ARTICLE

<https://doi.org/10.1057/s41599-019-0261-6>

OPEN

Early prediction of the outcome of Kickstarter campaigns: is the success due to virality?

Alex Kindler¹, Michael Golosovsky¹ & Sorin Solomon¹

ABSTRACT The spread of information, opinions, preferences, and behavior across social media is a crucial feature of the current functioning of our economy, politics, and culture. One of the emerging channels for spreading social collective action and funding of novelty in all these domains is Crowdfunding on various platforms such as Kickstarter, Indiegogo, Sella-band, and many others. The exact spreading mechanism of this collective action is not well-understood. The general belief is that virality plays a crucial role. Namely, the common hypothesis is that the information or behavior propagates through individuals affecting one another, presumably, through the links connecting them in social networks. The aim of our study is to find out the actual spreading mechanism in one particular case: spread of financial support for individual Kickstarter campaigns. To our surprise, our studies show that “virality” plays here only a minor role. We used this result to construct a simple behavior-grounded stochastic predictor of the success of Kickstarter campaigns which is not based on the viral mechanism. The crucial feature of the model underlying the prediction algorithm is that the success of a campaign depends less on the backers influencing one another (“virality”) but rather on the campaign appealing to a particular class of high-pledge backers. This appeal is usually revealed at the very beginning of the campaign and it is an excellent success predictor. The case of Kickstarter is consistent with a recently proposed generic hypothesis that popularity in social media arises more from independent responses by individuals belonging to a large homophily class rather than from percolation, self-exciting processes, and other cooperative mechanisms resulting from mutual influence between individuals. Thus, the very concept of “virality”, which implies contagion between participating individuals, plays only a minor role in the success mechanism proposed hereby. A more appropriate term for the mechanism underlying the social success in our model could be “social appeal” or “social fitness”.

¹The Racah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem, Israel. Correspondence and requests for materials should be addressed to A.K. (email: alexanderkindler@gmail.com)

Introduction and key findings

Predicting collective human behavior is a very difficult task because the causes at the individual level (reciprocal influences, groups of individuals with similar behavior) are often not directly recognizable from the systemic outcome.

Previous works uncovered a series of feedback mechanisms amplifying “microscopic”/individual inputs to the level of systemic transformations: multiplicative dynamics (Levy and Solomon, 1996), social percolation (Solomon et al., 2000), herding (Levy et al., 1994). Surprisingly, our findings for Kickstarter campaigns do not display such effects. We found that the success of a Kickstarter campaign depends on the arousal of a special type of backers’ behavior which can be inferred from the analysis of the statistical distribution of pledges. In Fig. 1 we show representative pledge trajectories for failed and successful Kickstarter campaigns, as a hint to the data features which we use for predicting the campaign outcome.

We measured the distribution of pledge sizes and compared the statistics of pledges for successful and failed campaigns. From the very first day, the successful campaigns display a “fat tail” distribution which follows a power-law dependence (straight line in the double-logarithmic plot of Fig. 2a). Conversely, the failed campaigns display much narrower exponential pledge distributions (straight lines in the semi-logarithmic plot of Fig. 2b).

One can attribute this huge difference in the shape of distributions to the backers’ behavior, as follows:

-In all campaigns there is a body of “rational” backers that pledge some “reasonable” fraction of the target sum leaving to the community of other backers to uphold the campaign. The latter have an exponential distribution of pledges characterized by a well-defined scale (up to 2% of the target sum per day [see Fig. 2]). However, the total pledge accumulated by this group of backers is mostly insufficient to reach the target sum in a

typical 30-day campaign. We call this group “backers of type I”.

-Successful campaigns, in addition to the “backers of type I”, have a small group of backers who are particularly congenial to the campaign idea and react differently and independently of “reasonable proportion”. We call these “backers of type II”. Their “beyond objective/reasonable evaluation criteria” behavior, originating in subjective “enthusiasm”/ “infatuation” with the project, leads to disproportionately high pledges.

Consequently, the distribution of pledges for successful campaigns spreads over a wide scale which is expressed in the power-law statistical distribution of pledges (Fig. 2). Notably, this power-law distribution is formed already in the first day of campaign. The early appearance of the power-law pledge distribution rules out the mechanisms that involve the influence of the previous pledges on the current ones, as it was the case in (Levy and Solomon, 1996; Levy et al., 1994).

One may already extract more general take-home insights from these findings:

-For a campaign to succeed it is usually not enough to have a community of backers reasonably interested in it, it is much more important to have a core of enthusiastic backers committed to contribute more than their “fair share” in the campaign goal. In future research one may try to find out which campaign characteristics evoke this kind of behavior.

