A Long-Term Study of a Crowdfunding Platform: Predicting Project Success and Fundraising Amount

Jinwook Chung and Kyumin Lee
Department of Computer Science
Utah State University
Logan, Utah 84322
{jinwook.chung@aggiemail.usu.edu, kyumin.lee@usu.edu}

ABSTRACT

Crowdfunding platforms have become important sites where people can create projects to seek funds toward turning their ideas into products, and back someone else's projects. As news media have reported successfully funded projects (e.g., Pebble Time, Coolest Cooler), more people have joined crowdfunding platforms and launched projects. But in spite of rapid growth of the number of users and projects, a project success rate at large has been decreasing because of launching projects without enough preparation and experience. To solve the problem, in this paper we (i) collect the largest datasets from Kickstarter, consisting of all project profiles, corresponding user profiles, projects' temporal data and users' social media information; (ii) analyze characteristics of successful projects, behaviors of users and understand dynamics of the crowdfunding platform; (iii) propose novel statistical approaches to predict whether a project will be successful and a range of expected pledged money of the project; and (iv) develop predictive models and evaluate performance of the models. Our experimental results show that the predictive models can effectively predict project success and a range of expected pledged money.

Categories and Subject Descriptors: H.3.5 [Online Information Services]: Web-based services

Keywords: crowdfunding; kickstarter; twitter; project success; fundraising amount

1. INTRODUCTION

Crowdfunding platforms have successfully connected millions of individual crowdfunding backers to a variety of new ventures and projects, and these backers have spent over a billion dollars on these ventures and projects [8]. From reward-based crowdfunding platforms like Kickstarter, Indiegogo, and RocketHub, to donation-based crowdfunding platforms like GoFundMe and GiveForwad, to equity-based crowdfunding platforms like CrowdCube, EarlyShares and Seedrs - these platforms have shown the effectiveness of funding projects from millions of individual users. The US Congress has encouraged crowdfunding as a source of capital for new ventures via the JOBS Act [2].

An example of successfully funded projects is E-paper watch project. The E-paper watch project for smartphones on a crowdfunding platform was created by Pebble Technology corporation on April 2012 in Kickstarter, expecting \$100,000 investment. Surprisingly, in 2 hours right after launching the project, pledged money was already exceeding \$100,000. In the end of the project period (about 5 weeks), the company was able to get investment over 10 million dollars [25]. This example shows the power of collective investment and a crowdfunding platform, and a new way to raise funding from the crowds.

Even though the number of projects and amount of pledged funds on crowdfunding platforms has dramatically grown in the past few years, success rate of projects at large has been decreasing. Besides, little is known about dynamics of crowdfunding platforms and strategies to make a project successful. To fill the gap, in this paper we are interested to (i) analyze Kickstarter, the most popular crowdfunding platform and the 373rd most popular site as of March 2015 [4]; and (ii) propose statistical approaches to predict not only whether a project will be successful, but also how much a project will get invested. Kickstarter has an All-or-Nothing policy. If a project reaches pledged money lower than its goal, its creator will receive nothing. Predicting a range of expected pledged money is an important research problem.

Specifically, we analyze behaviors of users on Kickstarter by answering following research questions: Are users only interested in creating and launching their own projects? or Do they support other projects? Has the number of newly joined users been increased over time? Have experienced users achieved a higher project success rate? Then, we analyze characteristics of projects by answering following research questions: How many projects have been created over time? What percent of project has been successfully funded? Can we observe distinguishing characteristics between successful projects and failed projects? Based on the analysis and study, we answer following research questions: Can we build predictive models which can predict not only whether a project will be successful, but also a range of expected pledged money of the project? By adding a project's temporal data (e.g., daily pledged money and daily increased number of backers) and a project creator's social media information, can we even improve performance of the predicative models further?

HT'15, September 1–4, 2015, Guzelyurt, TRNC, Cyprus.
© 2015 ACM. ISBN 978-1-4503-3395-5/15/09 ...\$15.00.
DOI: http://dx.doi.org/10.1145/2700171.2791045.

|Kickstarter projects|	151,608
|Kickstarter users|	142,890
|Kickstarter projects with temporal data|	74,053
|Kickstarter projects with Twitter user profiles|	21,028

Table 1: Datasets.

Toward answering these questions, we make the following contributions in this paper:

- We collected the largest datasets, consisting of all Kickstarter project pages, user pages, each project's temporal data and each user's Twitter account information, and then conducted comprehensive analysis to understand behaviors of Kickstarter users and characteristics of projects.
- Based on the analysis, we proposed and extracted four types of features toward developing project success predictors and pledged money range predictors. To our knowledge, this is the first work to study how to predict a range of expected pledged money of a project.
- Finally, we developed predictive models and thoroughly evaluated performance of these models. Our experimental results show that these models can effectively predict whether a project will be successful and a range of expected pledged money.

2. DATASETS

To analyze projects and users on crowdfunding platforms, and understand whether adding social media information would improve project success prediction and pledged money prediction rates, first we collected data from Kickstarter, the most popular crowdfunding platform, and Twitter, one of the most popular social media sites. The following subsections present our data collection strategy and datasets.

2.1 Kickstarter Dataset

Kickstarter is a popular crowdfunding platform where users create and back projects. As of March 2015, it is the 373rd most visited site in the world according to Alexa [4].

Static Data. Our Kickstarter data collection goal was to collect all Kickstarter pages and corresponding user pages, but Kickstarter site only shows currently active projects and some of the most funded projects. Fortunately, Kicktraq site[1] has archived all project page URLs of Kickstarter. Given a Kicktraq project URL[2], by replacing Kicktraq hostname (i.e, `www.kicktraq.com`) of the project URL with Kickstarter hostname (i.e., `www.kickstarter.com`), we were able to obtain the Kickstarter project page URL[3].

Specifically, our data collection approach was to collect all project pages on Kicktraq, extract each project URL, and replace its hostname with Kickstarter hostname. Then we collected each Kickstarter project page and corresponding user page. Note that even though Kickstarter do not reveal an old project page (i.e., a project's campaign duration was ended), if we know the project URL, we can still access the project page on Kickstarter.

Finally, we collected 168,851 project pages which were created between 2009 and September 2014. Note that Kickstarter site was launched in 2009. A project page consists of a project duration, funding goal, project description, rewards description and so on. We also collected corresponding 146,721 distinct user pages each of which consists of bio, account longevity, location information, the number of backed projects, the number of created projects, and so on. Among 168,851 project pages, we filtered 17,243 projects which have been either canceled or suspended, or in which the project creator's account has been canceled or suspended. Among 146,721 user pages, we filtered corresponding 8,679 user pages. Finally, 151,608 project pages and 142,890 user pages presented in Table 1, have been used in the rest of this paper.

