. For example, the team developed a targeted commenting functionality, which encouraged detailed participation, focused on specific aspects of the regulatory proposal. The design was supported by trained facilitators who guided the deliberation process. The RegulationRoom team demonstrated that using this two-pronged approach increased participation levels and facilitated more substantive citizen contributions. At the same time, the project underscores the fact that eliciting such participation requires the investment of considerable effort, both by citizens and by those who seek their informed policy input [18].

The RegulationRoom project sought to cope with the dilemmas of mass-commenting by trying to improve the epistemic value of each comment. Another approach, suggested by Michael A. Livermore, Vladimir Eidelman, and Brian Grom in a recent paper, is to extract more meaning from comments in their current form, using advanced information-processing and text analysis techniques [10]. Livermore and colleagues identified two key problems that regulatory agencies engaged in e-rulemaking are facing. The first is the “needle-in-the-haystack” problem: how to identify the comments with the most substantive content (those that possess the highest prospect of improving the proposed rule) in a large body of comments. The second is a “forest-for-the-trees” problem: how

²I take a broad view of rulemaking, which covers both national law, secondary rules produced by administrative or municipal authorities, and soft law produced by private institutions, such as universities or transnational standard-setting organizations [7].

³Some recent examples of proposed rules that received thousands of comments are the net neutrality rules proposed by the FCC in 2014 [14]; USDA’s proposed rules governing competitive foods in schools [15]; and the CDC quarantine rule [16].

⁴See Memorandum For The Heads Of Executive Departments And Agencies, Dec. 8, 2009, Open Government Directive, https://www.whitehouse.gov/sites/default/files/omb/assets/memoranda_2010/m10-06.pdf. The Trump administration has recently released the fourth iteration of the national action plan, also known as NAP4 (<https://open.usa.gov/assets/files/NAP4-fourth-open-government-national-action-plan.pdf>).

to extract overall trends or themes in large, unstructured collections of documents. The authors describe several techniques that can be used to solve these challenges computationally. For example, to identify the comments with the highest epistemic value (or gravitas), they use a method developed by FiscalNote, which ranks comments based on several identifying features, including length, complexity (or coarseness) of the language, whether the author is an organization, key person, or ordinary individual, and the number of cogent arguments expressed [10]. To identify general patterns in the data, they use topic modeling, a computational tool that can infer subject matter categories in an unstructured corpus and work by identifying patterns in the co-occurrence of words in documents [10]. The authors demonstrate this approach by an analysis of the comments received by the Environmental Protection Agency (EPA) in response to its proposed rule to limit greenhouse gas emissions by the electricity-generating sector, the Clean Power Plan.⁵

Both these approaches suffer from serious shortcomings. The RegulationRoom model, which relies heavily on human intermediation to improve the sophistication of comments, is costly, and therefore it cannot serve as a general template for e-rulemaking processes. There is also a deep tension between the fundamental ethos of directly deliberative democracy, which emphasizes the value of free discussion, unfettered by external intervention, and the emphasis on proactive human intermediation in the RegulationRoom model. Livermore, Eidelman, and Grom's model of computationally assisted regulatory participation, which focuses on optimizing the extraction of usable information from the existing imperfect body of public comments, reflects a more feasible approach [10]. From a democratic perspective, however, the approach of Livermore et al. raises two key problems: It leaves unresolved the question of how to improve the quality of the deliberation process [10], and the use of computational tools to analyze comments creates new challenges of accountability and control. As Livermore et al. have shown, these tools must be carefully calibrated to the particular corpus to which they are applied. The calibration process is highly discretionary and raises obvious agency problems [17].

In this article, I explore two additional responses to the challenge of e-rulemaking; one is based on the implementation of collaborative, wiki-styled tools; the other adopts a more futuristic approach and focuses on the potential development of autonomous, artificial democratic agents.

2 CHALLENGES IN THE IMPLEMENTATION OF COLLABORATIVE (WIKI) E-RULEMAKING

I argue that the approach of *collaborative e-rulemaking* deserves careful consideration primarily, because it is more consistent with the spirit of directly deliberative democracy than are other approaches [19].⁶ Conventional notice-and-comment platforms, such as Regulations.gov, leave no room for dialogue between citizens, and grant the power to consolidate and interpret the comments to the regulatory agency [2]. By contrast, collaborative writing platforms⁷ shift the responsibility to the participants by enabling users to jointly create a single text, and by allowing participants to talk to each other, usually through a dedicated side forum [20, 21].⁸ Various writers discussed the use of collaborative writing platforms as a medium for political deliberation with reference to wiki technology [22–24]. Beth Noveck used “Wiki-Government” as a general term to describe the concept of collaborative democracy that realizes the ideal of legitimate governance in the 21st Century [25].

Scholars of e-democracy have been toying with the idea of using wiki technology in democratic deliberation for quite some time, but such tools have rarely been used in actual regulatory consultations [26]. Prominent e-democracy platforms, such as the U.S. Regulations.gov, the Canadian Open Government Consultation portal,

⁵One of the first initiatives of the Trump administration was to repeal this rule. The new proposal is currently being debated using the same administrative processes. See Proposed EPA Rule on GHG Emissions from Coal-Fired Power Plants (83 FR 65424) <https://www.federalregister.gov/documents/2018/12/20/2018-27052/review-of-standards-of-performance-for-greenhouse-gas-emissions-from-new-modified-and-reconstructed>.

⁶Regulations.gov mimics the commenting structure of the U.S. Administrative Procedure Act 5 U.S.C. §§ 553 (2012) (<http://www.archives.gov/federal-register/laws/administrative-procedure/553.html> and <https://www.regulations.gov/#!faqs>).

⁷Examples include GoogleDoc, Etherpad, Mediawiki, and FidusWriter.

