

Integrating On-demand Fact-checking with Public Dialogue

Travis Kriplean^{1,4}, Caitlin Bonnar¹, Alan Borning¹, Bo Kinney², Brian Gill³

Dept. of Computer Science & Engineering and DUB Group, University of Washington¹,
Seattle Public Library², Dept. of Mathematics, Seattle Pacific University³, Invisible College LLC⁴

ABSTRACT

Public dialogue plays a key role in democratic society. Such dialogue often contains factual claims, but participants and readers are left wondering what to believe, particularly when contributions to such dialogue come from a broad spectrum of the public. We explore the design space for introducing authoritative information into public dialogue, with the goal of supporting constructive rather than confrontational discourse. We also present a specific design and realization of an archetypal sociotechnical system of this kind, namely an on-demand fact-checking service integrated into a crowdsourced voters guide powered by deliberating citizens. The fact-checking service was co-designed with and staffed by professional librarians. Our evaluation examines the service from the perspectives of both users and librarians.

Author Keywords

Deliberation; fact-checking; civic engagement; libraries; Value Sensitive Design

ACM Classification Keywords

H.5.3. Information Interfaces and Presentation (e.g. HCI): Group and Organization Interfaces

INTRODUCTION

Public deliberation can help uncover solutions to public policy questions and help establish the political will to act [14, 20, 22]. But public deliberation is challenging, requiring communicators to consider tradeoffs, listen to others, seek common ground, and be open to change given evidence [47]. Unfortunately, many people do not even want to talk with people holding different views [36]. In addition, current communication interfaces do not generally support the people who are interested. A few interfaces such as OpinionSpace [13] and ConsiderIt [29] have demonstrated that it is possible to create lightweight communication interfaces that encourage constructive interactions when discussing difficult issues. However, such systems also have a significant problem: the contributions often include factual claims, but participants have difficulty identifying which ones are trustworthy.

We describe a new approach for helping discussants decide which factual claims to trust. Specifically, we designed and

deployed a fact-checking service staffed by professional librarians and integrated into a crowdsourced voters guide. Overall, this research approach continues that used in our earlier work [29]: a field deployment of a real system, in a real election, with discretionary users who decide to participate to help with their electoral deliberations. It thus complements other research involving rigorously controlled experiments that, for example, use tasks presented in a laboratory experiment or paid participants. It also serves as an archetype of a more general class of systems that integrate the contributions of professionals and established institutions into what Benkler [4] calls “commons-based peer production.”

Much of our design effort went into crafting an interaction design that let the fact-checking enhance public deliberation rather than diminishing it. For example, one key design element is that the fact-checks are performed at the behest of discussion participants, rather than being imposed from outside. Our results indicate that participants welcomed the librarians as authoritative sources of information. We also found that librarians were willing and qualified to carry out the fact-checking. However, the fact-checking frame, emerging from journalistic practice, did not map smoothly into librarian reference practice. These findings inspire our design implications, where we discuss a shift to *interactive fact-seeking*.

BACKGROUND: FOSTERING PUBLIC DELIBERATION

To help explore design alternatives and evaluate this work, we turn to Value Sensitive Design (VSD), a methodology that accounts for human values in a principled and systematic way throughout the design process [6, 18]. As with prior work involving VSD in the civic realm [5], we distinguish between stakeholder values and explicitly supported values. Stakeholder values are important to some but not necessarily all of the stakeholders, and may even conflict with each other. For example, discussion participants may differently value environmental sustainability and property rights. While we do not build in support for stakeholder values, we provide a venue in which users can act on their stakeholder values.

Explicitly supported values, on the other hand, guide designers’ choices in the creation of the system: here, the values are democratic deliberation, respect, listening, fairness, and civility. Our explicitly supported values are only a small subset of the values at play in public dialogue. Deliberation in particular is just one normative ideal of political communication [3, 8, 16, 48]. Freelon [16] maps out two other prominent communicative ideals: liberal individualism [11] which values self-expression over listening, manifesting in monologues and flaming, while communitarianism [15, 45] emphasizes group identification with like-minded people for mutual support and to better organize against groups with whom they

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 ACM 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.

CSCW’14, February 15 - 19 2014, Baltimore, MD, USA.

Copyright 2014 ACM 978-1-4503-2540-0/14/02\$15.00.

<http://dx.doi.org/10.1145/2531602.2531677>

disagree. With respect to our explicitly supported values, unfortunately most commonly used communication interfaces, especially online comment boards, *implicitly* support liberal individualist and communitarian values through the particular interface mechanisms they provide. Kriplean et al. found that these interfaces make it far easier to speak into than to actively listen through creating an imbalanced communication model [30]. Because communicative behaviors are context sensitive and can be altered by interface design [35, 38, 44], we believe we can design lightweight interfaces that gently nudge people toward finding common ground and avoiding flame wars, thus supporting a deliberative environment while still accommodating some stakeholder values aligned with these other communicative paradigms (such as advocating for a belief).¹

We implemented ConsiderIt as an approach to this challenge [29]. In ConsiderIt, participants are first invited to create a pro/con list that captures their most important considerations about the issue. This encourages users to think through tradeoffs — a list with only pros or cons is a structural nudge to consider both sides. Second, ConsiderIt encourages listening to other users by enabling users to adopt into their own pro/con lists the pros and cons contributed by others. Third, ConsiderIt aggregates users' expressions of their primary considerations to help everyone better understand and explore facets of the issue being deliberated; they can even view the most important pros and cons for different voter segments (e.g. supporters), enabling users to find common ground. Finally, a per-point discussion facility was added in the 2011 deployment to help participants drill down into a particular pro/con point and have a focused conversation about it. We call special attention to this functionality because one component of our evaluation is an examination of how the focused discussion progressed before and after fact-checks.

Our primary deployment of ConsiderIt is the Living Voters Guide (LVG). The LVG is an ongoing collaboration started in Spring 2010 between our research team and the non-partisan Seattle CityClub, with the goal of inventing and piloting a new method for facilitating civic engagement [17, 28]. We identified deliberation on Washington State ballot measures as a high need area. In 24 states, these ballot measures can be added to the ballot directly by citizens (with the sufficient number of registered voters' signatures), while a legislative referendum process is allowed in all states. These measures are frequently hotly contested, such as Proposition 8 on same-sex marriage in California's 2008 election. They are also often opaquely worded and confusing. For example, in Washington's 2010 election, there were two irreconcilable measures for privatizing the sale of liquor. Unfortunately, there are few avenues, online or offline, for citizens to work through the various arguments and claims being made by campaigns and pundits, or hear the considerations of "everyday" people.

¹There have been many efforts at to create deliberation systems, most prominently argumentation visualization techniques (see [41, 42] for an overview). But these systems tend to require a high level of learning and investment from users and administrators [24, 40], which limits their scope and sustainability. We focus on interfaces that can be easily learned and used widely, even if they lack the semantic richness of these other approaches.

We have run the LVG for the past three elections in Washington State, with 30,000 unique visitors from over 200 Washington cities using LVG for nearly ten minutes on average. Our analysis of behavioral data has revealed a high degree of deliberative activity; for example, 41.4% of all submitted pro/con lists included both pros and cons [17, 28, 30]. Moreover, the tone of the community discussion has been civil: although CityClub actively monitored the site for hate speech and personal attacks, less than ten out of a total of 424 comments have been removed over the three deployments. We encourage readers to visit <http://livingvotersguide.org> to better understand the design of ConsiderIt and the LVG application.

