

Twitter bot surveys: A discrete choice experiment to increase response rates

Juan Pablo Alperin

Simon Fraser University
515 W. Hastings St.
Vancouver, BC V6B 5K3

juan@alperin.ca

Erik Warren Hanson

Simon Fraser University
515 W. Hastings St.
Vancouver, BC V6B 5K3

ewhanson@sfu.ca

Kenneth Shores

University of Pennsylvania
3700 Walnut Street
Philadelphia, PA 19104-6216

kshores@gse.upenn.edu

Stefanie Haustein

Université de Montréal
CP 6128, succursale
Centre-ville
Montreal, QC H3C 3J7

stefanie.haustein@umontreal.ca

ABSTRACT

This paper presents a new methodology—the Twitter bot survey—that bridges the gap between social media research and web surveys. The methodology uses the Twitter APIs to identify a target population and then uses the API to deliver a question in the form of a regular Tweet. We hypothesized that this method would yield high response rates because users are posed a question within the social media platform and are not asked, as is the case with most web surveys, to follow a link away to a third party. To evaluate the response rate and identify the most effective mechanism for increasing it, we conducted a discrete choice experiment that evaluated three factors: question type, the use of an egoistic appeal, and the presence of contextual information. We found that, similar to traditional web surveys, multiple choice questions, egoistic appeals, and contextual information all contributed to higher response rates. Question variants that combined all three yielded a 40.0% response rate, thereby outperforming most other web surveys and demonstrating the promise of this new methodology. The approach can be extended to any other social media platforms where users typically interact with one another. The approach also offers the opportunity to bring together the advantages of social media research using APIs with the richness of information that can be collected from surveys.

CCS Concepts

• Human-centered computing→Social media • Human-centered computing→Empirical studies in collaborative and social computing • Human-centered computing→Social content sharing

Keywords

social media; web surveys; survey methods; audience research; discrete choice experiment;

1. MOTIVATION

For many years, researchers have been taking advantage of the data offered by social media platforms to study and analyze a wide range of social phenomena. As social media provide easy access via application programming interfaces (APIs) to rich data about

people's interactions in various online environments, they have become a cost and time efficient alternative to surveys; however, the data that can be collected from these systems may not address specific questions that researchers may have—details that would have traditionally been explored through direct questions to participants.

Although the move towards online channels to distribute surveys offers some of the same advantages as data access over APIs (namely, global reach, speed and timeliness, and low cost), the current practice of directing individuals from email or social media sites to participate in an online survey hosted elsewhere on site such as Qualtrics or SurveyMonkey typically have low response rates. To this end, survey method researchers have systematically explored what factors increase response rates to avoid surveys being seen as junk mail and to avoid sample selection biases [11, 12, 19, 20]. Though these practices have been well documented in web surveys, the literature has not considered social media platforms themselves as a tool for conducting surveys, beyond using them to broadly disseminate traditional online surveys [5, 8].

More recently, some efforts have been made to combine access to social media users via APIs with surveys to have more direct and targeted interactions with individuals. In their review, Courtois and Mechant [10] highlight the possibility of using APIs to conduct social science research in combination with web surveys. They document a case study that uses the YouTube API to distribute links to surveys in the comments of relevant videos [9, 10]. Although this approach begins to connect data collected from social media APIs with the users' survey responses, the response rates for the surveys remain relatively low, between 11% and 16% [9, 10]. These response rates are similar to has been found in reviews of survey methods for comparable approaches [6, 15].

While the low response rates are discouraging, there have been recent efforts to explore ways of maximizing interactions with social media users. For example, Savage and colleagues point to different strategies for recruiting volunteers through Twitter bots like the one used in this survey [17]. In their study, Savage and colleagues identified how different motivating phrases and ways of having the bot identify itself affected users' interactions with their bot.

