When Subjects Interpret the Data: Social Media Non-use as a Case for Adapting the Delphi Method to CSCW

Eric P. S. Baumer[1,2], Xiaotong Xu[2], Christine Chu[2], Shion Guha[3], Geri K. Gay[1,2]
[1] Communication [2] Information Science [3] Mathematics, Statistics and Computer Science
Cornell University Marquette University
Ithaca, NY USA Milwaukee, WI USA
{ericpsb, xx247, cc2228, gkg1}@cornell.edu | shion.guha@marquette.edu

ABSTRACT
This paper describes the use of the Delphi method as a means of incorporating study participants into the processes of data analysis and interpretation. As a case study, it focuses on perceptions about use and non-use of the social media site Facebook. The work presented here involves three phases. First, a large survey included both a demographically representative sample and a convenience sample. Second, a smaller follow-up survey presented results from that survey back to survey respondents. Third, a series of qualitative member checking interviews with additional survey respondents served to validate the findings of the follow-up survey. This paper demonstrates the utility of Delphi by highlighting the ways that it enables us to synthesize across these three study phases, advancing understanding of perceptions about social media use and non-use. The paper concludes by discussing the broader applicability of the Delphi method across CSCW research.

Author Keywords
Delphi; methods; social media; Facebook; non-use.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Data do not speak for themselves. Rather, data must be interpreted to be made meaningful. Furthermore, data can have multiple different interpretations and meanings, not only for us as researchers, but also, in the case of human subjects research, for those people from whom the data are collected.

Researchers in CSCW and related fields have demonstrated interest in intellectual traditions that emphasize peoples' interpretations of their own activities. Such traditions include ethnomethodology [25,87], ethnography [18,24], action research [37,51], and others. Despite this interest, CSCW and HCI research rarely involves study participants or users in the interpretation of their own data. With a few exceptions [e.g., 16,97], it is uncommon to show a study participant analysis of data to which s/he contributed and ask that participant to interpret, corroborate, contradict, or otherwise react to the results. Doing so certainly has potential issues, such as ensuring that laypersons understand clearly the methods and results at hand, or potential violation of an appeal to scientific objectivity. However, just as participatory design benefits from incorporating perspectives of those who will use the system being designed, data analysis might benefit from incorporating the perspectives of those whose data are being analyzed.

As one means of doing so, this paper draws on the Delphi method [53]. Originally developed as a technique for facilitating reasoned discussion about contentious issues [30], the method has been used for a variety of purposes, including forecasting and predictions [22,55,77], crafting policy recommendations [28,45], opinion assessment [94], and others. This paper demonstrates how the Delphi method can be adapted as a means of incorporating study participants into the interpretation of their own data.

As a case study, we explore perceptions about the use and non-use of social media [5,6,81]. Technology non-use provides case well suited to examining study participants' perceptions and interpretations, particularly because of the varying degrees of visibility involved. While some people "quit Facebook in a huff" [69:1042], others' sudden absence may go virtually unnoticed [8] for a variety of reasons. Deactivating your account means that "your profile won't be visible to other people on Facebook" [https://www.facebook.com/help/214376678584711]. News feed curation may make absence less notable [23]. Human cognitive capacity may limit the number of friends one can remember [20]. At the same time, the majority of US internet users have a Facebook account [19], making the lack of an account more conspicuous. Thus, Facebook non-use offers an ideal setting in which to examine ways that perceptions and expectations both align with, and diverge from, data about social media usage.

We explore using the Delphi method in the context of a demographically representative survey about Facebook use

CSCW '17, February 25-March 01, 2017, Portland, OR, USA

ACM 978-1-4503-4335-0/17/03…$15.00
DOI: http://dx.doi.org/10.1145/2998181.2998182

and non-use. Respondents from this survey were subsequently asked first to predict the results from the survey, and then they were shown the actual results. The findings show that respondents had varying accuracy in their perception of others' social media use, often mediated by the relative visibility of the particular use (or non-use) activities in question. Furthermore, respondents' explanations for these results incorporate their perspectives back into the interpretation of the data. This case study offers a demonstration of how the Delphi method can be incorporated into standard survey research, and it offers examples of the kinds of insight that doing so can generate.

However, we also want to assess the validity of study participants' interpretations. To do so, we incorporated the iterative nature of the Delphi method to conduct a second, interview-based study that draws on the tradition of member checking [52,59]. These interviews not only helped confirm the validity of our first round study, but they also added a layer of complexity and subtlety to those first round interpretations. After discussing the results of this case study in detail, we consider both strengths and potential limitations for the incorporation of Delphi techniques into a variety of CSCW research areas.

RELATED WORK

The Delphi Method

At its core, the Delphi method is complex, multifaceted, and intentionally ill defined. Broadly speaking, Delphi is "a method for structuring a group communication process so that the process is effective in allowing a group of individuals, as a whole, to deal with a complex problem" [53:3]. Often, the group involved is a panel of experts, though not always. The communication process in question many times deals with making forecasts or predictions. Examples range from predicting novel plastics developments in materials science [22] to envisioning priorities for clinical nursing research [30].

However, "while many people label Delphi a forecasting procedure […], there is a surprising variety of other application areas" [53:3–4]. Previous work has employed the Delphi method for informing policy, estimating historical data, regional planning, educating participants, translating scientific findings into decision recommendations, and other areas [for these and other examples, see 53, Chapter III].

Despite this variety, Delphi studies adhere to four core elements:

- Controlled Feedback – Participants (either all participants or a subset thereof) are shown the results of the surveys or questionnaires they complete. These results shown are both selected and presented by researchers conducting the Delphi study.
- Anonymity – Participants often do not meet face to face, and individual responses are not associated with identifiable participants.
- Statistical Aggregation – Feedback is shown in the form of statistical analyses, which can help ensure anonymity. These statistics can also be compared with participants' estimates.
- Iteration – The results of in initial round of surveys and feedback are used to inform subsequent round(s), with the exact number of rounds varying by study, "though seldom goes beyond one or two iterations" [77:355].

Within HCI, Mankoff et al. [55] used Delphi to study views of sustainable researchers on the role evolving of HCI in their domain. Even without explicit directions, participants gave global attention in their answers, which allowed researchers to understand the breadth of issues and relative priorities of different technology-related issues in sustainability research. Jones et al. [42] used Delphi to study practices and strategies around personal information management. This example shows how Delphi can be used not only to predict the future but also to understand the present.

In some cases, applications of Delphi will deviate slightly from some of these key features. For example, one policy Delphi [28] incorporated "long-time residents, hoteliers, bar owners, real estate dealers, and civic officials as well as the usual 'experts'" [53:76]. One portion of the Mankoff et al. [55] study mentioned above included a synchronous video chat in which participants' identities were not anonymized. In yet another variant [94], doctors were asked to make estimates about the prevalence of absence from work due to illness. These estimates were then compared with data collected from workers' files to assess doctors' perceptions of employee absenteeism. These represent just a few variants that still adhere to the core Delphi structure.

At the same time, these differences draw attention to some of the critiques of Delphi. The variety of studies claiming the mantle of Delphi means that, if you ask, "What is the Delphi technique? No one, it seems, really knows" [29:196]. While anonymity may engender freer discussion, it may also grant impunity for socially sanctioned behaviors (e.g., stereotyping or prejudice) [56]. Researchers conducting the Delphi study may have undue influence over results through the selection and presentation of feedback to participants [30,92]. These and other critiques [79] have led to significant, some [e.g., 17,29] would say conclusive, debate supporting the validity of Delphi. Nonetheless, such potential critiques, as well as their particular relevance for CSCW research, are further considered in the discussion section.

Member Checking

In qualitative studies, researchers will at times bring their results back to the participants being studied. Doing so allows for checking researchers' interpretations against members' own understandings [52,59]. Such member checking can improve the validity and reliability of qualitative findings. Some HCI and CSCW research has employed related or similar techniques [e.g., 16,90,97].

Although bringing results and interpretations back to study participants can provide valuable insights, some argue that member checking places undue emphasis on arriving at a correct or truthful understanding [4,80]. Interpretivist approaches, this critique suggests, should focus less on establishing a veridical account of reality and instead focus on understanding interpretations of intersubjectively constructed social reality/ies [cf. 65]. Instead, they suggest a more holistic approach that emphasizes "fidelity to phenomena" and the "practical underpinnings of the inquiry [rather than] methodological criteria" [4:386–7].

This paper uses member checking primarily as a means for establishing validity of results, as is appropriate for a methodological contribution. Our member checking assesses the extent to which results obtained from our Delphi study ring true with the population from whom the data were drawn. Doing so helps justify the use of the Delphi method within CSCW. Member checking is also used here to assess the validity of our interpretation. That is, member checking helps ensure that we as researchers are interpreting the results in a manner that is both sensible and consistent with members' own interpretations.

Non-use

Since this paper examines perceptions of social media use and non-use as a case study, we provide here a brief background. Recent work in CSCW, HCI, and related areas has highlighted the importance of technology non-use [6,81,95]. Individual studies have taken a variety of approaches, including comparing users and non-users [1,8,36,44,78,85,88], examining the motivations for avoiding or resisting a given technology [2,5,69], arguing for the benefits of approaches that transcend a strict binary between use and non-use [2,6,10,44,50], and exploring the broader social, technical, and cultural milieux in which these practices unfold [15,43,66,82,83,96,98].