Our summative user studies also revealed the paramount limitation of the original LVG platform (and the motivation for the work reported here): in a crowdsourced voters' guide, participants have difficulty understanding what information to trust. Content analysis of the pro/con points in 2010 found that around 60% contained a verifiable statement of fact, such as a claim about what the ballot measure would implement, or included a reference to numerical data from an external source. Anyone can state a claim, but how do others know whether that claim is accurate?

GROUNDING DIALOGUE WITH AUTHORITATIVE INFO

Before delving into the design of our embedded, on-demand fact-checking service, we sketch the broader design space and touch on design-oriented related work.

What: Framing Authoritative Information

A basic design question for introducing authoritative information is its provenance and the format in which it is delivered.

Primary sources. Authoritative information can be provided directly using primary sources. However, suitable primary sources are often unavailable, and most deliberating bodies do not have the ability to commission a report from a dedicated organization like the Congressional Budget Office. Another option is to find experts in the relevant field, or advocates on either side of a public policy question, who are willing to discuss the subject. For example, face-to-face deliberation methodologies like Consensus Conferences and Citizen Juries incorporate expert Q/A as part of the process by which citizens come together to learn about, discuss, and jointly craft a policy recommendation [14]. Unfortunately, scalability is challenging because it is difficult to secure and coordinate time commitments of experts.

Fact-checking. Fact-checkers produce an evaluation of verifiable claims made in public statements through investigation of primary and secondary sources. A fact-check usually includes a restatement of the claim being investigated, some context to the statement, a report detailing the results of the report, and a summative evaluation of the veracity of the claim (e.g., Politifact's "Truth-O-Meter" ranging from "True" to "Pants-on-Fire"). Fact-checks provide a wealth of information and explicit guidance to the reader about whether to believe the claim or not. Unfortunately, establishing the legitimacy of fact-checks can be challenging because the format of a fact-check juxtaposes the claim and the fact-check, creating

an environment of competing claims [19]. This is problematic because those whose prior beliefs are threatened by the result of the fact-check are psychologically prone to dismiss counter-attitudinal information and delegitimize the source of the challenging information [31, 33, 34, 46], sometimes even strengthening belief in the misinformation [37].

Secondary source synthesis. Another approach is to synthesize information available in reliable secondary sources. This differs from fact-checking in that (1) the investigation does not create new interpretations of original sources and (2) the report does not explicitly rate the veracity of the claims. This approach may help reduce the confrontational nature of fact-checking, while still making information ready-at-hand (though less easily assessable by participants).

Who: Creating and Introducing Authoritative Information

To help ground dialogue, someone needs to identify, synthesize, and introduce the information. Design considerations regarding who does this work and how the process is organized impact the resulting trust in the information and the sustainability of the service.

Professional. A number of professions have developed practices for identifying, synthesizing, and introducing authoritative information. For example, modern political fact-checking emerged from the journalism profession [10], and many current fact-checking enterprises are housed in news organizations (e.g., *Tampa Bay Times* runs Politifact). The journalistic roots of fact-checking are apparent in the emphasis on investigating original sources and publishing an evaluation as a breaking report. Other professionals with relevant skills are reference librarians. One of the main roles of librarians is to help patrons find the information they seek amidst a sometimes overwhelming amount of source material. Librarians assess the content of resources for accuracy and relevance, determine the authority of these resources, and identify any bias or point of view in the resource [2]. As we explore in this paper, librarians can help introduce authoritative information into public dialogue, even though they typically do not provide their services outside of libraries.

Crowdsourcing. Crowdsourcing can also help identify, synthesize, and introduce authoritative information. Both of the aforementioned techniques (fact-checking and secondary source synthesis) have been crowdsourced. Reddit has crowdsourced political fact-checking, where members post claims, the community researches their accuracy, and forum moderators read through the discussion and evaluate the accuracy based on community consensus. Secondary source synthesis is crowdsourced by Wikipedia. While traditional encyclopedias rely on subject-matter experts, Wikipedia relies on many contributors' synthesis of secondary sources, encoded by Wikipedia policies like "No Original Research" [43]. Establishing trust in the results of these crowdsourced efforts is often a challenge [7, 9], but can be accomplished to some extent by transparency of process [27].

Where: Connecting Authoritative Info with Dialogue

Authoritative information can be tailored and introduced into a discussion in a number of ways.

Via participant link to dedicated entity. Authoritative information is usually published and promoted by dedicated entities. For example, fact-checking is often provided as a stand-alone service, as with Snopes, Politifact, and factcheck.org. The authoritative information is thus not explicitly tied to any particular ongoing discussion about the wider issue; only if discussants link to the resource does it potentially help ground a discussion. Even if participants are motivated to do so, other participants may call into question the motives of the individual who provided the link, the reliability of the source, the comprehensiveness of the given resource, and its applicability to the discussion.

Via embedded response. Professionals facilitating a discussion can shepherd authoritative information directly into a discussion. The authoritative information might be researched and introduced at the request of participants, or at the discretion of the professionals. We are unaware of any socio-technical system that embeds a fact-checking service into an ongoing dialogue. However, the aforementioned resource-intensive face-to-face deliberative methodologies often employ experts to help inform participants during a dialogue.

Via personal overlay. Systems like Dispute Finder [12] and Hypothes.is annotate claims at a sentence-by-sentence granularity at the point of consumption as users browse the web. This line of system building focuses on helping individuals be more responsible consumers of information they encounter, rather than grounding an ongoing discussion.

A FACT-CHECKING SERVICE FOR PUBLIC DIALOGUE

In this section, we present our primary socio-technical contribution: a design for introducing authoritative information by a trusted institution into ongoing public dialogue.

As discussed in the previous section, there are a number of potential institutional resources for a fact-checking service. We suggest that libraries and librarians have qualities that make them compelling for this activity (and more generally, as an institutional home for voter education and dialogue forums). The professional practice of librarianship, as embodied in the ALA's Code of Ethics [1], emphasizes the role of libraries in a democracy: "In a political system grounded in an informed citizenry, we are members of a profession explicitly committed to intellectual freedom and the freedom of access to information. We have a special obligation to ensure the free flow of information and ideas to present and future generations." The Code explicitly calls out such values as equitable access; accurate, unbiased, and courteous responses to all requests; and a clear distinction between personal convictions and professional duties. Further, in an era characterized by declining trust in most of our society's established institutions, public libraries continue to enjoy a high level of trust and respect by the public [25, 49]. Reference librarians in particular have well-developed techniques for finding, assessing, and communicating about reliable and useful sources of information [2]. For example, standard practice in addressing a request from a patron for information is the reference interview — as stated in the Reference & User Services Association guidelines, "A librarian treats the user as a collaborator

and partner in the information seeking process.” However, libraries and reference librarian practice have limitations, such as a prohibition on doing legal research. This is significant because the ballot measures being fact-checked in LVG have legal status and implications.

After the 2011 election we established a partnership with the Seattle Public Library (SPL) to provide on-demand fact-checking for the LVG. Any LVG user could request that a librarian fact-check any pro or con point contributed by another user on the site. Librarians reviewed these requests, identified verifiable claims, researched the accuracy of those claims, and posted a public evaluation and report about each claim. SPL sees its participation in the LVG as an opportunity to reimagine and expand its role in serving the public.

