

The happiness paradox: your friends are happier than you.

Johan Bollen^{1,3*}, Bruno Gonçalves², Ingrid van de Leemput^{1,3}, and Guangchen Ruan¹

¹Indiana University, Bloomington IN

²Center for Data Science, New York University, New York, NY.

³Wageningen University, 6700 AA, Wageningen, The Netherlands

*To whom correspondence should be addressed; E-mail: jbollen@indiana.edu

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 cluster assortatively by level of happiness. Our results indicate that the topology of online social networks and the distribution of happiness in some populations can cause widespread psycho-social effects

that affect the well-being of billions of individuals.

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) now use some form of social media which enables individuals to maintain virtual social networks that extend well 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 (6–9). 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 (10–12). This phenomenon, commonly referred to as the Friendship Paradox, has been attributed to an inherent structural bias in social network that favors popular individuals: they are by definition more likely to belong to someone’s social circle, thereby elevating local levels of popularity. If individuals equate popularity with prestige and compare their own popularity to that of their friends this may lead to increased levels of dissatisfaction (see Fig.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 (11, 13, 14), where most individuals are less happy than their friends on average (see Fig.1). In fact, the latter will 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 (15, 16), relative to that of

others (17, 18). 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.

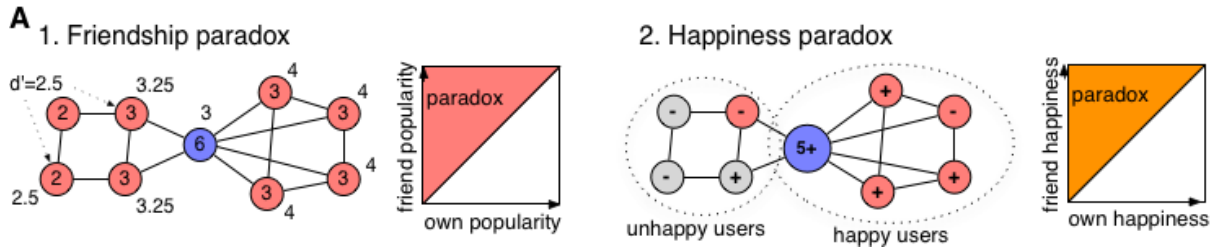


Figure 1: **A1:** Most social networks are characterized by very skewed degree distributions: a few individuals have very many connections, while most individuals have few connections. The number of connections are marked within each node. Those with many connections are by definition more likely to be someone’s friend. As a result their higher number of connections can increase the average degree of given friendship neighborhoods throughout the network (marked above each node) leading to a Friendship paradox (red nodes) in which most individuals nodes are less popular than the average of their friends. **A2:** When popular individuals are also more likely to be happy, their Happiness becomes more prevalent, raising average happiness levels throughout the friendship circles in the network. A Happiness paradox may result in which most individuals are less happy than their own friends on average. Individuals may cluster based on their Happiness or even the degree to which they experience a Happiness paradox.

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 (see Materials and Methods). We automatically assess each individuals’ Subjective Well-Being (SWB), on a scale of $[-1, +1]$ (See Materials and Methods), by applying a subjective sentiment analysis algorithm to their 3,200 most recent time-coded Tweets (19). Their “Happiness” (quantified as SWB) along with their “Popularity” (quantified as their number of in-network friends) is used to determine: 1) the fraction of individuals that has lower *popularity* than their friends on average (P_{pop}), 2) the fraction of individuals that has lower *happiness* than their friends on average (P_{hap}), and, finally, 3) the correlation between individual happiness and popularity

	Value (%) and 95% CI
P_{pop}	94.3%, [93.7, 94.9] N=5,000
F_{happ}	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

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).

(R(Happiness, Popularity)).

With these definitions, a Friendship or Happiness Paradox for our sample is indicated by P_{pop} and P_{hap} 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 simply calculate Pearson’s R correlation between the SWB values and log(degree) of all subjects in our cohort. The use log(degree) is meant to compensate for the very skewed distribution of degree values in our network.

The distribution of subjects and friendship relations in our social network is very unequal, so 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.

As shown in Table 1, we find that $P_{\text{pop}} = 94\%$ 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 $P_{\text{hap}} = 59\%$, 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 (Spearman’s $R = 0.100$).