⁸See the “talk page” in Wikipedia: https://en.wikipedia.org/wiki/Help:Using_talk_pages.

and the British e-petitions platform⁹ do not include collaborative writing tools [9, 27].¹⁰ This gap is reflected also in both the observational and experimental literature on digital democracy, which includes almost no discussion of collaborative tools, such as wikis [28–31].

The attempt to develop successful e-rulemaking programs faces two key challenges from the perspective of deliberative democracy: (a) how to achieve an epistemologically complex deliberation, and (b) how to attract a sufficient number of contributors actively involved in the deliberation.¹¹ The literature identifies several factors as potential causes for low participation levels and epistemologically shallow dialogue.

Social-structural factors. The first factor concerns the *public good* nature of participation in regulatory consultation. Deliberative e-rulemaking has the structure of a “give-some” public good game: participants are expected to contribute to the group by investing time and thought. The potential benefits of this contribution (e.g., improved rule or policy) are shared by the entire community, *regardless* of the level of effort invested by each member in the deliberative process [28, 32].¹² In this respect, e-deliberation resembles other digital settings, such as product reviews on Amazon, hotel recommendations on TripAdvisor, and editing in Wikipedia [33]. The public good feature of deliberative e-rulemaking means that people have strong incentive to sit on the fence and let others do the job. This generates a collective action dilemma, which can undermine participation. Empirical studies of collective action situations demonstrate that in many cases the public good is ultimately produced by a small subset of contributors (critical mass theory) [33–35]. The steep long-tail of participation in Wikipedia is a strong illustration of this phenomenon, which leaves open the question of the motivations or personality traits that drive individuals to participate in online deliberation, despite its public good nature; I return to this question below [28, 33, 36].

A second factor that may inhibit participation is external efficacy, which reflects people’s perceptions of the responsiveness of the regime [37]. If people do not believe that their contributions can make a difference, then they will hesitate to commit their limited time and cognitive resources to participating in public deliberation processes [3, 38]. This factor operates independently of the collective action factor, because the issue here is not the temptation to let others do the job, but the futility of the deliberation process as a whole.

The scale of the deliberation process constitutes a further potentially inhibitive factor. In large-scale participatory processes (e.g., Regulations.gov), participants do not know each other, and there is almost no opportunity for reciprocity, which exacerbates the collective action problem. In smaller communities (e.g., small town, university), there may be some scope for cross-domain reciprocity, and there may be other considerations, associated with group identity, that can provide stronger motivation for group members to contribute [39]. The size of the community also affects political efficacy: in smaller communities, the prospects of influencing the decision-making process are higher.

Another structural element that may influence the dynamics of e-deliberation processes is whether they take place in an anonymous or an identifiable setting. Anonymity can have conflicting effects. The literature has shown that anonymity can lead to uncivil forms of behavior through toxic disinhibition [40, 41], because it reduces people’s accountability for their behavior. Uncivil forms of communication, especially those that are interpersonal, can undermine people’s willingness to participate in deliberative processes, and prevent them from expressing their views freely and sincerely [42]. But anonymity can also have positive outcomes, because

⁹<http://www.regulations.gov/#!home>; <https://open.canada.ca/en>; <https://petition.parliament.uk/>.

¹⁰For further examples for consultation platforms that do not employ collaborative tools, see the recent consultation on the “Better Regulation Agenda” by the European Commission (“Stocktaking of the Commission’s ‘better regulation’ approach,” Ref. Ares (2018) 2332204–02/05/2018), and the European Medicines Agency consultation on its Regulatory Science 2025 policy (<https://etp-nanomedicine.eu/ema-regulatory-2025-consultation/>).

¹¹Additional challenges include the potential influence of business interests, the problem of mass e-mails, and the facilitation of regulatory responsiveness.

¹²The social dilemma underlying deliberative e-rulemaking differs from “take-some” dilemmas, in which there is a common resource from which each member of a group decides how much to harvest.

it allows people to express critical and controversial views without worrying about potentially adverse effects to themselves [43]. Research has shown that anonymity in comment boards expands the number of participants in the conversation and the range of views aired [44].

Individual personality factors. Some conspicuous factors that can influence levels of participation are technological competence, educational background, and the ability to engage in a complex deliberative discussion [45, 46]. In a recent paper, I examined the potential influence of another personality trait: participants' social value orientation (SVO) [47]. SVO reflects people's "stable preferences for certain patterns of outcomes for oneself and others" [48]. Three categories of SVO can be distinguished: prosocial, individualistic, and competitive [48]. *Prosocials* tend to maximize outcomes for themselves and others (reflecting a cooperative disposition); *individualists* tend to maximize their own outcomes, with little or no regard for others' outcomes; *competitors* tend to maximize their own outcomes relative to those of others, seeking an advantage over them [48].

The three SVOs are predictive of behavior in a variety of social dilemma situations. There is abundant research, in both laboratory and field settings, indicating that *prosocials* generally cooperate more and show greater concern for the consequences of their behavior on others and on the environment [32, 49] than do *proselfs* (*individualists* and *competitors*). For example, laboratory experiments of *give-some* public good games, found *prosocials* to be more cooperative than *proselfs* [49]. Field studies also indicate that a prosocial value orientation is positively related to cooperative behavior. Van Lange et al. found that *prosocials* are more likely to volunteer for psychological experiments than are individualists and competitors [50]. A series of experimental and field studies have found a positive correlation between prosocial disposition and pro-environmental behavior. Joireman et al. found that *prosocials* tend to show greater willingness to engage in pro-environmental behavior [51] than do *proselfs*. Other studies found that *prosocial* commuters exhibited greater preference for public transportation than did *proself* commuters [52], and that *prosocials* reported more energy conservation behavior than did *proselfs* [53].

