Facebook as a Research Tool for the Social Sciences

Article *in* American Psychologist · September 2015

DOI: 10.1037/a0039210

5 authors, including:

Sandra Matz
Columbia University

Samuel D Gosling
University of Texas at Austin

Vesselin Popov
University of Cambridge

David Stillwell
University of Cambridge

Facebook as a research tool for the Social Sciences: Opportunities, Challenges, Ethical Considerations, and Practical Guidelines

Michal Kosinski

Stanford University, Stanford, USA

Sandra C. Matz

University of Cambridge, Cambridge, UK

Samuel D. Gosling

University of Texas, Austin, USA

School of Psychological Sciences, University of Melbourne, Parkville, VIC, Australia

Vesselin Popov, David Stillwell

University of Cambridge, Cambridge, UK

Author Note

The authors would like to thank Lindsay Graham, Robert Wilson, Kristin Laurin and Wu Youyou for their feedback on this paper.

Abstract

Facebook is rapidly gaining recognition as a powerful research tool for the social sciences. It constitutes a large and diverse pool of participants, who can be selectively recruited for both online and offline studies. Additionally, it facilitates data collection by storing detailed records of its users' demographic profiles, social interactions, and behaviors. With participants' consent, these data can be recorded retrospectively in a convenient, accurate, and inexpensive way. Based on our experience in designing, implementing, and maintaining multiple Facebook-based psychological studies that attracted over 10 million participants, we demonstrate how to recruit participants using Facebook, incentivize them effectively, and maximize their engagement. We also outline the most important opportunities and challenges associated with using Facebook for research; provide several practical guidelines on how to successfully implement studies on Facebook; and finally, discuss ethical considerations.

Keywords: Facebook, research design, ethics, big data, myPersonality, snowball sampling

Facebook as a research tool for the Social Sciences: Opportunities, Challenges, Ethical Considerations, and Practical Guidelines

Facebook, the largest and most popular online social network, has become a significant part of daily life for nearly 1.4 billion people around the world (Facebook Inc., 2015). An increasing number of studies focus on Facebook's influence on individuals and societies (for a comprehensive review, see Wilson, Gosling, & Graham, 2012). However, its potential as a powerful research tool—for both online and offline research in the social sciences—has been largely overlooked. This is unfortunate, because Facebook provides a number of tools that can be used to inexpensively recruit large and diverse samples helping to address a major challenge in social science: its overreliance on relatively small, student,[1] and disproportionately WEIRD samples (Henrich, Heine, & Norenzayan, 2010). Furthermore, Facebook can be used as a powerful data-recording tool because it stores detailed demographic profiles and records of an enormous amount of users' *actual* behavior expressed in a natural environment. Investigators can, with users' consent, record their data retrospectively, greatly reducing the limitations associated with self-reported and laboratory-based studies (Paulhus & Vazire, 2007).

We discuss distinct opportunities and challenges offered by Facebook to researchers; provide a number of practical recommendations for effectively conducting research within this environment; and discuss several ethical considerations. We argue that using Facebook in research is often relatively straightforward and generally produces robust results. Contrary to a widespread belief, it rarely requires the development of a dedicated Facebook app, or substantial changes to existing research procedures. Benefiting from Facebook features could be as easy as

[1] Some estimates indicate that 85% of psychological studies are based on exclusively undergraduate samples (Gosling et al., 2004).

posting an advert on Facebook that targets a specific sub-population; or replacing an online or offline demographic survey with a button that enables participants to donate their behavioral and demographic data (stored on a Facebook profile) to the researchers.

However, to fully tap into the large samples and rich data offered by Facebook, researchers may need to rethink traditional research designs and acquire new skills, such as the basics of programming and web design. Similarly, because participants can easily (with just one click) abandon online studies, it is vital to focus on the needs and the experience of the participants. Finally, some of Facebook's advantages (such as easy access to large amounts of personal data) introduce serious ethical challenges that have yet to be addressed by the law and official ethical guidelines.

This paper is not a complete compendium of Facebook research methods, but we hope that it will encourage others to consider using Facebook as a research tool, benefiting from our experience and avoiding our mistakes. We heavily draw from our own experience in designing and maintaining several Facebook studies that have attracted more than 10 million participants. The most popular of our projects, myPersonality, was set up by David Stillwell in 2007 (Stillwell & Kosinski, 2015). It offered participants access to 25 psychological tests ranging from the 300-item IPIP proxy for the NEO-PI-R inventory, to the Satisfaction With Life Scale and a computer-adaptive proxy of Raven's Standard Progressive Matrices. The participants received immediate feedback on their results, and could volunteer to share their Facebook profiles with us; approximately 30% of them (i.e. over 2 million) decided to do so. Recognizing the generosity of the participants who donated their data for research purposes, with their consent, we have shared this resource with the academic community. The website, *http://mypersonality.org*, contains dozens of cleaned, pre-processed, and anonymized databases[2], a knowledge base, and examples of syntax useful in analysis. Nearly 200 researchers from over 100 academic institutions have used the data in their work.

Facebook has been the most popular social network globally since April 2008, but it will not necessarily remain so. The previously most popular social network, Friendster, was overtaken by Myspace, which was itself surpassed by Facebook (CNET, 2014). Nevertheless, although the platforms may replace one another, we predict that in the foreseeable future, there will still be online social networking. Some of our Facebook-specific recommendations might become outdated, but we expect that many will remain applicable to whichever platform takes Facebook's place.

Recruiting and retaining participants

The size and reach of the Facebook platform offers researchers an unprecedented opportunity to acquire large and diverse samples of participants. Evidently, the Facebook population is not perfectly representative; its users tend to be younger, better educated, and some groups might be entirely excluded (e.g., Amish, people without Internet access, or living in countries that block Facebook). However, the sheer size of Facebook's population implies that even the underrepresented populations are relatively large. For example, as of 2014, nearly 35% of America's older adults (over 65 years of age) were on Facebook, and their number was quickly growing (Duggan, Ellison, Lampe, Lenhart, & Madden, 2015; Pew Research Center, 2014). Furthermore, Facebook's global reach supports the investigation of cultural differences that are often overlooked in traditional studies (Heine, Lehman, Peng, Greenholtz, & Of, 2002). For instance, Facebook samples have been used to examine cultural differences in self-

[2] Information, such as names and Facebook IDs, were removed from the datasets and replaced by unique user IDs.

presentation in photographs (Huang & Park, 2013), as well as the establishment and maintenance of friendship networks (Peters, Winschiers-Theophilus, & Mennecke, 2013).

The following subsections discuss the ways in which Facebook can be used to recruit and retain large and diverse samples of participants.

Going viral: Recruiting participants using a snowball sampling approach

One of the least expensive ways of dipping into the Facebook participant pool is by snowball sampling (Goodman, 1961)—convincing Facebook users to invite their friends to join a study. If enough participants do so, the positive feedback loop may lead to self-sustaining studies with a rapid growth in sample size. For instance, the myPersonality project was originally only shared with the author's 150 Facebook friends, but it went viral and attracted over 6 million participants in four years.

Yet going viral is not easy. Only the most engaging studies, such as games or interesting psychological questionnaires that provide compelling feedback, are likely to go viral. This is because there are many other applications and websites (often very well-funded) that compete for users' attention. The popularity of myPersonality was certainly boosted by a novelty factor; it was the first application of its type to appear on Facebook. Similar personality testing applications released in the following years, often more advanced in terms of design and technology, attracted much less interest.

The chances of a study going viral can be increased by making "inviting friends" an integral part of the experience.[3] For example, myPersonality offered a 360-degree assessment feature, encouraging users to invite their friends to judge their personality. This resulted in a database of cross-ratings, while also helping to increase the virality of the application; those who were invited to rate their friends often proceeded to take a test themselves. Similarly, the *You Are What You Like* project[4] provides participants with feedback on their personalities, and also identifies friends that have similar personality profiles. Many of the participants share those reports on Facebook and tag their friends in the description, which attracts new participants. Note, however, that in an attempt to encourage application owners to spend money on paid advertising, Facebook can and does limit the exposure of posts created by apps and websites.

