Essays in Applied Microeconomics

by

Evan N. Wright

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2017

Doctoral Committee:

Associate Professor Ying Fan, Chair
Professor Daniel A. Ackerberg
Assistant Professor Andreas Hagemann
Assistant Professor A. Yeşim Orhun

Evan N. Wright
enwright@umich.edu
ORCID iD: 0000-0002-3660-1647

ACKNOWLEDGMENTS

I thank my committee, Ying Fan, Daniel A. Ackerberg, Andreas Hagemann, and A. Yeşim Orhun, for their guidance and support. Without their critical review this dissertation would be unfit for a napkin. I must further note my appreciation of the many professors at the University of Michigan whose classes I have had the pleasure of attending. Finally, I thank the Rackham Graduate School for its generous fellowship which has provided the luxury of several years of unencumbered research.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

SCM Synthetic Control Method

MSPE mean squared prediction error

ABSTRACT

Crowdfunding is a novel industry which facilitates the matching of consumers who want new interesting products with entrepreneurs who need capital to produce those products. Consumers may pledge to a crowdfunding project, essentially preordering the product, which allows entrepreneurs to obtain the minimum amount of capital necessary for production while simultaneously measuring demand for the product. The bulk of this dissertation explores how consumers behave when thinking about and pledging to crowdfunding projects.

Chapter 1 illustrates several notable empirical results with novel data from Kickstarter and Reddit. First, the number of new backers (consumers who pledge) per day accelerates up to the point when a project hits its goal then drops down—replicating previous findings in the literature with a different sample and more relaxed empirical assumptions. Second, "advertising" (Reddit posts) about projects has a positive effect on the number of new backers, and this effect is stronger when a project is close to or has already met its goal. These effects are shown to be both statistically and economically significant.

Chapter 2 proposes a dynamic theoretical model of consumer behavior in the context of crowdfunding. The model implies an acceleration in new backers per day as a project's probability of success nears one, offering one explanation for the acceleration identified in Chapter 2 and previous literature. In addition, the immediate effect of advertising on the number of backers is generally larger when the probability of success is higher, but there remains some ambiguity for probabilities close to one. The model relies on fairly simple assumptions about the preferences of consumers, but is only solved numerically.

Chapter 3 deviates from crowdfunding to discuss a particular econometric methodology, the synthetic control method (SCM). We highlight some data generating processes where the SCM could perform differently across units. However, we show inference based on the mean squared prediction error (MSPE) ratio is not substantially distorted. Additionally, we offer a word of warning about including all pretreatment outcomes as economic predictors in the selection of synthetic weights. Doing so could complicate inference based on the post/pretreatment MSPE ratio.

CHAPTER 1

An Empirical Analysis of Consumer Behavior and Advertising in Crowdfunding Markets

1.1 Introduction

A common problem faced by many entrepreneurs and small businesses is raising the capital necessary to launch a new product. The high risk nature of the endeavor means both debt and equity financing is expensive, if available at all. One solution has been offered by the relatively novel industry of crowdfunding. Crowdfunding firms provide a platform for entrepreneurs to advertise their project to consumers who "pledge" money to the project—called backers. If a predetermined goal is met within a time frame set by the entrepreneur, then he or she receives the full pledged amount. In return, backers are given a good after production takes place, equity in the firm, or nothing at all in the case of pure charitable projects. If the goal is not met, backers are refunded the pledged amount.[1]

While the economics of each type certainly overlap, this paper is primarily concerned with reward-based crowdfunding—where consumers are essentially pre-ordering a product and bearing the risk of production delays or failure. In the United States prior to May 2016, backers of equity-based crowdfunding projects were restricted to accredited investors by the SEC.[2] Reward-based crowdfunding is not subject to this restriction or securities-related advertising regulations. The most notable reward-based crowdfunding platforms are Kickstarter and Indiegogo. One example of an equity crowdfunding platform is Crowdfunder. Crowdfunding sites for purely charitable causes include Razoo and CrowdRise.

The budding industry of crowdfunding has continued to show strong growth in recent years. Kickstarter, the largest reward-based crowdfunding platform, has received pledges of $2.65 billion

[1] Some crowdfunding platforms do allow partial funding but are not the subject of this paper.

[2] Technically, this applies to the offering of any financial security, not just equity. There are still rules governing issuance, advertising, resale, and investment amounts. For an overview of current rules, see "Regulation Crowdfunding: A Small Entity Compliance Guide for Issuers" at https://www.sec.gov/info/smallbus/secg/rccomplianceguide-051316.htm

since its inception in 2009. $2.3 billion was pledged to over 110 thousand successful projects.[3] Given this growth and the recent relaxation of equity-based restrictions, crowdfunding will play an increasingly important role in the promotion of entrepreneurship and innovation.

Given the novelty of crowdfunding, the economics literature has been primarily exploratory in nature. Empirical results by Kuppuswamy and Bayus (2017), Mollick (2014), and others demonstrate the consumer behavior we study in this paper. In particular, Kuppuswamy and Bayus (2017) find significant concentrations of pledges at the beginning and end of a project's funding cycle, which we replicate in our empirical analysis. Concentration at the end of the cycle can simply be explained by some consumers discounting and pushing back the cost of pledging. Concentration at the beginning is likely due to friends, family, and/or fervent supporters of the project who knew about it ahead of time. Additionally, Kuppuswamy and Bayus (2017) find the number of new backers on a particular day is increasing in the cumulative progress (percentage of goal pledged) of the previous day, and they find a positive effect of Twitter mentions on the number of new backers.

Our contribution to the literature is twofold. First, we partially replicate results from Kuppuswamy and Bayus (2017) with a different dataset and relax a key methodological assumption. Second, we further explore the empirical effect of social media advertising on crowdfunding projects with data from Reddit; we interact our measure of advertising with the proportion of the goal raised to show the marginal effect of advertising varies. Advertising appears to have a stronger immediate effect when a project is close to or has exceeded its goal. However, we must rely on lags of advertising to identify these effects because we do not exogenously control the timing of advertising. Therefore, we caution against a causal interpretation.

1.2 Literature Review

Agrawal et al. (2013) provide a comprehensive overview of the economic issues and literature surrounding the crowdfunding industry including entrepreneur (dis)incentives, backer (dis)incentives, signaling effects, rules and regulations, and questions of social welfare. Agrawal et al. (2015) analyze investor behavior on Sellaband—an equity-based crowdfunding site for musicians—and find investors are more likely to back a project as the cumulative funds raised increases, but the effect is less pronounced for geographically local investors. Mollick (2014) describes some of the aspects which determine the success of reward-based crowdfunding projects including measures of quality, geography, and the social network of the creator. He uses data from Kickstarter's first three years of operation, and collected data on delays of production which plague more than 75% of successful projects. Kuppuswamy and Bayus (2017) also use Kickstarter data though from a later time period to find backers are concentrated at the beginning and end of a project's funding cycle,

[3] See https://www.kickstarter.com/help/stats for current statistics.

the number of new backers on a particular day is increasing in the project's prior success, and the number of new backers is positively affected by Twitter mentions. The empirical strategy in this paper is similar to Kuppuswamy and Bayus (2017) but allows previous values of the dependent variable (the number of new backers for a given project-day) to affect current period independent variables; whereas, Kuppuswamy and Bayus (2017) use a fixed effect Poisson approach.

In addition, Qiu (2013) uses data similar to Kuppuswamy and Bayus (2017) and includes Twitter mentions of projects—similar to our measure of "advertising" on Reddit. Qiu's (2013) empirical specification is different from ours but uses a similar dynamic panel GMM methodology. Our model has the number of new backers per day as an outcome variable, while Qiu (2013) uses the total number of backers as an outcome with a lagged outcome on the right-hand side as well. We further discuss the empirical differences between this paper and Qiu (2013) in Section 1.3.3. Li and Duan (2016) use a Bayesian hierarchical framework with a Poisson process determining the number of consumers arriving each day and a hazard function determining the probability of success, but unfortunately they have a more limited dataset than previous papers with only 577 projects. Both Qiu (2013) and Li and Duan (2016) find results similar to us: social media attention/advertising has a positive effect on the number of new backers pledging to a project.

Fundraising for charitable public goods is similar to reward-based crowdfunding in that consumers typically observe the progress of a project prior to donating, and we expect this observation to affect the probability they will donate. List and Lucking-Reiley (2002) use a field experiment to show that people are more likely to donate when they are told the project is closer to its goal. The magnitudes of the experiment were in line with typical crowdfunding projects—the goal amount was $3,000 and the average donation amounts were between $10 and $45. The data in our paper are not a randomized experiment; instead we observe a panel of projects over their lifetime.

Given the social aspects of Reddit, this paper can also be related to marketing literature which discusses word-of-mouth type advertising. For example, Liu (2006) uses online user reviews of movies to measure the effect of word-of-mouth on box office revenues. Interestingly, the valence (positive/negative sentiment) of the reviews did not have a statistically significant effect on revenues; the number of people talking about a film was far more important. In this paper, we do observe the body of comments but unfortunately do not have a decent measure of sentiment.[4]

[4]The observed score of posts (positive minus negative votes) is bounded below by zero, and posts are ranked based on score so low ranked posts are less visible. Additionally, Reddit may obfuscate the score of posts under certain circumstances to prevent voting manipulation.

1.3 Data and Empirical Strategy

1.3.1 Kickstarter

Since 2009, Kickstarter has successfully raised over $2 billion for more than 100 thousand projects. Kickstarter itself receives 5% of the successful dollars raised; unsuccessful projects pay no fee. Amazon payments also takes 3-5% of the final pledged amount for successful projects. Recall consumers are completely refunded if a project is not successful.[5]

Figure 1.1 shows an example of a Kickstarter project page. A viewer of this page does not see the full history of funding—only the current amount and days remaining. Further down the page, creators provide a full explanation of the project and rewards. Most projects have several reward tiers. For example, $45 may be the minimum amount to receive the item and $60 rewards a signed version of the item. Unfortunately, we do not observe rewards tiers—only the aggregate dollars pledged and number of backers at the daily level.

Data for 41,403 Kickstarter projects from August 2013 to November 2014 were scraped from Kickspy.com, which itself scraped from Kickstarter and is now shut down. The success rate for projects since inception is 36% but for our sample is 47%; this is probably because Kickstarter is more mature by the time our sample begins. We observe the amount raised and number of backers on a daily basis, as well as the goal, category, duration, and end date. Only projects in US dollars are considered. Missing observations result from Kickspy having difficulty scraping Kickstarter on that day for technical reasons; approximately 3.8% of project-days are missing.

In order to maintain consistent interpretation of coefficients in our regression specifications, we use only projects with funding cycles of exactly 30 days—about half our sample. Projects with missing observations, no backers, or a decrease in the number of backers at some point, are also dropped. If Kickstarter was experiencing difficulty on days with many missing observations, then there may have been fewer new backers on those days; however, we have no reason to believe features of projects caused their missingness.

The number of remaining projects is 10,798.[6] Table 1.1 provides descriptive statistics of these projects, and Figure 1.2 illustrates some histograms of various features at the project level.[7] Very few unsuccessful projects achieve more than 50% of their goal, and many successful projects receive only slightly more than their goal. This result is consistent with previous literature as noted in both Mollick (2014) and Kuppuswamy and Bayus (2017).

[5] Technically, the consumer is never charged, but the payment processor may place an authorization hold on the funds.

[6] Appendix A.1 provides comparable statistics and figures to this section with the full sample of projects.

[7] Except for "End date," the final bin includes all larger values.

Figure 1.1: Example Kickstarter project page

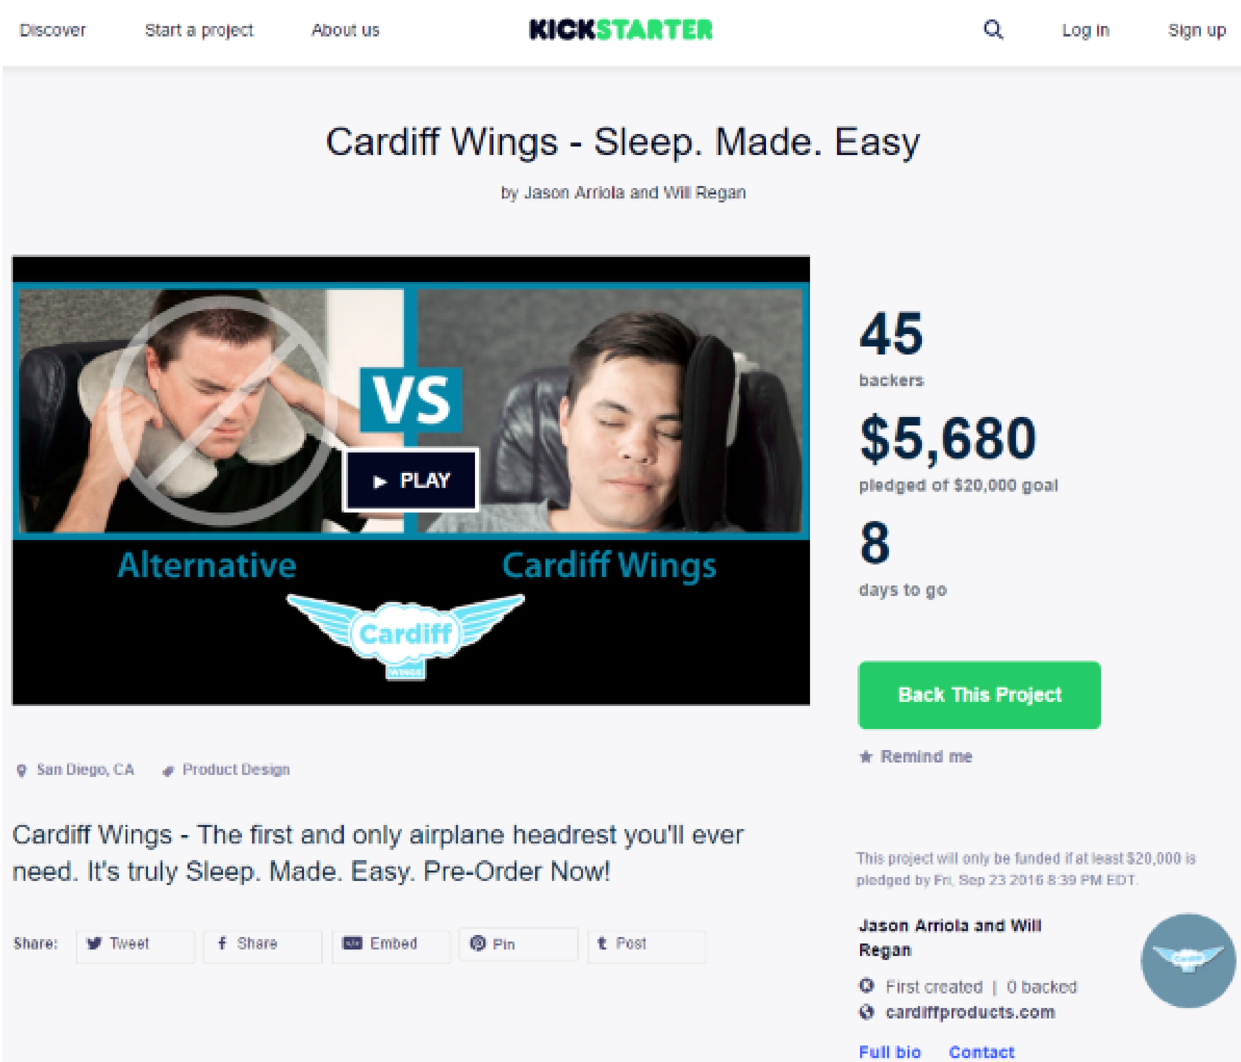

Table 1.1: Kickstarter projects descriptive statistics.

Variable	Min	Median	Max	Mean	Std. Dev.
Duration (days)	30.00	30.00	30.00	30.00	0.00
Goal ($)	1.00	6000.00	5000000.00	21210.31	106830.65
Goal if successful ($)	1.00	5000.00	1100000.00	11513.06	27044.58
Pledged ($)	1.00	1752.97	2410741.22	11524.91	56973.10
Pledged if successful ($)	1.00	6990.00	2410741.22	22204.56	80711.02
Prop of goal achieved	0.00	0.48	15804.00	4.99	216.77
Prop of goal achieved if suc	1.00	1.15	15804.00	10.39	314.76
Backers	1.00	28.00	22195.00	153.87	704.97
Backers if successful	1.00	95.00	22195.00	296.74	1000.12

1.3.2 Reddit

Reddit is a social media website where users submit links to other webpages (hereafter "posts") which are rated up or down and commented on by other users.[8] Posts are sorted by default based on score (upvotes minus downvotes) and time elapsed since posting (older posts are ranked lower). Posts are categorized by subject into "subreddits;" by default the front page combines several of the most popular subreddits. Figure 1.3 shows a snapshot of the Reddit front page. Reddit displays banner advertisements which are not the subject of this paper. Although we use the term "advertisement," we refer specifically to posts which link directly to Kickstarter projects. Users are effectively anonymous, so it is not possible to determine whether a link was submitted by a project's creator or not.

In 2013, Reddit averaged 48.5 million unique monthly visitors from the United States,[9] and it continues to rank as one of the most popular social media websites in the U.S. The data used for this paper include substantially all public Reddit posts which link to the domain kickstarter.com from August 2013 to November 2014.[10] The posts for all of Reddit were scraped during 2015 by Jason Baumgartner at pushshift.io and most were archived at that time—voting and commenting were no longer possible. The dataset was then uploaded to Google BigQuery by Felipe Hoffa, where we obtained posts linking to Kickstarter.[11] Comments for these posts were scraped by the author later in January 2016.

