June 2021

Reference-Dependent Choice on Digital Platforms

Joshua K. Kaisen
University of South Florida

Reference-Dependent Choice on Digital Platforms

by

Joshua K. Kaisen

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy in Economics
Department of Economics
College of Arts and Sciences
University of South Florida

Co-Major Professor: Bradley Kamp, Ph.D.
Co-Major Professor: Haiyan Liu, Ph.D.
Lu Lu, Ph.D.
Xin Jin, Ph.D.
Daniel Zantedeschi, Ph.D.
Andrei Barbos, Ph.D.

Date of Approval:
June 23, 2021

Keywords: Behavioral Economics, Income Targeting, Context Effects, Communication, Prosocial Behavior

DEDICATION

My path has been non-linear to say the least. This was due in part to circumstance, but mostly because my decisions are often as irrational as those of the individuals I study. Despite significant success, I chose to end my life as a physicist after growing disillusioned with the deification of science I lent on. I inherited much from my former life, particularly, the saying "On the Shoulders of Giants" continued to hold significance. The saying expresses the notion that the advancement of society comes only through the masterful work and dedication of inspired individuals who build atop one another. My own personal advancement was only possible through the inspiration provided by the work and dedication of these true masters which offered guidance in dark times. The seminal series of David Simon–*The Wire*–following an obsessive detective bent on achieving self-validation at any cost, and the films of Miloš Forman–*Amadeus* and *Man on the Moon*–depicting the lives of misguided artists, demonstrated what can be accomplished when one gives themselves over to passion. And the faculty at the University of South Florida (USF) provided a practicable example.

I dedicate this work to my committee–Andrei Barbos, Brad Kamp, Daniel Zantedeschi, Lu Lu, Haiyan Liu, and Xin Jin–the spectacular faculty who provided much appreciated guidance–Alfonso Sanchez-Penalver, Giula La Mattina, Mike Loewy, and Padmaja Ayyagari–and the department's unparalleled support staff–Diana Reese and Shay Ferrell.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Throughout this dissertation I explore a significant departure from the standard model of rationality known as reference-dependence. The theory reference-dependence asserts that an individual's choice is dependent on their frame of reference established through factors divorced from their rational cost and benefit. This behavior is inefficient as individuals fail to rationally optimize their payoffs. This behavior is understudied in natural settings where isolating specific stimuli which may establish a reference is challenging. However, digital platforms contain limited stimuli which are fully observable to researchers and present an ideal setting to study the theory of reference-dependence. Thus, I introduce three new datasets of human behavior in natural settings on digital platforms to show that references direct choice in scenarios where it is irrational for them to do so.

My first chapter explores workers who establish an income goal as reference for labor decisions and reduce production when this goal is achieved despite wages being unusually high. The worker perceives a loss while beneath their income goal which drives production such that once they achieve their goal they relax. However, the income goal achievement is accompanied by a wage increase wherein reducing production is highly inefficient intertemporally. This behavior is well-studied in a variety of industries; however, none possess a sufficient income nor wage shock to cleanly identify it. In this chapter I introduce a new setting, YouTube, where workers exhibit this reference-dependent behavior in response to an exogenous income and wage shock, a viral video.

My second chapter explores consumers who use neighboring menu options as reference for purchase decisions and are more likely to buy when the neighboring option induces a favorable comparison for the focal option. This leads consumers to prefer menu options with close price neighbors independent of the option's price and quality, converse to the standard model's predictions. This preference is not associated with the actual benefit of the product and leads to inefficient selection in our market of study, reward-based crowdfunding. Reward-based crowdfunding launches market innovations through consumer pre-purchase and philanthropic donations. This investment scheme outperforms traditional investment in terms of efficacy and demographical balance given rational choices by the consumers. When this rational choice is not upheld the efficacy crumbles and so it key to policy makers to understand the underlying mechanisms which inform choice on such platforms. I introduce evidence of this reference-dependent behavior in the decisions of consumers and further demonstrate that sellers learned to exploit it.

Finally, my third chapter explores beneficiaries who communicate references to their donors by setting goals which promise a prosocial benefit but lack an enforcement mechanism mandating the fulfillment of such benefits. Donors respond positively to this communication and contribute unconditional monetary support towards these prosocial goals despite significant incentives to deceive. This behavior is enabled through the existence of a penalty that the donor may impose provided the beneficiary deceives them. I construct a simple theory to envelop this behavior and derive policy guidelines to deal with such scenarios in general settings.

CHAPTER ONE.

POST-VIRAL DECLINE: EXPLORING PRODUCTION DECLINES DUE TO INCOME TARETING ON YOUTUBE

1.1 Introduction

Absent significant income effects, the standard model of rationality implies a positive wage elasticity. However, an extensive body of literature, starting with the seminal paper of Camerer et al., 1997 presented empirical evidence of instances of apparent negative wage elasticity. The common explanation offered for this failure of the standard model is that these workers may adopt income targeting. Income targeting is equivalent to a discontinuous drop in the marginal utility of income at some target income, which induces some workers to stop working precisely when their target income is met. For such workers, as their wage unexpectedly increases, their target income–which was formed with their expected wage in mind–is attained with less than ordinary work hours leading to a locally downward sloping labor supply when wage is exogenous and volatile.

In most work environments unexpected wage fluctuations are relatively rare and/or the workers have a limited scope for adjusting their labor supply in the short run, thus real-world evidence of such negative wage elasticities is rare. Most of the current empirical evidence comes from the New York City taxi industry where there is significant volatility in hourly wages and the cab drivers can easily adjust the amount of labor supplied in response to these wage fluctuations.

This paper presents evidence suggestive of negative wage elasticity from a different environment, the production of videos on YouTube by independent content creators (channels).

As a commercial environment, YouTube is a three-sided market between consumers (viewers), advertisers, and channels. The advertising is managed by a mostly automated system that fits advertisements to videos based on the demographics. When an ad is viewed, the advertiser pays YouTube who then splits the revenue with the channel. There are substantial search costs to viewers who desire specific heterogenous content. These costs are partially curbed by YouTube's subscriber feature–which allows viewers to subscribe to channels and be alerted when new content is uploaded by that channel–and by YouTube's algorithm for promoting videos on viewers' home page or side panel–which prioritizes uploads from recently viewed channels or on similar topics to recently viewed videos. These features are a key determinant of the wage fluctuations that we exploit in our study.

Using data from YouTube.com following 1,000 individually owned and operated channels over 6 months, this paper sheds light on the labor decisions of workers in response to a unique source of variation in wages, the viral video. A channel's viewership, which determines earnings, typically maintains a steady pattern, but if one of the channel's videos goes viral, their viewership proliferates through the features described above. Although the viewership spike driven by views on the viral video is typically only temporary, a fraction of the new viewers also subscribe to the channel which causes views to spill over into other videos uploaded by the channel and subsist to some degree over the lifetime of the channel, generating a permanent wage increase for the channel.[1] Moreover, even without subscribing, viewers of viral videos are more likely to see new

[1] We refer to wages as the monetary benefits for the additional effort exerted by a channel to create new videos over a period of time; these are different than the total earnings over a period, which are partially generated by viewership of videos previously created.

videos from that channel due to YouTube's algorithm which promotes videos from channels recently viewed or simply because these viewers may browse through other videos uploaded by a channel whose video they recently watched. As we also observe in our data, any additional video uploaded by the channel during a period following a video going viral is much more likely to be viewed and generate income. This implies a sharp but temporary increase in wages for a content creator following one of its videos going viral in addition to a smaller, but permanent, increase induced by new subscribers. Next, we discuss the theoretical implications of these wage dynamics on the short- and long-term productivity of the creator of a viral video.

According to the standard model of rationality, the long-term effect on a channel's productivity would be determined by their trade-off between the income effects generated by a permanent increase in income–which would depress productivity–and higher wages–which would increase it–with an ambiguous net effect. Importantly, the same model would predict a positive spike in productivity in the short-term since wages are significantly higher in the short-term both relative to the pre-treatment and the long-term levels. Our empirical analysis, which measures the weekly productivity of viral channels' creators over a period of 12 weeks centered at the time when the video went viral, reveals different short-term productivity dynamics. Specifically, productivity falls sharply in the week immediately following a channel's video becoming viral, and then gradually returns to the pre-treatment average over a period of several weeks.

A leading possible explanation for these observed productivity dynamics that lies outside of the standard model of rationality is that workers exhibit income targeting. If that is the case, then since the sudden increase in viewership due to the viral video likely pushes the respective periods' earnings above any plausible income target, the marginal benefit of effort may drop following the video going viral despite the higher wages. The channel may thus reduce his or her productivity in the short-term. In the medium-term, decreasing earnings due to the declining viewership of the viral video possibly combined with target adjustment may bring the period earnings below the income target, increasing the marginal benefit of effort and thus increasing productivity. These predictions are consistent with our empirically measured wage dynamics.

The theory underlying income targeting–reference-dependent expectation-based utility–was first proposed in the seminal work of Kahneman and Tversky (1979). Koszegi and Rabin (2006) expanded this theory to reconcile a host of seemingly contradictory results which showed positive wage/earnings elasticity in some contexts (Oettinger 1999; Fehr and Goette 2007) but negative wage/earnings elasticity in others (Carmerer et al. 1997; Farber 2008). Koszegi and Rabin's theory proposed a kink in worker utility around a target for earnings–when workers were below their income target, they felt disutility due to loss aversion which disappeared when passing the target leading discontinuous drop in the marginal utility of income. High anticipated earnings merit an equally high target such that the effects of the kink in the worker's utility is never felt these contexts, while unanticipated increases in earnings lead to a target established assuming ordinary wages being quickly surpassed which create an ideal stopping point for the worker at the kink.

Workers may target any facet of their job; Koszegi and Rabin (2006) proposed income as the primary target but left room for other targets. Crawford and Meng (2011) show empirically that workers target both income and hours worked. However, how workers determine income targets remains an open question. Thus far, the literature has assumed workers derive targets from the status quo of their career such as the average $8/hour earned by cab drivers (Farber 2008 and 2015; Crawford and Meng 2011). We forgo the assumption of a status quo target using large natural variation in earnings–generating income on the level of tens of thousands of USD–that would exceed any income target.

This chapter is broken into 5 sections. Section 1.2 discusses the institution background and relevant mechanisms. Section 1.3 details the data collection process and defines viral videos. Section 1.4 provides the main analysis and section 1.5 concludes.

1.2 YouTube

Launched in February 2005, YouTube was purchased by Google for $1.65 billion in 2006, who implemented a revenue scheme based on advertisements that support many workers with a livable wage highly variant in terms of earnings potential. As of 2021, YouTube is the most trafficked website in the world, with over 1.9 billion active users each month, 50 million independent workers managing their own channels, and 300 hours of videos uploaded every minute. The website created a new industry of entertainment centered on individuals with total creative control.

On YouTube, the revenue system is akin to advertiser-supported TV, albeit with one entity, YouTube, managing a three-sided market between viewers, advertisers, and channels. The advertising is managed by Google's AdSense which is a mostly automated system that fits advertisements to videos based on the demographics desired by the advertiser and the demographics of a channel's audience. When an ad is viewed in its entirety, the advertiser pays YouTube who splits the revenue with the channel (68%).[2] Videos of a larger length will have multiple ads put into their video as determined by Google's AdSense increasing revenue per view

[2] A study performed by InfluencerMarketingHub.com using their clients' earnings from YouTube showed the average earning for channels is $7 per 1000 views, while for smaller channels with less bargaining power such as those from our sample is at about $3 per 1000 views.

for the video; if a video surpasses the ten-minute mark a second ad will be included automatically, this is well-known to channels. In addition to the in-video ad, a banner ad generates revenue per click. These ads generate $2-3 per click but are a smaller source of revenue than the video ad.[3]

YouTube channels are paid for ads placed in their videos only after achieving partnership. In 2018-19 the requirement for applying for YouTube partnership was 4,000 viewer-hours in the past 12 months and 1,000 concurrent channel subscribers.[4] Once monetized, ads are selected by Google AdSense and can come in different forms with different payouts that produce heterogeneous revenue for channels. Channels are paid monthly on a day which is specific to each channel determined based on when they achieved partnership; we do not observe a channel's payday.

When opening YouTube, viewers are presented with a customized home screen recommending a list of videos that may be of interest to the viewer. Additionally, when a video is not watched in full-screen mode, the YouTube webpage also recommends a list of additional videos with corresponding preview images in a side panel. The videos on both lists are selected based on user preferences as determined by YouTube's algorithms. Videos from channels to whom the viewer subscribes or has viewed in the recent past are given priority, as are also videos which may be on similar topics with videos recently watched. Another factor that plays a role in the algorithm, especially for determining videos presented on the home page screen, is the "hotness" of a video, with viral videos being much more likely to be recommended if they are aligned with the viewer's preferences as revealed by their viewing history (The Guardian 2018).

[3] This estimate comes from a personal acquaintance, Ben Buchanan, the owner of HaruHome, who advertises through Google's AdSense.

[4] There is another level to YouTube's partnership program, specifically there are tiers of partnership which is based on number of subscribers. Beginning at Graphite (1-1k subs), then Opal (1k-10k), then Bronze (10k-100k), then enter silver (100k+), post silver the partnership tiers are determined on a case by case basis. These tiers include minor perks and a congratulations notice from YouTube. The effects of crossing partnership thresholds around the Bronze level were explored but yielded no significance and were omitted from this paper.

The key identification in this paper stems from a viral video, which is a video that achieves a large viewership in a short period of time through viewer to viewer sharing, with some videos reaching billions of views. The average viewership of the viral videos in our sample is 2,044,299, which equate to earnings of about $6,133 in a conservative estimate.[5] On top of this significant direct income shock, viral videos bring in views to future videos produced by the channel due to the subscriber feature and YouTube's algorithm for recommending new content to viewers (The Guardian 2018).[6] As discussed earlier, the channel's response to this secondary effect allows us to test in this labor environment the standard model of rationality against a model of income targeting.

1.3 Data

The data used for our analysis was collected by tracking a sample of 1,000 YouTube channels weekly on Mondays at 10am over a period of six months, between November 5th, 2018 and May 6th, 2019 (720,000 individual video-week observations). We collected data at the channel level as well as for the videos uploaded by the channels in our sample.

At video level, we retrieved the weekly total numbers of views, comments, likes and dislikes, as well as static characteristics such as the upload date, the video length, category (determined by the type of content: blogs, travel, or comedy). We then aggregated the video-week observations to the channel-week level to derive the total number of views and other metrics across a channel. The key outcome variables–the number of new videos uploaded by a channel each week and the amount of video-seconds uploaded–serve as a proxy to work hours.

[5] This estimate is based on the $3 per 1,000 views which is determined by the website Influencermarketinghub.com based on private data. This data also suggests that the average earnings per ad watch is $.18.
[6] Also to previous videos, but this effect is much weaker

For each channel, we also tracked the total number of subscribers on a weekly basis and, additionally, several static characteristics: the channel's age, determined by the days since their first upload, and the race and gender of the channel content creator which are collected through a manual review of their publicly posted channel information and collection of videos uploaded.[7]

1.3.1 Sample Selection

To select our sample, we used ChannelCrawler.com which can sort YouTube channels based on a variety of parameters of interest to the user. We chose the 1,000 channels that had the most recent video upload among the channels with 100 to 15,000 subscribers and who met two additional requirements: they were not associated with a business and they belonged to one of four categories, as defined by YouTube: people & blogs, entertainment, films & animation, and comedy. Creating videos in these categories typically requires a meaningful amount of effort, therefore satisfying workers face a trade-off between costs and benefits when deciding how much effort to put into creating these videos.

Out of the 1,000 channels we started out with, we dropped from our analysis a subset based on criteria we describe next. Since the web scrapper we employed can only track the 30 most recent videos of each channel, channels that ever uploaded more than 29 videos in a week were dropped. This ensures that we always correctly measure the weekly productivity for the channels left in our dataset. As the channels in the categories, we selected upload with low frequency (1.94 videos per week), the scope of this censoring is limited, with only 20 channels being dropped. Some channels upload live streams of live video games playing or cooking. The amount of effort required to create these videos is much lower than for the typical videos which would introduce a certain degree of

[7] 89% of channels were identified for gender, while the race was determined for 85% of the channels. Two separate persons performed the manual review.

undesirable heterogeneity. We removed channels with any video that has a length of more than eighty minutes, to address this.[8] 280 channels that had at least one video longer than this cutoff were eliminated from the sample. Some channels also upload copyrighted content such as movies or sports events and they were dropped (15 channels). Viral videos sometimes contain embarrassing or controversial content which inhibits future video uploads. This content is typically characterized by a removal of the like/dislike counter from a video. Channels that ever use this option were dropped (13 channels). Channels which were removed by YouTube or the channel owner during our data collection process, which had a break longer than 200 days between uploads, or for which the data collection was missed in more than 10% of the weeks[9] were also dropped (54 channels). Finally, we removed–as outliers–channels which, due to a very low long-term average channel productivity, recorded an increase in productivity more than 1000% relative to their average productivity when they did upload a video (47 channels, including two channels with viral videos). Our key insights yielded by the estimates of the average productivity dynamics after a video becomes viral are robust to this exclusion, but their graphical representations are smoother. After this sequential selection process, we are left with 572 channels for our analysis. A visual representation of the selection workflow follows.

[8] The cutoff takes into account that live streamed videos tend to have around 2-4 hours, but is otherwise arbitrary. The main results of the paper are not sensitive to the specific cut-off value.

[9] To prevent the inflation of views on videos using bots–programs that pull up a video multiple times with each registering as a view–YouTube monitors URL queries for each IP address that monitors successive access to videos and disallows access after certain thresholds are reached. We used the python module beautifulsoup which does not register as a view but still needs to make a request to access the information of each video's page as a bot would. In some instances, we scrapped a "blank page" showing that access to material was denied and missed that day's observation for some channels. Once this happens attempting to access that channel's content again would retrigger the prevention measure, so we set our code to pause and continue scrapping other channels after such an occurrence.

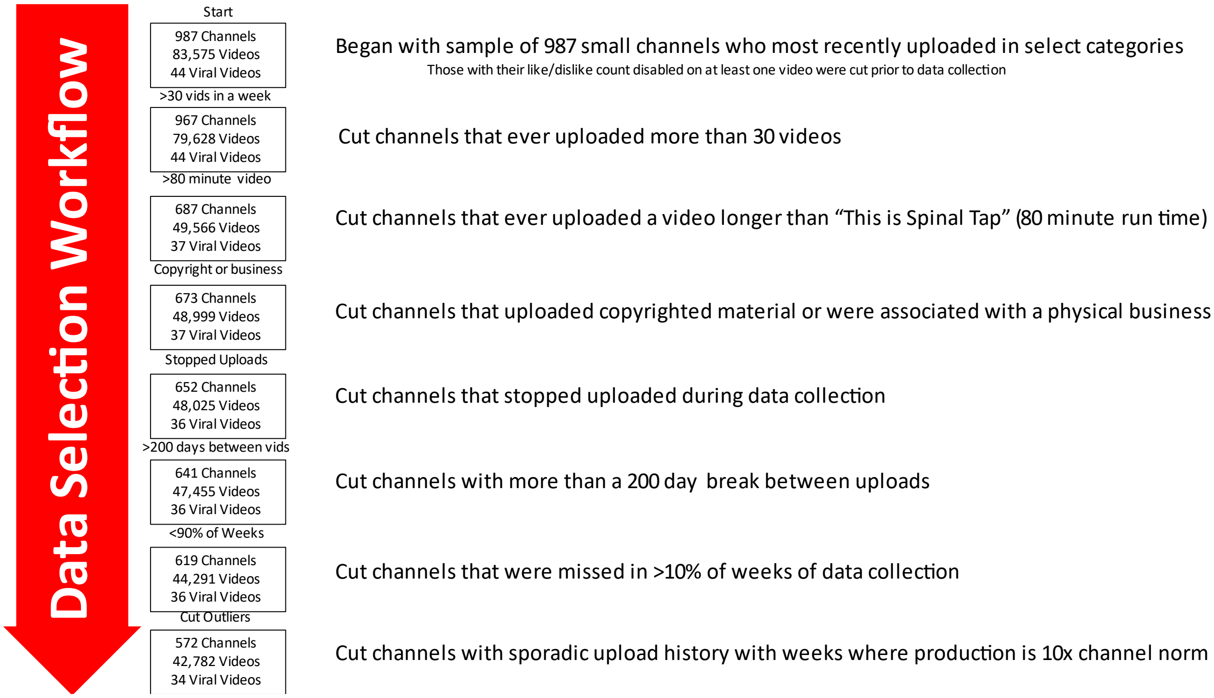

Figure 1.1: Data Selection Workflow

1.3.2 Classification of a Video as Viral

Viral videos provide the exogenous shock to wages that we employ to test the standard model of rationality in this labor environment. Simultaneously, viral videos generate an income shock that plausibly brings an individual’s income over their target if that worker’s labor decisions are as described by the alternative model of income targeting. In the latter case, we would not know the specific targets, the period over which these targets apply, or even which type of metric a channel targets. These do not constitute a problem if we are capturing with our definition of a viral video an event that: (1) increases the viewership of new videos posted by that channel and therefore the wages of the respective content creators, and (2) provides an income shock that is likely to move that individual’s income over any plausible target. To verify that our analysis and conclusion are robust to misspecifications, we will consider different measures and different cutoffs for the definition of a viral video.

We employ two distinct criteria for classifying a video as viral. Videos satisfying either criteria will determine a temporary increase in wages for that channel owner, but they address potentially different types targets that workers may adopt if non-rational. First, channels may target current views on a video since these impact current earnings. Alternatively, they may target the total number views on a video, which determine accumulated earnings. A sudden spike in current views also leads to a jump in total viewership and therefore the sets of viral videos as determined by the two classifications overlap significantly, and the key insights from our analysis are also similar between the two specifications.

In our main analysis, we employ a definition of a viral video that relies on the weekly views. Specifically, a video is defined as viral when it exceeds a threshold of 150,000 new weekly views, but we also considered other thresholds in the range of 10,000-400,000 weekly views with qualitatively similar insights. We call this the “Weekly Views" specification. With this method, we classify 45 videos as viral, belonging to 24 unique channels. Videos go viral an average of 66 days after upload with a median of 32 days.