The distribution of individual Happiness levels and mean neighborhood happiness in our sample is distinctly bi-modal. This is congruent with the observed distribution of Subjective Well-being across several cultures and nations (20). As shown in Fig. 2 this 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 (19, 21, 22). 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 Fig. 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 (See Suppl. Mat.) and demarcate both groups by simply determining whether a subject and its neighbors fall within 2 standard deviations from the center of one of the components (indicated by 2 ellipses in Fig. 2). Note that this procedure assumes a Gaussian density distribution which roughly matches the quantiles of the empirical density as shown by the contour lines of Fig. 2.

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

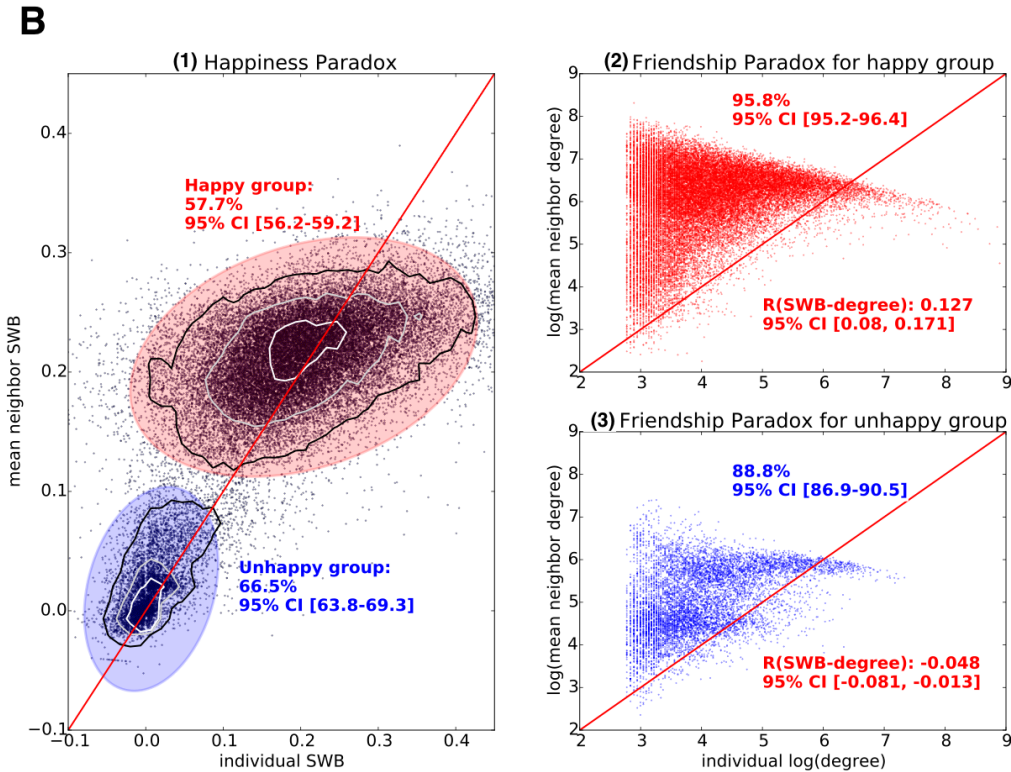


Figure 2: **(B1) Happiness Paradox:** Distribution of individual Happiness (x-axis) vs. average Happiness of one's friend's average (y-axis). Happiness is measured in terms of longitudinal Subjective Well-Being (SWB) scores. Subjects above the red paradox line experience lower happiness (SWB) than their friends' average. The distribution of SWB scores places a majority of subjects well above the diagonal Paradox line. Ellipses indicate the boundaries of 2 Gaussian Mixture Model components used to demarcate a Happy (red) and Unhappy (blue) groups of subjects. Paradox magnitudes are expressed in terms of the percentage of users who experience lower happiness than their friends. The 95% confidence intervals are calculated by a 5000-fold bootstrapping of a 10% sample to determine the sensitivity of our results to random network sampling variations. **(B2) and (B3) Friendship Paradox:** Distribution of individual Popularity (x-axis) vs. average Popularity of one's Friends (y-axis). Popularity is measured in terms of Log(degree) in the Friendship network. Subjects above the red paradox line experience lower popularity than their friends on average. As shown, we find significant Happiness and Friendship Paradoxes for all users, but Happy users experience a stronger Friendship Paradox whereas Unhappy users experience a stronger Happiness Paradox.