20.5% of Kickstarter projects have at least one Reddit post linking directly to the project page (13.5% for unsuccessful projects). Figure 1.4 shows the distribution of the number of posts per project. Most projects have very few posts, so we do not worry too much about creators spamming

[8]Users can also submit their own text, but this paper does not involve analysis of so-called self posts.

[9]https://redditblog.com/2013/12/31/top-posts-of-2013-stats-and-snoo-years-resolutions/

[10]We also include kck.st, which is Kickstater's url shortening domain.

[11]The data are still publicly available at https://bigquery.cloud.google.com/table/fh-bigquery:reddit_posts.full_corpus_201509

Figure 1.2: Histograms of project features.

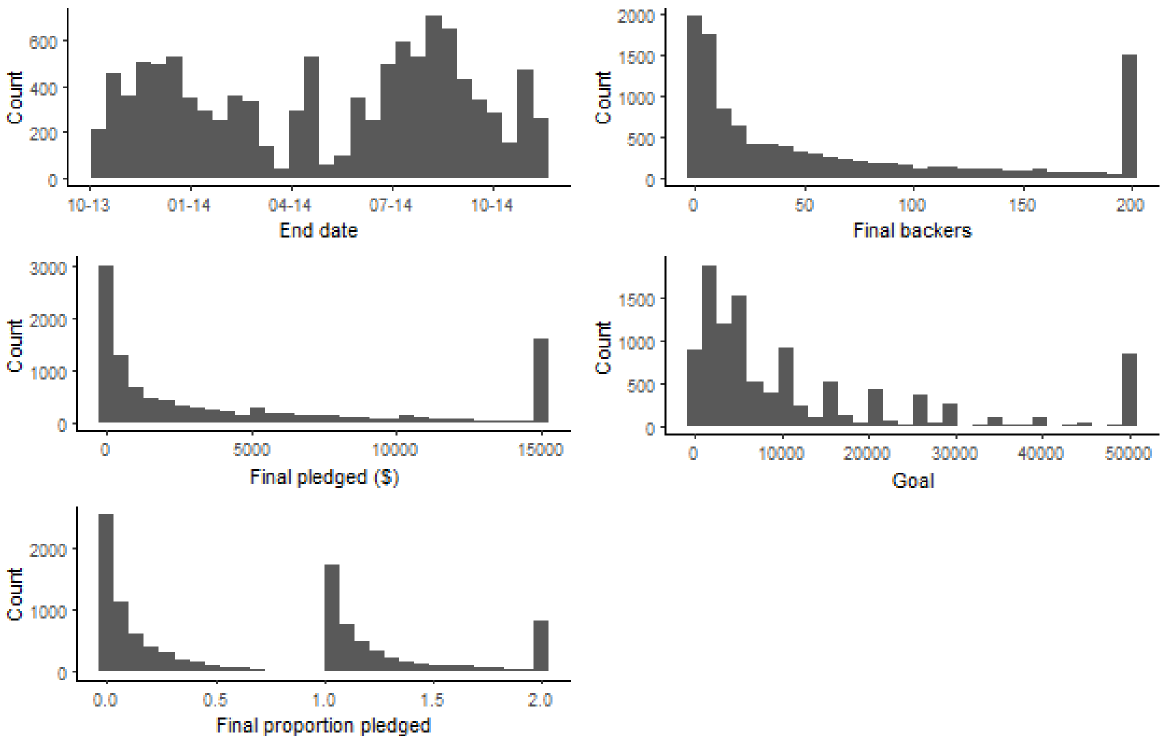

Reddit throughout the funding cycle. Many posts have no comments, but there is still enough variation to measure the effect of the number of comments in our empirical specifications. Figure 1.5 illustrates the distribution of posts across a project's funding cycle—both in time and proportion of goal raised. Most posts occur during the first week and under 50% raised.

In an ideal world, we would have randomized the timing of posts within each project to identify the effects of advertising.[12] The type of advertising considered is clearly endogenous in the sense that new backers to a project may go on Reddit to make a post soon after pledging or having decided to pledge. Therefore, an empirical specification including current period posts may be measuring this reverse effect. To mitigate this problem, we use lags of advertising to identify the effects of interest. However, there may still be other unobserved events which cause both an increase in new backers and the likelihood of a Reddit post. For example, suppose a project receives traditional media attention on day 1 which increases the number of new backers immediately and in subsequent days, 2-4. If the immediate increase in backers induces a Reddit post on day 1, we may attribute the abnormally large number of new backers on days 2, 3, and 4, to the Reddit post, even if it had no effect. Unfortunately, our only hope is these events are rare enough not to bias the results.[13]

[12]Unfortunately, there is no clear natural experiment in the data, either.

[13]Formally, we assume there are no such events. However, the empirical strategy does allow *previous period*

Figure 1.3: The Reddit front page. User-submitted posts are the subject of this paper (not banner ads).

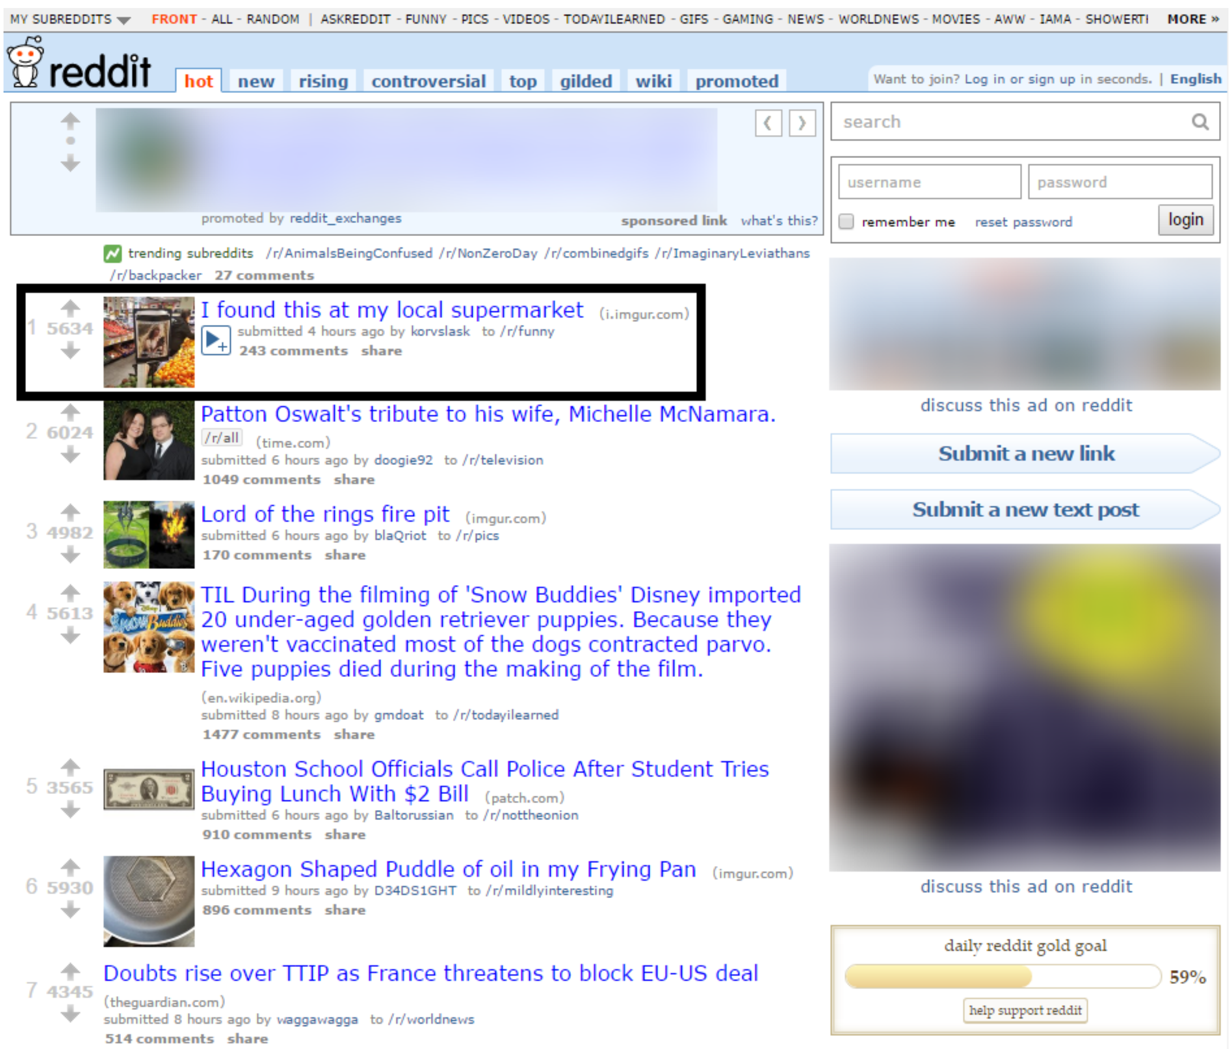

Figure 1.4: Histograms of the number of posts per project and number of comments per project-day with a new post.

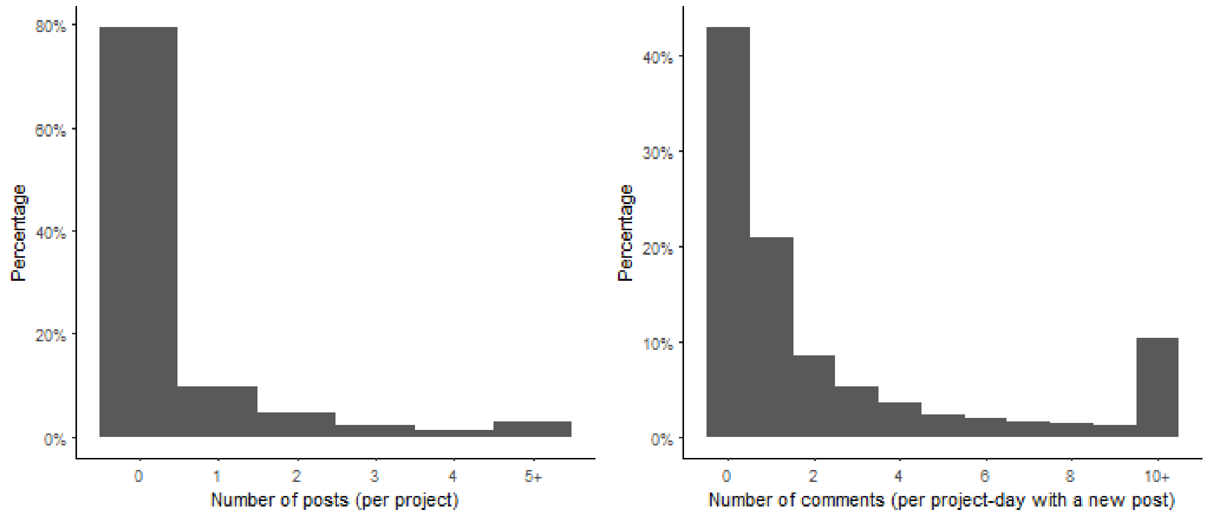

Figure 1.5: Distribution of Reddit posts across projects' funding cycles.

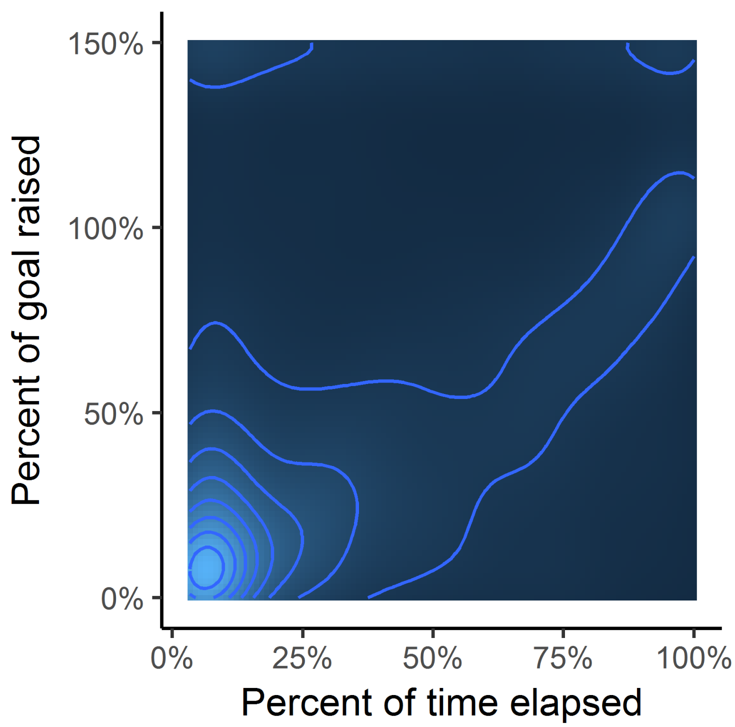

Reddit posts may also be directly correlated with other social media attention. Posts can be made by the entrepreneur themselves, so we might expect the timing of Reddit posts to be correlated with Facebook or Twitter, for example. In this case, a variable measuring Reddit activity is also capturing activity on other platforms. This is not particularly problematic, though, as our results can simply be interpreted as the effect of all social media attention to the extent it is correlated in time with Reddit posts.

1.3.3 Empirical Strategy

Our ultimate goal is to assess the effect of several variables on the number of new backers pledging to a project on a particular day: the proportion of goal raised, a dummy variable if there was a Reddit post about this project, and the interaction of the two. Two difficulties arise from this goal which prevent us from simply using OLS. First, projects have inherently different sizes in terms of the number of backers who would, or do, pledge. Consequently, an additive effect would not always be comparable across projects. For this reason, the empirical model must allow multiplicative effects and some "base" level of backers variation across projects. Therefore, we take the natural log of new backers (plus one to deal with zero backers) as our dependent variable, and include project fixed effects. Kuppuswamy and Bayus (2017) use a Poisson regression model with project fixed effects for similar reasons.

Second, each of the independent variables of interest is reasonably affected by previous period errors. That is, if there were a shock of more new backers in the past, we may expect today's proportion of the goal raised to be higher and the probability of having a Reddit post to be higher as well. However, Poisson with fixed effects assumes strict exogeneity—that previous period errors are uncorrelated with current period independent variables.[14] Therefore, we relax this assumption by using a dynamic panel methodology based on Arellano and Bond (1991).

The underlying regression equation for empirical analysis is

$$y_{it} = \beta_1 \mathbf{x}_{i,t} + \beta_2 \mathbf{z}_{i,t} + \eta_i + \varepsilon_{it}$$

where

- $y_{it} = \ln(NewBackers_{it} + 1)$ and $NewBackers_{it}$ is the net number of new backers pledging to project i on day t.
- $\mathbf{x}_{i,t}$ is a vector of predetermined variables. That is, $\mathbf{x}_{i,t}$ may be correlated with $\varepsilon_{i,1:t-1}$ but not $\varepsilon_{i,t:T}$.[15] $\mathbf{x}_{i,t}$ may include lags of predetermined variables as well, for example, the lagged

unobserved shocks to affect the likelihood of a Reddit post.

[14]See Wooldridge (1997) for an explanation.

[15]$\varepsilon_{i,1:t-1}$ is the vector $(\varepsilon_{i,1}, \ldots, \varepsilon_{i,t-1})$ and $\varepsilon_{i,t:T}$ is the vector $(\varepsilon_{i,t}, \ldots, \varepsilon_{i,T})$.

proportion of the goal raised.

- $\mathbf{z}_{i,t}$ is a vector of (possibly lagged) strictly exogenous variables. That is, $\mathbf{z}_{i,t}$ is uncorrelated with all $\varepsilon_{i,1:T}$. $\mathbf{z}_{i,t}$ will include dummy variables for the day of the week, for example.
- η_i is project-specific fixed effect.
- ε_{it} is idiosyncratic error.

The net number of new backers is generally nonnegative, but it is possible for a backer to cancel his or her pledge prior to the end of the funding cycle. Therefore, projects with at least one negative realization of new backers are dropped, as are projects with all zero outcomes.

We estimate (β_1, β_2) using generalized method of moments (GMM) with moment conditions very close to Arellano and Bond (1991) using the differenced equation,

$$\Delta y_{it} = \beta_1 \Delta \mathbf{x}_{i,t} + \beta_2 \Delta \mathbf{z}_{i,t} + \Delta \varepsilon_{it}.$$

The lagged levels of $\mathbf{x}_{i,t}$ are correlated with $\Delta \mathbf{x}_{i,t}$ but uncorrelated with the differenced errors $\Delta \varepsilon_{it}$, so

$$E[\mathbf{x}_{i,t-1} \Delta \varepsilon_{it}] = 0 \text{ for } t = 2, \dots T$$

which also holds for all lags of $\mathbf{x}_{i,t-1}$. However, $\Delta \varepsilon_{it}$ are serially correlated to some extent in our data, so we use lags of four periods or greater for moment conditions involving predetermined variables. Far lagged levels may also be weak instruments for the differenced term, so we limit the lags to five periods. Therefore, each predetermined variable contributes two moments for each available time period.

The strictly exogenous variables $\Delta \mathbf{z}_{i,t}$ instrument for themselves,

$$E[\Delta \mathbf{z}_i{}^T \Delta \varepsilon_i] = 0$$

and each contribute one moment.[16] The two-step estimator from Arellano and Bond (1991) is used with finite sample bias-corrected standard errors from Windmeijer (2005), unless otherwise stated.