To capture a situation where a channel may target total views, we also perform an analysis based on the definition of a video as viral once it exceeds a certain total viewership threshold. We use 500,000 as this threshold, but again, we also performed our analysis using different round numbers in the range 10,000-1,000,000 with similar insights. We call this the “Total Views” specification. Only 34 of the 44,290 videos in the sample are classified as viral with this method at an average age of 182 days with a median of 54.5 days. These viral videos were uploaded by 25 unique channels. There are 23 videos and 18 channels classified as viral using both methods.

We prefer the definition based on weekly views for two reasons. First, it is arguably more closely aligned with the common understanding of a viral video as being one whose viewership suddenly spikes. As such, it also makes it more likely to create the conditions under which a certain income target may be suddenly exceeded. On the other hand, a video classified as viral based on the total number of views may potentially gain viewership more gradually and may in theory push through several possible income targets, possibly leading to a decline in productivity over a longer period. The second reason for using the Weekly Views specification is that the pre-treatment trends of the treated and control groups are more similar than when we use the Total Views specification. However, the insights from both specifications are suggestive of income targeting, and so, we also report in the appendix the results obtained based on Total Views.

1.3.3 Descriptive Statistics

Table 1.1 shows the descriptive statistics at the video-level. These statistics are derived based on the characteristics of the videos as of the latest date at which they were observed. The table is split by viral (Treated) and non-viral (Untreated) videos based on the Weekly Views classification.

As expected, the viral videos have significantly higher numbers of views, likes, dislikes, and comments, which are all variables expected to be proportional to the viewership. On the other hand, the viral and non-viral videos are rather similar in terms of the like ratio (Likes/Dislikes)[10], time length, and age at the time when they exit the sample.

[10] The difference in the sample is statistically significant, but when weighting the observations by the number of views of the corresponding videos this difference is no longer statistically significant.

Table 1.1: Video Summary Statistics

Variable	Nonviral Mean	Viral Mean	Difference
Total Views	11,398.26	1,662,263	1,650,865***
	(6,046)	(316,541)	(218,143.9)
Video Length (Seconds)	478.27	573.91	95.64
	(2.39)	(54.63)	(86.13)
Video Age (Days)	113.60	111.26	-2.34
	(0.74)	(14.02)	(26.72)
Likes	156.4139	15,420.74	15,264.32***
	(34.89)	(3,139.80)	(1,260.83)
Dislikes	6.88	1,445.09	1,438.21***
	(2.02)	(338.20)	(73.28)
Comments	46.23	2,532.79	2,486.57***
	(2.07)	(737.14)	(77.23)
Like Ratio	0.950	0.886	-0.064***
	(0.0001)	(0.015)	(0.019)
Videos	44,195	34	44,229

Note. Standard errors are reported in parentheses. Video Age captures the age of the video when it exits the sample, which may happen either when that channel has 30 newer videos or when our data collection ends. Like ratio is by default 1 when more views are added there is more opportunity to deviate from this default value.

The channel-level statistics shown in Table 1.2 are recorded in the second week of our observation period before any of the channels had a viral video for our sample post-selection.[11] In the appendix, we also include a table with the summary statistics obtained by aggregating the averages for channels across the entire observation period. Treated channels have more weekly views and subscribers even before experiencing a viral video and are thus not identical to the control channels. However, the two types of channels are statistically indistinguishable in terms of the key variables that measure channel productivity, i.e., their weekly uploads and a related variable that measures the average time between uploads. The two types of channels are also

[11] Family run channels feature couples or families presenting videos of themselves. Only channels with race mutually identified several different persons are reported in the statistics.

mostly similar in terms of demographic characteristics, with the exception that content creators of viral videos are more likely to be female.

Table 1.2: Channel Summary Statistics

Variable	Nonviral Mean	Viral Mean	Difference
Weekly Uploads	2.37	3.10	0.72
	(0.14)	(0.82)	(0.72)
Time Between Uploads	6.03	5.19	-0.84
	(0.23)	(1.14)	(1.18)
Weekly Views	6,287.80	105,006.20	98,718.44***
	(992.53)	(52,126.19)	(11,174.22)
Subscribers	3,841.38	7,515.81	3,674.43***
	(127.35)	(897.49)	(674.96)
Channel Age	1,243.80	559	-684.80**
	(41.65)	(120.70)	(214.78)
Family Channel	0.053	0.143	0.0902
	(0.010)	(0.078)	(0.051)
Female	0.281	0.429	0.147
	(0.019)	(0.111)	(0.101)
Male	0.568	.286	-0.282*
	(0.021)	(0.101)	(0.110)
Black	.192	0.095	-0.097
	(0.017)	(0.065)	(0.087)
Asian	0.058	0.048	-0.010
	(0.010)	(0.048)	(0.052)
White	0.559	0.619	0.060
	(0.021)	(0.109)	(0.110)
Hispanic	0.049	0.095	0.046
	(0.009)	(0.066)	(0.049)
Channels	551	21	572

Standard errors reported in paratheses.
Note. All measures are calculated by the value at the beginning of data collection (in the second week to establish variables that need a reference week such as weekly uploads).

1.3.4 Trends in Viewership

The wage received by the content creator in a given period, which determines the level of effort in the standard model of rationality, is determined by the lifetime views recorded by the new videos created within that period (the wage is different than the earnings in a period, which are the monetary payments for the views recorded by all videos of that channel in that period). In this section, we describe how videos accumulate views over time to underlie the difference in viewership trends between viral and non-viral videos, and more importantly to also demonstrate the effect of a viral video on the viewership of the other videos uploaded by the same channel,[12] and thus on the wage dynamics faced by the content creator around the time a video goes viral.

For the non-viral videos uploaded by the non-viral channels in our sample, on average about 80% of the views are recorded within the first 3 weeks, likely due largely to the YouTube algorithm which tends to promote recent videos of channels that a user subscribed or watched videos of in the past. By week 9 it is expected that these non-viral videos receive 95% of their lifetime views.

On the other hand, viral videos have a different pattern of viewership which is depicted in Figure 1.2. The figure plots the weekly percentages of the total viewership of a viral video normalized around the week when the videos met for the first time the condition for being labeled as viral, using the definition based on the Weekly Views measure. By definition, viral videos experience a spike in the week when they go viral, which, as discussed earlier, can happen well into their lifespan. The viewership distribution for these videos takes an almost Gaussian shape, with a peak of about 30% of total views experienced during the week when they go viral. The increase in viewership begins about 3 weeks prior, followed by an exponential ascent generated by audience driven sharing. Viewership returns to normal 2-3 weeks after the spike week. This pattern is consistent with anecdotal accounts which describe viral videos as starting with a smaller

[12] The data are right side censored due to the fact that 30 video observation limit per channel and the six month observation period of channels and videos does not allow us to observe a video during its complete lifetime. However, we do observe all these non-viral videos during the first few weeks when most views are recorded.

audience who recommends them to friends, who recommend them to their friends, and so on until an explosion in viewership occurs (The Guardian 2012 and 2018).

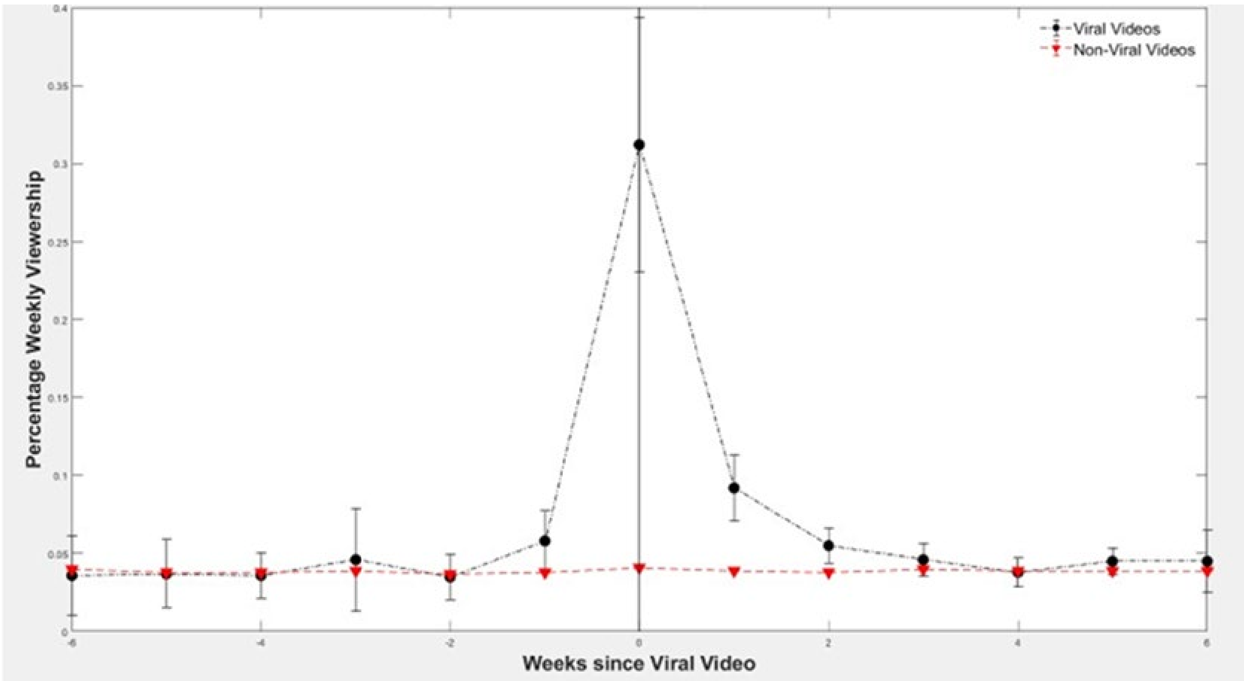

Figure 1.2: Viewership Trends of Viral Videos

Note. 95% confidence interval shown with error bars. Channels that go viral, achieve the majority of their viewership in the week of going viral as opposed to channels that never go viral which achieve the majority of their viewership in the first two weeks of life.

To elicit the wage dynamics faced by viral channels around the time when their videos go viral, Figure 1.3 describes the viewership patterns of the non-viral videos uploaded by these channels. In panel (a), the curve labelled with squares depicts the weekly views per video for the new non-viral videos uploaded during each week, where for each viral channel, the events are centered on the week when their video went viral (for instance, the value for week -3 would capture the average number of views per video recorded during week -3 for all non-viral videos uploaded during week -3). The curve labelled with triangles shows the weekly views per video for the old non-viral videos - these are all non-viral videos of the viral channels uploaded before that particular week. With a dotted flat line, we mark the average number of views per new video in the weeks 4-6 before a channel had a viral video, as a baseline for the number of weekly views of a new video before there were likely any spillover effects from the viral video. Starting from that baseline, the figure shows a stark increase in viewership over the next few weeks, which peaks during the week when the video of that channel went viral and continues at elevated values over the baseline for several weeks afterwards. Therefore, any video released by a viral channel during and immediately after their video went viral receives much more views during the week when they are uploaded than the videos of the same channel were normally receiving several weeks before their video went viral. The magnitude of this upward jump is significant with a peak increase in viewership in week 0 which is 5 times higher than the baseline. In week 1, when our subsequent regression analysis of the productivity dynamics will show a sharp decrease in video creation, the wages are still 4 times over the baseline. The similarly elevated levels of viewership for the old videos during the weeks after the video went viral suggest that this increase is persistent. Videos uploaded around the time when the video goes viral thus enjoy a significantly higher number of views over their lifetimes.

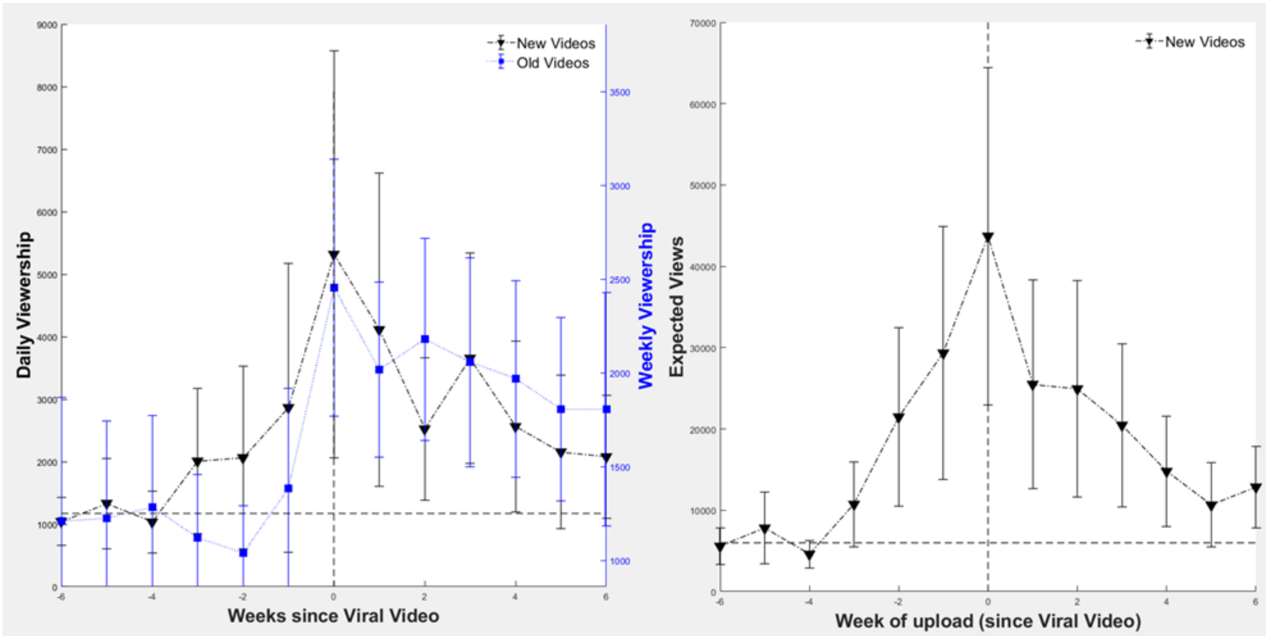

Figure 1.3: The Spillover Effect of Viral Videos

Note. 95% confidence interval shown with error bars. On the left (daily/weekly) viewership as a function of week since a channel went viral is plotted for videos (less than one week old/older than one week). On the right expected viewership (views after 3 weeks) for videos uploaded in each respective week around a channel going viral is shown.

To alleviate concerns that the insights from panel (a) may be driven by an intertemporal substitution of viewership towards the week when the videos are uploaded, in panel (b) of figure 1.3 we depict for each week, the total number of views per video of the non-viral videos uploaded in a given week, recorded during the week of the upload and the following two weeks (for instance, the value at week -3 would be the average weekly views of non-viral videos of the viral channels released in week -3, as recorded over the weeks -3, -2 and -1). Again, for each viral channel events are centered on the week when their video went viral. Because of the spillover effects, some of the videos released early will record more views than would be suggested by this figure. However, for the baseline weeks, since without another video of a channel becoming viral, typically most views (about 80%) of a video are recorded during the first 3 weeks after upload, the values in this figure provide a measure of the *expected* number of views of a new. Again, in this figure, the most informative comparison is between the baseline of weeks -6, -5, and -4, and the values immediately after the video going viral, with the figure suggesting a 4-5 times increase in expected viewership.

Since the wage a content creator expects to receive for effort exerted during a given week is proportional to the expected lifetime views of videos created in that week, the two panels in Figure 1.3 suggest the wage dynamics facing the content creator around the time when their video went viral. Specifically, they imply a 4 to 5 times increase in these wages compared to the weeks before the impact of the viral video. In the next section, we analyze the productivity dynamics for viral channels to test predictions of competing models about the relationship between wages and effort.

1.4 Analysis

The standard model of rationality yields the following predictions for productivity dynamics of a channel around the time one of its videos goes viral: (1) The productivity of the channel increases immediately following a video going viral. (2) After the initial increase in productivity, productivity gradually decreases.

On the other hand, a model of channels exhibiting income targeting would have the following implications on the dynamics of channel productivity: (1) The productivity of a channel decreases immediately following a video going viral. (2) After the initial drop in productivity following a viral video, productivity increases.

Our goal is to distinguish between the two theories by means of empirically testing their implications. Towards this end, as preliminary non-parametric evidence, we first plot in Fgure 1.4 weekly uploads in a 3-month window around each viral video and compare viral channels to non-viral channels. To account for the impact of differences in production scale between channels in discussion of results, where a video for one channel may incur higher effort than a video of another channel,[13] instead of using raw productivity numbers, we use as a measure of channel's weekly productivity, the number of uploads during that week divided by the average weekly uploads across our entire 6-month data-collection period. We refer to this measure as the channel's relative productivity.

The events are centered around the week when the video went viral, week 0. For non-viral channels, we computed the values represented in the figure using a method like the one employed in Figure 1.2. As an example, if videos went viral across our dataset in weeks 8, 12, and 17, then to compute the relative productivity value of the non-viral channels for one week after week 0, we computed the average relative productivities of the non-viral videos in weeks 9, 13, and 18. On the other hand, for the viral videos, the relative productivity for the same week is computed as the average of the relative productivities in week 9 of the channel whose video went viral in week 8, week 13, for the channel whose video went viral in week 12, and week 18 for the remaining viral channel.

Figure 1.4 shows a significant decline in production for the channels who had a viral video one week after their video went viral. The decline in production appears sustained in this simple figure. To understand the dynamics further we introduce a regression with controls for relevant factors affecting production beside the viral video.

[13] For example, a 1 video decline is significant for a channel that uploads 2 videos per week on average, but much less so for a channel that normally uploads 10 videos. To allow for more direct discussion of the impact of viral videos on productivity we scale our productivity measure to be mean 1 for all channels and measure percentage declines in production.

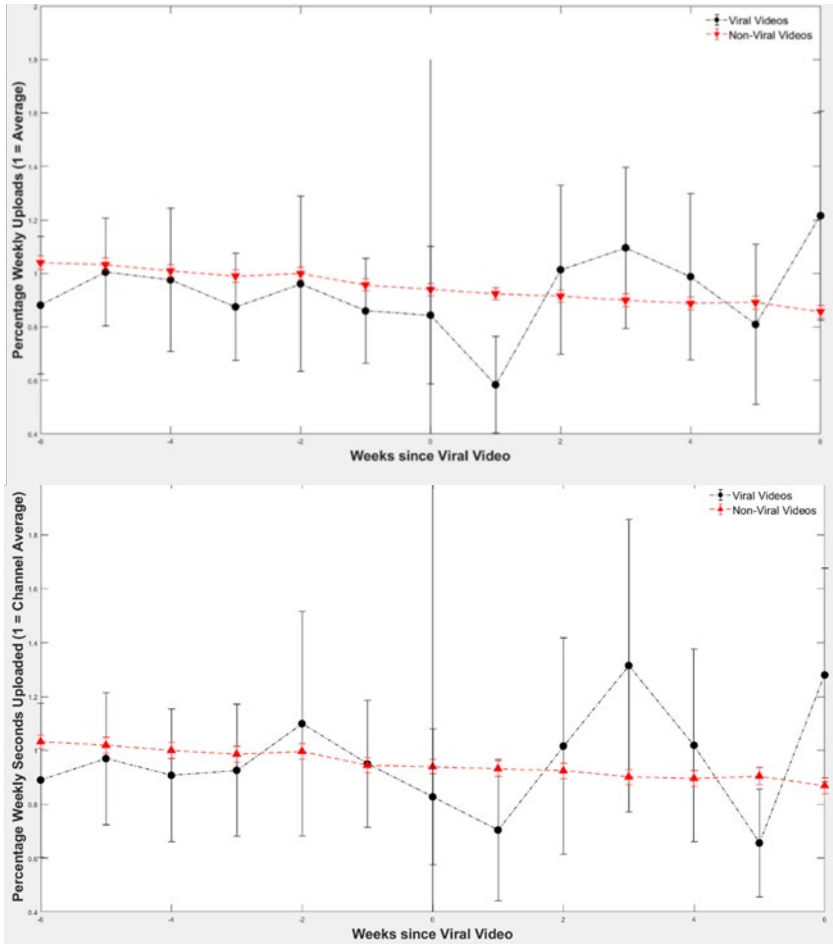

Figure 1.4: Model-Free Evidence of Income Targeting

Note. 95% confidence interval shown with error bars. A raw comparison of means of production around viral event is shown for channels that go viral in week 0 and non-viral channels. All values are scaled such that 1 represents a channel's average weekly production for tractability. There is a persistent downward trend because we condition our sample selection on activity in the first week observed but not for any week after.

We regress two productivity measures V_{tc}–weekly uploads or weekly video-seconds uploaded–on the timing of a viral video and controls. The impact of a viral video on productivity of a channel *c* depends on the temporal distance between the current week, *t*, and the week when the video of that channel went viral, which we denote by E_c. We use week level fixed effects, b_t and channel level fixed effects a_c. Stochastic controls, X_{tc}, are the number of subscribers and subscribers squared for a channel. These control for anticipated gains in earnings which increase production (Oettinger 1999; Fehr and Goette 2007). We design our dynamic event study as recommended by Sun and Abraham (2020) and exclude terms for week 0 and weeks <*-6* and >*6* to avoid multicollinearity.[14]

$$V_{tc} = \sum_{k=-6}^{-1}\left(\beta_k 1_{t=E_c+k}\right) + \sum_{k=1}^{6}\left(\beta_k 1_{t=E_c+k}\right) + \delta X_{tc} + a_c + b_t + \epsilon_{tc} \tag{1}$$

This measures the week-by-week effect of a viral video on a channel's production. The window is at 6 weeks before and after the treatment time since impacts largely dissipate beyond this point. The \beta is the effect net of anticipated increase in earnings of videos which has a positive impact and of the income target being met which has a negative impact.

Figure 1.5 shows the results of regression 1 graphically. The vertical axis is the \beta corresponding to week since a video went viral. The figure reveals a significant decline in production one week after a viral video for both productivity measures. This isolates the behavioral response to a viral video along each stage of the event. With controls in place, the pre-trend appears standard. Channels have a slow return to normalcy in the month following the viral video.

[14] Using the robust estimation recommended by Sun and Abraham yields the same results as a direct two-way fixed effects estimate so we only report the direct two-way fixed effect estimates. This is no surprise as the productivity decline is also demonstrated by a raw comparison of means for both productivity measures.