Much of this work focuses on social aspects. How does one individual's or group's (re)negotiation of their engagement with, or disengagement from, a technology impact others with whom they interact, and vice versa? Portwood-Stacer [69] describes a variety of reactions that her participants received after leaving Facebook, ranging from support, to confusion, to anger and disgust. Ugander et al. [89] find that the diversity of one's social network structure influences adoption of Facebook. A few studies examine how an individual's relationship with others, particularly romantic partners, may shape her/his technology use [15,27,99]. Baumer et al. [8] find that, when an individual leaves Facebook, others' reactions to that departure influence whether the individual subsequently returns to the site. Baumer et al. [5] also find evidence for a social contagion effect, in that survey respondents who knew someone else who had deactivated their Facebook account were significantly more likely to have deactivated their own account. In many ways, such work suggests that one's perceptions of how other people use (or do not use) social media may influence one's own use or non-use.

BACKGROUND

The work presented here builds on data collected during a prior study [7]. This section describes the data and sampling method for that prior study.

Prior Survey Design

Our prior survey included three groups of questions. First, a series of questions determined the type of non/user for each respondent. Second, existing, well-validated scales were used to measure four constructs that may influence types of non/use, as described below, as well as demographics. Third, the survey included several open-ended questions not analyzed here.

The survey included questions to collect four types of data about each respondent. First, we used the Facebook Intensity Scale (FBI) (8 items) [21] to assess overall intensity of Facebook usage. Second, the Bergen Facebook Addiction Scale (BFAS) (18 items) [3] assessed the degreed to which self-reported Facebook use resembles other instances of behavioral addiction [31,32]. Third, to measure questions around Privacy Behaviors and Experiences (PBE) (10 items), we drew heavily from Wang et al.'s [91] examination of regret and embarrassing experiences on Facebook. Finally, we a asked a series of demographic questions, including age, gender, household income, marital status, ethnicity, education, and political views.

Prior Survey Participants and Sampling

To acquire a representative sample of US Internet users, we contracted with a survey and sampling agency, Qualtrics. Their recruitment and sampling procedure is outlined on their website (https://www.qualtrics.com/online-sample/). Qualtrics' staff assembled a web panel of participants using demographic criteria derived in part from the internet omnibus survey conducted by Pew Research (http://www.pewinternet.org/datasets/january-2014-25th-anniversary-of-the-web-omnibus/). The demographic screening criteria used included gender, race/ethnicity, age, and income. Demographic questions were used to screen respondents. For example, once we received 89 respondents age 25-34 (i.e., 17.8% of our target sample size of 500 respondents), subsequent respondents age 25-34 did not pass the age criterion. Respondents who did not pass any of the demographic screening criteria were excluded.

Recruitment continued until we had accumulated sufficient numbers of respondents for each demographic category. Ultimately, we collected a web panel of N=515 participants, for which we paid $2,750. Of them, 379 participants either currently have or previously had a Facebook account. For comparison, we also recruited a convenience sample from Mechanical Turk (N=1000). Our analysis showed similar results across the two samples [7]. In supplementary material, we summarize and cross tabulate the two samples' demographics, showing that demographics of the subsample

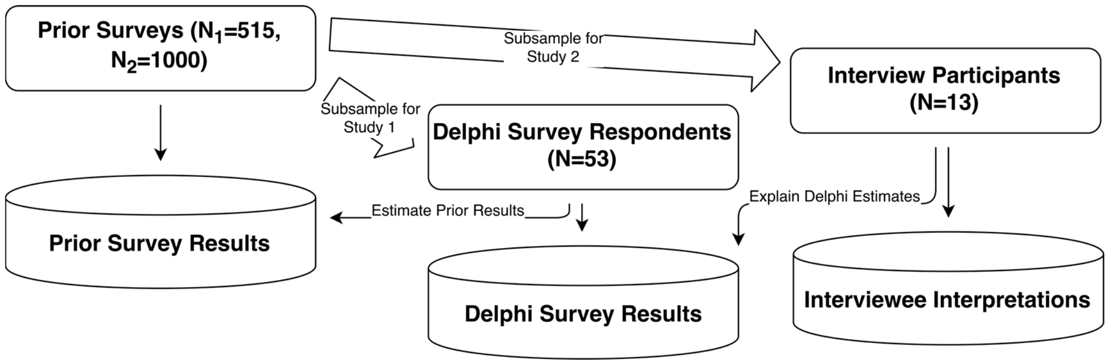

Figure 1 – Relationships among participants (rectangles) and data (cylinders) from our prior survey (left column), our Delphi survey (Study I, center column), and our member checking interviews (Study II, right column).

for Study I (described below) resemble those of our prior surveys. A full analysis of these data is currently under review [7].

STUDY I – DELPHI SURVEY

The results presented here from come an iterative Delphi study with three phases. Our prior survey, described above in the Background section, serves as the first phase. Participants who indicated at the end of that prior survey that they would be interested in follow-up studies were randomly recruited for Study I and for Study II. Study I, described in this section, asked participants to estimate the responses we received during our prior survey. Study II, described in the subsequent section, asked participants about their interpretations of the results from Study I. The study protocols were approved by our institution's IRB. Relationships among our prior survey, Study I, and Study II are depicted in Figure 1.

This design follows the core Delphi elements of controlled feedback, anonymity, statistical aggregation, and iteration [53]. It also incorporates two departures from common Delphi practice, both of which have been tested in prior work. First, rather than including only experts, we recruit a subsample of Facebook users and non-users. This approach resembles that of a policy Delphi study that included a variety of different stakeholders rather than strictly experts [28]. Second, rather than making predictions about the future, we ask about perceptions of present social media use and non-use. This aspect resembles a Delphi study in which doctors were asked to predict, and then shown data about, illness-related absence from work [94].

Study I Methods

Survey Design

We collected a series of results from our prior survey that pertained perceptions about different aspects of Facebook use and non-use. Possible questions included those that we believed that might provoke reactions, particularly because the results from our prior survey might be interesting, unexpected, or provocative. We also explored different means of presenting results from the same or similar questions to make each question as easy to understand as possible. Finally, to avoid overwhelming participants, we sought to make the survey brief while still comprehensive in its variety of questions. To do so, we pilot tested several different results and questions about them, as well as different means of presenting the same results. Ultimately, we chose five main results to include. This paper reports on responses to four of those, as described further below.

After gaining informed consent, the survey began by describing our prior survey. Respondents then answered which of a series of statements applied to them, including "I have, at least once, *deactivated* my Facebook account," "I have voluntarily *taken a break* from using Facebook for several weeks or more," and others (emphasis original) [8,44,72]. This was followed by a series of five results and accompanying sets of questions informed by our prior survey, all of which included some elements of a similar template. This paper reports on four out of five of those sets of questions, each on a different topic:

- *Deactivation* – "What percentage of people do you think have ever considered deactivating their Facebook account?" Responses were whole percentage between 0% and 100%.
- *Privacy* – "What percentage of Facebook users do you think fall into each of the following groups: Not familiar with privacy settings; Are familiar with privacy settings but have not changed them; and Are familiar with privacy settings and have changed them?" Responses for each category where whole percentages between 0% and 100%, and they were required to total 100% across the three categories.
- *Spend More Time* – "How often do you spend more time on Facebook than initially intended?" Responses were on a five-point Likert scale from "Very rarely" to "Very often."
- *Unsuccessful Cutting Down* – "How often have you tried to cut down on the use of Facebook without success?" Responses were on a five-point Likert scale from "Very rarely" to "Very often."

A fifth question compared the amount of time survey respondents spent on Facebook with whether they considered it a part of their every day activity. Since this question did not involve predicting prior survey responses, it differs from the other four. To adhere more closely to the Delphi method, for consistency of analysis, and in the interest both of space and of consistency of analysis, we focus on the four questions listed above.

For each prior survey result, respondents began by answering the above questions. For ordinal questions (Spend More Time and Unsuccessful Cutting Down), we then asked respondents how they think their response compares with others' from our previous survey. This was asked in the form, "I think my response is higher than ___% of survey respondents," which essentially asks respondents to estimate the percentile of their response. These steps capture the *iteration* key to Delphi [53]. After making their estimates, respondents were shown *controlled feedback* in the form of results from the prior survey. Prior results for percentile-based questions were shown using a donut chart. For example, Figure 3 shows that 62% of our prior survey respondents said they had considered deactivating their Facebook account. For ordinal questions, respondents were shown a graph indicating their percentile in comparison with results from the prior survey. For example, Figure 2 shows that respondents who indicated they "Sometimes" spend more time on Facebook than initially intended were between the 38th and 68th percentiles. That is, their response was higher than 38% and lower than 32% of all responses received. This *statistical aggregation* of prior survey results also provides *anonymity*, in that specific responses were not associated with individual respondents.

Thus, this study design captures the core elements of the Delphi method [53]. However, rather than Delphi's typical emphasis on consensus, we focus instead perceptions [cf. 94]. Thus, each of these figures is followed by the question "Do you find this result surprising?" with a five-item Likert

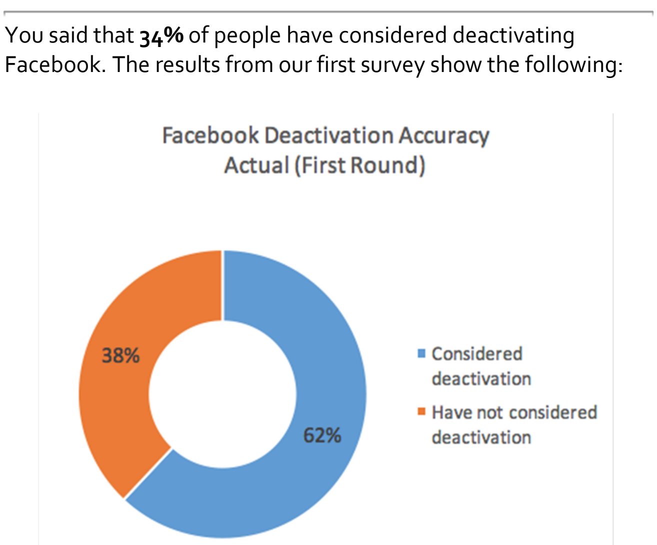

Figure 3 – Example presentation of percentage-based results from prior survey during Study I (Delphi survey). This image shows that 62% of prior survey respondents said they had considered deactivating their Facebook account, while the respondent guessed it would be 34%.

from "Not at all surprising" to "Very surprising." Respondents are then asked to elaborate why they did or did not find the results surprising.

