



The happiness paradox: your friends are happier than you

Johan Bollen^{1*}, Bruno Gonçalves², Ingrid van de Leemput³ and Guangchen Ruan¹

*Correspondence:

jbollen@indiana.edu

¹School of Informatics and
Computing, Indiana University,
Bloomington, IN, USA

Full list of author information is
available at the end of the article

Abstract

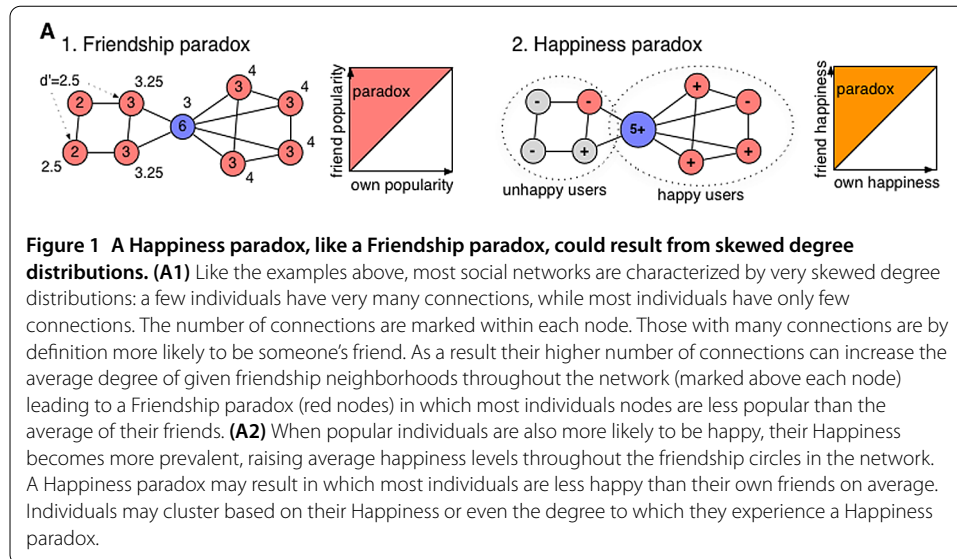
Most individuals in social networks experience a so-called Friendship Paradox: they are less popular than their friends on average. This effect may explain recent findings that widespread social network media use leads to reduced happiness. However the relation between popularity and happiness is poorly understood. A Friendship paradox does not necessarily imply a Happiness paradox where most individuals are less happy than their friends. Here we report the first direct observation of a significant Happiness Paradox in a large-scale online social network of 39,110 Twitter users. Our results reveal that popular individuals are indeed happier and that a majority of individuals experience a significant Happiness paradox. The magnitude of the latter effect is shaped by complex interactions between individual popularity, happiness, and the fact that users strongly cluster by similar level of happiness. Our results indicate that the topology of online social networks, combined with how happiness is distributed in some populations, may be associated with significant psycho-social effects.

Keywords: subjective well-being; social media; social network; friendship paradox; sentiment analysis; natural language processing; data science

1 Introduction

We are a profoundly social species [1]. The ability to establish face-to-face, physical, relationships in a rich social environment is paramount to our happiness and individual well-being [2–4]. However, technology is now playing an increasing role in forming our networks of social relationships. Nearly 1/7th of the world's population and over 2/3rd of the US population [5] use some form of social media which enables individuals to maintain virtual social networks that extend beyond geographical, economic, cultural, and linguistic boundaries.

Evidence has been accumulating that online social networking is associated with elevated levels of loneliness, anxiety, displeasure, and dissatisfaction [6–10]. The reason for this apparent contradiction is unknown, but it may be found in universal social network connectivity patterns. Surprisingly, measured in number of connections, most people will have fewer friends than their own friends do on average. This phenomenon, commonly referred to as the Friendship Paradox [11, 12], has been attributed to an inherent structural bias in social network that favors popular individuals: they are by definition more likely to belong to someones social circle, thereby elevating local levels of popularity. One may



speculate that if an individual compares their own popularity to that of their friends, this effect may, in some cases, lead to increased levels of dissatisfaction (see Figure 1).