The approach offered here seeks to combine the techniques used to conduct web surveys successfully, with the approach used by Savage and colleagues to identify successful strategies on social media. In doing so, it builds on a pilot study by Alperin [1, 2], by presenting a methodological approach that conducts a survey directly via a social media platform (namely, Twitter) and enriches survey responses with available user information. To the best of our

knowledge it represents the first to attempt to systematically use Twitter to administer and enrich a survey.

Our approach involves the creation of a Twitter bot, an automated program that interacts with targeted users in the same manner as they generally interact with each other. The bot uses regular tweets to ask specific questions directly targeted at selected users by using their Twitter handle (Figure 1). This methodology offers the advantages of being able to link survey responses with social media data and avoiding losing participants during a platform change, as the survey is directly carried out on Twitter. Due to the latter our approach is expected to have a higher response rates than traditional web surveys.

This paper presents a full description of the methodology of a bot-run survey on Twitter. It aims to determine which types and styles of questions lead to maximum response rates.

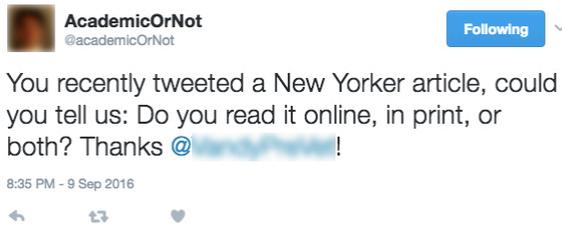


Figure 1. Example tweet survey question

2. METHODS

To understand the potential of surveying select users directly on Twitter, we carried out a case study targeting users who tweeted articles of the *New Yorker* magazine, asking them about their preferred format to read articles. For this study, we were not concerned about the sample selection, or the users’ eventual responses, as our focus was solely on measuring response rates.

We began by exploring strategies and approaches identified by web survey research [5, 19]. From here, we looked for factors that consistently improve response rates of surveys administered online. Although these strategies have not been studied in the context of social media surveys, they provide an adequate starting point for developing new online survey methods.

Based on the literature we identified three factors that are applicable to our Twitter surveys: question type (i.e., open ended, yes/no, or multiple choice) [5], personalizing the request (also known as an egoistic appeal) [16], and providing contextual information [14]. These best practices from web surveys led us to identify twelve question variants to test the most effective question formulations to solicit a high response rate (Table 1).

We formulated the variants in the following ways. We used one of three variants: an open-ended question (A: OE; “What format do you prefer to read it in?”), a yes/no question (A: YN; “Do you read it primarily online?”), and a multiple-choice question (A: MC; “Do you primarily read it online, in print, or both?”). The second factor, egoistic appeal, was tested using a generic appeal (B: No; “Please help us understand *New Yorker* readers”) and a statement that appeals to the user’s ego (B: Yes; “You recently tweeted a *New Yorker* article, could you tell us:”). To test the third factor, providing contextual information, we sent tweets as an @reply (C: Yes), posted as a response to and appearing below the user’s original tweet linking to the *New Yorker* article, and as an @mention (C: No), where the question tweet is not attached to the

user’s original tweet. In both cases, for questions using an @reply or @mention, an alert is automatically sent to the user notifying them of the question.

Table 1. Tweet factors attributes and levels

Variant	Factor ^a			Question text
	A ^b	B	C	
1	OE	No	Yes	@screenname Please help us understand New Yorker readers. What format do you prefer to read it in? Thanks!
2	OE	No	No	Please help us understand New Yorker readers. What format do you prefer to read it in? Thanks @screenname!
3	YN	No	Yes	@screenname Please help us understand New Yorker readers. Do you read it primarily online? Thanks!
4	YN	No	No	Please help us understand New Yorker readers. Do you read it primarily online? Thanks @screenname!
5	M C	No	Yes	@screenname Please help us understand New Yorker readers. Do you read it online, in print, or both? Thanks!
6	M C	No	No	Please help us understand New Yorker readers. Do you read it online, in print, or both? Thanks @screenname!
7	OE	Yes	Yes	@screenname You recently tweeted a New Yorker article, could you tell us: What format do you prefer to read it in? Thanks!
8	OE	Yes	No	You recently tweeted a New Yorker article, could you tell us: What format do you prefer to read it in? Thanks @screenname!
9	YN	Yes	Yes	@screenname You recently tweeted a New Yorker article, could you tell us: Do you read it primarily online? Thanks!
10	YN	Yes	No	You recently tweeted a New Yorker article, could you tell us: Do you read it primarily online? Thanks @screenname!
11	M C	Yes	Yes	@screenname You recently tweeted a New Yorker article, could you tell us: Do you read it online, in print, or both? Thanks!
12	M C	Yes	No	You recently tweeted a New Yorker article, could you tell us: Do you read it online, in print, or both? Thanks @screenname!