The influence of SVO on participation in web-based deliberation has not been explored in the literature, although these processes exhibit the features of a public good dilemma [39, 54].¹³ Participating in web-based deliberation, especially when it is combined with a collaborative writing tool, differs from other political behaviors, such as voting or signing a petition, in that it provides participants with wider opportunities to derive satisfaction from the activity itself. The unique features of e-deliberation produce two competing hypotheses regarding the potential effect of SVO on levels of participation.

First, the *give-some* public good aspect of deliberative e-rulemaking suggests that *prosocials* exhibit higher levels of participation than do *proselfs*. This hypothesis is consistent with the literature on SVO and public goods. The idea that citizens with a stronger disposition toward altruism exhibit higher participation levels also received some empirical support [57]. A second hypothesis focuses on the capacity of e-deliberation to provide participants with some incidental benefits. Individuals may derive some satisfaction from the opportunity to express their views and from their ability to shape the direction of the online discussion and of the collaborative product [33]. Because participatory platforms that use a collaborative writing tool provide participants with ample opportunities for self-expression and for influencing the discussion, one expects *individualists* as well to show high levels of participation.

3 EXPERIMENTAL STUDY OF A HYBRID COLLABORATIVE E-RULEMAKING SYSTEM

In a recent study [47], several colleagues and I examined the capacity of a hybrid collaborative deliberative e-rulemaking system to improve the epistemic quality of the deliberative process and its outcome. We tested the system in a field experiment conducted at Bar-Ilan University (BIU), in December 2014. We designed the experiment in a way that closely resembled a real regulatory consultation. We closely monitored the behavior

¹³Most of the studies that considered the influence of personality on participation have focused on general web-based forums, such as Facebook [39, 54–56].

of participants during the experiment. The focus of the deliberation process was the Bar-Ilan University Code Concerning Political Activity on Campus (Code). The BIU Code establishes procedures and rules for conducting political activities on campus. Different provisions of the Code seek to balance the values of freedom of expression, tolerance, and security. The Code is not unique to BIU. Of the seven research universities in Israel, six have similar codes.¹⁴ Furthermore, part of the issues discussed by the BIU Code, such as students' rights of expression and the freedom of association, are governed by formal legislation: the Student Rights Act, 2007.¹⁵ The issues discussed in our deliberation experiment were therefore of relevance to the entire student population in the country.

3.1 The Deliberation Platform

The platform included a forum on which participants could write posts and respond to posts by others, and an integrated collaborative writing tool that enabled participants to jointly contribute to (and edit) text that summarized the group's thinking about the BIU Code. The experiment included four groups of about 25 students. The design of the platform was based on the following principles [3, 17, 58].

1. *Simplicity.* We created a simple platform with a highly intuitive user interface, which allowed users to move easily between the wiki and the forum components. We expected most users in the age group of our participants (20-30) to feel comfortable posting and commenting in a forum, which was close in its structure to Facebook, but using a collaborative writing tool is less common. The tool we used had a simple structure that enabled simultaneous editing and was close in its design to the Google Docs platform. We integrated this tool into the main platform using a special button ("Editing the joint document").
2. *Multi-Layered Information Design.* A key challenge for e-deliberation is making sure that participants are provided with accessible yet sufficiently complex data about the debated topic. To meet the challenge posed by users' short attention span the platform included several layers of information. We created a concise version of the BIU Code, which included a brief summary of the main code provisions. This document also included a reference to parallel paragraphs in comparable codes at other universities in Israel and the U.S. (Hebrew University, Ben-Gurion University, Tel-Aviv University, and UC Berkeley). In addition, we also provided participants with the full text of the Code. We also provided participants with several articles from the press, which discussed various aspects of the issue of political activity on Israeli campuses.
3. *Motivating participants through rewards and realistic experimental design.* The public-good aspect of deliberative e-rulemaking creates a motivation challenge. The case of the BIU Code exemplifies this challenge especially, because it is somewhat distant from students' everyday concerns, which are grades, teaching quality, and price of food at the university cafeterias. We used two mechanisms to cope with this challenge. First, we offered students a reward of NIS 200 (\$50) for participating in the experiment. Students were notified that collecting the reward was contingent upon completion of both questionnaires and upon "active participation in forum discussions and in the preparation of the joint document." We did not stipulate any concrete quantitative requirements for "active participation," and ultimately all participants who demonstrated even a minimal level of activity received the reward. Second, because personal sense of efficacy has been shown to be an important component of the motivation to participate, we promised participants

¹⁴In addition to BIU: Ben Gurion University, University of Haifa, Hebrew University, Technion, and Tel Aviv University. The only institute without a code is the Weizmann Institute of Science.

¹⁵Article five of the Act states that "Without derogating from any other right stated by law, each student has the freedom to express his opinions, positions and views with regard to the content of study materials and the values expressed in them; nothing in this provision shall limit the capacity of an academic institution to regulate the way in which opinions, positions and views are expressed in order to protect the orderly function of teaching." Article six states that "Every student has the right to organize and demonstrate on any subject and topic, including on subjects pertaining to students and their rights, according to the rules governing this issue, included in the institution's code." The Student Rights Act gives Israeli universities significant discretion regarding the regulation of political activities.

that at the end of the experiment we would submit the joint document prepared by the group, without the students' names, to the Dean of Students, who is responsible for implementing the Code. In this way, we tried to give students the feeling that their voice would be heard beyond the context of the experiment. We indeed delivered on this promise, and after the experiment ended, we submitted a copy of the texts of all the four groups to the Dean and held a discussion with him about the important points included in the documents.