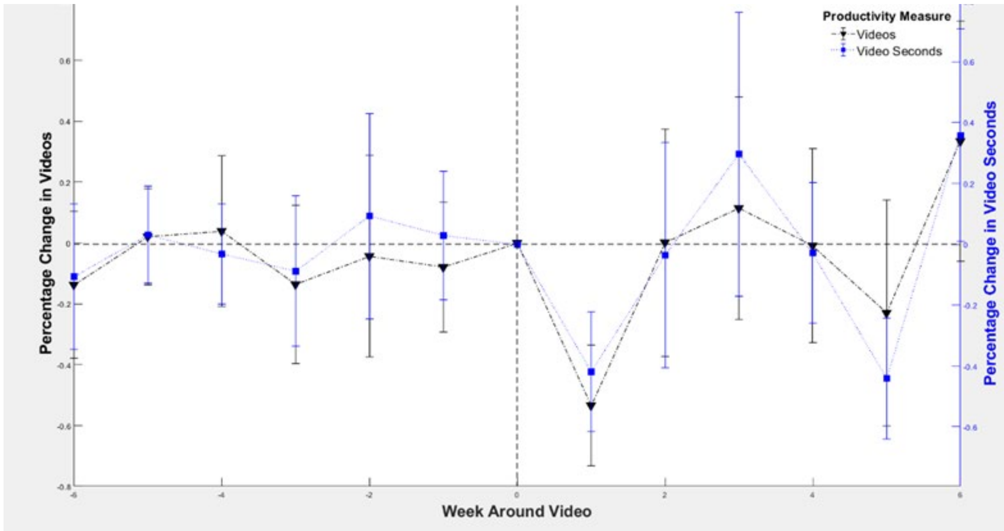

Figure 1.5: Evidence of Income Targeting

Note. 95% confidence interval shown with error bars. We regress with model 1 for two measures of productivity. The percentage change (by own channel's average) in videos uploaded (triangles and left axis) specification shows a significant 50% decline in production for the week after a channel goes viral. The percentage change in video seconds uploaded (squares and right axis) specification shows a smaller but still significant 40% decline in production.

Finally, we establish sensitivity bounds for the observed one-week production decline in our definition used for a viral video. Table 1.3 presents the value of β_1 for each definition of viral videos by production measure. We start with an overtly low definition of a viral video to differentiate between the income effects generated by a general increase in income predicted by the standard model of rationality and a discontinuous negative jump in one's marginal utility curve at some substantial income target predicted by the theory of income targeting. If our decline displayed in Figures 1.4 and 1.5 is due to income effects, we should see similar behavior when we institute a simple cut on small levels of viewership which separate channels by average variation around high and low performing weeks. However, we see no evidence of a decline for these small values which supports the claim of the decline measured in Figures 1.4 and 1.5 evidencing income targeting.

Table 1.3: Sensitivity Bounds in One-Week Productivity Decline

	1 week drop	
Viral Threshold	Videos	Video Seconds
Weekly Views>10k	0.0776	0.137
301 vids/119 chan	(0.0838)	(0.0893)
Weekly Views>25k	0.0445	0.142
186 vids/69 chan	(0.0872)	(0.115)
Weekly Views>50k	-0.150*	-0.0400
95 vids/45 chan	(0.0781)	(0.122)
Weekly Views>75k	-0.0900	-0.0499
64 vids/33 chan	(0.111)	(0.148)
Weekly Views>100k	-0.307***	-0.240*
59 vids/29 chan	(0.119)	(0.134)
Weekly Views>125k	-0.463***	-0.361**
52 vids/25 chan	(0.103)	(0.147)
Weekly Views>150k	-0.534***	-0.419***
41 vids/21 chan	(0.101)	(0.120)
Weekly Views>200k	-0.230*	-0.0577
28 vids/15 chan	(0.123)	(0.149)
Weekly Views>250k	-0.374***	-0.0675
22 vids/13 chan	(0.124)	(0.176)
Weekly Views>300k	-0.419***	-0.112
16 vids/12 chan	(0.147)	(0.294)
Weekly Views>350k	-0.326**	0.0528
14 vids/11 chan	(0.157)	(0.311)
Weekly Views>400k	-0.407**	0.0259
12 vids/10 chan	(0.168)	(0.357)

Standard errors in parentheses, clustered by channel.
Note. The above table presents a sensitivity analysis of β_1 from regression (1) by viral specification from a range of 10k to 400k weekly views.
* p<0.10, ** p<0.05, *** p<0.01

As we proceed down the columns of our sensitivity analysis, we increase the threshold for definition of viral videos and exclude more and more channels from the treated sample. This increases the positive wage shock but reduces the statistical power of our estimation which leads to a fallout in significance for the productivity decline in video-seconds but the productivity decline in videos remains negative and statistically significant throughout out sensitivity bounds (above 100k).

1.5 Conclusion

Jobs have rapidly begun an exodus from the salary worker in congregate offices to contract labor, ever so often pursued from a home office. The labor decisions of these workers are influenced by a variety of phenomena. They choose their working hours and are not under direct supervision. Many sources promote the use short-term goals to enhance production in these settings. The worker believes that they benefit from this goal but the theory of reference-dependent expectation-based utility claims otherwise. In this paper, we study workers with an income goal under volatility in wages. When wages are unexpectedly high, they work less and when wages are low, they work more due to the influence of set income goals. This behavior is intertemporally inefficient and should be avoided. We recommend policy makers and employers be aware of the inefficiencies presented by established labor goals and avoid instituting these whenever possible.

We use natural variation in wages–viral videos–which provide strong identification and ultimately reveal evidence of income targeting. This result challenges recent literature which found no evidence of income targeting and dispelled natural sources of variation used in past work in the process (Farber 2015). Our results are robust to modifications in definition and estimator.

Our methodology avoids the need to detail the establishment of targets, but this is a topic of necessary future study. Thus far the income targeting literature has focused on workers with no ability to grow average wages–taxi drivers, stadium vendors, and bike messengers. Previous results may fail to capture the behavior of individuals with the ability to grow in wage, yet still face significant wage uncertainty and control work hours such as entrepreneurs.

1.6 References

Camerer C, Babcock L, and Thaler R (1997) Labor supply of new york city cabdrivers: One day at a time. *Quarterly Journal of Economics* 112(2):407-41

Kahneman D, and Tversky A, (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263-91

Parkin S (2018) The youtube stars heading for burnout: “The most fun job imaginable became deeply bleak” *The Guardian*

Koszegi B, and Rabin M (2006) A model of reference-dependent preferences. *Quarterly Journal of Economics* 121(4):1133-65

Fehr E, and Goette L (2007) Do workers work more if wages are high? Evidence from a randomized field experiment. *American Economic Review* 97:298-317

Farber HS (2008) Reference-dependent preference and labor supply: The case of New York city taxi drivers. *American Economic Review* 98(3):1069-82

Farber HS (2015) Why you can’t find a taxi in the rain and other labor supply lessons from cab drivers. *Quarterly Journal of Economics* 130:1972-2026

Crawford VP, and Meng J (2011) New York city drivers’ labor supply revisited: reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review* 101(3):1912-32

Sun L, and Abraham S (2020) Estimating dynamic treatment effects in event studies with heterogenous treatment effects. *Journal of Econometrics*

Oettinger GS (1999) An empirical analysis of the daily labor supply of stadium vendors. *Journal of Political Economy* 107(2)

CHAPTER TWO.

STUCK IN THE MIDDLE: CONTEXT EFFECTS IN REWARD-BASED CROWDFUNDING MENUS

2.1 Introduction

The canonical model of rational choice provides precise predictions regarding substitutes: the introduction of a good should reduce demand for its substitutes. Moreover, substitute goods should have positive cross-price elasticity. However, existing empirical literature demonstrates that introducing a good may counter-intuitively increase demand for its substitutes. The common explanation is that menu options act as a reference to one another, and consumers are more likely to purchase options with a uniquely strictly dominated comparison (Huber et al. 1982, Ok et al. 2015).

This literature–broadly labeled as *context effects*–directly contradicts the weak axiom of revealed preferences which asserts that if in any choice set $A \subset B$ option y dominates option x, then for any $C \subset B$ where $C \ni \{x, y\}$, y dominates x. Huber et al. (1982) first described this possibility as the principle of asymmetric dominance; the introduction of a favorable comparison for x makes it more attractive to consumers relative to y. The empirical literature was pioneered by Simonson (1989) who showed that introducing a new option z, strictly dominated by x, leads to a reversal of preferences (i.e., x is ultimately preferred to y).

Firms may take advantage of these cognitive biases to maximize profits (Barbos 2010). Although this increases welfare by expanding firms' ability to discriminate (Ok et al. 2014; Carbajal and Ely 2016), markets reliant on rational choice and repeated consumer-seller interactions may see a decline in total welfare. Reward-based crowdfunding (RBC)–where entrepreneurs use pre-purchase capital to introduce new products and entrepreneurial interventions–is one such market. Through RBC, consumers directly decide the importance of new products, which yields an overall positive societal effect. Traditional investments often fail to reflect all demographics' will; gender bias and under-representation of minorities are well-documented in venture capitalism and angel investing (Mollick and Robb 2016). However, the efficacy of RBC is threatened when consumers make decisions based on features irrelevant to the product's value. Further, when sellers become known to manipulate their consumer base, prospective consumers are deterred from the market in general (Soubilere and Gehman 2020; Blaseg et al. 2020).

We introduce a new cross-sectional RBC dataset to investigate three research questions: (1) How do price neighbors affect consumer choice? (2) Do the marketplace policies modify these context effects? (3) Do sellers/buyers learn to exploit/overcome these changes?

We find that demand is higher for a focal option when the price of neighboring options (price neighbors) is closer.[15] A positive cross-price elasticity is present for the lower-priced neighbor (0.157) and a negative cross-price elasticity for the higher-priced price neighbor (-0.123). This contradicts the neoclassical model, which predict a positive cross-price elasticity for the higher-priced neighbor. The likely driver is the *compromise effect*–when an option z of higher price but only slightly higher quality than option x institutes a positive comparison which increases

[15] Our menu is always ordered from lowest-priced to highest-priced option.

preference for x an outside option–in this case, other products on the platform (Huber et al. 1982; Simonson 1989). We also find that framing an option as a sale increases demand by 78%. This behavior is due to a subset of the compromise effect–*price anchoring*–where z and x have the same quality, but z has a higher price. Our key findings are that the introduction of direct social learning tools for sellers generated an increase in the potency of price anchoring and diminishes other context effects. Finally, we find evidence of limited consumer attention–every option deeper in a menu has a subsequent drop of 9.42% in quantity demanded–which drives the long-term decline of context effects other than price anchoring as new features are added to menus and sales become more prevalent.

The application of context effects to real markets has recently come under question (Fredrick et al. 2014; Yang and Lynn 2014). Context effects are broadly explored in lab experiments with questionnaires and other hypothetical incentives. Among these experiments, context effects are almost exclusively observed in response to numerical stimuli, such as price or quantity.[16] In natural environments, context effects are studied via the addition of a new highest-priced or lowest-priced menu option (Doyle et al. 1999; Geyskens et al. 2010; Pinger et al. 2016; Wu and Cosguner 2018). Though these experiments directly identify context effects, they are not broadly applicable to real markets where dominated options rarely exist due to low profitability (Fredrick et al. 2014; Yang and Lynn 2014; Huber et al. 2014).[17] We discover context effects in a natural setting within menus (rather than isolated at the extrema) and document the decline of context effects as stimuli move from numerical (price) to text to extend this literature.

[16] See Huber et al. (2014) for a brief overview.
[17] However, price anchoring is a significant subset of the context effects literature which widely applied (Bonini and Rumiati 2002; Marzilli Ericson and Fuster 2011).

The impact of context effects is intertwined with the limited consumer attention. Tserenjigmid (2019) and Maltz and Rachmilevitch (2021) independently construct models where an increased number of references diminishes the potency of each and are supported empirically (Whyte and Sebenius 1997; Lichters et al. 2017). Liao et al. (2020) show farther priced menu options impose less potent context effects. RBC provides a natural setting to explore these interactions of limited cognition and context effects.

Many lab experiments reveal context effects in (hypothetical) RBC schemes. Weinmann et al. (2021) show the attraction effect–the contrapositive of the compromise effect–causes an increase in demand when low-priced "decoy" options are added to menus. Adam et al. (2019) show that sold-out on-sale options increase demand for available not on-sale options.

In general, the RBC literature is growing. The primary benefit of the market is the birth of innovations, which depends on the RBC venture's success. The literature investigating the determinants of RBC success throughout the campaign converges towards four classes of drivers: (a) product-related (description of the product, presence of videos), (b) customer-oriented related (number of updates, sentiment of the reply to comments/questions), (c) campaign specifics (funding goal amount, duration of the campaign, number of reward tiers) and (d) entrepreneurial demographics (previous successes, gender, and location). See Shneor and Vik (2020) for a comprehensive survey. It is important to notice that these drivers appear to have a differential effect through the campaign (Petitjean 2018).[18]

This paper uses data retrieved from Kickstarter, the most prominent reward-based crowdfunding platform for new product development. The data span from 2013 to 2020 and

[18] Beyond the campaign less is known. Viotta da Cruz (2018) investigates campaign features that can act as credible signals of the eventual product release: she finds that sentiment in comments, quality of spelling, number of backers, and social media likes are significant predictors for the product delivery. More generally, Moe and Fader (2002) find that pre-purchases forecast post-launch sales.

constitute the entire universe of 136,738 Kickstarter ventures created in the United States during this time. We use a flexible mixed-effects specification for analysis and use multiple partitions of the data to demonstrate the findings' robustness.

This chapter comprises seven sections. Section 2.2 introduces Kickstarter, its relevant features, and its policy changes throughout our data. Section 2.3 presents the data and section 4 the methodology. Section 2.5 and section 2.6 present the static and dynamic analysis, respectively. Section 2.7 concludes and provides direction for future research.

2.2 Crowd(fund) Surfing

2.2.1 Kickstarter

Kickstarter is a prominent platform where entrepreneurs can raise capital for new products through the pre-purchase of the product by buyers (consumers) and purely philanthropic donations by donors. The creator (entrepreneur) begins a Kickstarter project (new product) by setting a capital goal that must be met through buyers and donors (together called backers) and the time given to achieve said goal. If met, a contract between the entrepreneur and their backers is created, which, since October 2014, states that the creator must deliver or refund purchased rewards. Previously, projects were only contractually bound to deliver on promises made explicitly in their campaign, which did not always include delivery or the refund of pre-purchased rewards. If the goal is not met, the project fails, and backers' funds are returned in their entirety.

On Kickstarter, projects may only involve new products which do not exist in the market.[19] These products lack reputation, so backers lack prior knowledge of the project itself. However, pre-existing relationships between backers and creators (e.g., family and friends) drive surges in

[19] Projects seeking investment or purely donations are not allowed. See Kickstarter's terms of use for details.

donations on the first day of campaigns which may generate cascading behavior (Vismara 2016).[20] The prosocial behavior of this group of backers may limit the external validity of studies on Kickstarter. However, prosocial behavior is predominantly directed towards the achievement of projects' goals and there is little donor activity post-goal achievement (Dai and Zheng 2019). Overall, Belleflamme et al. (2013) report that most RBC backers are characterized as consumers who purchase based on product features and expect to receive product-related benefits.

Creators have access to a variety of features through the Kickstarter platform. A home page allows the overview of their product, a rewards page allows menu creation, and other features were introduced over time–detailed in the next section.

Menus feature multiple tiers (menu options), which are always arranged in order of price (lowest to highest) for all available tiers. Backers can purchase a maximum of one tier. Tiers offer a version of the product, an add-on (product branded t-shirts, stickers, etc.), recognition for donation, or some combination of these. Figure 2.1 shows an example of recognition for donation. These tiers are usually the lowest priced tier on the menu and allow for prosocial support without reward. For buyers, these may act purely as a *price anchor* since they are only interested in tiers offering the product but may attend to the irrelevant tier's price.

[20] See Cai et al. (2021) for a systematic review.

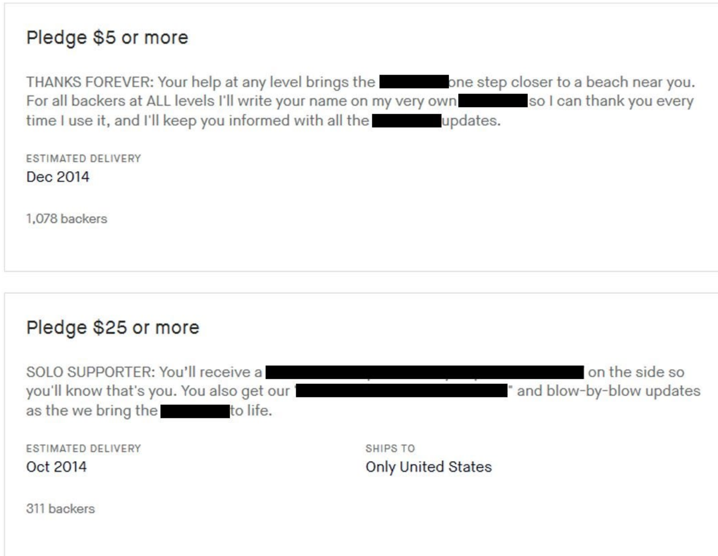

Figure 2.1: A Donation Tier

Note. The $5 tier offers recognition for donation without tangible reward

2.2.2 A Brief History of Kickstarter's Paradigm

Across our data Kickstarter implemented many feature introductions and updates, partnerships, and policy changes.[21] We discuss the relevant changes and their economic impact. Figure 2.2 presents a timeline and Table 2.1 presents the impact of selected changes.

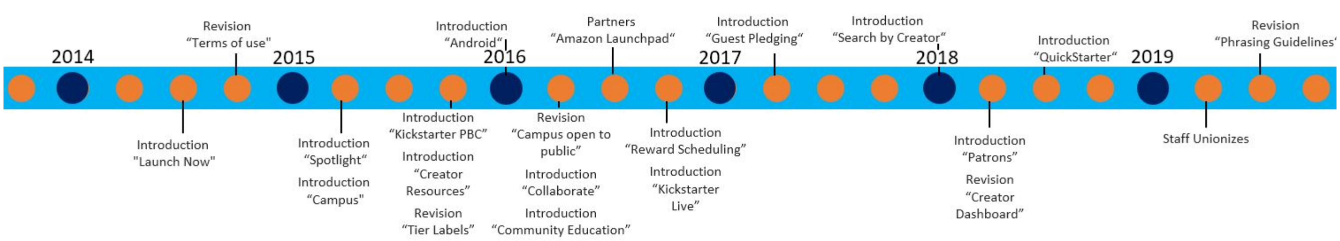

Figure 2.2: Timeline of Major Changes

Note. All major changes made by Kickstarter from October 2013 to December 2019 are shown by fiscal quarter.

[21] These changes are observed through Kickstarter's blog.

In June 2014, the addition of “Launch Now” simplified the project launch and abridged the rules for creating projects. A new cohort of creators could now enter the market increasing projects launched by 78% in the following month. In October 2014, the terms of use were revised to guarantee the delivery or refund of rewards and make delivery rules more flexible. Consumers gained more protection and entrepreneurs now had more time to deliver rewards after a successful campaign which increased the numbers of backers for each project.

In March 2015, the “Spotlight” feature let projects modify their Kickstarter page after a campaign. Creators could now advertise and coordinate backers through Kickstarter post-project. Sometime between April and October 2015, “Campus” introduced a private forum for creators.[22] Creators could now directly socially learn through Kickstarter. September 2015 was a turning point; Kickstarter became a public-benefit corporation. Public benefit corporations place poignant attention on their impact on society–in Kickstarter’s case, enabling entrepreneurs. In October 2015, Kickstarter launched a creator resources page–a list of companies' specialized fulfillment–and a creator handbook–a guide to creating and managing projects. Creators now had a uniform source of learning supplied by Kickstarter. In December 2015, titles were added to reward tiers.[23] Menus could now draw consumer attention to specific tiers using bold font text, increasing backer concentration in each campaign’s most popular campaign by 10%.

In January 2016, Kickstarter launched on Android. Consumers without computers could now participate in the market. In February 2016, the “Projects We Love” feature let Kickstarter draw attention to staff selected projects.[24] The demand for featured projects rose by 36.6%. In March 2016, the “Campus” feature was made public. New creators could now access the forums

[22] This introduction is not mentioned directly by a blog post but its date can be inferred from other posts.
[23] There was no blog post mention of this, but it is observed in our dataset
[24] This is an improvement on the “Staff Pick” feature which served the same purpose.

and archive to socially learn before launch and backers could also access the information. In May 2016, project management was enabled from multiple accounts. Previously only a single account could manage the project. Also, in May 2016, a community education program was launched. Kickstarter expanded learning initiatives for both creators and backers. In July 2016, Kickstarter partnered with Amazon Launchpad–which develops and delivers new products. In October 2016, "Reward Scheduling" let creators make limited-time sales and schedule sales throughout their campaign. Before, sales could only be limited by the number of backers.

Table 2.1: Summary of Effects of Chosen Changes on Kickstarter

Change	Summary	Metric	Effect
Launch Now June 3 2014	Eased restrictions to create projects	Projects	1514* (648.4)
Delivery Simplified October 29 2014	Made Delivery rules more flexible	Backers	28.18** (9.66)
Tier Titles December 1 2015	Gave tiers bold titles which direct attention	% of backers in most pop. tier	0.051*** (0.011)
Projects We Love February 2 2016	Made Staff Pick projects more noticeable	Backers	129.2 (94.93)
Search by Creator February 8 2018	Let backers search for specific creators	Backers†	143.2** (54.7)
Quickstarter June 18 2018	Prompted the creation of smaller projects	Small Projects	24 (54.7)

Standard errors in parentheses.
Note. The effect of each policy is shown through a comparison of means between 3 months after the policy and 3 months prior.
† The differences-in-differences estimate is shown where creators who have a previous project are the treated group.
* p<0.10, ** p<0.05, *** p<0.01

In April 2017, backers could pledge without making a Kickstarter account. The marginal cost to backing for new backers was then reduced. In February 2018, backers could now search for projects by creator. Seller reputation became more critical and the number of backers per project for repeat creators increased by 143 as a result. In March 2018, Kickstarter partnered with several institutional donors. These institutions offered large donations to specific projects, focused on societal betterment, directed by Kickstarter. In March 2018, the "Creator Dashboard" allowed all creators access to statistics and analysis on other projects. Creators could now indirectly socially learn from the success of others. In June 2018, Kickstarter challenged entrepreneurs to create small projects. The prompt encouraged projects with small goals (below $1,000) and low-value products (below $50). In March 2019, the Kickstarter staff unionized. Finally, in June 2019, phrasing guidelines were changed to limit outlandish project claims such as "World's Best."