The LVG fact-checking service was co-designed by our team and the participating SPL librarians. Librarians made recommendations regarding the functionality of the system, provided feedback on iterations of the design, and wrote web copy explaining the service to LVG users. As part of this preparation, they studied the pro/con points submitted in 2011, created criteria for the kinds of claims they would be able to fact-check, and wrote sample responses to a set of hypothetical fact-check requests based on the 2011 points. The specific guidelines that the librarians settled on were: (1) We will not conduct in-depth legal or financial analysis, but we will point users to research that has already been conducted; (2) We will not evaluate the merits of value or opinion statements, but we will evaluate their factual components; (3) We will not evaluate the likelihood of hypothetical statements, but we will evaluate their factual components.

Following a VSD approach, we consider both direct and indirect stakeholders in our analysis of the fact-checking service. Direct stakeholders are the users of the technology or its outputs. For this work, we call out the following direct stakeholders: (1) users of the Living Voters Guide, (2) authors of points that are fact-checked and (3) the reference librarians. We separate out the authors of points that are fact-checked because one can imagine a range of reactions on the part of someone whose contribution was fact-checked. While there are other possible categories of direct stakeholders (such as users who request the fact-checks), we suggest that the above list includes the most significant cases.

Indirect stakeholders are people or organizations that are significantly affected by the technology even though they do not use it directly. We call out the Seattle Public Library as an institution as an important indirect stakeholder because of our focus on the librarians’ role as fact-checkers, and since this activity represents a significant extension of its traditional role. There are numerous other indirect stakeholders, including voters in general, campaigns with interests in the measures being discussed, and government officials in charge of running the elections.

We now describe the fact-checking service’s sociotechnical design from the perspective of the direct stakeholders, and identify value tensions implicated by this design.

From the Perspective of Living Voters Guide Users

The primary design tension we faced was enabling LVG users to easily get a sense of which factual claims in the pro/con points were accurate (in support of deliberative values), while not making the fact-checking a confrontational, negative experience that would discourage contributors from participating again (or at all). This is a delicate design task because having one’s political beliefs challenged can actually reinforce the original belief. For example, Garrett [19] shows that even introducing a small delay between showing a false claim and a correction of the claim to someone whose prior beliefs are bolstered by the false claim significantly improves the participants’ receptivity to updating their prior belief and reduces their likelihood of rejecting the legitimacy of the fact-check. We now step through a few of the difficult design choices arising from this tension.

When: Origination of the Request

The service was on-demand: any registered user could request a fact-check, submitted with a brief description of what he or she wanted to have checked. The primary design alternative, as implemented by sites like Politifact, would have been to have librarians select which pro/con points to fact-check. By relying on LVG participants themselves to initiate a fact-check, we hypothesized that the degree of confrontation with an authority would be diffused. Further, we hoped that requests would come from supporters of the point to be checked, not just opponents — for example, a supporter (or even the author) might request a check as a way of bolstering the point’s credibility. This design decision is consistent with the values of librarians, as they provide information in response to requests rather than imposing it into discussion. It also enabled librarians to focus on questions that citizens actually cared about, potentially reducing their load and improving the sustainability of the service.

What: Substance of the Fact-Check

Each fact-check comprised (1) a restatement of each factual claim in the pro or con, (2) a brief research report for each claim, and (3) an evaluation of the accuracy of each claim. Dividing points into multiple claims provided an opportunity to address questions directly and reduce confrontation if some of the claims could be marked as accurate. Selecting the categories and labels for accuracy was difficult. We settled on a simple scheme of “accurate,” “unverifiable,” and “questionable.” Each of the evaluation categories was accompanied by a symbol representing the result: a checkmark for accurate, an ellipsis for unverifiable, and a question mark for questionable. The wording “questionable” was chosen to lessen the potential confrontational nature of the evaluation (e.g. in contrast to “wrong” with a red X symbol).

Where: Presentation of Results

There are two primary levels for communicating the result of the fact-checks: (1) when users are browsing multiple points at once and see only the summary of the point, choosing which ones to include in their list, and (2) when viewing the details of a pro/con point and discussing it.

The evaluations of claim accuracy primarily helped users quickly assess general trustworthiness as they browsed the

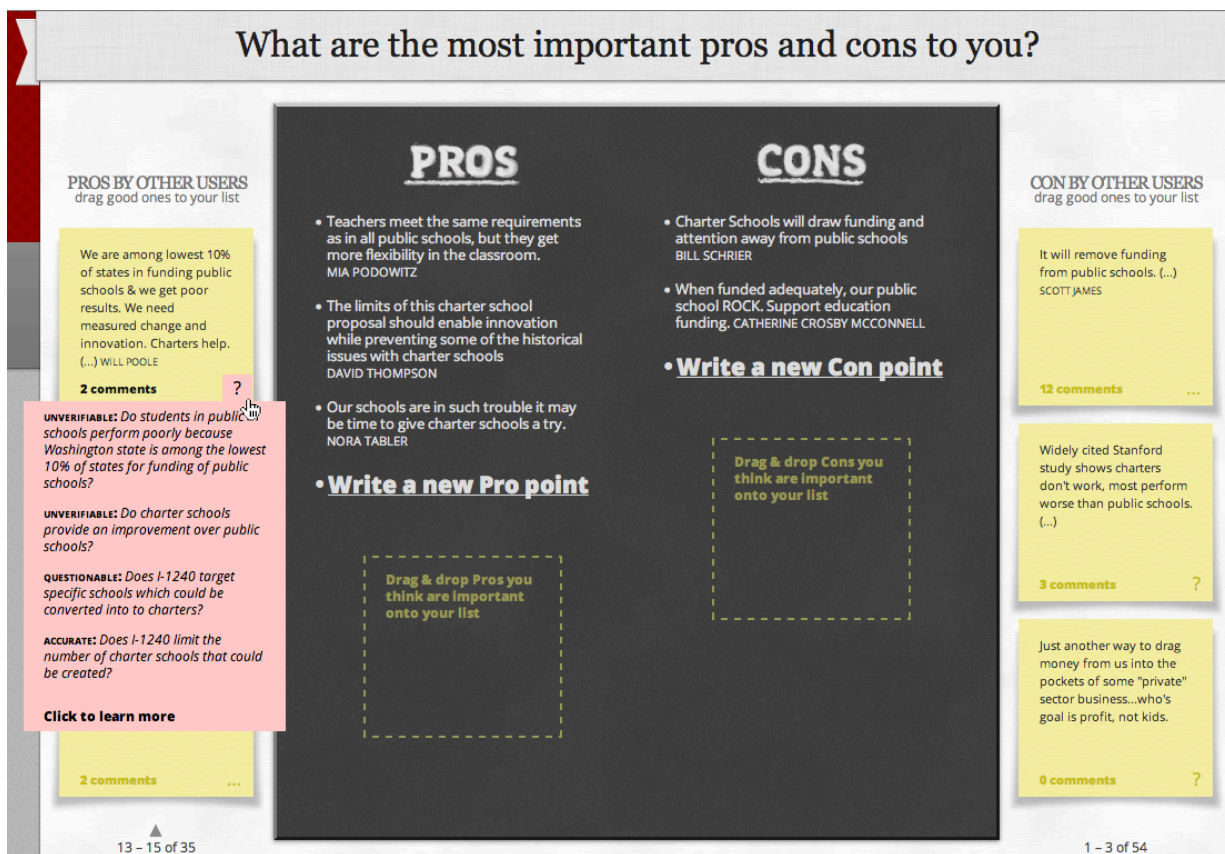


Figure 1. Users can browse the pros and cons contributed by other users while crafting their own pro/con lists. A summary verdict of each fact-checked point is shown in the lower-right corner. A question mark denotes that the point contains at least one inaccurate claim; an ellipsis means that the lowest rating is unverified; a checkmark indicates that all claims are accurate. Upon hovering on the icon, all the claims and their ratings are shown.

points. For each fact-checked pro or con, an icon representing the *most negative* evaluation was shown at the bottom. When a user hovered over the icon, a fact-check summary was shown (Figure 1). We hoped that putting a question mark for points that contained inaccurate claims would not be perceived as exceptionally confrontational to point authors, while being intriguing enough for general users that they would probe into what the question mark meant.