The estimation methodology we use is similar to that of Qiu (2013), but the form of the regression equation is different. Instead of $\ln(NewBackers_{it} + 1)$ on the left hand side, Qiu (2013) uses the cumulative outcome $Backers_{it}$ as an outcome with lagged $Backers_{it-1}$ on the right hand side.[17] Therefore, the interpretation of coefficients is additive, whereas our model is multiplicative—more similar to Kuppuswamy and Bayus (2017). Qiu (2013) must also assume the

[16] $\Delta \mathbf{z}_i$ is the vector $(\Delta \mathbf{z}_{i,1}, \dots, \Delta \mathbf{z}_{i,T})$ and $\Delta \varepsilon_i$ is the vector $(\Delta \varepsilon_{i,1}, \dots, \Delta \varepsilon_{i,T})$.

[17] Qiu (2013) also has a specification using the cumulative pledge amount in dollars in place of backers.

autoregressive process $\{Backers_{it}\}$ is stationary conditional on other covariates, which is quite strong in this application. Once a backer pledges they rarely cancel their pledge, so that additional backer is likely maintained in the level of $Backers_{it}$ forever.

To handle the high degree of persistence in the process and thus the weak instruments problem,[18] Qiu (2013) uses the system GMM approach from Blundell and Bond (1998) which adds additional moments based on lagged differences being correlated with current period levels. Although this approach reduces bias, it still requires stationarity, and it requires an additional assumption on the initial value of $Backers_{i1}$. Specifically, deviations of $Backers_{i1}$ from the unobserved heterogeneity, η_i, must be uncorrelated with the level of η_i conditional on other covariates.[19] For example, projects with a high unobserved quality must not be more likely to have abnormally high initial backers—an assumption unlikely to hold in this context.

Qiu's (2013) specification also does not directly measure the effect of the proportion of the goal raised, which is of particular economic interest, especially given the hard cutoff of one (success or failure). Moreover, Qiu (2013) includes current period measures of Twitter activity (analogous to Reddit posts) which introduces endogeneity concerns; consumers may post on Twitter as a consequence of backing a project, rather than the other way round. We are not immune to similar endogeneity concerns, but we limit them by lagging our measures of Reddit activity.

1.4 Empirical Results

We present two branches of empirical results. The first is an analysis of the dynamic behavior discussed earlier—that a project receives more new backers the closer it is to reaching its goal. We compare these dynamics between large and small projects and provide a partial replication of Kuppuswamy and Bayus's (2017) findings. The second branch of results reveals the positive effects of advertising on the the number of new backers pledging to a project as well as the interaction between the advertising effect and the proportion of the goal raised.

We include "day in funding cycle" fixed effects, so in order to maintain a consistent interpretation across projects, we use a subset of projects which last exactly 30 days.[20] In Appendix A.2, we provide a Poisson regression similar to Kuppuswamy and Bayus (2017) with similar results, so the selection of our subset is unlikely to have significantly distorted results.

[18] If $\{Backers_{it}\}$ has a root close to unity, lagged levels will be weakly correlated with current period differences.

[19] For details, see Section 4.2 of Blundell and Bond (1998).

[20] Of the full sample, projects which last 30 days comprise nearly half of all projects.

1.4.1 Without Advertising

In the first specification, we measure the effect of the proportion of the goal raised on the number of new backers pledging to a project (per day). The day of week, day in the funding cycle, number of other active projects, and number of new backers to other projects are included as controls. Specifically, the strictly exogenous variables in $\mathbf{z}_{it}$ are:

- Dummy variables for each day 2-29 in the funding cycle. Day 1 is removed because of lags for other variables. Day 30 is left out as the reference to avoid the dummy variable trap.
- Dummy variables for Monday-Saturday. Sunday is the reference.
- The number of other ($j \neq i$) active projects in i's category, $NumOtherProjects_{it}$.

The predetermined variables in $\mathbf{x}_{it}$ are:

- Dummy variables which bin the lagged proportion of goal raised, $PropGoal_{i,t-1}$, into five bins: 0.2-0.4, 0.4-0.6, 0.6-0.8, 0.8-1, and ≥ 1. The left out reference bin is 0-0.2.
- The number of new backers pledging to other ($j \neq i$) projects in i's category, $NumOtherNewbackers_{it}$.

Table 1.2 displays the results for all projects, projects with a goal smaller than \$10,000, and projects with a goal of at least \$10,000. Recall the left hand side of the regression equation is $\ln(NewBackers_{it} + 1)$, so the exponentiated coefficient ($\exp(\hat{\beta})$) for a variable is roughly the multiplicative effect of a one unit increase in that variable on $NewBackers_{it}$. For example, a project which was 60-80% funded yesterday may expect $\exp(0.356) = 1.43$ times as many new backers today as if it was 0-20% funded yesterday.[21] Consistent with previous literature, we see the largest increase in the number of new backers when a project is 80-100% towards its goal, and the effect drops off after the project has met its goal. This phenomenon exists both for small and large projects. We also see more backers pledge at the beginning and end of the funding cycle.

1.4.2 With Advertising

Now, we measure the effect of new Reddit posts on the number of new backers pledging to a project with lags of the dummy variable $RedditPosts_{it} > 0$ which is equal to one if and only if project i had at least one new Reddit post on day t.[22] The day of week, day in the funding cycle,

[21] Recall 0-0.2 is the reference bin for the dummy variable $0.6 \leq PropGoal_{i,t-1} < 0.8$.

[22] Reddit posts may be visible on the first few pages of a subreddit for several days, but this variable only measures posts actually written on day t. Reddit's default ranking algorithm ("Hot") is such that new posts are given considerably more weight. This is apparent in Figure 1.3 where none of the top seven posts are older than 10 hours.

Table 1.2: Arellano-Bond estimation of $\ln(\text{new backers } +1)$ on lagged proportion of goal raised and controls.

	(1)		(2)		(3)	
	All projects		Goal < \$10k		Goal $\geq$ \$10k	
Variables	coef	se	coef	se	coef	se
$t = 2$	-0.536***	(0.0119)	-0.501***	(0.0132)	-0.526***	(0.0200)
$t = 3$	-0.707***	(0.0145)	-0.649***	(0.0157)	-0.714***	(0.0247)
$t = 4$	-0.797***	(0.0159)	-0.724***	(0.0168)	-0.826***	(0.0280)
$\vdots$	$\vdots$	$\vdots$				
$t = 27$	-0.589***	(0.0109)	-0.555***	(0.0124)	-0.626***	(0.0198)
$t = 28$	-0.520***	(0.00923)	-0.493***	(0.0108)	-0.561***	(0.0168)
$t = 29$	-0.350***	(0.00783)	-0.351***	(0.00974)	-0.356***	(0.0134)
Monday	0.0653***	(0.00557)	0.0560***	(0.00623)	0.0954***	(0.00879)
Tuesday	0.0457***	(0.00726)	0.0437***	(0.00745)	0.0714***	(0.0116)
Wednesday	0.0488***	(0.00626)	0.0461***	(0.00670)	0.0710***	(0.0105)
Thursday	0.0294***	(0.00625)	0.0307***	(0.00652)	0.0481***	(0.0104)
Friday	0.0131***	(0.00407)	0.00796*	(0.00481)	0.0301***	(0.00727)
Saturday	-0.0384***	(0.00398)	-0.0355***	(0.00473)	-0.0467***	(0.00665)
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.461***	(0.0717)	0.259***	(0.0650)	0.548***	(0.156)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.349***	(0.0674)	0.175***	(0.0605)	0.435***	(0.152)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.356***	(0.0656)	0.174***	(0.0600)	0.518***	(0.147)
$0.8 \leq PropGoal_{i,t-1} < 1$	0.828***	(0.0650)	0.609***	(0.0620)	1.091***	(0.141)
$1 \leq PropGoal_{i,t-1}$	-0.0130	(0.0542)	-0.190***	(0.0474)	0.238*	(0.130)
$NumOtherProjects_{it}$	-0.0116***	(0.00142)	-0.00782***	(0.00153)	-0.0131***	(0.00193)
$NumOtherNewbackers_{it}$	0.000525***	(6.36e-05)	0.000386***	(7.90e-05)	0.000500***	(7.27e-05)
Observations	313,142		188,935		124,207	
Number of projects	10,798		6,515		4,283	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

number of other active projects, number of new backers to other projects, and the proportion of the goal raised are also included as controls. The lagged post indicators are treated as predetermined variables in $\mathbf{x}_{it}$.

Table 1.3 displays the results. We do see a positive effect of Reddit posts which is statistically and economically significant. Economic intuition suggests the positive effect of posts should be monotonically diminishing in the few days following a new post, but our estimates are not precise enough to make that determination. [23] If there was a Reddit post about project i three days ago, the expected number of new backers today is about 2.2 times what it would be without a post.

One may be worried the positive effect of posts occurs only for successful projects. Therefore, we split the sample into successful and unsuccessful projects, and see that Reddit posts have a significant and positive effect for both groups. The results appear in Table 1.4. Unfortunately, the estimates are not precise enough to determine whether the marginal effects are different across successful and unsuccessful projects.

Next, we consider the effect of the potential viewership on the relative effect of posts. Reddit is separated into "subreddits" which effectively categorize posts.[24] For example there is an Art subreddit, a Gaming subreddit, and even a Kickstarter subreddit. Each subreddit has a different level of viewership which can be measured by the number of "subscribers" to that subreddit. We sum the number of subreddit subscribers who could see posts on a given project-day and split this number into three categories: under 25 thousand, between 25 and 100 thousand, and over 100 thousand.[25] We lag this variable to $t-3$ and present the results in Table 1.5. There does not appear to be a statistically significant difference in effect size across subreddit viewership.

Activity of other users in a Reddit post may also affect the number of new backers, so we count the number of new comments made on posts linking to a given project each day.[26] The results are presented in Table 1.6. Our estimates are not precise enough to reject a zero effect of number of comments. Keep in mind comments may not be positive about a Kickstarter project, so we are only measuring activity, not whether consumers think the project is good.

Next, we consider the effect of Reddit posts about *other* projects in any category. Posts about other projects on Reddit may be expected to drive traffic to Kickstarter, which could in turn lead to more backers for the project of interest. The number of Reddit posts about other projects is treated as strictly exogenous and appears in $\mathbf{z}_{it}$. The results presented in Table 1.7 suggest the effect is positive but not statistically significant.

[23]Recall day 30 is the reference, so we expect the coefficients on the time fixed effects to be negative but closer to zero near the ends of the funding cycle.

[24]By default, a user visiting Reddit.com at the time of our sample viewed the Frontpage subreddit which is actually a compilation of the top posts from many different subreddits. At the time of our sample, those were essentially the largest subreddits which were "safe for work."

[25]These variables are treated as predetermined and appear in $\mathbf{x}_{it}$.

[26]These variables are treated as predetermined and appear in $\mathbf{x}_{it}$.

Table 1.3: Arellano-Bond estimation of $\ln$(new backers $+1$) on lagged Reddit posts and controls.

	coef	se
$t = 4$	-0.148***	(0.0393)
$t = 5$	-0.202***	(0.0350)
$t = 6$	-0.236***	(0.0319)
$\vdots$	$\vdots$	$\vdots$
$t = 27$	-0.434***	(0.00928)
$t = 28$	-0.413***	(0.00847)
$t = 29$	-0.282***	(0.00775)
Monday	0.0875***	(0.00477)
Tuesday	0.0775***	(0.00580)
Wednesday	0.0738***	(0.00547)
Thursday	0.0492***	(0.00545)
Friday	0.0164***	(0.00424)
Saturday	-0.0537***	(0.00405)
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.242***	(0.0685)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.423***	(0.0816)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.666***	(0.0924)
$0.8 \leq PropGoal_{i,t-1} < 1$	1.084***	(0.0935)
$1 \leq PropGoal_{i,t-1}$	0.733***	(0.106)
$NumOtherProjects_{it}$	-0.00267**	(0.00108)
$NumOtherNewbackers_{it}$	9.20e-05**	(4.18e-05)
$RedditPosts_{t-1} > 0$	1.619***	(0.298)
$RedditPosts_{t-2} > 0$	0.449*	(0.240)
$RedditPosts_{t-3} > 0$	0.785***	(0.170)
Observations	280,748	
Number of projects	10,798	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 1.4: Arellano-Bond estimation of $\ln(\text{new backers } +1)$ on lagged Reddit posts—separated into successful and unsuccessful groups.

	(Successful) coef	se	(Unsuccessful) coef	se
$t = 4$	-0.522***	(0.0332)	-0.222***	(0.0498)
$t = 5$	-0.630***	(0.0299)	-0.238***	(0.0488)
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$t = 28$	-0.654***	(0.0136)	-0.338***	(0.00867)
$t = 29$	-0.398***	(0.0114)	-0.241***	(0.00835)
Monday	0.150***	(0.00707)	0.0361***	(0.00375)
Tuesday	0.143***	(0.00834)	0.0307***	(0.00468)
Wednesday	0.132***	(0.00804)	0.0285***	(0.00431)
Thursday	0.0954***	(0.00788)	0.0227***	(0.00443)
Friday	0.0312***	(0.00691)	0.00685**	(0.00340)
Saturday	-0.0878***	(0.00626)	-0.0169***	(0.00327)
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.250***	(0.0266)	-0.178	(0.269)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.431***	(0.0286)	-0.0836	(0.364)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.598***	(0.0341)	0.878*	(0.456)
$0.8 \leq PropGoal_{i,t-1} < 1$	1.027***	(0.0372)	0.761	(0.661)
$1 \leq PropGoal_{i,t-1}$	0.249***	(0.0422)		
$NumOtherProjects_{it}$	-0.00238*	(0.00130)	-0.000749	(0.000933)
$NumOtherNewbackers_{it}$	7.38e-05	(4.93e-05)	5.55e-05	(3.42e-05)
$RedditPosts_{t-1} > 0$	0.432**	(0.214)	0.533**	(0.249)
$RedditPosts_{t-2} > 0$	0.133	(0.180)	0.268	(0.207)
$RedditPosts_{t-3} > 0$	0.494***	(0.135)	0.299**	(0.130)
Observations	133,094		147,654	
Number of projects	5,119		5,679	

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 1.5: Arellano-Bond estimation of $\ln$(new backers $+1$) on lagged Reddit posts by subreddit size.

	coef	se
$t = 4$	-0.132***	(0.0397)
$t = 5$	-0.190***	(0.0355)
$\vdots$	$\vdots$	$\vdots$
$t = 28$	-0.417***	(0.00857)
$t = 29$	-0.286***	(0.00770)
Monday	0.0905***	(0.00480)
Tuesday	0.0843***	(0.00581)
Wednesday	0.0804***	(0.00549)
Thursday	0.0567***	(0.00530)
Friday	0.0187***	(0.00418)
Saturday	-0.0532***	(0.00408)
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.264***	(0.0673)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.450***	(0.0805)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.693***	(0.0906)
$0.8 \leq PropGoal_{i,t-1} < 1$	1.131***	(0.0917)
$1 \leq PropGoal_{i,t-1}$	0.775***	(0.103)
$NumOtherProjects_{it}$	-0.00256**	(0.00109)
$NumOtherNewbackers_{it}$	0.000109**	(4.29e-05)
$PostWithSubsUnder25k_{i,t-3}$	1.997***	(0.421)
$PostWithSubs25kto100k_{i,t-3}$	0.761*	(0.390)
$PostWithSubsOver100k_{i,t-3}$	1.704***	(0.345)
Observations	291,546	
Number of projects	10,798	

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 1.6: Arellano-Bond estimation of $\ln(\text{new backers} + 1)$ on lagged Reddit posts and comments.

	coef	se
$t = 4$	-0.112***	(0.0362)
$t = 5$	-0.173***	(0.0327)
$\vdots$	$\vdots$	$\vdots$
$t = 28$	-0.414***	(0.00842)
$t = 29$	-0.285***	(0.00763)
Monday	0.0866***	(0.00453)
Tuesday	0.0776***	(0.00549)
Wednesday	0.0736***	(0.00523)
Thursday	0.0520***	(0.00509)
Friday	0.0167***	(0.00406)
Saturday	-0.0519***	(0.00387)
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.290***	(0.0643)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.473***	(0.0769)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.719***	(0.0870)
$0.8 \leq PropGoal_{i,t-1} < 1$	1.148***	(0.0880)
$1 \leq PropGoal_{i,t-1}$	0.795***	(0.0996)
$NumOtherProjects_{it}$	-0.00260**	(0.00102)
$NumOtherNewbackers_{it}$	9.31e-05**	(3.96e-05)
$RedditPosts_{t-1} > 0$	0.860***	(0.159)
$RedditPosts_{t-2} > 0$	0.425***	(0.135)
$RedditPosts_{t-3} > 0$	0.610***	(0.108)
$RedditComms_{t-1}$	0.00558	(0.00741)
$RedditComms_{t-2}$	0.00150	(0.00455)
$RedditComms_{t-3}$	-0.000144	(0.00148)
Observations	280,748	
Number of projects	10,798	

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 1.7: Arellano-Bond estimation of $\ln(\text{new backers} + 1)$ on Reddit posts about *other* projects.

	coef	se
$t = 2$	-0.536***	(0.0119)
$t = 3$	-0.707***	(0.0145)
$t = 4$	-0.797***	(0.0159)
$t = 5$	-0.853***	(0.0169)
$\vdots$	$\vdots$	$\vdots$
$t = 28$	-0.520***	(0.00922)
$t = 29$	-0.350***	(0.00783)
Monday	0.0634***	(0.00529)
Tuesday	0.0431***	(0.00684)
Wednesday	0.0460***	(0.00602)
Thursday	0.0274***	(0.00595)
Friday	0.0116***	(0.00411)
Saturday	-0.0388***	(0.00403)
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.458***	(0.0714)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.347***	(0.0672)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.354***	(0.0654)
$0.8 \leq PropGoal_{i,t-1} < 1$	0.826***	(0.0648)
$1 \leq PropGoal_{i,t-1}$	-0.0142	(0.0539)
$NumOtherProjects_{it}$	-0.0115***	(0.00141)
$NumOtherNewbackers_{it}$	0.000520***	(6.34e-05)
$NumOtherPosts_{it}$	0.000255	(0.000195)
Observations	313,142	
Number of projects	10,798	

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

1.4.3 Advertising Interacted With Funding

Reddit posts may have a varying effect depending on the project's funding state at the time of the post. In general, we intuit more consumers will pledge immediately if the project has already met its goal or is very likely to do so. We discuss more reasoning behind this idea in Chapter 2, but one simple explanation is that consumers care about the probability a project will be successful because they do not want to waste time creating an account and/or typing in a credit card number otherwise. To assess this effect in the data, we regress the number of new backers on interactions of the lagged funding state and a dummy variable for a new post.[27]

The results are presented in Table 1.8. The interaction terms for 80-100% and $\geq 100\%$ are positive and significant. While the other interactions are not significant, they still remain positive when combined with the post dummy variable, which suggests the total effect of Reddit posts measured earlier is not only because of the large effects above 80% funding. Reddit posts appear to have a significantly larger effect on days when a project is fully funded (or close) compared to days when it is not.