2.3 Data

Our data capture all 136,738 projects on Kickstarter between October 2013 and December 2019 created in the United States. Each project follows a standard html format through the Kickstarter website which enables direct web scrapping in python given the list of projects found on Kickstarter's API. Projects provide public information on their overall funds raised, total backers, goal, creator(s), and tiers. Tiers display backers, price, title (since December 2015), description, and a limit on backers if applicable. We compile our data from finished projects and record all values at the end of the project.

We clean the data by excluding projects with default currency other than US Dollars. We then remove projects with a free tier (a short-lived promotion on Kickstarter) and extreme values ($price \geq \$10{,}000$).[25]

To treat different tiers with solely cosmetic differences, we combine tiers with the same price and sum their backers to create a single tier at each price for each project.[26] We then match

[25] The mean tier price prior to this cut was $400 with a median of $60.
[26] Cosmetic features include color, character design, or format (USB, DVI, HDMI etc.)

tiers with titles and text descriptions and cut any tiers whose text description or title cannot be found. We are left with 97,333 projects with 797,676 tiers.

We use the titles and descriptions of tiers to match those that offer the same product at a different price. First, we force all characters into lowercase; then, we remove key phrases (shown in the appendix) associated with Kickstarter sales. After the text has been cleaned, we label tiers whose title or description exactly match as a sale set. Then we delete all tiers in this set aside from the lowest-priced member.[27] We create a binary variable, sale, which indicates if a tier is the lowest price member of a sale set. We remove 14,379 tiers through this process. Finally, we remove projects with too few (less than three) or too many (more than a hundred) tiers.

After cleaning the data, we generated our descriptive measures of interest. We set the price of lower and upper price neighbors by ordering each project's menu by price and taking the price of the previous (lower) and next (upper) tiers for each focal tier. We report the *sequential price gaps* below as the difference of the price of the upper neighbor and the focal tier.

Table 2.2: Descriptive Statistics by Year

Variable	Overall	2013	2014	2015	2016	2017	2018	2019
Tier Backers	19.70	15.04	12.83	22.09	21.31	21.82	22.18	28.54
	(332.48)	(151.02)	(170.24)	(601.28)	(337.29)	(200.14)	(180.14)	(262.25)
Tier Price	437.80	548.32	390.20	453.79	432.18	406.85	388.05	333.79
	(1308.05)	(1514.65)	(1405.05)	(1332.24)	(1303.03)	(1238.75)	(1207.70)	(1093.03)
Sequential Price	256.86	304.39	283.97	263.91	259.48	239.51	229.20	198.75
Difference	(850.32)	(906.48)	(889.81)	(865.92)	(863.72)	(815.12)	(805.32)	(742.05)
Sale (1 if on-sale)	0.014	0.003	0.004	0.006	0.030	0.019	0.017	0.019
	(0.12)	(0.055)	(0.064)	(0.079)	(0.169)	(0.138)	(0.129)	(0.138)
Tiers	703,345	13,021	196,015	129,114	129,153	108,416	66,351	61,152

Standard deviations shown in parentheses.

[27] We also record the discount from highest to lowest collapsed tier, but it is not used in analysis.

Before analysis, we remove tiers that are no longer available at the end of the campaign.[28] The final sample comprises 89,028 projects and 703,345 tiers. Table 2.2 shows the relevant summary statistics. On average, tiers have 20 backers, a price of $438, a sequential price difference of $257, and 1.4% are presented as sales.

2.4 Empirical Specification

We analyze quantity demanded which is directly observed on Kickstarter: the number of backers per tier. We parametrize a neoclassical demand function Q_{tj} of tier t in project j as $F\left(q_{tj}, p_{tj}, q_{t-1j}, p_{t-1,j}, q_{t+1j}, p_{t+1,j}\right)$ given price p_{tj} and quality q_{tj} of tier t and price and quality of neighbors $\mathrm{q}_{\mathrm{t-1j}}, \mathrm{p}_{\mathrm{t-1,j}}, \mathrm{q}_{\mathrm{t+1j}}, \mathrm{p}_{\mathrm{t+1,j}}$. However, quality is unobserved, so we express this function as a linear projection on prices:

$$\begin{aligned} Q_{tj} &= F(q_{tj}, p_{tj}, q_{t-1j}, p_{t-1,j}, q_{t+1j}, p_{t+1,j}) \\ &= \alpha p_{tj} + \lambda_1 p_{t-1,j} + \lambda_2 p_{t+1,j} \end{aligned} \tag{1}$$

The projection in equation 1 generates biases λ_1 and λ_2 on $p_{t-1,j}$ and $p_{t+1,j}$ respectively.

Theorem 1. *The neoclassical demand bias λ_2 is strictly positive.*

[28] We leave these tiers in the menu when we determine price neighbors because they were part of the menu for a portion of the campaign.

See the Appendix for the detailed proof. A profit-maximizing firm allocates tiers via price and quality to force consumers to self-select into optimal tiers. As firms isolate a tier, i.e. the firm reduces the lower price and raise the higher price, more consumers self-select into that tier. Note that the sign of bias λ_1 is ambiguous while α captures the price elasticity of demand tier t in project j which is also biased.[29]

We use Theorem 1 to inform the analysis of our model. We introduce the behavioral response of consumers to context effects δ_1 and δ_2, the anchoring effect of sales δ_3 and a measure of attrition in consumer attention δ_4. We use project-specific fixed effects ϵ_j to control for overarching project quality[30] and project-specific random effects v_j to mitigate potential heterogeneity sources due to unobserved tier quality.

$$Q_{tj} = \alpha p_{tj} + (\delta_1 + \lambda_1) p_{t-1,j} + (\delta_2 + \lambda_2) p_{t+1,j} + \delta_3 sale + \delta_4 order + v_j + \epsilon_j + u_{tj} \tag{2}$$

Model 2 allows us to explore the effect of menu design on consumer choice. Our setup allows the detection (but not point identification) of the compromise effect, when is δ_2 negative and overcomes positive bias λ_2. As the higher-priced tier comes closer in price to the focal tier, it becomes a stronger comparison, making the consumer more likely to purchase the focal tier over other projects on Kickstarter (their outside option). The latter terms δ_3 and δ_4 are point identifiable within model 2 as sale framing is unrelated to underlying quality, and the menu is strictly ordered through price on Kickstarter.

[29] Both these claims can be proved, but are tangential to the development of our analyses.
[30] The fixed-effects also capture time-specific shocks for each project and tier.

Altruism from backers may motivate prosocial behavior that is differentiated from the compromise effect but still yields a negative δ_2: backers with prosocial intentions may seek to overbuy tiers and purchase the next highest tier only when the difference is substantial. Prosocial behavior is most consequential (and thought of to be isolated) around a project's goal (Dai and Zheng 2019). Thus, we analyze model 2 while stratifying on the project's realized goal achievement. In our cross-sectional setup, we expect most prosocial behavior to come in projects that achieved just over 100% of their goal and the least in projects that well achieved their goal and well underachieved their goal. Thus, if the negative sign is specific to regions away from 100% goal attainment, prosocial behavior is separated from the actual driver.

The data offers several advantages: we observe an extended period across many policy changes directly impacting the menu structure. Thus, we extend this analysis through time to explore the impact of the significant policy changes discussed in Section 2.2.2. We estimate model 2 for each fiscal quarter beginning in 2013 Q4 and ending in 2019 Q4. These quarterly regressions demonstrate robustness across time, enable the analysis of policy impacts, and unveil any long-term learning.

2.5 Empirical Estimates: Stuck in the Middle

Table 2.3 provides the result of the log-log regression of equation 2. Specification 1 is the linear projection of the log of total backers on the log of price, the log of lower neighbor price, the log of upper neighbor price, a binary variable for sale, and the order the tier appears on the menu. Specification 2 adds project fixed effects to control for project quality. Specification 3 introduces random effects applied to both the intercept and the slopes of the model to account for variance across project tiers due to unobservable differences in tier quality. All estimations are performed in STATA using the mixed-effects package with default settings (StataCorp 2019).

Specification 1 shows an insignificant price elasticity coefficient due to tier order and the lower and upper cross-price elasticities introducing bias when project quality is not controlled. In specification 2, our results begin to take shape; here, consumers are inelastic in price, and the lower cross-price elasticity is positive as predicted by the neoclassical model. However, the upper cross-price elasticity is significantly negative, indicating a significant flow of consumers towards tiers closely neighbored in price either above or below. Such a result contradicts the neoclassical model, which predicts positive upper-cross price elasticity. If one inserted “decoy” tiers throughout the menu to halve the sequential price difference for all tiers, it would increase the demand in all options (apart from the decoys) by 6.89%.

Although the demand increase is significant, consumers have a larger response to sales driven by price anchoring with the discounted tier anchored by the original. Framing a tier as a sale increases demand by 78%. This is more than ten times as impactful as the aforementioned “decoy” effect. This discrepancy in magnitude drives our results in the following section, where creators deviate focus to sales. Finally, consumers exhibit large attrition with a 9.42% reduction in demand for each subsequent tier past the second on the menu.

Table 2.3: Demand Elasticity of Menu Features

Dependent Variable: Log(Tier Backers)			
	(1)	(2)	(3)
Log(Lower Price)	0.165***	0.143***	0.157***
	(0.0057)	(0.0049)	(0.0036)
Log(Upper Price)	-0.307***	-0.0576***	-0.123***
	(0.0064)	(0.0058)	(0.0047)
Log(Tier Price)	0.00385	-0.429***	-0.298***
	(0.0098)	(0.0089)	(0.0068)
Discounted Tier	0.787***	0.746***	0.780***
	(0.0337)	(0.0289)	(0.0210)
Tier Order	-0.0683***	-0.0821***	-0.0942***
	(0.0018)	(0.0023)	(0.0011)
Fixed Effects		X	X
Random Effects			X
Tiers	366,360	366,360	366,360
Projects	78,950	78,950	78,950

Standard errors clustered by project in parentheses
Note. Specification (1) shows menu pricing feature elasticity with standard errors clustered at the project level, Specification (2) introduces fixed effects, and Specification (3) introduces random effects. To analyze tiers must have a lower and higher price neighbor, thus the lowest and highest priced menu options are dropped.
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

To demonstrate that consumers rather than prosocial donors drive these results, we replicate the analysis by stratifying on the project's goal attainment: we bucket projects into those that are well beyond their goal (130%+), those that barely achieved their goal (10% bins from 100% to 130%), those that barely missed their goal (70%-100%), and those that missed their goal by a large margin (0-70%).

Table 2.4: Demand Elasticity by Goal Attainment

Log(Lower Price)	0.198***	0.203***	0.191***	0.154***	0.0478*	0.0321***
	(0.0077)	(0.0173)	(0.0129)	(0.0080)	(0.0245)	(0.0050)
Log(Upper Price)	-0.195***	-0.0206	0.0707***	0.137***	-0.106***	-0.0988***
	(0.0106)	(0.0218)	(0.0163)	(0.0098)	(0.0322)	(0.0061)
Log(Tier Price)	-0.280***	-0.507***	-0.606***	-0.653***	-0.231***	-0.0667***
	(0.0152)	(0.0322)	(0.0238)	(0.0145)	(0.0459)	(0.0090)
Discounted Tier	0.987***	0.608***	0.719***	0.473***	0.970***	0.742***
	(0.0322)	(0.0841)	(0.0558)	(0.0369)	(0.1012)	(0.0208)
Tier Order	-0.119***	-0.0952***	-0.0937***	-0.0767***	-0.0923***	-0.0768***
	(0.0020)	(0.0042)	(0.0033)	(0.0021)	(0.0072)	(0.0018)
Tiers	104,064	19,413	32,708	74,845	7,203	123,289
Projects	18,847	3,281	5,479	12,835	1,346	36,600

Standard errors clustered by project in parentheses
Note. Projects are partitioned based on their total funds raised divided by their initial project goal. The upper cross-price elasticity is negative for all 10% bins with 130%+ and so these are combined. Very few projects barely underachieve their goal and so we choose the bin 70%-100% to be wide to allow sufficient sample size.
* p<0.10, ** p<0.05, *** p<0.01

We see that regions where prosocial behavior is most prominent (100%-130%) fail to exhibit significant context effects. The negative upper cross-price elasticity changes its sign for projects that just barely achieved their goal. In these regions, donors actively push campaigns over their threshold (Dai and Zhang 2019). Donors are primarily interested in the project achieving its goal with consumption an auxiliary concern. Thus, the cross-comparison of tiers does not occur for donors. The positive biases drive the upper cross-price elasticity due to menu design; this positive coefficient may also indicate prosocial behavior that drives backers to overbuy into the next highest tier only when the difference is small. Prosocial behavior also diminishes sale elasticity, which is the smallest around the goal, further demonstrating backers are less concerned with consumption in these regions. Perhaps surprisingly, the price elasticity is higher for projects that barely achieved their goal. We speculate that such phenomenon can be explained by donors who purchase lower-priced tiers to propel projects towards their goal without incurring the cost of supplying additional higher-priced tiers (Dai and Zhang 2019).

In summary, we find consumers prefer tiers with close price neighbors. This is evidence of the compromise effect. We show that prosocial behavior does not play a significant role in driving the results. Cross-price elasticities driving this middle preference are small relative to sale elasticity. Sales can be used in the same way as "decoy" tiers by adding additional tiers compared to other menu options. Yet, sales are ten times as effective. This difference in magnitude is crucial as we explore these elasticities over time.

2.6 Dynamic Analysis: Evidence of Learning

We present an analysis of context effects by a quarter of project launch. We use model 2, which includes a project-level fixed effect that controls differences in quality between projects and controls for time-specific supply or demand shocks. Only changes that influence consumer choice or menu presentation should affect our elasticities.

Figure 2.3 shows the absolute value of elasticities for each quarter with the timeline of major changes overlaid. Table 2.5 presents the difference of elasticities three months before each change and three months after, where elasticities are calculated via model 2 for each period. The difference between the absolute values of elasticities are presented, so a positive always value represents an increase in significance.

The major events that affected menu presentation began in the second quarter of 2015 with "Campus," a forum that linked Kickstarter creators and enabled direct social learning. We see an increase in sale elasticity during this period but no significant change. Presumably, creators communally learned that anchoring consumer decisions through discounts yields high profits and began to focus on this mechanism. We also see a measured downward trend in both cross-price elasticities begin in this quarter. If a consumer is directed towards specific choices with the more powerful invocation, price anchoring through a sale, they have less attention to devote to other contexts. Thus, a decline in cross-price elasticity is explicable.

In the fourth quarter of 2015, Kickstarter released creator resources. These provided direct learning from experts. Again, we see an increase in sale elasticity but no significant changes. Later in the fourth quarter of 2015, tier titles were introduced, which helped draw attention to specific tiers with bolded labels. Sales were now easier to call attention to, which significantly increased sale elasticity by 0.266. Lower cross-price elasticity also increased significantly by 0.028 because donation tiers–which institute a lower price anchor–became more noticeable.

In the second quarter of 2016, the "Campus" forum was opened to the public, giving new creators a chance to participate before launching their first campaign and allowing backers to view creator discussions. This had a limited impact on the elasticities in the market. The lower cross-price elasticity significantly declined at this time; conceivably, backers became aware of the creator's use of donation tiers as a price anchor to drive sales. Consequently, this caused either a reduction of purchases from projects that use donation tiers or reduced donation tiers as creators wanted to distance themselves from this manipulative behavior after backers caught on. However, we lack sufficient evidence to prove these claims as Campus records are no longer available.

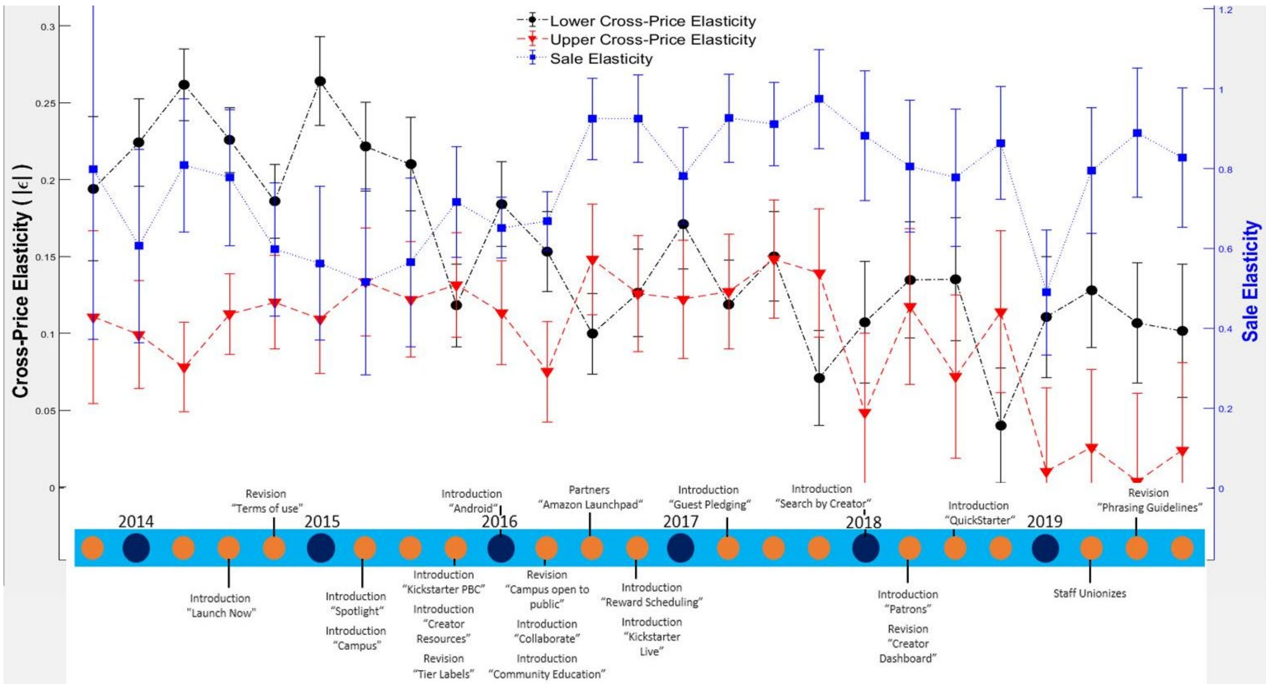

Figure 2.3: Demand Elasticities by Quarter

Note. We partition our sample by the quarter of launch and replicate specification 3 from Table 2.3. The absolute value of the cross-price elasticities is shown by the left axis, and the right axis shows the sale elasticity (percentage increase in demand when tier framed as a sale).
Error bars represent 95% confidence intervals based on standard errors clustered by project.

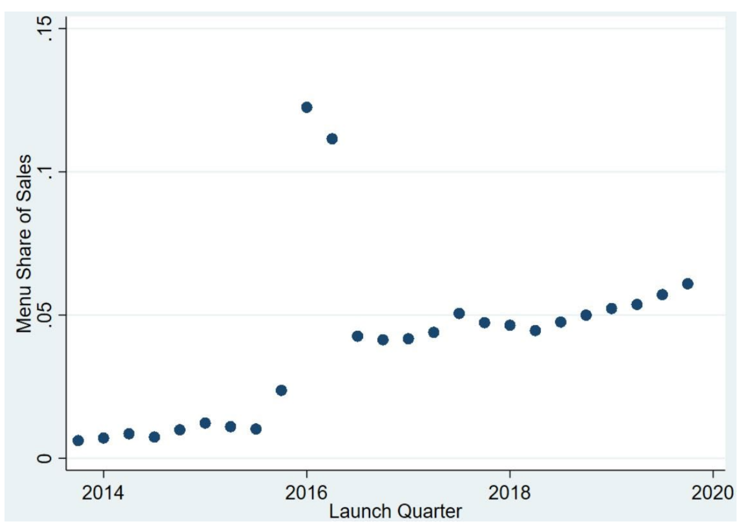

Figure 2.4: Trends in Sale Usage

Note. We determine menu share of sales as the number of tiers associated with a sale divided by total tiers for each quarter prior to any cuts. The large increase in sales in 2016 Q1 and Q2 is due to issues where created sale tiers became inaccessible and had to be remade but the old instances remained on menus.

Table 2.5: The Impact of Platform Changes on Elasticities

Change	Summary	Analysis	Change in Elasticity ($\Delta\|\epsilon\|$)			
			Sale	Lower	Upper	Price
Campus April 25 2015	Introduced a forum for creator discussion	Induced direct social learning among creators	0.313 (0.197)	-0.019 (0.017)	0.012 (0.021)	-0.031 (0.031)
Creator Resources October 29 2015	Launched a page with information for creators	Homogenized campaigns	0.139 (0.134)	-0.027 (0.017)	0.011 (0.021)	-0.046 (0.032)
Tier Titles December 1 2015	Gave tiers bold titles	Improved ability to direct consumer attention	0.266** (0.125)	0.028* (0.017)	-0.008 (0.021)	0.041 (0.029)
Public Campus March 2 2016	Publicly made available forum for creators	Let backers become cognisant of donation anchoring	0.001 (0.047)	-0.036** (0.016)	-0.033 (0.020)	-0.029 (0.029)
Reward Scheduling October 18 2016	Enabled scheduling of sales and limitation by time	Improved convenience but not efficiency of sales	0.065 (0.091)	-0.003 (0.017)	-0.011 (0.023)	-0.012 (0.033)
Dashboard March 28 2018	Launched a tool with statistics for creators	Induced indirect social learning of compromise effect for creators	0.132 (0.147)	0.030 (0.023)	0.065** (0.031)	-0.043 (0.044)

Pooled standard errors clustered by project in parentheses.
Note. The effect of each policy on each elasticity is calculated as the difference between the absolute value of the elasticity for the 3-month window after the policy and the 3-month window before the policy. These figures are initially calculated with model 2 for each window. Standard errors are pooled between both windows to derive those shown in the table. A summary of each policy and analysis of its effect are included.
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In the fourth quarter of 2016, Kickstarter instituted "Reward Scheduling," which made introducing sales into a campaign easier. Creators could now plot their full arsenal sales before the campaign's launch and limit them by time instead of a number of backers. Although Table 2.5 shows an insignificant effect, Figure 2.3 highlights a separation in cross-price and sale elasticity after this change; sales became more relevant and cross-price elasticities less significant.