The effects of this Friendship Paradox may extend beyond popularity. If popular individuals tend to be happier, their elevated happiness will become more prevalent as well. This may in turn lead to a Happiness Paradox, where most individuals are less happy than their friends on average (see Figure 1). In fact, the latter may contribute more directly to the negative psycho-social effects of social networking since it affects how individuals assess their own Subjective Well-being, i.e. general happiness or life satisfaction [13, 14], relative to that of others [15, 16].

At this point, however, it has not been established whether (1) popular individuals are indeed happier and (2) a Happiness Paradox does in fact occur in social networks. Given the magnitude of social media adoption, these are questions of global importance that may affect the well-being of billions of individuals.

2 Methods

Here we present the first large-scale longitudinal study of happiness and popularity levels for a network of 39,110 Twitter users that are connected by 'friendship' relations.

To generate a friendship network among Twitter users we start with an initial set of 4,844,430 randomly chosen users (years 2008-2009), for which we downloaded the full list of users that they 'follow' or that they are 'followed' by [17]. Reciprocal 'Follow' and 'Following' ties are taken as an indication of a friendship and mutual interaction relation between the two individuals [18]. We selected the Largest Connected Component of the resulting network. The resulting network consists of 102,009 individuals that share 2,361,547 edges.

We automatically assess each individual's Subjective Well-Being (SWB), on a scale of $[-1, +1]$ according to a procedure outlined in [17]. For each user we collected the 3,200 most recently submitted tweets, serving as a comprehensive longitudinal record pertaining to the individual user. We subject this entire timeline, not its individual tweets, to a sentiment analysis [17] based on the OpinionFinder (OF) subjectivity lexicon [19]. The procedure consists of counting the number of words that are members of either the set

of highly positive or highly negative words in the OF subjectivity lexicon, and calculating their fractional difference (number of positive words-negative words over all OF words in the timeline). Aggregating this information for all Tweets in an individual Twitter record, we determine the individual's overall Subjective Well-Being.

The OF toolkit, although different from the above described application of its subjectivity lexicon, was ranked 11th out of 24 tools in a large-scale survey of 2- and 3-way sentiment classification tasks against a variety of data sets [20]. Its accuracy in scoring individual tweets ranged from 57.60 to 80.77, mostly approximating, and sometimes exceeding, that of top rated tools such as VADER or AFINN, with coverage levels from 41.23 to 54.98, depending on the specific data set against which OF was tested. We stress that our procedure uses the OF Subjectivity Lexicon against an entire 6 month timeline of up to 3,200 tweets by the same individual Twitter user, not individual tweets. Hence the reported coverage values are expected to be a significant underestimation relative to our application. In addition, we exclude users with SWB values of exactly zero.

Finally, we restrict our analysis to individuals with more than 15 friends in order to exclude individuals with excessively low social activity and those that have non-zero SWB, i.e. that have shared some subjective information. This reduces our final cohort to 39,110 subjects that are connected by reciprocal friendship relations, have at least 15 friends, and have non-zero SWB values.

We quantify an individuals 'Happiness' as their SWB value and quantify their 'Popularity' as their number of network friends counted in our bidirectional social network of reciprocal connections. Simply put, we deem an individual 'Happy' when they have high SWB values and 'Popular' when they have many friends. The Happiness and Popularity values of all subjects are then used to determine:

1. the fraction of individuals that has lower *popularity* than their friends on average (P);
2. the fraction of individuals that has lower *happiness* than their friends on average (H);
3. the correlation between individual happiness and popularity, R (Happiness, Popularity).

We assess the magnitude of the Friendship Paradox in our network by calculating the fraction of the number of users $u_i \in U$ whose Popularity, denoted $D(u_i)$ (degree of u_i) is lower than the average Popularity \bar{D} of their nearest neighbors (or 'friends') $\mathcal{N}_{u_i} \subset U$ vs. the total number of individuals in the network $\|U\|$. This yields the magnitude of the Friendship Paradox as:

$$P = \frac{\|\{u_i \in U : D(u_i) < \bar{D}(\mathcal{N}_{u_i})\}\|}{\|U\|}. \quad (1)$$

The magnitude of the Happiness Paradox can be obtained in a similar way to how we measure the Friendship Paradox. We simply calculate the fraction of users $u_i \in U$ whose Happiness, denoted $H(u_i)$ (SWB value of u_i), is lower than the average Happiness \bar{H} of their nearest neighbors, \mathcal{N}_{u_i} , vs. the total number of individuals in the network $\|U\|$:

$$H = \frac{\|\{u_i \in U : H(u_i) < \bar{H}(\mathcal{H}_{u_i})\}\|}{\|U\|}. \quad (2)$$

A Friendship or Happiness Paradox for our sample is indicated by P and H values larger than 50%, i.e. a majority of individuals have lower Popularity or Happiness than their friends, on average.