4. *Soft human intermediation.* The platform was designed to allow participants to express their opinions and create a joint document that reflected the opinions of the group, providing *space for minority and dissenting views as well*. In contrast to other platforms, such as RegulationRoom, we did not use intensive intermediation to facilitate and mentor effective commenting [59–61]. We did not assist the participants in summarizing the discussion, nor did we intervene in the process of drafting and editing the joint document. We left it to each group to handle the task of moving the comments dispersed throughout the forum into the joint document. Nevertheless, we intervened in three ways to keep the discussion going: by keeping the participants updated about the conversation, alerting them about the timeline, and highlighting topics we thought were not receiving proper attention. The interventions were either posted to the forum or sent to the group from a dedicated email address used by the research team for the experiment. All the interventions were posted simultaneously at all four groups.
 - a. *Automatic interventions:* participants received an email when someone commented on their post. Participants also received updates at the end of each day, summarizing important daily activities.
 - b. *Procedural interventions:* these included various reminders (nudges) to participants, e.g., reminding them of the need to draft the joint document, or that the forum was about to close.
 - c. *Substantive interventions:* highlighting topics that did not receive the attention of the group.

3.2 RESULTS

Our findings reflected the potential democratic value of collaborative e-rulemaking. The platform generated significant activity. We found that the quality of the joint documents produced by each of the four groups participating in the experiment was high, both in the number of arguments, which covered various aspects of the Code, and in their epistemic robustness.

Our study provides intriguing results regarding the dynamics of deliberation and the antecedents of participatory action. First, conforming to the projections of critical mass theory, we found that the deliberative process on both the forum and the collaborative platform was dominated by relatively few participants. This finding is consistent with those of studies that looked at Wikipedia and other similar platforms, demonstrating the disproportional influence of a minority of elite editors on the production of content [34, 35, 62, 63]. The implications for deliberative e-rulemaking are mixed. On one hand, it demonstrates the dependence of online deliberation on “digital” leaders. On the other hand, in our case, digital leaders were not so rare as to prevent deliberation. Although the leaders demonstrated a significant presence, our results show that the discussion has also benefitted considerably from contributions from other participants. Our results can be interpreted as a reflection of two forms of leadership: one manifest in the contribution of new ideas to the conversation, the other in the transformation of the discussion on the forum into the joint text. The second form of leadership provides important “editorial services” to the community, without necessarily providing new content.¹⁶

Our study also offers some novel insights into the psychological antecedents of participation. We found that *individualists* were more involved in all activities than were *prosocials*. The apparent influence of SVO may have important policy implications. It may be easier to make participatory schemes appealing to individualists by emphasizing the self-expressive aspects of deliberative platforms. The literature on gamification seems especially promising in this regard [64, 65]. By contrast, making participatory schemes appealing to prosocials may require

¹⁶We found no evidence of a censoring effect, as reflected in references to minority opinions in the joint texts.

organizers to increase the political efficacy of such schemes, which may prove to be difficult because of political and regulatory hurdles. Our study fills a significant lacuna in the literature on e-rulemaking, because much of the research on e-deliberation has tended to be purely observational. A hybrid platform has several advantages over standard consultation platforms (e.g., Regulations.gov). By providing citizens with the opportunity to integrate their comments into a joint document through a collaborative process, the hybrid model creates a more democratic participation scheme, relative to the conventional “notice-and-comment” structure. Another advantage of our platform is that it used a relatively gentle form of intermediation, which made no use of proactive human mentoring (as was used, for example in the RegulationRoom), thus giving the participants greater freedom to shape the deliberation and the joint document. The hybrid model also has some cost-related advantages for the regulator, by eliminating the need for intensive intermediation and by assigning participants the task of summarizing the comment stream, releasing the regulatory agency from this task.

4 THE NEXT STEP: AI AND DEMOCRATIC DELIBERATION

The collaborative system suggested above is not suitable for every participatory setting. For example, using a collaborative tool may be less appropriate for eliciting responses from a broad and heterogeneous community. Such a setting may call for a different policy response that incorporates a wiki-style instrument as one element of its overall structure. For example, one option could be a multi-track e-commenting system, allowing users to choose between several participatory options, which would include the option to submit short comments at the one end of the spectrum (using a structured template) and the option of participating in a wiki-style deliberation at the other end. This multi-track model is based on the idea of “punctuated citizenship,” which recognizes that citizens can freely move between different states of political alertness. These states range from that of “passive follower” (where agents rely on various heuristics to make political decisions), to one of “like/dislike” (reflecting a more active but still epistemologically shallow form of citizenship), and “engaged citizenship” (reflecting a willingness to participate in more complex and demanding dialog) [17, 66]. A multi-track model of participation also conforms better to the multi-dimensional nature of human personality [67]. The wiki-collaborative track can use a random selection tool that allocates contributors to small deliberating groups of 20–25 people. It could also incorporate a voting mechanism that allows participants to choose between various proposals [68].

The concept of wiki-style e-consultation can be used more naturally in smaller-scale participatory processes, situated in urban environments or in professional communities. Two examples are the public consultation processes initiated by the Global Reporting Initiative (GRI) to discuss changes to its reporting standard, and the Co-Cities initiative, which creates a framework for collaborative government for urban communities.¹⁷

Although increasing the use of collaboration tools in e-rulemaking has great potential, it does not remove the main barriers to achieving a more inclusive and sophisticated deliberation. The public good and information overload aspects of deliberative e-rulemaking schemes continue to raise a formidable obstacle to participation. Livermore et al. [10] have argued that agencies should use advanced information processing and text analysis techniques to extract more meaning from the comments they receive. This approach provides a potential solution to what the authors described as the needle-in-the-haystack problem. The technological vision that undergirds the article by Livermore and colleagues focused on the information overload problem of regulatory agencies, but they did not consider the option of using AI to augment the deliberative capacities of citizens. In the concluding part of this article, I consider another option, based on the augmenting potential of AI technology. Citizens can entrust intelligent software agents to participate in e-rulemaking processes on their behalf. Intelligent agents can be defined as “software entities that carry out some set of operations on behalf of a user or another program with some degree of independence or autonomy, and in so doing, employ some knowledge or representation of the user’s goals or desires” [69]. In democratic deliberation, we are interested in a sub-category of such agents:

¹⁷More details on the [GRI consultation platform](https://www.globalreporting.org/standards/Pages/default.aspx) can be found here: <https://www.globalreporting.org/standards/Pages/default.aspx>; more details on Co-Cities can be found here <http://labgov.city/wp-content/uploads/sites/19/Protocol-.pdf>.

autonomous social bots. A social bot is “an *automated social actor*... [a] software designed to act in ways that are similar to how a person might act in the social space” [70]. I propose developing a new generation of social bots: intelligent e-democracy bots. An e-democracy bot can receive as input the political preferences and epistemic views of their principals, and on this basis participate on their behalf in digital consultation processes, exploiting sophisticated AI algorithms. Citizens can then choose in which dialogues to participate themselves, investing greater effort, and in which to entrust this task to their political e-agents. Under this definition, e-democracy bots would operate in the political sphere (e.g., post contributions to a consultation page) as clearly marked agents, without attempting to impersonate human users [71]. In a recent TED talk, César Hidalgo, Director of MIT Collective Learning group, presented a similar idea, referred to as “augmented democracy,” which he defined as using digital twins to expand the ability of people to participate directly in a large number of democratic decisions.¹⁸ A digital twin, software agent, or avatar is loosely defined as a personalized virtual representation of a human. It can be used, according to Hidalgo, to augment the ability of a person to make decisions by either providing information to support a decision, or by making decisions on behalf of that person.

Current social-bot technology is not sufficiently sophisticated to enable delegated participation in the way described above [72].¹⁹ Most bots that currently operate on social media are able only to perform simple functions like posting or reposting prepared content, favoring articles and posts, or following other accounts in an automated manner [72]. In a recent review of open bot technology, Assenmacher et al. identified a considerable gap between the postulated existence of intelligently acting bots, i.e., bots capable of producing original content, and available open software components [72]. But less sophisticated social bots, based on simple heuristics, seem to be within reach of current technology. For example, an environmentalist may provide the following guidance to its e-democracy agent: “Identify the most reputable environmental NGO that took part in the consultation and adopt its opinion, drawing on the introductory text or abstract submitted by that NGO.” To resolve the term “reputable,” the e-democracy bots could draw on various measures of network centrality [73].

The potential use of e-democracy bots raises a host of new challenges. Social-political bots have acquired a bad reputation in the past decade because of a series of malicious interventions in the political domain [74]. Particularly noteworthy were the use of malevolent social bots in the U.S. elections in 2016 [75], in the run-up to the French presidential election in 2017 [76], and in the period preceding the Brexit poll [77]. Political bots can hamper the democratic process through the propagation of false information, through demobilization of opposition, and by attacking the reputation of political candidates [74, 78]. Yet, this justified critique tends to underestimate the capacity of political bots to make a positive contribution to democracy, for example, by extending the capacity of boundedly rational citizens to participate in e-rulemaking processes [79].

The use of e-democracy bots raises various challenges that Hidalgo appears to disregard in his augmented democracy talk. First, we need a reliable authentication mechanism that would enable agencies to distinguish between legitimate bots, entrusted by authentic users, and malicious ones that are not associated with authentic citizens. Such authentication mechanism is critical to preserving the integrity of the e-deliberation process; discussing possible solutions to this problem is, however, beyond the scope of this article [80, 81].

Allowing people to entrust the deliberation task to artificial agents raises additional risks, however, even if we assume that the system is able to authenticate the software agents. The first risk arises from the fact that the use of software agents can make it easier for people to maliciously intervene in the deliberation process (e.g., by spreading false information), simply by lowering the costs of such intervention. A second risk is that the introduction of artificial agents into the deliberation process can deepen the polarization of the commenting space, without at the same time enhancing the epistemic sophistication of the comments. The reason is that to the extent that bots are instructed to faithfully represent their principals’ preferences, but are not, at the same

¹⁸See <https://www.peopledemocracy.com/>.

¹⁹Hidalgo comments on this fact in the FAQ section, id.: “honestly, I think we are far from software agents capable of creating good representations of us and mapping those representations to the goals and intent of policies.”

time constrained by some epistemic criteria, their best strategy could be to mimic an existing contribution that is consistent with their principals' preferences, even if the epistemic basis of the selected contribution is weak. Because bots are likely to lower significantly the barriers for participation, this mimicking strategy can generate a polarizing cascade. Although it is possible to incorporate epistemic constraints into the algorithms that guide the behavior of such bots, the structure of these constraints could be highly controversial, and it is not clear how a requirement of this type would be enforced.²⁰ The risk of polarization is a byproduct of the new human-machine environment that would be created by the introduction of new e-democracy bots [82]. The risks would be even greater if individuals not only entrusted agents to represent them in a given consultation but also allowed them to actively search participatory options.

The introduction of political bots into the e-rulemaking ecosystem is probably inevitable. The risks associated with this futuristic scenario make it necessary to develop a new regulatory framework that would cope with a new political space in which bots and humans interact.