Temporal Data. To analyze and understand how much each project has been pledged/invested daily and how many backers each project has attracted daily, whether incorporating these temporal data (i.e., daily pledged money and daily increased number of backers during a project duration) can improve project success prediction and expected pledged money prediction rates, we collected temporal data of 74,053 projects which were created between March 2013 and August 2014 and were ended by September 2014.

2.2 Twitter Dataset

What if we add social media information of a project creator to build predictive models? Can a project creator's social media information improve project success and expected pledged money prediction rates? Can we link a project creator's account on Kickstarter to Twitter? To answer these questions, we checked project creators' Kickstarter profiles. Interestingly 19,138 users (13.4% of all users in our dataset), who created 22,408 projects, linked their Twitter user profile pages (i.e., URLs) to their Kickstarter user profile pages. To use these users' Twitter account information in experiments, we collected their Twitter account information. Specifically, we extracted a Twitter user profile URL from each Kickstarter user profile, and then collected the user's Twitter profile information consisting of the basic profile information (e.g., a number of tweets, a number of following and a number of followers) and tweets posted during a project period. In a step of the Twitter user profile collection, we noticed that some of Twitter accounts had been either suspended or deleted. By filtering these accounts, finally, we collected 17,908 Twitter user profiles and tweets, and then combined these Twitter information with 21,028 Kickstarter project pages created by the 17,908 users.

3. ANALYZING KICKSTARTER USERS AND PROJECTS

In the previous section, we presented our data collection strategy and datasets. Now we turn to analyze Kickstarter users and projects.

3.1 Analysis of Users

Given 142,890 user profiles, we are interested in answering following research questions: Are users only interested in creating and launching their own projects? or Do they support other projects? Has the number of new users joined Kickstarter been increased over time? Do experienced users have a higher probability to make a project successful?

[1]`http://www.kicktraq.com/archive/`
[2]`http://www.kicktraq.com/projects/fpa/launch-the-first-person-arts-podcast/`
[3]`https://www.kickstarter.com/projects/fpa/launch-the-first-person-arts-podcast/`

Figure 1: Number of newly joined Kickstarter users in each month.

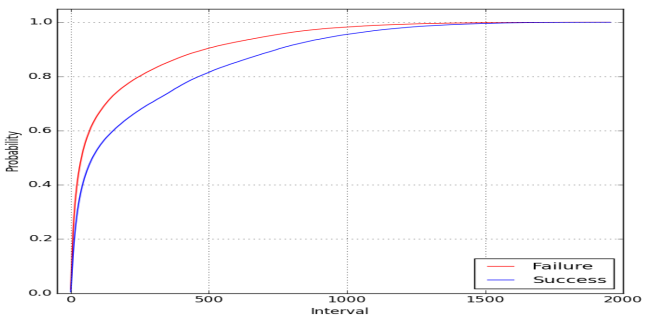

Figure 2: CDFs of intervals between user joined date and project creation date (Days).

First of all, we present general statistics of users in Table 2. The user statistics show that average number of backed projects and created projects are 3.48 and 1.19, respectively. It means that users backed larger number of projects and created less number of their own projects. Each user linked 1.75 websites on average into her profile so that she can get trust from potential investors. Examples of websites are company sites and user profile pages in social networking sites such as Twitter and YouTube. 13.4% Kickstarter users linked their Twitter pages, and 6.89% Kickstarter users linked their Youtube pages.

Next, we categorized Kickstarter users based on their project backing and creating activities. We found two groups of users: (i) all-time creator (AT creator), who only created projects and did not back other projects; and (ii) active user, who not only created her own projects but also backed other projects. As shown in Table 3, there are 66,262 (46.4%) all-time creators and 76,628 (53.6%) active users. Each all-time creator created 1.12 projects on average. These creators were only interested in creating their own projects and sought funds. Interestingly, the average number of created projects per all-time creator reveals that these creators created just one or two projects. However, each of 76,628 active users created 1.25 projects and backed 6.49 projects on av-

	Total
Total number of users	142,890
Number of backed projects per user	3.48
Number of created projects per user	1.19
Number of websites per user	1.75
Twitter connected	13.4% users
YouTube connected	6.89% users

Table 2: Statistics of Kickstarter users.

	Number	Avg. backed	Avg. created
AT creators	66,262	N/A	1.12
Active users	76,628	6.49	1.25

Table 3: Two groups of users: all-time (AT) creators and active users.

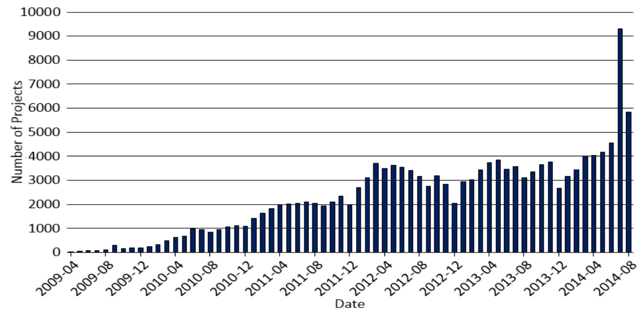

Figure 3: Number of created projects per month has been increased over time with some fluctuation.

erage. These active users created a little more projects than all-time creators, and backed many other projects.

Next, we analyze how many new users joined Kickstarter over time. Figure 1 shows the number of newly joined Kickstarter users per month. Overall, the number of newly joined users per month has been linearly increased until May 2012, and then has been decreased until June 2014 with some fluctuation. In July 2014, there was a huge spike. Note that we tried to understand why there was a huge spike in July 2014 by checking news articles, but we were not able to find a concrete reason. Interesting observation is that the number of newly joined users was the lowest during winter season, especially, December in each year. We conjecture that since November and December contains several holidays, people may delay to join Kickstarter.

A follow-up question is "Do experienced users achieve a higher project success rate?". We measured experience of a user based on when they create a project after joining Kickstarter. Figure 2 shows cumulative distribution functions (CDFs) of intervals between user joined date and project creation date in successful projects and failed projects. As we expected, successful projects had longer intervals. We conjecture that since users with longer intervals become more experienced and familiar with Kickstarter platform, their projects have become successful with a higher probability.

