MICHIGAN ROSS

Working Paper

A Dynamic Model of Crowdfunding

Saeed Alaei
Google, Inc.

Azarakhsh Malekian
Rotman School of Management
University of Toronto

Mohamed Mostagir
Stephen M. Ross School of Business
University of Michigan

Ross School of Business Working Paper
Working Paper No. 1307
February 2016

A Dynamic Model of Crowdfunding *

Saeed Alaei† Azarakhsh Malekian‡ Mohamed Mostagir§

Abstract

Crowdfunding is quickly emerging as an alternative to traditional methods of funding new products. In a crowdfunding campaign, a seller solicits financial contributions from a crowd, usually in the form of pre-buying an unrealized product, and commits to producing the product if the total amount pledged is above a certain threshold. We provide a model of crowdfunding in which consumers arrive sequentially and make decisions about whether to pledge or not. Pledging is not costless, and hence consumers would prefer not to pledge if they think the campaign will not succeed. This can lead to cascades where a campaign fails to raise the required amount even though there are enough consumers who want the product. The paper introduces a novel stochastic process — *anticipating random walks*— to analyze this problem. The analysis helps explain why some campaigns fail and some do not, and provides guidelines about how sellers should design their campaigns in order to maximize their chances of success.

1 Introduction

Crowdfunding is a new approach that allows businesses and entrepreneurs to decentralize the funding process by directly appealing to the end consumer as a funding source. In a crowdfunding environment, a seller solicits financial contributions from the crowd, usually in the form of pre-buying a still-unrealized product, and commits to producing the product if the total amount pledged is above a certain threshold. The threshold is usually chosen to cover the amount of funds required to start production, and the product is not funded if the campaign is unable to reach that threshold within an allotted amount of time. Crowdfunding platforms like Kickstarter and Indiegogo regulate these campaigns by a) vetting the seller and the proposed project, and b) ensuring that anyone who makes a pledge to buy actually follows through with their purchase if the campaign is successful.

One of the many successful examples of crowdfunding is the Pebble Watch, which at the time of its release in 2012 was a precursor to the current smart wearables market. Pebble, an unknown startup at the time, attracted enough attention through the Y-Combinator incubator program, but

*We thank Ozan Candogan, Robert Kleinberg, Alireza Tahbaz-Salehi, and seminar participants at Chicago Booth, NYU Stern, and Columbia Business School.

†Google Research. alaei@google.com

‡ Rotman School of Management, University of Toronto. Azarakhsh.Malekian@rotman.utoronto.ca

§ Ross School of Business, University of Michigan. mosta@umich.edu

failed to convert that into traditional funding from investors. As a response, the company launched a crowdfunding campaign with the goal of raising \$100,000 to start production. It ended up raising more than \$10,000,000 in under two months from consumers who liked the idea enough to support it. Pebble came back to crowdfunding in the first half of 2015 to raise funds for its Pebble Time Smartwatch, this time raising \$20,000,000 – substantially more than its target goal of \$500,000. This is just one of the many examples that illustrate how this decentralized funding paradigm can rectify some of the common inefficiencies found in the traditional supply and demand process, especially for niche products or new products that carry with them a high degree of uncertainty.[1] A seller does not commit to production until demand is observed and necessary capital is raised, and consumers who are interested in products that the market has deemed too risky or too specialized to produce get a chance to help bring these products to market.

This paper presents a dynamic model of crowdfunding where consumers arrive sequentially and make decisions on whether to back a product or not. Consumers have their own valuations for the product and make their pledging decision after observing the following: a) the price of the product, b) the aggregate contribution made by previous backers, c) the campaign funding target, and d) the duration of the campaign. Pledging is not costless: after making a pledge and until the campaign is over, a backer may have to pass on opportunities to use their money in anticipation of using it to pay for their pledge if the campaign is successful. However, if the campaign fails, then not only does the backer not get the product, but he has also forfeited the option to use the funds in the interim. Because of this, consumers need to estimate the chances that a campaign will succeed before they decide to make a pledge. This can lead to cascades where the absence of earlier pledges makes those who arrive later pessimistic about the chances of campaign success, and therefore discourages them from pledging, leading to a vicious cycle. Conversely, a pattern may emerge where consumers estimate the success probability to be high and create a virtuous cycle through pledging. The goal of this paper is to understand how and when do these different dynamics arise and how the campaign designer can influence them through her choices (of product price, number of required pledges for success, and campaign duration) in order to maximize her chances of success.

Contribution This paper contributes on three dimensions. First, on the modeling side, it provides a simple model that allows us to explain a sharp pattern observed in the data on crowdfunding campaigns, namely, the bimodal distribution of campaign outcomes. Campaigns either succeed, or fail to generate any substantial fraction of their funding goal, i.e. it is quite unlikely that a campaign would fail by raising, for example, 85% of its target funding. The possible outcomes are overwhelmingly tilted towards either success or dismal failure. Figure 1 shows that out of all projects on Kickstarter, 49% failed to generate more than 20% of their goal while 40% were successfully funded. As we show, this distinct difference can be explained by the cascade effects we described earlier.

The second contribution is methodological. We analyze the behavior of consumers through intro-

[1]Most of the literature that attempts to mitigate these inefficiencies usually rely on contracts between suppliers, manufacturers, and retailers (see for example, Cachon (2004); Cachon and Lariviere (2005), and the book by Cachon and Terwiesch (2009)) and consider the end consumer as a source of uncertainty rather than a potential asset.

0%	$1-20\%$	$21-40\%$	$41-60\%$	$61-80\%$	$81-99\%$	$\geq 100\%$
20,871	68,546	12,153	4,629	1,577	798	73,530

Figure 1: Number of campaigns and the percentage of funding goal they received on Kickstarter (as of 11/12/2014)

ducing a novel random walk that we call an *anticipating random walk*, which has the property that the probability of transition at every step of the walk depends on future transitions and on the probability of the walk eventually reaching a particular *success* state after a certain number of time steps. This random walk is complicated to analyze because it anticipates the success probability at every time step, and this anticipation takes into consideration future anticipations as well, leading to a recursive structure. We circumvent this difficulty by introducing a related dynamic process whose analysis leads to a solution to the anticipating random walk and the associated crowdfunding problem. We believe that this random walk and its analysis can find wider applications in other contexts and dynamic problems where the decisions made by current actors are not only based on history but also on how they will affect the decisions of future actors.[2]

The third contribution is a set of concrete recommendations for the design of crowdfunding campaigns. On the technical side, the main theorem in the paper provides a quick tool to design campaigns that have very high success probabilities. On the managerial side, the insight behind these recommendations stems from the recognition that uncertainty about the outcome –particularly, uncertainty about other consumers– is one of the salient features that drives pledging behavior. Essentially, the simplest question that a consumer asks himself when making a pledging decision is whether there are enough people who value the product enough to send it to production, and the success of the campaign depends on the choices that the campaign designer makes to alleviate this uncertainty and help the consumer answer this question. To this end, we find that uncertainty increases as the number of pledges required to make a campaign succeed increases. Because of this, and in order to make pledging more attractive, the price should be discounted to compensate consumers for the uncertainty they face. Taken together, this implies that the seller's optimal design should try to raise the target goal using as few backers as possible. This reduces uncertainty and allows the seller to keep the price high, maximizing the total expected pledge.

The analysis in the paper assumes that the seller constrains the campaign to be successful only if it raises a pre-determined, exogenous goal. This goal can be thought of as the minimum target fund the seller has to raise in order to cover the various costs associated with production, transportation, etc., and hence it is sensible for the seller to not move forward with production if that minimum goal is not met. Relaxing this constraint leads to some findings that on first look might appear a bit pardoxical. In particular, it can sometimes be beneficial for the seller to set an artificial goal that is lower than the

[2]A setting with a similar flavor is how donations are made to candidates running for office (for example, Mutz (1995)). Donors anticipating that a particular candidate will not generate a lot of donation money can be less inclined to put in their own money. Similarly, Andreoni (2006) discusses how a large donor might back a candidate in anticipation of how this backing can signal future potential donors to contribute as well.

actual funding target required for production and as a result, ends up raising more money than if she had set the target equal to the actual goal. The reason, as alluded to in the previous paragraph, is that a lower target reduces consumers' uncertainty about the probability of campaign success (since for example, all being equal, it is easier to raise $10K than $20K) and hence they are more willing to pledge and create a good cascade. One example is the Coolest Cooler campaign on Kickstarter, which debuted in late 2013 with a goal of $125,000 and failed to reach that target. The campaign relaunched again a few months later, with a virtually identical product and campaign specification, but this time with a funding goal of only $50,000. It raised more than $13,000,000, making it the largest funded project in Kickstarter history at the time. Another example is the Spark Core, which is a wi-fi enabled board that consumers can use to connect various appliances to the Internet. The initial campaign asked for $250,000 and was unable to reach that target. The campaign relaunched a short time after, dropping the goal to $10,000 and successfully raising more than $500,000.

Of course, the downside to the strategy of shading the real target is that it is possible that the campaign ends up raising enough money to cover the artificial target –and hence 'succeed'– but not the actual one. This leaves the seller with a commitment to deliver a product that it does not have enough means of producing. As we discuss in Section 4, this predicament is often coupled with costly outcomes like litigation and loss of reputation. Because of this, it is natural to ask if there are better ways in which the seller can manage uncertainty. Section 5 introduces dynamic and menu pricing, with the idea that pledge amounts can adjust during the course of the campaign to account for the different uncertainties faced by consumers.

Related Literature The potential for using crowds to improve the decisions and operations of firms has attracted a lot of recent attention. Araman and Caldentey (2016) and Marinesi and Girotra (2013) explore the idea of using a crowd voting mechanism to gauge demand and interest in a product before committing to production. Voting is costless and does not come with a commitment to purchase, and hence their focus is on how the firm can use this information to update its estimates of demand and adjust its funding target. Huang et al. (2014) use data from Dell's IdeaStorm to describe how a crowd can be used to potentially influence the actual designs and offering of the firm, and Lobel et al. (2015) study how customer referrals can be used as a marketing tool for the firm. On the empirical front, Mollick (2014) provides a detailed analysis of data from $\sim 48,500$ Kickstarter projects. Similar to Figure 1, he finds that projects either succeed, or they fail by large margins, with failed campaigns raising only 8% of their funding target on average. Kuppuswamy and Bayus (2013) use data to try and empirically understand the dynamics of crowdfunding, and find that the propensity of backers to contribute is influenced by how much has already been pledged. Similar effects are found by Wu et al. (2014) in the context of group buying. This is consistent with the empirical and experimental literature on public goods. In a randomized experiment, List and Lucking-Reiley (2002) find that donors give more when they are told that the required funding for a project is near its goal. Similarly, Vesterlund (2003) finds that announcing contributions generates more contributions. This self-reinforcing behavior is also observed in our dynamic model, where backers' decisions are influenced by the current contribution level. Bagnoli and Lipman (1989) and Varian (1994) study if public goods can be curated through collecting private donations in an all-or-nothing mechanism. In contrast to their work, the product in our model is a private good produced through aggregating the actions of several agents, and no free-riding can happen since consumers only get the good if they pay for it. Additionally, prices and the number of people required for success are set by the campaign designer instead of letting donors decide their own contribution levels.