The survey concluded with a 10-item personality inventory [73] and a series of demographic questions, including gender, age, income, marital status, self-described ethnicity, education level, and political ideology.

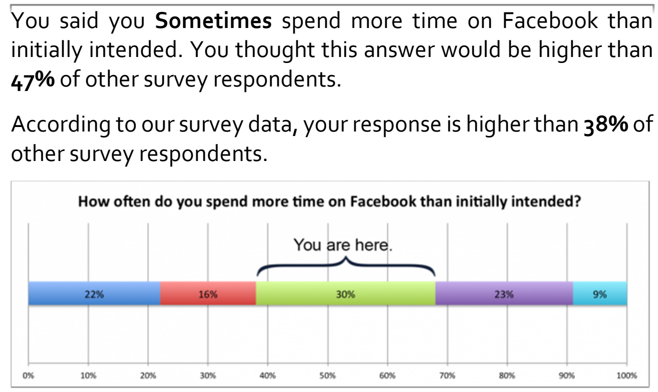

Figure 2 – Example presentation of ordinal-based results from prior survey during Study I (Delphi survey). This image shows that prior survey respondents who indicated they "Sometimes" spend more time or Facebook than initially intended were between the 38th and the 68th percentiles, while the respondent guessed the 47th percentile.

Survey Analysis

Our analysis involved three main questions. First, how accurate were respondents in estimating results from our prior survey? For percent-based questions (Deactivation and Privacy), accuracy was calculated as the raw difference between the respondent's estimate and the actual result; an overestimate would lead to a positive difference, and an underestimate to a negative difference. For Likert-style questions (Spend More Time and Unsuccessful Cutting Down), data from our prior survey were used to determine percentile-based bins associated with each of the five Likert response points. For instance, Figure 2 shows that respondents who indicated they "Sometimes" spend more time on Facebook than initially intended were between the 38th and 68th percentile. If such a respondent estimated that s/he would fall in the 30th percentile, the response was coded with an accuracy of –1, i.e., s/he underestimated the relationship of her/his response with others by one percentile-based bin.

Second, how surprising were prior survey results to the Delphi respondents? To address this question, we simply tabulated responses to the question asking Delphi respondents how surprising they found the results.

Third, which Delphi respondents were most accurate in their predictions? To do so, we employed exploratory regression modeling to examine relationships between attributes of each respondent and their accuracies for each question. To develop these models, we began with a set of potential predictors, including demographics (age, gender, income, etc.), responses on the prior survey (FBI scores [21], BFAS scores [3], etc.), and responses on the Delphi survey (whether respondent has previously taken a break from Facebook, whether using or not using Facebook is a voluntary choice, brief personality inventory [73], etc.). We used a process of iterative step-wise forward model selection. That is, we began by testing all models with only one predictor, selecting the best model based on AIC and variance explained. Using that one-predictor model as a basis, we then tested all two-predictor models that included the best single predictor. This process was repeated until adding predictors neither decreased AIC nor increased the variance explain. This process yielded a significant model only for respondents' accuracy in estimating the proportion of Facebook users who had considered deactivation.

Study I Results

We recruited a total of 63 respondents. Due to an administrative oversight, we were unable to connect ten of these responses with those from the prior survey, leaving N=53 for the main analysis (18 female, 35 male, 0 other; age 23 to 63, *M*=36.4, *Mdn*=34).

For each of the four items shown here, we present results using a similar structure. First, we compare respondents' answers with data from our prior survey to determine how accurate their responses were. Overall, respondents were most accurate in comparing their own Facebook usage to others, while they were least accurate in estimating others' consideration of deactivation and familiarity with privacy settings. Second, we asked respondents how surprised they were by the actual results. Overall, respondents showed relatively little surprise, even when their estimates were fairly inaccurate. Third, we used predictive regression modeling to assess the various factors influencing both respondents' accuracy and their surprise. A significant model converged only for Considered Deactivating, we only discuss regression results for that question.

Considered Deactivating

To reiterate, we asked Delphi survey respondents, "What percentage of people do you think have ever considered deactivating their Facebook account?"

Accuracy – In our prior study, 62% of respondents indicated they had considered deactivating. However, respondents in our Delphi survey estimated, on average, that only 42.45% or respondents had considered deactivating (see Figure 4), a significant underestimate.

Surprise – Despite their fairly low accuracy, only 30.2% found this result "moderately surprising" or "very surprising," while 60.4% of respondents described this result as only "a little surprising" or "not surprising at all." We suggest at least two possible explanations. First, Delphi respondents' may have had little confidence in their estimates, such that having them discredited would not be surprising. Second, upon seeing the actual results, respondents' may have formulated explanations that they then viewed as more credible than those informing their estimates. This second explanation is revisited below.

Predicting Accuracy – In the linear regression model for respondents' accuracy (Adj. R^2=0.371, p=0.001), only one variable was statistically significant (p=0.04). Those respondents who had "voluntarily taken a break from using Facebook for several weeks or more" were more accurate in their estimates; those who had not taken a break made estimates that averaged 19.6% lower (see Figure 5).

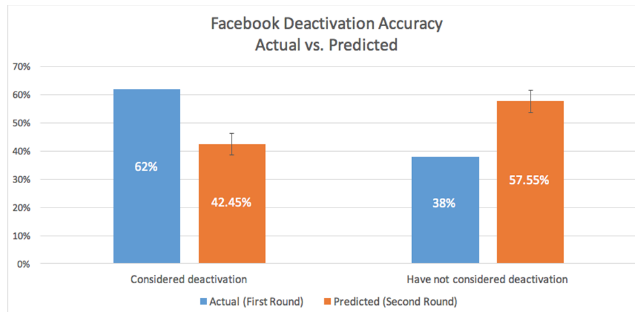

Figure 4 – Difference between proportion of respondents in prior survey considered deactivation (blue) and the average proportion of Facebook users that Delphi respondents thought have considered deactivation (orange). Delphi respondents underestimated how many users have considered deactivation.

Privacy Settings

As described above, Delphi respondents were asked what percentage of Facebook users fall into each of three categories: not familiar with privacy settings, familiar with privacy but have not changed them, and familiar with privacy settings and have changed them.

Accuracy – In our prior study, the majority (65%) of respondents report being familiar with and having changed their privacy settings on Facebook, with the fewest (17%) reporting not being at all familiar. However, our Delphi survey respondents estimated a more even distribution across these three categories (see Figure 6). Thus, as with consideration of deactivation, respondents were fairly inaccurate at assessing other Facebook users' familiarity and experience with privacy settings.

Surprise - Of our respondents, only 18.9% found this result "moderately surprising" or "very surprising," while 47.2% found it only "a little surprising" or "not surprising at all." As with considering deactivation above, there seems a discrepancy between the inaccuracy of respondents' estimates and their surprise at the actual findings.

Spending More Time than Intended

Delphi survey respondents answered a BFAS item [3] asking how often they spent more time on Facebook than they intended to spend (five-point Likert from "Very Rarely" to "Very Often"). They were then asked to estimate the percentile of their response, i.e., "I think my response is higher than ___% of survey respondents."

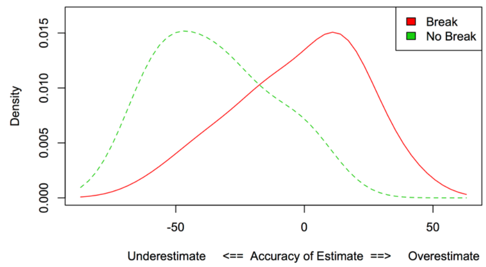

Figure 5 – Respondents' accuracy estimating the proportion of Facebook users who had considered deactivating their account was based in part on respondents' own prior non-use. Those who had previously taken a break from Facebook (solid red) made significantly more accurate estimates than those who had not previously taken a break (dashed green).

Accuracy – Overall, respondents were fairly accurate in comparing their own (over-)use of Facebook with that of others. Figure 7 shows that the majority respondents' percentile estimates were within the correct bin, with a slight skew to overestimating their own use (or, put differently, underestimating how often others spend more time on Facebook than they intended).

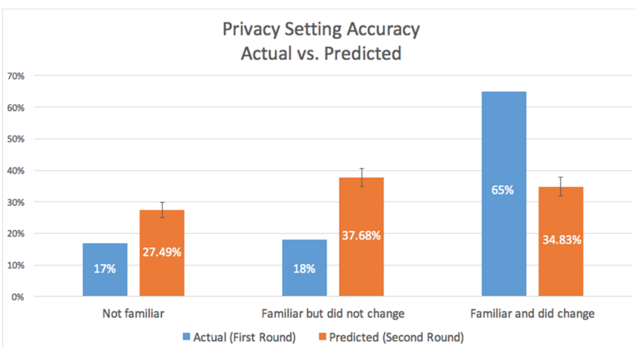

Figure 6 – Differences between familiarity of respondents in prior survey with Facebook privacy settings (blue) and the average familiarity that Delphi respondents thought Facebook users have with privacy settings. Delphi respondents overestimated the proportion of users who are unfamiliar with or have not changed their privacy settings, and they underestimated the proportion of users who have changed their privacy settings.