To assess the correlation between Happiness and Popularity we calculate Pearson's R between the SWB values and $\log(\text{degree})$ of all subjects in our cohort. The use of $\log(\text{degree})$ is meant to compensate for the very skewed distribution of degree values in our network.

We assess the robustness of our results by performing a bootstrapping procedure in which we randomly sample 10% of subjects and their network connections with replacement 5,000 times to assess the distribution of our paradox indicators for different samples of our network. Furthermore, we validate the statistical significance of our results by comparing them to a null-model where we reshuffle the SWB values across all individuals in our network. In this way, we are able to maintain the same identical distribution of SWB values and network structure, while completely eliminating any possible correlation that might be present. The null-model was bootstrapped 20,000 times and, as expected, it eliminated the Happiness paradox.

Finally, we determine the sensitivity of our results to reductions of the sample due to the 'minimum friends threshold' by recalculating all Happiness Paradox magnitudes for values of the threshold ranging from 1 to 200.

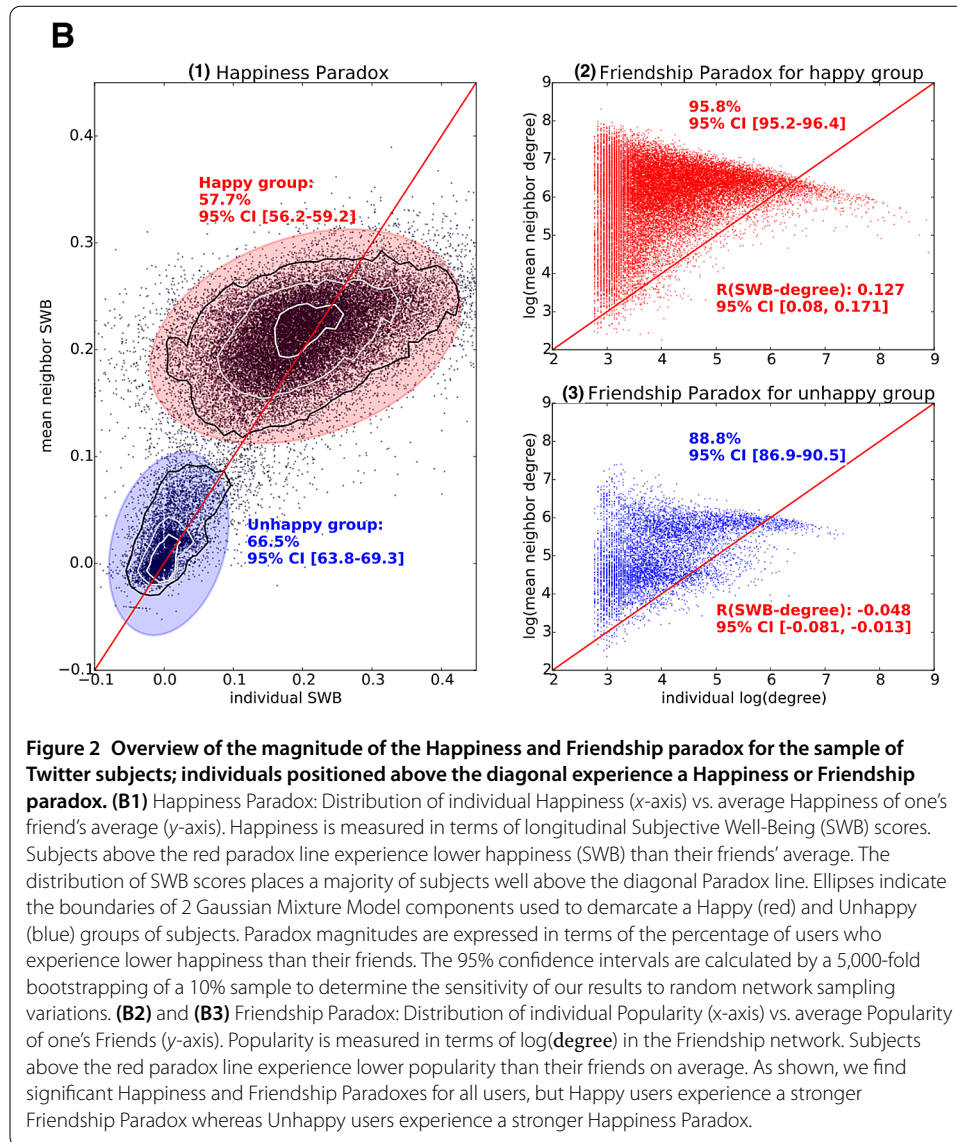
3 Results