In addition to the work referenced above, there are several empirical studies that deal with various facets of crowdfunding. Mollick and Kuppuswamy (2014) examine the outcomes of crowdfunding in terms of delivery delays, revenues, and the viability of crowdfunding as a successful entrepreneurship platform. Agrawal et al. (2011) discuss how geography affects crowdfunding outcomes despite the notion that one of the main advantages of crowdfunding is breaking down standard geographical barriers. In Agrawal et al. (2013), the same authors discuss how equity-based crowdfunding markets operate when standard market constructs like reputation and transaction costs are upended by this new platform.

On the theory side, Belleflamme et al. (2013) develop a static model to explore when the campaign creator should offer a pre-ordering scheme (like the one explored in this paper) or an equity-based scheme where backers become investors. Anand and Aron (2003) also use a static model to examine when group-buying is preferred to a posted-price mechanism and the timing of pricing and production decisions when there is uncertainty about the size of the market. Jing and Xie (2011) show that social interactions between informed and uninformed consumers can improve the efficiency of group buying. In more recent work, Hu et al. (2015) adopt a product-line design approach to the crowdfunding problem. They develop a two-period model and show that offering slightly-differentiated products and an accompanying menu of prices is optimal. We briefly touch on menu pricing in Section 5, but our main interest is in understanding how to influence consumer behavior in a dynamic setting through the choices of campaign duration, price, and number of pledges required to reach an *exogenous* goal. Hu et al. (2013) compare sequential and simultaneous group-buying mechanisms and show that sequential mechanisms outperform simultaneous mechanisms in a two-period model when the deal succeeds only if both consumers sign up. They extend this to an N period model when success again requires all N consumers to sign up. In contrast, our paper considers the number of periods, the number of consumers required to sign up (for success), as well as the price they sign up at to all be design variables that the seller selects, and lays out how these variables are connected and how they should be used together to maximize the chances of success.

This paper is also related to the social learning literature, e.g. Bikhchandani et al. (1992); Acemoglu et al. (2011), and the survey in Acemoglu and Ozdaglar (2011), where agents arrive sequentially and observe other agents' choices and update their beliefs about an underlying state of the world. The important difference is that the agents in the environment we consider in this paper do not act just to maximize their current utility (for example, by choosing a better retailer or a better restaurant), but also because of the consequence of their choices on the actions of those agents who follow, whose own choices affect the agents' utility through the dependence of the payoff on campaign success. In that sense, the underlying state in crowdfunding can be thought of as the probability of campaign success, and hence is time- and action- variant compared to the fixed underlying state in the social learning literature. Because of this, the underlying dynamics of these two environments are quite different.

The paper is organized as follows. The benchmark model is introduced in Section 2, and Section 3 introduces anticipating random walks and uses them to analyze backer behavior. Section 4 outlines general guidelines for how a seller should design the campaign, and Section 5 examines various extensions to the model. Section 6 concludes the paper and discusses several directions for future research.

2 Model

Our benchmark model has a seller who offers a (still-unrealized) product to consumers through a crowdfunding platform.[3] The seller is interested in raising an exogenous minimum target fund G to cover production costs, and chooses a pledge amount, or price p, associated with the product being sold. The seller also chooses a duration n for the campaign to raise the target G. Consumers arrive sequentially, with one consumer arriving in each time period, which for simplicity we call a day.[4] We refer to the consumer arriving on day $i = 1, ..., n$, as consumer i. Consumers can pledge p to the project upon their arrival or exit without pledging. If by the deadline the total number of pledges is at least equal to $k = \left\lceil \frac{G}{p} \right\rceil$, then the campaign is successful: backers are charged p each, funds are allocated to the seller, and production commences. Conversely, if the campaign falls short of its goal, no one is charged and no money changes hands. We assume that consumer i has Bernoulli valuation $v_i \in \{v^H, v^L\}$ for the product being sold, and that π, the probability of having a high valuation v^H, is known. We further assume that v^L is equal to zero so that consumers with low valuations never pledge, i.e. a consumer is either interested in the product or not. If consumer i decides to pledge, he gets utility equal to $v_i - p = v^H - p$ if the campaign is successfully funded. As discussed in the introduction, consumers who pledge anticipate paying the pledge amount at the conclusion of the campaign and might therefore forfeit the opportunity to use the funds in the interim. If the campaign fails, then backers would have tied up their money and got nothing in return. This scenario is captured by having an opportunity (uncertainty) cost $c > 0$ that a backer incurs in case of campaign failure.

The utility of backer i can thus be written as

$$u_i = \begin{cases} v_i - p & \text{if funded successfully,} \\ -c & \text{if not funded.} \end{cases} \tag{1}$$

A consumer therefore would only pledge if he is reasonably sure that the campaign will succeed, i.e. that there will be at least k pledges made. It is important to stress that the question of whether a

[3]We discuss several extensions to the model in Section 5.

[4]We use the terms 'consumers' and 'backers' interchangeably.

campaign will succeed or not is much more involved than simply asking whether at least k out of n consumers have high valuations. This is because a consumer's problem is *time* and *path*-dependent: it matters not only how many consumers have high valuations, but also the times at which these consumers arrive and what happened up to that time. A campaign that is 30 days long and needs 5 more pledges to reach its funding goal can look very different to consumer number 20 than it looks to consumer number 10 or 15.

Based on this discussion, a consumer's decision to pledge depends on the valuation for the product as well as the price, but also crucially hinges on their belief about whether the campaign will succeed in reaching its goal k or not. This belief changes from one backer to the next depending on the current state of the campaign (i.e. how many pledges have been placed and how much time is left). Let α be the success probability that consumer i estimates, then using (1) to write the expected utility $E[u_i]$, we get

$$E[u_i] = \alpha(v_i - p) - (1 - \alpha_i)c,$$

which means that backer i will pledge only if his valuation v_i satisfies

$$v_i \geq \left(\frac{(1-\alpha)c}{\alpha} + p \right)$$

or equivalently, if the estimated success probability satisfies

$$\alpha \geq \frac{c}{v_i - p + c} = \frac{c}{v^H - p + c} = \beta \tag{2}$$

Equation (2) indicates that a consumer will pledge if they estimate the probability of success to be at least equal to a probability threshold β. For backers to make a decision then, they need to be able to estimate the success probability α. The sequential nature of the problem lends itself to a dynamic programming formulation and an associated stochastic process that is analyzed in Section 3.

The risk-neutral seller is interested in maximizing the discounted total pledge. To this end, the seller selects the duration n, number of pledges k, and price p that solve

$$\begin{aligned} &\max_{n,k,p} \mathbb{E}\left[\delta^n p \sum_{i=1}^{n} I_i(n,k,p) \right] \\ &\text{s.t.} \quad kp \geq G \end{aligned} \tag{3}$$

where $\delta < 1$ is a discount factor for the seller and $I_i(n,k,p)$ is an indicator function that is equal to 1 when consumer i pledges and 0 otherwise. The functions $I_i, i = 1, ..., N$ are correlated, since per our earlier discussion, the pledging decision of a backer is interlinked with past and future decisions. The values that the seller chooses for k and p are such that any successful campaign should raise at least the funding target G. In Section 4 we examine the problem when this constraint is relaxed.

Discussion of Model The benchmark model makes several assumptions. First, it assumes that consumer valuations only take one of two values instead of being drawn from a continuum. We believe that this assumption captures a simple and fundamental question that consumers ask themselves when deciding to make a pledge, namely "are there enough people interested in this product?", which in our model translates to "are there enough people with high valuations?" It is possible that some consumers think about the degree of how much other consumers value the product. This introduces mathematical intricacies that, as we show using simulations in Section 5, do not appear to generate new insights.

We assume that backers incur an uncertainty cost when they make a pledge to a campaign that still has not raised enough money to meet its threshold. As mentioned in the introduction, this cost captures the possibility that a backer may have to pass on opportunities to use their money in anticipation of using that money to pay for their pledge if the campaign is successful. One can imagine though that this cost might depend on when a consumer arrives. A consumer who arrives on the very last day for example faces no uncertainty and does not tie up their money if they make a pledge. It might then be reasonable to have the uncertainty cost depend on arrival time, perhaps through some inverse discounting that makes the price of the product appear higher if a pledge is made earlier in the campaign. We again examine this scenario using simulation and find no qualitative difference in the outcomes we obtain under the benchmark model.

Similar to the literature on sequential decision making cited earlier, we assume that at each time step a new consumer arrives and makes a decision. In particular, consumers either pledge or exit, and do not have the option to wait before making a pledge. Kickstarter's recent addition of a "Remind Me" button to campaign web pages – a button that users click so that they can be reminded to check the campaign again in the future– seems to suggest that backers checking a product's campaign and not returning is not uncommon. However, it is possible that there are consumers who are interested in the campaign but monitor it without pledging, and only pledge when they think the campaign is likely to succeed. The main technical difficulty with this scenario is that the (hidden) state variable associated with the number of consumers who are waiting makes the backers' decision problem intractable as the state space becomes very large. We discuss more possibilities to extend the work in this paper in Section 6.

3 Backer Behavior and Anticipating Random Walks

This section analyzes how backers behave as a function of the parameters of the problem. As discussed in the previous section, backers estimate success probabilities in order to decide whether to pledge or not. The following dynamic program computes the probability of success at each time step. Let s_i^j denote the probability of success conditioned on having j pledges by end of time i, i.e. s_i^j is the success probability that consumer i estimates given that j pledges have been made up to and including his own pledge. Recall that π is the probability of having a high valuation and that a campaign succeeds if at least k pledges are made by time n. The exact value of s_i^j is computed recursively as follows:

$$s_i^j = \begin{cases} 0 & i = n,\, j < k \\ 1 & i = n,\, j \geq k \\ s_{i+1}^j & i < n,\, s_{i+1}^{j+1} < \beta \\ \pi_{i+1} s_{i+1}^{j+1} + (1 - \pi_{i+1}) s_{i+1}^j & j < n,\, s_{i+1}^{j+1} \geq \beta \end{cases} . \tag{4}$$

We are interested in computing the campaign success probability s_0^0. To this end, the following section introduces *anticipating random walks* and shows how they can be used to provide a quick and verifiable characterization of successful and failed crowdfunding campaigns.