Surprise – Of our respondents, 22.6% found this result "moderately surprising" or "very surprising," while 66.0% described it as "not at all surprising." Taken alone, this low degree of surprise makes sense, since respondents were fairly accurate in comparing their own responses to others'. It becomes more difficult to interpret, though, in light of the above results about surprise. Respondents were far more accurate with their responses to this question than those to the questions above about privacy settings or considering deactivation. However, their degree of self-reported surprise remains quite similar. As above, we suspect that relatively low confidence in initial estimates may have meant that respondents were less surprised when they discovered their inaccuracy.

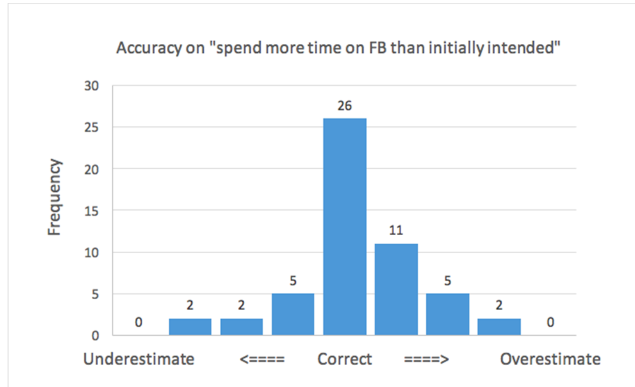

Figure 7 – Delphi respondents' estimates of the relationship between how often they spend more time on Facebook than initially intended and how often others do. The majority of respondents correctly estimated their relationship of their response with others'.

Attempt to Cut down without Success

Delphi survey respondents answered a BFAS item [3] asking how often they had tried to cut down on the use of Facebook without success (five-point Likert from "Very Rarely" to "Very Often"). As above, they were then asked to estimate the percentile of their response.

Accuracy – As above, data from our prior survey were used to determine percentile-based bins, which were then used to compare respondents' answers to actual data. Overall, respondents were fairly accurate in comparing their own ability to cut down on Facebook use with that of others'. Figure 8 shows that most respondents' percentile estimates were within the correct bin, with a slight skew to overestimating their own percentile (or, put differently, underestimating how often others have tried unsuccessfully to cut down on Facebook use).

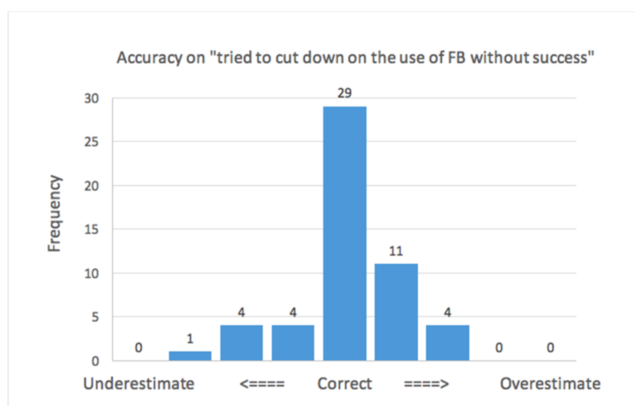

Figure 8 – Delphi respondents' estimates of the relationship between how often they spend more time on Facebook than initially intended and how often others do. The majority of respondents correctly estimated their relationship of their response with others'.

Surprise – Of our respondents, 15.1% found this result "moderately surprising" or "very surprising," while 73.6% described it as "not at all surprising." A similar results occurs here as in the question about spending more time than intended above. Respondents' accuracy would conceivably lead to reduced levels of surprise, but the levels of surprise are not vastly lower than with questions where respondents were far more *in*accurate. Again, we suggest that relatively low confidence in initial estimates may have led to less surprised when respondents were shown discrepancies between their estimates and actual results.

Study I Summary

Three key findings emerge from Study I. First, most respondents were fairly accurate at comparing their ability to control their FB use with that of others. Second, most respondents were fairly *inaccurate* when estimating others' consideration of non-use as well as others' privacy settings. Third, self-reported levels of surprise did not vary drastically among the different questions.

These findings show how the Delphi method can provide insights about users' perceptions of others' social media use and non-use. The consistently low levels of surprise suggest that these perceptions are knowingly rough guesses rather than firmly held beliefs. Since they are knowingly imperfect, these perceptions of others' social media use and non-use may have less sway over individual behavior. Furthermore, such perceptions are neither uniformly accurate nor uniformly inaccurate. Rather, accuracy varies depending on the particular aspect of social media under consideration. We discuss further below potential explanations for these differences in accuracy.

STUDY II – MEMBER CHECKING INTERVIEWS

To assess the validity of results from Study I, we conducted a second study. Study II blends the iteration that is a central tenet of Delphi [53] with the qualitative technique of member checking [52,59] as a means of establishing validity. To do so, we collected the results from Study I described above and formatted them to present to a panel of interviewees. Interview participants were recruited from the same pool as the respondents for the survey (i.e., those who had completed our initial survey and indicated that they were interested in follow-up studies). While often used for qualitative data, techniques similar to member checking have also been employed for quantitative results [e.g., 16].

Study II Methods

Interview Design

As with Study I, we pilot tested different sets of results to present, different means of presenting those results, and different questions to ask about them. We selected a subset of results based on responses during these pilot interviews. Results that struck pilot participants as confusing or obvious were avoided in favor of results that were more comprehensible and prompted more discussion.

These results were then presented to interviewees via Skype screen sharing and were grouped into three areas: (1) accuracy in estimating the percentage of users who have *considered deactivation*, (2) the best *factors predicting accuracy* in respondents' estimates, and (3) accuracy in estimating familiarity with *privacy settings*. Within each category, we described both the question from our prior survey and the question asked during the Delphi survey. We then showed respondents the actual data from our prior survey, followed by Delphi survey respondents' estimates.

Considered Deactivation – Interview participants saw Figure 3, which depicted results about the number of prior survey respondents who had considered deactivating their Facebook account. Interviewees were asked to guess how accurately respondents from the Delphi survey had predicted the results from the prior survey. We asked participants for their guess about general trends, rather than concrete numbers, as well as the reasoning behind their guesses. Participants were then shown Figure 9, which depicts Delphi respondents' estimates. Interviewees were asked about the ways that Delphi respondents' answers both fit with and diverged from their expectations, as well as their explanations.

Factors Predicting Accuracy – Interview participants were shown the relationship between Delphi respondents' accuracy and those respondents' prior non-use (Figure 5). Interviewees were asked first whether the result made sense and was consistent with their experiences. They were then asked for their interpretation, that is, what might explain the differences in accuracy we found. We also offered our own explanations, asking respondents to describe which they found most and least plausible, and why.

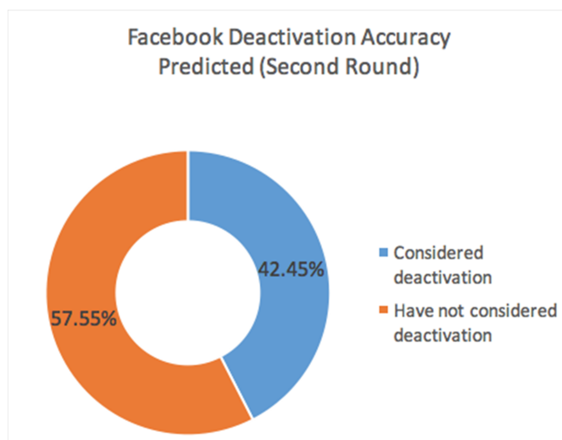

Figure 9 – Interview participants saw depictions of results from our Delphi survey. For instance, this chart, shown to interviewees along side Figure 3, depicts Delphi respondents' estimates for the proportion of Facebook users who have considered deactivating their account.

Privacy Settings – This question followed a similar format to that of *considered deactivation*. Interviewees were shown results from our prior study (the blue bars in Figure 6 presented in a style similar to Figure 9) and asked to guess how accurately they expected Delphi respondents would estimate this result. They were then shown the Delphi respondents' estimates and asked why they thought those estimates might have differed from the prior survey results.

Interview Analysis

These interviews were transcribed and analyzed in a qualitative, inductive manner [54]. As described above, our approach here is influenced by the qualitative technique of member checking [52,59], giving us two goals. First, we sought to establish validity, i.e., to confirm that the results obtained during the Delphi study rang true for interviewees. Second, we sought to assess our interpretations and explanations of the Delphi results by comparing them with interviewees' interpretations and explanations.

Study II Results

We recruited 13 interviewees (6 female, 6 male, 1 other; age 24 to 51, M=35.4, Mdn=38). By presenting interview participants' explanations for and interpretations of results from the Delphi survey, this section serves two validating functions. First, it validates results from the Delphi survey through confirmation by interviewees. Second, it validates our interpretations of those results by showing their alignment with interviewees' interpretations. The results presented here are organized around the three categories of questions used during the interviews. All participants are listed with number (P##), gender (F/M/O), and age (##).

Considered Deactivation

Most interviewees guessed that Delphi respondents would underestimate how many Facebook users had considered deactivation (9 out of 13). Interview participants thought Delphi respondents' estimates might be around 40%, compared to the prior survey result of 62%. The proximity of this guess to the actual average estimate of 42.45% lends credence to the finding that people underestimate the prevalence of considering deactivation.

Justifications for this underestimate were multifaceted but often incorporated personal experiences. Interviewees expected that others would not see many people whom they personally know deactivating their Facebook accounts, i.e., that deactivation has relatively low visibility. These interviewees also tended to have considered deactivating their own account. In mentioning this thought, they often brought up negative topics and experiences, such when "your timeline [is] filled with ads constantly" (P11, M, 38), feeling "ridiculously addicted" to Facebook (P13, F, 39), or "surveillance stuff, [how] they are censoring everyone's newsfeeds" (P9, M, 45). The few interview participants who guessed that Delphi respondents would overestimate the proportion who had considered deactivation focused on the idea of avoiding "drama," for example, "they go through a break up they don't want to see the old significant other on there" (P8, M, 28). Thus, interviewees corroborated both our results about accuracy in estimating the prevalence of deactivation and our interpretation of those results.