^a A: Question Type; B: Egoistic appeal; C: Context information

^b OE: Open Ended; YN: Yes/No; MC: Multiple Choice

To identify effect sizes for each of the factor types (question type, personalization and contextual information) and variants within

type¹ we conducted a discrete choice experiment (DCE). DCEs are widely used in health economics and other disciplines to discern user preferences with multiple factors (e.g., to identify what treatment patients prefer given different potential side effects) [4]. In our case, user preference was indicated by their choice to respond or not. By randomly assigning factor variants to respondents, it is possible to recover independent estimates for each factor type and variant [13]. When the effect sizes of each factor are unknown, DCE experiments commonly use a rule of thumb of running at least 100 trials [3]; we ran our experiment until we had asked each variant 100 times, leading to total of 1,331 questions. The outcome variable of interest is whether the user responded to the prompt.

Our sample is constructed by identifying targeted users who had recently shared (tweeted or retweeted) a *New Yorker* article on Twitter, we entered the search terms “*newyorker.com filter: links*” into the “*search/tweets*” Twitter API endpoint between July 5 and July 13, 2016. This query returned both original tweets and retweets linking to the *New Yorker* website. All of the tweet data and the associated user information were subsequently stored in a local database, which allows us to subsequently link this user information to their responses.

To automatically survey the targeted users via the Twitter bot, we set up a Twitter account, @academicOrNot, which identified itself as a bot used for research purposes in the Twitter bio, including a link to a description of the research project and a picture of the first author. Survey responses were collected for a period of 31 days by intermittently searching the “*statuses/mentions_timeline*” of the @academicOrNot account, which captures all of the tweets that mention our bot (by default on Twitter, all replies begin with the original poster’s screen name).

3. RESULTS

Of the 1,331 Twitter users who were asked a question, 309 users provided a valid response, yielding an overall response rate of 23.2%, which exceeds the usually expected response rates of web surveys [6, 15]. However, there was significant variability in the response rates of each variant, ranging from 9.9% (variant 4) to 40% (variant 11), highlighting a direct effect of question type and style on user participation. At respective response rates of 40.0% and 35.5%, multiple choice questions with an egoistic appeal with (variant 11) and without context information (variant 12) attracted the highest user participation. (Table 2).

We also see variability between the response rates associated with each of the three factors (Table 3). Among question types, open-ended questions (19.9%) and yes/no questions (20.2%) showed comparable and much lower response rates than multiple-choice questions (29.9%). Using an egoistic appeal (29.9%) almost doubled user participation compared to questions without a personalized request (16.7%). The difference between providing and not providing contextual information was lower, as replies (25.6%) received, on average, a slightly higher response rate than mentions (21.0%).