REFERENCES

- [1] M. Lodge and K. Wegrich. 2015. Crowdsourcing and regulatory reviews: A new way of challenging red tape in British government? *Regul. Gov.* 9, 1 (2015), 30–46.
- [2] M. Herz. 2016. E-rulemaking's democratic transformation: Anticipated, actual, and potential. *Revue Int. Gouv. Ouverts* 3 (2016), 195–208.
- [3] J. Stromer-Galley, N. Webb, and P. Muhlberger. 2012. Deliberative e-rulemaking project: Challenges to enacting real-world deliberation. *J. Info. Technol. Politics* 9, 1 (2012), 82–96.
- [4] S. Royo, A. Yetano, and B. Acerete. 2014. E-participation and environmental protection: Are local governments really committed? *Public Admin. Rev.* 74, 1 (2014), 87–98.
- [5] J. Shkabatur. 2011. Cities@ crossroads: Digital technology and local democracy in america. *Brooklyn Law Rev.* 76, 4 (2011), 1413.
- [6] G. D. Jones. 2010. Electronic rulemaking in the new age of openness: Proposing a voluntary two-tier registration system for regulations. *Gov. Admin. L. Rev.* 62, 1261.
- [7] O. Perez. 2015. Fuzzy law: A theory of quasi-legal systems. *Can. J. Law Jurisprud.* 28, 2 (2015), 343–370.
- [8] G. Mucciaroni and P. J. Quirk. 2010. Rhetoric and reality: Going beyond discourse ethics in assessing legislative deliberation. *Legisprudence* 4, 1 (2010), 35–52.
- [9] C. R. Farina, M. J. Newhart, and J. Heidt. 2013. Rulemaking vs. democracy: Judging and nudging public participation that counts. *Mich. J. Environ. Admin. Law*, 2, 1 (2013) 124.
- [10] M. A. Livermore, V. Eidelman, and B. Grom. 2017. Computationally assisted regulatory participation. *Notre Dame Law Rev.* 93 (2017), 977.
- [11] S. W. Shulman. 2009. The case against mass e-mails: Perverse incentives and low quality public participation in U.S. federal rulemaking. *Policy Internet* 1, (2009), 23–53.
- [12] G. Moss and S. Coleman. 2014. Deliberative manoeuvres in the digital darkness: E-democracy policy in the UK. *Brit. J. Politics Int. Rel.* 16, 3 (2014), 410–427.
- [13] Y. Benkler, H. Roberts, R. Faris, A. Solow-Niederman, and B. Etling. 2015. Social mobilization and the networked public sphere: Mapping the SOPA-PIPA debate. *Politic. Commun.* 32, 4 (2015), 594–624, 2015/10/02.
- [14] J. A. Obar. 2016. Net neutrality| closing the technocratic divide? activist intermediaries, digital form letters, and public involvement in FCC policy making. *Int. J. Commun.* 10 (2016), 24.
- [15] L. M. Dinour and A. Pole. 2017. Potato chips, cookies, and candy oh my! Public commentary on proposed rules regulating competitive foods. *Health Edu. Behav.* 44, 6 (2017), 867–875.
- [16] D. R. Reiss. 2018. Regulating in the era of fake news: Anti-vaccine activists respond to the CDC quarantine rule. *Univ. Pitt. Law Rev.* 79 (2018), 675.
- [17] O. Perez. 2013. Open government, technological innovation and the politics of democratic disillusionment: (E-)democracy from socrates to obama. *I/S: J. Law Policy Info. Soc.* 9, 1 (2013), 61–138.
- [18] C. R. Farina, C. L. Blake, M. Newhart, and C. Nam. 2017. Digital democracy in a globalized world. In *Digital Democracy in a Globalized World*, Corien Prins, C. Cuijpers, P. L. Lindseth, and M. Rosina (Eds.). Edward Elgar Publishing, 201–228.
- [19] D. A. Farber and A. J. O'Connell. 2014. The lost world of administrative law. *Tex. Law Rev.* 92 (2014), 1137–2211.
- [20] U. Cress and J. Kimmerle. 2008. A systemic and cognitive view on collaborative knowledge building with wikis. *Int. J. Comput.-Support. Collab. Learn.* 3, 2 (2008), 105–122.

²⁰I mentioned above possible epistemic criteria that could distinguish between comments on the basis of length, the complexity of the language used, and the external prestige of the author, but these criteria are controversial.