Factors Predicting Accuracy

This interpretation was further reinforced when interviewees were shown the relationship between having taken a break from Facebook and accuracy in estimating the proportion of users who have considered deactivation (i.e., Figure 5). All interview participants found the relationship reasonable. When asked to explain the relationship, interviewees commonly believed that people who had taken a break understood better the motivations that might lead someone to leave Facebook, such as a "big work assignment or school assignment" (P10, F, age 25) or due to job hunting. In contrast, interviewees believed that those who never took a break themselves were more likely users who "post *everything*" on Facebook and would hardly be able to imagine others taking a break. One interviewee speculated that those individuals did not realize that pervasive advertising bothered others. As another interviewee put it:

> *People who have never considered taking a break, they are not even conceptualizing of a world without FB. I mean it is so integrated in their life that -- taking a break, that would be a step toward deactivation, but they don't even consider that path something they would want to think about.* (P9, M, age 45)

These comments confirm both our findings about considering deactivation and our interpretation thereof. That is, the interviews reinforce the role that personal experiences play in perceptions of low visibility social media activities, such as considering deactivation.

Privacy Settings

When shown results about privacy settings (a donut chart version of blue bars in Figure 6), interview participants guessed that Delphi respondents' estimates (the orange bars in Figure 6) would follow the following general outline. First, they believed respondents would *overestimate* the proportion *not familiar* with privacy settings (respondents did actually overestimate this proportion). Second, they believed Delphi respondents would be mostly *accurate* for *familiar but not changed* (respondents actually overestimated this proportion). Third, interviewees believed that Delphi respondents would *underestimate* for the proportion *familiar and did change* (respondents did actually underestimate this proportion). This section discusses interview participants' reactions to Delphi respondents' estimates for each category in turn.

In our prior survey, 17% of Facebook users were *not familiar with privacy settings*, while our Delphi respondents estimated that about 27.5% of users would fall into this category. Although our interview participants accurately anticipated Delphi respondents' overestimates here, they were surprised at the magnitude of the overestimation. However, interviewees still came up with multiple explanations for this overestimate, such as "as the non-technical, oh-my-gosh-how-do-I-turn-on-the-computer people. […] I guess there's a slightly larger subset of non-technically inclined people than I would've guessed" (P3, M, age 24). Interview participants also suggested non-frequent or apathetic users: "if you rarely use Facebook you're probably not going to mess around with privacy settings, or the people who […] simply don't care" (P3, M, age 24). However, interview participants offered little explanation for why Delphi respondents would believe such apathy would be more prevalent than our prior survey results indicated.

Interviewees were most surprised by the results in the *familiar with but not changed* section. Many expected that estimates for this category would be relatively accurate, around 18% to 20%. They were surprised when it emerged as the largest estimated portion of the three (37.68%). Interview participants found interpreting this result somewhat difficult, as it differed so drastically from their expectations. "If you're familiar with it [privacy settings] I'd at least almost certainly think you'd at least change it a bit" (P4, M, age unknown). Some even suggested a blurry boundary between this category and the *not familiar* category. "It's amazing how much people don't know how to use anything, [it's] more a lack of knowledge" than a conscious choice (P13, F, age 39). Others suggested apathy, saying that it was a "who cares who has my data kind of thing" (P3, M, age 24) that drove Delphi respondents' estimates. One interviewee suggest that, "cause they're used to it [not changing privacy settings], they would just assume well everybody else is doing it and they're fine so I can too" (P1, F, age 44). Again, we see interviewees citing the interplay between personal experiences and perceptions of others' social media use.

Finally, interview participants correctly guessed that Delphi respondents' would underestimate the proportion of users who were *familiar with and did change* their privacy settings. Common justifications included, "I just feel that they would think less people are aware of you know what they have the ability to change" (P5, F, age 31). "A lot of people don't even know what all the settings are and wouldn't even find them on Facebook. Or," reiterating the apathy argument, "they're not worried about their privacy" (P1, F, age 44).

Study II Summary

The findings from Study II demonstrate two important points. First, in all cases, interviewees' explanations and interpretations focused less on what they expected Delphi respondents' answers to be and more on the actual results from our prior survey. This pattern aligns with the finding from Study I that, even when highly inaccurate, Delphi respondents did not report being surprised. Similarly, if interview respondents had been highly confident in their initial guesses, they may have tried to amend those explanations when they differed from actual results or perhaps even refute the results. Rather than explaining their guesses, they focused more on explaining the results.

Second, these interviews serve to help validate Study I. The interview results about perceptions of considering deactivation align very closely with our results from Study I, validating both the findings and our interpretations. The results about privacy were only partially validated, with interviewees agreeing with Delphi respondents in some places but disagreeing in others. Regardless, the justifications provided by interview participants help with overall interpretation of the results across the three phases of this study.

DISCUSSION

Each phase of our Delphi study presented above revealed interesting insights. This discussion describes how particular aspects of the Delphi method enable us to synthesize across the different phases and to form a coherent explanation for the results. We then raise some considerations and suggestions on the use of Delphi more broadly within CSCW research.

How Delphi Illuminates the Link between Visibility and Perceptions of Social Media Non/use

Prior work has highlighted the importance of visibility in perceiving social behavior [33,46,49,70,86]. Our use of the Delphi methods shows how those findings apply to perceptions of social media use. Broadly speaking, we find that the *more visible* an activity is, the more perceptions derive primarily from *observing others*, even when those highly visible behaviors may not actually be representative. In contrast, the *less visible* an activity is, the more perceptions derive from one's *own experiences*, regardless of the relationship between those experiences and others' actions.

The social media activities involved in our study – ability to control Facebook usage, changing privacy settings, and considering deactivation – all have vary in their level of visibility. As the activities in question become less visible, Delphi respondents' estimates become less accurate. Below, we describe the affordances of the Delphi method that highlighted each case of this relationship.

Ability to Control Facebook Usage

Our Delphi respondents were most accurate at comparing their own ability to control their Facebook usage with others' ability to do so. In many cases, difficulty in limiting one's own usage is relatively more visible. Frequency of both posting and even lurking [63,71] can be readily seen, the latter via status awareness. Thus, compared to other computer mediated environments [49], it is easier to see how often others are or are not using Facebook.

This relationship was identified largely through the iterative nature of Delphi feedback, first asking participants their own answer then asking them to compare their answer with others'. Doing so uniquely juxtaposes respondents' own activities and their perceptions of others' activities.

Privacy Settings

In contrast, Delphi respondents were less accurate at estimating others' familiarity with privacy settings. More judicious about their privacy settings make embarrassing, compromising, or otherwise undesirable content less visible. However, Facebook's privacy settings default to making most content fairly public [58,64]. The highly salient nature of content viewed by an unintended audience [57] makes such practices disproportionately more visible. Furthermore, media rhetoric around social media privacy [e.g., 14] may lead to a perception that indiscrete privacy settings are highly normative [12,40,74,76]. Delphi's iterative comparison of respondents' predictions with statistical aggregation of previous survey results highlights the inaccuracy of these perceptions.

Our member checking interviews confirm this interpretation. When explaining and interpreting the results (both those from our prior survey and predictions made by Delphi respondents), interviewees made reference to others' behavior: those who lack tech savvy, are apathetic, or simply use Facebook less. Interviewees often volunteered specific examples of such individuals without prompting. Regardless of the representivity of these individuals, their high visibility gave them significant influence over interviewees' perceptions.

Considering Deactivation

Considering non-use has even less visibility than changing one's privacy settings. While it is generally difficult to see what actions others are considering, deactivation has unique visibility characteristics. Only a user's friends can see when her/his account has been deactivated. Popular media reports often cast such quitters as exceptions rather than as exemplars of a much more common phenomenon [5,67,68]. The Delphi-enabled comparisons confirm that respondents were fairly inaccurate at gauging how many Facebook users had considered deactivating their account.

However, not all respondents were equally inaccurate. The statistical aggregation presented during the member checking interviews identified personal experience with deactivation as most predictive of accuracy. In following Delphi's dictum for controlled feedback, we eschewed the full details of the regression model in favor of presenting interviewees a pair-wise relationship (Figure 5). While explaining this result, interviewees struggled to think of individuals they personally knew who had deactivated, and these examples were never offered without prompting. When asked for their interpretation, interviewees suggested that Delphi respondents who feel "ridiculously addicted" and have not considered leaving Facebook would think it unlikely that others would consider leaving. Thus, the member checking interviews confirmed that perceptions of considering deactivation arose primarily from one's own experiences.

Thus, the core Delphi elements (iteration, statistical aggregation, anonymity, and controlled feedback) organize and synthesize results across multiple study phases, showing how varying levels in visibility of social media activities either increase or decrease both the accuracy of perceptions and the basis for forming those perceptions.

Limitations

When interpreting the results presented above, some caveats should be kept in mind. First, respondents in our Delphi survey and interviews are not as demographically representative as the full sample from which they are drawn. That said, responses to many questions in the Delphi survey about Facebook use closely resemble responses from the full sample. Furthermore, our regression modeling indicated that demographics did not significantly predict respondents' accuracy. Future work should examine how demographics influence norm perception in other computationally mediated contexts.