1.4.4 Back-of-the-envelope Calculations

Using the coefficient estimates from Table 1.3, some back-of-the-envelope calculations reveal the magnitude of the effect of Reddit posts in dollars raised by Kickstarter projects. The following calculations are based on correlations in our sample and would require truly exogenous variation in the timing of posts (through an experiment, for example) for a causal interpretation. For projects with only one post-day, suppose that post-day is removed. Then, the number of new backers on the following day would be lower by a factor of $\exp(-1.619) \approx 0.2$ on average, using the specification in Table 1.3. If we assume the dollar pledge by each backer is the same, we may estimate the change in dollars raised without a post. If we sum the effects from the three days following a post, the median loss in dollars raised is \$354 across projects.

Those fewer dollars raised may be carried forward to possibly change the proportion of goal raised, $PropGoal_{i,t}$, in subsequent periods. If we account for this full dynamic effect, the median fewer dollars raised across projects is \$435. Removing posts resulted in 86 successful projects in this group (of 398) becoming unsuccessful.

The same process may be carried out for projects with no post-days by adding a post. If we add a post on day 3 for projects which had no posts, the median extra dollars raised is \$166. Part of this large deviation from the numbers above may be due to the different features (size, for example) of projects which had a Reddit post in our sample and those which did not. Additionally, the effect of a post we estimate is only multiplicative when in reality it may be additive (or both); multiplying

[27] All of the interactions are treated as predetermined and appear in $\mathbf{x}_{it}$.

Table 1.8: Arellano-Bond estimation of $\ln(\text{new backers} + 1)$ on lagged interaction of Reddit posts and funding state.

	coef	se
$t = 2$	0.241***	(0.0423)
$t = 3$	0.0500	(0.0369)
$t = 4$	-0.0638*	(0.0329)
$t = 5$	-0.141***	(0.0307)
$\vdots$	$\vdots$	$\vdots$
$t = 28$	-0.412***	(0.00843)
$t = 29$	-0.284***	(0.00770)
Monday	0.0922***	(0.00447)
Tuesday	0.0828***	(0.00536)
Wednesday	0.0796***	(0.00511)
Thursday	0.0594***	(0.00505)
Friday	0.0195***	(0.00399)
Saturday	-0.0516***	(0.00376)
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.285***	(0.0632)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.474***	(0.0756)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.717***	(0.0853)
$0.8 \leq PropGoal_{i,t-1} < 1$	1.143***	(0.0859)
$1 \leq PropGoal_{i,t-1}$	0.775***	(0.0981)
$NumOtherProjects_{it}$	-0.00233**	(0.000981)
$NumNewbackersOther_{it}$	7.02e-05*	(3.86e-05)
$RedditPosts_{t-1} > 0$	-0.0586	(0.421)
$(0.2 \leq PropGoal_{i,t-1} < 0.4) * (RedditPosts_{t-1} > 0)$	0.909	(0.568)
$(0.4 \leq PropGoal_{i,t-1} < 0.6) * (RedditPosts_{t-1} > 0)$	0.487	(0.574)
$(0.6 \leq PropGoal_{i,t-1} < 0.8) * (RedditPosts_{t-1} > 0)$	0.722	(0.596)
$(0.8 \leq PropGoal_{i,t-1} < 1) * (RedditPosts_{t-1} > 0)$	1.771***	(0.590)
$(1 \leq PropGoal_{i,t-1}) * (RedditPosts_{t-1} > 0)$	1.481***	(0.522)
Observations	302,344	
Number of projects	10,798	

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

an already small number (net dollars raised on the relevant days) does not result in a large absolute increase. Only 77 (of 7,653) unsuccessful projects were made successful by adding a Reddit post. This is not particularly surprising given that very few unsuccessful projects finish close to their goal—see Figure 1.2.

1.5 Discussion and Future Research

The primary contribution of this paper is an empirical analysis of backer behavior in reward-based crowdfunding markets. We found results similar to Kuppuswamy and Bayus (2017) with a separate sample: backers are concentrated at the beginning and end of the funding cycle, and backers are more likely to pledge the closer a project is to its goal—but only until the project is fully funded. We also assessed the impact of social media advertising via Reddit posts. There are more new backers pledging to a project in the days following a Reddit post about the project, and the effect appears stronger if the project has already met its goal. Prior to reaching the goal, Reddit posts have about the same (multiplicative) effect for unsuccessful and successful projects. However, potential viewership of the post and the number of comments on a post do not have a statistically significant effect on the number of new backers.

Given the novelty of the crowdfunding industry, the bulk of literature has been descriptive in nature. However, a deeper understanding of consumer behavior will be valuable as the industry grows, especially now that equity-based crowdfunding regulations have been relaxed. One way to move forward is to structurally model consumers' behavior and analyze their beliefs more carefully. In Chapter 2, we discuss a theoretical model of consumer behavior which may move us toward such a goal. Unfortunately, this also requires finer data, in particular accurate measures of traffic to crowdfunding project pages to disentangle a change in traffic from a change in the likelihood of pledging.

Given the recent relaxation of equity-based crowdfunding in the United States, future research on the effects of advertising in this context, especially the extent to which it is persuasive versus informative, will prove valuable. Regulators face trade-offs when determining exactly how and to whom entrepreneurs can advertise investment their new product or business. Promoting innovation by enabling entrepreneurs to acquire funding is important, but consumers are necessarily exposed to substantial risk—which could exceed that of traditional publicly traded securities. We hope this paper contributes to moving forward our economic understanding of the consumer behavior underlying these issues.

CHAPTER 2

A Theoretical Analysis of Consumer Behavior and Advertising in Crowdfunding Markets

2.1 Introduction

In this chapter, we propose a model of consumer behavior which implies some of the empirical results described in Chapter 1. We focus on the consumer's beliefs and decisions primarily because the data from Chapter 1 admit analysis only of consumers' decisions. Although the entrepreneur's choices of project goal, length, and advertising decisions are interesting, we consider those choices to be exogenous for our purposes. Moreover, those decisions depend upon the preferences and behavior of consumers, so this is a natural starting point.

Given the growth of the reward-based crowdfunding industry and the recent relaxation of equity crowdfunding regulation in the United States, a deeper understanding of how consumers think and behave is valuable to entrepreneurs, crowdfunding platforms, and regulators. Entrepreneurs and platforms want to know how and when consumers pledge to encourage an optimal matching of projects and backers. Regulators need to know how advertising affects consumers' decisions when evaluating policy which restricts advertising. The goal of this chapter is thus to move forward our theoretical understanding of consumer behavior in these markets.

The primary contribution of this chapter is to show that simple consumer preferences can imply the acceleration in the number of project backers as a project nears its goal. A key finding in the data, though, is that the number of new backers per day drops back down (not to zero but to a lower level) after the project meets its goal. The model presented implies this effect. If the model was purely one of higher funding signaling higher quality, we would expect the rate of new backers to continually increase even past 100% funding.

An additional contribution is the model implies the immediate effect of advertising is generally increasing in the proportion of the goal raised (really the probability of success), which is consistent with the empirical results from Chapter 1. Intuitively, consumers learn about the project from advertising but may or may not pledge immediately. A higher funding percentage increases the likelihood that they pledge immediately, which gives the result.

With better data than what we have in Chapter 1, the dynamic model could be structurally estimated. One valuable outcome of structural estimation would be a measure of the relative time cost that consumers spend to mechanically pledge versus how much they spend to return to the project page in the future (which one could interpret as a cost of attention as well). The implications of the theoretical model suggest the this difference in cost can have an impact on whether a project is funded or not. If the relative cost of returning in the future is high, then consumers are more likely to pledge immediately or leave forever, which in turn affects the amount of funding raised and the probability of success of the project.

2.2 Literature Review

Theoretical economic literature has not focused much on crowdfunding specifically, but the general concepts are quite similar to the economics of public good provision which has seen extensive work. In the crowdfunding context, the success of the project (i.e. reaching the goal) is itself a public good, while the product delivered is typically a private good.

Vesterlund (2003) proposes a theoretical model of charitable fundraising with imperfect information and differences in project quality. In her model, the charity is able to reveal initial contributions or not, and the revelation of large contributions signals a high quality project to remaining donors. She shows there exist equilibria where revealing initial contributions is optimal and improves the total funds raised for high quality charities relative to the perfect information case. In our context, all contributions are revealed and consumers have imperfect information, but consumers consider only one project at a time.

Andreoni (2006) alters Vesterlund's (2003) model by allowing more variation of project quality, and allows the leading donor(s) to self select. Heterogeneity among donors in the cost of learning the project's quality results in a signaling effect where high quality projects are revealed by receiving more contributions early on. While this does not exactly explain the dynamics we observe in crowdfunding, it provides a nice foundation for the reasoning that consumers will be more likely to pledge as a project is closer to its goal. Our theoretical model focuses more on consumers' belief about the probability of success, which quality may enter implicitly. Consumers in our model also have heterogeneous costs of contributing while the believed probability of success does not vary across consumers.

In the Andreoni (2006) and Vesterlund (2003), the outcome (amount of contributions) is continuous, but crowdfunding has an important binary component in reaching the goal or not. Palfrey and Rosenthal (1984) consider a simultaneous-move game where the public good is binary—either provided or not—and the cost of contributing is constant across players. In cases with and without refunds there are many pure and mixed strategy equilibria, but the mixed strategy equilibria become close to pure strategies as the number of players increases. Cadsby and Maynes (1999) consider a similar model but with continuous contributions which may vary across players. They experimentally test the effect of different features of the game on the likelihood of provision, finding continuous contributions increase the likelihood of provision compared to a binary contribution amount. Refunding contributions if provision fails also improves the likelihood of provision. In our model, the consumers have a trinary decision to pledge with heterogeneous non-refundable costs of doing so, to wait, or to leave.[1]

Makris (2009) propose a model with some portion of players who have a dominant strategy to contribute to the public good—warm-glow altruists. All players are uncertain about the number of total players and the number of warm-glow altruists. Notably, if the total number of players is Poisson distributed, the number of equilibrium contributors is uniquely determined. Moreover, the expected equilibrium number of contributors converges to that of the symmetric mixed-strategy equilibrium of the no uncertainty game as the expected total number of players increases. This idea of uncertainty about the number and composition of other players is analogous to crowdfunding, but the game in Makris (2009) is still simultaneous-move. The key source of uncertainty in our dynamic model is the number of consumers who will learn about the project in future periods.

Bliss and Nalebuff (1984) consider the provision of a threshold public good in a dynamic context where consumers have private information about their own cost. In the model, consumers are able to supply the good themselves and choose an optimal waiting time in equilibrium, which is decreasing in cost. In an equilibrium with finite consumers, the good is not immediately (and thus not efficiently) supplied unless some consumer has zero cost; however, the expected time until the good is supplied goes to zero as the number of consumers goes to infinity.[2] Gradstein (1992) studies a model similar to Bliss and Nalebuff (1984) but with a production function of the public good which is increasing in the number of contributions.[3] With this change, the inefficient waiting time does not go to zero as the number of consumers increases. Our paper has similar ideas to these in the sense that consumers do wait to pledge, but they do so because they are unsure about the total number of consumers who will be aware of the project at the final time period.

Marx and Matthews (2000) consider a dynamic public good provision game with a fixed, but possibly infinite, time horizon. Consumers may make continuous contributions at each stage, and

[1]Consumers also receive utility for "eating" the good which may be thought of net the dollar pledge amount. Consumers either eat the good and pay the dollar amount or do not eat the good and do not pay the dollar amount—because they did not pledge or the project was not successful. Therefore, the only relevant cost is a non-refundable cost of time to pledge.

[2]See the paper for details of additional assumptions.

[3]The production function also has diminishing marginal product. The "production function" in Bliss and Nalebuff (1984) is binary: the good is provided if and only if at least one person contributes.

are motivated to contribute prior to the final period by a threat of punishment in the form of lower contributions by other consumers in the future. Marx and Matthews (2000) show that given a sufficiently high discount factor and sufficiently long time horizon, projects will be completed in finite time under relatively general assumptions. In our model, consumers are motivated to pledge early to avoid some cost of waiting, and any given consumer waiting does not affect the decisions of other consumers (as long as that consumer is around in the last period).

There has been some crowdfunding-specific theoretical work as well. Qiu (2013) proposes a two period theoretical model to explain the type of acceleration in new backers we observe in Chapter 1. In the model, consumers have the option of delaying their pledge to observe what others do, and can also expend some effort advertising so that the project has a higher chance of success. In our model, consumers are small in the sense they do not consider their own pledge having a direct effect on the probability of success. Li and Duan (2016) propose a model more similar to our dynamic model. Consumers arrive according to a Poisson process and realize a heterogenous valuation at each time period. In our case, consumers draw some individual cost of pledging and then possibly choose to remain aware of the project for future time periods; whereas, consumers in Li and Duan (2016) are "new" each period.

2.3 Model Definition

In the context of reward-based crowdfunding, backers only receive a reward if the project is successful, and consumption of the reward may happen months after the fundraising period while monetary payment happens at the end of the fundraising period.[4] Moreover, consumption of the reward happens at the same time for all backers regardless of when the backer pledged during the fundraising period.[5] However, backers still pay some cost of time to go through the pledging process, which may include creating an account and typing in a credit card number.[6] Given these facts, it is not surprising to see some backers wait until the very end of a project's funding cycle to pledge, even if the project is already successful—see Table 1.4 in Chapter 1 for some evidence. One simple explanation for this behavior is those backers are discounting the cost of time to pledge, so they prefer to pledge at the last moment even if the project is already successful. If consumers are able to discount the cost of pledging but not the benefit which is realized at some fixed point in the future, then they naturally wait as long as possible to pledge. But of course, there are many

[4]For large pledges, the payment processor may place a hold on the backer's credit card, so in that case the backer could be temporarily credit constrained.

[5]In practice, a few projects do have multiple production runs and prefer earlier backers for the earlier production runs. We ignore these types of projects in this paper.

[6]Kickstarter.com reports that as of September 2016, about 69% of backers only back one project, so creating an account is a frequent requirement.

other consumers who pledge earlier in the funding cycle, and it is the behavior of these consumers were are concerned with explaining.

One other cost consumers may be faced with is a cost of attention. If a consumer does not pledge today but wishes to possibly pledge in the future, he or she must remember to return to the project's web page at some time in the future (and actually do so). In the model we propose, this cost of attention plays a key role in the consumer's decision to pledge now, to wait and return tomorrow, or to forget about the project completely. When the consumer faces this decision, there is some uncertainty about whether the project will eventually be successful or not. Therefore, the fundamental trade-off is to pledge now and assume the risk of project failure (thus wasting the time cost of pledging) or wait to gain more information (or leave completely).

For an individual consumer, define

- $c_i > 0$ as the cost of time (in dollars) for consumer i to pledge, e.g. create an account and enter payment information,
- $v_i > 0$ as consumer i's net benefit of consuming the reward, i.e. marginal benefit of "eating" the good minus the dollar pledge amount, and
- $\lambda > 0$ as the cost of attention (in dollars) to return one period in the future. Note λ does not vary across consumers, but the model could easily be changed to allow additional heterogeneity.

We now make some assumptions to simplify further discussion. First, normalize $v_i = 1$ and assume c_i is drawn from a distribution F_c on $(0, 1)$, i.e. we only consider consumers who would pledge to the project when it is certain to be successful. Second, restrict λ to also be less than 1; if it were equal to or greater than one, it would never be optimal to wait. Third, assume consumers are risk-neutral and do not discount future periods. Finally, assume consumers are infinitesimal and thus do not have an individual effect on the probability the project is successful.