As shown in Table 1, we find that $P = 94.3$, 95% CI [93.7, 94.9], indicating a very significant Friendship Paradox across all subjects, meaning that the great majority of users are less popular than their friends are on average. We also find a modest but robust value of $H = 58.5\%$, 95% CI [57.2, 59.8], indicating the presence of a Happiness Paradox. Hence a majority of subjects is indeed less happy than their friends on average. Our null-model indicates the absence of a Happiness paradox when the effects of network structure on Happiness levels are removed by random re-assignment. The lower magnitude of the Happiness Paradox could result from the rather low, yet robust, correlation between Happiness and Popularity (Pearson's $R = 0.100$, 95% CI [0.076, 0.140]).

As shown in Figure 2 the joint distribution of individual Happiness levels and mean neighbor happiness in our sample is distinctly bi-modal. There exists some evidence that Subjective Well-Being itself can be distributed in a bimodal fashion across several cultures and nations [21] matching previous observations by [17]. In this case, bi-modality also occurs at the level of our friendship network which separates subjects into 2 distinct groups: *Happy subjects with Happy friends* (the 'Happy' group) and *Unhappy subjects with Unhappy friends* (the 'Unhappy' group). This result follows earlier reports of happiness being homophilic or assortative in social networks [17, 22, 23]. Note that this phenomenon

Table 1 Magnitude of Friendship Paradox, Happiness Paradox (compared to null-model produced by randomly re-assigning SBW values across all subjects), and Happiness-Popularity correlation coefficient (Pearson's R) for all subjects ($N = 39,110$)