- [21] K. R. Parker and J. T. Chao. 2007. Wiki as a teaching tool. *Interdisc. J. Knowl. Learn. Objects* 3 (2007), 57.
- [22] N. J. Klemp and A. T. Forcehimes. 2010. From town-halls to wikis: Exploring Wikipedia's implications for deliberative democracy. *J. Public Delib.* 6, 2 (2010), 4.
- [23] T. Heikka. 2015. The rise of the mediating citizen: Time, space, and citizenship in the crowdsourcing of Finnish legislation. *Policy Internet* 7, 3 (2015), 268–291.
- [24] C. Iaione. 2016. The CO-City: Sharing, collaborating, cooperating, and commoning in the city. *American J. Econ. Sociol.* 75, 2 (2016), 415–455.
- [25] B. S. Noveck. 2009. *Wiki Government: How Technology Can Make Government Better, Democracy Stronger, and Citizens More Powerful*. Brookings Institution Press, New York.
- [26] T. Aitamurto and K. Chen. 2017. The value of crowdsourcing in public policymaking: Epistemic, democratic, and economic value. *Theory Pract. Legis.* 5, 1 (2017), 55–72.
- [27] S. W. Yackee. 2013. Assessing regulatory participation by health professionals: A study of state health rulemaking. *Public Admin. Rev.* 73, s1 (2013), S105–S114.
- [28] C. W. Phang, A. Kankanhalli, and L. Huang. 2014. Drivers of quantity and quality of participation in online policy deliberation forums. *J. Manage. Info. Syst.* 31, 3 (2014), 172–212.
- [29] K. Strandberg. 2015. Designing for democracy? an experimental study comparing the outcomes of citizen discussions in online forums with those of online discussions in a forum designed according to deliberative principles. *Eur. Politic. Sci. Rev.* 7, 3 (2015), 451–474.
- [30] G. Smith, P. John, and P. Sturgis. 2013. Taking political engagement online: An experimental analysis of asynchronous discussion forums. *Politic. Studies* 61, 4 (2013), 709–730.
- [31] S. Himmelroos and H. S. Christensen. 2014. Deliberation and opinion change: Evidence from a deliberative mini-public in finland. *Scand. Politic. Studies* 37, 1 (2014), 41–60.
- [32] D. Balliet, C. Parks, and J. Joireman. 2009. Social value orientation and cooperation in social dilemmas: A meta-analysis. *Group Process. Intergroup Rel.* 12, 4 (2009), 533–547, 1, 2009.
- [33] N. B. Peddibhotla and M. R. Subramani. 2007. Contributing to public document repositories: A critical mass theory perspective. *Organiz. Studies* 28, 3 (2007), 327–346, 1, 2007.
- [34] T. Graham and S. Wright. 2014. Analysing 'super-participation' in online third spaces. In *Analyzing Social Media Data and Web Networks*, M. Cantijoch, R. Gibson, and S. Ward (Eds.). Palgrave Macmillan UK, 197–215.
- [35] T. Graham and S. Wright. 2014. Discursive equality and everyday talk online: The impact of superparticipants. *J. Comput.-Mediat. Commun.* 19, 3 (2014), 625–642.
- [36] M. Xia, Y. Huang, W. Duan, and A. B. Whinston. 2012. Research note-to continue sharing or not to continue sharing? An empirical analysis of user decision in peer-to-peer sharing networks. *Info. Syst. Res.* 23, 1 (2012), 247–259.
- [37] S. W. Yackee. 2015. Participant voice in the bureaucratic policymaking process. *J. Public Admin. Res. Theory* 25, 2 (2015), 427–449.
- [38] S. Coleman and J. G. Blumler. 2009. *The Internet and Democratic Citizenship: Theory, Practice and Policy*. Cambridge University Press.
- [39] N. Sun, P. P.-L. Rau, and L. Ma. 2014. Understanding lurkers in online communities: A literature review. *Comput. Hum. Behav.* 38, 110–117, 9//, 2014.
- [40] A. D. Santana. 2013. Virtuous or vitriolic. *Journ. Practice* 8, 1 18–33, 2014/01/02, 2013.
- [41] I. Rowe. 2014. Civility 2.0: A comparative analysis of incivility in online political discussion. *Info. Commun. Soc.* 18, 2 (2014), 121–138, 2015/02/01.
- [42] E. S. Smith and A. Bressler. 2013. Who taught you to talk like that?: The university and online political discourse. *J. Politic. Sci. Edu.* 9, 4 (2013), 453–473, 2013/10/01.
- [43] R. F. Jørgensen. 2006. *Human Rights in the Global Information Society*. MIT Press.
- [44] M. McCluskey and J. Hmielowski. 2012. Opinion expression during social conflict: Comparing online reader comments and letters to the editor. *Journalism* 13, 3 (2012), 303–319.
- [45] A. Chadwick. 2009. Web 2.0: New challenges for the study of e-democracy in an era of informational exuberance. *I/S: J. Law Policy Info. Soc.* 5, 1 (2009), 9–41.
- [46] R. Lindner and U. Riehm. 2011. Broadening participation through e-petitions? an empirical study of petitions to the german parliament. *Policy Internet* 3, 1 (2011), 1–23.
- [47] O. Perez, J. Bar-Ilan, T. Gazit, N. Aharoni, Y. Amichai-Hamburger, and J. Bronstein. 2018. The prospects of e-democracy: An experimental study of collaborative E-rulemaking. *J. Info. Technol. Politics* 15, 3 (2018), 278–299.
- [48] P. Van Lange, E. De Bruin, W. Otten, and J. A. Joireman. 1997. Development of prosocial, individualistic, and competitive orientations: Theory and preliminary evidence. *J. Personal. Soc. Psychol.* 73, 4 (1997), 733.
- [49] S. Bogaert, C. Boone, and C. Declerck. 2008. Social value orientation and cooperation in social dilemmas: A review and conceptual model. *Brit. J. Soc. Psychol.* 47, 3 (2008), 453–480.
- [50] P. Van Lange, M. Schippers, and D. Balliet. 2011. Who volunteers in psychology experiments? An empirical review of prosocial motivation in volunteering. *Personal. Individ. Diff.* 51, 3 (2011), 279–284.