Second, respondents in our Delphi survey included those from our prior survey who indicated interest in follow-up studies and responded to our recruitment calls. Such an interest in social media may impact perceptions of Facebook use and non-use. Our findings about the aspects of Facebook usage that respondents perceived more accurately (e.g., ability to control one's own usage) and less accurately (e.g., frequency of considering deactivating) should be tested in other contexts with other populations. For instance, one could embed within a survey a variant of the FBI [21] or BFAS [3] that asks how the respondent believes certain questions might be answered by other social groups, such as their family members, other members of the same (or a different) political party, others of the same (or a different) religious faith, others of the same (or a different) economic standing, etc. This approach offers numerous possibilities ripe for future work. It also points to ways that Delphi methods could be incorporated more broadly into CSCW research.

Methodological Reflections and Implications

Other Methods for Incorporating Participants' Interpretations

As a field, CSCW has drawn on several diverse methodological and disciplinary traditions, from computer science, to sociology, to organizational studies, to anthropology. Many of those incorporate, in various ways, the perspectives and interpretations of the people being studied. The Delphi method provides a complement to this current panoply, offering a means of incorporating study participants into the analytic and interpretive processes that differs from methods currently common in the field.

For example, ethnomethodology [25] is concerned with members' methods. That is, how do people who are doing something know what it is that they are doing? This approach often involves very fine grained analysis of video data [e.g., 38,87]. Rarely, though, does such work explicitly ask research participants for their own interpretations. Nor does it present the results of initial study phases back to participants, as in a Delphi study. Thus, ethnomethodology may be better suited to studying mumbers' interpretive processes *in situ*, while Delphi may be better suited to calling those interpretations to conscious attention.

Ethnography offers another example, making a distinction between etic and emic perspectives [24:Ch. 2]. An etic perspective uses categories and concepts that come from outside the social group or culture being studied, often from researchers. For instance, an anthropologist might note that a stone adze is used as an agricultural tool, as a weapon, and as a religious implement [47]. An emic perspective, in contrast, interprets behaviors in terms that are meaningful for members of the social group or culture being studied. That is, the people the aforementioned anthropologist is studying might not distinguish in the same way among the categories of agriculture, warfare, and religion. While Delphi emphasizes study participants' interpretations, it does not carry the same kind of commitment to capturing the subjective experiences of participants, i.e., it neither commits to, nor strictly opposed, an interpretivist epistemology [cf. 13,65].

As a final example, action research [51] places "partnerships with research participants at the center of the process of inquiry" [37:3]. In this approach, participant-researchers often help craft research questions, develop a study design, collect and analyze data, interpret the results, and make recommendations. In contrast, Delphi focuses participant involvement at specified points in the research process, leaving the selection of research questions and design of the study in the hands of researchers. This difference can be seen most prominently in the emphasis on *controlled* feedback, that is, results selected and presented by researchers. Thus, Delphi may be less appropriate in situations where study participants should set the agenda and research questions, rather than simply aiding with interpretation.

These represent only a few notable examples. Similar comparisons could be drawn with participatory design [62], cultural probes [26], or a number of other approaches. Nonetheless, these examples help situate Delphi within a broader class of techniques used to incorporate study participants into the analysis and interpretation of data.

Potential Challenges

The histories of CSCW and HCI research contain prominent examples of methodological adaptation with less than desirable results [13,18,84]. Informed by our first-hand experiences, we articulate a few challenges with which CSCW research will need to grapple in adapting Delphi.

As mentioned in prior critiques, the selection of feedback potentially grants researchers the ability to influence Delphi results [30,92]. In the work presented here, our study designs were informed by pilot testing. For controlled feedback, we selected the questions that provoked the strongest responses during our pilots. It is conceivable that a researcher could select certain subsets of results or present them in such a way so as to guide Delphi respondents in a particular direction. Similar issues emerge in the reporting of results from scientific research more generally. However, papers reporting scientific results are the subject of expert peer review. Delphi participants may or may not have enough expertise in, say, the visual presentation of quantitative information to know when results are presented in a misleading way.

One simple means of addressing this challenge involves researchers explicating their process of selecting controlled feedback during Delphi iterations, as done in this paper. A more complex line of future could intentionally manipulate Delphi feedback in different ways or to varying degrees. At what points do Delphi respondents refuse to believe the results they are shown? Or might trust in scientific researchers [cf. 34,60] lead respondents to believe grossly falsified data as true? While not terribly informative in the study presented above, asking respondents about levels of surprise may provide one means of exploring this question.

One final challenge deals with the flexibility of Delphi as a method. While its originators praise the method's flexibility as a value and encourage creativity in designing Delphi studies [53], others suggest moderation in exercising that creativity [e.g., 29]. A survey of Delphi studies and critiques [17] attempts to synthesize a working definition by identifying traits common among Delphi studies: anonymity of participants (though this is sometimes violated [11]); control of information flow and feedback by researchers; iteration and repetition; presentation of interim results using data and statistical analysis. The astute reader will likely notice that these resemble quite closely the main components of Delphi's original formulation [53]. Future work in CSCW using Delphi, as outlined below, would benefit from adhering to these core elements.

Future Directions

Interpreting Big Data – With computational approaches for analyzing social data becoming more prevalent [48], interpretation of results becomes a more complex challenge [35]. Recent work has argued for the value of comparing, or in some cases combining, computational analysis with qualitative methods [9,39,61,75]. Delphi studies may provide a compelling means of creating such combinations. An initial iteration of computational analysis followed by surveys and/or interviews would allow for incorporating subjects whose data are being analyzed into the interpretive process [cf. 16]. Subsequent iterations could then adjust the computational analysis based on subjects' interpretations. Such an approach would certainly carry unique challenges, e.g., explaining to participants how a Naïve Bayes classifier works or what odds ratios mean in a logistic regression model. Notwithstanding, this approach may provide an attractive hybrid between computational and qualitative methods.

Norms – Significant work has investigated norms in computer mediated contexts [e.g., 49,70,86]. communication scholarship on norms distinguishes between descriptive norms (what people *actually* do) and perceived norms (what we *believe* people do) [40,41,46,74]. Disconnects between these two types of norms can be highlighted using the Delphi method. Indeed, it may be informative to ask participants to reconcile differences between their perceived norms and data-based descriptive norms.

Citizen Science – Most citizen science projects involve non-expert members of the public primarily for data collection or content processing [93]. Far less often do such citizens participate the processes of interpreting data or results. The Delphi method could provide a means of incorporating non-expert interpretations into the analysis process. As suggested above, doing so could provide particular value in the case of interpreting social data where the Delphi participants are selected from the same population from whom the data were collected.

CONCLUSION

The same data can mean different things to different people. As researchers, we are interested in particular questions, concepts, and theories. We collect and analyze data about groups of people interacting with computational technologies to test these theories, build these concepts, and answer these questions. In so doing, our interpretations are guided by our interests as researchers. The people whom we are studying, though, may have different interests or perspectives that enable them to corroborate, challenge, deny, or propose alternatives to our interpretations. CSCW research can benefit when these voices enter the conversation. As demonstrated in this paper, the Delphi method provides one compelling means for doing so.

ACKNOWLEDGMENTS

This material is based upon work supported by the NSF under Grant No. IIS-1110932. Work presented here was conducted while all authors were affiliated with Cornell University. Thanks to our survey and interview participants for their time and insights, and to the anonymous reviewers for valuable suggestions.

REFERENCES

1. Alessandro Acquisti and Ralph Gross. 2006. Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook. In *Proc Privacy Enhancing Technologies Workshop*, 36–58.
2. Morgan G Ames. 2013. Managing Mobile Multitasking: The Culture of iPhones on Stanford Campus. In *Proc CSCW*, 1487–1498. https://doi.org/10.1145/2441776.2441945
3. Cecilie Schou Andreassen, Torbjørn Torsheim, Geir Scott Brunborg, and Ståle Pallesen. 2012. Development of a Facebook Addiction Scale. *Psychological Reports* 110, 2: 501–517. https://doi.org/10.2466/02.09.18.PR0.110.2.501-517
4. Maureen Jane Angen. 2000. Evaluating interpretive inquiry: Reviewing the validity debate and opening the dialogue. *Qualitative Health Research* 10, 3: 378–395. https://doi.org/10.1177/104973200129118516
5. Eric P. S. Baumer, Phil Adams, Vera D. Khovanskaya, Tony C. Liao, Madeline E. Smith, Victoria Schwanda Sosik, and Kaiton Williams. 2013. Limiting, Leaving, and (re)Lapsing: An Exploration of Facebook Non-Use Practices and Experiences. In *Proc CHI*, 3257–3266. https://doi.org/10.1145/2470654.2466446
6. Eric P. S. Baumer, Jenna Burrell, Morgan G. Ames, Jed R. Brubaker, and Paul Dourish. 2015. On the Importance and Implications of Studying Technology Non-use. *interactions* 22, 2: 52–56. https://doi.org/10.1145/2723667
7. Eric P. S. Baumer, Shion Guha, and Geri K. Gay. All Non-users are Not Created Equal: Predictors Vary for Different Forms of Facebook Non-use. *under review*.
8. Eric P. S. Baumer, Shion Guha, Emily Quan, David Mimno, and G. K. Gay. 2015. Missing Photos, Suffering Withdrawal, or Finding Freedom? How Experiences of Social Media Non-Use Influence the Likelihood of Reversion. *Social Media + Society* 1, 2: 1–14. https://doi.org/10.1177/2056305115614851
9. Eric P. S. Baumer, David Mimno, Shion Guha, Emily Quan, and Geri Gay. Comparing Topic Modeling and Grounded Theory: Extreme Divergence or Unlikely Convergence? *JASIST* in press.
10. Eric P. S. Baumer. 2015. Usees. In *Proc CHI*, 3295–3298.
11. Wendell Bell. 2003. *Foundations of Futures Studies: Human science for a new era: History, purposes and knowledge*. Transaction Publishers, Piscataway, NJ.
12. Ulf Böckenholt and Peter G. M. van der Heijden. 2007. Item Randomized-Response Models for Measuring Noncompliance: Risk-Return Perceptions, Social Influences, and Self-Protective Responses. *Psychometrika* 72, 2: 245–262. https://doi.org/10.1007/s11336-005-1495-y
13. Kirsten Boehner, Janet Vertesi, Phoebe Sengers, and Paul Dourish. 2007. How HCI interprets the probes. In *Proc CHI*, 1077–1086. https://doi.org/10.1145/1240624.1240789
14. John Brownlee. 2012. This Creepy App Isn't Just Stalking Women Without Their Knowledge, It's A Wake-Up Call About Facebook Privacy. *Cult of Mac*. Retrieved from http://www.cultofmac.com/157641/this-creepy-app-isnt-just-stalking-women-without-their-knowledge-its-a-wake-up-call-about-facebook-privacy/
15. Jed R. Brubaker, Mike Ananny, and Kate Crawford. 2014. Departing glances: A sociotechnical account

of "leaving" Grindr. *New Media & Society* 18, 3: 373–390. https://doi.org/10.1177/1461444814542311