The full fact-check was presented immediately below the text of the point when users drilled into the point details and discussion page (Figure 2). This made the fact-check easily accessible, at the expense of not integrating well into the ongoing dialogue about that point.

From the Perspective of SPL Librarians

While high-level SPL management participated in planning the fact-checking service, the service was implemented by eight reference librarians and managed by a single librarian who also helped perform fact-checks. Librarians took shifts to cover each day during the three weeks leading up to election that the service was deployed. We initially suggested a prioritization scheme for the fact-check requests, but in keeping with the professional values of equitable service policies and access [1] the librarians chose instead to deal with all the requests, in the order received, and committed to responding within a 48-hour period (which they did).

Every fact-check request generated a notification e-mail to the librarians. Librarians then logged into a custom fact-checking dashboard. The dashboard listed each point where a user had requested a fact-check, showing the fact-check status (e.g. completed), which librarian was taking responsibility for it, and the result of the fact-check (if applicable). Librarians could claim responsibility for new requests, and then conduct the fact-check.

The fact-checking page enabled librarians to restate each factual claim made in a pro or con as a researchable and verifiable question, and then answer the research question. The librarians would spend up to two hours researching each of these claims and writing a brief report on their findings (preferably including references to the information sources drawn upon). Finally the librarians would evaluate whether each claim was “accurate,” “unverifiable,” or “questionable.” The fact-checking team decided that the librarians should always identify as many researchable, verifiable questions as a pro or con point contained, even if the request was only for a very specific claim to be checked. This helped scale the service because each point then only had to be checked once.

Every fact-check was a collaboration between at least two librarians. One librarian would write the initial fact-check, which was then reviewed by a second librarian before being published. Additionally, since they worked in close proxim-

Charters can and do deny enrollment to anyone they wish, while claiming public funds needed desperately to serve the majority of kids.

LIBRARIAN'S FACT CHECK

? **Could charter schools deny enrollment to anyone they wish?** *Questionable.*

Our research does not support this claim. According to the Explanatory Statement prepared by the Washington Attorney General (visible above), charter schools "would be open to all students, and could only limit admission based on age group, grade level, or capacity of the school."

We prepared this in response to a user request. We hope the information will be useful to this discussion!

Sincerely,

The Seattle Public Library

[Read about the fact-checking process](#)

Discussion

The article cited discusses for-profit schools, whereas the proposed initiative explicitly requires NON-profit schools which cannot deny admission on any grounds except for age, grade level or capacity

Figure 2. If users open a pro or con point summary, they access a page with additional point details and a discussion space for that point. If that point has been fact-checked, the librarians' report is displayed immediately below the point author's description. In this case, the librarians only identified a single claim and marked it as questionable.

ity, librarians often discussed the fact-checks. This communication facilitated learning and coherence for the service, and also drove functionality changes during the early stage of the pilot.

After each fact-check was published, a notification e-mail was sent out to the requester of the fact-check, the author of the pro/con point, and any other interested party (such as people who were participating in a discussion on that point). Each fact-check was signed collectively as "The Seattle Public Library," not as the individual librarian, in order to protect the librarian from possible harassment. As many of the librarians pointed out in interviews, this was also a way to distinguish personal opinion from a professional response.

METHODS

Following the Value Sensitive Design approach, we used a variety of methods to help us evaluate the fact-checking service from the point of view of each of the main stakeholder groups identified in our conceptual investigations. For users, we relied on three sources. The first was the log files from the LVG deployment. In the 2012 implementation of LVG, this data captures activity by 549 registered users who wrote 332 pro/con points and 424 comments on these points. Seven out of eight of the ballot measures had at least one point that was fact-checked; we therefore constrain our analysis to the 294 pro/con points written for these seven ballot measures. The second was a survey sent out by our collaborator CityClub. The survey asked users about the general usefulness of the LVG and also included a section of six questions specifically about the fact-checking facility. Respondents were asked five questions about the impact and importance of the fact-checking facility. They were scored using a 5-point Likert scale ranging from "strongly disagree" to "strongly agree." The other question asked if they had read at least one of the

librarians' fact-checks. The survey overall received 55 responses (10.7% response rate); out of these, 33 (60%) indicated that they read at least one of the librarians' fact-checks and 14 provided additional comments on the fact-checking facility. The third was another survey, sent to the 40 authors of points that were fact-checked. This survey included both fixed and open-ended questions. We received 16 responses (40% response rate). For the librarians, the second author interviewed seven of the eight reference librarians who did the fact-checking, using a value-oriented semi-structured interview [26]. The interviews ranged from 25 to 60 minutes ($\mu = 36$ minutes). Questions were divided into two categories: individual (e.g., whether the librarians felt qualified to perform this service, and how they felt they represented the SPL's values and roles in the community), and organizational (e.g., how they felt this project affected SPL's reputation, and ways in which this project differed from traditional reference practice). Finally, for the Seattle Public Library as an institution, we interviewed two SPL managers, using the above organizational questions as well as a final question regarding the library's future involvement with the LVG.

Each of these data sources is partial, but together, we believe they help enrich our understanding of the strengths and weaknesses of the fact-checking service, even if the limitations of the data preclude us from making strong generalized claims. The next section weaves this data together to communicate about fruitful next directions.

RESULTS

Our goal is to enhance public dialogue by providing authoritative information from a trusted third party. We first answer a fundamental question: whether or not there was demand for this service. Before we launched, we did not know whether users would be interested in requesting fact-checks, and if




Figure 3. Cumulative time series showing total pro/con points, fact-check requests and completed fact-checks aggregated across all ballot measures that received at least one fact-check. The fact-checking service launched midway on 10/17/2012. Total requests are slightly greater than the completed fact-checks because sometimes several people requested the check before the librarians could publish the fact-check.

they were, would act on that interest. After establishing that there was significant demand for the service, we then analyze three aspects of the goal of enhancing dialogue with authoritative information: first, whether the library was perceived as a legitimate, welcomed, authoritative participant; second, whether the particular way in which authoritative data was introduced helped invite wider, deeper participation; and third, whether the service did in fact inform public dialogue. We undertake this analysis first from the viewpoint of LVG users, and then from the viewpoint of the librarians and the library.

Broad Trends

Demand

When we launched the fact-checking pilot, we did not know whether users would request fact-checks at all, or conversely, whether they would request fact-checks about nearly every point and overburden the librarians.

In this deployment, users demonstrated demand for fact-checks and the willingness to submit requests for them. Through the course of the 3 weeks the fact-checking service was live, SPL librarians responded to 55 requests by 30 distinct users to fact-check 47 distinct pro/con points related to 7 of the 8 statewide ballot measures (Figure 3).

Users thus requested fact-checks of 14.2% of the pros and cons submitted. We consider this a respectable, even high, level of participation, for a number of reasons. First, many pros and cons do not contain checkable claims because they are statements of principle, such as the vast majority of pros and cons submitted to discussion around 2012’s ballot measure for legalizing gay marriage. In an informal analysis, librarians found that approximately half of all submitted pros and cons in 2011 contained claims checkable by their criteria. Second, not all pros and cons are examined with the same degree of scrutiny by other users. For example, some are poorly worded or even incoherent. Because of ConsiderIt’s PointRank algorithm [29], these points rapidly fall to the bottom and are only seen by the most persistent users. Finally, because of an inopportune illness, the fact-checking interface was launched midway through the deployment, necessarily limiting participation. We were pleased therefore that many of the prominent points with factual claims on popular measures had fact-checks requested, suggesting that users were overall parsimonious with their requests.