Let time periods be indexed by $t \in \{0, \dots, T\}$ with finite T, and suppose

- an M_t mass of consumers learn about the project at each t, with M_t drawn from the distribution $F_m(m)$.
- Z_t is the number of consumers who are aware of the project at t, including M_t and all previous consumers who have not left completely (i.e. they could have previously chosen to pledge or to wait). Z_t will depend on both $\{M_t\}$ realizations and the policy function of consumers, defined in a moment.
- B_t is the number of consumers who have pledged to the project before period t; B_t does not include consumers who will choose to pledge at t.

- G is the goal number of consumers. The project is considered successful if $B_{T+1} \geq G$.

If a consumer is aware of the project during period t, he or she observes $M_1, \ldots, M_t$ and will be able to compute Z_t. Suppose in the final period all aware consumers who have not already pledged choose to pledge if and only if $Z_T \geq G$. In other words, if everyone at T could jointly pledge to make the project successful, they will do so. If the project would not meet the goal even if everyone at T pledged, nobody pledges. Note that Z_T depends not only on the realizations of M_t but also the choices of consumers at each period. Now, the probability the project will be successful given information available at t is

$$s_t = \mathbb{P}_t(Z_T(M_1, \ldots, M_T, h_1, \ldots, h_T) \geq G | M_1, \ldots, M_t, h_1, \ldots, h_t). \tag{2.1}$$

where h_t is the policy function defined in Equations 2.6 and 2.7. Before writing the form of the function $Z_t(M_1, \ldots, M_t, h_1, \ldots, h_t)$, let us define

$$\phi(c, a, b, M_1, \ldots, M_b) = \begin{cases} 1 & \text{if none of } h_a(M_1, \ldots, M_a; c), \ldots, h_b(M_1, \ldots, M_b; c) \text{ are “leave”} \\ 0 & \text{if any of } h_a(M_1, \ldots, M_a; c), \ldots, h_b(M_1, \ldots, M_b; c) \text{ are “leave”} \end{cases} \tag{2.2}$$

where $b \geq a$. In other words, $\phi(c, a, b, M_1, \ldots, M_b)$ is equal to one if and only if a consumer with $c_i = c$ who first learned about the project at period a chose not to leave at any period from a to b, inclusive. Now we may write, Z_t, the number of consumers aware of the project at t:

$$Z_t(M_1, \ldots, M_t, h_1, \ldots, h_t) = M_t + \sum_{k=1}^{t-1} M_k \int_0^1 \phi(c, k, t-1, M_1, \ldots, M_{t-1}) f_c(c) dc \tag{2.3}$$

where $f_c(c)$ is the density which c_i are drawn from.

The value function of a consumer with cost c_i aware of the project at $t < T$ is therefore

$$V_t(M_1, \ldots, M_t; c_i) = \max \begin{cases} s_t - c_i & \text{pledge now} \\ \mathbb{E}_t[V_{t+1}(M_1, \ldots, M_{t+1}; c_i)] - \lambda & \text{wait} \\ 0 & \text{leave} \end{cases} \tag{2.4}$$

$$V_T(M_1, \ldots, M_T; c_i) = \begin{cases} 1 - c_i & \text{if } Z_T \geq G \\ 0 & \text{if } Z_T < G \end{cases} \tag{2.5}$$

and the policy function is analogous:

$$h_t(M_1,\ldots,M_t;c_i)=\text{argmax}\begin{cases} s_t-c_i & \text{pledge now}\\ \mathbb{E}_t[V_{t+1}(M_1,\ldots,M_{t+1};c_i)]-\lambda & \text{wait}\\ 0 & \text{leave}\end{cases} \tag{2.6}$$

$$h_T(M_1,\ldots,M_T;c_i)=\begin{cases}\text{pledge now} & \text{if } Z_T\geq G\\ \text{leave} & \text{if } Z_T<G\end{cases} \tag{2.7}$$

Notice how s_t depends not only on $M_1,\ldots,M_t$ but also the primitives of the model (T, F_m, F_c, λ) and the policy function of consumers. The probability is really over future realizations of M_t because consumers know the policy function given c and the distribution of c, so they can compute how many people will pledge/wait/leave given some M_t realization. The policy function also depends on s_t, so equilibrium is reached when Equations 2.1-2.7 hold for every t and every possible $\{M_1,\ldots,M_T\}$ realization. There may be multiple equilibria depending on the primitives of the model.

2.4 Implications of the Model

It may be useful to first visualize a policy function in s and c space in equilibrium. Suppose there are only two periods, then the policy function in the first period is

$$\begin{aligned} h_1(M_1;c_i)&=\text{argmax}\begin{cases} s_1-c_i & \text{pledge now}\\ \mathbb{E}_t[V_2(M_1,M_2;c_i)]-\lambda & \text{wait}\\ 0 & \text{leave}\end{cases}\\ &=\text{argmax}\begin{cases} s_1-c_i & \text{pledge now}\\ s_1(1-c_i)-\lambda & \text{wait}\\ 0 & \text{leave}\end{cases}\\ &=\begin{cases}\text{pledge now} & \text{if } c_i\leq\dfrac{\lambda}{1-s_1} \text{ and } c_i\leq s_1\\ \text{wait} & \text{if } c_i\leq\dfrac{s_1-\lambda}{s_1} \text{ and } c_i>\dfrac{\lambda}{1-s_1}\\ \text{leave} & \text{if } c_i>s_1 \text{ and } c_i>\dfrac{s_1-\lambda}{s_1}\end{cases}\end{aligned}$$

h_1 is illustrated in Figure 2.1 for two different values of λ. In each of the plots, the dark

Figure 2.1: Plots of h_1 with different λ.

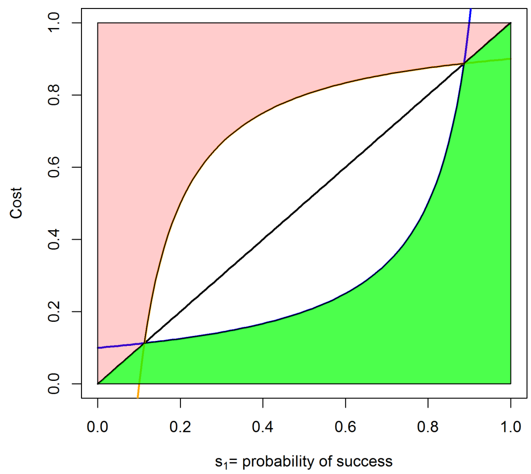

(a) $\lambda = 0.1$

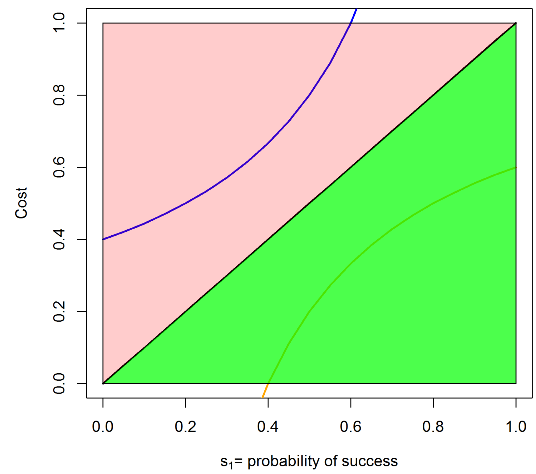

(b) $\lambda = 0.4$

green shaded area represents s_1, c_i combinations where consumer i will pledge now (in period 1). Likewise the light red area shows "leave" and the unshaded area shows "wait." h_1 is computed for each probability of success (on the x-axis) as if the primitives of the model were such that s_1 and h_1 were in equilibrium. Notice how given a fixed c_i a consumer moves from "leave" to "wait" to "pledge now" as the probability of success increases. Also, given a fixed s_1 the policy function moves from "leave" to "wait" to "pledge now" as cost decreases. The unshaded area representing "wait" becomes smaller and eventually empty as λ increases. Just as we might expect, waiting becomes suboptimal for any c, s as the cost of waiting increases.

2.4.1 Without Advertising

One interesting phenomenon found in previous literature and in Chapter 1 is that the number of consumers pledging per day increases as a project nears its goal but then drops down afterwards. In the context of our model, we might expect to see this if some consumers are waiting around and the probability of success increases over time. If λ is low enough, some proportion of consumers at each time period will choose to wait instead of pledge now or leave forever. As the project's probability of success gets closer to one (and likewise, the proportion of the goal raised nears one), more and more of the consumers who were waiting will choose to pledge now. As a result, the number of consumers pledging per period will increase even if M_t is not necessarily increasing. In the few periods prior to meeting its goal, it may even be possible for the number of consumers pledging per period to exceed M_t. After the project meets its goal, everyone will pledge immediately, so the number of people pledging per day will be exactly M_t.

Given the all-or-nothing nature of the policy function in the final period, this effect cannot be shown with only a two period model. Therefore, we present a three period example which illustrates this phenomenon. With discrete distributions of M_t and c_i we can solve for the policy function iteratively, working backwards from the final period. In the penultimate period, we guess some policy function for each c_i and compute the probability of success. Given the probability of success, we may then calculate the policy function. This process converges to a fixed point quite quickly; however, it is sensitive to initial conditions for some parameterizations. For example, if the initial guess of the policy function is for everyone to leave, it may stay at the equilibrium where everyone leaves immediately, even if there is a more efficient equilibrium (in the sense that $s_{T-1} > 0$) available. Therefore, we initialize the policy function to assume everyone pledges immediately, so the process is less likely to converge to an inefficient equilibrium.

Our example uses the following parameterization:

- $G = 3$
- $\lambda = .1$
- $M_1 = 1$
- M_2 and M_3 are drawn from a discrete uniform distribution with 10 possible values equally spaced between 0.9 and 1.5, inclusive.
- c_i are drawn from a discrete uniform distribution with 20 possible values equally spaced between 0.05 and 0.95, inclusive.

After iterating to find the optimal policy function, the unconditional probability of success at the first period is found to be $s_1 = 0.47$. The first period policy function, h_1 is of the following form: pledge if $c_i \leq c^a$; wait if $c^a < c_i \leq c^b$; leave otherwise; where $c^a \approx 0.192$ and $c^b \approx 0.761$. Proportionally, 20% pledge now, 60% wait, and 20% leave. The second period policy function, h_2, has a similar form with the values of c^a and c^b depending on the exact realization of M_2. To illustrate the effect of interest, we will focus on the particular realization $M_2 = 1.2\overline{3}$, which is the lowest realization of M_2 such that $s_2 \geq s_1$. With $M_2 = 1.2\overline{3}$, the h_2 has $c^a \approx 0.287$ and$c^b \approx 0.856$. Notice that because c^a increased from h_1 to h_2, some of the consumers who waited from period 1 will now pledge in period 2. Therefore, we not only have a larger proportion of M_2 consumers pledging, but also some fraction of previously aware consumers pledging at period 2 as well. This is precisely the type of acceleration we are looking for.

For higher M_2 such that $s_2 = 1$, we see that all types of consumers pledge at $t = 2$, so counting the consumers who waited over from period 1, the total number pledging at $t = 2$ is greater than

M_2. In the next period $t = 3$, there are M_3 consumers who all pledge. Now in this particular example $t = 3$ is the final period, but it is apparent that in a model with more time periods, all future periods would have all M_t consumers pledging once $s_t = 1$ even if t is not the final period.

There remains one small disconnect between the model and the data patterns we observe. When discussing the model, we have consumers' exact beliefs about the probability of success and this probability determines their choices. However, in the data we instead observe the proportion of the goal raised. We must assume the proportion of the goal raised and consumers' belief about the probability of success are positively correlated with one another, at least holding time remaining in the funding cycle constant, if we are to use the real-world data patterns as evidence of the model's applicability. This seems quite a reasonable assumption.

2.4.2 With Advertising

Another interesting data pattern shown in Chapter 1 is the immediate effect of advertising appears to be stronger when a project is more highly funded. In this section, we show that "advertising" taking the form of an unanticipated exogenous increase in M_t will typically have a larger immediate effect on the number of pledging consumers when s_t is higher. In addition, we discuss the optimal timing of advertising.

First, note the optimal policy function at any period is to pledge now if the probability of success is one. If $s_{t-1} = 1$, then no consumers will be waiting over to period t. Therefore, an m_a unit increase in M_t from advertising will be met with exactly an m_a unit increase in the number of consumers pledging at t.[7]

However, when s_t is less than one some portion of the consumers aware of the project at t will choose to wait or leave altogether. An m_a increase in M_t in this case has an ambiguous effect on the number of consumers pledging at t. The larger M_t induces a higher probability of success which may cause some consumers to flip from wait or leave to pledge now, so the effect of m_a is magnified (call this the "flipping" effect). However, some portion of the m_a new consumers will choose to wait or leave, so the immediate effect of m_a is diminished (call this the "leaving" effect). These two conflicting effects imply the total immediate effect of advertising when s_t is less than one is ambiguous. One special case does exist where the total immediate effect is certain to be at least m_a. If s_t is less than one without advertising but equal to one with advertising, then all of the m_a consumers will pledge now. In addition, some of the M_t consumers who would have left or waited with $s_t < 0$ will now pledge immediately. In general though, we would expect the leaving effect to overpower the flipping effect for small s_t.

[7]Throughout this section, we assume the m_a advertised consumers draw from the same distribution of costs, F_c, as all other consumers. One could imagine consumers who learn about a project from advertising have a different distribution of costs, but we do not delve into this problem here.

Table 2.1: Effect of 0.1 Unit of Advertising on the Probability of Success

Advertising Period	Probability of Success at $t = 1$
None	0.47
$t = 1$	0.67
$t = 2$	0.57
$t = 3$	0.48

Table 2.2: Effect of 1 Unit of Advertising on the Probability of Success; $M_1 = 0.1$

Advertising Period	Probability of Success at $t = 1$
None	0.00
$t = 1$	0.67
$t = 2$	0.70
$t = 3$	0.00

Entrepreneurs are likely also concerned about how the timing of advertising affects the project's overall probability of success. In light of this fact, we return to the previous numeric example and explore the effect of advertising at different time periods on the unconditional probability of success, s_1. To be clear, consumers behave as before (as if $s_1 = 0.47$) until the period where advertising happens—then they update their beliefs and policy function. Table 2.1 displays the effect of 0.1 unit of advertising on the probability the project is successful. It appears at first glance that earlier advertising is better, but this is only one parameterization of a three period model.

In Table 2.2 we present similar results but with $M_1 = 0.1$ and the amount of advertising equal to 1 unit. Under this parameterization, the probability of success is zero without any advertising. The project still fails when the advertising happens at the last period because consumers in earlier periods do not anticipate there being any hope of success, and therefore all leave, even though M_2 and M_3 could be high enough to reach the goal number of consumers. Table 2.2 also shows that earlier advertising is not unequivocally better.

2.5 Discussion

In this paper we have presented a dynamic model of consumer behavior in crowdfunding markets. The model allows consumers to consider future periods—both in terms of future utility gains and changes in the probability of success of the project. Unfortunately, the complexity does not allow a simple analytic solution to a project's funding path. However, simulations with reasonable primitives do suggest the dynamic model implies the types of behavior we see in the real world—in particular, the number of new backers per day accelerates up until the project meets its goal, then drops back down.

For future research one could estimate the dynamic model structurally, but this requires finer data than what we have access to in Chapter 1. One interesting parameter to estimate is λ, which measures the time cost of consumers returning to check on the project in the future. The ideal data tracks individual consumers over a project's life cycle and includes visits to the project's web page as well as views of advertising or social media attention elsewhere. However, the model could still be estimated with data which disentangles traffic to a project's page from the number of pledges, which could still be aggregated to the daily level, for example. In Chapter 1, we have some measures of traffic (posts on Reddit, pledges to other projects, day of week, etc.), but the data are not fine enough to separate traffic from the propensity to pledge.

CHAPTER 3

Issues with Inference Using the Synthetic Control Method

3.1 Introduction

Comparative case studies with few (even one) treated units and aggregate data require selection of control units for estimation and inference of the treatment effect. Selection of control units has typically been a subjective procedure which complicates inference. Moreover, conditional on a selection of control units, the comparability of the controls and treated units relative to a "placebo" selection is often unclear.

The Synthetic Control Method (SCM) is designed to mitigate these problems. It was first used by Abadie and Gardeazabal (2003) and later formalized in a seminal paper by Abadie et al. (2010). The general idea of the SCM is to estimate a weighted average of potential control units in a data-driven way using pretreatment data. This admits at least three significant benefits: (1) the weight of each potential control unit is made explicit, (2) the selection procedure can be objectively repeated in a permutation test for each placebo, (3) the relative fit of the synthetic controls to the pretreatment realizations is made explicit for the treated unit and placebos.

Abadie et al. (2010) show that as long as "true" weights exist and the data conform to a linear factor model, the absolute bias of the SCM goes to zero as the number of pretreatment periods increases. They also show the unbiasedness of the SCM for an autoregressive model with time-varying coefficients requiring only one pretreatment period.

In this paper, we highlight some data generating processes for which the synthetic control method exhibits differential performance across units in finite samples. In particular, if the outcome variable of interest is correlated within groups of units, SCM may result in poor fit of units in small groups (relative to units in larger groups). If group structure is unobserved, it may be difficult for the researcher to infer whether the unit of interest is being fit well relative to other units, especially if the treated unit lies is a small group. However, we show inference based on the mean squared prediction error ratio is not substantially distorted for units in small groups. Additionally, we offer word of warning about including all pretreatment outcomes as economic predictors in the selection of synthetic weights. Doing so could invalidate inference based on the post/pretreatment mean squared prediction error (MSPE) ratio if the researcher would have chosen not to report results when pretreatment MSPE for the treated unit was high.