Importantly, snowball sampling methods do not meet the gold standard of randomized sampling, because the method can introduce biases (Kurant, Markopoulou, & Thiran, 2011). The first participants (seeds) are likely to disproportionately affect the composition of the sample, because people tend to interact with others similar to themselves (McPherson, Smith-Lovin, & Cook, 2001). Furthermore, people with many friends are more likely to be recruited into the sample (Kurant, Markopoulou, & Thiran, 2010). Yet most psychological research does not use randomized samples, instead relying on ad-hoc samples, often of undergraduate Psychology students or professional test-takers recruited on Amazon's Mechanical Turk (AMT) (Chandler, Mueller, & Paolacci, 2014). Even supposedly representative samples, obtained from panel data, are affected by self-selection biases because only certain types of people sign up to these panels (Blumberg & Luke, 2009).

In this context, Facebook samples provide an inexpensive and relatively high-quality alternative (Baltar & Brunet, 2012). For example, the age distribution of the myPersonality snowballed sample is not more biased than traditional paper-and-pencil studies published in the *Journal of Personality and Social Psychology* (Gosling, Vazire, Srivastava, & John, 2004).

---

[3] Facebook provides advice and numerous case studies that help to understand how to make one's content viral. It also offers robust documentation and samples of code ready to be copied and pasted into one's study. See: http://developers.facebook.com.

[4] http://www.youarewhatyoulike.com

Additionally, the size and diversity of the Facebook population help to minimize the disadvantages of snowball sampling. First, given enough participants, the representativeness of the population can be improved by weighting. For example, participants over 55 years of age are underrepresented in the myPersonality sample, but it still contains nearly 100,000 of them (Stillwell & Kosinski, 2015). Second, employing virality based on an intrinsic motivation means that people share and participate in a study out of personal interest rather than financial gain, which is thus likely to yield better data of higher quality (see the section: *Improving the quality of the self-reports*). Third, social network information provided by Facebook can be used to control for the fact that the sample was collected in a non-random way. For example, respondent-driven sampling provides a mathematical model that weights the sample to compensate for the fact that it was collected in a non-random way (Salganik & Heckathorn, 2004). Finally, as discussed in the next section, researchers can reduce biases by starting with a diverse set of initial participants (David L., 2008, pp. 816–817).

Targeting: Recruitment through Facebook advertising

Facebook's advertising platform[5] provides an alternative to snowball sampling, albeit at a price. Researchers can create and publish an advert that promotes their offline study, website, game, online questionnaire, or Facebook application. These adverts can be targeted at users that are defined by a wide range of preferences (e.g. liking *"getting up early in the morning")*; behaviors (liking *"running"*); and demographic variables including location, education, language, political views, ethnicity, sexual orientation, income, and many more.

This approach is particularly useful when seeking participants that are otherwise hard to reach, including those stigmatized in the offline world (Mangan & Reips, 2007); or those who are hesitant to meet researchers face-to-face (Batterham, 2014). For example, an advert could be targeted at 40 to 50-year-old men of Hispanic origin, who live in the U.S., and who are interested in relationships with other men. At the time of writing, the Facebook advertising platform reported 11,000 individuals matching this profile, who could be reached with an estimated cost-per-click of between $0.40 and $0.80.

Moreover, targeted advertising can help protect participants' anonymity and decrease the number of (potentially sensitive) questions that have to be asked. In the example presented above, only users matching the targeted profile would see and be able to follow the advert in the first place. Consequently, there is no need to identify the participants, record their Facebook data, or pose questions related to their sexual orientation, race, gender, or age.

Previous research shows that Facebook advertising can help reduce the costs of targeted participant recruitment for online surveys (Batterham, 2014) and offline studies (Close, Smaldone, Fennoy, Reame, & Grey, 2013), as well as build long-term participant pools (K. J. Johnson, Mueller, Williams, & Gutmann, 2014). Not only did Facebook ads outperform traditional methods such as postal surveys (Batterham, 2014), but they were also more cost-efficient than Google advertising, online newsletters, and emails (Carlini, Safioti, Rue, & Miles, 2014). Based on 10 recent studies (identified on *Scopus.com* at the time of writing) that reported using the Facebook advertising platform, the average cost-per-participant was $13.75, with a range of $1.51 (Batterham, 2014) to $33 (Richiardi, Pivetta, & Merletti, 2012). Given the significant range of costs, we offer the following advice on how to keep Facebook advertising costs to a minimum.

Advertisement type. Depending on the type of advert, the costs of advertising on Facebook vary considerably. Directing users to internal Facebook pages is generally less expensive than directing users to external websites, because the social network has an interest in

[5] See http://www.facebook.com/ads

keeping users within the platform. Furthermore, promoting a Facebook page, rather than an external one, makes it easier for participants to share the study with their friends. Batterham (2014) was able to reduce the cost-per-participant from $9.82 to $1.51 by redirecting users to a Facebook page that contained a link to the external online survey, rather than promoting the external link directly.

Bidding. Facebook asks advertisers to “bid” or set the maximum price that they are willing to pay for each click or each impression. This value is used by a bidding mechanism that selects adverts to be published to users; adverts accompanied by higher bids are more likely to be published. Facebook suggests a bid range guaranteeing an effective campaign in a given targeted group. Bidding below that suggested range still usually results in adverts being published, but it may take longer for an ad campaign to be completed. In our experience, it is possible to receive a desired number of clicks with half of the lowest recommended bid value. Besides, if the recruitment speed proves to be unsatisfactory, the bid can always be increased at a later time.

Target group. The cost-per-click (and therefore cost-per-participant) substantially varies across targeting criteria such as geographical location, demographic traits, or behaviors and preferences. For example, the suggested bid per click for women with a doctorate degree located in the U.S. was $0.43, compared to $0.35 in the UK. Thus, whenever feasible, researchers may consider recruiting participants from populations that are cheaper to reach.

Tracking participants across studies

Another Facebook feature useful in recruitment is the ability to use a Facebook identification number to track participants across time and studies. This feature could be also used to protect privacy, as researchers can efficiently recognize participants without necessarily storing any personally identifiable information.

Participant tracking could also reduce errors and help save participants' time by minimizing the need to frequently re-enter the same information (e.g. basic demographic facts) in related studies. Additionally, participant tracking turns repeated submissions—a challenge in traditional web-based studies, as noted by Birnbaum, 2004—into an opportunity for collecting longitudinal data. For example, myPersonality's implementation of the IPIP proxy for NEO-PI-R was *retaken* over 1.15 million times, and over 5,000 participants answered this questionnaire more than 10 times. This kind of retest information can be used in longitudinal studies to examine substantive changes (e.g. in personality), and to evaluate the psychometric properties of the questionnaire, such as its test-retest reliability or the degree of misrepresentation.

Staying connected: building an ex-participant community

Facebook features can be used to remain in contact with ex-participants, enabling researchers to build useful communities. For instance, users who "Liked" the Facebook page of the study or laboratory will occasionally[6] see the updates published by the researchers. Over the years, over 150,000 people have Liked myPersonality's Facebook page. The posts published there attracted considerable attention—in just a matter of hours, it was possible to recruit tens of thousands of users to participate in a new questionnaire or experiment. Although such groups are unavoidably biased (i.e. are comprised of people most interested in researchers' work), they may create a powerful starting sample for a broad viral recruitment. Additionally, users' comments on the researchers' posts provide important feedback on the design of and potential problems with the studies, which proved very useful in the myPersonality project.

Researchers should be aware that the action of joining a Facebook group or liking a Facebook page is visible to respondents' friends. This could have adverse effects; for example, by joining an adoption study community, users may inadvertently reveal that they were adopted. Therefore, researchers should discuss with their Institute's Review Board whether it is appropriate to set up a public community.