The long-term increase in sales significance is heralded by the rise in sale frequency among all projects. Figure 2.4 shows a consistent positive trend in sale usage throughout the entire sample. The most significant increase occurs after tier titles are introduced. During this time, creators learned to effectively use sales to garner consumer attention, diverting consumer attention away from non-sale tiers, which diminished context effects but did not eliminate them. Menus also generally became crowded with sales which made the direct comparison of price neighbors more challenging.

Despite this, in the second quarter of 2018, the “Creator Dashboard” allowed creators to view aggregate statistics on other projects and easily see which tiers performed well. This caused indirect social learning and led to the accidental discovery of the compromise effect amongst creators who began to use more narrowly distributed prices on menus and significantly increased the (absolute value of) upper cross-price elastic by 0.065.

Overall, we see a trend of increased significance of price anchoring driven by Kickstarter's social learning tools and helped by tools that made directing consumer attention with text easier and the initiation of sales easier. During this time, sale use became more widespread. This led to a decline in consumer attention to neighboring tiers' prices which diminished context effects on the platform but failed to dissolve them completely.

2.7 Conclusion

We find that backers allocate towards tiers with close price neighbors. For a focal good, the lower price neighbor has a positive cross-price elasticity, while the higher price neighbor has a negative cross-price elasticity. We conclude this middling preference is evidence of context effects; when either price neighbor acts as a close comparison, consumers are more likely to purchase the focal option over an outside option. This middling preference diminishes as distractions are introduced to menus. Through the introduction of new tools that promote direct social learning and enable sellers to better direct consumer attention, sellers learned to exploit sales–which act as price anchors–and draw consumer attention away from price neighbors.

Our results support Fredrick et al. (2014) and Yang and Lynn (2014), who argue that context effects diminish as stimuli aside from numerical values become more readily available. However, we also support the counterargument of Huber et al. (2014) by providing evidence that despite the introduction of powerful non-numerical stimuli, context effects persist.

Our findings are directly applicable to the natural setting studied, reward-based crowdfunding. We show that despite the large impact of sales, context effects are still significant. Importantly, we show attrition that reduces demand for tiers deeper in menus. An RBC entrepreneur may be able to increase demand by adding sale tiers and “decoy” tiers to menus. However, diminishing returns may eventually lead to a loss with the introduction of new tiers.

We introduce a new dataset for further exploration of reward-based crowdfunding and consumer choice and offer some possible directions for future research.

Donation tiers may act as “no-purchase” options which have been shown to diminish context effects (Lichters et al. 2017). We were unable to identify donation tiers with sufficient precision and left this aspect of the market unexplored. The interaction of context effects and the “phantom” effect of Adam et al. (2019) merits further exploration. Finally, the interplay between pre-purchase sales, cognitive biases, and the prediction of post-launch sales remains a relevant area of investigation (Moe and Fader 2002).

2.8 References

Adam M, Wessel M, Benlian A (2019) Of early birds and phantoms: How sold-out discounts impact entrepreneurial success in reward-based crowdfunding. Review of *Managerial Science* 13:545–60.

Barbos A (2010) Context effects: A representation of choices from categories. *Journal of Economic Theory* 145(3):1224–1243.

Belleflamme P, Lambert T, Schwienbacher A (2013) Individual crowdfunding practices. *Venture Capital* 15(4):313–333.

Blaseg D, Schulze C, Skiera B (2020) Consumer protection on kickstarter. *Marketing Science* 39(1):211–33.

Bonini N, Rumiati R (2002) Acceptance of a price discount: the role of the semantic relatedness between purchases and the comparative price format. *Journal of Behavioral Decision Making* 15(3):203–20.

Cai W, Polzin F, Stam E (2021) Crowdfunding and social capital: A systematic review using a dynamic perspective. *Technological Forecasting and Social Change*.

Carbajal J, Ely J (2016) A model of price discrimination under loss aversion and state-contingent reference points. *Theoretical Economics* 11(2):455–485.

Dai H, Zhang DJ (2019) Prosocial goal pursuit in crowdfunding: Evidence from Kickstarter. *Journal of Marketing Research* 56(3):498–517.

Doyle JR, O'Connor DJ, Reynolds GM, Bottomley PA (1999) The robustness of the asymmetrically dominated effect: Buying frames, phantom alternatives, and in-store purchases. *Psychology & Marketing* 16(3):225–243.

Fredrick S, Lee L, Baskin E (2014) The limits of attraction. *Journal of Marketing Research* 51(4):487–507.

Geyskens I, Gielens K, Gijsbrechts E (2010) Proliferating private-label portfolios: How introducing economy and premium private labels influences brand choice. *Journal of Marketing Research* 47(5):791–807.

Huber J, Payne J, Puto E (1982) Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research* 9(1):90–8.

Huber J, Payne J, Puto E (2014) Let's be honest about the attraction effect. *Journal of Marketing Research* 51(4):520–525.

Liao J, Chen Y, Lin W, Mo L (2020) The influence of distance between decoy and target on context effect: Attraction or repulsion? *Journal of Behavioral Decision Making*.

Lichters M, Bengart P, Sarstedt M, Vogt B (2017) What really matter in attraction effect research: When choices have economic consequences. *Marketing Letters* 28:127–38.

Maltz A, Rachmilevitch S (2021) A model of menu-dependent evaluations and comparison-aversion. *Journal of Behavioral and Experimental Economics* 91.

Marzilli Ericson KM, Fuster A (2011) Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments. *The Quarterly Journal of Economics* 126(4):1879–1907.

Moe WW, Fader PS (2002) Fast-track article: Using advance purchase orders to forecast new product sales. *Marketing Science* 21(3):347–364.

Mollick E, Robb A (2016) Democratizing innovation and capital access: The role of crowdfunding. *California Management Review* 58(2):72–87.

Ok E, Ortoleva P, Riella G (2014) Theory of product differentiation in the presence of the attraction effect.

Ok E, Ortoleva P, Riella G (2015) Revealed (p)reference theory. *American Economic Review* 105(1):299–321.

Petitjean M (2018) What explains the success of reward-based crowdfunding campaigns as they unfold? Evidence from the french crowdfunding platform kisskissbankbank. *Finance Research Letters* 26:9–14.

Pinger P, Ruhmer-Krell I, Schumacher H (2016) The compromise effect in action: Lessons from a restaurant's menu. *Journal of Economic Behavior & Organization* 128:14–34.

Shneor R, Vik A (2020) Crowdfunding success: A systematic literature review 2010-2017. *Baltic Journal of Management* 15:149–182.

Simonson I (1989) Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research* 16(2):158–74.

Soubilere JF, Gehman J (2020) The legitimacy threshold revisited: How prior successes and failures spillover to other endeavors on kickstarter. *Academy of Management Journal* 63(2).

StataCorp (2019) Stata Statistical Software: Release 16. College Station, TX: StataCorp LLC.

Tserenjigmid G (2019) Choosing with the worst in mind: A reference-dependent model. *Journal of Economic Behavior & Organization* 157:631–52.

Viotto da Cruz J (2018) Beyond financing: crowdfunding as an informational mechanism. *Journal of Business Venturing* 33(3):371–393.

Vismara S (2016) Equity retention and social network theory in equity crowdfunding. *Small Business Economics* 46:579–90.

Weinmann M, Mishra A, Kaiser L, vom Brocke J (2021) The attraction effect in digital markets: Evidence from four experiments on reward-based crowdfunding.

Whyte G, Sebenius J (1997) The effect of multiple anchors on anchoring in individual and group judgement. *Organizational Behavior and Human Decision Processes* 69(1):75–87.

Wu C, Cosguner K (2018) Profiting from the decoy effect: A case study of the online diamond marketplace.

Yang S, Lynn M (2014) More evidence challenging the robustness and usefulness of the attraction effect. *Journal of Marketing Research* 51(4):507–13.

CHAPTER THREE.

WHEN BELIEVING IS A STRETCH: STRATEGIC COMMUNICATION OF PROSOCIAL GOALS IN REWARD-BASED CROWDFUNDING

3.1 Introduction

Donations and prosocial behavior can bridge many economic chasms to enhance societal welfare. However, donors often lack clarity in the value of their support (to themselves and/or their cause) which drives inefficiencies in prosocial markets. Donors rely on communication from fellow donors and their beneficiaries to alleviate this uncertainty which provides opportunity for deception. Individual donors may wish to free-ride and deceive others into donating their share which breaks down credible donor-to-donor communication in theory (Arrow 1950). And beneficiaries may inflate or misrepresent prosocial goals to reap personal benefit (McFadden 1992).

Despite these incentives to deceive, actionable donor-to-donor communication is well-documented (Beersma et al. 2020) yet beneficiary-to-donor communication has proved challenging to study. In this paper, we explore the communication of prosocial goals by a beneficiary–a reward-based crowdfunding entrepreneur–to his donors, whether it is believed and acted upon, and how it should be handled by policy makers.

Reward-based crowdfunding (RBC) is an investment structure which enables the introduction of market innovations through consumer pre-purchase and philanthropic donations.

RBC entrepreneurs set an initial goal, that if met establishes a contract for the creation of a new product and delivery to those that pre-purchased. These products often fill markets gaps which provides a prosocial incentive for philanthropists (donors) to donate (Pew Research Center 2015). However, RBC entrepreneurs often establish secondary goals (stretch goals) that, unlike initial goals, offer improvements on the original product that are not contractually bound. This communication lacks credibility, yet these *stretches* are commonplace in RBC and other natural settings (e.g., CEOs commonly communicate stretch goals to investors). We introduce a panel dataset of all 720 projects in the technology category on the largest RBC platform, Kickstarter, from May 2020 to December 2020 to study this communication.

There are four broad classifications of credible strategic communication---cheap talk, verifiable message, Spence signaling, and Bayesian persuasion---which introduce realism to economic theory by allowing agents to communicate as they may in natural settings. Cheap talk allows parties to communicate directly and, when incentives align, coordinate such that they mutually maximize payoffs in one-off games (Crawford and Sobel 1982); despite realism added to specific games such as the battle of the sexes, cheap talk fails to yield realistic predictions in a variety of settings such as the stag hunt game proposed by Aumann (1990). A verifiable message allows parties to directly communicate such that there is no ability to deceive (Milgrom 1981); this theory adds realism to scenarios where communication is enforceable such as manufacturer warranties (Grossman 1981). Spence (1973) originated signaling games where heterogeneous players communicate through costly signals which result in separating Perfect Bayes-Nash Equilibria wherein each player informatively communicates their message; Spence signaling adds realism in scenarios where communication is informative despite incentive to deceive through the assumption of signal cost. And the most recent communication protocol, Bayesian persuasion lets parties communicate informatively with probabilistically deceptive messages which yield benefit to both parties despite the known deception (Kamenica and Gentzkow 2011); Bayesian persuasion adds realism to scenarios where parties with competing incentives must communicate in such a way that underlying strategies are known (either through repeated interactions or verification of strategy) such as a prosecutor with incentive to convict and a judge who desires to correctly rule. Communication is not easily quantified and has thus been underexplored in natural settings. Our communication---the stretch goal---is a combination of Spence signaling---we find separating use of the communiqué---and cheap talk---however, the signal is costless---and enables informative exchange despite incentives to deceive.

We develop a multi-goal contribution game to characterize the communication of multiple unverifiable secondary goals in general settings. In this sequential game, over a limited number of periods, beneficiaries communicate multiple stretch goals to donors who then respond by either fulfilling the goal or ending the game. Throughout the game, donors send a noisy signal of their value of a potential prosocial goal---the donor is uncertain of the form of the goal and thus can only offer a noisy estimate of its value. In any period, beneficiaries can wait and accumulate more information on the donor's value, set an honest stretch goal or set a deceptive one. However, he is punished by a penalty (imposed by donors in our case) if a deceptive stretch goal is fulfilled.[31] This game results in an equilibrium where only honest beneficiary-to-donor communication occurs but is truncated by the penalty for deception. The larger this penalty, the higher the equilibrium payoffs for beneficiaries and donors. This framework applies to all multi-goal contribution games where communication is unverified yet deception may be punished. We recommend policy makers

[31] This penalty is a key feature of our game and is discussed in detail in section 3.3.4.

support prosocial activity by maximizing penalties for deception and enabling precise channels for communication between beneficiaries and donors.

We use our new dataset on RBC entrepreneurial communication of stretches to test our model's implications and unveil specific attributes of this communication in RBC---which is in itself a valuable market of study---through model-free presentation of the data and reduced-form analysis.

3.1.1 Related Literature

When individuals act prosocially, they are prone to cognitive biases (Huck et al. 2015; Hutchison-Quillian et al. 2018). Limited cognition is pronounced in prosocial behavior because the individual lacks full clarity on the value of their actions (Yoruk 2012; Scherhag and Boenigk 2013). Prosocial settings are reliant on social learning to overcome these deficits and donating is known to be a sociable event (Vesterlund 2003; Andreoni 2006). Beersma et al. (2020) provide a systematic review of the social donation literature. We extend the literature of prosocial behavior to a new natural setting and introduce meaningful communication between beneficiaries and donors.

In RBC, donors are a highly social group who actively contribute to product improvement through repeated communication of features they value to entrepreneurs (Candogan et al. 2021). These donors have strong inclinations to achieve the initial capital goals set by entrepreneurs (Dai and Zhang 2019; Deb et al. 2021). However, RBC donors are often deceived by entrepreneurs through fake or misrepresented products (Kuppuswamy and Bayus 2017; Leamy 2018). Mechanisms to address this deception have been proposed; up-front fees reclaimable on positive product delivery is one possible solution (Agrawal et al. 2015). However, little is known beyond the achievement of the initial goal and eventual creation of the product to inform these mechanisms. Viotto da Cruz (2018) investigates campaign features that can act as credible signals of the eventual fulfillment of campaign promises: she finds that sentiment in comments, quality of spelling, number of backers, and social media likes are significant predictors. Significantly, Blaseg et al. (2020) show that price often changes after the campaign and conclude entrepreneurs deceive backers by misrepresenting prices during the campaign. We extend the RBC literature beyond the achievement of the initial goal and introduce the first study of stretch goal communication by entrepreneurs.

This chapter is broken into 5 section. Section 3.2 develops a simple sequential game which describes beneficiary to donor communication of prosocial goals. Section 3.3 details a relevant institutional background, Kickstarter, where this communication is commonplace. Section 3.4 details our panel data from Kickstarter to provide both model-free evidence and reduced form analysis of the relevant underlying mechanisms of this communication. And section 3.5 concludes.

3.2 A Multi-Goal Contribution Game

We create a simple sequential game to capture the behavior of a beneficiary (he) as he communicates stretch goals with his donor (she); these stretches are unverified communications of goals which if met promise some prosocial benefit (e.g., "give me \$200 and I will plant 200 trees"). Over T periods, the beneficiary acts first then the donor responds. There are T possible stretches $m = 1,2,\ldots,T$ which the beneficiary may set in order from 1 to T. For each stretch $m = 1,2,\ldots,T$ the beneficiary incurs known cost $C_m \sim U(0,C)$ drawn from a uniform distribution---which are strictly ordered from lowest to highest such that $C_1 \leq C_2 \leq \cdots \leq C_T$.[32] These costs

[32] This assumption limits strategies where the beneficiary will incur a loss in one period to collect a gain later

represent the net cost of enacting a prosocial benefit absent donations.[33] For each stretch $m = 1,2,\ldots,T$ the donor knows her value W_m, where $E(W_m) = W$, and both parties know the penalty δ that the donor can impose if the beneficiary deceives her with a fake prosocial goal. Each period $t = 1,2,\ldots,T$ begins with a fuzzy signal of the donor's value for the beneficiary $w_m = W_m + \epsilon$, where $\epsilon \sim N(0,\sigma)$. We think of this signal as the donor's direct communication of her value of a potential prosocial benefit with only a loose idea of what that benefit could entail. The donor then learns of her true value W_m of goal m once it is announced.

In every period the beneficiary has three actions---*stretch*, *fake*, or *sit*. To *stretch* the beneficiary announces a stretch goal s which if met incurs cost C_m and iterates m to $m + 1$. To *fake* the beneficiary announces a goal s which if met incurs penalty δ and iterates m to $m + 1$.[34] And in the *sit* action, the beneficiary remains inactive until the next period when he receives an additional signal on the current stretch m.

The donor cannot distinguish between a *stretch* or a *fake* and has two actions in response---*donate* or *pass*. If the donor *donates* she contributes s to meet stretch goal m which immediately yields payoffs $W_m - s$ and $s - C_m$ for the donor and beneficiary respectively if m was a *stretch* and payoffs $-s$, $s - \delta$ if m was a *fake*, and the game continues in the next period $t + 1$. If the donor *passes* then the game ends and no further actions are taken by either player. If the beneficiary *sits* then the donor also *sits* and takes no action for the period.

A visual representation of the game follows.

[33] In the RBC case, the loss from the cost of the stretch minus the gain from the demand expansion due to the product improvement. Though these net costs can be negative, the game is more generalized with strictly positive costs and there is no change to model intuitions with either assumption.
[34] The donor is unaware of deception until the full game as concluded

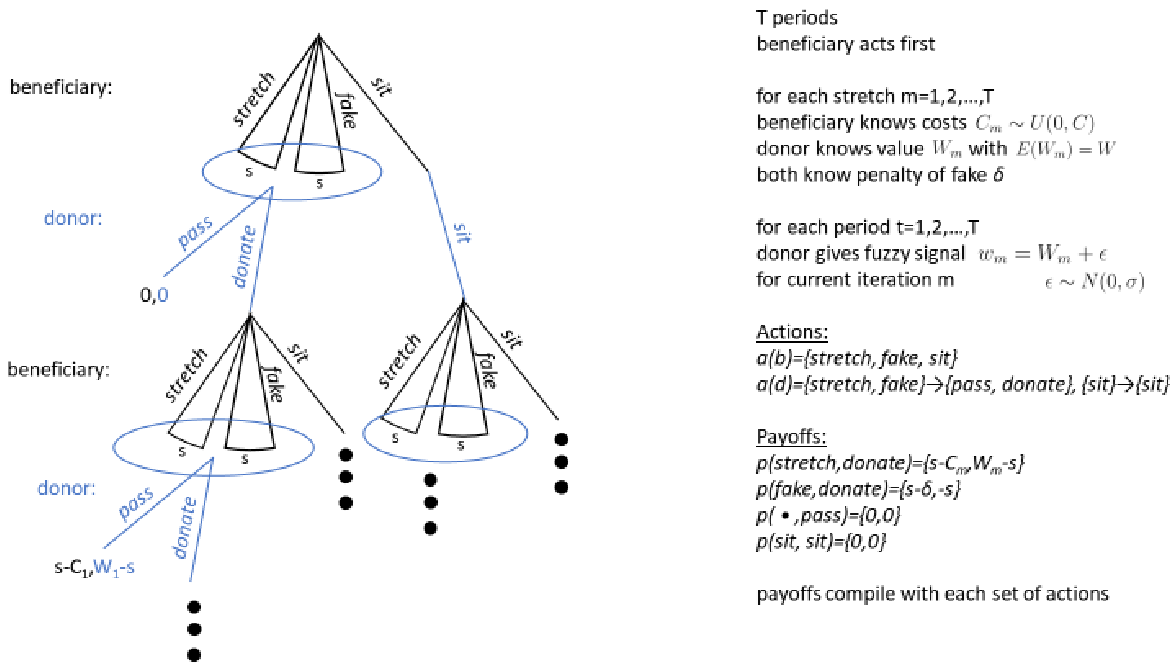

Figure 3.1: Visual Representation of Sequential Game

We first consider the beneficiary's problem. If the beneficiary *stretches* or *fakes* and the donor *passes* he forgoes the payoffs that could be gained were the game to continue. Instead, if the beneficiary *sits* he gains additional information on the donor's value and improves his knowledge on the distribution of noise in the donor's signal ϵ by $N\left(0, \frac{\sigma}{\sqrt{n_m}}\right) \rightarrow N\left(0, \frac{\sigma}{\sqrt{n_m+1}}\right)$ leading into the next period, where n_m is the number of signals the beneficiary has received for value W_m.[35] It is vital to define the value of a period $V\left(u \middle| \widehat{W_m}, n_m, \underline{C}_m\right)$ where $u = T - t$ is the number of remaining periods, $\widehat{W_m}$ is the average of received signals on W_m (W if there are no signals), n_m is the number of signals received on W_m, and $\underline{C}_m = \{C_m, C_{m+1}, \ldots, C_T\}$ is the vector of remaining stretch costs. The function is defined below.

[35] This is a simple transformation made possible by the assumption of a normally distributed ϵ.

$$
\begin{aligned}
V(u|\hat{W}_m, n_m, \underline{C}_m) \equiv \quad & \max\{stretch, fake, sit\} \\
stretch = \quad & \max_s \quad [s - C_m + V(u-1|W, 1, \underline{C}_{m+1})]Pr(donate, W_m > s|\hat{W}_m, n_m) \\
fake = \quad & \max_s \quad [s - \delta + V(u-1|W, 1, \underline{C}_{m+1})]Pr(donate, W_m > s|\hat{W}_m, n_m) \\
sit = \quad & V(u-1|\hat{W}_m, n_m + 1, \underline{C}_m) \\
where & \\
V(-1|\cdot) \equiv \quad & 0
\end{aligned}
\tag{1}
$$

The value of the optimal s, s^*, depends directly on C_m for *stretches* and δ for *fakes* and whether the donor would *donate* given such an s, but the exact value is irrelevant to the game's solution and desired intuition. More importantly, the value of *stretching* and *faking* is increasing in both $\widehat{W_m}$ and n_m. The beneficiary's strategy is simply to play the best action in each period.