3.1 Anticipating Random Walks

An anticipating random walk is a stochastic process that is closely connected to the problem of computing success probabilities in crowdfunding and whose analysis characterizes the outcomes of crowdfunding campaigns, but is also more generally applicable to settings where sequential decisions stochastically depend on future outcomes. Formally

Definition 1. Given a probability $\beta \in (0, 1)$, a target $k \in \mathbb{N}$, and a sequence of independent Bernoulli random variables $X_1, \ldots, X_n$ with known means $\pi_1, \ldots, \pi_n \in [0, 1]$, an *anticipating random walk* with the aforementioned parameters is given by a sequence of positions on the real line denoted by $Y_0, Y_1, \ldots, Y_n$ defined as follows:

$$Y_i \mid Y_{i-1}, X_i = \begin{cases} Y_{i-1} + X_i & \text{if } \mathbf{Pr}[Y_n \geq k \mid Y_i = Y_{i-1} + X_i] \geq \beta \\ Y_{i-1} & \text{if } \mathbf{Pr}[Y_n \geq k \mid Y_i = Y_{i-1} + X_i] < \beta \end{cases}$$

$$Y_0 = 0$$

A realization of an anticipating random walk is called *successful* if it advances at least a total of k units by the end of time n, that is $Y_n \geq k$.

Informally, an anticipating random walk advances one unit at time i if $X_i = 1$ and if the probability of eventual success conditioned on current position at time i is at least β. We are interested in computing the probability of success at time 0.

Connection to crowdfunding Consider a scenario with n backers where the parameters of the problem are such that a consumer with high value pledges only if he believes that the campaign will be funded with probability at least β (c.f. Equation (2)). This can be interpreted as an anticipating random walk which each time step corresponding to one backer, with π_t being the probability of backer at time step t having a high value, and k being the minimum number of pledges required for successful funding. The success probability of the anticipating random walk is exactly equal to the success probability of the corresponding crowd funding campaign.

This section has three main results. Theorem 1 is a general result that implies the existence of a phase transition in probability of success as a function of $\sum_i \pi_i, k$, and β, i.e. a bimodal outcome for anticipating random walks where the probability of success is either bounded below by β or is equal to zero. In order to prove this theorem, we prove a closely-related result in Theorem 2 and use that to show that the characterization in Theorem 1 holds.

A slight perturbation of the result in Theorem 1 leads to Theorem 3 which shows that an anticipating random walk either succeds with probability $1-\epsilon$ or with probability zero. Besides their theoretical interest, these results map to crowdfunding campaigns to show that they are either funded with very high probability or with probability zero, depending on the problem's parameters. The conditions of the theorem provide a quick test for whether a campaign is likely to be successful or not, and therefore gives general guidelines into how a campaign designer should set these parameters.

The first result on anticipating random walks delineates the conditions on the two main outcomes for success probability.

Theorem 1. *An anticipating random walk as given in Definition 1 has a phase transition around $\bar{k} \in \mathbb{N}$ such that the following two properties:*

- *if $k \leq \bar{k}$, then $\mathbf{Pr}[Y_n \geq k] \geq \beta$,*
- *if $k > \bar{k}$, then $\mathbf{Pr}[Y_n \geq k] = 0$.*

Furthermore,

$$\frac{\sum_i \pi_i}{\ln(1/(1-\beta)) + 1} \leq \bar{k} \leq \frac{\sum_i \pi_i}{\ln(1/(1-\beta))}$$

.

Theorem 1 states that there is a threshold $\bar{k}$ such that if the number of pledges required for success is above that threshold then the campaign will not get funded. On the other hand, the campaign will be funded with probability at least β if the number of required pledges is less than the threshold. The theorem also gives lower and upper bounds on the threshold value. The next section is devoted to the proof of the previous theorem.

3.2 A "Moving Sand" Process

This section analyzes a process that is closely related to the dynamic program in (4) and that enables us to obtain bounds on s_i^j and in particular, on s_0^0, the probability of success of the anticipating random walk and the corresponding crowdfunding campaign. The process and the intuition are explained here before launching into the proof. Consider an unlimited number of initially empty buckets with unit capacity each. Each time step an amount of sand –a fraction between 0 and 1– is added to the leftmost bucket (bucket 1). If after a few time steps the amount of sand in that bucket is more than a parameter β, then a fraction of sand in that bucket is moved to the next bucket (bucket number 2), and the process is recursively repeated for all buckets: once the amount of sand in a bucket is above the parameter β, a fraction of sand in that bucket is moved to the next bucket. Thus a bucket will not start filling with sand unless the bucket preceding it has at least β of sand in it. The dynamics of the process are designed such that the amount of sand in bucket k at time n is equal to the success probability of the equivalent anticipating random walk and crowdfunding campaign, i.e. the probability that an anticipating random walk makes k advances in n time steps or that k pledges are made during n consumer arrivals. We then argue that depending on which side k is on compared to the threshold $\bar{k}$, the amount of sand in the k^{th} bucket is either at least equal to β or is equal to zero. These correspond to the two cases in Theorem 1. The following is a formalization of these arguments.

Definition 2 (A Moving Sand Process)**.** Given a sequence $q_1, \ldots, q_n \in (0, 1]$ and an infinite sequence of initially empty buckets indexed by integer numbers, the following takes place at every time step $t \in \{1, \ldots, n\}$:

- Simultaneously for all buckets $j \in \mathbb{N}$: if the amount of sand in bucket j is at least β, a fraction q_t of the sand in that bucket is moved to bucket $j + 1$, otherwise a fraction q_t is discarded.
- q_t units of sand is added to bucket 1.

Formally the amount of sand in bucket $j \in \mathbb{N} \cup \{0\}$ at the end of time step $t \in \{0, \ldots, n\}$ is given by b_j^t defined as follows:

$$b_j^t = \begin{cases} q_t b_{j-1}^{t-1} + (1 - q_t) b_j^{t-1} & b_{j-1}^{t-1} \geq \beta \\ 0 & b_{j-1}^{t-1} < \beta \\ 0 & t = 0, j \geq 1 \\ 1 & j = 0 \end{cases} \tag{5}$$

Connection to anticipating random walk An anticipating random walk with parameters β, k, and $\pi_1, \ldots, \pi_n$ is related to a moving sand process with parameters $q_1, \ldots, q_n$ where $q_t = \pi_{n-t+1}$. In particular, s_i^j as defined in (4) is given by

$$s_i^j = b_{k-j}^{n-i}. \tag{6}$$

The following theorem summarizes the main result of this section. Once this result is proved, Theorem 1 follows through the connection in Equation (6).

Theorem 2. *A moving sand process as defined in Definition 2 satisfies the following two properties for any $k \in \mathbb{N}$:*

- *if $\sum_t q_t \geq k\left(\ln(1/(1-\beta)) + 1\right)$, then $b_k^n \geq \beta$*
- *if $\sum_t q_t < k \ln(1/(1-\beta))$, then $b_k^n = 0$.*

The rest of this section is devoted to proving the previous theorem. The idea is the following: we call the non-empty bucket with the highest index the 'Charging Bucket', and then we track the amount of sand in all buckets over time in order to show that at time n, if the condition in the theorem holds then the index of the charging bucket must at least be equal to $k+1$. But if that is the case, then by the definition of the moving sand process, bucket k has to have at least β amount of sand in it, which, from the interpretation of the amount of sand in a bucket provided before Definition 2, is exactly what we want to show.

Definition 3 (Charging Bucket)**.** For each $t \in \{1, \ldots, n\}$, the *charging bucket* at time t, whose index is denoted by λ^t, is the bucket with largest index that contains non-zero amount of sand at the end of time t. Formally

$$\lambda^t = \max\left\{j \mid b_j^t > 0, j \in \mathbb{N}\right\} \qquad \forall t \in \{1, \ldots, n\}. \tag{7}$$

We also define B^t to denote the total amount of sand in all buckets (except bucket 0) at time t, i.e.

$$B^t = \sum_{j \in \mathbb{N}} b_j^t \qquad \forall t \in \{0, \ldots, n\}. \tag{8}$$

Definition 4 (Bucket Charging Time)**.** For each $j \in \mathbb{N}$, the charging time of bucket j, denoted by τ_j, is the first time at which bucket j becomes the charging bucket:

$$\tau_j = \min\left(\left\{t \mid \lambda^t = j\right\}, \infty\right) = \min\left(\left\{t \mid b_j^t > 0\right\}, \infty\right) \tag{9}$$

Observe that bucket j is the charging bucket only within the times $t \in \{\tau_j, \ldots, \tau_{j+1} - 1\}$, assuming $\tau_{j+1} < \infty$.

The following lemma simply states that the amount of sand in a bucket increases over time and that buckets are filled left-to-right with any bucket having at most an amount of sand that is equal to the amount of sand in the bucket to its left, i.e. a bucket with a lower index. The proof of the lemma is deferred to the appendix.

Lemma 1. *The value of b_j^t is (weakly) decreasing in j and (weakly) increasing in t. Furthermore, $b_{j'}^{t'} \geq \beta$ for all $j' < j$ and all $t' \geq \tau_{j+1} - 1$, and $b_{j'}^{t'} < \beta$ for all $j' \geq j$ and all $t' < \tau_{j+1} - 1$.*

We next bound the amount of sand in a bucket at any time t after that bucket starts charging.

Lemma 2. *The amount of sand in bucket $j \in \mathbb{N}$ at time $t \geq \tau_j$ is at most*

$$b_j^t \leq 1 - \prod_{t'=\tau_j}^{t} (1 - q_{t'}). \tag{10}$$

The proof is by induction on t. The base case is $t = \tau_j$ for which the statement holds trivially because

$$\begin{aligned} b_j^{\tau_j} &= q_{\tau_j} b_{j-1}^{\tau_j - 1} + (1 - q_{\tau_j}) b_j^{\tau_j - 1} && \text{by (5)} \\ &\leq q_{\tau_j} && \text{by } b_{j-1}^{\tau_j-1} \leq 1 \text{ and } b_j^{\tau_j-1} = 0. \end{aligned}$$

To prove the induction step for $t > \tau_j$ observe that

$$\begin{aligned}
b_j^t &= q_t b_{j-1}^{t-1} + (1-q_t) b_j^{t-1} && \text{by (5)} \\
&\le q_t + (1-q_t)\left(1 - \prod_{t'=\tau_j}^{t-1} (1-q_{t'})\right) && \text{by } b_{j-1}^{\tau_j - 1} \le 1 \text{ and induction hypothesis} \\
&= 1 - \prod_{t'=\tau_j}^{t} (1-q_{t'}).
\end{aligned}$$

■

The next lemma simply finds an expression for the total increase in sand between two consecutive time steps.