Though the librarians were challenged at times to respond to the requests, the demand matched capacity. Librarians were

able to respond to each fact-check request within the 48-hour time frame that they promised. In 2012, we were pleased to have hit a sweet spot between capacity and demand. However, in the future, the service could be overwhelmed. One factor mitigating this possibility is a structural resiliency in ConsiderIt that disincentivizes political gaming of the fact-checking service: if a strong advocate decides to request fact-checks of all the points he or she disagrees with, the claims could be evaluated as accurate and end up bolstering the opposing side’s case. The more likely risk in overloading the librarians stems from pranksters operating without a political agenda. This could be handled by triaging requests based on the user requesting the fact-check.

Requesting Fact-Checks

Fact-check requests demonstrated a variety of information needs. According to a post-mortem report by the librarian who managed the fact-checking project, requests usually took one of five forms: (1) a simple paraphrase or copy and paste of the original claim, (2) a request for a specific aspect of the claim to fact-check, (3) a request based on implicit assumptions made within a point, (4) a request on the reliability of a source cited by the author of the pro or con point, or (5) a question that was related to the pro or con point.

We expected that users would request fact-checks of points that opposed their own positions (e.g. a proponent of a ballot measure requesting a fact-check of a con point). This was borne out, with 28 (75.7%) of the requests targeted to position-opposing points, eight requests (21.6%) targeting position-supporting points and one request submitted by a neutral user. (Thirteen requests were made by users who did not register their positions, which were excluded from this analysis.) The 21.6% of requests that targeted position-supporting points is notable because it is potentially evidence of normatively desirable deliberative behavior, with users questioning and probing considerations posted in favor of their own positions. (An alternative explanation is that these requests reflected a desire to bolster the point’s legitimacy in order to better persuade other users, which still seems normatively desirable, in that it is seeking to ground deliberation in fact.)

Librarian Responses

Overall, the librarians demonstrated that they were able to evaluate complex requests and find specific information to help evaluate the accuracy of claims made on the LVG.

Librarians divided the fact-check requests of 47 points into 65 claims. The librarians’ investigations of those claims resulted in 19 claims marked as “accurate,” 34 marked as “unverifiable,” and 12 marked as “questionable.” One reason for the large number of “unverifiable” claims is that, for the political domain, what one might think are straightforward factual questions turn out to be more nuanced on closer examination. Another reason was SPL’s policy about not doing legal research; this meant that claims asking whether the ballot measure actually stipulated some specific change had to be marked “unverifiable,” even though the librarian was often able to find and report specific relevant quotes in the text of the measure. Additionally, about half of the “unverifiable” claims were hypothetical statements. Regardless of whether

the accuracy of the claim could be established, librarians were always able to write an informative response to each fact-check request.

From the Living Voters Guide Users' Perspective

Legitimacy and Welcome

We hypothesized that librarians would be perceived as a qualified, credible source of information and be welcomed by users. To assess these hypotheses, we draw upon our survey results. Users generally agreed with the librarian's analysis and found value in it (8.6% strongly agreed, 65.7% agreed, 25.7% neutral, 0% disagreed, 0% strongly disagreed). People who were fact-checked felt that librarians generally assessed their points in a fair manner (62.5% "fair," 0% "unfair," 37.5% "neither").

Users who were fact-checked were divided on how accurate they found the resulting fact-check to be (53.3% answered "accurate," 6.7% "inaccurate," 40.0% "partly"). All respondents who answered that the resulting fact-check was either "partly accurate" or "inaccurate" received ratings of either "unverifiable" or "questionable," while all who received "accurate" ratings responded that the fact-check on their point was "accurate." Four respondents whose points had been labeled "questionable" responded that they felt the fact-check was "accurate" — an encouraging result.

Users did express a desire for better communication with the librarians. One user, for example, stated:

But you could never ask a second question or counter the fact checking.
... I'd say it was a good service to have, we just didn't seem to know how to properly communicate with it.

Perhaps most importantly, the fact-checking service improved users' perception of the overall credibility of the LVG (65.7% strongly agreed, 25.7% agreed, 8.6% neutral, 0% disagreed, 0% strongly disagreed). Even those who were fact-checked stated they were either more likely to return (56.3%) or neutral (37.5%), suggesting that the fact-checking at least did not harm the group most likely to perceive the service as illegitimate or unwelcome. The extensive positive press coverage that the service received also suggests that the legitimacy of LVG increased. For example, a *Seattle Times* column [21] praised the librarian's contribution to the LVG, stating that "there's something refreshing in such a scientific, humble approach to information."

Inviting of wider, deeper participation

It is difficult to assess the impact of the fact-checks on users' overall propensity to participate. Since the LVG received more pros and cons than in the previous two elections, it did not appear to significantly deter users from contributing new pros and cons; furthermore, the favorable press coverage of the fact-checking service likely attracted additional users. However, we can begin to evaluate the impact fact-checking had on the discussions rooted to each pro/con point, though we cannot make authoritative causal claims.

We examined the data to determine if pro or con points that were fact-checked tended to receive fewer comments after the fact-check compared to points that did not receive a fact-check. Our dataset contains timestamped comments, as well as timestamped views of each pro/con point, by registered

users. We can therefore operationalize the commenting rate as the number of comments per view over a given period of time. We expected the commenting rate to drop slightly after the point was fact-checked because a fact-check should eliminate the need for many of the comments people might make, such as asking a question of the point author or challenging the claims. However, we expected this effect to be partially offset by the fact-check spurring further discussion about the result of the fact-check.

To properly analyze the impact of fact-checks generally on the commenting rate, the statistical analysis needs to control for a major temporal factor: *users in general, even for non-fact-checked points, were less likely to comment on points as the election drew closer*. In other words, LVG users as a whole become more passive consumers of the voters guide as the election neared. This temporal effect also holds for point inclusions, as discussed later.

We therefore conducted a permutation test to assess whether the decrease in the commenting rate from before the fact-check to after was significantly larger than the change that would be expected due to temporal effects alone. If the fact-check had no effect on the commenting rate, then we would expect to see similar before/after differences if we simply randomly reassign the fact-checks and their timestamps to a different set of points. To conduct the permutation test, Monte Carlo simulation was used to repeatedly randomly reassign the 47 fact-checks and their original timestamps to 47 of the 294 points. In the randomization process, each fact-check was assigned at random to a point which had at least one view prior to the timestamp at which the original request for the fact-check was submitted. For each of these new simulated groups of 47 "fact-checked" points, we then computed the before-after difference in comment rates. We repeated this permutation process 10,000 times, recording these differences for each permutation. We then obtained a Monte Carlo estimated p-value for the test by determining the proportion of the permutations that produced differences as extreme or more extreme than the difference observed in our actual data.

In our data, 3.1% of users who viewed a fact-checked point prior to the actual fact check commented on the point; after the fact-check this comment rate decreased to only 0.5%. The permutation test revealed that this before-after decrease in the comment rate for fact-check points was significantly larger ($p=.024$) than would be expected based purely on the temporal effect. The observed decrease in commenting rates before-after fact-checking was an estimated 1.57 times as large as would be expected from temporal effects alone.