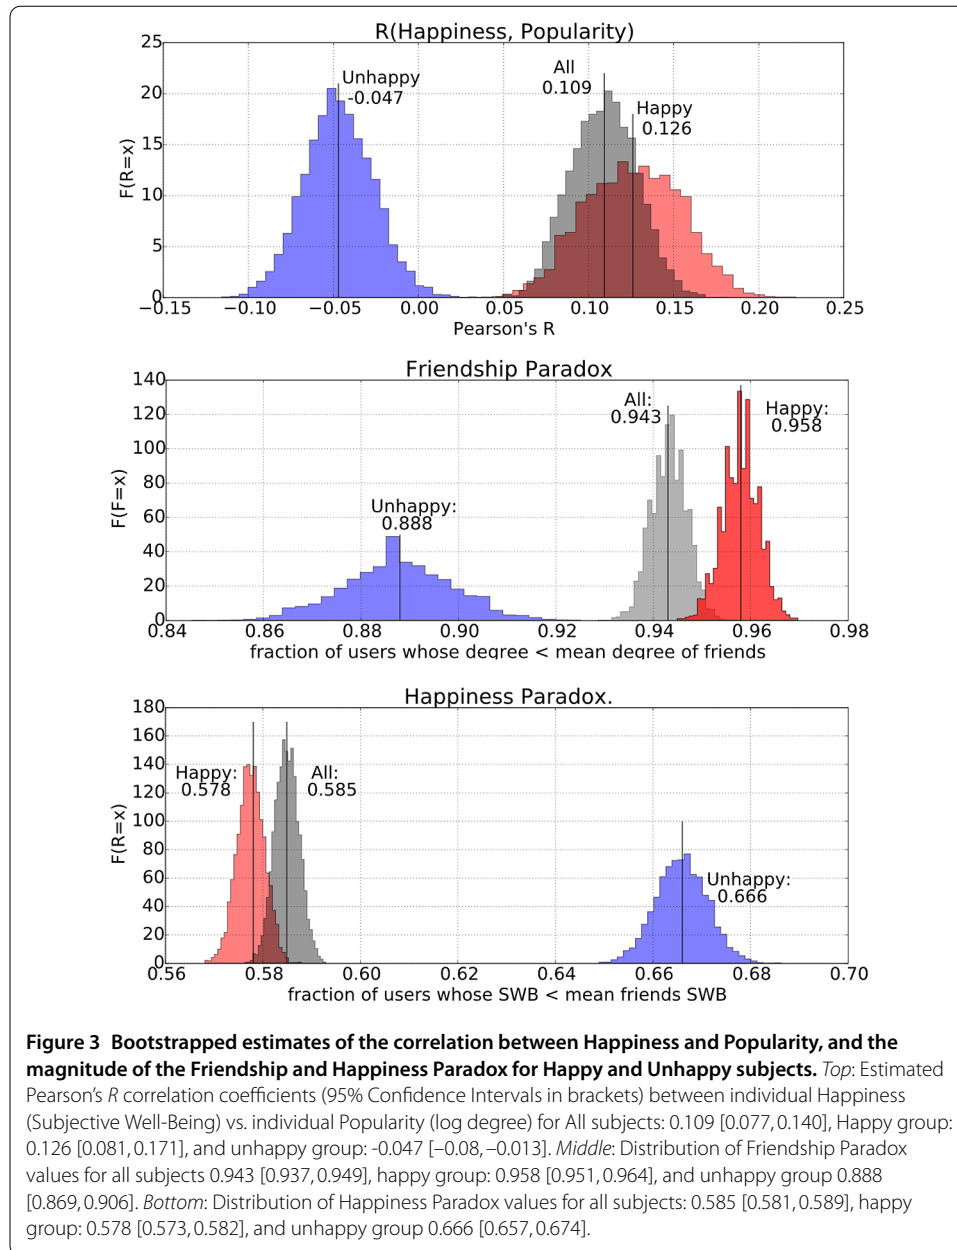
	Value (%) and 95% CI
P	94.3%, [93.7, 94.9], $N = 5,000$
H	58.5%, [57.2, 59.8], $N = 5,000$
Null-model	50.1%, [49.4, 50.8], $N = 20,000$
R (Happiness, Popularity)	0.109, [0.076, 0.140], $N = 5,000$



is not the result of the distribution of our SWB values, but is determined by network connections; i.e. unhappy people tend to have unhappy friends and happy people tend to have happy friends.

Since a Happiness Paradox specifically compares individual happiness to the average happiness of one's friends, this homophilic bi-modality must be factored into our analysis. By performing a separate analysis for Happy and Unhappy groups of users, we attempt to equalize the effects of neighbor happiness across the two groups.

As shown in Figure 2 we use a Gaussian Mixture Model (GMM) to demarcate our Happy and Unhappy groups. We determine the location and distribution of two separate Gaussian components in the distribution of individual happiness vs. mean friend happiness and demarcate both groups by simply determining whether the SWB value of a subject and the mean SWB values of their neighbors fall within 2 standard deviations from the center of either one of the components (illustrated by the 2 ellipses in Figure 2). We thereby split our sample in 2 groups of subjects: a group of Happy individuals with Happy friends



($N = 29,033$) and a group of Unhappy individuals with Unhappy friends (8,018). This procedure assumes a Gaussian density distribution which roughly matches the quantiles of the empirical density as shown by the contour lines of Figure 2, so that both Gaussian components capture 95% of our sample. These groups are well-separated in the underlying social network. Of all edges adjacent to individuals in the Happy group only 0.682% connect to individuals in the Unhappy group. Vice versa, of all edges adjacent to individuals in the Unhappy group only 3.103% connect to individuals in the Happy group.

We re-run our analysis for the Happy and Unhappy groups separately. The results are summarized in Figures 2 and 3.

These results reveal that the Happy group experiences a strong Friendship Paradox but a weak, yet very robust Happiness Paradox. The Unhappy group experiences a weaker

Friendship Paradox, but a significantly stronger Happiness Paradox than the Happy group, in spite of subjects being surrounded by less Happy friends.