Collecting Facebook profile data

The previous section focused on the benefits that Facebook offers in the context of recruiting and communicating with participants. Here we introduce another strength of this environment: access to an extraordinary amount of information about individuals and groups, recorded in a natural environment.

Facebook profile information includes self-reported information (e.g., schools attended, current workplace, age or gender); traces of behavior (e.g. status updates or Likes); and data contributed by the others (e.g. photo tags or comments on a user's wall). These data can be recorded retrospectively, and thus overcome shortcomings of participants' memory and biases during their participation in the study. For example, even the most motivated participants are unlikely to have enough time, attention, and knowledge to reliably report on past events attended (e.g. used in: Han et al., 2012); the natural language used in their day-to-day conversations (e.g. used in: Schwartz et al., 2013); or the shape of their own egocentric network. The latter often encompasses hundreds of agents and thousands of friend connections (e.g. used in: Aral & Walker, 2012; Arnaboldi, Guazzini, & Passarella, 2013).

The full list of variables available to researchers through Facebook constantly changes as the platform expands and modifies its policies. At the time of writing, the following categories of information can be recorded with users' permission:

[6] In 2012, each post on a Facebook page was viewed by 16% of the page's fans, but in 2014 this has reduced to 6.5%. This is partly due to an increase in the number of Facebook pages, and partly due to Facebook prioritizing paid news feed promotions. See: http://techcrunch.com/2014/04/03/the-filtered-feed-problem

1. **Demographic profile,** comprising a unique user ID, full name, profile picture, age, gender, relationship status, romantic interests, geographical location, place of origin, work and education history, biography, link to personal website, time zone, political and religious views, general interests; and lists of favorite music, movies, TV shows, books, quotes, and sports.
2. **User-generated content** consisting of status updates, photos, videos, comments on other people's content or pages, links, and notes published by users or their friends. Each piece of content also contains metadata, such as the positions of people present in the picture, date of publication, list of people who Liked it, its privacy settings, etc.
3. **Social network structure** containing a list of friends and the type of users' connections. Connection types include friendships, family links (e.g. spouse, siblings, parents, or children), and followers.[7]
4. **User preferences and activities** comprising their Likes,[8] group memberships, attended events, installed applications, and tags in photos or posts.[9]
5. **Information about users' friends**, such as demographic details and friends' activities that are visible to a given user.
6. **Private messages** between users, usually written and sent through the instant messenger feature.

Facebook collects other forms of data in addition to those listed above, such as users' visits on others' profiles and the time that the user signed up to Facebook. However, at present, it

[7] A one-way connection, similar to following on Twitter, where the follower receives updates about the followed person, but not the other way round.

[8] Likes can be used by Facebook users to endorse content such as status updates, comments, photos, links shared by friends, advertisements, Facebook pages, or external websites. Endorsements also result in users receiving updates on a given piece of content, such as comments on a Liked status update or news published by a Liked Facebook page. Likes were introduced on February 9, 2009.

does not share this information with users or third parties. The full list of accessible variables is available at: http://developers.facebook.com/docs/graph-api/reference.

Interestingly, Facebook data offer insights reaching well beyond the boundaries of the platform. Liking content, posting photos, and interacting with friends represent behavioral patterns that are unlikely to be limited to Facebook. Adding a Facebook friend is, in most cases, merely a shadow of an interaction in the real world. Similarly, liking a website, musician, or activity is usually a proxy for behavior in other environments, such as visiting websites, listening to songs, or engaging in those activities. Moreover, Facebook stores data explicitly reaching into other online and offline environments. Examples include lists of real-world events attended; check-ins at geographical locations; and the Likes generated by clicking a button located on an external website or by scanning a QR code.

Data collected from Facebook profiles can be further enriched using external models and services. For example, Pennebaker's Linguistic Inquiry and Word Count (Pennebaker, Francis, & Booth, 2001) can be used to tag status updates with their emotional valence; and www.applymagicsauce.com (Kosinski, Stillwell, & Graepel, 2013) can translate Facebook Likes into estimates of personality, intelligence, satisfaction with life, gender, age, political and religious views, and a number of other traits. Such procedures require skills and effort, but they can save participants' time and reveal information that is difficult to access otherwise.

Pitfalls of using Facebook profile data

Recording Facebook profile data offers an unprecedented opportunity to observe people in realistic environments. However, a major challenge in studying observational data, such as Facebook profile data, is to draw conclusions that are acceptably free from influences by overt

[9] Users can label or "tag" each other in photos and posts. Users can review/remove tags related to them (i.e. "untag").

biases, and to assess the impact of potential hidden biases. Two major issues have to be considered before drawing conclusions from Facebook profile data.

First, some parts of the Facebook profile are self-reported and other parts (e.g. behavioral traces) can be selectively removed by the profile owner. Thus, the quality of Facebook profile data may be affected by user-induced biases typical for self-reports, such as social desirability and intentional misrepresentation. Empirical evidence, however, indicates that Facebook profiles contain valid information about their creators. For example, Back et al. (2010) showed that Facebook profiles reflect the actual and not self-idealized personality profile of their owners. Also, Kosinski et al. (2013) showed that other elements of Facebook (such as gender, age, political views, religion, and Facebook Likes) show consistent and meaningful relationships within Facebook profiles. Finally, in a subset of myPersonality participants (n=28,628), where both the Facebook profile and self-reported information about gender were available, they matched in 98.8% of the cases. A somewhat lower accuracy (95%; n=18,321) was recorded for age.[10] Those results suggest that biases typical for self-reports are not particularly strong in this environment. This could stem from the fact that Facebook profiles are generated outside of a research context. Also, as Facebook friendships are usually preceded by real-world interactions (Lampe, Ellison, & Steinfield, 2006; Ross et al., 2009), inaccurate or enhanced profile invalid information may be difficult to maintain in a network of friends who can challenge false assertions (Lampe, Ellison, & Steinfield, 2007; Pempek, Yermolayeva, & Calvert, 2009).

Second, Facebook profile data is, to some extent, affected by the mechanics of the platform. The user experience on Facebook is highly personalized. This personalization includes

---

[10] Several of the questionnaires published by myPersonality included questions related to participants' age and gender. Neither of these sources can be treated as a "ground-truth" data, but potential inconsistencies provide a lower-bound estimate of the invalidity: inconsistent data are inevitably invalid. Note that providing matching yet

the selection of news feed stories, advertisements, and even friend suggestions. Algorithms regulating individual experience are constantly evolving and, most likely, function differently for different users.[11] As users are more likely to interact with content and people suggested to them by Facebook, their behavior is driven not only by their intrinsic goals and motivations, but also (to some unknown extent) by the Facebook algorithms constantly adjusting their exposure to content and friends. For example, friends' photos that appear on a given user's Facebook news feed are clearly more likely to be Liked. Essentially, largely unknown effects of personalization represent a general class of confounding variables characteristic for observational research and deserve further study.[12] Results likely to be affected by Facebook algorithms should be interpreted with caution.

Furthermore, profile entries such as job position or political views are recorded using open text input fields, which are often equipped with an auto-completion mechanism (in the form of a pop-up list) that suggests potential entry values based on the first few letters typed in. This may improve the quality of the data (e.g. by discouraging spelling errors), but it can also introduce biases. For example, a user might have intended to identify himself as being a *Social Psychologist*, but may settle on *Social Scientist* if Facebook suggests the latter based on the first few letters. Consequently, a mere change in an auto-completion mechanism could affect the relative frequencies of the entries, and be misinterpreted as a real-world phenomenon (i.e. a sudden decrease in people's self-identification as social psychologists).

---

incorrect values is unlikely and pointless, given that those who did not want to reveal their real age or gender could simply skip relevant questions or hide the information from their Facebook profiles.