3.2 Analysis of Projects

So far we have analyzed user profiles. We now analyze Kickstarter projects. Interesting research questions are: How many projects have been created over time? What percent of projects has been successfully funded? Can we observe clearly different properties between successfully funded projects and failed projects? To answer these questions, we analyzed Kickstarter project dataset presented in Table 1.

Number of projects and project success rate over time. Figure 3 shows how the number of projects has been changed over time. Overall, the number of created projects per month has been increased over time with some fluctuation. Interestingly, lower number of projects in December of each year (e.g., 2011, 2012 and 2013) has been created. Another interesting observation was that the largest number of projects (9,316 projects) were created in July 2014. The

	Success	Failure	Total
Percentage (%)	46	54	100
Classified project count	69,448	82,160	151,608
Duration (days)	33.21	36.2	34.83
Project Goal (USD)	8,364.34	35,201.89	22,891.15
Final money pledged (USD)	16,027.96	1,454.18	8,139.37
Number of images	4.63	3.37	3.95
Number of videos	1.18	0.93	1.04
Number of FAQs	0.84	0.39	0.6
Number of rewards	9.69	7.49	8.5
Number of updates	9.59	1.59	5.26
Number of project comments	77.52	2.45	36.89
Facebook connected (%)	61.00	59.00	60.00
Number of FB friends	583.48	395.15	481.54
Number of backers	211.16	19.34	107.33

Table 4: Statistics of Kickstarter projects.

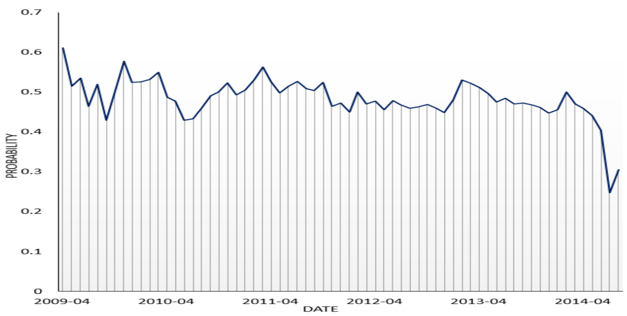

Figure 4: Project success rate in each month.

phenomena would be related to the number of newly joined users per month shown in Figure 1 in which less number of users joined Kickstarter during Winter season, especially in December in each year, and many users joined in July 2014.

Next, we are interested in analyzing how project success rate has been changed over time. We grouped projects by their launched year and month. Interestingly, the success rate has been fluctuated and overall project success rate in each month has been decreased over time as shown in Figure 4. In July 2014, the success rate was dramatically decreased. We conjecture that since many users joined Kickstarter in July 2014, these first-time project creators caused the sharp decrease of success rate.

Statistics of successful projects and failed projects. Next, we analyze statistics of successful projects and failed projects. Table 4 presents the statistics of Kickstarter projects. Overall, percentage of the successful projects in our dataset is about 46%. In other words, 54% of all projects was failed. We can clearly observe that the successful projects had shorter project duration, lower funding goal, more active engagements and larger number of social network friends than failed projects.

Figure 5 shows more detailed information about how project success rate was changed when a project duration was increased. This figure clearly shows that project success rate was higher when a projet duration was shorter. Intuitively, people may think that longer project duration would be helpful to get more fund, but this analysis reveals the opposite result. To show how many projects have what duration, we plotted Figure 6. 39.7% (60,191 projects) of all projects had 30 day duration and then 6.5% (9,784 projects) of all projects had 60 day duration. We conjecture that since 30 day duration is the default duration on Kickstarter, many users just chose 30 day duration for their projects.

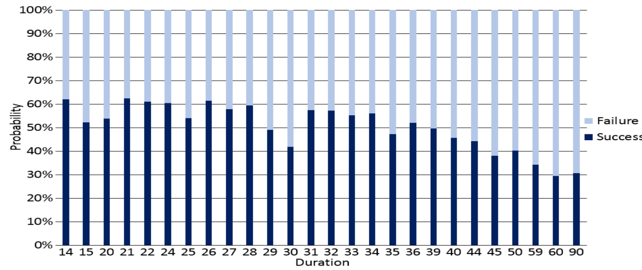

Figure 5: Project success and failure rates according to a duration that more than 1,000 projects has.

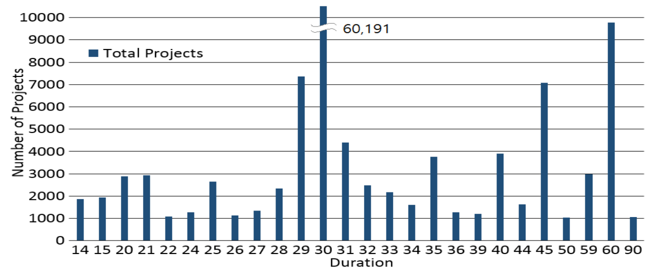

Figure 6: Number of Projects according to a duration that more than 1,000 projects has.

While the average project goal of successful projects was 3 times less than failed projects, the average pledged money of successful projects was 10 times more than failed projects. Project creators of successful projects spent more time to make better project description by adding a larger number of images, videos, FAQ and reward types. The creators also frequently updated their projects. Interestingly, project creators of the successful projects had a larger number of Facebook friends. It means that the creators' Facebook friends might help for their project success by backing the projects or spreading information of the projects to other people [19].

When a user creates a project on Kickstarter, she can choose a category of the project. Does a category of a project affect a project success rate? To answer this question, we analyzed project success rate according to each category. As you can see in Figure 7, projects in Dance, Music, Theater, Comics and Art categories achieved between 50% and 72% success rate which is greater than the average success rate of all projects (again, 46% success rate).

Location. A user can add location information when she creates a project. We checked our dataset to see how many

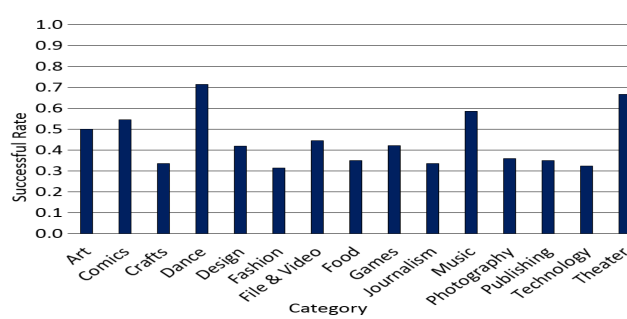

Figure 7: Project success rate under each of 15 categories.