16. Munmun De Choudhury and Michael Massimi. 2015. "She said yes!" – Liminality and Engagement Announcements on Twitter. In *Proc iConference*. Retrieved from http://hdl.handle.net/2142/73658

17. Holly M. Donohoe and Roger D. Needham. 2009. Moving Best Practice Forward: Delphi Characteristics, Advantages, Potential Problems, and Solutions. *International Journal of Tourism Research* 11, 5: 415–437. https://doi.org/10.1002/jtr.709

18. Paul Dourish. 2006. Implications for design. In *Proc CHI*, 541–550. https://doi.org/10.1145/1124772.1124855

19. Maeve Duggan, Nicole B. Ellison, Cliff Lampe, Amanda Lenhart, and Mary Madden. 2015. *Pew Social Meida Report 2015*. Washington, D.C. Retrieved from http://www.pewinternet.org/2015/01/09/social-media-update-2014/

20. R.I.M. Dunbar. 1993. Coevolution of neocortical size, group size and language in humans. *Behavioral and Brain Sciences* 16, 4: 681–735. https://doi.org/10.1017/S0140525X00032325

21. Nicole B. Ellison, Charles Steinfield, and Cliff Lampe. 2007. The Benefits of Facebook "Friends:" Social Capital and College Students' Use of Online Social Network Sites. *Journal of Computer-Mediated Communication* 12, 4: 1143–1168. https://doi.org/10.1111/j.1083-6101.2007.00367.x

22. Selwyn Enzer. 1975. Plastics and Competing Materials by 1985: A Delphi Forecasting Study. In *The Delphi Method*, Harold A Linstone and Murray Turoff (eds.). Addison-Wesley, Reading, MA, 189–203.

23. Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. "I always assumed that I wasn't really that close to [her]": Reasoning about Invisible Algorithms in News Feeds. In *Proc CHI*, 153–162. https://doi.org/10.1145/2702123.2702556

24. David Fetterman. 2010. *Ethnography: Step-by-Step*. Sage Publications, Thousand Oaks, CA.

25. H. Garfinkel. 1967. *Studies in Ethnomethodology*. Prentice-Hall, Englewood-Cliffs, NJ.

26. William W. Gaver, T. Dunne, and E. Pacenti. 1999. Cultural Probes. *interactions* 6, 1: 21 – 29. https://doi.org/10.1145/291224.291235

27. Ilana Gershon. 2011. Un-Friend My Heart: Facebook, Promiscuity, and Heartbreak in a Neoliberal Age. *Anthropological Quarterly* 84, 4: 865–894. https://doi.org/10.1353/anq.2011.0048

28. Peter G. Goldschmidt and Andrew W. Dahl. 1976. Demoflush: Estimating Population in Seasonal Resort Communities. *Growth & Change* 7, 2: 44–48. https://doi.org/10.1111/j.1468-2257.1976.tb00305.x

29. Peter G. Goldschmidt. 1975. Scientific Inquiry or Political Critique - Remarks on Delphi Assessment, Expert Opinion, Forecasting, and Group Process by H. Sackman. *Technological Forecasting and Social Change* 7, 2: 195–213. https://doi.org/10.1016/0040-1625(75)90059-1

30. Claire M. Goodman. 1987. The Delphi technique: a critique. *Journal of Advanced Nursing* 12: 729–734. https://doi.org/10.1111/j.1365-2648.1987.tb01376.x

31. Jon E. Grant, Marc N. Potenza, Aviv Weinstein, and David A. Gorelick. 2010. Introduction to Behavioral Addictions. *Am J Drug Alcohol Abuse* 36, 5: 233–241. https://doi.org/10.3109/00952990.2010.491884.Introduction

32. Jon E. Grant. 2008. *Impulse Control Disorders*. W. W. Norton & Company, New York.

33. Shion Guha and Jeremy Birnholtz. 2013. Can you see me now?: location, visibility and the management of impressions on foursquare. In *Proc MobileHCI*, 183–192. https://doi.org/10.1145/2493190.2493209

34. Craig Haney, Curtis Banks, and Philip Zimbardo. 1973. Interpersonal Dynamics in a Simulated Prison. *International Journal of Criminology and Penology* 1: 69–97. https://doi.org/10.1037/h0076835

35. E. Hargittai. 2015. Is Bigger Always Better? Potential Biases of Big Data Derived from Social Network Sites. *The ANNALS of the American Academy of Political and Social Science* 659, 1: 63–76. https://doi.org/10.1177/0002716215570866

36. Eszter Hargittai. 2008. Whose Space? Differences Among Users and Non-Users of Social Network Sites. *Journal of Computer-Mediated Communication* 13, 1: 276–297. https://doi.org/10.1111/j.1083-6101.2007.00396.x

37. Gillian R. Hayes. 2011. The relationship of action research to human-computer interaction. *ACM Transactions on Computer-Human Interaction* 18, 3: 1–20. https://doi.org/10.1145/1993060.1993065

38. C. Heath and P. Luff. 1992. Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *Computer Supported Cooperative Work* 1, 1-2: 24 – 48. https://doi.org/10.1007/BF00752451

39. Brent Hecht, Lichan Hong, Bongwon Suh, and Ed H. Chi. 2011. Tweets from Justin Bieber's Heart: The Dynamics of the "Location" Field in User Profiles. In *Proc CHI*, 237–246. https://doi.org/10.1145/1978942.1978976
40. Michael A. Hogg and Scott A. Reid. 2006. Social identity, self-categorization, and the communication of group norms. *Communication Theory* 16, 1: 7–30. https://doi.org/10.1111/j.1468-2885.2006.00003.x
41. Christine Horne. 2001. Sociological Explanations of the Emergence of Norms. In *Social Norms*, Michael Hechter and Karl-Dieter Opp (eds.). Russell Sage, New York, 3–34.
42. William Jones, Anne Diekema, Jaime Teevan, Manuel Pérez-Quiñones, Jesse David Dinneen, and Bradley Hemminger. 2015. "For Telling" the Present: Using the Delphi Method to Understand Personal Information Management Practices. In *Proc CHI*, 3513–3522. https://doi.org/10.1145/2702123.2702523
43. Ronald Kline. 2003. Resisting Consumer Technology in Rural America: The Telephone and Electrification. In *How Users Matter: The Co-construction of Users and Technology*, Nelly Oudshoorn and Trevor Pinch (eds.). MIT Press, Cambridge, MA, 51–66.
44. Cliff Lampe, Jessica Vitak, and Nicole Ellison. 2013. Users and Nonusers: Interactions between Levels of Facebook Adoption and Social Capital. In *Proc CSCW*, 809–819. https://doi.org/10.1145/2441776.2441867
45. Jon Landeta. 2006. Current validity of the Delphi method in social sciences. *Technological Forecasting and Social Change* 73, 5: 467–482. https://doi.org/10.1016/j.techfore.2005.09.002
46. Maria Knight Lapinski and Rajiv N. Rimal. 2005. An explication of social norms. *Communication Theory* 15, 2: 127–147. https://doi.org/10.1093/ct/15.2.127
47. Bruno Latour. 1993. Ethnography of a High Tech Case: About Aramis. In P. Lemonnier (ed.). Routledge, London, 372–398.
48. David Lazer, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, Tony Jebara, Gary King, Michael Macy, Deb Roy, and Marshall Var Alstyne. 2009. Life in the network: the coming age of computational social science. *Science* 323, 5915: 721–723. https://doi.org/10.1126/science.1167742
49. Amanda B. Lenhart. 2005. Unstable Texts: An Ethnographic Look at How Bloggers and Their Audience Negotiate Self-Presentation, Authenticity, and Norm Formation.
50. Karen Levy. 2015. The User as Network. *First Monday* 20, 11. https://doi.org/10.5210/fm.v20i11.6281
51. Kurt Lewin. 1946. Action Research and Minority Problems. *Journal of Social Issues* 2, 4: 34–46. https://doi.org/10.1111/j.1540-4560.1946.tb02295.x
52. Yvonna S. Lincoln and Egon G. Guba. 1985. *Naturalistic Inquiry*. Sage Publications, Newbury Park, CA.
53. Harold A Linstone and Murray Turoff. 1975. *The Delphi Method*. Addison-Wesley, Reading, MA.
54. J. Lofland and L. Lofland. 1995. *Analyzing Social Settings: A Guide to Qualitative Observation and Analysis*. Wadsworth, Belmont, CA.
55. Jennifer Mankoff, Jennifer A Rode, and Haakon Faste. 2013. Looking Past Yesterday's Tomorrow: Using Futures Studies Methods to Extend the Research Horizon. In *Proc CHI*, 1629–1638. https://doi.org/10.1145/2470654.2466216
56. Joseph P. Martino. 1970. The consistency of Delphi forecast. *Futurist* 4, 2: 63–64.
57. Alice E. Marwick and danah boyd. 2010. I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. *New Media & Society* 13, 1: 114–133. https://doi.org/10.1177/1461444810365313
58. Matt McKeon. 2010. The Evolution of Privacy on Facebook. Retrieved January 11, 2016 from http://mattmckeon.com/facebook-privacy/
59. S.B. Merriam. 1998. *Qualitative research and case study applications in education*. Jossey-Bass, San Francisco, CA.
60. S Milgram. 1963. Behavioral Study of Obedience. *Journal of Abnormal Psychology* 67, 4: 371–378. https://doi.org/10.1037/h0040525
61. Michael Muller, Shion Guha, Eric P. S. Baumer, David Mimno, and N. Sadat Shami. 2016. Machine Learning and Grounded Theory Method: Convergence, Divergence, and Combination. In *Proc. GROUP*.
62. Michael J Muller and Allison Druin. 2012. Participatory Design: The Third Space in HCI. In *The Human-Computer Interaction Handbook*, Julie Jacko (ed.). Lawrence Erlbaum Associates, Hillsdale, NJ, 1–70.