[11] For example, due to experiments conducted by a Facebook administrator to improve user experience or investigate their behavior (referred to as A/B testing).

[12] Researchers could also attempt to collaborate or discuss their work with internal Facebook research teams who may be better informed about the functioning of Facebook algorithms and could help to avoid misinterpretations.

Also, Facebook itself has changed considerably over the years, driving changes in users' behavior. Likes are now a staple of the platform, despite only being introduced in 2009. Originally, users would list their favorite "activities", "interests", and other preferences in free-text fields; after 2009, these entries were converted into Facebook Likes. Another change concerns status updates, which were originally preceded with the user's name (i.e. "*Alice Miller is...*"), encouraging updates written in the third person. Nowadays, status updates are free to contain any text.

Additionally, users can use privacy settings to limit access to parts of their profiles, and the platform-wide privacy rules are constantly evolving. Initially, users could see and interact with everyone in the same university network, whereas today the visibility and interactions are usually limited—by default—to users' immediate circle of friends. As with changes in auto-completion mechanisms, there is a risk that evolving privacy rules of the platform could be misinterpreted as a real-world phenomenon. For example, one may conclude that a certain group of users are deciding to make more of their content private, when in fact the default visibility setting of the content may have been altered for that group by Facebook, as part of a batch-release of new features.

Finally, it is easy to create a fake profile on Facebook, and use it to participate in the research. However, such profiles are relatively easy to detect (Yang et al., 2011). For example, real users accumulate their friends and Likes over a long period of time, but fake profiles are likely to be filled with Likes and friends added in a single burst of activity. Also, egocentric networks of real users are relatively dense (i.e. their friends tend to know each other) and clustered in terms of geography and institutions (e.g. schools or employers). Friends of the fake users, on the other hand, are unlikely to know each other or be collocated. It is also relatively unlikely that participants will make the effort to use a fake account to participate in a study, when it is possible to simply adjust the privacy settings to prevent third parties, including researchers, from recording their data.

Collecting self-reports from Facebook users

Collecting self-reports from Facebook users is similar to collecting such data in other online environments. In fact, an existing online survey or questionnaire can be integrated with Facebook merely by adding a fragment of HTML code ("Log in with Facebook" button[13]), allowing researchers to identify participants by their Facebook user ID and, with their consent, record selected data from their Facebook profile.

Using Facebook to collect self-reports offers a number of advantages. As discussed before, participants can be reliably identified and tracked across studies, without necessarily revealing any personal details apart from their Facebook user ID. Moreover, many of the typical questions (e.g. related to demographics) can be skipped entirely, because much data can be obtained directly from the Facebook profile or inferred from the targeting approach used to promote the link to the study.[14] Such an approach saves time and effort for both participants and researchers; helps to retain participants' attention; and prevents typing errors or other unintentional mistakes. Even if the desired information is not directly available through the Facebook profile data, it can often be inferred. For example, myPersonality did not contain any questions related to views on environmental issues, but such information can be extracted from other variables, including participants' Likes or status updates (e.g. liking "*Stop Global Warming*" or posting about it). Similarly, rather than collecting participants' responses to personality items such as "*I make friends easily*", researchers could measure the size and growth of their Facebook

[13] See: http://developers.facebook.com/docs/facebook-login
[14] See the section: *Targeting: Recruitment through Facebook advertising*

friendship networks; analyze the emotional valence of friends' comments on their activities; or count the number of times they were tagged in other people's pictures (Kosinski, Bachrach, Kohli, Stillwell, & Graepel, 2013; Kosinski, Stillwell, et al., 2013; Youyou, Kosinski, & Stillwell, 2015).

Finally, it is important to recognize that participants may have disabilities (e.g. vision impairment) and use a variety of devices when accessing the study. For example, a growing number of people access the Internet (and Facebook) using mobile devices. It is thus important to make sure that a study presents itself well on large as well as small screens, and that the text and buttons are large enough (or scalable) to accommodate participants with disabilities. Naturally, such a requirement applies to all research, not just studies mediated via the Internet. In fact, Internet studies, which do not require participants to physically come to a research venue, are often more accessible for disabled individuals.

Improving the quality of the self-reports

Facebook offers a number of advantages when collecting self-reported data, but it shares some of the limitations of online research. First, low barriers to access the study, ease of responding, and instant feedback may encourage the participants to rush through the study without paying as much attention (J. A. Johnson, 2005; Kurtz & Parrish, 2001). Second, researchers have little or no control over the circumstances in which the study is being accessed, so it is possible that some participants simultaneously engage in other activities. Third, the lack of face-to-face contact increases the psychological distance between the researcher and the participants, which may decrease participants' feeling of accountability (Caspi & Gorsky, 2006; Gosling et al., 2004). Finally, because Facebook participants come from diverse backgrounds, they may misunderstand instructions or test questions due to linguistic or cultural differences (J. A. Johnson, 2004, 2005).

In our experience, nevertheless, Facebook samples produce self-reported data of very high quality. The reliability[15] of the psychometric scales published on the myPersonality application was similar to that reported in their manuals or respective standardization samples. For example, the average reliability of our implementation of the 100-item version of the IPIP proxy for NEO-PI-R equaled 0.91 (Stillwell & Kosinski, 2015), compared with 0.89 reported for the standardization sample (Goldberg, 1999). The same effect was observed for the 300-item version of this questionnaire (0.84, compared with 0.80). Similarly good results were obtained for the discriminant validity[16] (John & Benet-Martinez, 2000), which was similar to the one obtained in paper-and-pencil samples (0.16, compared with 0.20; John & Srivastava, 1999; Stillwell & Kosinski, 2015). High reliability and discriminant validity in the myPersonality sample are even more remarkable given its diversity and international character. Though standardization samples are usually composed of carefully selected and highly motivated participants, myPersonality users were not controlled in any way and for many of them, English was not their first language. Finally, the quality of the self-reported data collected by myPersonality is also supported by the quality of the insights they offer. For example, Youyou et al. (2015) used self-reported personality scores to build a model predicting personality from respondents' Likes. The resulting personality predictions were more accurate than those made by

[15] Reliability conveys the consistency of response patterns across participants. There are many types of reliability coefficients, but here we report Cronbach's alphas. Fisher r-to-z transformation is used when averaging the reliability coefficients.

[16] Discriminant validity measures the degree to which scales are independent of each other. In most cases, scales should be independent and the response data breaching this assumption might indicate dishonest respondents being driven by factors such as social desirability. Discriminant validity is usually expressed as an average absolute Pearson product-moment correlation between scales of an instrument; the lower the value of this correlation, the higher the discriminant validity.

participants' spouses and family members.[17] The high predictive accuracy of the computer model could not have been achieved if the self-reports used in its training were not of high accuracy.

The high quality of the data collected by the myPersonality project does not guarantee that all other studies will be equally successful. However, the size of our sample, wide range of questionnaires used, and the longevity of the study suggest that it was not an accident. The example of myPersonality shows that participants recruited on Facebook, working in a non-controlled environment and motivated solely by obtaining feedback on their scores, can provide researchers with high-quality data. Below we present several of the factors that are often undervalued by social scientists but, in our opinion, were key to obtaining high-quality data in our projects.

Choosing the appropriate incentives. In our experience, providing feedback on the scores or performance is one of the most efficient ways of compensating participants for their time and effort. Participants motivated by financial rewards or credit points—as commonly studied undergraduate Psychology students or AMT workers are—might be inattentive, respond dishonestly, act unnaturally, or perform half-heartedly. Such behavior invalidates not only the results of the study, but also the feedback. Therefore, rewarding them with feedback alone promotes attentive and honest participation, improving the validity of the results.