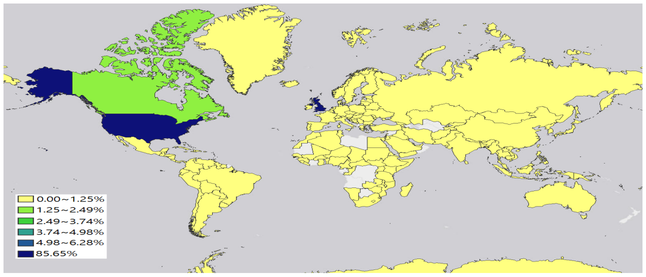

Figure 8: Distribution of projects in the world.

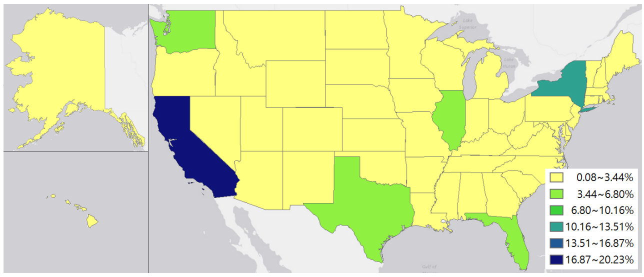

Figure 9: Distribution of projects in US.

projects contain location information. Surprisingly, 99% project pages contained location information. After extracting the location information from the projects, we plotted distribution of projects on the world map in Figure 8. 85.65% projects were created in US. The next largest number of projects were created in the United Kingdom (6.23%), Canada (2.20%), Australia (1%)and Germany (0.92%). Overall, the majority of projects were created in the western countries. The project distribution across countries makes sense because initially only US based projects on Kickstarter were created, and then the company allowed users in other countries to launch projects since October 2012. Since over 85% projects were created in US, we plotted distribution of the projects on US map in Figure 9. Top 5 states are California (20.23%), New York (12.93%), Texas (5.45%), Florida (4.57%) and Illinois (4.03%). This distribution mostly follows population of each state.

A follow-up question is how project distribution across states in US is related to projects success rate. To answer this question, we plotted project success rate of each state in Figure 10. Top 5 states with the highest success rate are Vermont (63.81%), Massachusetts (58.49%), New York (58.46%), Rhode Island (58.33%) and Oregon (53.56%). Except New York state, small number of projects were created in the four states. To make a concrete conclusion, we measured Pearson correlation between distribution of projects and project success rate. The correlation value was 0.25 which indicates that they are not significantly correlated.

Analysis of Kickstarter Temporal Data. As we presented in Table 1, we collected temporal data of 74,053 projects (e.g., daily pledged money and daily increased number of backers). Using these temporal data, we analyzed what percent of total pledged money and what percent of backers each project got over time after launching a project. Since each project has different duration (e.g., 30 days or 60 days), first, we converted each project duration to 100 states (time slots). Then, in each state, we measured percent of pledged money and number of backers.

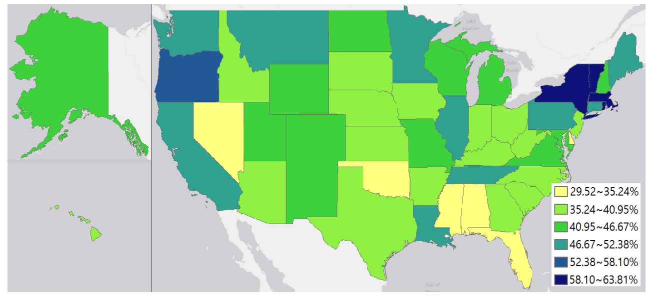

Figure 10: Project success rate across states in US.

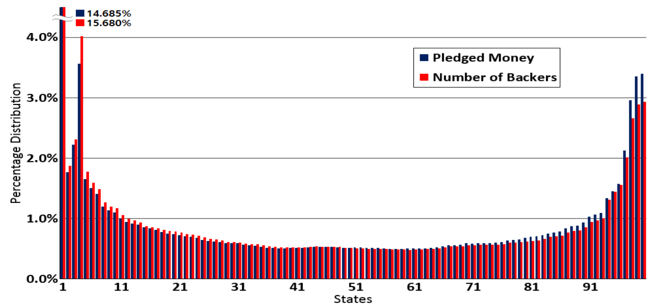

Figure 11: Percentage distribution of pledged money and number of backers per state.

Figure 11 shows the percentage distribution of pledged money and number of backers per state over time. One of the most interesting observations is that the largest amount of money was pledged in the beginning and end of a project. For example, 14.69% money was pledged and 15.68% backers were obtained in the first state. Other researchers also observed the same phenomena in smaller datasets [13, 15].

Another interesting observation is that there is another spike after the first spike in the beginning of project durations. We conjecture that the first spike was caused by a project creator's family and friends who backed the project [6], and the second spike was caused by other users who noticed the project and heard of a trend of the project.

The other interesting observation is that after 60th state, the number of backers and the number of pledged money have been exponentially increased. Especially, people rushed investing a project, as a project was heading to the end of the project duration. The phenomenon is called the Deadline effect [21, 24]. Even amount of invested money has been increased more quickly than the number of backers. This may indicate that people tend to purchase more expensive reward item. They may want to make sure a project become successful, achieving higher amount of pledged money than a project goal[4]. In another case, they knew that other people already supported the project with a large amount of money which motivated them to back the project with high trust.

4. FEATURES AND EXPERIMENTAL SETTINGS

In the previous section, we analyzed behaviors of Kickstarter users and characteristics of projects. Based on the analysis, in this section we propose features which will be

[4]Kickstarter has an All-or-Nothing policy. If a project reaches at or over its goal, its creator will receive pledged fund. Otherwise, the project creator will receive nothing.

used to develop a project success predictor and an expected funding range predictor. We also describe our experimental settings which are used in Sections 5 and 6.

4.1 Features

We extracted 49 features from our collected datasets presented in Table 1. Then, we grouped the features to 4 types: (i) project features; (ii) user features; (iii) temporal features; and (iv) Twitter features.

4.1.1 Project Features

From a project page, we generated 11 features as follows:

- Project category, duration, project goal, number of images, number of videos, number of FAQs, and number of rewards.
- SMOG grade of reward description: To estimate the readability of the all rewards text.
- SMOG grade of main page description: To estimate the readability of the main page description of a project.
- Number of sentences in reward description.
- Number of sentences in the main description of a project.

The SMOG grade estimates the years of education needed to understand a piece of writing [17]. The higher SMOG grade indicates that project and reward descriptions were written well. To measure SMOG grade, we used the following formula:

$$1.043\sqrt{|polysyllables| \times \frac{30}{|sentences|}} + 3.1291$$

, where the number of Polysyllables is the count of the words of 3 or more syllables.