One possible explanation is that the fact-check resolved the questions which subsequent commenters might have had. This user's quote conveys this perspective:

I accepted the fact checking as the end of the discussion. Can't fight the facts ... it definitely felt like once something had been fact checked, any speculations about something were put to rest.

However, we need more data to know how common was this view and whether we can attribute causality. At this point there are other possible explanations, based on factors that our statistical test could not control for, such as the partic-

ular persuasive style these points employed. Interestingly, the points that were eventually fact-checked generated significantly more comments before the fact-checking service launched than points that were not fact-checked ($p=.0098$).² By the last week before the election, there was virtually no commenting on any of the points. It is possible that the commenting had largely run its course and would have died out by late October on all points even if there were no fact-checking, and that the sorts of points that had fact-checks had a shorter “commenting half life” — that is, they attracted a stronger initial burst of discussion that then died off more quickly.

Informing Public Dialogue

Though our high-level goal is to foster a deliberative environment in which slow, long-term opinion change might take place, we still want to know whether the librarian’s fact-check made an impact on a person’s decision about which pro or con points were important to their decision, and also whether the fact-check resulted in a more informed decision (even if the stance remained the same). Our first hypothesis was that the fact-checking service would help people make their decisions about whether to vote for a ballot measure or not. Secondly, we expected that points labeled “accurate” would be included in people’s pro/con lists at a higher frequency after the fact-check and that points given a label of “questionable” would be included less frequently.

According to our survey data, general users overwhelmingly self-reported that they thought library fact-checks increased their understanding of the pro/con points as well as helped them become more generally informed about the ballot measure itself (31.4% strongly agreed, 48.6% agreed, 11.4% neutral, 8.6% disagreed, 0% strongly disagreed). The same held true for the users who were fact-checked (68.8% said fact-checks were “informative,” 18.8% “confusing,” 6.3% “neither,” 6.3% gave no response). As one fact-checked user shared, “*I learned things about several issues that I did not know. In one case, my vote was changed.*” One author whose point was fact-checked stated that the fact-check caused him to think more critically about validity of his source:

I found the most interesting part to be that the Librarian fact checker not only checked the fact I submitted, but also the source from whence it came. I had accepted and regurgitated this think-tank’s fact and when the library cautioned that the think-tank source was “unverifiable” in its trustworthiness, this really forced me to think critically about the source of the fact that I had thoughtlessly used.

However, in practice, we found no evidence that the fact-checks changed participants’ likelihood of including a pro or con point, regardless of whether the point contained accurate or inaccurate claims.

²There were 248 points with at least one view prior to the launch of the fact-check system, 44 of which would later be fact-checked. Under a null hypothesis that there is no difference in comment rates for points that would later be fact-checked and points that would not later be fact-checked, we conducted a permutation test in which we randomly assigned 44 of the 248 points as the “fact-checked” points, repeating this process 10,000 times. For each permutation, we computed the difference between the comment rate for the points randomly assigned as “fact-checked” points and the comment rate for the points not assigned to be “fact-checked.” This permutation test showed that there is a significant difference in comment rate for points that would eventually be fact-checked with a Monte Carlo p-value of .0098.

As with the analysis of commenting rates, we employed a Monte Carlo permutation test with 10,000 repetitions of the simulation to answer the question of whether the change in the rate of inclusions before/after fact-checking is different from what would be expected based on temporal effects alone. The Monte Carlo estimated p-value of .160 suggests that the change in inclusion rates is consistent with what would be expected based on temporal effects alone. Moreover, we find no evidence that the evaluation of the factual accuracy affected inclusion rates. Thus, we found no support for our hypotheses regarding the impact of fact-checks on users’ decisions about which points are important.

There are several possible reasons why fact-checks may not have broadly influenced inclusion rates. First, people’s values and self-perception may not be aligned with their actual behavior. And second, opinions may have changed, but only for those whose prior beliefs were not threatened by the fact-check [19]. On the whole, it appears that there was little evident short-term opinion change, at least as exhibited through the adoption of other users’ pro/con points.

From the Librarian’s Perspective

We now analyze the fact-checking service from the perspective of the librarians, and their experience in acting as authorities in the service of public dialogue.

Legitimacy and Welcome

There are three aspects that contributed to the librarians’ view of whether their contributions were perceived as legitimate: (1) if the provision of this service fit within the library’s core mission and goals, (2) if the librarians felt qualified to fact-check, and (3) if the librarians felt welcomed. To answer these questions, we draw on our one-on-one interviews with the participating librarians and managers, and an internal report compiled by the librarian who acted as project manager.

We first address whether the design of the fact-checking facility aligns with the library’s core mission and values, traditional reference practice, and SPL’s goals to innovate. On the innovation side, both librarians and managers alike saw this project as a way to reach out to the community on a larger scale, as well as highlight one of the core services that the library can provide. As one manager put it,

I think we get lots of advantage from putting our answers to people’s questions in a public setting. Foremost of which is that other people can see those answers too. The information is not limited to this one-on-one exchange that we have with a person . . . it starts from that, but you know, we know that we had fifty-some requests that we responded to, but we don’t have any idea of how many other people read those responses and also learned something from that exchange. So that’s something that we don’t get from the traditional experience which I think is really valuable.

Some librarians liked the ability to provide information that might aid voter decisions and felt that they could contribute balanced (if not neutral) information while maintaining their non-partisan stance. However, some librarians raised concerns about the ability of the library to maintain their reputation as a neutral institution. These two viewpoints are epitomized by the following quotes from two different librarians:

I think ultimately what we are doing is helping people find information, and I don’t think . . . we want to strive for absolute neutrality at

the expense of informing people. . . . it's hard, it's a difficult thing to do well, but I do think that it is something that is possible to do.

We're a neutral territory. Anyone could come in here, regardless of your political persuasion and so forth and we try to give you equitable service and so on. I think when we get into the realm of political discourse around campaigns and initiatives and so forth, we can get pulled one way or another, and then become labeled as an advocate for that position or an advocate for that position.

However, despite the concerns about the library maintaining a non-partisan stance, the librarians felt that they were able to keep their own political positions separated from their responses. Some mentioned the fact that they already handle this in their day-to-day question answering, and that they are professionally trained to separate their opinion from their work. Others also described that by reading and discussing each others' fact-checks they were able to further mitigate the risk of injecting their own viewpoints.

Five of the librarians reported that they felt qualified, claiming that providing fact-checks was similar to their everyday reference work. One librarian felt his lack of legal background prevented him from properly evaluating some claims. Other librarians who gave an affirmative response brought this up as well, stating that they had to be more cautious because of the legal implications of many points, but felt that they were overall able to uphold the guidelines. This mix of feelings is well captured here:

In a sense, I did feel qualified but then in [another] sense, I didn't because I didn't have the expertise in government or the types of things that they were looking for in those types of questions. But after reading on the different things that they were trying to pass, all the initiatives, all the measures, and after answering a few of them over time, I felt more and more comfortable and more qualified to be able to answer or point them to different resources that they could look at.

All seven librarians felt qualified to check whether a source was consistent with a claim; however, many cited the fact that in traditional reference practice they leave the final analysis to the patron. Comfort with this varied. Some felt they had the skills and the experience to perform it even if the analysis was not normal. One felt that the "unverifiable" category gave him the opportunity to leave the analysis up to the reader if he did not feel they could provide a definitive answer. Another felt that she should not be conducting any analysis at all.