Lemma 3. *The increase in the total amount of sand in all buckets at time step $t \in \{1, \ldots, n\}$ is*

$$B^t - B^{t-1} = q_t \cdot (1 - b_{\lambda^t}^{t-1}) \tag{11}$$

Except for the first bucket and the charging bucket (the one at index λ^t), sand may only move from one bucket to another. This means the total change is equal to the amount of sand that enters the first bucket which is q_t minus the amount of sand that is discarded from the charging bucket which is $q_t \cdot b_{\lambda^t}^{t-1}$. ■

With the previous lemmas in tow, we prove the following result, which is the last lemma required to prove Theorem 2. Roughly, this lemma gives a lower bound on the amount of sand in all buckets during the time when bucket j is the charging bucket.

Lemma 4. *The increase in the total amount of sand in all buckets during the times that bucket $j \in \mathbb{N}$ is charging is at least*

$$B^{t^\dagger} - B^{\tau_j - 1} \ge \beta + (1-\beta)\left(\sum_{t=\tau_j}^{t^\dagger} q_t - \ln(1/(1-\beta))\right)^+ \tag{12}$$

where $t^\dagger \le \tau_{j+1} - 1$. The shorthand notation $(x)^+$ means $\max(x, 0)$.

First observe that if $\sum_{t=\tau_j}^{t^\dagger} q_t \le \ln(1/(1-\beta))$, the right hand side of eq. (12) is β, so the inequality holds trivially because

$$\begin{aligned}
B^{t^\dagger} - B^{\tau_j - 1} &\ge \sum_{j'} b_{j'}^{t^\dagger} - b_{j'}^{\tau_j - 1} \\
&\ge b_j^{t^\dagger} - b_j^{\tau_j - 1} && \text{by monotonicity of } b_{j'}^{t'} \text{ in } t' \text{ (see Lemma 1)} \\
&\ge \beta && \text{because } b_j^{\tau_j - 1} = 0 \text{ and } b_j^{t^\dagger} \ge \beta \text{ by Lemma 1}
\end{aligned}$$

For the rest of the proof we assume without loss of generality there exists a $t^* \in \{\tau_j, \ldots, t^\dagger\}$ and

$\Delta \in [0, q_{t^*}]$ such that $\sum_{t=\tau_j}^{t^*-1} q_t + \Delta = \ln(1/(1-\beta))$. Observe that by definition of τ_j for all times $t \in \{\tau_j, \ldots, t^\dagger\}$, the charging bucket is j, that is $\lambda^t = j$. Therefore

$$
\begin{aligned}
B^{t^\dagger} - B^{\tau_j - 1} &= \sum_{t=\tau_j}^{t^\dagger} B^t - B^{t-1} \\
&= \sum_{t=\tau_j}^{t^\dagger} q_t \cdot (1 - b_j^{t-1}) && \text{by Lemma 3 and } \lambda^t = j \\
&= \sum_{t=\tau_j}^{t^*-1} q_t \cdot (1 - b_j^{t-1}) + \Delta \cdot (1 - b_j^{t^*-1}) \\
&\quad + \overbrace{(q_{t^*} - \Delta)(1 - b_j^{t^*-1}) + \sum_{t=t^*+1}^{t^\dagger} q_t \cdot (1 - b_j^{t-1})}^{\Psi} \\
&\geq \sum_{t=\tau_j}^{t^*-1} q_t \cdot \prod_{t'=\tau_j}^{t-1} (1 - q_{t'}) + \Delta \cdot \prod_{t'=\tau_j}^{t^*-1} (1 - q_{t'}) + \Psi && \text{by Lemma 2} \\
&= 1 - (1 - \Delta) \prod_{t'=\tau_j}^{t^*-1} (1 - q_{t'}) + \Psi && \text{by refactoring} \\
&\geq 1 - e^{-\Delta - \sum_{t'=t}^{t^*-1} q_{t'}} + \Psi && \text{by } e^{-x} \geq 1 - x \\
&= \beta + \Psi && \text{by definition of } t^* \text{ and } \Delta \\
&= \beta + (q_{t^*} - \Delta)(1 - b_j^{t^*-1}) + \sum_{t=t^*+1}^{t^\dagger} q_t \cdot (1 - b_j^{t-1}) && \text{by expanding } \Psi \\
&\geq \beta + (q_{t^*} - \Delta)(1 - \beta) + \sum_{t=t^*+1}^{t^\dagger} q_t \cdot (1 - \beta) && \text{by Lemma 1} \\
&= \beta + (1 - \beta) \left(\sum_{t=\tau_j}^{t^\dagger} q_t - \ln(1/(1-\beta)) \right) && \text{by definition of } t^* \text{ and } \Delta
\end{aligned}
$$

■

We are now ready to prove Theorem 2. The idea of the proof is to show that if the condition in the theorem holds by the end of time n, then the total amount of sand in all buckets must be greater than k, but since each bucket has unit capacity, it must be that there are at least $k+1$ non-empty buckets. But by Definition 2 and Equation (5), bucket $k+1$ will only start filling up if bucket k has more than β amount of sand in it, and so the result follows.

Proof of Theorem 2 We start by proving the first statement. We show that the index of the charging bucket at time n must be at least $k+1$ which then implies $b_k^n \geq \beta$. The proof is by contradiction.

Suppose the index of the charging bucket at time n is $j^* \leq k$, then

$$\begin{aligned}
B^n &= \sum_{j=1}^{j^*} B^{\min(\tau_{j+1}-1,n)} - B^{\tau_j - 1} \\
&= \sum_{j=1}^{j^*} \beta + (1-\beta)\left(\sum_{t=\tau_j}^{\min(\tau_{j+1}-1,n)} q_t - \ln(1/(1-\beta))\right)^+ && \text{by Lemma 4} \\
&\geq j^* \cdot \beta + (1-\beta)\left(\sum_{t=1}^{n} q_t - j^* \ln(1/(1-\beta))\right)^+ \\
&> j^* \cdot \beta + (1-\beta)\left(k \cdot \ln(1/(1-\beta)) + k - j^* \cdot \ln(1/(1-\beta))\right) && \text{by theorem's hypothesis} \\
&\geq k.
\end{aligned}$$

Therefore the total amount of sand in all buckets at the end of time n must be at least k. However, given that each bucket may contain at most one unit of sand, there must be non-zero amount of sand in bucket $k+1$ which contradicts our assumption of bucket j^* being the charging bucket.

The proof of the second statement is also by contradiction. Consider an instance with $q_t = \frac{k \cdot \ln(1/(1-\beta))}{n}$, and $n \to \infty$. It is easy to show that the index of the charging bucket is strictly less than k by the end of time n. ■

The proof of Theorem 1 follows immediately through the equivalence in Equation (6).

Proof of Theorem 1 The claim follows from Theorem 2 and noticing that $s_0^0 = b_k^n$. ■

The next and last result of this section shows that a slight perturbation of the first case in Theorem 1 leads to the campaign getting funded with probability close to 1. Taken together, this concludes the results of this section by showing that campaign outcomes are largely swayed towards one of two extremes depending on which condition in the statement of the theorem holds.

Theorem 3. *For any ϵ, there exists δ such that if $\sum_i \pi_i \geq k(\ln \frac{1}{1-\beta} + 1)(1+\delta)$, the campaign is funded with probability $1-\epsilon$.*

The proof of the theorem follows from concentration arguments and is deferred to the appendix.

4 Seller's Problem

Having obtained bounds on the probability of campaign success as a function of the parameters of the problem, we now turn our attention to how the seller should set these parameters. The results of this section are summarized in the following points:

- Uncertainty about outcomes drives consumers' behavior. The more people that the outcome depends on, i.e. the more people required to pledge in order for the campaign to succeed, the more uncertain the outcome is.

- Price serves a dual purpose. In addition to making the campaign reach its goal through accumulating pledges, it is set in such a way to offset the uncertainty faced by the consumers, i.e. consumers get an 'uncertainty discount' to make pledging more appealing.
- Taking the two above points together, the seller's optimal strategy is to generate as much funding from as few consumers as possible. This ensures that uncertainty is kept at a low level (because not many people are required to pledge) and price is kept at a high level since the discount handed out to consumers is minimal.

The next example shows that Theorem 1 does not just describe the possible outcomes as a function of campaign parameters, but it is also a powerful tool that can be used to quickly design a campaign that succeeds with very high probability, or to determine the feasibility of a campaign , i.e. whether it is possible at all to raise that much funding in that much time.

Example 1. Consider two campaigns of duration 30 each. The campaigns are identical in that consumers have a high valuation equal to 10, probability of a high valuation consumer is $\pi = 0.5$, and uncertainty cost $c = 3$. The only difference is that the first campaign has a goal of $\$50$ while the second has a goal of $\$100$. We compute $\beta = \frac{c}{v-p+c} = \frac{3}{13-p}$. For the first campaign, assume the seller sets the price to $p = \$7$. This means that to satisfy the constraint $kp \geq 50$, k has to be at least equal to 8. Plugging this to the lower bound on $\bar{k}$ in Theorem 1, we get $\frac{30(0.5)}{\ln(1/(1-\frac{3}{13-7}))+1} = 8.85$. Thus, since $k = 8 < 8.85$, by Theorem 3 this campaign is funded with probability close to 1. Indeed, explicitly solving the dynamic program with $p = \$7$ and $k = 8$ leads to success probability equal to 0.996. What if we increase the price to $\$8$ and decrease k to 7 (so that the constraint still holds, as well as the bound in Theorem 1)? The funding probability goes up to 0.9990. Further increasing p to $\$9$ and decreasing k to 6 results in probability 0.9997. This interaction between k and p is explored in Proposition 1.

Consider now the second campaign that wants to raise $\$100$. If the price is again set to $\$7$, then k has to be set to at least 15 in order to satisfy $kp \geq 100$. The lower bound on $\bar{k}$ in Theorem 1 does not change, but now the lowest value possible for k is higher and since $15 > 8.85$, the condition in the theorem does not hold. Indeed, this campaign has success probability zero, and changing the p and k variables gives no combination that simultaneously satisfies the condition in the theorem as well as the constraint $kp \geq 100$, indicating that it is not possible to design a successful campaign of duration 30 with these parameters. Consequently, the designer has to increase the third variable, the campaign duration, in order for the first condition in the theorem to apply again and for the campaign to succeed.

The analysis of the seller's problem proceeds by examining the case when n is fixed and the seller has to decide on the number of pledges k and price p. While there are no closed-form solutions for any of these quantities, the associated dynamic programs allow us to glean important structural properties that are discussed in the following results. The seller's total expected pledge for a given set of variables can be obtained by taking the output of the dynamic program in Equation (4) – the success probabilities at each time step– and using these probabilities to solve another dynamic program that computes the total expected pledge. The latter dynamic program is presented in Appendix B.