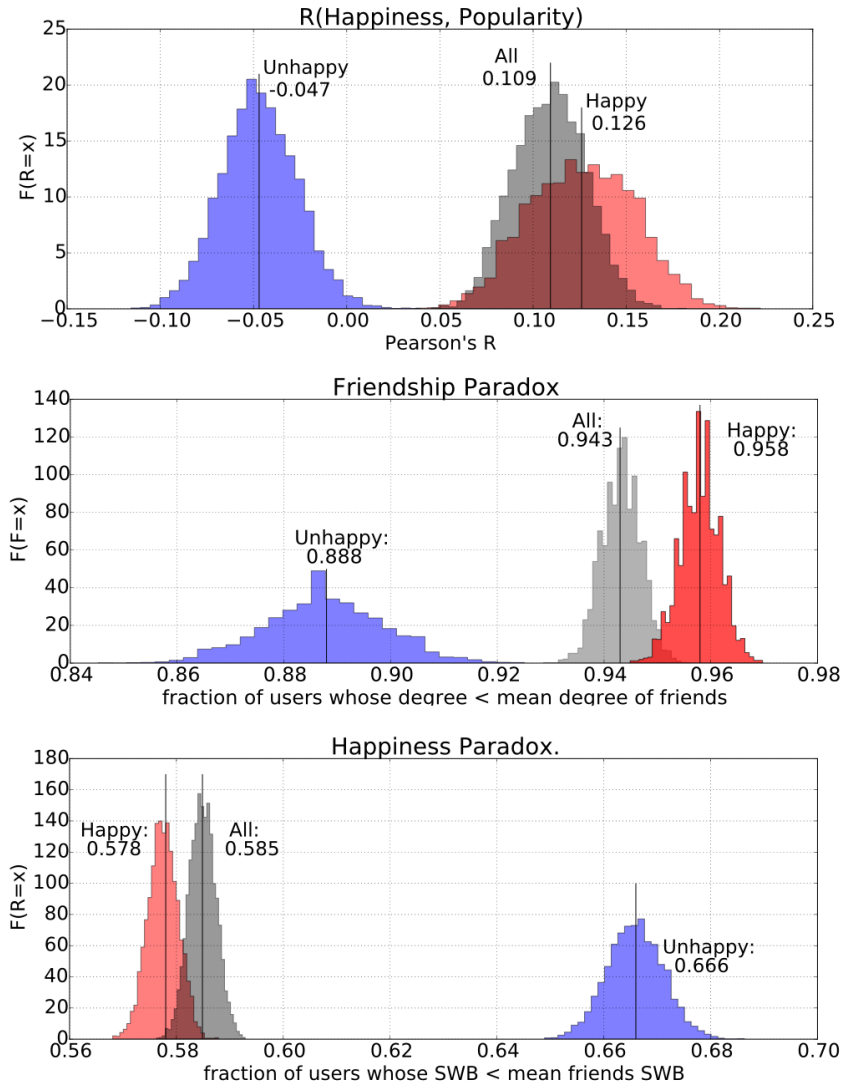


Figure 3: **Top:** Estimated Pearson's R correlation coefficients (95% Confidence Intervals in brackets) between individual Happiness (Subjective Well-Being) vs. individual Popularity (log degree) for All subjects: 0.109 [0.077, 0.140], Happy group: 0.126 [0.081, 0.171], and unhappy group: -0.047 [-0.08, -0.013]. **Middle:** Distribution of Friendship Paradox values for all subjects 0.943 [0.937, 0.949], happy group: 0.958 [0.951, 0.964], and unhappy group 0.888 [0.869, 0.906]. **Bottom:** Distribution of Happiness Paradox values for all subjects: 0.585 [0.581, 0.589], happy group: 0.578 [0.573, 0.582], and unhappy group 0.666 [0.657, 0.674].