Lastly, we consider how the librarians regarded their participation and whether they felt welcomed. As previously discussed, users were receptive but wished there were better ways to communicate with the librarians. Librarians shared this desire. Librarians often felt that they needed to make assumptions about what the requester of the fact-check was asking or what the original poster intended. Six of the seven librarians, as well as both managers, mentioned that this was at odds with their day-to-day reference practice where they converse with patrons in order to get at the heart of the question and make sure they are providing the correct information.

. . . sometimes it would have helped to get a little more clarification or to have some of that back and forth. When you have a face-to-face interaction with someone, that is a big part of what we do, sort of a conversation before we start looking for sources is figuring out what is this person really looking for, what is going to be most useful to them.

The lack of a communication mechanism also prevented librarians from knowing how their fact-checks were received by the author of the point, the requester(s) of the fact-check, and anyone reading the fact-check, leading to a general disconnect between librarian and user.

Inviting of Wider, Deeper Participation

All librarians and managers interviewed felt that they provided fact-checks in a responsive, non-discriminatory manner that was inviting to the public. The project manager decided on a 48-hour timeframe for answering a fact-check request. Librarians felt this was reasonable, and were able to complete all fact-checks within that time period. Only a few felt they could have used more time on a particular fact-check. And while the lack of communication mechanisms as discussed above was problematic, librarians sometimes still found ways to respond to concerns. In one instance, a user posted on the LVG's Facebook page a concern with the evaluation that the librarians had made on a fact-check relating to same-sex parenting. The librarian in charge of the fact-check went back and reworded the claim to provide more clarity, as well as adding additional explanation of the assessment. This suggests that librarians were not only willing to listen to user's concerns, but that they also took their suggestions seriously.

Informing Public Dialogue

Both managers and librarians felt that their specific skill sets were conducive to providing reliable and helpful information. As mentioned previously, many librarians see themselves as navigators of the vast amount of information found not only within the library, but also on the Web.

Our job is helping people become informed in whatever realm that means, so certainly we can contribute information to political discussion. . . . I think LVG showed actually several different ways that we can participate in that, you know, just by inserting ourselves into a conversation and saying, 'here is some information that might help,' or . . . just by letting people know that we are available to help them find information is making a contribution to the discussion.

However, some felt that the labels associated with the fact-checks — "accurate," "questionable," and "unverifiable" — could sometimes be misleading. For example, since points sometimes contained multiple claims, librarians broke fact-checks into separate claims, rephrased them as a question and provided a label for each. The interface then displayed the lowest rating associated with the point.

Yeah, I mean, one thing that might have been confusing for the people using it was the fact that you have an original post, which then somebody flags, and then we restate it, and then we reflag it with one of the [labels]. And so when we flag it as one of the three, what we were doing was flagging that in response to the way we restated the question, which worked I think because we were able to very carefully do that. I think if the person wasn't very carefully figuring that out, they could see our flag and read it as back to the original point.

In summary, librarians felt able to provide authoritative, relevant information to the public. They felt that this project was not only a good way to showcase the skills that they possess in terms of providing answers to complex questions, but also a way to reach a wider audience. Their concerns centered primarily around the way the service was manifested on the LVG, which we believe can be mitigated by redesign.

DISCUSSION AND DESIGN IMPLICATIONS

We suggest three takeaways from this work. First, crowdsourcing systems can benefit from reintroducing professionals and institutions that have been up to now omitted. They can provide their professional skills and the authority of their institutions as design assets, for example in support groups for ongoing health conditions in which health care professionals may occasionally participate [23]. Conversely, some of these organizations see engagement in public dialogue and other crowdsourcing systems as a way to help cultivate a new or evolved identity. As socio-technical architects, we should be combining old and new to create entirely new possibilities.

Second, users welcomed the librarians' contributions, even those whose statements were challenged. Our perspective is that correcting widely held, deep misperceptions is something that cannot be quickly fixed (e.g., with a link to a Snopes article), but is rather a long process of engagement that requires a constructive environment [32]. Our initial data indicates that participants were willing to consider the library as an authoritative source of information, even if our data does not suggest that the fact-checks influenced participants' short-term opinions. This willingness to welcome SPL lays the groundwork for long-term opinion change and cultural shifts within an environment for constructive public dialogue.

Third, we found that the journalistic fact-checking frame did not map smoothly to librarian reference practice. Librarianship is rooted in guiding people to sources rather than evaluating claims. Discomfort stepping outside these bounds was magnified by the lack of opportunities for librarians to communicate with users to clarify request and the original point. This points to an evolution of our approach to introducing authoritative information into public dialogue that we call *interactive fact-seeking*. In the remainder of this section, we outline ideas for this new approach.

Interactive Creation. In our implementation, there was a glaring lack of communication among the librarians, the user whose point was being fact-checked, and the fact-check requester. Librarians were hamstrung in knowing with confidence whether they were addressing requesters' needs; they also did not have a way to ask clarifying questions of the author of the point. Participants, on the other hand, could not readily ask followup questions. Straightforward solutions are possible, such as providing mechanisms for private communication between librarians and users, and the ability for users to thank librarians for fact-checks they found useful. A more comprehensive solution might enable the interactive creation of fact-checks. Each fact-check request could open a new forum to which the author of the point, the requester of the fact-check, and anyone else following the point are invited. Librarians can use the forum to ask clarifying questions, post sources they are consulting, and seek feedback. After the fact-check is completed, the forum would serve as a transparent log of the fact-check, helping to establish the legitimacy and participatory nature of the process. The principal challenge with this approach is likely the management of conflicts between the point author and the requester of the fact-check.

Integrate Fact-Checks with Other Comments. Fact-checks are displayed immediately below the pro/con point (Figure

2). While this maximizes visibility, it also highlights competing claims [19], overshadows the discussion section, and obfuscates when the fact-check happened with respect to the discussion. In fact, this could have been a factor in the drop in commenting rates after a fact-check occurred. An interactive fact-seeking approach might situate the fact-check in the flow of the comments as just one part of a larger effort by discussants to ground the discussion and build from it. This design gives librarians the opportunity to frame the fact-checks in a conversational (though professional) manner that maximizes receptivity (e.g. by framing the fact-check in an affirming manner that heightens the chance that people will respect the authoritative information and update their prior beliefs [19, 32]).

Separate Research from Evaluation. Some librarians felt they were overstepping their professional practice when evaluating the veracity of many claims. This problem could be mitigated by creating categories better suited to librarianship, such as evaluating a claim's consistency with reputable sources, or enabling librarians to signal their confidence in the evaluation. However, the problem might be eliminated if the librarians do not evaluate the claims, and instead only seek out and synthesize reliable information on the claim. The evaluation of the claim could be carried out by some other process, such as by an independent panel of trusted community members formed ahead of the election, or crowdsourced by community vote. This latter option is intriguing from an interactive fact-seeking perspective because it invites users to participate in assessing truth, which in turn may encourage opinion change grounded in accurate information because opinion change is most likely to occur if users synthesize and re-articulate information themselves [39].

Encourage Self Fact-Check Requests. Ideally, participants would experience being fact-checked as a learning opportunity and not be embarrassed or offended. While not always possible, the design can do more to foster this fact-seeking attitude. One design direction might encourage participants who are writing a pro or con to actually request a fact-check of themselves, such as by making a "double-check this please" option salient in the authoring process. If it becomes routine to request a fact-check of one's own contribution, the potential confrontational aspect of fact-checks could be lessened.

As of this writing (October 2013), we are in the midst of our second election with the fact-checking service. We implemented many of the above changes in collaboration with SPL librarians, including (1) a facility for the reference librarians to email the fact-check requester or the author of the original point; (2) revised evaluation categories that state whether the claim is consistent with authoritative sources, rather than accurate or inaccurate; and (3) incorporated the results of the fact check chronologically in the comment thread, with a short note at the top that links to the main fact check.