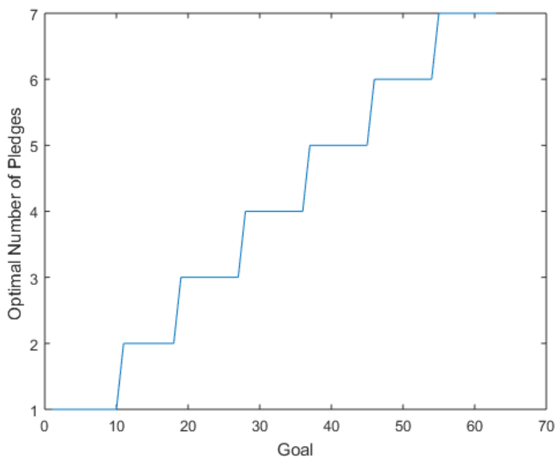

(a) k^* vs. Goal

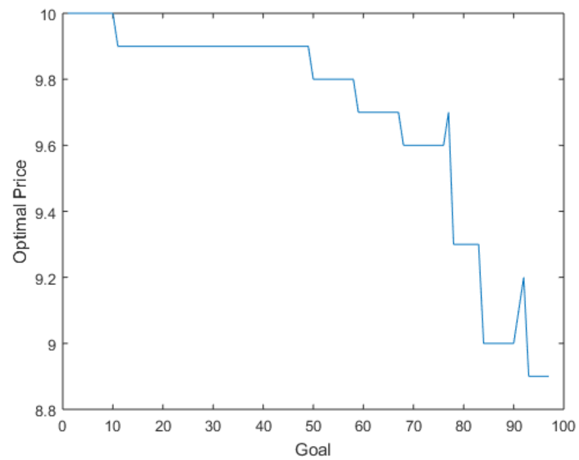

(b) p^* vs. Goal G

Figure 2: Relationship between optimal price, optimal number of Pledges, and Goal for $v^H = 10, \pi = 0.5$, and $c = 3$

While the effect of k on the expected pledge is clear from Theorem 1, the effect of the price p is ambiguous. An increase in k leads to more uncertainty, since the success of the campaign now depends on more people pledging, and hence leads to a higher chance of campaign failure and lower expected pledge. The role of the price though depends on k: if there is little or no uncertainty, then higher prices approaching consumers' valuations increase the total expected pledge, whereas if k is such that there is substantial uncertainty then decreasing the price lowers the probability threshold that a consumer uses to decide whether or not to pledge (c.f. Equation (2)), i.e. it makes pledging more appealing, which in turn increases the probability of success and the number of pledges. This interplay between k and p and their relationship to the goal G is discussed in the next result, which describes the structure of the solution when n is fixed. This result is also illustrated in Figure 2.

Proposition 1. *Fix n and let G_1 and $G_2 > G_1$ be two goal values and let (k_1^*, p_1^*) and (k_2^*, p_2^*) be the optimal solutions corresponding to G_1 and G_2, respectively, then*

- $k_1^* \leq k_2^*$
- *If $k_1^* = k_2^*$, then $p_1^* \leq p_2^*$*
- *If $k_1^* < k_2^*$, then $p_1^* \geq p_2^*$*

Proposition 1 suggests that a higher goal amount can only lead to an increase in the optimal k, as can be seen in Figure 2a, but the price may go up or down depending on whether the optimal k is the same or different for the goal values being considered.

Figure 3a magnifies and superimposes a section of the two graphs in Figure 2. One can see the different cases in Proposition 1: when the number of required pledges increases, prices always drop (for example when k goes from 5 to 6, 7 to 8, and 8 to 9) or stay constant (when k goes from 6 to 7).

Sometimes the price increases for two goal values that have the same k, as they do for goal values in the range from 70 to 75 where the price increases near the end of the interval. The intuition behind why the price decreases when the number of required pledges increase is the following. Let v^H be the high valuation and assume that a consumer pledges if he is indifferent, i.e. if his utility is zero. Now consider the trivial case that G is so low that it lies in $(0, v^H]$, so that a single pledge is all that is required to meet the goal, then the optimal solution is to set $k = 1$ and $p = v^H$. This is because there is no uncertainty in this scenario – a consumer with a high valuation knows that he will get the product with probability 1 if he pledges, and hence can be charged the full amount or as close as possible to that valuation. Consider now a goal $G > v^H$ that cannot be reached through a single pledge. The only way to reach that goal is to increase the number of pledges in order to make $kp \geq G$. But increasing k introduces uncertainty about the outcome, since success now depends on the combined action of consumers instead of a single individual decision, and so the price has to be adjusted downwards to balance that effect and still make pledging an attractive option to consumers. For example, if the goal is in $(v^H, 2v^H]$, then setting the price to v^H as before will lead to zero pledges, since any value of $\pi < 1$ implies that $\alpha < 1$, and the expected return for the first high-value consumer at that price is equal to $\alpha(v^H - v^H) - (1 - \alpha)c < 0$, and thus this consumer never pledges and consequently no other high-value consumers who follow will pledge either. Essentially, the optimum price can be thought of as having two components. One component is the valuation of the consumer, and the other is a discount that consumers get to compensate them for the uncertainty that they face. Since this uncertainty increases as the number of required pledges increases, the discount becomes higher (i.e. the price decreases) as the goal increases.

To see that the price is nondecreasing when the optimal k for two different goals is the same, consider two goals G_1 and $G_2 > G_1$. Any solution (k, p) that is feasible for G_2 is also feasible for G_1. This implies that if the solutions corresponding to the two goals are different, then it must be that the expected pledge under the goal G_2 is less than the expected pledge under the solution for G_1, because if not, then the solution for G_1 cannot be optimal since we can replace it with the solution for G_2 and obtain a higher expected pledge. Thus, if the optimal k is the same for both goal values and the prices are different, it must be that the price p_2 is greater than p_1, since if not, then p_1 is feasible for G_2 and can be used for that goal value to obtain a higher expected pledge. Therefore $p_2 \geq p_1$. This can be observed in the occasional price spikes in Figure 2b and Figure 3a. This happens because the difference between the two goals is not enough to warrant adding an extra required pledge and increase k –with the extra uncertainty and reduction in expected pledge that this entails– and can instead be managed by a (usually slight) price increase at the expense of a smaller reduction in the expected pledge. Notice the gradual decrease in expected pledge in Figure 3b around the $G = 74$ to 75 mark, as opposed to the sharper drops for the other transitions in the graph. This corresponds to the price increase in the $G = 74$ to 75 range.

It follows from the above arguments that

Corollary 1. *Denote by z_i^* the optimal expected amount raised for goal G_i, then $z_i^* \geq z_j^*$ whenever $G_i < G_j$.*

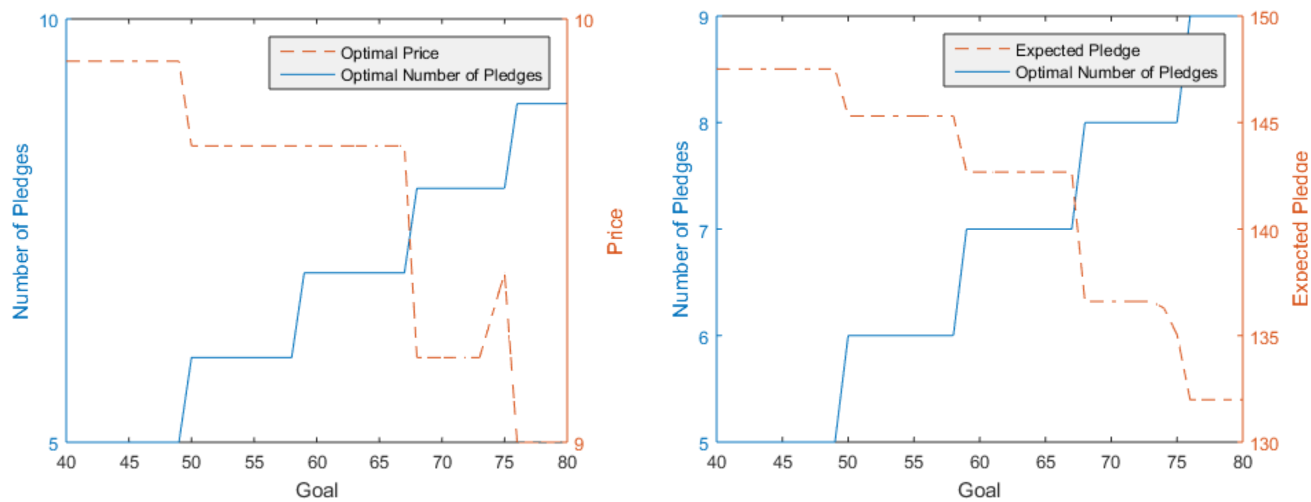

(a) Optimal k and p for Duration $n = 100$, $v^H = 10$, $\pi = 0.5$, and $c = 3$.. (b) Optimal k and Expected Pledge as a function of Goal

Figure 3: Relationship between Number of Pledges, Price, Expected Pledge, and Goal

An immediate implication of Proposition 1 and the above corollary is that the seller's optimal strategy is to *try and generate as much funding from as few backers as possible*, since every time the seller increases the number of pledges required for success the total expected pledge decreases as a result of the increased uncertainty. Moreover, because a consumer with a high valuation makes a pledge if he arrives after the goal is met, having this happen sooner than later increases the total amount pledged.

We now turn our attention to how the expected pledge varies with the campaign duration n. Regardless of the values of k and p, the success probability is always nondecreasing in n. However, because of discounting, and conditional on choosing k and p optimally, there is a unique duration that maximizes the discounted total pledge, as can be seen in Figure 4a. Naturally, for any fixed goal G, there will be values of n that are too low, so that there are not enough consumers to back the campaign. This is the case in the figure when n is less than or equal to 6, and it leads to zero funding probability and zero expected pledge. As n increases, the probability of funding and the expected pledge amount increase, until discounting makes it less profitable to extend the campaign beyond a certain time. The same effects allow us to understand how the optimal duration n^* varies with the goal, which it does in a way that is similar to how k varies with G, i.e. for every range of G, there is a single corresponding optimal value of n. Once the goal value moves out into another range, the optimal n increases. The duration being constant over a range of goal values is because this is the optimal duration that balances the probability of success and expected pledge on one hand with the effect of discounting on the other. As the goal value increases beyond that range, the same duration leads to lower probabilities and lower expected pledges and therefore, similar to Example 1, it is best to increase the duration in order to counter that effect. This can be seen in Figure 4b, where the optimal duration has a staircase shape. For example, for goal values that can be generated from a single pledge (values less than a consumer's high valuation of 10 in this example), n^* is equal to 9.