3.2 Literature Review

There has been a lot of recent econometric work refining and studying the SCM. Abadie et al. (2014) study the 1990 German reunification and consider explicitly restricting the number of potential control units to limit overfitting. In that particular application, this restriction appears to have minimal effect on the inferential outcome, except for the case of only one potential control unit (which also corresponds to substantially poorer pretreatment fit).

Ferman and Pinto (2017) show results similar to ours, albeit more formally. In particular, they show the SCM can be asymptotically biased when common factors are nonstationary, and smaller groups of units have a larger proportion of misallocated weights, even in the stationary case. Kaul et al. (2016) note that including each pretreatment outcome as economic predictors can significantly bias the SCM estimator when there exist other covariates which are meaningful for counterfactual prediction; the source of the problem is that the SCM will not place any weight on these covariates. Despite this fact, Kaul et al. (2016) note several published papers using the SCM do use pretreatment outcomes as predictors.[1] We show that even without other meaningful covariates, including pretreatment outcomes could lead to invalid inference if the researcher chooses to use the SCM or not based on pretreatment fit of the treated unit.

Firpo and Possebom (2016) formalize the inferential procedures of Abadie et al. (2010) and introduce confidence sets for the SCM. They also compare the performance of the SCM to similar diff-in-diff procedures and find the SCM offers improvements, at least for the data generating processes (DGP) they consider. Doudchenko and Imbens (2016) propose relaxing some of the restrictions imposed by the SCM, for example that the weights sum to one, and applying regularization techniques such as elastic net to deal with the large number of potential control unit weights.

A number of papers use the SCM in applied work as well; we note only a few of these below. Billmeier and Nannicini (2013) apply the SCM to measure the effect of 30 country-level economic liberalization events on per capita real GDP. Bilgel and Galle (2015) use the SCM to assess tax incentives for kidney donations. Cavallo et al. (2013) measure the effect of large natural disasters on GDP growth.

[1] It is important to note Abadie et al. (2010) only use three of 19 available pretreatment period outcomes and do not suggest one should use all available outcome, but they do not warn against doing so, either.

3.3 Synthetic Control Method Setup and Estimation

Suppose there are $J + 1$ units (e.g. US states) indexed by i where unit $i = 1$ is the only treated unit. Further suppose we observe T time periods indexed by t, and let T_0 be the last untreated time period. The observed outcome variable of interest is Y_{it}. Y_{it}^N is the outcome with no treatment, and Y_{it}^I is the outcome with the treatment. Assume a model with linear additive treatment effects which may vary over time:

$$Y_{it} = Y_{it}^N + \alpha_{it} D_{it}$$

where α_{it} are the treatment effects of interest, and D_{it} is a dummy variable equal to one if and only if unit i has received the treatment by t. For $t > T_0$, we observe only Y_{it}^I for the treated unit $i = 1$, and Y_{it}^N for all other units.

To motivate the SCM, Abadie et al. (2010) use a linear factor model:

$$Y_{it}^N = \delta_t + \boldsymbol{\theta}_t \boldsymbol{Z}_i + \boldsymbol{\lambda}_t \boldsymbol{\mu}_i + \varepsilon_{it}$$

where

- δ_t is an unknown common factor with constant loadings
- $\boldsymbol{Z}_i$ is a vector of observed covariates (length r)
- $\boldsymbol{\theta}_t$ is a vector of unknown parameters
- $\boldsymbol{\lambda}_t$ is a vector of unobserved common factors with loadings $\boldsymbol{\mu}_i$ and length F, and
- ε_{it} is a mean zero shock, independent across units and time.

Estimation of Y_{1t}^N, $t < T_0$, for the treated unit is necessary to estimate the treatment effects α_{it}. Abadie et al. (2010) propose constructing a weighted average of untreated units' outcomes which would closely follow the counterfactual (untreated) outcome of the treated unit during post-treatment periods.

Now consider weights $\boldsymbol{W} = (w_2, \ldots, w_{J+1})'$ with the restrictions $w_j \geq 0$ and $w_2 + \ldots + w_{J+1} = 1$. Suppose there exists $\boldsymbol{W}^*$ such that

$$\sum_{j=2}^{J+1} w_j^* Y_{j1} = Y_{11}, \sum_{j=2}^{J+1} w_j^* Y_{j2} = Y_{12}, \ldots, \sum_{j=2}^{J+1} w_j^* Y_{jT_0} = Y_{1T_0}$$
$$\sum_{j=2}^{J+1} w_j^* \boldsymbol{Z}_j = \boldsymbol{Z}_1.$$

Then Abadie et al. (2010) (Appendix B) show

$$\mathbb{E}\left|Y_{1t}^N - \sum_{j=2}^{J+1} w_j^* Y_{jt}\right|$$

is bounded by function which goes to zero as the number of pretreatment periods, T_0, increases—with standard assumptions on ε_{it}. Therefore, an estimator of $\hat{\alpha}_{1t}$ is $Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt}$.

However, $\boldsymbol{W}^*$ is not observed. The core of the SCM lies in the estimation of $\hat{\boldsymbol{W}}$. Define the $(k \times 1)$ vector

$$\boldsymbol{X}_1 = (\boldsymbol{Z}_1', \bar{Y}_1^{(1)}, \ldots, \bar{Y}_1^{(M)})$$

where $\bar{Y}_1^{(m)}$ are linear combinations of pretreatment outcomes. Define $\boldsymbol{X}_0$ similarly as a $(k \times J)$ matrix for the untreated units $j = 2, \ldots, J+1$. Now choose $\hat{\boldsymbol{W}}$ to minimize

$$\sqrt{(\boldsymbol{X}_1 - \boldsymbol{X}_0\boldsymbol{W})'\boldsymbol{V}(\boldsymbol{X}_1 - \boldsymbol{X}_0\boldsymbol{W})}$$

where $\boldsymbol{V}$ is a $(k \times k)$ symmetric positive-semidefinite matrix of weights on the variables in $\boldsymbol{X}$. Abadie et al. (2010) choose $\boldsymbol{V}$ to minimize the mean squared prediction error of the outcome in pretreatment periods.

The resulting estimator of the treatment effect, α_{1t}, is

$$\hat{\alpha}_{1t} = Y_{1t} - \sum_{j=2}^{J+1} \hat{w}_j Y_{jt}$$

where $\sum_{j=2}^{J+1} \hat{w}_j Y_{jt}$ is called the "synthetic control."

3.4 Synthetic Control Method Inference

Abadie et al. (2010) propose a type of permutation test to infer whether $\hat{\alpha}_{1t}$ is large relative to what might be expected due to noise alone—both from the estimation of the weights and the underlying data-generating process.[2] Abadie et al. (2010) apply the SCM to each untreated unit in the sample to obtain estimates $\hat{\alpha}_{jt}$ for $j \neq 1$. Intuitively, one may compare $\hat{\alpha}_{1t}$ during posttreatment periods to the distribution of $\hat{\alpha}_{jt}$, $j \neq 1$, during posttreatment periods to see if $\hat{\alpha}_{1t}$ is abnormal and thus implies a significant non-zero treatment effect.

In practice, some of the units may have poor pretreatment fit (i.e. large $\hat{\alpha}_{jt}$, $t \leq T_0$) relative to other units, which suggests either the SCM performs poorly for these units or these units are

[2] Although Firpo and Possebom (2016) introduce confidence sets, we stick with Abadie et al.'s (2010) original inference procedure for this paper.

simply more noisy. Abadie et al. (2010) handle this problem in two ways.

First, they consider throwing out units with poor pretreatment fit and then visually comparing $\hat{\alpha}_{1t}$ to $\hat{\alpha}_{jt}$, $j \neq 1$ during the posttreatment periods. Specifically, they throw out units with a pretreatment MSPE some multiple (say, five times) higher than the treated unit of interest.[3]

Second, they consider the ratio of posttreatment MSPE to pretreatment MSPE:

$$R_i = \left(\frac{1}{(T - T_0)} \sum_{t=T_0}^{T} \hat{\alpha}_{it} \right) \Bigg/ \left(\frac{1}{T_0} \sum_{t=1}^{T_0} \hat{\alpha}_{it} \right).$$

If the ratio for the treated unit, R_1, is large relative to the distribution of ratios of all untreated units, R_j, $j \neq 1$, one may infer there is a significant treatment effect. More formally, if the treatment label were assigned to a unit at random, the probability of obtaining a ratio for the labeled unit at least as high as R_1 is

$$\frac{\sum_{j=1}^{J+1} \mathbb{I}(R_j \geq R_1)}{J + 1}.$$

This is conceptually similar to Fisher's exact test.

3.5 DGPs with Potentially Poor Synthetic Control Performance

We consider a very basic model which allows Z_{it} (now scalar) to vary over time, but only Y_{it} is now observable:

$$Y_{it}^{N} = Z_{it} + \varepsilon_{it}$$

Let $g(i) = 1, \ldots, G$ index the group which contains unit i, and suppose Z_{it} is equal for units in the same group, $Z_{it} = Z_{jt}$ if $g(i) = g(j)$. Importantly, $g(i)$ is unobserved to the econometrician. Now suppose $Z_{g(i)t}$ is an AR(1) process. ε_{it} are independent and identically distributed. Then, the outcome variable Y_{it}^{N} is an AR(1) process which is correlated with other units within its group.

Many outcome variables in practice may be correlated across units within groups, and may also have unit roots (so we will allow $Z_{g(i)t}$ to have a unit root). For example, the outcome variable in Abadie et al. (2010) is aggregate state-level cigarette consumption per capita. It would be reasonable to expect variation in the observed dependent variable which affects certain regions (such as the Northeast) but not others (such as the Southwest), but the linear factor model requires all units to be affected in the same proportion, through $\boldsymbol{\theta}_t$.

Although the model may seem overly simplistic, Kaul et al. (2016) note that including pretreatment outcomes at each period as predictors in the selection of synthetic weights (which we do) implies any other covariates (observed independent variables, for example) contribute nothing

[3] Pretreatment MSPE for unit i is $\frac{1}{T_0} \sum_{t=1}^{T_0} \hat{\alpha}_{it}$

to the selection of weights but may still impact counterfactual prediction. Despite this fact, Kaul et al. (2016) notes many researches have done this in practice. In our model, there are no other observed covariates which could impact counterfactual prediction, so we are in some sense operating in an "easier" scenario. Adding observed covariates would likely only magnify our results.

As a result of the correlation within groups, one intuitively expects the SCM to choose higher weights for units in the same group. Additionally, the synthetic control of unit i may fit the realized outcome of unit i during the posttreatment period better if more units exist in its group. If this is true, units in small groups may also be more likely to exhibit a high MSPE ratio—which could complicate the inferential procedure. If the treated unit is in a small group (relative to other units), it may appear to have a high MSPE ratio even if the treatment effect is zero.

We now explore these concerns with simulations of the following parameterization:

- $Z_{g(i)t} = \beta Z_{g(i)t-1} + u_{g(i)t}$ where $u_{g(i)t} \sim N(0, 1)$ is iid across g and t.
- $\varepsilon_{it} \sim N(0, 1)$ is iid across i and t.
- $T = 16$, $T_0 = 8$, and $\alpha_{it} = 0\ \forall i, t$.

Figure 3.1 displays an example realization of Y_{it} with all units on the same plot, and Figure 3.2 displays the same realization with units grouped according to $g(i)$. In this example the treated unit has one other unit in its group, $J = 51$, $\beta = 1$, and all other groups are of size 10. Note how it is difficult to see simply from Figure 3.1 whether there is any clustering, let alone which units are in which cluster. More importantly, it is difficult to infer how many control units the treated unit has in its group.[4]

We now consider the difference between simulated and nominal rejection probabilities for many different variations of the above parameterization. For the estimation of weights, we use each pretreatment period outcome, so
$\boldsymbol{X}_1 = (Y_{11}, Y_{12}, \ldots, Y_{1T_0})$ and V is the identity matrix I_J. We use an R implementation of the SCM provided by Abadie et al. (2011). We reject the null hypothesis of no treatment effect if the "treated" unit's MSPE ratio is the fifth highest (or higher), so the exact nominal rejection rates depend on the number of units and range from $8.\overline{3}\%$ to 9.6%.. The number of control units outside the treated unit's group is 50 for each parameterization, but the number of control units within the treated unit's group varies from one to nine. β takes one of the values in $\{0.5, 0.7, 1\}$.

For each parameterization, we simulate 1000 independent panels. Figure 3.3 displays the results. The plots on the left display the difference in rejection proportions for the treated units, while the plots on the right display the same for control units. There appears to be no strong strong

[4]Keep in mind only periods 1–8 should be used for this inference in practice because the treatment occurs in period 9.

Figure 3.1: Example realization of Y (treated unit in bold)

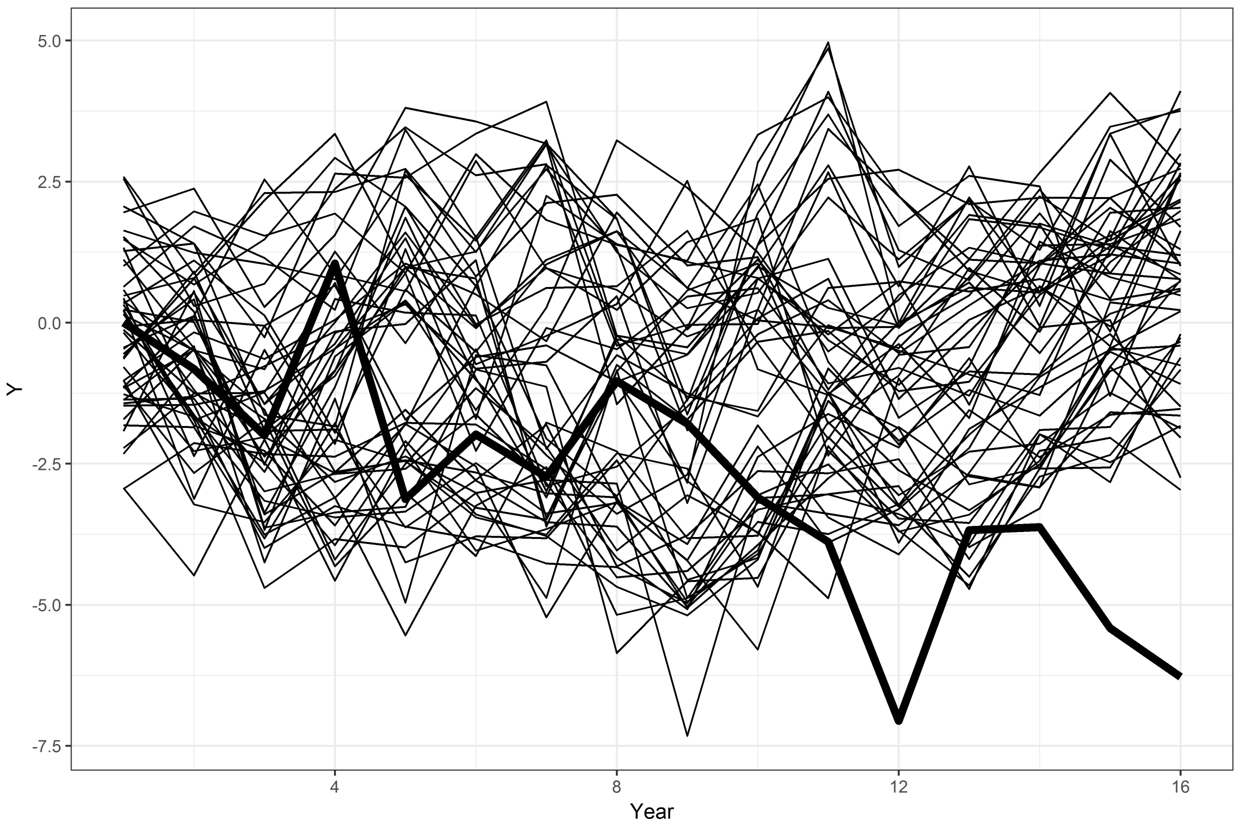

Figure 3.2: Example realization of Y (grouped; treated unit in group 0)