The donor's strategy follows from the fact that s^* depends on C_m for *stretches*. The donor will *donate* if a given goal s is less than her value W_m and the penalty δ and *pass* otherwise. Given this donor strategy, the beneficiary's strategy reverts to the following.

First we identify the ideal *stretch* given costs C_m. Consider the lowest cost $C_m^{crit} < \delta$ such that $s^* > \delta$, in this state a payoff of $\delta - C_m > 0$ for a *stretch* s infinitesimally smaller than δ is preferable to a payoff of 0 yielded by a *stretch* $s = s^* \geq \delta$. This relationship holds for all $C_m \in \left[C_m^{crit}, \delta\right)$ such that all those with costs in this range will choose a *stretch* s infinitesimally smaller than δ. Those with costs $C_m \in \left[0, C_m^{crit}\right)$ will choose a *stretch* $s = s^*$. And those with costs $C_m \in [\delta, C]$ cannot meaningfully choose a *stretch* s such that the donor will *donate* and yield positive payoff. Likewise, any meaningful *fake* s such that the donor will *donate* and yield positive payoff does not exist in equilibrium. Finally, beneficiaries choose to *sit* as before, now dependent on the augmented value of *stretching* and *faking*.

The equilibrium strategy is shown below.

$$
\begin{aligned}
&Donor: \\
&\qquad s < \delta \rightarrow donate \\
&\qquad s \geq \delta \rightarrow pass \\
&Beneficiary: \\
V(u|\hat{W}_m, n_m, \underline{C}_m) \equiv &\quad \max\{stretch, fake, sit\} \\
stretch = &\quad \max_{s<\delta} \quad [s - C_m + V(u-1|W, 1, \underline{C}_{m+1})]Pr(W_m > s|\hat{W}_m, n_m) \\
fake = &\quad 0 \\
sit = &\quad V(u-1|\hat{W}_m, n_m + 1, \underline{C}_m)
\end{aligned}
\tag{2}
$$

Proof of Subgame Perfect Nash Equilibrium. For the beneficiary, the optimal strategy maximizes expected payoffs across the whole game given the donor's strategy. If the donor only *donates* when $s < \delta$ then the beneficiary will only ever meaningfully choose to *stretch* with $s < \delta$ or *sit*. Thus, the beneficiary will take the optimal actions governed by the expected value of each action shown in equation 2 when $C_m < \delta$. When $C_m \geq \delta$ he will act randomly since he is indifferent to every action (they all yield 0 payoff) but will never influence the donor since any *stretch* or *fake* $s < \delta$ incurs negative payoff. Then the donor knows given an $s < \delta$ *donating* yields positive payoff $W_m - s$ if $W_m > s$ and always *donates* in such scenarios and *passes* otherwise. Were the donor to modify their strategy to donate at some cutoff greater than δ then the beneficiary's optimal response modifies to *fake* when his $C_m > \delta$ and the donor gains negative payoff for supporting any additional range of s beyond δ. Likewise, chooses a cutoff below δ excludes positive payoff gained from beneficiaries with costs between the new cutoff and δ. Thus, the optimal strategy for the donor is to play a cutoff strategy with δ and the beneficiary's best response is to play the strategy in equation 2 when $C_m < \delta$ and inconsequential otherwise. ■

3.2.1 Model Implications

This game provides tractable structure which fits general scenarios where a beneficiary may endogenously set prosocial goals with incentive to deceive through fake or misrepresented goals. In equilibrium some beneficiaries stretch multiple times while others may never stretch or act randomly. In addition to this insight, this model provides two testable implications: (1) Donations stop if a stretch is set too high (2) Beneficiaries set realistic goals which allow time to act in future periods.

We proceed by describing RBC, a prosocial market where this behavior is a regular occurrence and detail relevant mechanisms in section 3.3. We then introduce an RBC panel data and explore donor interactions with stretch goals in section 3.4.

3.3 Kickstarter

Kickstarter is an RBC platform where an entrepreneur (beneficiary) can raise capital for a new product through the pre-purchase of the product by buyers and purely philanthropic donations by donors. The entrepreneur begins a Kickstarter project by setting a capital goal that must be met through the funds of backers (buyers and donors combined). Backers then have a limited time set by the entrepreneur to fund the goal. If met, the project succeeds and a contract between the entrepreneur and their backers is created, which states that the entrepreneur must deliver or refund purchased rewards. If the goal is not met, the project fails, and backers' funds are returned in their entirety.

Entrepreneurs on Kickstarter maintain a homepage which describes the project, an "updates" feature which allows direct communication with backers, and a rewards page which contains the price menu for the offered rewards. All pages can be modified throughout the campaign. The only static components of the project are the initial goal and deadline to meet said goal.

Buyers purchase rewards to derive value as a typical consumer and donors provide pure donations to projects without reward or overpay for a chosen tier to derive utility from a sense of contribution to society (Dai and Zhang 2019; Deb et al. 2021). This paper's focus is the contribution game played between the donor and the entrepreneur. We proceed by outlining the motivations of crowdfunding donors and the tools used by the entrepreneur for communication with this group---goals and stretch goals.

3.3.1 Donor Profile

RBC donors have been explored in depth in previous works. Dai and Zhang (2019) show that prosocial behavior driving donations is focused on achieving project success. Deb et al. (2021) describe how the donor manipulates buyer's perception of project success to increase sales revenue and reduce the amount needed in donations for project success. And Kuppuswamy and Bayus (2018) provide evidence that buyers may also act as donors and pay in excess of purchased tiers to ensure the project succeeds such that they receive their purchase.

Pew Research Center conducted an online poll of American adults in 2015 regarding their interactions with crowdfunding.[36] The survey found that 22% of adults have contributed to a crowdfunding campaign. Within those that have contributed there are two distinct groups---13%

[36] This survey was not exclusive to reward-based crowdfunding and included equity-based crowdfunding platforms such as GoFundMe.

of those that have participated in crowdfunding are serial donors who have supported more than five projects and the remainder are casual donors who have supported five or less projects. 14% of serial donors donate more than $500 per donation while only 3% of the full population contribute this amount on average.

Donor motivations were also polled. They found that 63% of casual donors backed projects to aid personal acquaintances, friends, or family members and 44% donated to fund a new product, invention, or business they believed was socially beneficial. The serial donor group differs in motivation with 85% donating to fund a new product, invention, or business they believed was socially beneficial. For both groups, motivations beyond investment in a person or social good are significant, 87% of all respondents claimed supporting projects built a sense of community and they felt personally connected to the projects they supported.

There are more family and friends that act as casual donors, but donation revenue is largely generated from the serial donor group. These individuals are willing to donate large sums to support projects they believe are socially beneficial and they act communally. We characterize our donors as this serial donor group and consider all donors as a single entity with utility derived from a Kickstarter project's success.

3.3.2 Initial Goals

The most direct measure of success for Kickstarter projects is initial goal achievement. A project's prosocial benefit cannot be realized if this initial goal is not met. Most of the existing RBC literature focuses on success as a binary for whether a project's initial goal was met. The most in-depth description of behavior around initial goals comes from Deb et al. (2021) who describe how donors dynamically influence buyers to always buy such that project success is achieved with minimal donations.

Deb et al. (2021) present strong evidence of this behavior, specifically a missing mass plot showing a distinct lack of projects that finished just below their initial goal. They show convincingly that these projects get pushed above the threshold by donors. However, projects that succeed tend to meet their goal in the first quarter of their campaign on average and the literature thus far has failed to explore what occurs in the remaining three quarters of successful campaigns. We continue to discuss the goals that come after the initial goal, stretch goals.

3.3.3 Stretch Goals

Success can go beyond the initial goal. Kickstarter focuses on bringing new products to market. These products lack existing form and have room for expansion in design, function, and/or scope (e.g., additional colors, alternative designs, compatibility with extra devices, improved battery). Backers are highly involved in this improvement process and offer feedback throughout the campaign (Kuppuswamy and Bayus 2017; Candogan et al. 2021). We define a product improvement as a project's addition of more to their product than envisioned at the campaign's onset. These product improvements improve the social benefit of the proposed product by expanding the product's appeal to more consumers or increasing its value for those to whom it already appeals.

It is well observed that projects implement new features to their products throughout their campaign.[37] These are especially prevalent after the initial goal has been achieved. An example is

[37] These can be encouraged through backer to entrepreneur communication (Candogan et al. 2021).

shown below in Figure 3.2 where an entrepreneur announced a product improvement through Kickstarter's updates feature.

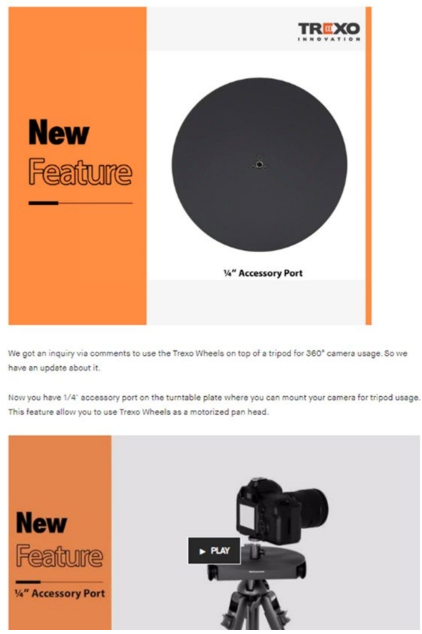

Figure 3.2: An Example of a Product Improvement

Product improvements are not always feasible. New features are costly to implement, and this cost cannot always be overcome through an accompanying demand increase. Instead, entrepreneurs often appeal to donors to support these improvements to overcome positive net costs through a *stretch goal*. Stretches are set throughout a campaign and indicate capital goals that must be met to implement the accompanying product improvement. However, these stretch goals are unlike initial project goals and lack any contractual enforcement mandating the promised product improvement. If a stretch goal is achieved there is no obligation for the entrepreneur to implement the new feature. This provides ample opportunity to cheat and mislead backers as has been well documented behavior on Kickstarter since its inception; we will discuss this in more depth in subsection 3.3.4.

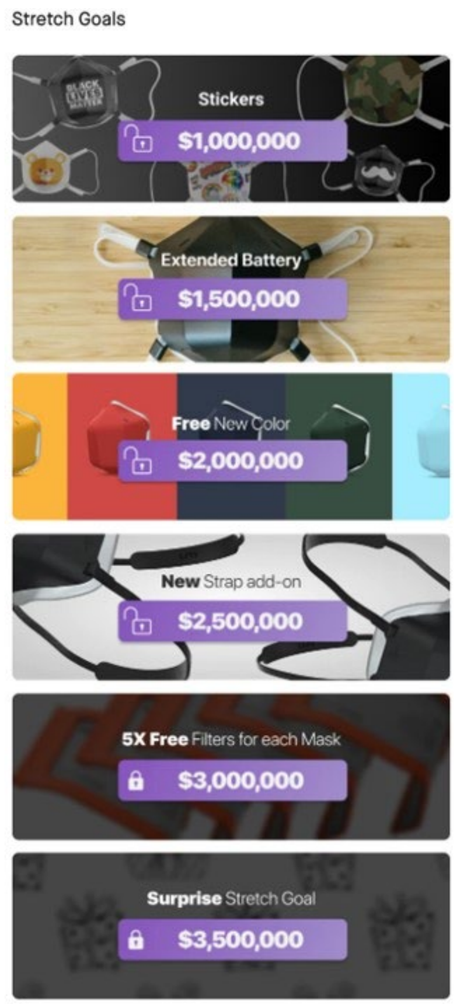

Figure 3.3: An Example of a Pre-Stretch Set of Stretch Goals

A sample set of stretch goals is shown in Figure 3.3. Stretch goals can be created during a campaign or before the campaign starts. Typically, a stretch goal created during a campaign occurs after or near initial goal achievement. We call those stretch goals set prior to the beginning of a campaign, pre-stretches and the entrepreneurs that implement them pre-stretchers. These pre-stretches tend to come in sets of potential improvements whereas stretch goals set during the campaign are more commonly singular.

Stretch goals are usually announced through the “updates” feature or graphics uploaded to a project's homepage. Though most commonly stretch goals offer product improvements, they can also promise add-ons to already purchased rewards or be simple monetary figures proposed by the entrepreneur (e.g., “Let’s see if we can hit $2 million”). We do not distinguish between these types of stretches. Offered add-ons pose the same contribution to society as a direct product improvement might, although, buyers may react differently in the two cases. And solely monetary goals pose no social contribution but instead create a target whose achievement may give donors some positive utility. We perform a separate analysis for donor response to simple round number targets (e.g., $100k or $1m) to distinguish that behavior from the prosocial behavior driven by product improvements or add-ons in section 3.4.

3.3.4 Retaliation Against Deception

Due to the nature of entrepreneurship, RBC projects may fail to deliver their product even after a successful campaign. Kickstarter reports that over 9% of successful projects fail to deliver rewards. In the event of this failure, backers’ funds are not returned so long as the entrepreneur provides “sufficient” evidence of exclusive use of funds towards product creation.[38] Unfortunately, this policy provides ample space for entrepreneurs to deceive backers into believing they have properly used their funds and make off with pure profit. The biggest RBC scandals have provided evidence of this deceptive behavior where entrepreneurs report fake problems which prevent further development of products (e.g., iBackPack and CST-01 which each raised over $1 million),

[38] This policy is outlined in Kickstarter's terms of use.

design a fake product to begin with (e.g., Skarp Laser Razor which raised over $4 million), or disappear altogether/file for bankruptcy after collecting backers' funds (e.g., Zano Autonomous Drone which raised over $3 million).

However, retaliatory measures exist to combat this deception. Kuppuswamy and Bayus (2018) show that backers may mutually withdraw funds if they find evidence of entrepreneurial deception prior to the end of the campaign and that they socially learn to detect this. Though, many campaigns slip past this backer-policing, in response, backers have two additional channels to penalize deception. Backers can engage in class action lawsuits to punish entrepreneurs who failed to deliver on the promises made during their campaign. And Blaseg et al. (2020) show that backers who feel deceived flood products with negative reviews which hurt long-term sales.

Lawsuits are the most direct retaliation method but are costly to backers and usually only manage to recoup a fraction of their funds after a long wait. And negative reviews provide no recompense to backers but have proved successful in penalizing entrepreneurs (Viotto da Cruz 2018).

Backers can impose a penalty on deceptive entrepreneurs; however, it is limited in scope and costly to both parties. The limited degree to which backers can penalize also limits the stretches they can support. However, in practice, the existence of these penalties does not prevent all deception as our model predicts. We can address this concern by introducing uncertainty in the penalty that can be imposed in our model. When donors are uncertain of the monetary penalty they can impose they risk supporting deceptive goals yet would still donate in the confines of our model.

In the following section we further detail the contribution game between entrepreneurs and donors with a panel dataset from Kickstarter and reduced-form analysis.

3.4 Happenings After the Goal: Answer from the Great Beyond

"It's the real performers who can hold people's attention and keep them from walking away." –Andy Kaufman

We introduce a panel dataset of RBC projects on Kickstarter to explore how the campaign functions beyond the initial goal has been met. This phase is centered around the stretch goal communication described above and is characterized by the game developed in section 3.2. We use our data to evidence our model and test its implications.

3.4.1 Data Collection

From May 2020 to December 2020, we collected data from 720 Kickstarter projects in the technology category through a handmade python web scraping script and accumulated 33,085 project-days of data. For each project-day we scraped the amount raised, number of backers, price and associated backers in each reward tier, and project level attributes (e.g., goal, deadline, location). Information on when a project stretched and the goal amount was added by hand.[39] If a project announced a monetary goal in their updates or on their homepage it was recorded as a stretch goal.\footnote{data entry and scraping occurred at 7am each morning} If projects set a sequence of goals as shown in Figure 3.3, we consider their current goal as the lowest goal not yet achieved for the purpose of all graphs and analysis.[40]

We used our data to separate buyer and donor revenue. First, we compiled the amount spent on rewards to get the buyer revenue by taking the price of each reward tier and multiplying it by

[39] These goals often appear in pictures and thus no efficient scraper could collect them so we had to regress to old techniques with hand entry.
[40] Projects that posted a set of stretch goals in the first day of their campaign were marked as pre-stretchers.

the number of new backers in that tier each day. Then we deduct this buyer revenue from the total amount raised on any given day to obtain the donor revenue.[41]

We cleaned our data by removing projects that were cancelled prior to their deadline and projects that had excessive pull out of donor or buyer revenues (more than a $1,000 drop in either value in at least one day, 6% of our sample). Backers often learn of deception by the entrepreneur during the campaign and withdraw funds enmasse (Kuppuswamy and Bayus 2018). Though this is evidence of the penalty noted in section 3.3.4, we separate it from the rest of our statistics and analysis because it is a separate behavior from the normal contribution game.

We use percentage daily revenue as a key outcome measure for figures---daily revenue divided by total revenue by the end of the campaign---to limit the influence of over-achieving projects without excluding them. And to differentiate the number of daily buyers and donors, we use arriving backers and the daily buyer and donor revenue. Because buyers can be to some extent a donor, we assume arriving buyers and donors to be a continuum rather than discrete. We observe the total number of backers that arrived each day not individual arrivals, so we lack sensitivity in the exact contribution of each individual. We instead extrapolate by taking the daily backers and multiplying it by the ratio of daily donor revenue and daily total revenue to calculate the number of arriving donors. The number of arriving buyers is daily backers minus daily donors. For example, if 5 backers arrived and $10,000 was raised, $1,000 of which came from donations we would measure $5 \times \frac{1{,}000}{10{,}000} = 0.5$ arriving donors and 4.5 arriving buyers.

[41] Shipping costs count towards amount raised but not the buyer revenue which may introduce error to our donor revenue, however U.S. to U.S. based shipping is free through Kickstarter and our sample only includes U.S. based projects, so this is of negligible concern.

3.4.2 Model-Free Exploration

We first examine behavior around the initial goal to demonstrate consistency with past results. Our data collection occurred during the COVID-19 pandemic so it is important verify consistent behavior. We see simple targeting behavior around initial goals in Figure 3.4. There is distinct missing mass prior to 100% goal attainment in a probability distribution function of all project's goal attainment at their deadline which is consistent with past work and indicates donors target initial goal achievement (Dai and Zhang 2019; Deb et al. 2021).

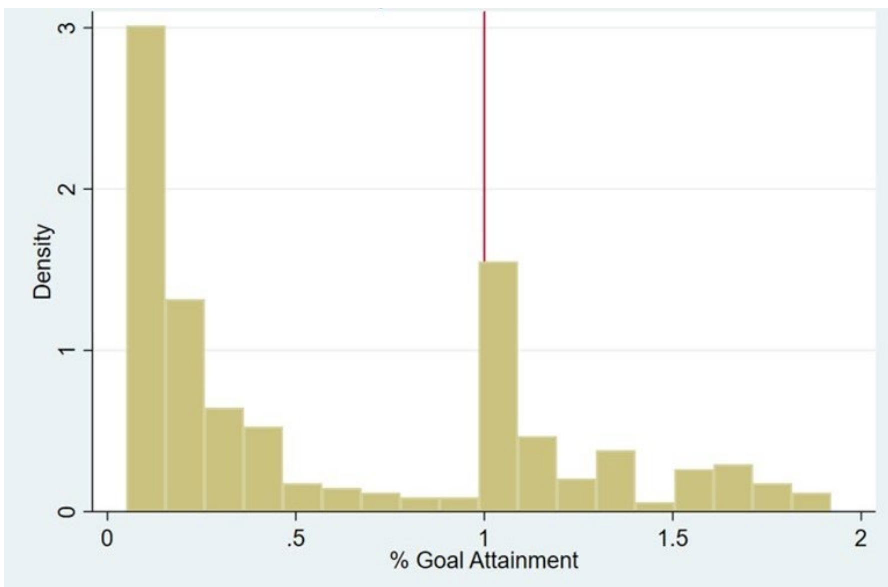

Figure 3.4: Histogram around Goal Attainment

Note. Projects are sorted by initial goal attainment at the end of their campaign. Density is shown to be larger than 1 because bins are horizontally smaller than 1.

Although donors evidently target initial goal achievement and donate primarily to push projects over this threshold, donations continue well beyond the initial goal's achievement albeit in diminished capacity. Much like their interaction with buyers, the donor tends to signal her value of improvement to entrepreneurs as soon as the campaign hits its initial goal. We argue that the donor broadcasts her wealth to encourage mutually beneficial stretches. Figure 3.5 shows daily donor revenue over the relative time of the campaign (1 is the deadline) for projects without an active stretch goal after they have achieved their initial goal. Because time is limited, yet the potential for product improvement is unlimited post-goal achievement, donors value time and provide a large signal of value to encourage stretch goals which decays as less time remains to successfully stretch and the entrepreneur fails to announce a stretch indicating they lack the ability to do so. Donors may also communicate directly with entrepreneurs through the "comments" feature on Kickstarter which is akin to our noisy signals of donor value throughout the campaign. In application value communication is a mixture of the costly donation signal shown and Figure 3.5 and direct communication through "comments."

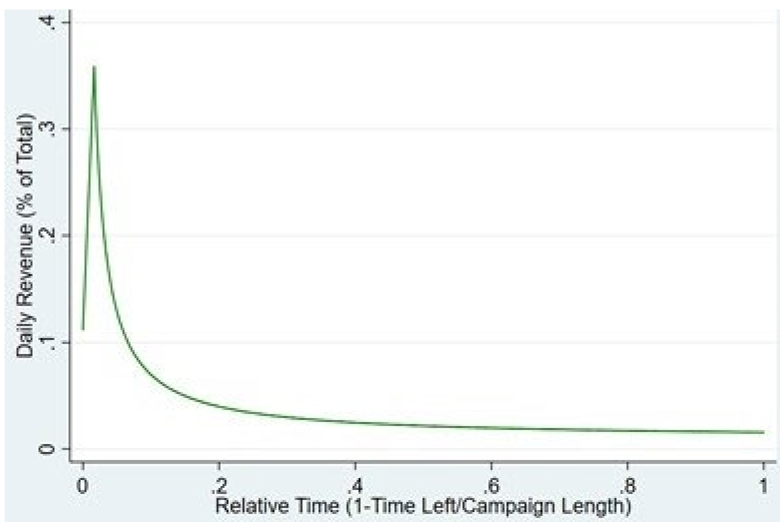

Figure 3.5: Donation Signaling after the Goal

Note. The polynomial fit of donations for projects that never stretch and have exceeded their goal by the relative time remaining in the campaign.