4.1.2 User Features

From a user profile page and the user's previous experience, we generated 28 features as follows:

- Distribution of the backed projects under the 15 main categories (15 features): what percent of projects belongs to each main category.
- Number of backed projects, number of created projects in the past, number of comments that a user made in the past, number of websites linked in a user profile, and number of Facebook friends that a user has.
- Is each of Facebook, YouTube and Twitter user pages connected? (3 features)
- SMOG grade of bio description, and Number of sentences in a bio description.
- Interval (days) between a user's Kickstarter joined date and a project's launched date.
- Success rate of the backed projects by a user.
- Success rate of the projects created by a user in the past.

4.1.3 Temporal Features

As we mentioned in Section 2, we collected 74,053 projects' temporal data consisting of daily pledged money and number of daily increased backers. First, we converted these temporal data points (i.e., daily value) to cumulated data points. For example, if a project's daily pledged money for 5 days project duration are 100, 200, 200, 100 and 200, cumulated data point in each day will be 100, 300, 500, 600 and 800. Since each project has various duration, we converted a duration to 100 states (time slots). Then, we normalized cumulated data points by 100 states. Finally, we generated two time-series features:

- Cumulated pledged money over time.
- Cumulated number of backers over time.

4.1.4 Twitter Features

As we mentioned in Section 2, 17,908 users linked their Twitter home pages to their Kickstarter user pages. From our collected Twitter dataset, we generated 8 features as follows:

- Number of tweets, Number of followings, Number of followers and Number of favorites.
- Number of lists that a user has been joined in.
- Number of tweets posted during active project days (e.g., between Jan 1, 2014 and Jan 30, 2014).
- Number of tweets containing word "Kickstarter" posted during active project days.
- SMOG grade of aggregated tweets which are posted during active project days.

The first five features were used for any project created by a user. The rest three features were generated for each project since each project was active in different time period.

Finally, we generated 49 features from a project and a user who created the project.

4.2 Experimental Settings

We describe our experimental settings which are used in the following sections for predicting project success and expected pledged money range.

Datasets	|Projects|	|Features|
KS Static	151,608	39
KS Static + Twitter	21,028	47
KS Static + Temporal + Twitter	11,675	49

Table 5: Three datasets which were used in experiments.

Datasets. In the following sections, we used three datasets presented in Table 5. Each dataset consists of a different number of projects and corresponding user profiles as we described in Section 2. Two datasets (KS Static + Twitter, and KS Static + Temporal + Twitter) contained Twitter user profiles as well.

We extracted 39 features from KS Static dataset (i.e., project features and user features), 47 features from KS Static + Twitter dataset (i.e., project features, user features and Twitter features), and 49 features from KS Static + Temporal + Twitter (i.e., all four feature groups). Note that in this subsection we presented the total number of our proposed features before applying feature selection.

Predictive Models. Since each classification algorithm might perform differently in our dataset, we selected 3 well-known classification algorithms: Naive Bayes, Random Forest, AdaboostM1 (with Random Forest as the base learner). We used Weka implementation of these algorithms [11].

Feature Selection. To check whether the proposed features were positively contributing to build a good predictor, we measured χ^2 value [23] for each of the features. The larger the χ^2 value is, the higher discriminative power the corresponding feature has. The feature selection results are described in following sections.

Evaluation. We used Accuracy as the primary evaluation metrics and Area under the ROC Curve (AUC) as the secondary metrics, and then built and evaluated each predictive model (classifier) by using 5-fold cross-validation.

5. PREDICTING PROJECT SUCCESS

Based on the features and experimental settings, we now develop and evaluate project success predictors.

5.1 Feature Selection

First of all, we conducted χ^2 feature selection to check whether the proposed features were all significant features. Since we had three datasets, we applied feature selection for each dataset. All features in KS Static dataset had positive distinguishing power to determine whether a project will be successful or not. But, in both of KS Static + Twitter dataset and KS Static + Temporal + Twitter, "Is each of Facebook, YouTube and Twitter user pages connected" features were not positively contributing, so we excluded them. Overall, some of project features (e.g., category, goal and number of rewards), some of user features (e.g., number of backed projects, success rate of backed projects, number of comments), some of Twitter features (e.g. number of lists, number of followers and number of favorites), and all temporal features were the most significant features.

5.2 Experiments

Our experimental goal is to develop and evaluate project success predictors. We build project success predictors by using each of the three datasets and evaluate performance of the predictors.

Classifier	Accuracy	AUC
Naive Bayes	67.3%	0.750
Random Forest	75.2%	0.827
AdaboostM1	**76.4%**	**0.838**

Table 6: Experimental results of three project success predictors based on Kickstarter static features.

Using KS Static dataset. The first task was to test whether only using Kickstarter static features (i.e., project and user features) would achieve good prediction results. To conduct this task, we converted Kickstarter static dataset consisting of 151,608 project profiles and user profiles to feature values. Then, We developed project success predictors based on each of 3 classification algorithms – Naive Bayes, Random Forest and AdaboostM1. Finally, we evaluated each predictor by using 5-fold cross-validation. Table 6 shows experimental results of three project success predictors based on Kickstarter static features. AdaboostM1 outperformed the other predictors, achieving 76.4% accuracy and 0.838 AUC. This result was better than 54% accuracy of a baseline which was measured by a percent of the majority class instances in Kickstarter static dataset (54% projects were unsuccessful). This result was also better than the previous work in which 68% accuracy was achieved [10].

Using KS Static + Twitter dataset. What if we add Twitter features to Kickstarter static features? Can we even improve performance of project success predictors? To answer these questions, we compared performance of predictors without Twitter features with performance of predictors with Twitter features. In this experiment, we extracted Kickstarter static features from 21,028 projects and corresponding user profiles, and Twitter features from corresponding Twitter user profiles. As you can see in Table 7, AdaboostM1 classifier with Twitter features achieved 75.7% accuracy and 0.826 AUC, increasing accuracy and AUC of AdaboostM1 classifier without Twitter features by 2.5% ($= \frac{75.7}{73.9} - 1$) and 3.5% ($= \frac{0.826}{0.798} - 1$), respectively.