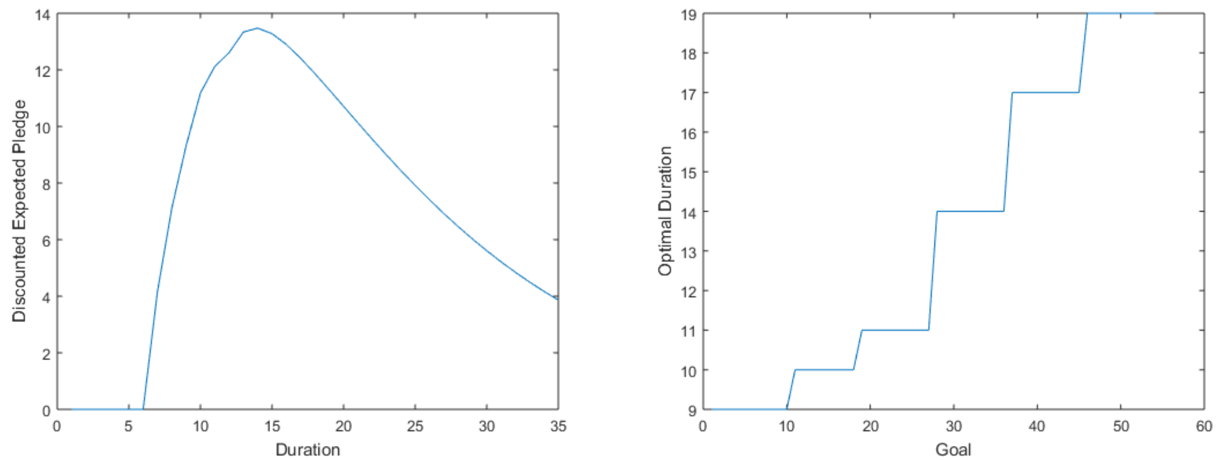

(a) Expected Pledge vs. Duration for $G = 28, v^H = 10, \pi = 0.5$, and $c = 3$.

(b) Optimal Duration n^* as a function of the Goal G

Figure 4: Relationship between Duration, Expected Pledge, and Goal

This is the value that best balances the loss from discounting with the extra pledges obtained. The value of n increases for goal value outside this range in order to increase the probability of success at the cost of more discounting. This discussion is summarized in the following proposition. The proof follows from the arguments in the text and is therefore omitted.

Proposition 2. *Let n_i^* be the optimal duration for a campaign with goal G_i, then conditional on choosing k and p optimally, $n_i^* \leq n_j^*$ whenever $G_i < G_j$.*

We have so far examined the problem under the constraint that the seller wants every successful outcome to raise at least G, i.e. that every consumer arrival sequence that is considered a success should raise at least the amount of funds required to start production. On the other hand, Corollary 1 states that lower goals (equivalently, lower k) can lead to an increase in success probability and an increase in the number of pledges. It is therefore natural to ask whether the seller should purposefully set the goal at a lower value than the one required to start production. It is possible that this can be a beneficial strategy for the seller because by setting an artificial goal $G' < G$, consumers face less uncertainty about the success prospects of the campaign, leading to more contributions, which in turn can lead to achieving the actual goal G. In considering this option, the seller assumes a risk of raising more than G' but less than G, i.e. the campaign is successful in raising the artificial goal but is short of the actual desired goal required to start production, leaving the seller with an obligation that may not be possible to fulfill. Depending on the costs associated with this scenario, the seller may elect (or not) to pursue such a strategy. In order to analyze this situation, we generalize the seller's problem to include (3) as a special case. Denote by $z = \sum_{n=1}^{n} I_n(n, k, p)$ and let M be a penalty associated with a campaign that raises what it asked for but not enough to cover the actual costs of production. The seller's objective can then be written as

$$\max_{n,k,p} \delta^n pE\left(z|z \geq G\right) - M \Pr(kp \leq z < G) \tag{13}$$

where $\Pr(kp \leq z < G)$ is the probability that the funds raised are between the reported/artificial goal and the actual target. The case where $M \to \infty$ is equivalent to the problem we started with, where the seller would never set $kp < G$. For small values of M however, a risk-neutral seller may indeed benefit from setting values for k and p that can sometimes lead to a successful campaign that is unable to deliver what it promised. The following example illustrates this.

Example 2. Consider the second campaign in Example 1 again. There is no solution for the constrained problem that leads to positive expected pledge. Now assume that the constraint $kp \geq 100$ is removed and set $k = 1$ so that all uncertainty about production is removed from the consumer's decision problem.[5] When p takes its maximum value of 10, the total expected pledge is equal to the price multiplied by the expected number of high valuation consumers conditioned on the event that at least 10 high-value consumers arrive, i.e. $E[z|z \geq 10]$, which in this scenario is equal to $\sum_{i=10}^{30} i\binom{30}{i}(0.5)^{30}$. Multiplying this by the price 10 we get \$148. To evaluate the objective in Equation (13), we need to figure out the probability of all the 'successful' outcomes that raise less than 100. This is the probability of the set of outcomes where less than 10 consumers, excluding zero, pledge, and is equal to the cumulative binomial probability of choosing 1 through 9 successes out of 30 trials at $\pi = 0.5$. This probability is equal to 0.0214. The objective value is then equal to $148 - 0.0214M$ and is nonnegative for $M \leq\approx 6915$.

It is worthwhile to note that some crowdfunding campaigns that failed on their first launch were able to succeed when the campaign was run again but with a lower target. The new campaigns were able to generate much more funds than what they asked for in their original failed campaigns. It is possible that one reason for this is the mitigation of uncertainty that is achieved from the combination of lowering the target and the previous contribution level attained in the failed campaign. As mentioned in the introduction, The Coolest Cooler campaign initially failed by raising slightly more than $\$100,000$ out of its $\$125,000$ goal, but raised millions of dollars when the campaign relaunched a few months later with a comparatively paltry goal of $\$50,000$ – a goal that was lower than what the first campaign was able to raise. Spark Core, also mentioned in the introduction, had a similar trail, going from an initial goal of $\$250,000$ for their failed campaign to a mere $\$10,000$ (and raising more than half a million dollars in the process). An interesting empirical question is verifying the effect hypothesized in this paper about how the goal level affects the outcome through the uncertainty governing consumer behavior.

On the flip side, The problem with shading the goal downwards is that the penalty M can be hard to estimate, and is typically quite large, as it is potentially associated with outcomes like litigation

[5]The case $k = 1$ provides the simplest exposition. Of course, the seller can choose k optimally in a way that best trades off the gains in expected pledge with the losses from the expected penalty.

and loss of reputation. Recent examples include the espresso machine ZPM, the iPad holder Han-Free, and the mini drone Zano (c.f. Lewis-Kraus (2015), Shontell (2015), Cellan-Jones (2015)). More recently, the Federal Trade Commission went after its first ever crowdfunding case when the seller did not deliver the promised product and instead spent the money on other expenses (c.f. Schifferle (2015)). Because of this, it is desirable for the seller to explore other options that would mitigate the uncertainty without having potentially insurmountable costs. Some of these options are described in the next section.

5 Simulations and Extensions

This section shows how sensitive the main insights of the papers are to some of our assumptions. The primary difference between the cases we discuss here and the benchmark model analyzed earlier manifests itself in the fact that the threshold for pledging can vary from one consumer to the next.

Valuations: We first check the case when consumer valuations are taken from a continuous distribution. Let v_i be the valuation of consumer i and assume that valuations for different consumers are independent and identically distributed with distribution $F(v)$. Like before, if consumer i estimates success with probability α, then he pledges if his valuation satisfies

$$v_i \geq \left(\frac{(1-\alpha)c}{\alpha} + p \right)$$

Denote the right hand side of the above inequality by $\theta_i(\alpha)$, then Backer i pledges when his valuation is at least equal to $\theta_i(\alpha)$, i.e. he pledges with probability

$$q_i(\alpha) = 1 - F(\theta_i(\alpha)), \tag{14}$$

where as mentioned F is the valuation distribution of the i^{th} backer. In the Bernoulli case $q_i(\alpha)$ is simply equal to π when $\theta_i(\alpha) \leq v^H$ and 0 otherwise.

Let α_i^l denote the probability of successful funding that backer i estimates, given that l pledges are made in $1, ..., i$, then α_i^l can be written recursively as

$$\alpha_i^\ell = \begin{cases} q_{i+1}(\alpha_{i+1}^{\ell+1}) \cdot \alpha_{i+1}^{\ell+1} + (1 - q_{i+1}(\alpha_{i+1}^{\ell+1})) \cdot \alpha_{i+1}^{\ell} & i < n \\ 1 & i = n, \ell \geq k \\ 0 & i = n, \ell < k \end{cases} \tag{15}$$

We examine how the success probability varies with the number of required pledges k as the other parameters are held constant. Figure 5a shows that the bimodal structure is conserved: campaigns either succeed with probability 1 or zero, with a narrow transition band between both outcomes. One important point is how the price affects the comparative sizes of the success and failure regions. As the price increases the size of the success region, in terms of k, shrinks while the size of the failure

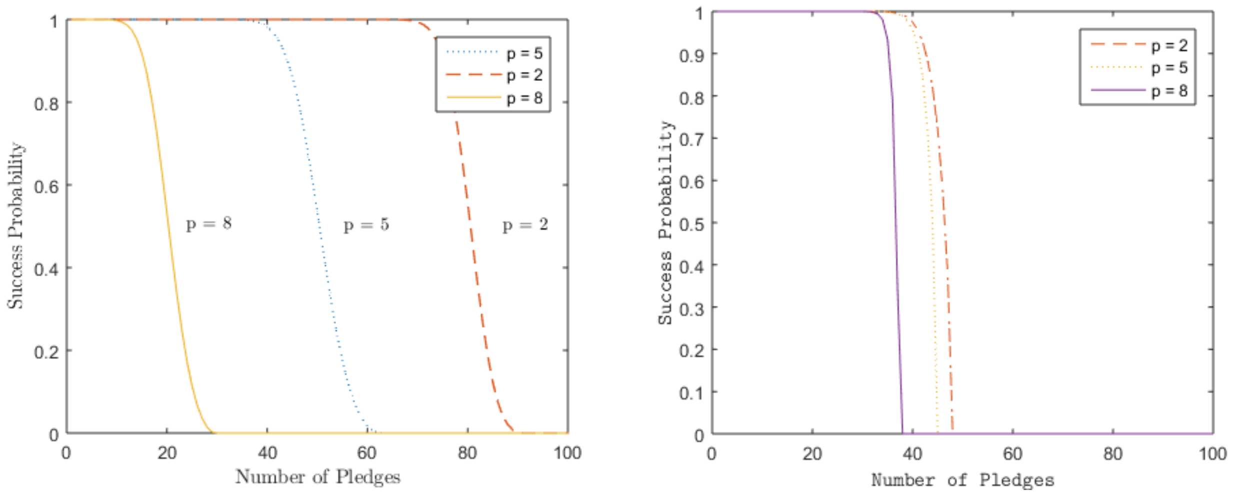

(a) $v_i \sim U(0, 10)$ (b) $v_i = 10$ with $\pi = 0.5$

Figure 5: Success Probability vs. Number of Required Pledges k for Duration $n = 100$, $c = 3$, and prices $p = \$2, \$5$, and $\$8$.

region increases. Contrast this with the benchmark model in Figure 5b, where the size of the two regions is not as sensitive to price. This is a simple consequence of the fact that in the benchmark model price is alway above the high valuation, and the probability of a consumer with valuation below this price does not change as the price increases further. Price adjustments only change the number of consumers who pledge based on how their high valuations adjust downward because of uncertainty. In the continuous case, the probability of a consumer having a valuation below a certain price changes with the price, and each price has a range of consumers who will not pledge, regardless of uncertainty. Other consumers take this into account and whenever the price increases, the number of potential backers decreases as does the probability of success, leading to a larger size of the failure region. The opposite reasoning applies when the price increases.