-Our data suggest that unlike “viral propagation” whereby connected individuals transmit their enthusiasm one to another, the success of a Kickstarter campaign is a mere result of the fact that the project elicits interest from a number of independent individuals with outstanding behavior. See also (Muchnik et al., 2013) for an earlier observation congenial to the present one.

In any case, our data do not indicate any significant influence of the previous pledges on the current pledges, as it occurs in wealth evolution (Solomon and Richmond, 2001) or citations dynamics, which are governed by the multiplicative or self-exciting processes (Golosovsky and Solomon, 2013). This observation limits the predictive power of the time-pattern fitting techniques used in the past (Etter et al., 2013; Chung and Lee, 2015; Chen et al., 2013; Rao et al., 2014) for prediction success of the Kickstarter campaigns. By contrast, the mere detection of the “backers of type II” in daily pledge distributions is an efficient and very early predictor of success, as we explain below.

Indeed, consider the probability density functions (PDF) of the accumulated pledge distribution (Fig. 3). The PDF for failed campaigns, $\rho_f(q)$, has most of its weight at small daily pledge q and drops exponentially at high q -s. This contrasts the PDFs for successful campaigns, $\rho_s(q)$, that exhibit a maximum at certain $q_{\max}$ that depends on time. Most of the weight of the PDF for successful campaigns is located around this maximum and the weight at low pledges is strongly diminished. Thus, the PDFs of the failed and successful campaigns are distinctly different and occupy different portions of the p - q diagram (Fig. 3) and this prompts us to introduce a threshold. Should the two distributions have a large overlap, one could not find a threshold separating effectively the sets of successful and failed campaigns.

Conversely, if the PDFs are disjointed one could find a threshold (in fact more than one) that effectively separates the distributions for successful and failed campaigns. The actual situation is close to the second possibility: one sees in Fig. 3 that the blue (successful) and red (failed) PDFs intersect. This allows one to choose the threshold indicated by a vertical arrow. By plotting the position of the threshold for everyday of the

Fig. 1 Pledge dynamics of three successful Kickstarter campaigns (S1, S2, S3) and three failed campaigns (F1, F2, F3). All pledges are divided by the fundraising goal. The duration of all campaigns is 30 days. Trajectories of the successful campaigns display one or more jumps of size much larger than the average daily pledge. By contrast, the daily pledges of the failed projects are all of the same order of magnitude. The red line indicates a threshold that separates very neatly the successful trajectories S1, S2, S3 from the failed trajectories F1, F2, F3. A detailed discussion of the threshold and how it is derived appears later in the text

Fig. 2 Cumulative distribution functions (CDF) of accumulated pledges up to day 1, 10, and 20 of the Kickstarter campaigns with 30 days duration. **a** Double-logarithmic scale. **b** Semi-logarithmic scale. The successful project CDFs (S1, S10, S20) exhibit the power-law tail with the exponent -1.35 shown by straight lines in **a**. The failed projects (F1, F10, F20) have CDFs with exponential tails shown as straight lines in **b**

Fig. 3 The number of successful, $\rho_s N_s$, and failed, $\rho_f N_f$, Kickstarter campaigns as a function of accumulated normalized pledge q by the day 10 of the campaign. N_s and N_f are the total number of successful and failed campaigns (from the set of 7141 campaigns), while $\rho_s(q)$ and $\rho_f(q)$ are the corresponding probability density functions. The clear distinction between the two plots allows a prediction criterion based on a threshold: the campaigns that accumulated pledges above the threshold are likely to succeed while those that didn't manage to accumulate pledges above the threshold are likely to fail. The threshold is chosen in such a way as to minimize the total error, namely the sum of the number of campaigns below threshold that eventually succeeded and the number of campaigns above the threshold that eventually failed. In the text we show that this criterion corresponds to the choice of threshold at the intersection of two curves

campaign one obtains the red threshold line in Fig. 1. This threshold is a basis for our predictor of success.

Background information on Kickstarter

A typical crowdfunding campaign is initiated by an entrepreneur proposing his project through the Internet platform. The entrepreneur indicates two most important parameters: the target investment necessary for performing the project and the deadline for reaching this target. Once the project is published on the Internet, any individual can become a backer by pledging funds to support the campaign. The number of backers and the amount of daily pledges are publicly available. If the deadline comes and the total amount of pledges equals or exceeds the target investment, then the campaign is considered as a success and its implementation begins. If the target investment is not achieved, the campaign is considered a failure and is discontinued.