Table 2. Question variant response rates

Variant	Question count	Response count	Response rate
1	119	20	16.8%
2	112	14	12.5%
3	116	20	17.2%
4	111	11	9.9%
5	106	24	22.6%
6	112	24	21.4%
7	110	31	28.2%
8	109	24	22.0%
9	114	33	29.0%
10	105	26	24.8%
11	110	44	40.0%
12	107	38	35.5%
Total	1,331	309	23.2%

Table 3. Response rate by factor variants
Question type

Open-ended	19.9%
Yes/no	20.2%
Multiple choice	29.9%
Egoistic appeal	
No	16.8%
Yes	29.9%
Reply/mention	
Reply	25.6%
Mention	21.0%

To determine whether these effect sizes are statistically significant, we estimated the following logistic regression equation:

$$\Pr(Y_i = 1 | \mathbb{X}) = F(\alpha + \beta_1 A_2 + \beta_2 A_3 + \beta_3 B + \beta_4 C + \epsilon_i)$$

Where $F(z) = \frac{e^z}{1+e^z}$, i.e. is the logistic function. Y_i is equal to 1 when respondents replied; A_2 is a variable equal to 1 if respondent received a question type that was “Yes/no”; A_3 is a variable equal to 1 if respondent received a question type that was “Multiple choice.” B is a variable equal to 1 if respondent received an “egoistic appeal.” C is a variable equal to 1 if respondent received a “mention” (as opposed to a reply). (In all cases, where variable is not equal to 1, it is set to 0). Results are shown in Table 4.

Results indicate that multiple choice items increase response rate by 0.56 log-points (with an odds ratio of 1.75) and are significantly different from the omitted category of open-ended response types. Egoistic appeals increase response rates by 0.76 log-points (with an

¹ Specifically, Factor types A includes variants OE, YN, and MC; factor type B includes variants Yes and No; factor type C includes variants Yes and No.

odds ratio of 2.14) and are significantly different from non-egoistic appeals. Sending the message as simply a mention reduces response rates by nearly 0.27 log- points (with an odds ratio of 0.76) compared to threading the response in a person's Twitter feed (as a reply).

Table 4: Point Estimates for Item Variants

Variable	Beta
Question Type "Yes/No"	0.020 (0.169)
Question Type "Multiple choice"	0.557*** (0.160)
Egoistic Appeal: Yes	0.760*** (0.135)
Contextual Information: No (@mention)	-0.271* (0.133)
Constant	-1.685*** (0.156)
Number of respondents	1331

Note: Standard error in parenthesis

* for $p < .05$, ** for $p < .01$, and *** for $p < .001$

4. DISCUSSION AND CONCLUSION

The results of this pilot study show the viability of surveying users directly on a social media platform like Twitter. At the same time, they highlight that question type and style can have a measurable effect on response rates. While the overall response rate for all question variants was 23.2%, two variants led to response rates of above 35%. This is a significant improvement over previous attempts at recruiting users on social media, which only had 10-15% of users follow a link to a survey [9]. It is also an improvement over other types of web surveys, such as "pop-up" surveys, which typically have response rates of around 22% [7], and of web surveys generally which have response rates around 11% [15].

While the response rates overall were comparable or higher than other web surveys, the most successful variants are consistent with the strategies shown to be successful in other forms. Overall, multiple choice questions had a higher response [18], as did egoistic appeals [16], and providing context [14]. Combining these variants yielded the highest response rate (variant 11, 40.0%), which demonstrates that these effects are additive. The time of day a question is sent might be an additional factor helping to maximize response rates, as might skipping accounts thought to be automated (i.e., bot) accounts themselves (estimated between 9-15% of active Twitter accounts) [21].

Although asking questions on Twitter is limited to short questions and responses, the information gathered could be useful for identifying a target audience for further surveys or detailed quantitative analysis. Sending longer multiple choice questions in the form of attached pictures might be a possibility to bypass the 140 character limitation.

The results of our case study suggest that the presented methodology is promising for conducting social media surveys, as it leads to higher overall response rates than regular web surveys. Optimizing the way in which questions are asked helps to further increase response rates.

Moreover, the approach of embedding the survey into the social media environment facilitates the enrichment of user responses information about their social media behavior, obtained from the particular platform. This approach thus allows us to gain all of the advantages of social media research and to complement it with the user details that can only be gleaned from a survey. By linking all of the data from the user accounts with user responses, this method provides a better and more complete understanding of the users behind the social media accounts. In the case of Twitter, we can map their tweeting behavior and tweet contents with their responses to questions about their motivations, affiliations, personality, opinions, etc.