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 Friend-

ship 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, such as contagion or increased homophily, we examine the relation between individual happiness and the average happiness of one's neighbors. As indicated by the distribution of individuals in Fig. 2, the strength of the relationship between a subjects' Happiness and the average Happiness of their friends differs significantly between the Happy and Unhappy group. A linear regression analysis indicates that individual Happiness and average friends' Happiness are more strongly related within the Unhappy group ($b=+0.9439$, $F=2.222e+05$, $p<0.001$) than within the Happy group ($b=+0.4459$, $F=34665$, $p<0.001$). This result suggests that unhappy users are 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.

Conclusion

This work constitutes the first direct measurement of a Happiness paradox in social networks, rather than its theoretical derivation from hypothetical network attributes and properties. Our results suggest that previous observations of decreased happiness among social media users may result directly from 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 are 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 one's 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. 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 (23–25) and potentially verbal commiseration and mirroring.

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, no indication 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 expect that these environments may induce longitudinal changes in the public's social behavior and may over time alter the very nature of social relations themselves (26). 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.

In spite of these limitations our results provide a strong indication that widespread social media use may lead to increased levels of social dissatisfaction and unhappiness since individuals will be prone to unfavorably compare their own happiness and popularity to that of others.

Happy social media users will likely 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. We caution against the widespread use of social media given the likelihood that it decreases the happiness and well-being of particularly the most vulnerable groups in society.

Acknowledgements

BG thanks the Moore and Sloan Foundations for support as part of the Moore-Sloan Data Science Environment at NYU.

References and Notes

1. C. L. Apicella, F. W. Marlowe, J. H. Fowler, N. A. Christakis, *Nature* **481**, 497 (2012).
2. R. Dunbar, *Neurosci Biobehav Review* **34**, 260 (2010).
3. A. Steptoe, A. Shankar, P. Demakakos, J. Wardle, *Proceedings of the National Academy of Sciences* **110** (2013).
4. S. Pinker, *The Village Effect: Why Face-to-Face Contact Matters* (Atlantic Books, 2015).
5. A. Perrin, *Report: Pew Research Center* (2015).
6. M. Burke, C. Marlow, T. Lento, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10 (ACM, New York, NY, USA, 2010), pp. 1909–1912.
7. E. Kross, *et al.*, *PLoS ONE* **8**, e69841 (2013).
8. M. Tromholt, M. Lundby, K. Andsbjerg, M. Wiking, The facebook experiment: does social media affect the quality of our lives?, *Tech. rep.*, The Happiness Institute (2015).

9. L. Mitchell, M. R. Frank, K. D. Harris, P. S. Dodds, C. M. Danforth, *PLoS ONE* **8**, e64417 (2013).
10. S. L. Feld, *American Journal of Sociology* **96**, 14641477 (1991).
11. N. O. Hodas, F. Kooti, K. Lerman, *Proceedings of the Seventh International AAI Conference on Weblogs and Social Media* (AAAI Press, 2013).
12. F. Kooti, N. O. Hodas, K. Lerman, *Proceedings of the Eighth International AAI Conference on Weblogs and Social Media* (AAAI Press, 2014).
13. H.-H. Jo, Y.-H. Eom, *Phys. Rev. E* **90**, 022809 (2014).
14. Y.-H. Eom, H.-H. Jo, *Sci. Rep.* **4** (2014).
15. A. J. Oswald¹, S. Wu, *Science* **327**, 576.
16. A. B. Krueger, A. A. Stone, *Science* **346**, 42 (2014).
17. D. R. Posel, D. M. Casale, *Social Indicators Research* **104**, 195 (2011).
18. A. de la Garza, G. Mastrobuoni, A. Sannabe, K. Yamada, The Relative Utility Hypothesis With and Without Self-reported Reference Wages, *Carlo Alberto Notebooks 159*, Collegio Carlo Alberto (2010).
19. J. Bollen, B. Gonçalves, G. Ruan, H. Mao, *Artificial Life* **17**, 237 (2011).
20. B. Stiller, *et al.*, Measuring national well-beingproposed well-being indicators, *Tech. rep.*, The Commission on Measuring Well-being, Japan (2010).
21. M. McPherson, L. Smith-Lovin, J. M. Cook, *Annual Review of Sociology* **27**, 415 (2001).
22. J. H. Fowler, N. A. Christakis, *British Medical Journal* **337** (2008).