The impact of accurately predicting whether the campaign succeeds (or to what extent it succeeds) serves many purposes. First, among the multitude of similar campaigns, such as designer shoes or computer games of a specific genre, early classification may help promoters direct their energies towards campaigns that need it most. From the perspective of potential backers, it is also of interest to evaluate the stakes at hand when pledging: the consequences of pledging for a campaign that is likely to succeed even without their pledge, or for a one that is more likely to fail, are different.

It is thus natural to consider predictor algorithms that treat crowdfunding as a dynamical stochastic or causal process whose behavior can be studied, understood, predicted, and influenced. Our goal is to develop such algorithm. In contrast to machine-learning algorithms, we base our predictor on the deep understanding of the dynamics of crowdfunding. The present study is important because it can be extended to similar problems arising in very different contexts such as elections, betting, economic and financial markets, etc. Recall in this context the self-referential dynamical aspects of the famous Keynesian beauty contest game (Keynes, 2018).

Previous studies of crowdfunding are summarized in (Kuppuswamy and Bayus, 2015). Chung and Lee (2015), suggested a set of predictors of success based on pledge time-series, tweets, and campaign/backer graphs. These predictors achieve spectacular 76% accuracy after 4 h of campaign launch. In addition, Etter et al., 2013, constructed an Internet site which can run such predictor in real time. Greenberg et al. (2013) employed and trained support vector machine (SVM), and thus created a classifier system for predicting campaign success with a 68% accuracy based solely on initial conditions. Qiu (2013) showed that posting the campaign on the Kickstarter homepage has a positive effect on receiving pledges. Mitra and Gilbert (2014) analyzed the contents of campaign web-pages and showed that their language may be used to improve prediction. Chung and Lee (2015) developed models which predict success and total amount of expected pledged money. The above models operate with time-resolved Kickstarter and Twitter data and yield approximately 90% accuracy when the campaign achieves 30% stage (the time window between the launch of the project and its deadline). Chen et al. (2013) created and trained an SVM to predict the probability of success. This predictor achieves 67% accuracy at the launch phase, and approximately 90% accuracy when the campaign proceeds to the 40% stage. Rao et al. (2014) used decision-tree models to investigate the extent to which simple inflows and first-order derivatives can predict campaign success. Basing on the initial 15% of money inflows they could predict success with 84% accuracy.

Data analysis

We analyzed the publicly available data on the Kickstarter campaigns using the database reported in (Etter et al., 2013). The time-series data covering $N_0 = 16,043$ campaigns were assessed on a daily basis (daily pledges, daily number of backers), without access to individual pledges or backer specifications. First, we analyzed campaigns with $t_0 = 30$ day duration as the in-sample database to determine the parameters that discern campaign success from failure. These included $N_s = 3177$ successful and $N_f = 3964$ failed campaigns, whereas the rest (4562 successful and 4339 failed campaigns) were used as the out-of-sample database on which we tested our success predictor.

We denote by $Q(t)$ the accumulated pledge by day t of the campaign and by Q_0 we denote the goal, namely, the target pledge. Since the projects differ greatly in their target pledge, to put all campaigns on the same standing we considered the reduced pledge, $q(t) = Q(t)/Q_0$. We found that the dynamics of pledge accumulation depends more on $q(t)$ rather than on the target pledge Q_0 , hence this reduction is justified.

We divided all campaigns onto successful and failed ones and studied their statistics separately. To this end we built statistical distribution of the accumulated pledges for every stage (day) of the campaign and characterized them by CDF, $\Pi(q, t) = \int_q^\infty \rho(q', t) dq'$, where ρ is the PDF. The following relation holds $\rho_0 N_0 + \rho_s N_s + \rho_f N_f$ where ρ_s and ρ_f are the PDFs for successful and failed projects correspondingly, and ρ_0 is the PDF for all projects together.

Though we only have data on the total daily pledges rather than individual backer pledges, we can use a large daily pledge signal as a proxy for a large individual backer pledge. Indeed, our analysis of the number of backers that contributed in each day revealed that the days with large total pledges did not display significantly more backers than ordinary days.

From Figs. 2 and 3 one sees that the set of successful campaigns has from the very start a signature that is quite easy to resolve: this is the presence of fat tails in the pledge distribution of successful campaigns which is indicative of type II backers.