5. REFERENCES

- [1] Alperin, J.P. 2015. Moving beyond counts: A method for surveying Twitter users. *Altmetrics15 Workshop* (Amsterdam, Netherlands, Oct. 2015).
- [2] Alperin, Juan Pablo 2015. *The public impact of Latin America's approach to open access*. Stanford University.
- [3] de Bekker-Grob, E.W. et al. 2015. Sample Size Requirements for Discrete-Choice Experiments in Healthcare: a Practical Guide. *The Patient*. 8, 5 (2015), 373–384.
- [4] Clark, M.D. et al. 2014. Discrete choice experiments in health economics: a review of the literature. *PharmacoEconomics*. 32, 9 (Sep. 2014), 883–902.
- [5] Couper, M.P. 2008. *Designing effective Web surveys*. Cambridge University Press.
- [6] Couper, M.P. 2000. Review: Web Surveys: A Review of Issues and Approaches. *Public Opinion Quarterly*. 64, 4 (Feb. 2000), 464–494.
- [7] Couper, M.P. 2001. The promises and perils of web surveys. *The challenge of the Internet*. 35, (2001), 56.
- [8] Couper, M.P. and Miller, P.V. 2008. Web Survey Methods Introduction. *Public Opinion Quarterly*. 72, 5 (Dec. 2008), 831–835.
- [9] Courtois, C. et al. 2010. Teenage Uploaders on YouTube: Networked Public Expectancies, Online Feedback Preference, and Received On-Platform Feedback. *Cyberpsychology, Behavior, and Social Networking*. 14, 5 (Nov. 2010), 315–322.
- [10] Courtois, C. and Mechant, P. 2014. An Evaluation of the Potential of Web 2.0 APIs for Social Research. *Audience research methodologies: between innovation and consolidation*. G. Patriarche et al., eds. Routledge. 212–224.
- [11] Das, M. et al. 2011. *Social and behavioral research and the internet: advances in applied methods and research strategies*. Routledge.
- [12] Evans, J.R. and Mathur, A. 2005. The value of online surveys. *Internet Research*. 15, 2 (Apr. 2005), 195–219.
- [13] Hainmueller, J. et al. 2014. Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments. *Political Analysis*. 22, 1 (Jan. 2014), 1–30.
- [14] Liu, M. et al. 2015. Impact of Web Survey Invitation Design on Survey Participation, Respondents, and Survey Responses. *Social Science Computer Review*. (Sep. 2015), 894439315605606.
- [15] Manfreda, K.L. et al. 2008. Web surveys versus other survey modes. *International Journal of Market Research*. 50, 1 (Jan. 2008), 79–104.
- [16] Pedersen, M.J. and Nielsen, C.V. 2016. Improving Survey Response Rates in Online Panels Effects of Low-Cost Incentives and Cost-Free Text Appeal Interventions. *Social Science Computer Review*. 34, 2 (Apr. 2016), 229–243.

- [17] Savage, S. et al. 2016. Botivist: Calling Volunteers to Action using Online Bots. *Microsoft Research*. (Feb. 2016).
- [18] Schwanda-Sosik, V. et al. 2014. Online Microsurveys for User Experience Research. *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2014), 889–892.
- [19] Tourangeau, R. et al. 2013. *The science of web surveys*. Oxford University Press.
- [20] Tourangeau, R. et al. 2013. “Up Means Good” The Effect of Screen Position on Evaluative Ratings in Web Surveys. *Public Opinion Quarterly*. 77, S1 (Jan. 2013), 69–88.
- [21] Varol, O. et al. 2017. Online Human-Bot Interactions: Detection, Estimation, and Characterization. *arXiv:1703.03107 [cs]*. (Mar. 2017).