Some studies do not produce interesting feedback, but researchers can include elements that do so. Such an approach not only provides motivation for the participants, but also allows researchers to collect additional, potentially useful data. For example, while developing a new general intelligence test, researchers can mix new test items with ones from an established

[17] The accuracy of human personality judgments in this study was consistent with previous research (Connelly & Ones, 2010).

measure. Consequently, they can reward participants with the feedback on the scores and collect valuable cross-validation data.[18]

Another approach to motivating participants was tested in the case of the longest measure implemented on myPersonality: a 300 item-long personality inventory that required up to one hour of a participant's attention. To get access to the questionnaire, participants were required to either fill out a series of uninteresting marketing surveys or make a donation ($4 USD) toward the cost of project hosting. This step severely limited (and presumably biased) the number of participants in this sub-population, but the resulting data had excellent psychometric properties (Stillwell & Kosinski, 2015).

In general, we believe that financial incentives—although popular among social scientists and boosting the sample size (Doody & Sigurdson, 2003)—have numerous disadvantages. First, they do not reward people for responding honestly or behaving naturally, but merely for participating in the study. This reward structure can be detrimental to the quality of the data, especially considering that self-reports can easily be filled with dishonest or random information. Second, financial incentives are not scalable and become prohibitively expensive if the required sample size grows. Third, financial incentives can attract semi-professional participants, who may have previous knowledge of the experimental design or materials used, and who have most likely participated in similar studies before (e.g. samples recruited through AMT or university students studying Psychology) (Chandler et al., 2014). It is often easier to use the grant money to pay the participants than to design an engaging study, but we believe that financial incentives should be avoided whenever possible, due to their potentially detrimental effect on data quality.

[18] Other examples of online studies offering enjoyable feedback include: http://www.beyondthepurchase.org; http://www.outofservice.com; http://www.celebritytypes.com/big-five; http://labs.five.com; and http://www.personalitylab.org

Avoiding any coercion. When designing studies, researchers are too often completely preoccupied with their goals, and tend to forget about the participant's experience. Consequently, participants are often forced to follow strict guidelines, have little control over the procedure, and are often presented with boring and tedious tasks. This approach may work relatively well in controlled laboratory settings, but online studies offer significantly less direct control over participants and their environment. Consequently, some of the established research practices may have to be dropped or redesigned.

One research practice that is particularly difficult to enforce in an online environment is limiting access to the study to only the chosen population (e.g. women). There are many ways in which researchers can try to limit access to a study, for example by rejecting participants that fail a pre-screening question about their gender. However, we believe that such an approach is counterproductive. Some especially motivated members of other groups (e.g. males) could lie or use a fake Facebook account, creating an unnecessary source of invalid data that may be difficult to filter out. In practice, it is usually easier to remove participants that do not meet the target criteria before the analysis, rather than trying to filter them out at the data-collection stage.

Another practice that is likely to trigger dishonest responses or behaviors is when the researcher prevents participants from skipping questions or tasks. Even a small number of such invalid records can have a detrimental effect on the results of the study, so it is advisable to allow the participants to skip tasks. Afterwards, incomplete protocols may be removed or the missing data imputed (Schafer & Graham, 2002). Additionally, we suggest that personal data or consent to record profile information should be requested at the end, rather than at the beginning of a study. Having acquainted themselves with the design of the study and the feedback it offers, participants might be more willing to provide honest responses and to trust researchers with access to their profiles.

Keeping it short and minimizing number of questions. The length of the study is another crucial issue to be considered when conducting online research. Decreased direct control over the participants increases the risk of them providing invalid responses or abandoning the study altogether. Thus, we believe that participants should be engaged gradually, for example by distributing incentives across the study. In the case of myPersonality, respondents could choose to receive their reward (basic feedback on their personality) after answering as few as 20 out of 100 personality questions. This ensured that the immediate barrier to participation was low and encouraged many of the participants to answer more questions or proceed to other questionnaires. Even if such an approach has several disadvantages (e.g. 20 questions did not provide enough information for many of the analyses, and early feedback could affect respondents' further answers), we believe that it is better than encouraging participants to submit random responses or abandon the study prematurely.

Ethical considerations

In the previous sections, we have argued that Facebook and other online environments have made it easier than ever before to observe individuals. Although this has the potential to greatly boost social science research, it also introduces new ethical challenges. The repercussions of misconduct in online Human Subjects Research (HSR) could be far greater than ever imagined (Barchard & Williams, 2008; Hall & Flynn, 2001; Molokken-Ostvold, 2005). This was recently illustrated by the public reaction to manipulating users' Facebook news feeds to study emotional contagion (Kramer, Guillory, & Hancock, 2014).[19] Nevertheless, the protocols and

[19] See, for example, http://www.theatlantic.com/technology/archive/2014/06/even-the-editor-of-facebooks-mood-study-thought-it-was-creepy/373649/

guidelines related to designing online studies, storing data, and analyzing results are scarce and often contradictory (Solberg, 2010; Wilson et al., 2012). At the time of writing, the American Psychological Association's website[20] lists only three documents (Frankel & Siang, 1999; Hewson, 2003; Kraut et al., 2003) containing guidelines related to research on the Internet, the most recent being from 2003 (nearly two years before Facebook was founded). In the UK, the most recent edition of the Economic and Social Research Council Framework for Research Ethics (2012) simply states that Internet research presents new ethical dilemmas, without suggesting any guidelines beyond requiring that any research conducted on the Internet should be subject to a full review by the appropriate Institutional Review Board (IRB).

The lack of clear guidelines is exacerbated by ever-accelerating technological progress; both researchers and IRB members may over or underestimate the threats to participants, thereby hindering benign projects or approving malignant ones. Both factors discourage social scientists from carrying out online research or submitting studies for review (Hall & Flynn, 2001; Singer & Vinson, 2001). As a result, an increasing fraction of HSR is carried out by computer scientists, who are often unconcerned about or unfamiliar with the ethical and social implications of HSR (Buchanan, Aycock, Dexter, Dittrich, & Hvizdak, 2011; Buchanan & Ess, 2009). Additionally, it discourages non-academic researchers (e.g. in commercial companies) from seeking any kind of ethical clearance or sharing their results with other scientists and the general public.

This trend is disconcerting, and not only because Facebook constitutes a powerful research tool and an important area of interest for social sciences. Thus, we hope to encourage the relevant bodies like IRBs; federal agencies like U.S. Department of Health and Human Services; and the American Psychological Association Ethics Committee to increase their focus on new research tools and environments, including Facebook.

[20] Source: http://www.apa.org/research/responsible/human/index.aspx

Below we discuss several aspects of consent and data collection that are, in our experience, especially relevant in the Facebook environment.

Participants' control over data

Facebook and other online environments have made it easier than ever before to share private information, in both greater detail and larger volume. Content that has traditionally been considered intensely private, and worthy of staunch legal protection, is now—at the click of a button—openly broadcast to one's network of friends. Moreover, participants in Facebook-based studies seem to feel comfortable sharing those extensive records with researchers (Dwyer, Hiltz, & Passerini, 2007). For example, over 30% of myPersonality participants decided to volunteer the contents of their Facebook profiles, together with their personality, intelligence, and other psychometric scores (Stillwell & Kosinski, 2015).

Facebook users' increased openness to share their personal information with researchers may be driven by several factors. First, Facebook users seem to have far greater control over their data than is usually assumed. Facebook profile data are shared with and scrutinized by hundreds of their friends and acquaintances. Information that a given user considers to be overly intimate, or that which casts him/her in a bad light, is unlikely to remain on the profile and thus would be unavailable to the researcher. Second, Facebook's privacy settings enable users to revoke or limit access to data after it has been granted. Third, Facebook provides an additional layer of privacy by requiring users to confirm the consent given to applications (see Figure 1). Finally, traditional studies frequently involve personal interaction with the researcher in an unnatural laboratory setting. This can often make respondents anxious and self-aware, decreasing the chances of them being honest and behaving naturally (Smyth & Pearson, 2011). By contrast, studies based on Facebook and other online environments usually enable people to participate in a time and place of their own choice, favoring generalizability and external validity of the results (Reips, 2000).