Fig. 4 **a** CDFs of accumulated pledges for successful campaigns versus time from the launch of campaign. Only first 7 days are displayed. All pledges are divided by the fundraising goal. The shades of gray for each value of the reduced accumulated pledge q indicate what fraction of campaigns have accumulated pledges that are below q . The red line is the threshold. It is clearly seen that there are only few successful campaigns that at any time fall below the threshold. **b** CDFs of accumulated pledges for the failed campaigns. It is clearly seen that there are almost no failed campaigns above the threshold. We show in the text how to use this fact to make early prediction of the campaign outcome

In Fig. 4 one sees that there exists a line (colored in red in the figure) such that the vast majority of successful campaigns lie above it while the failed campaigns lie mostly below it. Indeed, one sees that the red line is placed:

-In Fig. 4a it is above the white region which contains only the 10% successful campaigns with the lowest pledges.

-In Fig. 4b it is below the white region which contains only the 10% failed campaigns with the highest pledges.

On this basis one can predict success quite precisely on each day by establishing whether the accumulated pledges at that day are above or below the threshold.

In what follows we explain how we chose the threshold q_0 (red curve in Figs. 1 and 4) and why it provides a useful criterion for predicting the outcome of a campaign. Following Fig. 3, we consider the following function

$$F(q_0) = N_f \cdot \Pi_f(q_0, t) + N_s [1 - \Pi_s(q_0, t)] \quad (1)$$

The first term represents the total number of failed campaigns that by time t garnered total pledge exceeding q_0 , while the second term represents the number of successful campaigns that by time t garnered total pledge below q_0 . If we consider $q_0(t)$ as a predictor of success after stage t , the first term in Eq. (1) is the number of false positive events and the second term represents the number of false negative events. Our prediction algorithm minimizes the number of false events at each time t , namely, for each stage of the campaign we find $q_0(t)$ that minimizes $F(q_0)$.

Fig. 5 Prediction accuracy $A(q_0, t)$ of our study as compared with others. Our data are for campaigns with 30 days duration (the vast majority of Kickstarter projects)

To measure the success of our prediction algorithm we define the prediction accuracy as:

$$A(q_0, t) = \frac{N_s \cdot \Pi_s(q_0, t) + N_f [1 - \Pi_f(q_0, t)]}{N_0} \quad (2)$$

Namely, for each stage of the campaign we add up the number of false positive predictions (the total number of failed campaigns above the threshold in Fig. 4b) and false negative predictions (the total number of successful campaigns below the threshold in Fig. 4a). Then, we divide this sum by the total number of campaigns.

To allow comparison with previous studies we plot on Fig. 5 the prediction accuracy $A(q_0, t)$ reported in each of these studies as a function of the campaign stage. We observe that the accuracy of our predictor exceeds that of other groups, especially in the crucial days at the beginning of the campaigns.

The threshold is a binary predictor, offering the best estimate of whether a given project is more likely to succeed or fail, regardless of how far above or below the threshold the project's pledge trajectory is. In what follows we consider a more accurate tool which calculates the probability of a project to succeed. We determine the probability of success of a campaign that accumulated the normalized pledge q at time t as follows:

$$P_{\text{success}}(q, t) = \frac{\rho_s N_s}{\rho_s N_s + \rho_f N_f} = \frac{\rho_s(q, t) N_s}{\rho_0(q, t) N_0} \quad (3)$$

Figure 6 plots this probability for different days of campaign. In fact, Fig. 6 can be used as a graphic tool similar to a “table” with continuum number of columns to read directly the probability of success as the equal probability line on which the x -coordinates and y -coordinates, that correspond to the t and q of the campaign, meet.

As an example of the use of this map, suppose that at a campaign stage corresponding to 70% of the allocated time, the campaign accumulated 30% of the target pledge. We observe that the corresponding dashed straight lines intersect at the green contour line representing success probability 0.5.

As it turns out, when the method for deriving Fig. 6 is applied to the out-of-sample set of campaigns, one obtains that only 0.5% of the predictions are off by more than 10%.

Discussion

Our analysis shows a striking difference between the funding dynamics of successful and failed Kickstarter campaigns, even at a very early campaign stage. Figures 2 and 3 show that these funding patterns are fundamentally different, since failed campaigns display exponential pledge distribution, while successful campaigns display a heavy-tail pledge distribution.

Fig. 6 Probability of success according to reduced accumulated pledge q at each campaign stage. Y-axis is the accumulated pledge normalized by the fundraising goal (target pledge), X-axis is the project stage normalized to the total campaign duration (since most of the campaigns had duration of 30 days we also labeled the X-axis from day 1 to day 30 on the top X-axis). The colored curves indicate the contours of equal probability of success

For successful campaigns, the tail of the pledge distribution follows the power-law dependence with the exponent close to unity (1.35), implying a spread of pledge sizes over a very wide scale.