ACKNOWLEDGMENTS

Many thanks to our other collaborators at the University of Washington, Seattle CityClub, and the Seattle Public Library; and in particular to the reference librarians at SPL who took part in the experiment, and to all the Washington voters who

participated by using LVG. This research was supported in part by the National Science Foundation under Grant No. IIS 0966929 and in part by a Google Research Award.

REFERENCES

- American Library Association. Code of ethics of the American Library Association, 2008. Available at <http://www.ala.org/advocacy/proethics/codeofethics/codeethics>.
- American Library Association. Reference and user services association guidelines, 2013. Available at <http://www.ala.org/rusa/resources/guidelines>.
- Anderson, C. Deliberative, agonistic, and algorithmic audiences. *Int. J. of Communication* 5 (2011).
- Benkler, Y. *The Wealth of Networks*. Yale University Press, 2006.
- Borning, A., Friedman, B., Davis, J., and Lin, P. Informing public deliberation: Value sensitive design of indicators for a large-scale urban simulation. In *Proc. 9th European Conference on Computer-Supported Cooperative Work* (Paris, Sept. 2005).
- Borning, A., and Muller, M. Next steps for value sensitive design. In *CHI* (2012).
- Castillo, C., Mendoza, M., and Poblete, B. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, WWW '11, ACM (New York, NY, USA, 2011), 675–684.
- Dahlberg, L. Democracy via cyberspace. *New Media & Society* 3, 2 (2001).
- Denning, P. J., Horning, J., Parnas, D. L., and Weinstein, L. Wikipedia risks. *Communications of the ACM* 48, 12 (2005), 152.
- Dobbs, M. The rise of political fact-checking: How Reagan inspired a journalistic movement: A reporter's eye view. New America Foundation, Feb. 2012. Available at http://newamerica.net/sites/newamerica.net/files/policydocs/The_Rise_of_Political_Fact-checking.pdf.pdf.
- Dryzek, J. *Liberals, Critics, Contestations*. Oxford Press, 2002.
- Ennals, R., Trushkowsky, B., and Agosta, J. M. Highlighting disputed claims on the web. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, ACM (New York, NY, USA, 2010), 341–350.
- Faridani, S., Bitton, E., Ryokai, K., and Goldberg, K. OpinionSpace: A scalable tool for browsing online comments. In *CHI* (2010).
- Fishkin, J. Strategies of public consultation. *Integrated Assessment* 6, 2 (2006).
- Fraser, N. A contribution to the critique of actually existing democracy. *Social Text*, 25/26 (1990).
- Freelon, D. Analyzing online political discussion using three models of democratic communication. *New Media & Society* 12, 7 (2010).
- Freelon, D., Kriplean, T., Morgan, J., Bennett, W. L., and Borning, A. Facilitating diverse political engagement with the living voters guide. *Journal of Information Technology & Politics* 9, 3 (2012), 279–297.
- Friedman, B., Kahn Jr., P. H., and Borning, A. Value Sensitive Design and information systems: Three case studies. In *Human-Computer Interaction and Management Information Systems: Foundations*. M.E. Sharpe, Armonk, NY, 2006.
- Garrett, R. K., and Weeks, B. The promise and peril of real-time corrections to political misperceptions. In *CSCW* (2013).
- Gastil, J. *Political Communication and Deliberation*. Sage, 2008.
- Guzmán, M. Seattle library fact check experiment risky, but valuable. *Seattle Times* (Oct. 2012). See <https://wash.livingvotersguide.org/home/media> for this and other links to press coverage.
- Habermas, J. *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. MIT Press, 1991.
- Huh, J., Patel, R., and Pratt, W. Tackling dilemmas in supporting “the whole person” in online patient communities. In *CHI* (2012).
- Isenmann, S., et al. IBIS: a convincing concept ... but a lousy instrument? In *DIS* (1997).
- Jaeger, P. T., and Fleischmann, K. Public libraries, values, trust, and e-government. *Information Technology and Libraries* 26, 4 (2007), 34–43.
- Kahn Jr., P. H. *The Human Relationship with Nature: Development and culture*. MIT Press, 1999, ch. 5: Structural-developmental methods, 77–93.
- Kittur, A., Suh, B., and Chi, E. H. Can you ever trust a wiki?: Impacting perceived trustworthiness in Wikipedia. In *CSCW* (2008).
- Kriplean, T. *Encouraging Reflective Discussion on the Web*. PhD thesis, Dept. of Computer Science & Engineering, University of Washington, Dec. 2011.
- Kriplean, T., Morgan, J., Freelon, D., Borning, A., and Bennett, W. L. Supporting reflective public thought with ConsiderIt. In *CSCW* (2012).
- Kriplean, T., Toomim, M., Morgan, J., Borning, A., and Ko, A. Is this what you meant? promoting listening on the web with reflect. In *CHI* (2012).
- Kuklinski, J. H., Quirk, P. J., Jerit, J., Schwieder, D., and Rich, R. F. Misinformation and the currency of democratic citizenship. *Journal of Politics* 62, 3 (2000), 790–816.
- Lewandowsky, S., Ecker, U., Seifert, C., Schwarz, N., and Cook, J. Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest* 13 (2012), 106–131.
- Lord, C. G., Ross, L., and Lepper, M. R. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37, 11 (1979), 2098–2109.
- MacCoun, R. J., and Paletz, S. Citizens' perceptions of ideological bias in research on public policy controversies. *Political Psychology* 30, 1 (2009), 43–65.
- MacKuen, M., et al. Resolute partisanship or reflective deliberation. *Am. J. of Pol. Sci.* 54, 2 (2010).
- Mutz, D. *Hearing the Other Side: Deliberative versus Participatory Democracy*. Cambridge, 2006.
- Nyhan, B., and Reifler, J. When corrections fail: The persistence of political misperceptions. *Political Behavior* 32, 2 (June 2010), 303–330.
- Park, S., et al. Newscube: Delivering multiple aspects of news to mitigate media bias. In *CHI* (2009).
- Price, V., Nir, L., and Cappella, J. N. Normative and informational influences in online political discussions. *Communication Theory* 16, 1 (Feb. 2006), 47–74.
- Shipman, F., et al. Formality considered harmful. *JCSCW* 8, 4 (1999).
- Shum, S. *Visualizing Argumentation*. Springer, 2003.
- Shum, S. Cohere: Towards web 2.0 argumentation. In *COMMA* (2008).
- Stvilia, B., Twidale, M. B., Smith, L. C., and Gasser, L. Information quality work organization in Wikipedia. *Journal of the American Society for Information Science and Technology* 59, 6 (2008), 983–1001.
- Sukumaran, A., et al. Normative influences on thoughtful online participation. In *CHI* (2011).
- Sunstein, C. *Republic.com 2.0*. Princeton Univ. Press, 2009.
- Taber, C. S., and Lodge, M. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* 50, 3 (July 2006), 755–769.
- Thompson, D. Deliberative democratic theory and empirical political science. *Annual Review of Political Science* 11 (2008).
- Wahl-Jorgensen, K. Letters to the editor as a forum for public deliberation. *Crit. Stud. in Media Comm.* 18, 3 (2001).
- Zickuhr, K., Rainie, L., and Purcell, K. *Library Services in the Digital Age*. Pew Internet & American Life Project, 2013. Available at <http://libraries.pewinternet.org/2013/01/22/library-services/>.