63. Blair Nonnecke and Jenny Preece. 2001. Why lurkers lurk. In *Americas Conference on Information Systems, Boston*, Paper 294. Retrieved April 2, 2011 from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.92.183&rep=rep1&type=pdf

64. Kurt Opsahl. 2010. Facebook's Eroding Privacy Policy: A Timeline. *Electronic Frontier Foundation*. Retrieved January 11, 2016 from https://www.eff.org/deeplinks/2010/04/facebook-timeline

65. Wanda J. Orlikowski and Jack J. Baroudi. 1991. Studying Information Technology in Organizations: Research Approaches and Assumptions. *Information Systems Research* 2, 1: 1–28. https://doi.org/10.1287/isre.2.1.1

66. Ethan R. Plaut. 2015. Technologies of avoidance: The swear jar and the cell phone. *First Monday* 20, 11. https://doi.org/10.5210/fm.v20i11.6295

67. Laura Portwood-Stacer. 2012. How We Talk About Media Refusal, Part 1: "Addiction." *Flow* 16, 3.

68. Laura Portwood-Stacer. 2012. How We Talk about Media Refusal, Part 2: "Asceticism." *Flow* 16, 6. Retrieved from http://flowtv.org/2012/09/media-refusal-part-2-asceticism/

69. Laura Portwood-Stacer. 2013. Media Refusal and Conspicuous Non-Consumption: The Performative and Political Dimensions of Facebook Abstention. *New Media & Society* 15, 7: 1041–1057. https://doi.org/10.1177/1461444812465139

70. Tom Postmes, Russell Spears, and Martin Lea. 2000. The Formation of Group Norms in Computer-Mediated Communication. *Human Communication Research* 26, 3: 341–371. https://doi.org/10.1111/j.1468-2958.2000.tb00761.x

71. Jenny Preece, Blair Nonnecke, and Dorine Andrews. 2004. The top five reasons for lurking: improving community experiences for everyone. *Computers in Human Behavior* 20, 2: 201–223. https://doi.org/10.1016/j.chb.2003.10.015

72. Lee Rainie, Aaron Smith, and Maeve Duggan. 2013. *Coming and Going on Facebook*. Washington, D.C. Retrieved from http://www.pewinternet.org/~/media//Files/Reports/2013/PIP_Coming_and_going_on_facebook.pdf

73. Beatrice Rammstedt and Oliver P. John. 2007. Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German. *Journal of Research in Personality* 41, 1: 203–212. https://doi.org/10.1016/j.jrp.2006.02.001

74. Rajiv N. Rimal and Kevin Real. 2003. Understanding the influence of perceived norms on behaviors. *Communication Theory* 13, 2: 184–203. https://doi.org/10.1111/j.1468-2885.2003.tb00288.x

75. Mattias Rost, Louise Barkhuus, Henriette Cramer, and Barry Brown. 2013. Representation and communication. In *Proc CSCW*, 357–362. https://doi.org/10.1145/2441776.2441817

76. J.A. Roth and J.T. Scholtz (eds.). 1989. *Taxpayer compliance: Social science perspectives*. University of Philadelphia Press, Philadelphia, PA.

77. Gene Rowe and George Wright. 1999. The Delphi technique as a forecasting tool: issues and analysis. *International Journal of Forecasting* 15: 353–375. https://doi.org/10.1016/S0169-2070(99)00018-7

78. Tracii Ryan and Sophia Xenos. 2011. Who uses Facebook? An investigation into the relationship between the Big Five, shyness, narcissism, loneliness, and Facebook usage. *Computers in Human Behavior* 27, 5: 1658–1664. https://doi.org/10.1016/j.chb.2011.02.004

79. Harold Sackman. 1975. *Delphi Assessment, Expert Opinion, Forecasting, and Group Processes*. Rand Corporation, Santa Monica, CA.

80. Margarete Sandelowski. 1993. Rigor or rigor mortis: The problem of rigor in qualitative research revisited. *Advances in Nursing Science* 16, 2: 1–8.

81. Christine Satchell and Paul Dourish. 2009. Beyond the user: use and non-use in HCI. In *Proc OZCHI*, 9–16. https://doi.org/10.1145/1738826.1738829

82. Sarita Yardi Schoenebeck. 2014. Giving up Twitter for Lent: How and Why We Take Breaks from Social Media. In *Proc CHI*, 773–782. https://doi.org/10.1145/2556288.2556983

83. Neil Selwyn. 2003. Apart from technology: understanding people's non-use of information and communication technologies in everyday life. *Technology in Society* 25, 1: 99–116. https://doi.org/10.1016/S0160-791X(02)00062-3

84. Shilad Sen, Margaret E. Giesel, Rebecca Gold, Benjamin Hillmann, Matt Lesicko, Samuel Naden, Jesse Russell, Zixiao "Ken" Wang, and Brent Hecht. 2015. Turkers, Scholars, "Arafat" and "Peace": Cultural Communities and Algorithmic Gold Standards. In *Proc CSCW*, 826–838. https://doi.org/10.1145/2675133.2675285

85. Stefan Stieger, Christoph Burger, Manuel Bohn, and Martin Voracek. 2013. Who commits virtual identity suicide? Differences in privacy concerns, internet addiction, and personality between facebook users and quitters. *Cyberpsychology, Behavior, and Social Networking* 16, 9: 629–34. https://doi.org/10.1089/cyber.2012.0323

86. J. Stromer-Galley and R. M. Martey. 2009. Visual spaces, norm governed places: the influence of spatial context online. *New Media & Society* 11, 6: 1041–1060. https://doi.org/10.1177/1461444809336555

87. L.A. Suchman. 1987. *Plans and Situated Actions: The problem of human-machine communication*. Cambridge University Press, Cambridge.

88. Zeynep Tufekci. 2008. Grooming, Gossip, Facebook and MySpace. *Information, Communication & Society* 11, 4: 544–564. https://doi.org/10.1080/13691180801999050

89. Johan Ugander, Lars Backstrom, Cameron Marlow, and Jon Kleinberg. 2012. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences* 109, 16: 5962–5966. https://doi.org/10.1073/pnas.1116502109/-/DCSupplemental.www.pnas.org/cgi/doi/10.1073/pnas.1116502109

90. Ron Wakkary and Karen Tanenbaum. 2009. A sustainable identity: the creativity of an everyday designer. In *Proc CHI*, 365–374. https://doi.org/10.1145/1518701.1518761

91. Yang Wang, Saranga Komanduri, Pedro Giovanni Leon, Gregory Norcie, Alessandro Acquisti, and Lorrie Faith Cranor. 2011. "I regretted the minute I pressed share": A Qualitative Study of Regrets on Facebook. In *Proc SOUPS*. https://doi.org/10.1145/2078827.2078841

92. Gordon Welty. 1971. A Critique of the Delphi Method. In *The Joint Statistical Meeting of American Statistical Association*, 377–382.

93. Andrea Wiggins and Kevin Crowston. 2015. Surveying the Citizen Science Landscape. *First Monday* 20, 1. https://doi.org/10.5210/fm.v20i1.5520

94. John W. Williamson and Mart G. van Nieuwenhuijzen. 1974. Health Benefit Analysis: An Application in Industrial Absenteeism. *Journal of Occupational Medicine* 16, 4: 229–233.

95. Sally Wyatt. 2003. Non-Users Also Matter: The Construction of Users and Non-Users of the Internet. In *How Users Matter: The Co-construction of Users and Technology*, Nelly Oudshoorn and Trevor Pinch (eds.). MIT Press, Cambridge, MA, 67–79.

96. S. Wyche and E. P. S. Baumer. 2016. Imagined Facebook: An exploratory study of non-users perceptions of social media in Rural Zambia. *New Media & Society*. https://doi.org/10.1177/1461444815625948

97. Susan P. Wyche, Paul M. Aoki, and Rebecca E. Grinter. 2008. Re-placing faith: reconsidering the secular-religious use divide in the United States and Kenya. In *Proc CHI*, 11–20. https://doi.org/10.1145/1357054.1357057

98. Susan P. Wyche, Sarita Yardi Schoenebeck, and Andrea Forte. 2013. "Facebook is a Luxury": An Exploratory Study of Social Media Use in Rural Kenya. In *Proc CSCW*, 33–43. https://doi.org/10.1145/2441776.2441783

99. Xuan Zhao, Victoria Schwanda Sosik, and Dan Cosley. 2012. It's Complicated: How Romantic Partners Use Facebook. In *Proc CHI*, 771–780. https://doi.org/10.1145/2207676.2207788