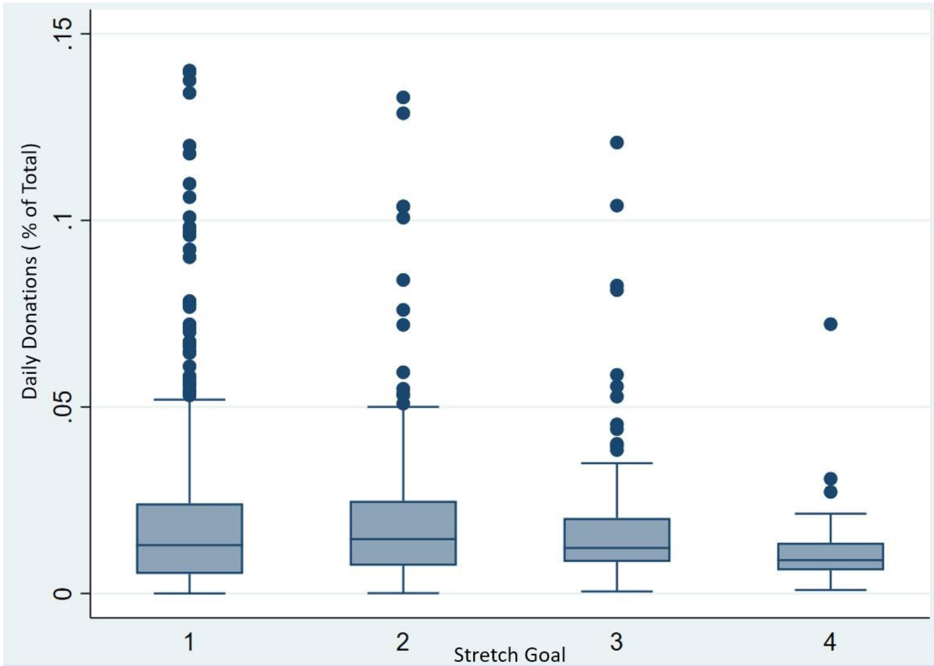

Figure 3.6: Attrition in Donor Activity by Number of Stretches

Note. Projects that stretch are sorted by number of their current stretch goal and a box and whiskers plot is constructed for each set of daily observations

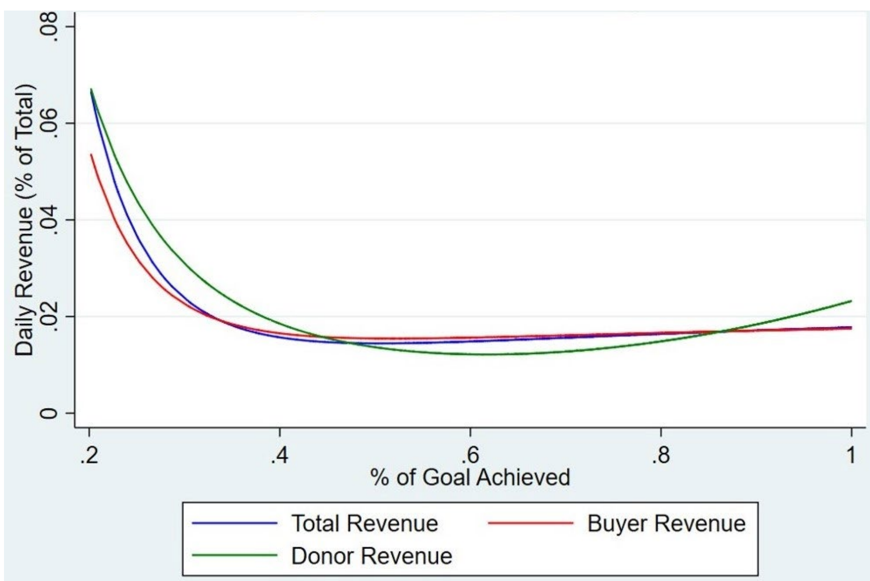

Figure 3.7: Trends in Daily Revenue approaching Stretch Goal

Note. This figure presents a polynomial fit of daily revenue (as a percentage of the highest observed value for each revenue type) by the percentage attainment of the current stretch goal.

Entrepreneurs commonly launch multiple stretch goals (and the first stretch goal itself is the second goal set in the campaign), Figure 3.6 shows attrition in support for sequential stretch goals. The average daily donation towards a stretch goal declines as the order of appearance increases. The variance in the distribution also declines. In statistical mechanics, a population's measure of energy---the inverse of entropy---is increasing in variance of individual states. Thus, a decrease in variance measures a decline in energy, also known as attrition (Sen 2002). This attrition is interesting in its own right but also supports our model's assumption of strict ordering in the benefits of stretching.

Figure 3.7 shows the trend in percentage daily revenue for buyers and donors separately as projects move from a small percentage of stretch goal attainment, measured as $\frac{amount\ raised}{stretch\ goal}$, towards the stretch goal's threshold. We use a polynomial fit of our data to show these trends. Donors clearly react to stretch goals and increase donations to meet them. Whereas buyers do not respond.

3.4.3 Summary Statistics

Table 3.1 presents our project level summary statistics. These statistics are reported after the removal of projects which had mutual fund withdrawal. These constitute 6% of our projects, 30% of which stretched supporting the observations of Kuppuswamy and Bayus (2018) and evidencing a penalty for deception. Of the remaining projects, we find that only 16% stretch. We separate this group into those that stretch late into a campaign and those that pre-stretch. Charness (2000) shows that individuals are more trusting of communication when talk occurs before action and thus we suspect pre-stretchers should have more trust from backers.

Projects that never stretch have lower success rates than projects that stretch at some point and those that pre-stretch due to innate quality differences. And projects that never stretch also seem to have a higher share of donations than late stretchers but not pre-stretchers. We describe donation share below with two values, the donation revenue is simply the average value of donation revenue across all projects for each stretch type and donor share is the average value of donation revenue divided by total revenue across all projects. The latter removes skewness caused by high performing projects. We use donor revenue relative to total revenue as our primary measure to discuss donations.

Table 3.1: Summary Statistics: Projects

Stretch Goal			Goal		Total	Donor	
Type	Amount	N	Amount	Success	Revenue	Revenue	Share
No Stretch			38,303	0.26	17,898	2,137	0.35
623 Projs			(65,262)	(0.44)	(128,353)	(7,178)	(0.34)
Stretch	363,583	1.99	20,951	0.90	276,754	20,346	0.16
49 Projs	(577,715)	(1.81)	(22,500)	(0.31)	(733,824)	(44,383)	(0.17)
Pre-Stretch	260,419	3.30	24,310	0.54	71,114	8,736	0.31
48 Projs	(334,911)	(3.23)	(26,052)	(0.50)	(211,513)	(21,517)	(0.31)

Standard deviations in parentheses

We are interested in behavior beyond the initial success of projects; thus, we condition our statistics on project success as shown in Table 3.2. In this frame, projects that stretch suffer a smaller decline in donations relative to those that do not. We expect the general decline in donations because they are heavily motivated by initial goal achievement and if their support is not needed for this, there is less incentive to donate.

Table 3.2: Summary Statistics: Successful Projects

Stretch Goal			Goal	Total	Donor		Rel. Time
Type	Amount	N	Amount	Revenue	Revenue	Share	Met Goal
No Stretch			10,715	64,633	6,778	0.261	0.370
163 Projs			(13,577)	(245,423)	(12,678)	(0.242)	(0.386)
Stretch	356,353	1.74	17,884	307,454	22,452	0.159	0.183
44 Projs	(647,929)	(1.75)	(21,398)	(769,194)	(46,411)	(0.174)	(0.292)
Pre-Stretch	300,203	1.90	14,524	127,053	15,461	0.244	0.305
26 Projs	(362,471)	(2.49)	(16,635)	(277,381)	(27,696)	(0.239)	(0.362)

Standard deviations in parentheses

Here we find evidence of our model's first testable implications, that donations are sensitive to the penalty they can impose and if stretches exceed this value they will simply cease. After projects succeed if they set a stretch too high donations stop. For 12% of our sample that stretch, donations drop to $0 after a stretch. Further, the largest donations towards an initial goal was $283,411 (for a non-stretching project) but for a stretch goal the maximum donated was only $46,942 (donor revenue contributed after the initial goal).

Stretchers appear to set lower initial goals when compared across general project summary statistics. These projects likely seek to minimize their chance of failure and plan to later induce donations through stretch goals. Although this behavior is interesting, our model solely focuses on behavior after the goal and fails to capture any subtly here. This would be an interesting expansion on our work. Though, this pre-emptive behavior is not of large concern, as of successful projects, those that never stretch have the lowest goals.

Projects that stretch in the campaign achieve their initial goal much sooner than those that never stretch and pre-stretchers. This gives these campaigns longer to gather information on their backers and plan stretches. Further, an unexpected early achievement of ones' goal indicates an underestimation of demand. If demand is higher than expected it is more feasible to implement product improvements as the net cost of these improvements declines in demand.

Table 3.3: Summary Statistics: Daily

	Daily						
Stretch Goal Type	Total Revenue	Donor Revenue	Backers	Buyers	Donors	Post Goal	Post Stretch
No Stretch 29,731 Days	397 (7,659)	52 (518)	3.46 (66.2)	3.05 (64.75)	0.41 (5.04)	0.12 (0.32)	
Stretch 1,793 Days	7,684 (46,823)	628 (3,789)	45.06 (254)	41.39 (249)	3.67 (14.4)	0.72 (0.45)	0.09 (0.28)
Pre-Stretch 1,751 Days	2,648 (14,694)	359 (2,571)	20.8 (115)	18.3 (102)	2.49 (15.1)	0.38 (0.49)	0.05 (0.22)

Standard deviations in parentheses

We further explore our data with daily summary statistics shown in Table 3.3. They provide evidence for our model's second testable implication that beneficiaries set realistic goals which allow time to act in future periods. Those that stretch spend much more time after achieving their goal as shown in Table 3.2, but they spend similar time past their stretch goals as projects that never stretch spend beyond their initial goal (roughly 10% of the campaign's duration). The 10% limit seems to be the crucial final opportunity to act and entrepreneurs will set goals up to this point.

Table 3.4: Before and After Goal Difference

Stretch Goal Type	Daily Total Revenue	Daily Donor Revenue	Daily Backers	Daily Buyers	Daily Donors
No Stretch	-383.1***	56.08	-5.50***	-4.69***	-0.815***
	(145.63)	(56.84)	(1.13)	(0.98)	(0.24)
Stretch	-1,995.3*	-383.3*	-15.73*	-12.36*	-3.36**
	(1,052)	(190.38)	(8.5)	(7.14)	(1.58)
Pre-Stretch	-766.3	230.4	-9.85**	-7.63**	-2.22*
	(878.2)	(527.0)	(4.6)	(3.49)	(1.25)

Standard errors in parentheses

Table 3.4 provides a comparison of means for successful projects before and after their initial goal has been met. For all projects daily revenue, arriving backers, buyers, and donors drop, with the largest drop coming in projects that stretch. However, donor revenue only decreases for projects that stretch later in their campaign and actually increases for those that pre-stretch or never stretch. Those that stretch during a campaign have less trust of donors and the existence of an active stretch eliminates the need of a donor signal in this stage. It is of note that for some of these projects, the donor does eventually donate to meet the stretch goal premature of the campaign's deadline but if the project never nears their stretch goal then the donor does not participate and this drives the observed decline.

Next, we use a simple reduced-form analysis to study each backer groups' response to the announcement of mid-campaign stretch goals and nearing the goal in amount raised (for both stretches and pre-stretches). We contrast these results with backer response to whole numbers that may act as targets in the same framework. Further we use two reduced-form analysis to test our model's implications. To gauge sensitivity in donations to some penalty we test whether donations stop once a stretch is set too high. And to measure entrepreneur's goal setting to time we test whether goals set earlier in a campaign are more likely to be smaller than goals set later.

3.4.4 Reduced-Form Analysis

To explore backer reaction to stretch goals we look at the change in daily total, buyer, and donor revenue to the announcement of a stretch goal $stretch_{it}$ and the proximity to set stretch goals $\left(1 - \frac{funds_{it}}{stretch\ goal_{it}}\right)^{-1}$ The latter term is a function with a discontinuous negative jump at $funds_{it} = stretch\ goal_{it}$, so if one group donates more to push projects over these thresholds we should see a significantly negative coefficient β_2. We use the following fixed effects structure to study this relationship. We control for relative time and time squared during campaign T (e.g., 10% of the campaign remaining and 10% squared), day of the year fixed effects γ_t, project level fixed effects α_i and u_{it} is econometric error.

$$Y_{it} = \beta_1 stretch_{it} + \beta_2 (1 - \frac{funds_{it}}{stretch\ goal_{it}})^{-1} + \delta T + \alpha_i + \gamma_t + u_{it} \quad (3)$$

We estimate the model with an Arellano-Bond estimator to eliminate the inter-temporal endogeneity of the second term in STATA with default settings (StataCorp 2019). Results are shown below in Table 3.5. Donors significantly react to the announcement of stretch goals and their proximity. Initiating a stretch goal increases daily donations by 1.9% while the stretch is active and around the goal achievement, donations jump up to a 66% increase on average.[42] Buyers do not have a significance response to stretch goals.

Table 3.5: Revenue at Stretch Goal Approach

Dependent Variable: % Daily Revenue			
	Total	Buyer	Donation
Announced Stretch Goal	0.00887	-0.00206	0.0192**
	(0.0073)	(0.0073)	(0.0083)
$(1 - \frac{Funds}{Stretch\,Goal})^{-1}$	-5.78e-08***	-3.77e-08	-1.27e-07***
	(1.28e-08)	(3.25e-08)	(3.43e-08)
% of Time Left	-0.474***	-0.419***	-0.507***
	(0.0151)	(0.0161)	(0.0162)
% of Time Left Squared	0.329***	0.285***	0.355***
	(0.0106)	(0.0117)	(0.0114)
Observations	33,085	25,869	33,085
Adjusted R^2	0.177	0.145	0.175

Standard errors in parentheses
Note. Sales revenue has less observations than total and donation because some projects never sold a single product and where thus dropped by the fixed effect analysis.
* p<0.10, ** p<0.05, *** p<0.01

We extend this estimation framework to test if backers have a response to whole numbers in general and not specifically to stretch goals (which are commonly set at whole numbers). We create "placebo" stretches. which we call targets, for all projects at $10k, $50k, $100k, $500k, and $1m. At this stage it is useful to remind readers that our analysis uses the lowest yet to be achieved

[42] This figure was determined by taking the average change in $\left(1 - \frac{funds_{it}}{stretch\ goal_{it}}\right)^{-1}$ as projects passed the threshold and multiplying by the reported β_2.

goal as the current goal for all analysis. If a project has currently raised $40k their assigned target will be $50k the day they surpass $50k to let's say $57k they will still have their target registered as $50k for the purpose of this analysis which captures the discontinuity around goal attainment. The day after the target has been achieved it is registered as the next highest, $100k in this context.

Table 3.6: Revenue at Approach of Whole Numbers

Dependent Variable: % Daily Revenue			
	Total	Buyer	Donation
1/(1-%Target)	-1.85e-06	-2.77e-06	-1.21e-06
	(-3.22e-06)	(-5.87e-06)	(-2.35e-06)
% of Time Left	-0.473***	-0.419***	-0.505***
	(0.0150)	(0.0161)	(0.0161)
% of Time Left Squared	0.329***	0.285***	0.354***
	(0.0106)	(0.0117)	(0.0114)
Observations	33,085	25,869	33,085
Adjusted R^2	0.177	0.145	0.174

Standard errors clustered by project in parentheses
Note. The fixed effect regression of the response of the overall amount raised for projects (Total), the buyer revenue raised for projects (Buyers), and the donations raised (Donors) on random targets to examine potential response to achieving whole numbers rather than stretch goals. Targets used: $10k, $50k, $100k, $500k, $1m. Projects with stretch goals were excluded.
* p<0.10, ** p<0.05, *** p<0.01

We also exclude the indicator $stretch_{it}$, since we now assume all projects have the same set of targets and are always under their influence. The estimation results in Table 3.6 reveal no significant whole number affixation from any group. All β_2 are far from significant whereas before the coefficient was an order of magnitude higher for donors relative to buyers and highly significant.

These results show that donors respond positively to the announcement of a stretch goal and seek to contribute towards that goal specifically. We next explore our model's testable implications through reduced-form analysis.

Table 3.7: Likelihood of Donations Stopping by Stretch Goal Amount

Dependent Variable: Probability of No Donations	
Stretch Goal Amount	-3.71e-6***
	(6.7e-5)
Stretch Goal Amount Squared	8.17e-13***
	(7.3e-12)
% of Time Left	-1.711**
	(0.7456)
% of Time Left Squared	3.341***
	(0.8366)
Observations	2,188

Standard errors clustered by project in parentheses
Note. A logit regression of the likelihood of donations being less than or equal to 0 in a given period on the size of active stretch goal and the size squared with controls for campaign time.
* p<0.10, ** p<0.05, *** p<0.01

Our model predicts that donations will cease if a stretch goal is set too high (above the penalty). We explore this with a logit model regression of the likelihood of receiving zero or less donations (if donors withdraw funds) on the stretch goal amount and the amount squared controlling for relative time remaining in the campaign. The results, presented in Table 3.7, reveal that donations are more likely to stop as the stretch goal increases to very high values which is consistent with our model.

Table 3.8: Size of Stretch Goals by Campaign Time Remaining

Dependent Variable: Stretch Goal Amount	
% of Time Left	-268,292.4***
	(8.85e+04)
% of Time Left Squared	123,863.9
	(9.08e+04)
Buyer Revenue	1.542***
	(0.1030)
Donor Revenue	-1.133
	(1.0216)
Stretch Goals	264
Adjusted R^2	0.617

Standard errors clustered by project in parentheses
Note. A regression of stretch goal amount on time left in the campaign controlling for initial goal amount, buyer and donor revenue.
* p<0.10, ** p<0.05, *** p<0.01

Our model also predicts that the size of stretch goals should increase over time. This is due to shrinkage in the value of remaining time in the game which acts as a marginal cost when entrepreneur evaluate a stretch goal amount. The more time remaining in a campaign, the more valuable is free action so entrepreneurs set lower stretch goals initially to ensure the ability to act in the future. To capture this, we perform the following linear regression of the stretch goal's amount on the relative time remaining in the campaign shown in a quadratic form and controlling for the project's initial goal amount and their buyer and donor revenue prior to the goal being set. The resulting regression, shown in Table 3.8, reveals an increase in the stretch goal amount set with time as expected. The linear term is negative meaning as campaign time passes, the percentage of campaign time remaining decreases and the average stretch goal set increases by \$26.8k for every 10% of the campaign that elapses. The quadratic term is insignificant. It also interesting to note that stretch goals are typically set 1.5x above the current buyer's revenue and irrespective of current donor revenue. This indicates a lack of interest in donors for funds already invested to achieve previous goals, meaning each new goal seems to offer benefits independent of the last to donors.

This reduced-form evidence supports our model and gives us freedom to use derived interpretations to guide policy. In the next section we express this policy guidance and conclude.

3.5 Conclusion

We establish a simple theory model to explore cheap talk communication of charitable goals. We find that this cheap talk communication is informative and welfare enhancing when there exists a large enough penalty for deception. We use data from RBC to further explore this type of communication; called stretch goals.

The data show a significant positive response of donors to this communication, attrition in donor attentiveness to multiple goals, evidence of a penalty for deception and validate our model's testable implications. These results shed light on beneficiary-to-donor communication when there is incentive to deceive and extend the RBC literature to the contribution game beyond the goal.

3.5.1 Policy Implications

Our model demonstrates that costless communication between donors and their beneficiaries is beneficial to both parties when the punishment for deception is sufficiently high. Increasing this penalty strictly increases the payoffs of all players and is an admissible strategy for policy makers. This advice applies beyond RBC to general setting where similar communication exists. The cost of misleading those with prosocial intentions needs to be increased drastically such that the behavior is allowed to exist in higher sums and enable more prosocial ventures. This can be accomplished with additional laws, fees, or policing which restrict a beneficiary's ability to deceive his donors or penalizes him for failing to satisfy donors.

We also show that more information flow is beneficial to both parties. In our data we show that donors actively pay costly signals to demonstrate wealth to beneficiaries despite the beneficiary's incentive to deceive. However, donors are limited by not just their ability to communicate, but also their knowledge of the value of prospective prosocial ventures. The accessibility and popularity of crowdfunding ventures as a whole exists to fill this gap and provide knowledge to a wide swath of individuals as to a variety of prosocial projects that they may support. Kickstarter specifically has worked to make the communication of beneficiaries more feasible; but in several cases this has backfired (Blaseg et al. 2020; Kaisen and Zantedeschi 2021).

A combination of communication expansions and penalties for deception is the best practice for policy makers dealing with charitable ventures.

3.5.2 Extensions

Communication similar to stretch goals exist in entrepreneurial communications to shareholders (also called stretch goals), company advertisements to consumers (challenges to purchase some amount of the product to trigger some mutual benefit), and charitable events (secondary goals are common in charity auctions, fundraisers, and quarterly reports). These capacities likely exhibit similar behavior to that observed in online RBC and would expand the external validity of claims derived from our results.

3.6 References

Agrawal A, Catalini C, and Goldfarb A (2015) Crowdfunding: Geography, social networks, and the timing of investment decisions. *Journal of Economics & Management Strategy* 24(2):253–274.

Andreoni J (2006) Chapter 18 philanthropy. In Kolm S-C, and Ythier JM editors, Applications, volume 2 of *Handbook of the Economics of Giving, Altruism and Reciprocity* 1201–1269.

Arrow KJ (1950) A difficulty in the concept of social welfare. *Journal of Political Economy* 58(4):328–46.

Aumann R (1990) Nash equilibria are not self-enforcing. Economic Decision-Making: Games, *Econometrics and Optimisation* 201–6.

Beersma B, Bekkers R, and van Teunenbroek C (2020) Look to others before you leap: A systematic literature review of social information effects on donation amounts. *Nonprofit and Voluntary Sector Quarterly* 49(1):53–73.