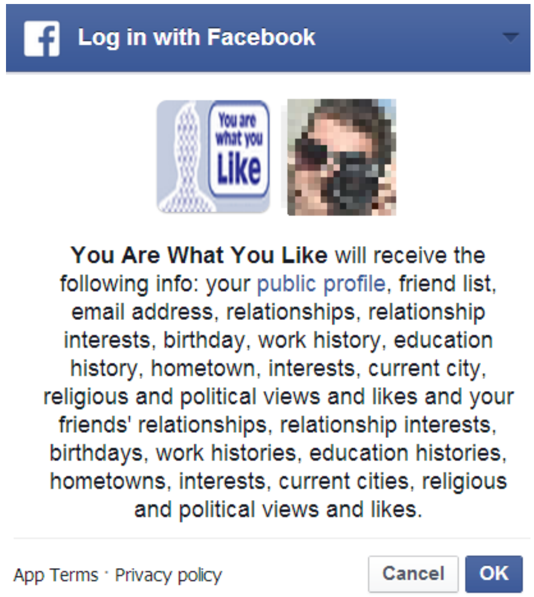

Figure 1. An example of a Facebook consent dialog box, shown after offering access to one's Facebook data to a third party.

Facebook offers participants a relatively high degree of control over their data, but it is the researcher's responsibility to weigh the costs and benefits of collecting and using personal user information (and to defer to an IRB in case of doubt). The mere availability of data and participants' willingness to share them does not grant researchers an automatic right to record and use them freely. The boundaries of processing the data should always be agreed on with the participant; this concern is especially acute in the context of findings showing that even a seemingly innocuous digital footprint, such as Facebook Likes, can be used to infer highly intimate details, such as intelligence, personality, political views, and sexual orientation (Kosinski, Bachrach, et al., 2013; Kosinski, Stillwell, et al., 2013; Youyou et al., 2015). Fortunately, the Facebook environment enhances researchers' ability to treat their participants as collaborators in research, rather than subjects. This could be achieved by ascertaining that the participants:

1. Obtain clear details about what information is going to be extracted from their data, and who will be able to access it;
2. Obtain feedback on their results or performance (where possible), regardless of whether they have consented to their data being stored;
3. Can review and retract any piece of information that researchers have collected based on their consent;
4. Can request to receive a notification about the publication of any results based on their data. (We believe that, whenever possible, researchers should publish in open-access journals, and also endeavor to write and share non-technical reports of results, such as blogs or articles in popular press.)

Granting participants the aforementioned rights would not only minimize the chances of researchers attempting to use data in an unacceptable way, but might also secure participants' well-being, engagement, and trust.

Consent forms

Consent forms used in social sciences are aimed at protecting participants and researchers. However, in the online environment, already replete with license and consent forms, users have been conditioned to click “Agree” without first reading the fine print (Böhme & Köpsell, 2010). This puts both the participants and the researchers at risk. Thus, researchers should take advantage of the design flexibility offered by the online environment to craft consent forms in a participant-friendly way.

The most important information could be clearly laid out in few sentences (or bullet points), using a large accessible font, with hyperlinks or pop-up dialog boxes enabling interested participants to obtain more details on each of the points. Such an approach could result in participants paying more attention to what they consent to, and making a truly informed decision as to whether they want to participate or not. We hope that researchers and relevant institutions—including IRBs, federal agencies, and the American Psychological Association Ethics Committee—will investigate the feasibility of more participant-friendly formats of informed consents for online usage.

Importantly, as mentioned in the *Participants' control over data* section, we believe that people should be allowed to participate in the study (e.g. take the questionnaire or play a game) even if they declined to share their data or results with the researchers. Note that in some cases, consent could be requested (or re-requested) at the end of the participation, or along with the feedback; participants might be more open to sharing their data and results after experiencing the study.

Boundaries of individual consent

Perhaps the most challenging aspect of recording Facebook data relates to the vague boundaries between information that belongs solely to participants, and that which can be accessed with their consent, but belongs to or relates to other individuals. Examples of such data include pictures or videos featuring non-participants; messages or comments received from other people; or information about participant's friends. These can all, at present, be recorded on Facebook with users' consent.

We propose that it should be generally acceptable to use data generated by or containing references to non-participants, provided that the analyses are aimed *exclusively* at those directly participating in the study. For example, non-participants' demographic profiles and network connections (that can be observed with participants' consent) could be used to establish the parameters of participants' egocentric social networks, or gender ratio among their friends. Similarly, non-participants' activity, such as photo tags, comments, or Likes, could be analyzed in an aggregated form, to extract knowledge about participants, such as their popularity or social activity.

Collecting public Facebook profile data

Some basic profile information is publicly available and even indexed by external search engines. Social scientists disagree on whether collecting publicly available data falls within the regulatory definition of HSR (Economic and Social Research Council, 2012), and whether it therefore requires participants' consent and IRB approval (Schultze & Mason, 2012; Solberg, 2010; Wilson et al., 2012). Some scholars point out that the border between public and private is not determined by accessibility, but by social norms and practices (Frankel & Siang, 1999; Waskul, 1996). This point is illustrated by the example of a small town in which everyone knows intimate details about everyone else, but people pretend not to know those facts that would be considered personal (Schultze & Mason, 2012). Others argue that mining public data is an equivalent to conducting archival research, a method frequently employed in disciplines such as history, art criticism, and literature, which rarely involve rules for the protection of human subjects (Bruckman, 2002; Herring, 1996).

We agree with the latter argument and propose that using public Facebook profile data (e.g. data that are publicly available to any Facebook user and not only to one's Facebook friends) should not require participants' consent if the following conditions are met:

1. It is reasonable to assume that the data were knowingly made public by the individuals;
2. Data are anonymized after collection and no attempts are made to de-anonymize them;
3. There is no interaction or communication with the individuals in the sample; and
4. No information that can be attributed to a single individual, including demographic profiles and samples of text or other content, is to be published or used to illustrate the results of the study.

If any one of the above conditions is not met, the study should be closely scrutinized by an IRB.

Discuss ethical considerations and privacy consequences when publishing results

Finally, following (Schultze & Mason, 2012), we encourage researchers conducting HSR on Facebook (and other online environments) to include a discussion of ethical considerations related to the design of a given study in their publications. Furthermore, we believe that the study should discuss the ethical implications of the findings. Such an approach would ensure that the authors have properly considered ethical aspects of their own work, and would support the evolution of standards and norms in the quickly changing technological environment.

Conclusions

A growing proportion of human activities, such as social interactions, entertainment, shopping, and gathering information, are now mediated by digital devices and services such as Facebook. Such digitally-mediated behaviors can easily be recorded and analyzed, fueling the emergence of *computational social science* (Lazer et al., 2009; Markowetz, Blaszkiewicz, Montag, Switala, & Schlaepfer, 2014). They also facilitate the transition from small-scale experiments and observational studies to large-scale projects based on thousands or millions of individuals (Backstrom, Boldi, Rosa, Ugander, & Vigna, 2012; Eichstaedt et al., 2015; Kosinski, Bachrach, et al., 2013; Kosinski, Stillwell, et al., 2013; Kramer et al., 2014; Ugander, Karrer, Backstrom, & Marlow, 2011; Youyou et al., 2015). Observing or experimenting with large samples enables scientists to minimize the problem of sampling error typical to social science, and to detect patterns that might not be apparent in smaller samples. It also offers unprecedented insights into the dynamics and organization of individual behavior and social systems, with the potential to radically improve our understanding of human psychology (Lazer et al., 2009). The same processes and technologies, which drive the emergence of computational social science, are also rapidly transforming the human environment. This drives the need to re-examine the relevance of established social science theories, and use modern technology to develop new ones.