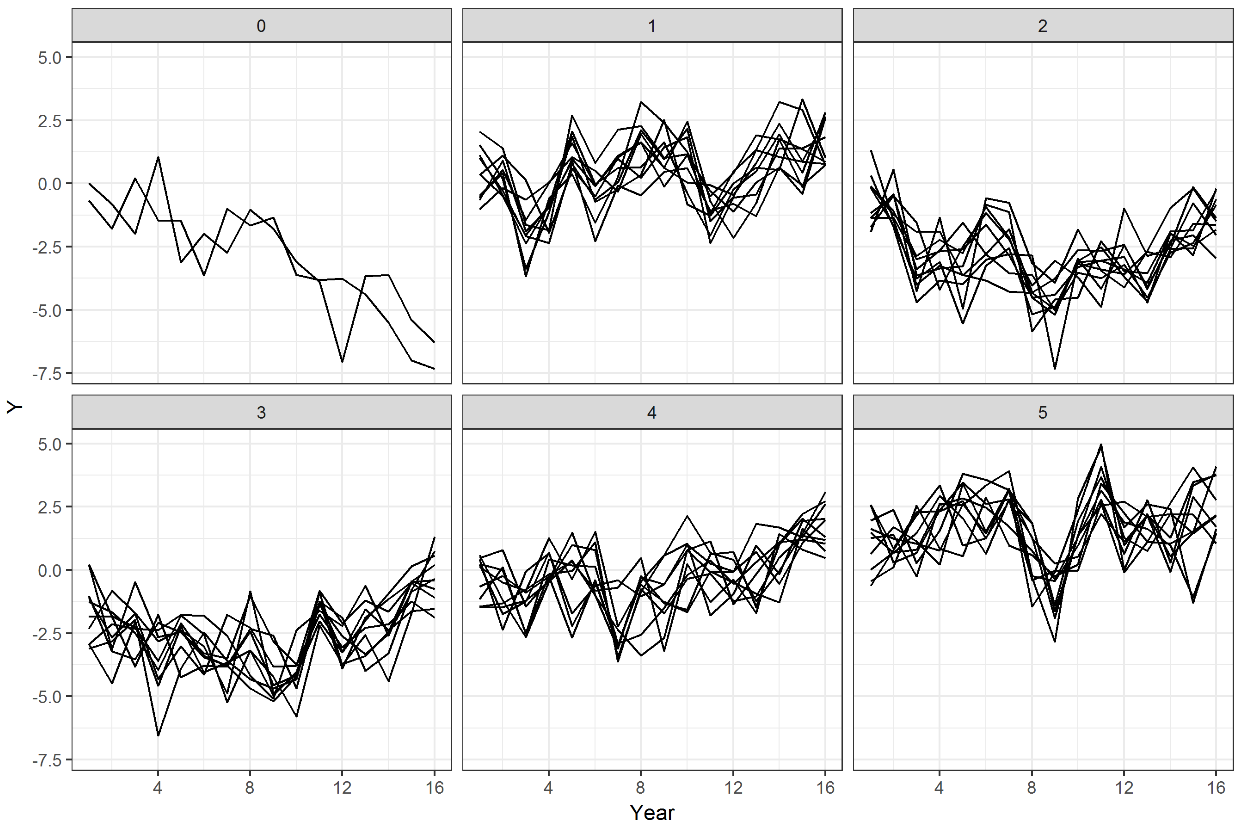

Table 3.1: MSPE Statistics for Treated Unit Group Size = 2, Control Unit Group Size = 5

T = Treated, C = Control, $5/52 \approx 9.6\%$ nominal rejection rate

	$\beta = 1$		$\beta = 0$		$\beta = 0.5$	
	T	C	T	C	T	C
Proportion Rejected	0.084	0.104	0.085	0.093	0.083	0.106
Median MSPE Ratio	11.24	9.92	10.94	12.32	6.02	7.00
Median Pretreatment MSPE	0.428	0.292	0.316	0.230	0.378	0.265
Median Posttreatment MSPE	3.98	2.85	3.46	3.00	2.19	1.89

distortions except for the treated unit when its group size (2) is much less than that of the control groups (10), but even in that case, inference based on the MSPE is conservative.

Although there do not appear to be substantial distortions in rejection probabilities, units in smaller groups do tend to have worse fit (higher MSPE) both pre- and posttreatment. Table 3.1 displays some additional statistics for the parameterization with a treated unit group size of 2 and a control unit group size of 5. For the three values of β considered, the treated unit has a higher median MSPE both pre- and posttreatment than a control unit.

3.6 Selection of Analyses Based on Pretreatment MSPE of the Treated Unit

In some cases the pretreatment MSPE of the treated unit is large relative to other units, so the researcher may choose not to use the SCM (or not to report any results)—in other words, the optimal weights $\boldsymbol{W}^*$ are presumed not to exist. Although Abadie et al. (2010) do not directly say to do this, they write the following passage which could be misapplied:

> So for each particular application, the analyst can decide if the characteristics of the treated unit are sufficiently matched by the synthetic control. In some instances, the fit may be poor and then we would not recommend using a synthetic control.

In Abadie et al.'s (2010) explanation, "characteristics" include several other covariates in addition to pretreatment outcomes. However, if a researcher primarily (or only) considered pretreatment fit of outcomes *and* would not report the statistical analysis if the fit was poor, then analyses with low pretreatment MSPE are implicitly being selected for. Therefore, inference based on the MSPE ratio, R_1, is likely to be affected. Recall pretreatment MSPE of the treated unit is in the denominator of R_1, so an abnormally low pretreatment MSPE may imply an abnormally high R_1 even in the absence of a non-zero treatment effect.[5]

[5] By "abnormally" low/high we mean relative to the distribution implied by the underlying DGP and application of the SCM without any selection of analyses.

Figure 3.3: Plots of Deviations from Nominal Rejection Rates

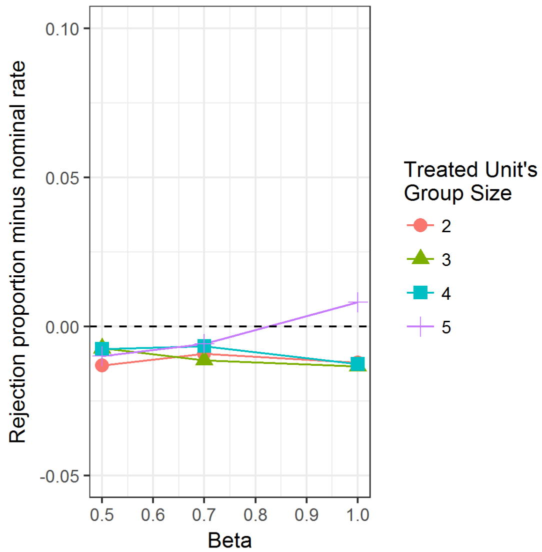

(a) Treated Units; Control Groups Size 5

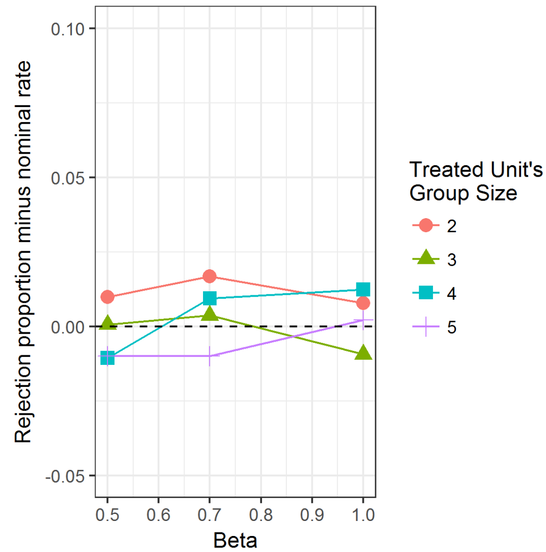

(b) Control Units; Control Groups Size 5

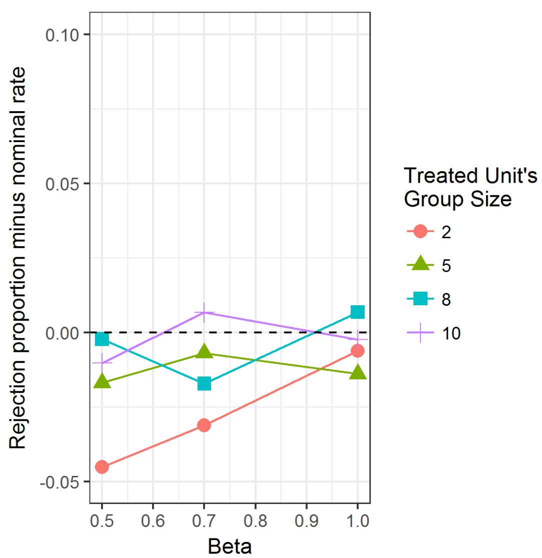

(c) Treated Units; Control Groups Size 10

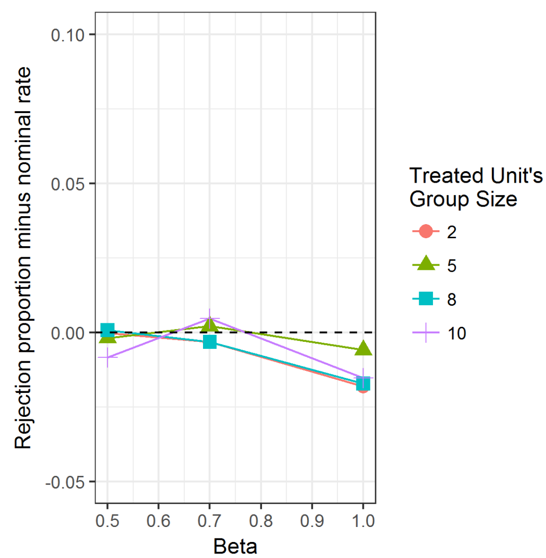

(d) Control Units; Control Groups Size 10

Table 3.2: MSPE Statistics for Panels with "Good" Treated Unit Fit; Treated Unit Group Size = 2, Control Unit Group Size = 5

T = Treated, C = Control, $5/52 \approx 9.6\%$ nominal rejection rate

	$\beta = 1$		$\beta = 0$		$\beta = 0.5$	
	T	C	T	C	T	C
Proportion Rejected	0.133	0.0978	0.123	0.0967	0.132	0.100
Median MSPE Ratio	22.97	9.67	22.28	12.86	12.00	7.10
Median Pretreatment MSPE	0.241	0.308	0.163	0.236	0.198	0.274
Median Posttreatment MSPE	4.20	2.74	3.41	3.12	2.18	1.96

For the same sets of 1000 panels in Section 3.5, we consider only those where the pretreatment MSPE of the treated unit lies below the 75th percentile of all control units' pretreatment MSPE; there are around 600–700 such panels depending on the exact parameterization. Figure 3.4 and Table 3.2 displays the results. As one might expect, rejection probabilities for the treated unit are now significantly higher than nominal rates, while those of the control units are close to nominal. The intuition here is that researchers may conclude the SCM is performing well based on the goodness of fit of the treated unit during the pretreatment periods (relative to control units), but this fact may also correspond to a higher MSPE ratio (relative to control units) even without a treatment effect. The 75th percentile is not a particularly aggressive cutoff, either; more aggressive requirements for pretreatment fit will exacerbate the problem. The simple solution is, as Abadie et al. (2010) suggest, to consider how close other covariates are to their synthetic counterparts instead of only the pretreatment MSPE of the treated unit.

3.7 Discussion

In this paper, we have shown that if outcome variables of units are correlated within groups, the SCM can have varying performance across group sizes. In particular, if the treated unit of interest belongs to a small group (relative to other units' group sizes), then the fit of the synthetic control to the realized outcome will be poorer for the treated unit both pre- and posttreatment. This could cause problems in interpretation of SCM performance if the researcher does not observe group structure. However, we do not see substantial distortions in inference using the MSPE ratio, and in fact, small groups may have relatively conservative rejection probabilities. Ferman and Pinto (2017) make progress in theoretical evaluation of these claims by showing non-stationary common factors can lead to asymptotic bias, although they use a slightly different DGP.

Our second contribution is a word of warning that applying the SCM only in situations where the unit of interest has "good" pretreatment fit may complicate inference based on the MSPE ratio, if pretreatment outcomes at each period are used as predictors (in $\boldsymbol{X}$) in the selection of synthetic

Figure 3.4: Plots of Deviations from Nominal Rejection Rates (for Panels with "Good" Treated Unit Fit)

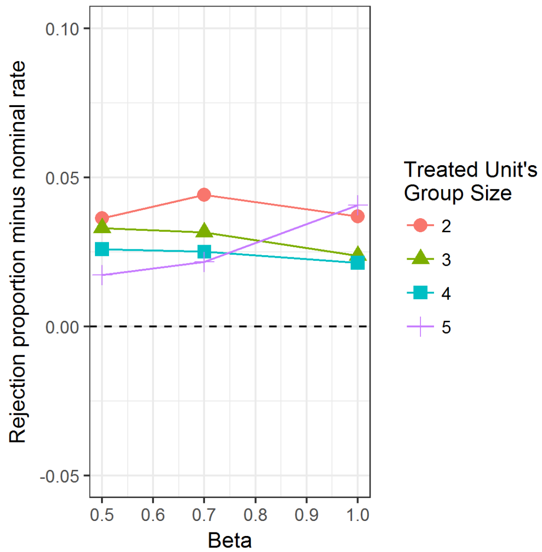

(a) Treated Units; Control Groups Size 5

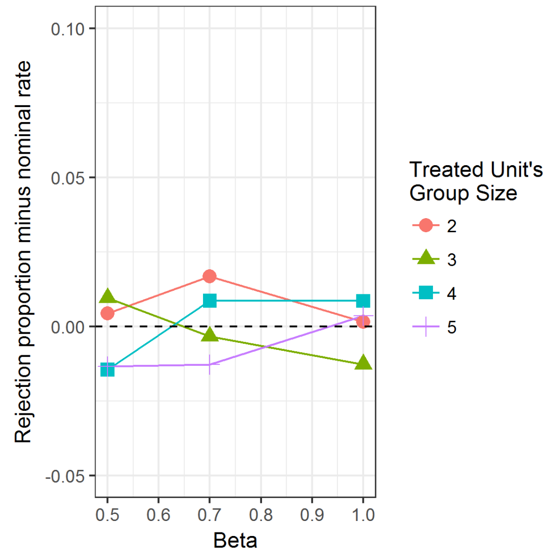

(b) Control Units; Control Groups Size 5

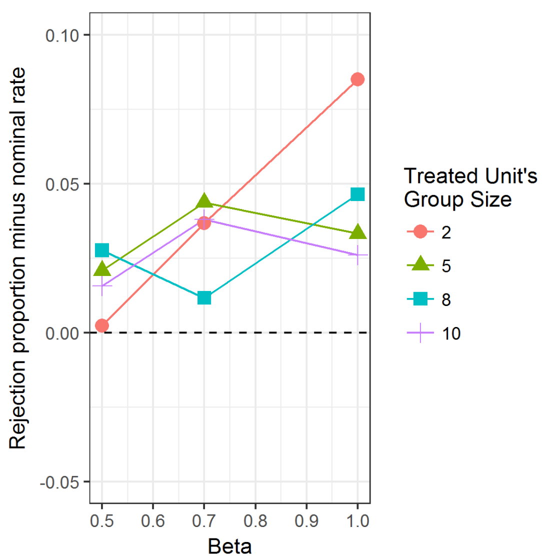

(c) Treated Units; Control Groups Size 10

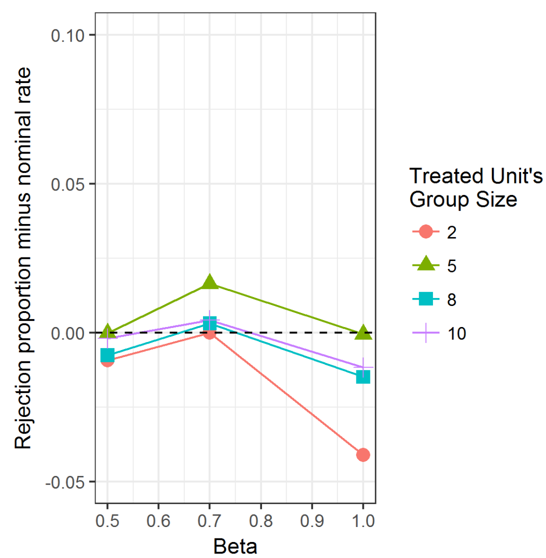

(d) Control Units; Control Groups Size 10 weights. It is important to note the inferential results described by Firpo and Possebom (2016) are not contradictory to what we show here; in our case, Assumption 2 of Firpo and Possebom (2016) is violated because the labeling of the treated unit is no longer random. Kaul et al. (2016) shows that even when other covariates (which are meaningful for counterfactual prediction) are included as predictors, the SCM will select weights such that only the pretreatment outcomes matter. So including other covariates does not solve our problem. As Kaul et al. (2016) propose, one should include only the mean or at most one pretreatment realization of the outcome variable. And given our results, one should not throw away analyses based solely on the pretreatment MSPE of the treated unit if the post/pretreatment MSPE ratio is to be used for inference.

APPENDIX A

Additional Statistics and Empirical Results

A.1 Descriptive Statistics for the Full Sample

Except for the dates projects are active, the descriptive statistics for the full sample are substantially similar to the subset used for the empirical analysis.

Table A.1: (Full Sample) Kickstarter projects descriptive statistics.

Variable	Min	Median	Max	Mean	Std. Dev.
Duration (days)	1	30.00	84.00	32.09	10.48
Goal ($)	1	6000.00	100000000.00	31895.02	897252.41
Goal if successful ($)	1	5000.00	1500000.00	10550.05	27467.91
Pledged ($)	0	1568.00	13285226.36	11276.07	97238.34
Pledged if successful ($)	1	5766.22	13285226.36	22096.55	142277.29
Prop of goal achieved	0	0.41	15804.00	2.63	126.81
Prop of goal achieved if suc	1	1.16	15804.00	5.58	186.96
Backers	0	26.00	105857.00	138.09	856.73
Backers if successful	1	80.00	105857.00	270.21	1246.09

A.2 Poisson With Fixed Effects à la Kuppuswamy and Bayus (2017)

For the subset of projects which last exactly 30 days, we estimate Poisson regression with project fixed effects. The results presented in Table A.2 are similar to Kuppuswamy and Bayus (2017) which demonstrates the set of projects we use is substantially comparable to that of previous literature.

Figure A.1: (Full Sample) Histograms of project features.