Classifier	Accuracy	AUC
Kickstarter		
Naive Bayes	60.3%	0.722
Random Forest	72.8%	0.790
AdaboostM1	**73.9%**	**0.798**
Kickstarter + Twitter		
Naive Bayes	56.5%	0.724
Random Forest	73.4%	0.800
AdaboostM1	**75.7%**	**0.826**

Table 7: Project success predictors based on Kickstarter static features vs. based on Kickstarter static features and Twitter features.

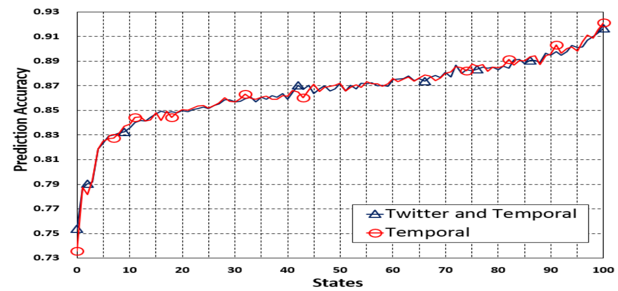

Figure 12: Project success prediction rate of predictors based on Kickstarter static and temporal features with/without Twitter features.

Using KS Static + Temporal + Twitter dataset. What if we replace Twitter features with Kickstarter temporal features? Or what if we use all features including Kickstarter static, temporal and Twitter features? Would using all features give us the best result? To answer these questions, we used KS Static + Temporal + Twitter dataset consisting of 11,675 project profiles, corresponding user profiles, Twitter profiles and project temporal data. Since each project has a different project duration, we converted each project duration to 100 states (time slots). Then we calculated temporal feature values in each state. Finally, we developed 100 predictors based on KS Static + Temporal features and 100 predictors based on KS Static + Temporal + Twitter features (each predictor was developed in each state). Note that in the previous experiments AdaboostM1 consistently outperformed the other classification algorithms, so used AdaboostM1 for this experiment. Figure 12 shows two project success predictors' accuracy in each state. In the beginning, KS Static + Temporal + Twitter features based predictors were slightly better than KS Static + Temporal features based predictors, but both of approaches performed similarly after 3rd state because temporal features became more significant. Overall, accuracy of predictors has been sharply increased until 11th state and then consistently increased until the end of a project duration. In 10th state (i.e., in the first 10% duration), the predictors achieved 83.6% accuracy which was increased by 11% ($=\frac{83.6}{75.3}-1$) compared with 75.3% accuracy when a state was 0 (i.e., without temporal features). The more a state value increased, the higher accuracy a predictor achieved.

In summary, we developed project success predictors with various feature combinations. A project success predictor based on Kickstarter static features achieved 76.4% accuracy. Adding social media features increased the prediction accuracy by 2.5%. Adding temporal features consistently increased the accuracy. The experimental results confirmed that it is possible to predict a project's success when a user creates a project, and we can increase a prediction accuracy further with early observation after launching the project.

6. PREDICTING AN EXPECTED PLEDGED MONEY RANGE OF A PROJECT

So far we have studied predicting whether a project will be successful or not. But a project's success depends on a project goal and pledged money. If pledged money is equal to or greater than a project goal, the project will be successful. On the other hand, even though a project received a lot of pledged money (e.g., $99,999) , if a project goal (e.g., $100,000) is slightly larger than the pledged money, the project will be failed. Remember the All-or-Nothing policy. If we predict how much a project will get invested in advance, we can set up a realistic project goal and make the project successful. A fundamental research problem is "Can we predict expected pledged money? or Can we predict a range of expected pledged money of a project?". To our knowledge, no one has studied this research problem yet. In this section, we propose an approach to predict a range of expected pledged money of a project.

6.1 Approach and Feature Selection

In this section, our research goal is to develop predictive models which can predict a range of pledged money of a project. We defined the number of classes (categories) in two scenarios: (i) 2 classes; and (ii) 3 classes. In a scenario of 2 classes, we used a threshold, $5,000. The first class is $\leq \$5,000$, and the second class is $> \$5,000$. In other words, if pledged money of a project is less than or equal to $5,000, the project will belong to the first class. Likewise, in a scenario of 3 classes, we used two thresholds, $100 and $10,000. The first class is $\leq \$100$, the second class is $\$100 < project \leq \$10,000$ and the third class is $> \$10,000$. Now we have the ground truth in each scenario.

Next, we applied feature selection to our datasets. In 2 classes, "Is Youtube connected" feature was not a significant feature in KS Static and KS Static + Temporal + Twitter datasets. "Is Twitter connected" feature was not a significant feature in KS Static + Twitter and KS Static + Temporal + Twitter datasets. In 3 classes, "Is Twitter connected" feature was not a significant feature in KS Static + Twitter and KS Static + Temporal + Twitter datasets.

6.2 Experiments

We conducted experiments in two scenarios – prediction in (i) 2 classes and (ii) 3 classes.

Using KS Static dataset. The first experiment was to predict a project's pledged money range by using KS Static dataset (i.e., generating the static features – project features and user features). A use case is that when a user creates a project, this predictor helps the user to set up an appropriate goal. We conducted 5 fold cross-validation in each of the two scenarios. Table 8 shows experimental results in 2 classes. AdaboostM1 outperformed Naive Bayes and Random Forest, achieving 86.5% accuracy and 0.901 AUC. When we compared our predictor's performance with the baseline – 74.8% accuracy (percent of the majority class, assuming selecting the majority class as a prediction result) –, our approach increased 11.5% ($=\frac{86.5}{74.8}-1$).

Classifier	Accuracy	AUC
Naive Bayes	75.9%	0.780
Random Forest	85.6%	**0.906**
AdaboostM1	**86.5%**	0.901

Table 8: Experimental results of pledged money range predictors based on Kickstarter static features under 2 classes.

Classifier	Accuracy	AUC
Naive Bayes	49.4%	0.713
Random Forest	73.3%	**0.817**
AdaboostM1	**74.2%**	0.811

Table 9: Experimental results of pledged money range predictors based on Kickstarter static features under 3 classes.

We also ran another experiment in 3 classes. Table 9 shows experimental results. Again, AdaboostM1 outperformed the other classification algorithms, achieving 74.2% accuracy and 0.811 AUC. When we compared its performance with the baseline – 63.1% –, it increased 17.6% ($=\frac{74.2}{63.1}-1$). Regardless of the number of classes, our proposed approach consistently outperformed than the baseline. The experimental results showed that it is possible to predict an expected pledged money range in advance.