However, researching psychological phenomena in the digital environment requires skills that are relatively uncommon among social scientists, such as recording, storing, processing, and analyzing large databases. For example, the combined size of all databases published on the myPersonality project website exceeds 50GB, and some of the tables contain hundreds of millions of rows. Processing or analyzing tables of this size and complexity requires at least a basic knowledge of scientific programming (e.g. R or Python) and database management (such as MySQL). Additionally, it is necessary to truly understand the online environment under investigation. A degree of familiarity can be achieved by studying the literature, but this is no substitute for personal experience with the platform or, in other words, becoming an active Facebook user.

Another example of a potential drawback is that Facebook allows users to list their family members, but instead some users choose to list their close friends as brothers or sisters rather than, or in addition to, their actual family members. Researchers who are not active

Facebook users are prone to invalidating their results by overlooking or misunderstanding peculiarities such as this.

As social scientists are relatively slow in embracing the necessary practical skills (Lazer et al., 2009), data-driven human subjects research is increasingly ceded to computer scientists and engineers, who often lack the appropriate theoretical background and ethical standards (Buchanan et al., 2011; Hall & Flynn, 2001). We strongly encourage our fellow social scientists to not only train themselves in modern computation methods, but to also immerse themselves in new human environments, including Facebook. These digital platforms are more than an object of study. Being as rich and diverse as the complex human environments from which they emerge, digital platforms offer new opportunities for social science research and new challenges for researchers in their own right. With proper training, traditional social science studies can be conducted online at a lower cost and larger scale than ever before.

References

Aral, S., & Walker, D. (2012). Identifying influential and susceptible members of social networks. *Science*, *337*(6092), 337–341.

Arnaboldi, V., Guazzini, A., & Passarella, A. (2013). Egocentric online social networks: Analysis of key features and prediction of tie strength in Facebook. *Computer Communications, 36*(10-11), 1130–1144.

Back, M. D., Stopfer, J. M., Vazire, S., Gaddis, S., Schmukle, S. C., Egloff, B., & Gosling, S. D. (2010). Facebook Profiles Reflect Actual Personality, Not Self-Idealization. *Psychological Science*, *21*(3), 372.

Backstrom, L., Boldi, P., Rosa, M., Ugander, J., & Vigna, S. (2012). Four degrees of separation. In *Proceedings Of the Web Science Conference* (pp. 33–42).

Baltar, F., & Brunet, I. (2012). Social research 2.0: virtual snowball sampling method using Facebook. *Internet Research*, *22*(1), 57–74.

Barchard, K. A., & Williams, J. (2008). Practical advice for conducting ethical online experiments and questionnaires for United States psychologists. *Behavior Research Methods*, *40*(4), 1111–1128.

Batterham, P. J. (2014). Recruitment of mental health survey participants using Internet advertising: Content, characteristics and cost effectiveness. *International Journal of Methods in Psychiatric Research*.

Birnbaum, M. H. (2004). Human research and data collection via the Internet. *Annual Review Of Psychology*, *55*, 803–832.

Blumberg, S. J., & Luke, J. V. (2009). Reevaluating the need for concern regarding noncoverage bias in landline surveys. *American Journal of Public Health*, *99*(10), 1806–1810.

Böhme, R., & Köpsell, S. (2010). Trained to accept? In *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10* (p. 2403). New York, New York, USA: ACM Press.

Bruckman, A. (2002). Studying The Amateur Artist: A Perspective On Disguising Data Collected In Human Subjects Research On The Internet. *Ethics And Information Technology*, *4*(3), 217–231.

Buchanan, E. A., Aycock, J., Dexter, S., Dittrich, D., & Hvizdak, E. (2011). Computer science security research and human subjects: Emerging considerations for research ethics boards. *Journal of Empirical Research on Human Research Ethics*, *6*(2), 71–83.

Buchanan, E. A., & Ess, C. (2009). Internet research ethics and the institutional review board: Current practices and issues. *ACM SIGCAS Computers and Society*, *39*(3), 43–49.

Carlini, B. H., Safioti, L., Rue, T. C., & Miles, L. (2014). Using Internet to Recruit Immigrants with Language and Culture Barriers for Tobacco and Alcohol Use Screening: A Study Among Brazilians. *Journal of Immigrant and Minority Health*, 1–8.

Caspi, A., & Gorsky, P. (2006). Online deception: Prevalence, motivation, and emotion. *CyberPsychology & Behavior*, *9*(1), 54–59.

Chandler, J., Mueller, P., & Paolacci, G. (2014). Nonnaïveté among Amazon Mechanical Turk workers: consequences and solutions for behavioral researchers. *Behavior Research Methods*, *46*(1), 112–30. doi:10.3758/s13428-013-0365-7

Close, S., Smaldone, A., Fennoy, I., Reame, N., & Grey, M. (2013). Using information technology and social networking for recruitment of research participants: Experience from an exploratory study of pediatric klinefelter syndrome. *Journal of Medical Internet Research*, *15*(3), e48.

CNET. (2014). Remember Friendster? Thought so. Retrieved from http://www.cnet.com/news/remember-friendster-thought-so/

Connelly, B. S., & Ones, D. S. (2010). An other perspective on personality: meta-analytic integration of observers' accuracy and predictive validity. *Psychological Bulletin*, *136*(6), 1092–122. doi:10.1037/a0021212

David L., M. (2008). *The SAGE Encyclopedia of Qualitative Research Methods* (pp. 816 – 817). SAGE Publications, Inc.

Doody, M., & Sigurdson, A. (2003). Randomized trial of financial incentives and delivery methods for improving response to a mailed questionnaire. *American Journal of …*. Retrieved from http://aje.oxfordjournals.org/content/157/7/643.short

Duggan, M., Ellison, N. B., Lampe, C., Lenhart, A., & Madden, M. (2015). Social Media Update 2014. *Pew Research Center*.

Dwyer, C., Hiltz, S. R., & Passerini, K. (2007). Trust and Privacy Concern Within Social Networking Sites: A Comparison of Facebook and MySpace. In *Proceedings Of The Americas Conference On Information Systems* (p. 339).

Economic and Social Research Council. (2012). ESRC Framework for Research Ethics 2010 (revised September 2012).

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., … Seligman, M. E. P. (2015). Psychological Language on Twitter Predicts County-Level Heart Disease Mortality. *Psychological Science*, 0956797614557867–. doi:10.1177/0956797614557867

Facebook Inc. (2015). Key Facts About Facebook. Retrieved from http://newsroom.fb.com/key-facts

Frankel, M. S., & Siang, S. (1999). *Ethical And Legal Aspects Of Human Subjects Research On The Internet*. Washington DC: American Association For The Advancement Of Science.

Goldberg, L. R. (1999). A Broad-Bandwidth, Public Domain, Personality Inventory Measuring The Lower-Level Facets Of Several Five-Factor Models. *Personality Psychology In Europe*, *7*, 7–28.

Goodman, L. A. (1961). Snowball Sampling. *The Annals of Mathematical Statistics*, *32*(1), 148–170. doi:10.1214/aoms/1177705148

Gosling, S. D., Vazire, S., Srivastava, S., & John, O. P. (2004). Should we trust web-based studies? A comparative analysis of six preconceptions about internet questionnaires. *American Psychologist*, *59*(2), 93.

Hall, T., & Flynn, V. (2001). Ethical issues in software engineering research: A survey of current practice. *Empirical Software Engineering*, *6*(4), 305–317.

Han, J., Niu, J., Chin, A., Wang, W., Tong, C., & Wang, X. (2012). How online social network affects offline events: A case study on douban. In *Proceedings - IEEE 9th International Conference on Ubiquitous Intelligence and Computing and IEEE 9th International Conference on Autonomic and Trusted Computing, UIC-ATC 2012* (pp. 752–757).

Heine, S. J., Lehman, D. R., Peng, K., Greenholtz, J., & Of, D. (2002). What’s wrong with cross-cultural comparisons of subjective Likert scales? The reference-group problem. *Journal of Personality and Social Psychology (JPSP)*, 903–918.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The WEIRDest People In The World. *Behavioral And Brain Sciences*, *33*(2-3), 61–83.