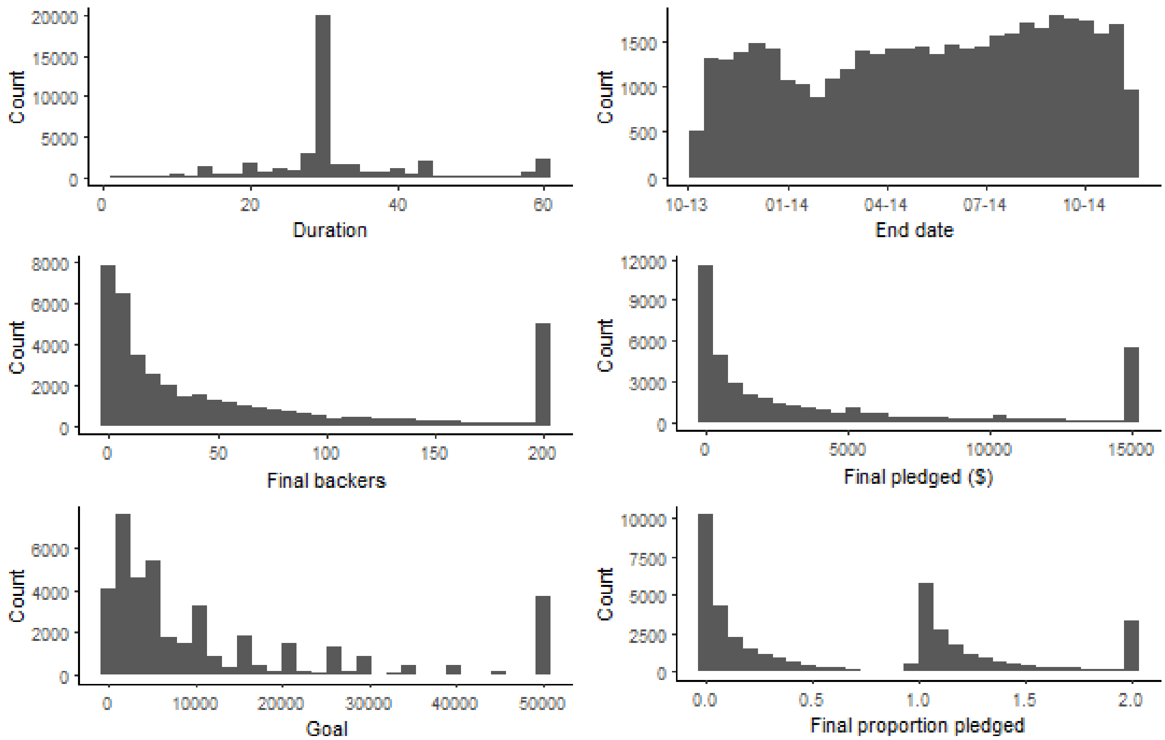

Figure A.2: (Full Sample) Histograms of the number of posts per project and number of comments per project-day with a new post.

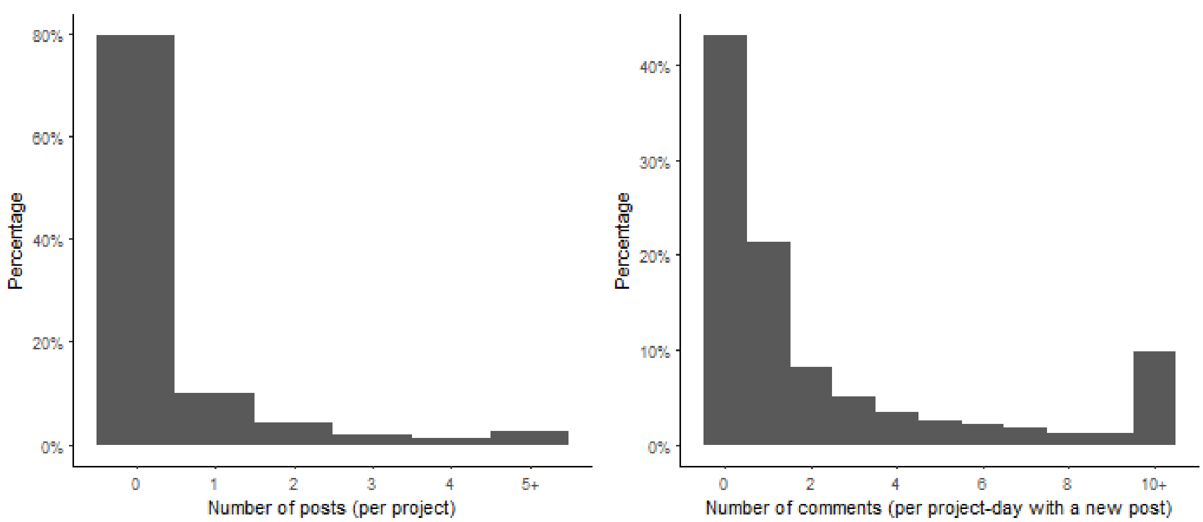

Table A.2: Poisson regression of new backers on controls

	(1)		(2)	
	coef	se	coef	se
$0.2 \leq PropGoal_{i,t-1} < 0.4$	0.316***	(0.0497)	0.134***	(0.0417)
$0.4 \leq PropGoal_{i,t-1} < 0.6$	0.582***	(0.0723)	0.366***	(0.0547)
$0.6 \leq PropGoal_{i,t-1} < 0.8$	0.781***	(0.0870)	0.591***	(0.0737)
$0.8 \leq PropGoal_{i,t-1} < 1$	1.000***	(0.104)	0.779***	(0.0827)
$1 \leq PropGoal_{i,t-1}$	0.494***	(0.122)	0.193*	(0.103)
$t = 2$	1.418***	(0.0713)		
$t = 3$	1.082***	(0.0859)		
$t = 4$	0.783***	(0.0700)	0.583***	(0.0592)
$t = 5$	0.671***	(0.0653)	0.526***	(0.0522)
$t = 6$	0.470***	(0.0439)	0.380***	(0.0360)
$t = 7$	0.372***	(0.0408)	0.316***	(0.0373)
$t = T$	1.411***	(0.0514)	1.436***	(0.0503)
$t = T-1$	0.955***	(0.0473)	0.982***	(0.0430)
$t = T-2$	0.529***	(0.0625)	0.559***	(0.0575)
$t = T-3$	0.371***	(0.0677)	0.387***	(0.0613)
$t = T-4$	0.242***	(0.0639)	0.261***	(0.0629)
$t = T-5$	0.136***	(0.0467)	0.155***	(0.0462)
$t = T-6$	0.116***	(0.0435)	0.137***	(0.0391)
Monday	0.222***	(0.0159)	0.212***	(0.0186)
Tuesday	0.235***	(0.0209)	0.229***	(0.0200)
Wednesday	0.246***	(0.0252)	0.226***	(0.0209)
Thursday	0.253***	(0.0277)	0.200***	(0.0220)
Friday	0.128***	(0.0232)	0.101***	(0.0236)
Saturday	-0.0723***	(0.0180)	-0.0698***	(0.0212)
$NumOtherProjects_{it}$	-0.000189	(0.00170)	-0.00120	(0.00181)
$NumOtherNewbackers_{it}$	5.86e-06	(2.77e-05)	-1.61e-05	(1.72e-05)
$RedditPosts_{t-1} > 0$			0.357***	(0.0593)
$RedditPosts_{t-2} > 0$			0.195***	(0.0386)
$RedditPosts_{t-3} > 0$			0.168***	(0.0376)
Observations	306,646		279,666	
Number of projects	10,574		10,358	

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure A.3: (Full Sample) Distribution of Reddit posts across projects' funding cycles.

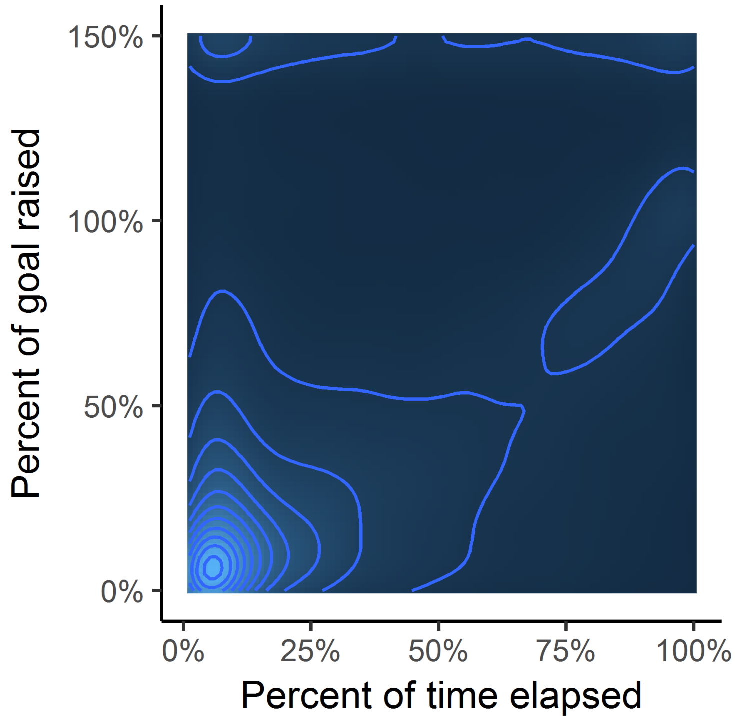

A.3 Estimation of the Probability of Success at t

In this section, we estimate the probability of success of a project based on the number of days remaining and the proportion of the goal raised so far, $\hat{s}(T-t, g_t)$. For the subset of projects which last exactly 30 days, we take a random day from each project as an observation. Then, we perform a logistic regression with:

$$\mathbb{P}[Success_i|X_i] = (1 + \exp(-\alpha - X_i\beta))^{-1}$$

where X_i includes the funding proportion and days remaining as 2nd degree polynomials and the interaction between the two. The results are presented in Table A.3.[1] Figure A.4 illustrates the predicted probabilities for a grid of values.

[1]Table created by the stargazer package.
Hlavac, Marek (2015). stargazer: Well-Formatted Regression and Summary Statistics Tables. R package version 5.2.
http://CRAN.R-project.org/package=stargazer

Table A.3: Logistic regression for predicting the probability of success

g_i	22.859^{***}
	(1.166)
$DaysLeft_i$	0.391^{***}
	(0.034)
g_i^2	-11.322^{***}
	(0.901)
$DaysLeft_i^2$	-0.006^{***}
	(0.001)
$g_i \times DaysLeft_i$	-0.187^{***}
	(0.043)
Constant	-8.956^{***}
	(0.397)
Observations	14,670
Log Likelihood	$-3,728.185$
Akaike Inf. Crit.	7,468.370
Note:	*p<0.1; **p<0.05; ***p<0.01

Figure A.4: Logistic prediction of probability of success

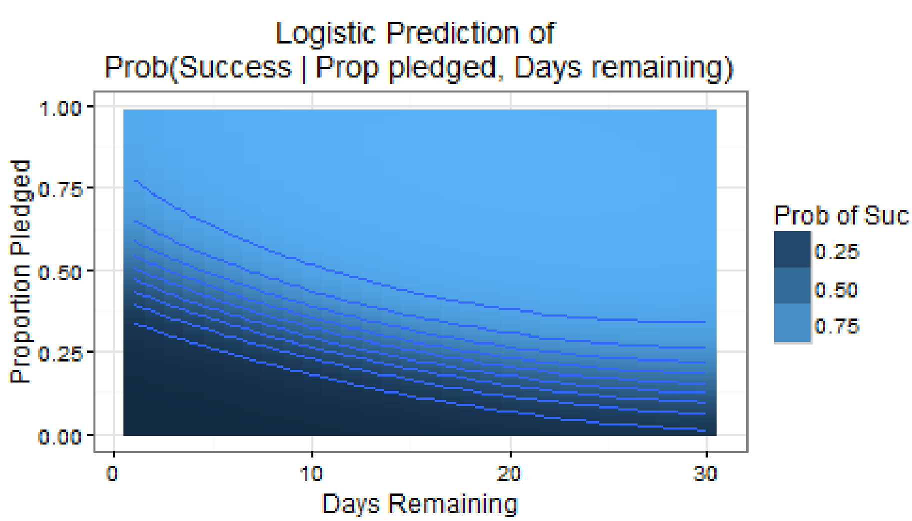

BIBLIOGRAPHY

ABADIE, A., A. DIAMOND, AND J. HAINMUELLER (2010): "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 105, 493–505.

——— (2011): "Synth: An R Package for Synthetic Control Methods in Comparative Case Studies," *Journal of Statistical Software*, 42, 1–17.

——— (2014): "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, 59, 495–510.

ABADIE, A. AND J. GARDEAZABAL (2003): "The Economic Costs of Conflict: A Case Study of the Basque Country," *The American Economic Review*, 93, 113–132.

AGRAWAL, A. K., C. CATALINI, AND A. GOLDFARB (2013): "Some Simple Economics of Crowdfunding," *NBER Working Paper Series*, `http://www.nber.org/papers/w19133`.

——— (2015): "Crowdfunding: Geography, Social Networks, and the Timing of Investment Decisions," *Journal of Economics & Management Strategy*, 24, 253–274.

ANDREONI, J. (2006): "Leadership Giving in Charitable Fund-Raising," *Journal of Public Economic Theory*, 8, 1–22.

ARELLANO, M. AND S. BOND (1991): "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations," *The Review of Economic Studies*, 58, 277–297.

BILGEL, F. AND B. GALLE (2015): "Financial Incentives for Kidney Donation: A Comparative Case Study Using Synthetic Controls," *Journal of Health Economics*, 43, 103–177.

BILLMEIER, A. AND T. NANNICINI (2013): "Assessing Economic Liberalization Episodes: A Synthetic Control Approach," *The Review of Economics and Statistics*, 95, 983–1001.

BLISS, C. AND B. NALEBUFF (1984): "Dragon-Slaying and Ballroom Dancing: The Private Supply of a Public Good," *Journal of Public Economics*, 25, 1–12.

BLUNDELL, R. AND S. BOND (1998): "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models," *Journal of Econometrics*, 87, 115–143.

CADSBY, C. B. AND E. MAYNES (1999): "Voluntary Provision of Threshold Public Goods with Continuous Contributions: Experimental Evidence," *Journal of Public Economics*, 71, 53–73.

CAVALLO, E., S. GALIANI, I. NOY, AND J. PANTANO (2013): "Catastrophic Natural Disasters and Economic Growth," *The Review of Economics and Statistics*, 95, 1549–1561.

DOUDCHENKO, N. AND G. W. IMBENS (2016): "Balancing, Regression, Difference-in-differences, and Synthetic Control Methods: a Synthesis," *NBER Working Paper 22791*, 1–33.

FERMAN, B. AND C. PINTO (2017): "Revisiting the Synthetic Control Estimator," *Working Paper*, 1–45.

FIRPO, S. AND V. POSSEBOM (2016): "Synthetic Control Estimator: A Generalized Inference Procedure and Confidence Sets," *Working Paper*, 1–45.

GRADSTEIN, M. (1992): "Time Dynamics and Incomplete Information in the Private Provision of Public Goods," *Journal of Political Economy*, 100, 581–597.

KAUL, A., S. KLÖSSER, G. PFEIFER, AND M. SCHIELER (2016): "Synthetic Control Methods: Never Use All Pre-Intervention Outcomes as Economic Predictors," *Working Paper*.

KUPPUSWAMY, V. AND B. L. BAYUS (2017): "Does My Contribution to Your Crowdfunding Project Matter?" *Journal of Business Venturing*, 32, 72–89.

LI, Z. AND J. A. DUAN (2016): "Network Externalities in Collaborative Consumption: Theory, Experiment, and Empirical Investigation of Crowdfunding," *Working Paper*, Available at SSRN: https://ssrn.com/abstract=2506352.

LIST, J. A. AND D. LUCKING-REILEY (2002): "The Effects of Seed Money and Refunds on Charitable Giving: Experimental Evidence from a University Capital Campaign," *Journal of Political Economy*, 110, 215–233.

LIU, Y. (2006): "Word of Mouth for Movies: Its Dynamics and Impact on Box Office Revenue," *Journal of Marketing*, 70, 74–89.

MAKRIS, M. (2009): "Private Provision of Discrete Public Goods," *Games and Economic Behavior*, 67, 292–299.

MARX, L. A. AND S. A. MATTHEWS (2000): "Dynamic Voluntary Contribution to a Public Project," *Review of Economic Studies*, 67, 327–358.

MOLLICK, E. (2014): "The Dynamics of Crowdfunding: An Exploratory Study," *Journal of Business Venturing*, 29, 1–16.

PALFREY, T. R. AND H. ROSENTHAL (1984): "Participation and the Provision of Discrete Public Goods: A Strategic Analysis," *Journal of Public Economics*, 24, 171–193.

QIU, C. (2013): "Issues in Crowdfunding: Theoretical and Empirical Investigation on Kickstarter," *Working Paper*, Available at SSRN: http://ssrn.com/abstract=2345872.

VESTERLUND, L. (2003): "The Informational Value of Sequential Fundraising," *Journal of Public Economics*, 87, 627–657.

WINDMEIJER, F. (2005): "A Finite Sample Correction for the Variance of Linear Efficient Two-Step GMM Estimators," *Journal of Econometrics*, 126, 25–51.

WOOLDRIDGE, J. M. (1997): "Multiplicative Panel Data Models without the Strict Exogeneity Assumption," *Econometric Theory*, 13, 667–678.