Blaseg D, Schulze C, and Skiera B (2020). Consumer protection on kickstarter. *Marketing Science* 39(1):211–33.

Candogan ST, Cornelius PB, Gokpinar B, Korpeoglu E, and Tang CS (2021). Product development in crowdfunding: Theoretical and empirical analysis.

Charness G (2000) Self-serving cheap talk: A test of aumann's conjecture. *Games and Economic Behavior* 33(2):177–194.

Crawford VP, and Sobel J (1982) Strategic information transmission. *Econometrica* 50(6):1431–1451.

Dai H, and Zhang DJ (2019) Prosocial goal pursuit in crowdfunding: Evidence from kickstarter. *Journal of Marketing Research* 56(3):498–517.

Deb J, Ory A, and Williams KR (2021). Aiming for the goal: Contribution dynamics of crowdfunding.

Grossman SJ (1981) The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics* 24(3):461–83.

Huck S, Rasul I, and Shephard A (2015). Comparing charitable fundraising schemes: Evidence from a natural field experiment and a structural model. *American Economic Journal: Economic Policy* 7(2):326–69.

Hutchison-Quillian J, Reiley DH, and Samek A (2018) Hassle costs and workplace charitable giving: Field experiments with google employees.

Kaisen J, and Zantedeschi D (2021) Stuck in the middle: Context effects in reward-based crowdfunding menus.

Kamenica E, and Gentzkow M (2011) Bayesian persuasion. *American Economic Review* 101(6):2590–2615.

Kuppuswamy V, and Bayus BL (2017) Does my contribution to your crowdfunding project matter? *Journal of Business Venturing* 32(1):72–89.

Kuppuswamy V, and Bayus BL (2018) Crowdfunding Creative Ideas: The Dynamics of Project Backers. *The Economics of Crowdfunding* 151–182.

Leamy E (2018) Backer beware: What to know before you support a crowdfunding product. *Washington Post*.

McFadden RD (1992) Small-change donations going to vendors, not charities, abrams charges. *The New York Times*.

Milgrom PR (1981) Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics* 12(2):380–391.

Scherhag C, and Boenigk S (2013) Different or equal treatment? donor priority strategy and fundraising performance assessed by a propensity score matching study. *Nonprofit Management and Leadership* 23(4):443–472.

Spence M (1973) Job market signaling. *The Quarterly Journal of Economics* 87(3):355–74.

StataCorp (2019) Stata Statistical Software: Release 16. College Station, TX: StataCorp LLC.

Vesterlund L (2003) The informational value of sequential fundraising. *Journal of Public Economics* 87(3):627–657.

Viotto da Cruz J (2018) Beyond financing: crowdfunding as an informational mechanism. *Journal of Business Venturing* 33(3):371–393.

Yoruk BK (2012) Do fundraisers select charitable donors based on gender and race? evidence from survey data. *Journal of Population Economic*, 25:219–43.

APPENDICES

Appendix A: Viral Videos by Total Views

We replicate the procedures of the main paper for an alternate definition of viral video, by total views, to demonstrate further robustness of results. We observe the same one week decline with this definition as before but the result is less robust to changes in definition. Each measure captures a different treated sample of channels (with only 50% mutual channels) and those that are mutually identified are captured as going viral in different periods. Figures A4 and A5 reproduce figures 1.4 and 1.5 for total views and table A3 includes sensitivity analysis for both measures.

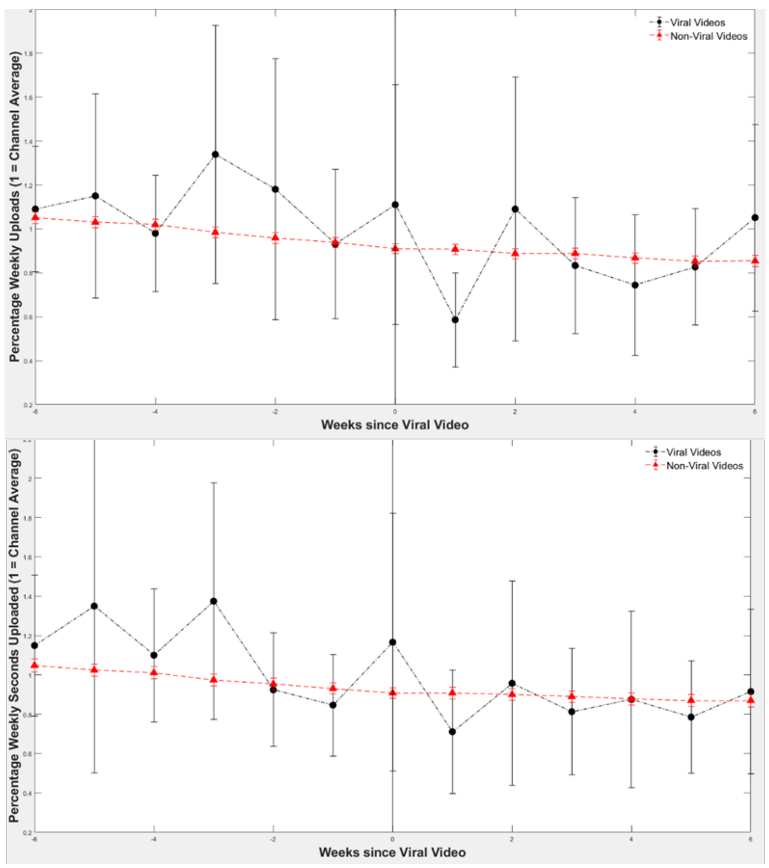

Figure A4: Model-Free Evidence of Income Targeting

Note. A raw comparison of means of production around viral event is shown for channels that go viral in week 0 and non-viral channels. All values are scaled such that 1 represents a channel's average weekly production for tractability. There is a persistent downward trend because we condition our sample selection on activity in the first week observed but not for any week after.
95% confidence interval shown with error bars.

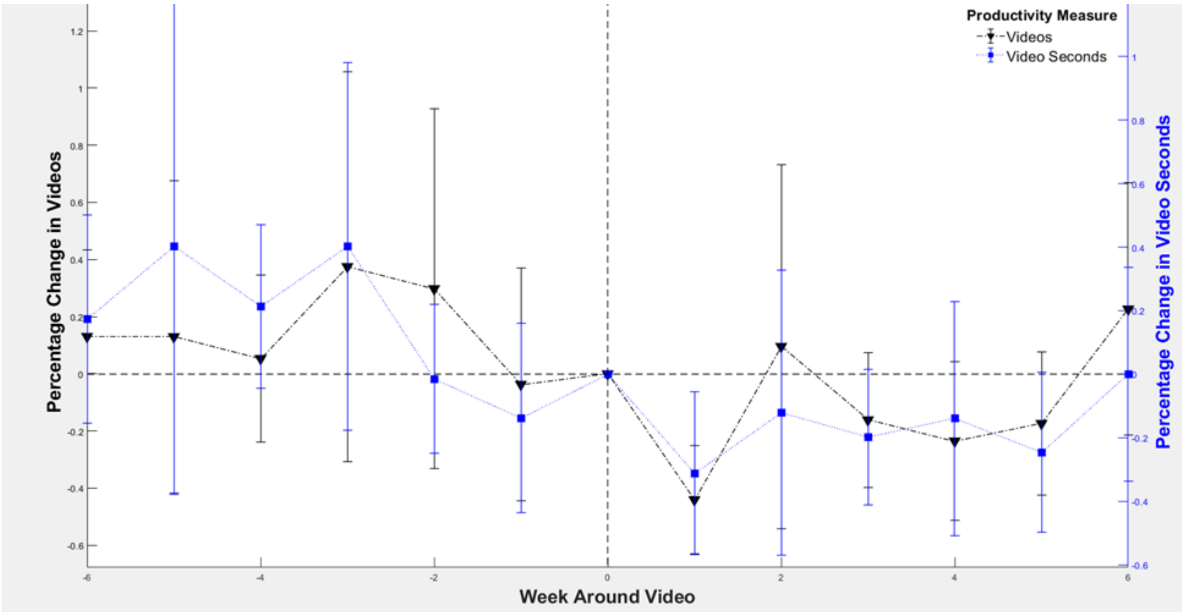

Figure A5: Evidence of Income Targeting

We regress with model 1 for two measures of productivity. The percentage change (by own channel's average) in videos uploaded (triangles and left axis) specification shows a significant 50% decline in production for the week after a channel goes viral. The percentage change in video seconds uploaded (squares and right axis) specification shows a smaller but still significant 40% decline in production.
95% confidence interval shown with error bars.

Table A3: Sensitivity Bounds in One-Week Productivity Decline

Weekly Views			Total Views		
	1 week drop			1 week drop	
Viral Threshold	Videos	Video Seconds	Viral Threshold	Videos	Video Seconds
Weekly Views>10k	0.0776	0.137	Total Views>10k	0.105**	0.110*
301 vids/119 chan	(0.0838)	(0.0893)	771 vids/210 chan	(0.0490)	(0.0580)
Weekly Views>25k	0.0445	0.142	Total Views>50k	-0.00970	-0.0235
186 vids/69 chan	(0.0872)	(0.115)	241 vids/96 chan	(0.0571)	(0.0684)
Weekly Views>50k	-0.150*	-0.0400	Total Views>100k	-0.0383	0.0541
95 vids/45 chan	(0.0781)	(0.122)	159 vids/76 chan	(0.0789)	(0.100)
Weekly Views>75k	-0.0900	-0.0499	Total Views>200k	-0.190	-0.149
64 vids/33 chan	(0.111)	(0.148)	76 vids/41 chan	(0.135)	(0.145)
Weekly Views>100k	-0.307***	-0.240*	Total Views>300k	-0.0449	0.00839
59 vids/29 chan	(0.119)	(0.134)	57 vids/32 chan	(0.128)	(0.168)
Weekly Views>125k	-0.463***	-0.361**	Total Views>400k	0.141	0.111
52 vids/25 chan	(0.103)	(0.147)	48 vids/30 chan	(0.140)	(0.146)
Weekly Views>150k	-0.534***	-0.419***	Total Views>500k	-0.441***	-0.312**
41 vids/21 chan	(0.101)	(0.120)	31 vids/ 22 chan	(0.0969)	(0.156)
Weekly Views>200k	-0.230*	-0.0577	Total Views>600k	-0.219	-0.0252
28 vids/15 chan	(0.123)	(0.149)	29 vids/20 chan	(0.151)	(0.229)
Weekly Views>250k	-0.374***	-0.0675	Total Views>700k	-0.255	-0.390**
22 vids/13 chan	(0.124)	(0.176)	24 vids/17 chan	(0.174)	(0.159)
Weekly Views>300k	-0.419***	-0.112	Total Views>800k	0.143	0.396
16 vids/12 chan	(0.147)	(0.294)	21 vids/14 chan	(0.175)	(0.420)
Weekly Views>350k	-0.326**	0.0528	Total Views>900k	-0.447**	-0.172
14 vids/11 chan	(0.157)	(0.311)	17 vids/13 chan	(0.193)	(0.298)
Weekly Views>400k	-0.407**	0.0259	Total Views>1m	0.108	0.322
12 vids/10 chan	(0.168)	(0.357)	14 vids/11 chan	(0.280)	(0.442)

Standard errors clustered by channel in parentheses

Note. The above table presents a sensitivity analysis of β_1 from regression (1) by viral specification from a range of 10k to 400k weekly views.

* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Appendix B: Demand Bias in Neoclassical Menus

$$\begin{aligned} Q_{tj} &= F(q_{tj}, p_{tj}, q_{t-1j}, p_{t-1,j}, q_{t+1j}, p_{t+1,j}) \\ &= \alpha p_{tj} + \lambda_1 p_{t-1,j} + \lambda_2 p_{t+1,j} \end{aligned} \tag{1}$$

Theorem 1. *The neoclassical demand bias λ_2 is strictly positive.*

We consider optimal and non-optimal menus; in both scenarios we prove λ_2 is strictly positive.

B.1 Optimal Menus

An optimal menu allocates price and quality options such that the more isolated an option in price the more consumers self-select into it which generates a positive λ_2. When only prices are observable and firms consider both price and quality when designing menus, the more isolated a menu option from its higher priced neighbor the more demand for said option in a neoclassical framework.

Theorem 2. *In a three option menu an endogenous increase in the highest-priced option increases demand for the middle-priced option.*

B.1.1 preparatory material. We follow the set-up of Mussa and Rosen (1978) but deviate to a menu of discrete options rather than a continuous spectrum. Consider a profit-maximizing firm with three menu options---high quality (H), medium quality (M), and low quality (L)---and marginal cost a convex function of quality $c(q)$ where $c'(q) > 0$ and $c''(q) > 0$.[43] Consumer utility is given by $U(\theta, q, p) = \theta q - p$ where θ is heterogeneous preference for quality whose distribution---known to firms---$F(\theta)$ is continuous and strictly increasing for all positive values of θ, and 0 otherwise. We define $f(\theta) \equiv \frac{\partial F(\theta)}{\partial \theta} > 0$ over all positive values of θ.

[43] Only a convex cost function produces a menu with divergent menu options. If costs are concave, one option at $q = \infty$ uniquely maximizes profit (and utility). If costs are linear, the choice of quality and price are arbitrary.

Firms choose a set of quality options $q = \{q_L, q_M, q_H\}$, and accompanying prices $p = \{p_L, p_M, p_H\}$ to maximize profit. This selection is subject to strict ordering of price and quality from low to high such that $q_L < q_M < q_H$ and $p_L < p_M < p_H$.[44]

Lemma 2.1. *If consumers can only choose one menu option, then they self-select into utility maximizing options such that consumers with* $\theta \in \left[\theta_{HM} \equiv \frac{p_H - p_M}{q_H - q_M}, \infty\right)$ *purchase the high quality option, those with* $\theta \in \left(\theta_{ML} \equiv \frac{p_M - p_L}{q_M - q_L}, \theta_{HM} \equiv \frac{p_H - p_M}{q_H - q_M}\right)$ *the medium quality option, those with* $\theta \in \left[\theta_L \equiv \frac{p_L}{q_L}, \theta_{ML} \equiv \frac{p_M - p_L}{q_M - q_L}\right]$ *the small quality option, and all others do not purchase.*

Proof of Lemma 2.1. Trivial.

The firm's problem is shown below.

$$\begin{aligned} \max_{p,q} \quad & [1 - F(\theta_{HM})](p_H - c(q_H)) + [F(\theta_{HM}) - F(\theta_{ML})](p_M - c(q_M)) + [F(\theta_{ML}) - F(\theta_L)](p_L - c(q_L)) \\ s.t. \quad & 0 \leq p_L < p_M < p_H, \quad 0 \leq q_L < q_M < q_H \end{aligned} \tag{1}$$

The first order conditions for price yield the following relationships

$$\begin{aligned} p_L = \quad & q_L[\frac{1 - F(\theta_L)}{f(\theta_L)}] + c(q_L) \\ p_M = \quad & (q_M - q_L)[\frac{1 - F(\theta_{ML})}{f(\theta_{ML})}] + q_L[\frac{1 - F(\theta_L)}{f(\theta_L)}] + c(q_M) \\ p_H = \quad & (q_H - q_M)[\frac{1 - F(\theta_{HM})}{f(\theta_{HM})}] + (q_M - q_L)[\frac{1 - F(\theta_{ML})}{f(\theta_{ML})}] + q_L[\frac{1 - F(\theta_L)}{f(\theta_L)}] + c(q_H) \end{aligned} \tag{2}$$

The first order conditions for quality yield the following relationships

$$\begin{aligned} c'(q_L)[F(\theta_{ML}) - F(\theta_L)] = \quad & \frac{\theta_{ML}}{q_M - q_L} f(\theta_{ML})[(p_L - c(q_L)) - (p_M - c(q_M))] + \frac{\theta_L}{q_L} f(\theta_L)[p_L - c(q_L)] \\ c'(q_M)[F(\theta_{HM}) - F(\theta_{ML})] = \quad & \frac{\theta_{HM}}{q_H - q_M} f(\theta_{HM})[(p_M - c(q_M)) - (p_H - c(q_H))] \\ & + \frac{\theta_{ML}}{q_M - q_L} f(\theta_{ML})[(p_M - c(q_M)) - (p_L - c(q_L))] \\ c'(q_H)[1 - F(\theta_{HM})] = \quad & \frac{\theta_{HM}}{q_H - q_M} f(\theta_{HM})[(p_H - c(q_H)) - (p_M - c(q_M))] \end{aligned} \tag{3}$$

[44] Though we specify this constraint in the maximization problem to follow, because we assume convex costs, the boundary conditions are never binding.

B.1.2 proof of positive demand bias in optimal menus. The first order conditions by price shown in equation 2 yield

$$\theta_{HM} = \quad \frac{p_H - p_M}{q_H - q_M} = \frac{c(q_H) - c(q_M)}{q_H - q_M} + \frac{1 - F(\theta_{HM})}{f(\theta_{HM})} \tag{4}$$

We now introduce demand for the high quality option $D_H = [1 - F(\theta_{HM})]$, demand for the middle quality option $D_M = [F(\theta_{HM}) - F(\theta_{ML})]$, and demand for the low quality option $D_L = [F(\theta_{ML}) - F(\theta_L)]$ to combine equation 2 and equation 3 with result 4 completes our set of tools as follows.

$$\begin{aligned} \theta_L = &\quad \frac{D_H c'(q_H) + D_M c'(q_M) + D_L c'(q_L)}{D_H + D_M + D_L} \\ \theta_{ML} = &\quad \frac{D_H c'(q_H) + D_M c'(q_M)}{D_H + D_M} \\ \theta_{HM} = &\quad c'(q_H) \end{aligned} \tag{5}$$

Proof of Theorem 2.

$$\begin{aligned} \frac{\partial D_H}{\partial p_H} = &\quad -f(\theta_{HM})\frac{\partial \theta_{HM}}{\partial p_H} \\ \frac{\partial D_M}{\partial p_H} = &\quad f(\theta_{HM})\frac{\partial \theta_{HM}}{\partial p_H} - f(\theta_{ML})\frac{\partial \theta_{ML}}{\partial p_H} \end{aligned} \tag{6}$$

Using result 5.

$$\begin{aligned} \frac{\partial \theta_{HM}}{\partial p_H} = &\quad c''(q_H)\frac{\partial q_H}{\partial p_H} \\ \frac{\partial \theta_{ML}}{\partial p_H} = &\quad \frac{[\frac{\partial D_H}{\partial p_H}c'(q_H) + \frac{\partial D_M}{\partial p_H}c'(q_M) + D_H c''(q_H)\frac{\partial q_H}{\partial p_H} + D_M c''(q_M)\frac{\partial q_M}{\partial p_H}]}{1 - F(\theta_{ML}) - \theta_{ML} f(\theta_{ML})} \end{aligned} \tag{7}$$

Then

$$\begin{aligned} \frac{\partial D_M}{\partial p_H} = &\quad \frac{1}{1 - F(\theta_{ML}) - \theta_{ML} f(\theta_{ML}) + f(\theta_{ML})}\{f(\theta_{HM})[1 - F(\theta_{ML}) - \theta_{ML} f(\theta_{ML})]c''(q_H)\frac{\partial q_H}{\partial p_H} \\ &- f(\theta_{ML})[1 - F(\theta_{HM}) - \theta_{HM} f(\theta_{HM})]c''(q_H)\frac{\partial q_H}{\partial p_H} - f(\theta_{ML})[F(\theta_{HM}) - F(\theta_{ML})]c''(q_M)\frac{\partial q_M}{\partial p_H}\} \end{aligned} \tag{8}$$

Without loss of generality, we assume a uniformly distributed $F(\theta) \sim U(0,1)$. Given the assumption, equation 8 simplifies to

$$\frac{\partial D_M}{\partial p_H} = \quad \frac{[\theta_{HM} - \theta_{ML}](c''(q_H)\frac{\partial q_H}{\partial p_H} - c''(q_M)\frac{\partial q_M}{\partial p_H})}{2 - 2\theta_{ML}} \tag{9}$$

From here we know $2 - 2\theta_{ML} > 0$ since θ_{ML} is now restricted to $[0,1)$ and there must be some $\theta_{HM} > \theta_{ML}$ by construction. And $\left(c''(q_H)\frac{\partial q_H}{\partial p_H} - c''(q_M)\frac{\partial q_M}{\partial p_H}\right) > 0$ because $c''(q_H) > c''(q_M)$ since $c(q)$ is convex and $q_H > q_M$ by construction and $\frac{\partial q_H}{\partial p_H} \geq \frac{\partial q_M}{\partial p_H}$ since the converse implies there exists some p_H such that $q_M > q_H$.[45] Therefore $\frac{\partial D_M}{\partial p_H}$ is strictly positive. ■

B.2 Non-Optimal Menus

Non-optimal menu creation generates a positive λ_2 as backers may only select one tier such that each tier is a (forced) substitute for others on the menu.[46]

Theorem 3. *A randomly designed Kickstarter menu has positive cross-price elasticity for each combination of tiers when consumers are rational.*

Proof of Theorem 3. Tiers within menus offer similar products by construction. Consumers may only select one of these tiers. Therefore, all tiers are substitutes. When consumers are rational, the neoclassical model predicts strictly positive cross-price elasticity for substitutes. ■

B.3 Proof of the Main Result

Proof of Theorem 1. Theorem 2 proves that λ_2 is positive when firms design optimal menus and Theorem 3 proves that λ_2 is positive when firms use non-optimal menus. If the marketplace contains a mixture of these types of firms then the composite λ_2 is positive. ■

[45] There is no long-lived convergence in menu options as $p_H \to \infty$ also $q_H \to \infty$ such that $c'(q_H) \to \infty$ thus $\theta_{HM} \to 1$ and the price menu becomes dual-optioned s.t. $\frac{\partial q_M}{\partial p_H} \to 0$.

[46] In this case λ_2 is not a bias.

Appendix C: Auxiliary Details

We match sales by eliminating words associated with sales in the specific order shown to avoid mismatching. Keywords: 48 hours only, 24 hours only, 48 hour, 24 hour, black friday, cyber monday, flash, sale, super early bird special, super early bird discount, super early bird, early bird special, early bird discount, early bird, superearlybird, super earlybird, earlybird, kickstarter special, kickstarter, discount, ks special, ks, seb, eb, -, :, |, \, /, (,), !, “, ', commas and remaining white-space.