Figure 5 shows that even limiting our method to a binary choice for the purpose of comparability with other works, our method offers better accuracy already in the first three days of the campaign. Moreover, our method is transparent and may be easily applied by anyone (especially by reading the success probability from the Fig. 6). Possibly, these results may also help developing strategies to promote campaigns by intervening in the pledge process (Muchnik et al., 2013).

While the data allow us to exclude multiplicative random walk or self-exciting processes in the case of Kickstarter campaigns it is still unclear what is the origin of the heavy-tail pledge distribution and especially how to induce it. This is a crucial point to understand and reproduce successful campaigns: for Kickstarter campaigns the very name “viral”, that implies contagion between participating individuals, is put under question by the present data which display heavy-tail distributions from the very beginning of the campaign. Thus, the Kickstarter campaigns better fit the “homophily” hypothesis (Muchnik et al., 2013), which postulates that the reason for individuals responding to the same stimulus is not contagion over their connections but the mere fact that they have similar preferences.

The hope to understand the mechanism responsible for the success of a Kickstarter campaign is further encouraged by the fact that in contrast to the “black-box” approach of machine learning or statistical inference, our analysis exploits the observed behavior of the relevant groups involved in the process.

Conclusions

We have constructed an easy and accurate tool for predicting success of a Kickstarter campaign, even at its early stage. We found almost no correlation between the pledges made at

different days that would indicate the influence of early pledges on the later ones. Thus, our predictive tool is not based on time-series analysis. Our empirical findings favor an a priori intrinsic property of the community of backers rather than the communication/reciprocal influence among them as the key to campaign success. It is left for future studies to check if this “non-viral” collective response is behind social success of other classes of processes such as YouTube views, Twitter tweets, Facebook, and blog likes.

Data availability

The datasets analyzed in this study have been retrieved from the Sidekick.com repository, <http://sidekick.epfl.ch/data>, collected by V. Etter, M. Grossglauser, and P. Thiran.

Received: 23 October 2018 Accepted: 2 May 2019

Published online: 14 May 2019

References

- Chen, K, Jones B, Kim I, Schlamp B (2013) Kickpredict: Predicting kickstarter success. Technical Report, California Institute of Technology
- Chung J, Lee K (2015) A long-term study of a crowdfunding platform: predicting project success and fundraising amount. In Proceedings of the 26th ACM Conference on Hypertext and Social Media. ACM. pp 211–220
- Etter V, Grossglauser M, Thiran P (2013) Launch hard or go home!: predicting the success of kickstarter campaigns. In Proceedings of the first ACM conference on Online social networks. ACM. pp 177–182
- Golosovsky M, Solomon S (2013) The transition towards immortality: non-linear autocatalytic growth of citations to scientific papers. *J Stat Phys* 151 (1–2):340–354
- Greenberg MD, Pardo B, Hariharan K, Gerber E (2013) Crowdfunding support tools: predicting success and failure. In CHI’13 Extended Abstracts on Human Factors in Computing Systems. ACM. pp 1815–1820
- Keynes JM (2018) The general theory of employment, interest, and money. Springer, Cham
- Kuppuswamy V, Bayus BL (2015) A review of crowdfunding research and findings. handbook of new product development research (Forthcoming). <https://ssrn.com/abstract=2685739>. Accessed 28 Oct 2015

- Levy M, Solomon S (1996) Power laws are logarithmic Boltzmann laws. *Int J Mod Phys C* 7(04):595–601
- Levy M, Levy H, Solomon S (1994) A microscopic model of the stock market: cycles, booms, and crashes. *Econ Lett* 45(1):103–111
- Muchnik L, Aral S, Taylor SJ (2013) Social influence bias: a randomized experiment. *Science* 341(6146):647–651
- Mitra T, Gilbert E (2014) The language that gets people to give: phrases that predict success on kickstarter. In Proceedings of the 17th ACM conference on Computer supported cooperative work and social computing. ACM. pp 49–61
- Qiu C (2013) Issues in crowdfunding: theoretical and empirical investigation on Kickstarter. <https://ssrn.com/abstract=2345872>. Accessed 27 Oct 2013
- Rao H, Xu A, Yang X, Fu WT (2014) Emerging dynamics in crowdfunding campaigns. In International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction. Springer, Cham, pp 333–340
- Solomon S, Weisbuch G, de Arcangelis L, Jan N, Stauffer D (2000) Social percolation models. *Phys A* 277(1–2):239–247
- Solomon S, Richmond P (2001) Power laws of wealth, market order volumes and market returns. *Phys A* 299(1–2):188–197

Additional information

Competing interests: The authors declare no competing interests.

Reprints and permission information is available online at <http://www.nature.com/reprints>

Publisher’s note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019