- [51] J. A. Joireman, T. P. Lasane, J. Bennett, D. Richards, and S. Solaimani. 2001. Integrating social value orientation and the consideration of future consequences within the extended norm activation model of proenvironmental behaviour. *Brit. J. Soc. Psychol.* 40, 133 (2001).
- [52] M. Van Vugt, P. A. Van Lange, and R. M. Meertens. 1996. Commuting by car or public transportation? A social dilemma analysis of travel mode judgements. *Eur. J. Soc. Psychol.* 26, 3 (1996), 373–395.
- [53] B. Sütterlin, T. A. Brunner, and M. Siegrist. 2013. Impact of social value orientation on energy conservation in different behavioral domains. *J. Appl. Soc. Psychol.* 43, 8 (2013), 1725–1735.
- [54] J. Preece, B. Nonnecke, and D. Andrews. 2004. The top five reasons for lurking: Improving community experiences for everyone. *Comput. Hum. Behav.* 20, 2 201–223.
- [55] Y. Amichai-Hamburger, T. Gazit, J. Bar-Ilan, O. Perez, N. Aharoni, J. Bronstein, and T. Sarah Dyne. 2016. Psychological factors behind the lack of participation in online discussions. *Comput. Hum. Behav.* 55, Part A, 268–277, 2//, 2016.
- [56] M. Aristeidou, E. Scanlon, and M. Sharples. 2017. Profiles of engagement in online communities of citizen science participation. *Computers in Human Behavior* 74 (2017), 246–256. DOI: <https://doi.org/10.1016/j.chb.2017.04.044>
- [57] M. Kube, D. Hilgers, G. Koch, and J. Fuller. 2015. Explaining voluntary citizen online participation using the concept of citizenship: An explanatory study on an open government platform. *J. Bus. Econ.* 85, 8 (2015), 873–895.
- [58] R. E. Kraut, P. Resnick, S. Kiesler, M. Burke, Y. Chen, N. Kittur, J. Konstan, Y. Ren, and J. Riedl. 2012. *Building Successful Online Communities: Evidence-based Social Design*. MIT Press.
- [59] D. Epstein, C. Farina, and J. Heidt. 2014. The value of words: Narrative as evidence in policy making. *Evid. Policy: J. Res. Debate Pract.* 10, 2 243–258, //, 2014.
- [60] D. Epstein and G. Leshed. 2016. The magic sauce: Practices of facilitation in online policy deliberation. *J. Public Delib.* 12, 1 2016.
- [61] C. R. Farina, D. Epstein, J. B. Heidt, and M. J. Newhart. 2013. Regulation room: Getting “more, better” civic participation in complex government policymaking. *Transform. Gov.: People Process Policy* 7, 4 (2013), 501–516.
- [62] A. Halfaker, O. Keyes, and D. Taraborelli. 2013. Making peripheral participation legitimate: Reader engagement experiments in Wikipedia. In *Proceedings of the Conference on Computer Supported Cooperative Work*. 849–860.
- [63] D. M. Wilkinson. 2008. Strong regularities in online peer production. In *Proceedings of the 9th ACM Conference on Electronic Commerce*. 302–309.
- [64] S. H. Hsu, J.-W. Chang, and C.-C. Lee. 2013. Designing attractive gamification features for collaborative storytelling websites. *Cyberpsychol. Behav. Soc. Netw.* 16, 6 (2013), 428–435.
- [65] S. L. Kimbro. 2015. Increasing online engagement between the public and the legal profession with gamification. Available at SSRN 2586150 (2015).
- [66] S. Coleman. 2011. Making the e-citizen: A socio-technical approach to democracy. *Connecting Democracy: Online Consultation and the Future of Political Communication* S. Coleman and P. Shane, eds. MIT Press, Boston, 379.
- [67] O. Perez. 2003. Electronic democracy as a multi-dimensional praxis. *N. Car. J. Law Technol.* 4, 275–306, 2003.
- [68] C. Premat. 2018. Can the french republic be digital? lessons from the last participatory experience on the law-making process. *Policy Analytics, Modelling, and Informatics: Innovative Tools for Solving Complex Social Problems*, J. R. Gil-Garcia, T. A. Pardo and L. F. Luna-Reyes, eds. Springer International Publishing, Cham, 247–264.
- [69] S. Franklin and A. Graesser. Is It an agent, or just a program?: A taxonomy for autonomous agents. In *Proceedings of the International Workshop on Intelligent Agents III Agent Theories, Architectures, and Languages*. 21–35.
- [70] N. Abokhodair, D. Yoo, and D. W. McDonald. 2015. Dissecting a social botnet: Growth, content and influence in Twitter. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work and Social Computing*. ACM, 839–851.
- [71] A. Murgia, D. Janssens, S. Demeyer, and B. Vasilescu. Among the machines: Human-bot interaction on social Q&A websites. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1272–1279.
- [72] D. Assenmacher, L. Adam, L. Frischlich, H. Trautmann, and C. Grimme. 2019. Openbots. *Arxiv Preprint Arxiv:1902.06691*.
- [73] O. Perez, R. Cohen, and N. Schreiber. 2019. Governance through global networks and corporate signaling. *Regul. Govern.* 13, 4 (2019), 447–469.
- [74] S. C. Woolley. 2016. Automating power: Social bot interference in global politics. *First Monday* 21, 4 (2016).
- [75] A. Bessi and E. Ferrara. 2016. Social bots distort the 2016 U.S. presidential election online discussion. *First Monday* 21, 11 (2016).
- [76] E. Ferrara. 2017. Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday* 22, 8 (2017).
- [77] M. T. Bastos and D. Mercea. 2019. The brexit botnet and user-generated hyperpartisan news. *Soc. Sci. Comput. Rev.* 37, 1 (2019), 38–54.
- [78] S. Zannettou, M. Sirivianos, J. Blackburn, and N. Kourtellis. 2018. The web of false information: Rumors, fake news, hoaxes, clickbait, and various other shenanigans. *Arxiv Preprint Arxiv:1804.03461*.
- [79] W. C. Ellis, J. T. Ripberger, and C. Swearingen. 2017. Public attention and head-to-head campaign fundraising: An examination of U.S. senate elections. *Amer. Rev. Politics* 36, 1 (2017), 30–53.
- [80] V. Augoye and A. Tomlinson. 2018. Analysis of electronic voting schemes in the Real World. In *Proceedings of the UK Academy for Information Systems Conference*. 14. <https://aisel.aisnet.org/ukais2018/14>.

- [81] F. Þ. Hjalmarsson, G. K. Hreiðarsson, M. Hamdaqa, and G. Hjalmtýsson. Blockchain-based e-voting system. In *IEEE 11th International Conference on Cloud Computing (CLOUD'18)*. IEEE, 983–986.
- [82] G. Teubner. 2018. Digital Personhood? *The Status of Autonomous Software Agents in Private Law*. Available at SSRN: <https://ssrn.com/abstract=3177096>.

Received April 2019; accepted July 2019