Uncertainty cost: The assumption for the benchmark model was that the uncertainty cost c is the same for every one. The effect of c on how success probability varies with k is similar to the effect of price: As c increases, the size of the success region shrinks while the size of the failure region increases. Under the interpretation of c as an opportunity cost, it is reasonable to relax the above assumption and instead make the uncertainty cost dependent on the time of arrival – since consumers who arrive early face a longer duration where their funds are tied up compared to consumers who arrive later. Thus, instead of assuming that c is the same for all consumers, we instead consider the case where consumers inversely-discount time, so that pledging early is more costly than pledging later. Let γ be the discount factor for consumers, then the utility that backer i gets for pledging is now given by

$$u_i = \begin{cases} v_i \, \gamma^{n-i} - p & \text{if funded successfully} \\ p \, \gamma^{n-i} - p & \text{if not funded} \end{cases}$$

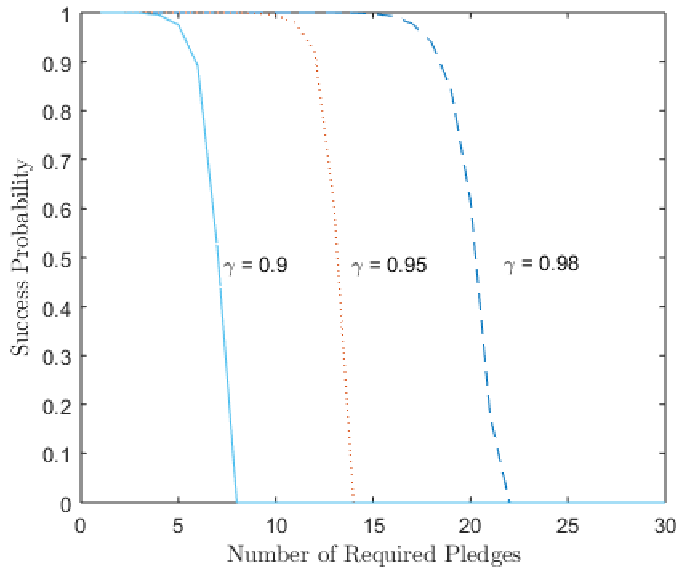

Figure 6: Success Probability vs. Number of Required Pledges k for Duration $n = 100$, $v_i \sim U(0, 10)$, Price $p = \$4$, and discount factors $\gamma = 0.9, 0.95$, and 0.98.

As before, if Backer i believes the campaign will succeed with probability α, then his expected utility is

$$u_i = \alpha \cdot (v_i - p) \cdot \gamma^{n-i} - p \cdot (1 - \gamma^{n-i}).$$

which means that backer i will pledge only if

$$v_i \geq p \cdot \left(\frac{\left(\frac{1}{\gamma}\right)^{n-i} - 1}{\alpha} + 1 \right)$$

If the right hand side above is again denoted by $\theta_i(\alpha)$, then the expressions for the probability of pledging and the ensuing recursion are identical to Eqns. (14) and (15). Figure 6 shows that the overall structure is again maintained, with the effect of different discount factors similar to the effect of different uncertainty costs. The closer the discount factor is to 1, the bigger the success region is. One noticeable difference is that the failure region is much larger compared to the case when the cost does not depend on time. This is because consumers who arrive early in a very long campaign face such a high discount that they never pledge. It is possible that this is one of the reasons why Kickstarter has changed its rules about campaign duration, enforcing the rule that no campaign can be longer than 60 days.

Essentially, the primary difference between these variations and the benchmark model is that consumer heterogeneity –whether it is introduced through different valuations, different uncertainty costs, or different discounts tied to arrival times– changes the right hand side of the threshold for pledging. While this significantly complicates the theoretical analysis of these cases, numerically solving the resulting dynamic programs does not lead to qualitatively different insights than the ones developed in the benchmark model.

Dynamic Pricing: Section 4 discussed the possibility of the seller shading the goal in order to reduce consumers' uncertainty and obtain better outcomes in expectation. Given the risk associated with setting a funding goal lower than the actual target and the difficulty of assigning a monetary value to negative outcomes, it is natural to ask whether there is a scheme that can improve campaign outcomes within our simple framework and without resorting to such measures. One such scheme can be arrived at through observing the disparity in uncertainty between consumers and the effect of this disparity on campaign success. For example, consider the last consumer to arrive. This consumer's decision problem is trivial, as he observes the current total pledge and knows immediately whether the campaign will succeed or not. In contrast, the first consumer's problem is fraught with uncertainty about what type of consumers will arrive and their contribution decisions. Consumers who arrive at different times between those two extremes face different uncertainties as well depending on what happened so far and their time of arrival. Despite this asymmetry, all consumers face the same probability threshold for pledging, despite some consumers bearing more of the brunt of getting the campaign up and running. Based on this, a natural candidate to improve campaign outcomes is to compensate consumers based on the uncertainty they have to deal with. Because uncertainty is heterogenous over time, a discount that adjusts based on uncertainty can be implemented through dynamic pricing. The following is a simple example that highlights how dynamic pricing can improve outcomes.

Example 3. Consider an environment with $n = 3, v = 10, \pi = 0.5$, $c = 3$, and $G = 15$. The optimal solution under the single price scheme we have studied so far is to set $p = \$9$ and $k = 2$. This gives probability of success equal to 0.375 and expected pledge equal to 7.875. Basically, this campaign cannot succeed unless the first consumer has high valuation. If this does not happen, then the second consumer will not pledge at any price, even if he has high valuation. This is because the uncertainty that he faces is worse than what the first consumer faces. For the first consumer, assuming he has a high valuation, the probability of success depends on whether one of the remaining two consumers will also have high valuation, i.e. the probability of success is equal to $1 - (0.5)^2 = 0.75$ whereas a second consumer arriving after no pledge has been made by the first consumer faces a probability of success that is equal to the third consumer having high valuation = $\pi = 0.5 < 0.75$.

If a dynamic pricing scheme is used however, the latter scenario can be corrected as follows. If the first consumer does not pledge at price $\$p = 9$, then the price decreases to $\$7$ for the second consumer and goes up to $\$10$ for the third consumer. The second consumer's utility is thus equal to $0.5(10 - 7) - 0.5(3) = 0$ and he pledges. The last consumer is charged $\$10$ and he pledges as well since there is no uncertainty. This means that the goal will be reached whenever two high-valuation consumers arrive, regardless of *when* they arrive (as opposed to the single price setting where the goal cannot be reached unless the first consumer has high valuation). Consequently, the probability of success in this case is equal to $0.5 > 0.375$. In addition to the increase in probability, there is also an increase in the expected pledge, since if the first consumer has high valuation and pledges, then consumers 2 and 3 are charged the full valuation $\$10$.

One important feature of dynamic pricing observed in the previous example is that it can be nonmonotonic. The price may go up and down over time depending on the compensation that consumers get for the uncertainty they face. Since uncertainty is nonmonotonic, so are the prices. Such a scheme is unlikely to be utilized in practice: consumers who have pledged at a high price in an early period may rightly be offended that the pledge value has decreased and that consequent consumers are getting a better deal. In fact, Kickstarter explicitly prevents such scenarios from occurring by not allowing sellers to change the pledge value once a consumer has made a pledge.[6]

An indirect way to approximate dynamic pricing is through using a menu of prices. It is possible that consumers are given several price options that are available at all times. A consumer would then choose the price that maximizes their expected utility. This is not necessarily the lowest price because each price brings with it a tradeoff: while a higher price may decrease the net value a consumer gets from obtaining the product, it can also increase the probability of getting that product as the campaign edges closer to its goal. To continue with the example above, assume that a consumer can choose from one of two pledge options, $\$7$ or $\$10$, then the success probability matches that obtained by dynamic pricing, i.e. funding is successful whenever two or more out of the three consumers have high valuations, regardless of arrival times. Consider the first consumer and assume she has high valuation, then she chooses to pledge $\$7$. If the second consumer has high valuation as well, her options are to either pledge $\$7$ and wait to see if the third consumer is a high-valuation consumer (since the total combined pledge is still short of the $\$15$ goal), or to pledge $\$10$ and remove any uncertainty by making the campaign succeed. Both of these options have the same expected utility of zero (we assume ties are broken by choosing to pledge, or alternatively, the pledge amounts can be made equal to $\$7$ and $\$10 - \epsilon$ in order to make the payoff from selecting the second option strictly positive). If the first consumer does not pledge and the second consumer has high-valuation, then she pledges $\$7$ – exactly the same as in the case of dynamic pricing. The last consumer pledges $\$10$ in this case to push the campaign over the success threshold. While the probability of success under this particular menu matches that of dynamic pricing, the expected pledge is less, since once a campaign is successful (for example if the first two consumers pledge) then any other high-valuation consumer that follows will choose the lowest pledge value.

The last point can be addressed by limiting or capping the number of pledges that can be made at the different levels. For example, it is possible to have only one 'slot' for the $\$7$ pledge and two for the $\$10$ pledge. This ensures that the cheaper pledge is used only once (either by the first or second consumer, since this is all that it is required in this example depending on the realization of values) and therefore the case above, where the last consumer will always select the cheaper $\$7$ option, cannot happen since that option will be 'sold-out' by then. This method closely mirrors how crowdfunding –

[6]More generally, Kickstarter campaigns usually have different tiers for contribution at different pledge levels. Once someone pledges in a tier, the price associated with that tier cannot be changed.

and particularly, Kickstarter– campaigns are organized, with different tiers of pledge values that each have different capacities. Designing these menus in terms of the appropriate prices and capacities of each price tier is an interesting problem for future work.

6 Final Remarks

The success of a crowdfunding campaign hinges on understanding and accurately predicting the behavior of potential backers. A model that accounts for the myriad behavioral and other idiosyncratic factors that go into a backing decision is likely to be analytically intractable. In this paper, we isolate what we believe is a salient feature of these campaigns – backers' uncertainty about the outcome, and show that a simple model that incorporates this uncertainty and in which backers make decisions sequentially produces predictions that agree with observed empirics. We use this model to provide general guidelines about campaign design. In particular, we show that it is best to try and raise as much funding from as few backers as possible through charging a high price instead of charging a lower price and relying on more consumers to raise the required capital. The optimal price has two components: the first extracts the valuation of the consumers, and the second is a discount that consumers get as a compensation for the uncertainty they face.