23. A. D. I. Kramer, J. E. Guillory, J. T. Hancock, *Proceedings of the National Academy of Sciences* **111** (2014).
24. E. Ferrara, Z. Yang, *PLoS ONE* **10** (2015).
25. E. Ferrara, Z. Yang, *PeerJ Computer Science* **1:e26** (2015).
26. R. Dunbar, *Philos Trans R Soc Lond B Biol Sci* **367**, 2192 (2012).
27. B. Gonçalves, N. Perra, A. Vespignani, *PLoS One* **6**, e22656 (2011).

Supplementary materials

The Friendship Network

To generate a friendship network among Twitter users we started with an initial set of 104,115 users, for which we downloaded the full list of users that they “follow” or that they are “followed” by. Reciprocal “Follow” and “Following” ties are taken as an indication of a friendship relation between the two individuals (27). This resulted in a network of 104,115 nodes connected by 23,640,257 reciprocal edges. We further remove subjects with less than 15 friends in order to improve the reliability of calculating the mean degree and mean SWB values of an individual’s friends. This reduces our final cohort to 39,110 subjects connected by reciprocal friendship relation.

Friendship Paradox

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

$$P(D) = \frac{\|\{u_i \in U : D(u_i) < \mathcal{D}(NN(u_i))\}\|}{\|U\|} \quad (1)$$

When the magnitude of $P(D) > 0.5$ we conclude that the majority of users experiences a Friendship.

Happiness Paradox

For each of the $N = 39,110$ users that fulfill all the requirements listed above, we further collect their complete Twitter history based on which we can assess the SWB of each individual. With this information hand, the magnitude of the Happiness Paradox can be obtained in a way similar

to the way in which we measure the Friendship Paradox. We simply calculate the fraction of users $\|u_i \in U\|$ whose Happiness, denoted $H(u_i)$ respectively, is lower than the average Happiness of their nearest neighbours, denoted \mathcal{H} , vs. the total number of individuals in the network $\|U\|$, or, mathematically:

$$P(H) = \frac{\|\{u_i \in U : H(u_i) < \mathcal{H}(NN(u_i))\}\|}{\|U\|} \quad (2)$$

Bootstrapping

To determine the significance of our results we employ a bootstrapping procedure in which we repeatedly re-sample the set of individuals in our data with replacement and re-calculate our indicators to assess the variance of results resulting from random changes in the underlying population. This procedure allows us to obtain Confidence Intervals for all indicators by determining the 5th and 95th percentile of the results obtained for each of 5000 sub-samples with replacement over the entire set of individuals.

Null model

As mentioned in the text we verify the importance of popularity-happiness correlations by comparing the results we obtained in our dataset with those of a simple null-model. We keep the structure of the network and SWB distributions intact by simply resampling the complete set of SWB values with replacement and re-calculating Eq. 2 and Eq.1. This procedure is performed 20,000 times. We report the 95% confidence intervals for the resulting distribution of paradox values.

Gaussian Mixture Components

Our data contains two data points for each user:

- their own Popularity or Happiness
- the average Popularity or Happiness of their friends

Each user can then be described as a point on a 2-dimensional euclidean plane P spanned by their own popularity or happiness (x) and the average happiness of their group of friends (y).

In this plane, users cleanly separate in 2 clusters in P according to matching levels of popularity or happiness. To determine an objective demarcation criterion we use a Gaussian Mixture Model (GMM) to identify membership in either of the 2 groups. The GMM is trained from our empirical data by means of a standard Expectation-Maximization procedure to identify two 2D Gaussian distributions that are each characterized by a center μ_c and co-variance σ_c to best match the distribution of individuals in P . Each components carries a weight w with which to mix the 2 components to match the probability density function of the data, but we are only concerned with their location to demarcate the two groups of individuals. The gaussian parameter values obtained using the Scikit-learns `sklearn.mixture` package without any constraints on the covariance model are:w

Component	μ_c	σ_c
1 (Happy group)	(0.2037652, 0.21266452)	(0.00186789, 0.00046923)
2 (Unhappy group)	(0.00704093,0.0182976)	(0.00046923, 0.0018294)