To determine whether the strong Happiness Paradox for the unhappy group, in spite of its lower correlation between Popularity and Happiness, may be related to interpersonal effects, we examine the relation between individual happiness and the average happiness of ones neighbors. As visually indicated by the distribution of individuals in Figure 2, the strength of the relationship between a subjects' Happiness and the average Happiness of their friends may differ between the Happy and Unhappy group. The results of a linear regression predicting mean neighbor SWB from a user's own SWB values match the visual tilt of the separate GMM components. For the Happy group we find $F(1, 29,031) = 6,580$, $p < 0.001$ with an adjusted $R^2 = 0.185$. The resulting regression equation is Mean Neighbor SWB = $0.1768 + 0.1815$ (own SWB). For the Unhappy Group we find $F(1, 8,018) = 3,274$, $p < 0.001$ with an adjusted $R^2 = 0.290$. Here, the resulting regression equation is Mean Neighbor SWB = $0.0135 + 0.5716$ (own SWB). An ANCOVA analysis to predict Average Neighbor SWB from the interaction between a user's own SWB and their membership of the Happy or Unhappy group results in $F(3, 39,106) = 70,040$, $p < 0.001$ with an adjusted $R^2 = 0.843$. On the basis of this last result we reject the null-hypothesis that the regression slopes between own SWB and mean neighbor SWB are equal between the Happy and Unhappy group. This outcome suggests that unhappy users may be more strongly affected by the lower happiness of their friends, possibly explaining why this group exhibits a stronger Happiness Paradox in the absence of a strong correlation between Popularity and Happiness. However, we caution that further investigation is warranted on this matter.

Finally, to determine the sensitivity of our analysis to the choice of our minimum neighbors threshold (the minimum number of neighbors a user must have to be included in our sample), we recalculate Happiness paradox values for all threshold values between 1 and 200. We use the same GMM component locations for every calculation to provide an equal basis for comparison. As shown in Table 2 and Figure 4, we find that requiring a minimum threshold of 15 neighbors for each individual in our data constitutes a signif-

Table 2 Percentage of sample remaining after applying a minimum number of neighbors threshold

\geq Neighbors	Sample percentage
1	100
5	74.9
10	57.1
15	46.6
20	39.7
30	30.5
40	24.7
50	20.8
60	17.7
70	15.4
80	13.6
90	12.1
100	10.8
125	8.5
150	6.9
175	5.7
200	4.8

For the analysis reported in this paper we used a threshold of 15, i.e. each user included in the analysis should have at least 15 neighbors.

Figure 4 Minimum neighbor threshold vs. percentage of sample retained. In our calculation of the Happiness Paradox we apply a minimum neighbor threshold for each user, i.e. a minimum number of neighbors the individual must have to be included in our sample, to (1) reduce the chances of including defunct, automated, or particularly socially inactive users and (2) ensure that a mean neighbor degree or SWB is calculated on the basis of no less than 15 data points. This threshold leads to a reduction of our sample. This graph visualizes the magnitude of the reduction at chosen threshold value of 15 and all other values between 1 and 200. Values provided in Table 2.


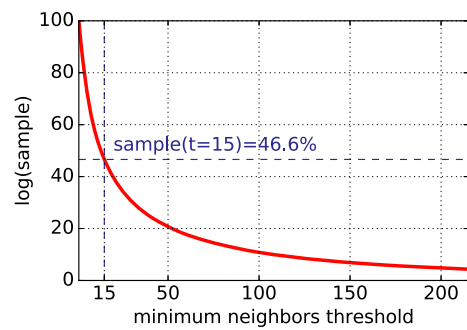


Figure 5 Happiness paradox values for Happy and Unhappy group at different minimum neighbor threshold values, with 95% confidence intervals. Overlaid are the remaining sample size for the same threshold values. The gray area indicates paradox value below 50%, i.e. no happiness paradox. We find a significant happiness paradox, under all but the most extreme values of the minimum neighbor threshold.

icant reduction of our sample, but we still retain more than 45% of the entire sample. In exchange for this reduction, we attempt to increase the chances of removing noise caused by socially inactive, and possibly automated or defunct accounts. The value of this threshold is notably much lower than recent estimations of the online Dunbar number [18].

The results of our sensitivity analysis, shown in Figure 5, indicate that the magnitude of the happiness paradox is largely *independent* from our choice of threshold, with the exception of extremely large values, i.e. where less than 10% of the original sample remains and only individuals with more than 150 friends are included, thus approaching the Dunbar limit [18]. We find a significant happiness paradox for the Unhappy group for all values of the minimum neighbor threshold indicating that this result is largely independent from our choice of threshold.