Herring, S. (1996). Linguistic And Critical Analysis Of Computer-Mediated Communication: Some Ethical And Scholarly Considerations. *The Information Society*, *12*(2), 153–168.

Hewson, C. (2003). Conducting research on the internet. *Psychologist*, *16*(6), 290–293.

Huang, C.-M., & Park, D. (2013). Cultural influences on Facebook photographs. *International Journal of Psychology*, *48*(3), 334–343.

John, O. P., & Benet-Martinez, V. (2000). H. T. Reis And C. M. Judd (Eds.) Handbook Of Research Methods In Social And Personality Psychology. In H. T. Reis & C. M. Judd (Eds.), (pp. 339–369). New York, NY: Cambridge University Press.

John, O. P., & Srivastava, S. (1999). The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of Personality: Theory and Research*, *2*, 102–138.

Johnson, J. A. (2004). The impact of item characteristics on item and scale validity. *Multivariate Behavioral Research*, *39*(2), 273–302.

Johnson, J. A. (2005). Ascertaining the validity of individual protocols from web-based personality inventories. *Journal Of Research In Personality*, *39*(1), 103–129.

Johnson, K. J., Mueller, N. L., Williams, K., & Gutmann, D. H. (2014). Evaluation of participant recruitment methods to a rare disease online registry. *American Journal of Medical Genetics, Part A*.

Kosinski, M., Bachrach, Y., Kohli, P., Stillwell, D. J., & Graepel, T. (2013). Manifestations Of User Personality In Website Choice And Behaviour On Online Social Networks. *Machine Learning*, *95*(3), 1–24.

Kosinski, M., Stillwell, D. J., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(15), 5802–5805.

Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 201320040.

Kraut, R., Olson, J., Banaji, M., Bruckman, A., Cohen, J., & Couper, M. (2003). Psychological Research Online: Opportunities and Challenges. *Psychological Research*, *412*, 268–7694.

Kurant, M., Markopoulou, A., & Thiran, P. (2010). *On the bias of BFS (Breadth First Search)*. Retrieved from http://arxiv.org/abs/1004.1729

Kurant, M., Markopoulou, A., & Thiran, P. (2011). Towards Unbiased BFS Sampling. *IEEE Journal on Selected Areas in Communications*, *29*(9), 1799–1809. doi:10.1109/JSAC.2011.111005

Kurtz, J. E., & Parrish, C. L. (2001). Semantic response consistency and protocol validity in structured personality assessment: The case of the NEO-PI-R. *Journal Of Personality Assessment*, *76*(2), 315–332.

Lampe, C., Ellison, N., & Steinfield, C. (2006). A Face(book) in the Crowd: Social Searching vs. Social Browsing. *Proceedings of the 2006 20th Anniversary Conference on Computer-Supported Cooperative Work CSCW '06*, 167–170.

Lampe, C., Ellison, N., & Steinfield, C. (2007). A Familiar Face (book): Profile Elements as Signals in an Online Social Network. *Technology*, 435–444.

Lazer, D., Pentland, A., Adamic, L. A., Aral, S., Barabási, A.-L., Brewer, D., … Van Alstyne, M. (2009). Life in the network: the coming age of computational social science. *Science*, *323*(5915), 721–723.

Mangan, M. A., & Reips, U.-D. (2007). Sleep, sex, and the Web: Surveying the difficult-to-reach clinical population suffering from sexsomnia. *Behavior Research Methods*, *39*(2), 233–236. doi:10.3758/BF03193152

Markowetz, A., Blaszkiewicz, K., Montag, C., Switala, C., & Schlaepfer, T. E. (2014). Psycho-Informatics: Big Data shaping modern psychometrics. *Medical Hypotheses*, *82*(4), 405–411.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 415–444.

Molokken-Ostvold, K. (2005). Ethical Concerns when Increasing Realism in Controlled Experiments with Industrial Participants. *Proceedings Of The Annual Hawaii International Conference on System Sciences*, 264a–264a.

Paulhus, D. L., & Vazire, S. (2007). The self-report method. *Handbook of Research Methods in Personality Psychology*, 224–239.

Pempek, T. A., Yermolayeva, Y. A., & Calvert, S. L. (2009). College students' social networking experiences on Facebook. *Journal of Applied Developmental Psychology*, *30*(3), 227–238.

Pennebaker, J. W., Francis, M. E., & Booth, R. J. (2001). *Linguistic Inquiry and Word Count: LIWC 2001* (p. 71). Mahway: Lawrence Erlbaum Associates.

Peters, A., Winschiers-Theophilus, H., & Mennecke, B. (2013). Bridging the digital divide through Facebook friendships: A cross-cultural study. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW* (pp. 249–254).

Pew Research Center. (2014). *Older Adults and Technology Use*.

Reips, U.-D. (2000). The Web experiment method: Advantages, disadvantages, and solutions. *Psychological Experiments On The Internet*, 89–117.

Richiardi, L., Pivetta, E., & Merletti, F. (2012). Recruiting study participants through Facebook. *Epidemiology*, *23*(1), 175.

Ross, C., Orr, E. S., Sisic, M., Arseneault, J. M., Simmering, M. G., & Orr, R. R. (2009). Personality And Motivations Associated With Facebook Use. *Computers In Human Behavior*, *25*(2), 578–586.

Salganik, M. J., & Heckathorn, D. D. (2004). Sampling and Estimation in Hidden Populations Using Respondent-Driven Sampling. *Sociological Methodology*, *34*(1), 193–240. doi:10.1111/j.0081-1750.2004.00152.x

Schafer, J. L., & Graham, J. W. (2002). Missing data: Our view of the state of the art. *Psychological Methods*, *7*(2), 147–177.

Schultze, U., & Mason, R. O. (2012). Studying Cyborgs: Re-Examining Internet Studies As Human Subjects Research. *Journal of Internet Technology*, *27*(4), 301–312.

Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., … Ungar, L. H. (2013). Personality, Gender, and Age in the Language of Social Media: The Open-Vocabulary Approach. *PLoS ONE*, *8*(9).

Singer, J., & Vinson, N. (2001). Why and how research ethics matters to you. Yes, You! *Empirical Software Engineering*, *6*(4), 287–290.

Smyth, J. D., & Pearson, J. E. (2011). Social and Behavioral Research and the Internet. In M. Das, P. Ester, & L. Kaczmirek (Eds.), (pp. 11–44). New York: Routledge.

Solberg, L. (2010). Data mining on Facebook: A free space for researchers or an IRB nightmare? *University of Illinois Journal of Law, Technology & Policy*, *2*.

Stillwell, D. J., & Kosinski, M. (2015). myPersonality Project Website. Retrieved from http://mypersonality.org

Ugander, J., Karrer, B., Backstrom, L., & Marlow, C. (2011). The anatomy of the Facebook social graph. *arXiv Preprint arXiv:1111. 4503*.

Waskul, D. (1996). Considering The Electronic Participant: Some Polemical Observations On The Ethics Of On-Line Research. *The Information Society*, *12*(2), 129–140.

Wilson, R. E., Gosling, S. D., & Graham, L. T. (2012). A Review Of Facebook Research In The Social Sciences. *Perspectives On Psychological Science*, *7*(3), 203–220.

Yang, Z., Wilson, C., Wang, X., Gao, T., Zhao, B. Y., & Dai, Y. (2011). Uncovering social network sybils in the wild. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference - IMC ’11* (p. 259). New York, New York, USA: ACM Press. doi:10.1145/2068816.2068841

Youyou, W., Kosinski, M., & Stillwell, D. J. (2015). Computer-based personality judgements are more accurate than those made by humans. *Proceedings Of The National Academy Of Sciences (PNAS)*, *112*(4), 1036–1040.