Classifier	Accuracy	AUC
Kickstarter		
Naive Bayes	70.6%	0.759
Random Forest	81.4%	0.889
AdaboostM1	**82.5%**	**0.896**
Kickstarter + Twitter		
Naive Bayes	70.7%	0.763
Random Forest	83.1%	0.904
AdaboostM1	**84.2%**	**0.910**

Table 10: Experimental results of pledged money range predictors based on Kickstarter static features and Twitter features under 2 classes.

Using KS Static + Twitter dataset. What if we add Twitter features? Will these improve a prediction accuracy? To answer this research question, we used KS Static + Twitter dataset in each of 2 classes and 3 classes. Experimental results under 2 classes and 3 classes are shown in Tables 10 and 11, respectively. In case of 2 classes, AdaboostM1 with Twitter features increased 2.1% ($=\frac{84.2}{82.5}-1$) compared with a predictor without Twitter features, achieving 84.2% accuracy and 0.91 AUC. In case of 3 classes, AdaboostM1 with Twitter features also increased 1.8% ($=\frac{77.2}{75.8}-1$) compared with a predictor without Twitter features, achieving 77.2% accuracy and 0.843 AUC. The experimental results confirmed that adding Twitter features improved prediction performance.

Using KS Static + Temporal + Twitter dataset. What if we add temporal features? Can we find a sweet spot

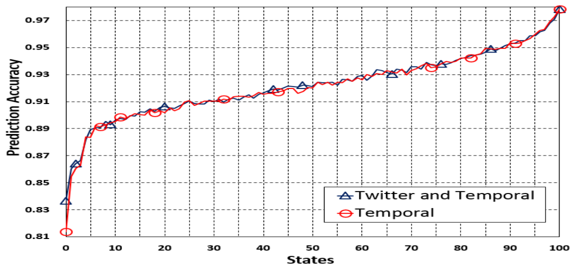

(a) Under 2 classes

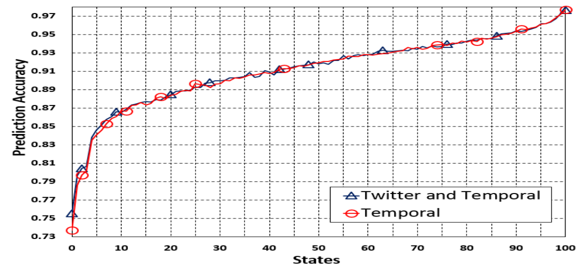

(b) Under 3 classes

Figure 13: Pledged money range prediction rate of predictors based on Kickstarter static and temporal features with/without Twitter features under 2 and 3 classes.

Classifier	Accuracy	AUC
Kickstarter		
Naive Bayes	48.6%	0.677
Random Forest	74.2%	0.829
AdaboostM1	**75.8%**	**0.830**
Kickstarter + Twitter		
Naive Bayes	48.8%	0.668
Random Forest	75.4%	0.841
AdaboostM1	**77.2%**	**0.843**

Table 11: Experimental results of pledged money range predictors based on Kickstarter static features and Twitter features under 3 classes.

where we can reach to a high accuracy in a short period? To answer these questions, we used KS Static + Temporal + Twitter dataset. Again, each project duration was converted to 100 states (time slots). Figure 13 shows how accuracy of predictors has been changed over time under 2 classes and 3 classes. Prediction accuracy of AdaboostM1 classifiers with all features (project features + user features + temporal features + Twitter features) has been sharply increased until 5th state in 2 classes and 10th state in 3 classes. The classifiers reached to 90% accuracy in 15th state under 2 classes, and in 31st state under 3 classes.

What if we do not use Twitter features? In both 2 and 3 classes, adding Twitter features slightly increased prediction accuracy until 3rd state in 2 classes, and 9th state in 3 classes compared with predictors without Twitter features.

In summary, our proposed predictive models predicted a project's expected pledged money range with a high accuracy in 2 classes and 3 classes. Adding Twitter and Kickstarter temporal features increased a prediction accuracy even higher than only using Kickstarter static features. Our experimental results confirmed that predicting a project's expected pledged money in advance is possible.

7. DISCUSSION

In previous section, we described our proposed approaches with a list of feature, and showed experimental results. In this section, we discuss other features that we tried to use but finally excluded because of degrading performance of our predictive models.

7.1 N-gram Features

In the literature, researcher have generated and used n-gram features from texts such as web pages, blogs and short text messages toward building models in various domains like text categorization [1], machine translation [16] and social spam detection [14].

We extracted unigram, bigram and trigram features from Kickstarter project descriptions after lowercasing the project descriptions, and removing stop words. Then, we conducted χ^2 feature selection so that we could only keep n-gram features which have positive power distinguishing between successful projects and failed projects. Finally, we added 22,422 n-gram features to our original feature set (i.e., project features, user features, temporal features and Twitter features) described in Section 4. Then, we built and tested project success predictors. Unfortunately, adding n-gram features deteriorated performance of project success predictors compared with only using the original feature set described in Section 4. The experimental results were the opposite of our expectation because other researchers [18] reported that using n-gram features improved their prediction rate in their own Kickstarter dataset. We conjecture that the researchers used smaller dataset which might give them some improvements. But, given the larger dataset containing all Kickstarter projects, using n-gram features decreased a prediction rate.

7.2 LIWC Features

We were also interested in using the Linguistic Inquiry and Word Count (LIWC) dictionary, which is a standard approach for mapping text to psychologically-meaningful categories [20], to generate linguistic features from a Kickstarter project main description, reward description and project creator's bio description. LIWC-2001 defines 68 different categories, each of which contains several dozens to hundreds of words. Given a project's descriptions, we measured linguistic characteristics in the 68 categories by computing a score of each category based on LIWC dictionary. First we counted the total number of words in the project description (N). Next we counted the number of words in the description overlapped with the words in each category i on LIWC dictionary (C_i). Then, we computed a score of a category i as C_i/N. Finally, we added 68 features to the original features described in Section 4. Then we built project success predictors and evaluated their performance. Unfortunately, the predictors based on 68 linguistic features and the original features were worse than predictors based on only the original features.

8. RELATED WORK

In this section we summarize crowdfunding research work in three categories: (i) analysis of crowdfunding platforms; (ii) analysis of crowdfunding activities and backers on social media sites; and (iii) project success prediction.

Researchers have analyzed crowdfunding platforms [5, 8, 9, 12]. For example, Kuppuswamy and Bayus [13] examined the backer dynamics over the project funding cycle. Mollick [19] studied the dynamics of crowdfunding, and found that personal networks and underlying project quality were associated with the success of crowdfunding efforts. Xu et al. [22] analyzed the content and usage patterns of a large corpus of project updates on Kickstarter.