4 Conclusion

We find that a majority of both Happy and Unhappy social media users will experience a significant Popularity and Happiness Paradox, i.e. they will be less popular and less happy than their friends on average. One may speculate that previous observations reporting decreased happiness among social media users may be associated with a widespread inflated perception of the happiness of one's friends. Although happy and unhappy groups of subjects are both affected by a significant happiness paradox, unhappy subjects seem most

strongly affected. This is counter-intuitive for two reasons. First, the correlation between happiness and popularity is lowest for individuals in the unhappy group. A happiness paradox can result from a friendship paradox when popularity and happiness are correlated, since more popular and thus more prevalent individuals will increase the average happiness of ones circle of friends. As a result, the unhappy group, with the lowest correlation between popularity and happiness, should experience the lowest happiness paradox. Second, the strong assortativity of happiness in our social network reduces the prevalence of happy subjects in the social network circle of unhappy subjects. Therefore, it should be easier for individuals in this group to surpass the average happiness of their friends. Our results show that neither is the case. We speculate that a possible explanation may lie in the stronger relation between the happiness of individuals in this group and the overall happiness of their friends. This effect may point to an alternate origin for the occurrence of a Happiness paradox; instead of resulting from the greater prevalence of popular and happy individuals, in some cases, a happiness paradox may result from the complex social interactions between individuals and their friends, e.g. through mood contagion [24–26] and potentially verbal commiseration and mirroring.

Our results suggest that social media use may be associated with increased levels of social dissatisfaction and unhappiness if individuals are subject to unfavorable comparison between their own happiness and popularity to those of their friends. Happy social media users may think their friends are much more popular and slightly happier than they are while unhappy social media users will likely have unhappy friends that will still seem much happier and more popular than they are on average.

However, our study has limitations. First, the assessment of Subjective Well-Being from social media using text analysis algorithms may not be perfectly reliable. However, given the large number of individuals in our dataset, the absence of consistent directional bias, and the magnitudes of the observed effects, we expect this will not affect the validity of our observations. Future improvements in sentiment and mood analysis, and ground truth obtained from user surveys, may increase the reliability of our SWB estimates. Second, given the large role that social media plays in the social lives of billions of individuals, we speculate that these environments may induce longitudinal changes in the public social behavior and may over time alter the very nature of social relations themselves [27]. Further analysis will be required to determine the extent and significance of these changes, and how they affect the propensity of online users to experience the effects of a Friendship and Happiness Paradox over time. Lastly, we stress that our study is, by definition, observational in nature. We do not and can not make any claims with respect to the causal nature of the relationships between Subjective Well-Being, popularity, mood contagion, and social network use. This will require an extensive program of experimental verification over extended periods of time which we are planning.

Acknowledgements

We thank Marten Scheffer and Filippo Raddichi for their feedback on earlier versions of this manuscript.

Funding

Bruno Gonçalves thanks the Moore and Sloan Foundations for support as part of the Moore-Sloan Data Science Environment at NYU. Johan Bollen thanks the Defense Advance Research Projects Agency (NGS2 2016 #D17AC00005) and the National Science Foundation (SMA-SBE: 1636636) for their support.

Competing interests

The authors declare that they have no competing financial interests.

Authors' contributions

JB designed and executed the experimental methodology, performed the data analysis, and wrote the manuscript. IL performed data analysis, advised on methodological and analytical issues, and co-authored the manuscript. BG collected the data, advised on methodological and analytical issues, and co-authored the manuscript. GR collected and processed the raw data. All authors reads and approved the final manuscript.

Author details

¹School of Informatics and Computing, Indiana University, Bloomington, IN, USA. ²Center for Data Science, New York University, New York, USA. ³Environmental Sciences, Wageningen University, Wageningen, The Netherlands.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 14 July 2016 Accepted: 5 May 2017 Published online: 18 May 2017