In another research direction, researchers have studied social media activities during running project campaigns on crowdfunding platforms. Lu et al. [15] studied how fundraising activities and promotional activities on social media simultaneously evolve over time, and how the promotion campaigns influence the final outcomes. Rakesh et al. [3] used a promoter network on Twitter to show the success of projects depended on the connectivity between the promoters. They developed backer recommender which recommends a set of backers to Kickstarter projects.

Predicting the success of a project is one of important research problems, so researchers have studied how to predict whether a project will be successful or not. Greenberg et al. [10] collected 13,000 project pages on Kickstarter and extracted 13 features from each project page. They developed classifiers to predict project success. Their approach achieved 68% accuracy. Etter et al. [7] extracted pledged money based time series features, and project and backer graph features from 16,000 Kickstarter projects. Then, they measured how prediction rate has been changed over time. Mitra et al. [18] focused on text features of project pages. They extracted phrases and some meta features from 45,810 project pages, and then showed that using phrases features reduced prediction error rates.

Compared with the previous research work, we collected the largest datasets consisting of all Kickstarter project pages, corresponding user pages, each project's temporal data and each user's social media profiles, and conducted comprehensive analysis of users and projects. Then, we proposed and extracted comprehensive feature sets (e.g., project features, user features, temporal features and Twitter features) toward building project success predictors and pledged money range predictors. To our knowledge, we are the first to study how to predict a range of expected pledged money of a project. Since the success of a project depends on a project goal and the amount of actually pledged money, studying the prediction is very important. This research will complement the existing research work.

9. CONCLUSION

In this paper we have presented comprehensive analysis of users and projects in Kickstarter. We found that 46.4% users were all-time creators and 53.6% users were active users who not only created their own projects but also backed other projects. We also found that project success rate in each month has been decreasing as new users jointed Kickstarter and launched projects without enough preparation and experience. When we analyzed temporal data of our collected projects, we noticed that there were two peaks in the beginning of a project duration and there was the deadline effect, rushing to invest the project as the project was heading to the end of its duration. Then, we proposed 4 types of features toward building predictive models to predict whether a project will be successful and a range of pledged money. We developed the predictive models based on various feature sets. Our experimental results have showed that project success predictors based on only static features achieved 76.4% accuracy and 0.838 AUC, by adding Twitter features, increased accuracy and AUC by 2.5% and 3.5%, respectively. Adding temporal features consistently increased the accuracy. Our pledged money range predictors based on the static features have achieved up to 86.5% accuracy and 0.901 AUC. Adding Twitter and temporal features increased performance of the predictors further.

10. REFERENCES

[1] *N-Gram-Based Text Categorization*, 1994.

[2] Jumpstart our business startups act. http://www.gpo.gov/fdsys/pkg/BILLS-112hr3606enr/pdf/BILLS-112hr3606enr.pdf, 2012.

[3] *What Motivates People to Invest in Crowdfunding Projects? Recommendation using Heterogeneous Traits in Kickstarter*, 2015.

[4] Alexa. kickstarter.com site overview - alexa. http://www.alexa.com/siteinfo/kickstarter.com, March 2015.

[5] P. Belleflamme, T. Lambert, and A. Schwienbacher. Crowdfunding: Tapping the Right Crowd. *SSRN Electronic Journal*, 2012.

[6] Economist. The new thundering herd. http://www.economist.com/node/21556973, 2012.

[7] V. Etter, M. Grossglauser, and P. Thiran. Launch hard or go home!: Predicting the success of kickstarter campaigns. In *COSN*, 2013.

[8] E. M. Gerber and J. Hui. Crowdfunding: Motivations and deterrents for participation. *ACM Trans. Comput.-Hum. Interact.*, 20(6), Dec. 2013.

[9] E. M. Gerber, J. S. Hui, and P.-Y. Kuo. Crowdfunding: Why people are motivated to post and fund projects on crowdfunding platforms. In *CSCW*, 2012.

[10] M. D. Greenberg, B. Pardo, K. Hariharan, and E. Gerber. Crowdfunding support tools: Predicting success & failure. In *CHI Extended Abstracts*, 2013.

[11] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: An update. *SIGKDD Explor. Newsl.*, 11(1):10–18, Nov. 2009.

[12] J. S. Hui, M. D. Greenberg, and E. M. Gerber. Understanding the role of community in crowdfunding work. In *CSCW*, 2014.

[13] V. Kuppuswamy and B. L. Bayus. Crowdfunding Creative Ideas: The Dynamics of Project Backers in Kickstarter. *Social Science Research Network Working Paper Series*, Mar. 2013.

[14] K. Lee, J. Caverlee, and S. Webb. Uncovering social spammers: Social honeypots + machine learning. In *SIGIR*, 2010.

[15] C.-T. Lu, S. Xie, X. Kong, and P. S. Yu. Inferring the impacts of social media on crowdfunding. In *WSDM*, 2014.

[16] J. B. Mariòo, R. E. Banchs, J. M. Crego, A. de Gispert, P. Lambert, J. A. R. Fonollosa, and M. R. Costa-jussà. N-gram-based machine translation. *Comput. Linguist.*, 32(4):527–549, Dec. 2006.

[17] H. G. McLaughlin. SMOG grading - a new readability formula. *Journal of Reading*, pages 639–646, May 1969.

[18] T. Mitra and E. Gilbert. The language that gets people to give: Phrases that predict success on kickstarter. In *CSCW*, 2014.

[19] E. Mollick. The dynamics of crowdfunding: An exploratory study. *Journal of Business Venturing*, 29(1):1–16, 2014.

[20] J. Pennebaker, M. Francis, and R. Booth. *Linguistic Inquiry and Word Count*. Erlbaum Publishers, 2001.

[21] A. E. Roth, J. K. Murnighan, and F. Schoumaker. The deadline effect in bargaining: Some experimental evidence. *The American Economic Review*, pages 806–823, 1988.

[22] A. Xu, X. Yang, H. Rao, W.-T. Fu, S.-W. Huang, and B. P. Bailey. Show me the money!: An analysis of project updates during crowdfunding campaigns. In *CHI*, pages 591–600. ACM, 2014.

[23] Y. Yang and J. O. Pedersen. A comparative study on feature selection in text categorization. In *ICML*, 1997.

[24] M. Yildiz. Optimism, deadline effect, and stochastic deadlines. 2004.

[25] N. Zipkin. The 10 most funded kickstarter campaigns ever. http://www.entrepreneur.com/article/235313, March 2015.