References

1. Apicella CL, Marlowe FW, Fowler JH, Christakis NA (2012) Social networks and cooperation in hunter-gatherers. *Nature* 481:497-501
2. Dunbar R (2010) The social role of touch in humans and primates: behavioural function and neurobiological mechanisms. *Neurosci Biobehav Rev* 34(2):260-268
3. Steptoe A, Shankar A, Demakakos P, Wardle J (2013) Social isolation, loneliness, and all-cause mortality in older men and women. *Proc Natl Acad Sci USA* 110:5797-5801
4. Pinker S (2015) *The village effect: why face-to-face contact matters*. Atlantic Books, London
5. Perrin A (2015) *Social media usage: 2005-2015*. Report, Pew Research Center
6. Burke M, Marlow C, Lento T (2010) Social network activity and social well-being. In: *Proceedings of the SIGCHI conference on human factors in computing systems, CHI '10*. ACM, New York, pp 1909-1912
7. Kross E, Verduyn P, Demiralp E, Park J, Lee DS, Lin N, Shaback H, Jonides J, Ybarra O (2013) Facebook use predicts declines in subjective well-being in young adults. *PLoS ONE* 8(8):e69841
8. Tromholt M, Lundby M, Andsbjerg K, Wiking M (2015) *The Facebook experiment: does social media affect the quality of our lives?* Technical report, The Happiness Institute
9. Mitchell L, Frank MR, Decker Harris K, Dodds PS, Danforth CM (2013) The geography of happiness: connecting Twitter sentiment and expression, demographics, and objective characteristics of place. *PLoS ONE* 8(5):e64417
10. Lin L, Sidani JE, Shensa A, Radovic A, Miller E, Colditz JB, Hoffman BL, Giles LM, Primack BA (2016) Association between social media use and depression among U.S. young adults. *Depress Anxiety* 33(4):323-331
11. Feld SL (1991) Why your friends have more friends than you do. *Am J Sociol* 96:1464-1477
12. Hodas NO, Kooti F, Lerman K (2013) Friendship paradox redux: your friends are more interesting than you. In: *Proceedings of the seventh international AAAI conference on weblogs and social media*. AAAI Press, Menlo Park
13. Oswald AJ, Wu S (2010) Objective confirmation of subjective measures of human well-being: evidence from the U.S.A. *Science* 327(5965):576-579
14. Krueger AB, Stone AA (2014) Progress in measuring subjective well-being. *Science* 346(6205):42-43
15. Posel DR, Casale DM (2011) Relative standing and subjective well-being in South Africa: the role of perceptions, expectations and income mobility. *Soc Indic Res* 104:195-223
16. de la Garza A, Mastrobuoni G, Sannabe A, Yamada K (2010) The relative utility hypothesis with and without self-reported reference wages. *Carlo Alberto Notebooks* 159, Collegio Carlo Alberto
17. Bollen J, Gonçalves B, Ruan G, Mao H (2011) Happiness is assortative in online social networks. *Artif Life* 17:237-251
18. Gonçalves B, Perra N, Vespignani A (2011) Modeling users' activity on Twitter: validation of Dunbar's number. *PLoS ONE* 6:e22656
19. Wilson T, Hoffmann P, Somasundaran S, Kessler J, Wiebe J, Choi Y, Cardie C, Riloff E, Patwardhan S (2005) Opinionfinder: a system for subjectivity analysis. In: *Conference on empirical methods in natural language processing*, pp 34-35
20. Ribeiro FN, Araújo M, Gonçalves P, Gonçalves MA, Benevenuto F (2016) Sentibench - a benchmark comparison of state-of-the-practice sentiment analysis methods. *EPJ Data Sci* 5(1):23
21. Stiller B, Bocek T, Hecht F, Machado G, Racz P, Waldburger M (2010) *Measuring national well-being: proposed well-being indicators*. Technical report, The Commission on Measuring Well-being, Japan
22. McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: homophily in social networks. *Annu Rev Sociol* 27(1):415-444
23. Fowler JH, Christakis NA (2008) Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham heart study. *Br Med J* 337:a2338
24. Kramer ADI, Guillory JE, Hancock JT (2014) Experimental evidence of massive-scale emotional contagion through social networks. *Proc Natl Acad Sci USA* 111:8788-8790
25. Ferrara E, Yang Z (2015) Measuring emotional contagion in social media. *PLoS ONE* 10:e0142390
26. Ferrara E, Yang Z (2015) Quantifying the effect of sentiment on information diffusion in social media. *PeerJ Comput Sci* 1:e26
27. Dunbar R (2012) Social cognition on the Internet: testing constraints on social network size. *Philos Trans R Soc Lond B, Biol Sci* 367(1599):2192-2201