The paper's central contribution is the introduction of anticipating random walks. These stochastic processes have the property that their stopping times depend on what might happen in the future, which complicates their analysis. Theorem 1 provides a general result about the probability of success of these walks, and can be applied to a broad range of settings. For example, one application is understanding how candidates raise money from donors in election campaigns, or how voters make decisions in sequential voting processes. The theorem provides an insight into the inner workings of these and similar stochastic situations and we hope that it finds applicability outside of the setting we employ it for in the paper.

Our model makes some simplifying assumptions to focus on the uncertainty aspect and to achieve tractability. While one can imagine multiple customers arriving and making a decision within a day, the sequential, one-by-one arrival of customers that we employ in our model can be thought of as the limit case when each time period is small enough that no more than one customer arrives within a period. While consumers may have a continuum of valuations, the effect of this heterogeneity will be reflected in choosing the appropriate price that balances a consumer's valuation with the probability of future arrivals of consumers whose valuations are above a certain threshold. An extension of the model in this paper would incorporate a full price menu that aims to exploit both the exogenous variation in valuations and the endogenous variation in uncertainty that consumers face.

The decisions taken by consumers in our model are informed by what happened in the past and by their estimation of what will happen in the future (and how they can affect that through their own actions). The decisions were linked together through the dependence of the outcome on the aggregation of individual actions. It is possible that current pledges affect not just the future decisions of pledging, but also future valuations. For example, a campaign that receives so much funding at some point may convert consumers who were initially not interested in the product or who have had low valuations to desiring the product and, consequently, to pledging themselves. There are several mechanisms that may lead to this (for example, a campaign may have some early runaway success that gets it press coverage and therefore exposes it to a wider section of the population, a high current pledge can be used as a proxy for quality, etc.). While analyzing this problem is important and might result in better campaign management strategies, it is a scenario that is difficult to capture in a dynamic programming framework because the state space is exorbitantly large.

Another interesting aspect worthy of exploration is the information disclosure policy of the campaign. In particular, should the campaign designer have a real time update of how much the campaign has raised or should this amount not be made available? It is easy to think about the pros and cons of these extreme strategies. If the current funding level is always observable, as is the current practice on all crowdfunding platforms, then it is possible that consumers start pledging as soon as they see that the campaign is doing well. Conversely, if the campaign is off to a slow start, then maybe that will put off potential consumers from pledging. Not revealing the pledge amount, while it may circumvent the later scenario to a certain extent, also brings about uncertainty and ambiguity about how the campaign is going, and as a result may end up delivering undesirable outcomes as well. One potential candidate for an optimal information disclosure policy that combines the benefits of the aforementioned extremes may by a periodic disclosure policy, where the designer commits to announcing the current pledge amount every T units of time.

Finally, crowdfunding systems are characterized by extreme decentralization, information asymmetries, and complicated dynamics of interaction between campaign owners, funders, and platforms. As such, no one model is likely to provide a singular explanation for all possible outcomes. As decentralized funding becomes more pervasive, the need to study the different aspects of these systems theoretically and empirically becomes increasingly important. We hope that this paper and the different future directions outlined herein serve as a step in this direction.

Appendix

A Additional Proofs

Proof of Lemma 1 We first show that b_j^t is decreasing in j, using induction on t. For $t = 0$, $b_0^0 = 1$ and $b_j^0 = 0$ for all $j > 0$. We next assume that for all $t' < t$, $b_j^{t'} \geq b_{j+1}^{t'}$, for all j, and we prove that $b_j^t \geq b_{j+1}^t$, for all $j < \lambda^{t-1}$. Using (5), we have

$$b_j^t = q_t b_{j-1}^{t-1} + (1 - q_t) b_j^{t-1} \geq q_t b_j^{t-1} + (1 - q_t) b_{j+1}^{t-1} = b_{j+1}^t,$$

where the first inequality holds by induction. Moreover, for $j = \lambda^{t-1}$, by definition $b_j^t > b_{j+1}^t = 0$, and for $j > \lambda^{t-1}$, $b_j^t = b_{j+1}^t = 0$, completing the inductive proof of the first part. Using this property, we next show that b_j^t is weakly increasing in t. For all buckets with $b_j^{t-1} > 0$, using (5), we have

$$b_j^t = q_t b_{j-1}^{t-1} + (1 - q_t) b_j^{t-1} = b_j^{t-1} + q_t (b_{j-1}^{t-1} - b_j^{t-1}) \geq b_j^{t-1},$$

where the last inequality follows from b_j^t being weakly decreasing in j.

Finally by definition, at τ_{j+1}, $b_{j+1}^{\tau_{j+1}} > 0$ for the first time, which can only happen if $b_j^{\tau_{j+1}-1} \geq \beta$. Using the first two parts of this lemma, it implies that for all $j' \leq j$, and for all $t' \geq \tau_{j+1} - 1$, $b_{j'}^{t'} \geq \beta$. Similarly, at $t' < \tau_{j+1}$, $b_{j+1}^{t'} = 0$, which implies that $b_j^{t'-1} < \beta$. Combining this with the first two parts of the lemma we have, for all $t' < \tau_{j+1} - 1$, and for all $j' \geq j$, $b_{j'}^{t'} < \beta$, completing the proof.

■

Proof of Theorem 3 Without loss of generality we assume that valuations come from the same Bernoulli distribution so that the term $\sum_i \pi_i$ is replaced by $n\pi$, where π is the mean of the Brenoulli distribution. We first show that $Y_i = Y_{i-1} + X_i$ for all $0 \leq i \leq n_0 = n\frac{\delta}{1+\delta}$, i.e. a consumer will pledge if he has high valuation. Let $\theta = \ln \frac{1}{1-\beta}$, and $\bar{\pi} = \frac{k}{n}(1+\theta)$. Assuming $n_0 = n\frac{\delta}{1+\delta}$, we have

$$(n - n_0)\pi \geq n(1 - \frac{\delta}{1+\delta})\bar{\pi}(1+\theta)(1+\delta) = n\bar{\pi}(1+\theta).$$

Using Theorem 1, this would imply that at step $i \leq n_0$, the probability of $[Y_n \geq k] \geq \beta$, therefore by Definition 1 we have $Y_i = Y_{i-1} + X_i$.

We next show that for all $i \geq n_0$, if $\sum_{j=0}^{i} X_j \geq \frac{ki}{n} - k\delta(1 - \frac{i}{n})$, then $Y_i = Y_{i-1} + X_i$. The proof is by induction. Assume that until time i if the condition holds, then $Y_j = Y_{j-1} + X_j$ for $j < i$. We then show that consumer i will also pledge. Note that when deciding at time i, assuming that all consumers with high type pledge, we have

$$\begin{aligned} P[Y_n \geq k] = Pr[Y_{n-i} \geq k - \sum_{j=0}^{i} X_j] \geq Pr[Y_{n-i} \geq k - \frac{ki}{n} + k\delta(1 - \frac{i}{n})] \qquad (16) \\ = Pr[Y_{n-i} \geq \frac{k}{n}(1+\delta)(n-i)] \geq \beta. \end{aligned}$$

Note that

$$\pi(n-i) = \frac{k}{n}(1+\delta)(n-i)(1+\theta),$$

and combining this with Theorem 1, gives us Eq. (16).

Let ε_i denote the event that at time i, $\sum_{j=1}^{i} X_j \leq \frac{ki}{n} - k\delta(1 - \frac{i}{n})$. We show that for any $i \geq n_0$, $P[\varepsilon_i] \leq \epsilon$.

Using Hoeffding's inequality we have,

$$
\begin{aligned}
Pr[\varepsilon_i] &= Pr\left[\sum_{j=1}^{i} X_i \leq \frac{ki}{n} - k\delta(1-\frac{i}{n})\right] = Pr\left[\frac{1}{i}\sum_{j=0}^{i} X_i \leq \frac{k}{n} - \frac{k\delta}{i}(1-\frac{i}{n})\right] \\
&= Pr\left[(E_i[X] - \frac{1}{i}\sum_{j=0}^{i} X_i) \geq \frac{k}{n}(\delta + \theta + \delta\theta + \frac{\delta(n-i)}{i})\right] \\
&\leq e^{-2i\bar{\pi}^2(\delta+\theta(1+\delta)+\delta\frac{n-i}{i})^2}.
\end{aligned}
$$

We are interested in the probability of the event $\bigcap_{i=n_0}^{n} \bar{\varepsilon_i}$. We then have

$$
Pr\left[\bigcap_{i=n_0}^{n} \bar{\varepsilon_i}\right] = 1 - \bigcup_{i=n_0}^{n} \varepsilon_i \geq 1 - \sum_{i=n_0}^{n} e^{-2i\bar{\pi}^2(\delta+\theta(1+\delta)+\delta\frac{n-i}{i})^2} \geq 1 - e^{-2n_0\bar{\pi}^2\theta^2}.
$$

Therefore having $\delta \geq \frac{\ln(\frac{1}{\epsilon})}{2n\bar{\pi}^2\theta^2}$, the campaign will be funded with probability at least $1-\epsilon$. ■

Proof of Proposition 1 Partition R^+ into intervals $G_1, G_2, \ldots$ such that $g \in G_i < g' \in G_j$ whenever $i < j$.[7] We start by proving a slightly stronger variant of the statement $k_1^* \leq k_2^*$: that k_i^* is constant over G_i and is the smallest feasible k for any $g \in G_i$, i.e. that there is no price p such that $kp \geq g$ if $k < k_i^*$. The proof is by induction. Fix n and let v^H and $\pi \in (0,1)$ be the high valuation and the probability of having a high valuation, respectively. For the base case, consider the first interval $G_1 = (0, v^H]$. For any $g \in G_1$, it is optimal to set k_1^* to 1 and p_1^* to v^H, for an expected pledge of $\pi n v^H$. Setting $k_1 = 0$ yields $0 < g$ and is thus infeasible, and setting $k_1 > 1$ (while keeping $p_1 = v^H$) leads to zero expected pledge. To see this, let u be the expected utility of the first consumer to pledge and assume that this consumer estimates the probability of campaign success to be equal to α. Because $k_1 > 1$, $\alpha < 1$, and that consumer's expected utility is equal to $u = \alpha(v^H - p) - (1-\alpha)c < 0$. Therefore this consumer will never pledge and the total pledge expected under $p_1 = v^H$ is equal to zero. Since u increases with decreasing p, the point at which that consumer pledges and expected revenue increases above zero –if it exists– must happen at $p_1' < p_1^* = v^H$. But if that is the case then the expected pledge under this solution is bounded above by $\pi n p_1'$, which is less than $\pi n v^H$, and hence $k_1^* = 1$ and $p_1^* = v^H$ is the optimal solution for any $g \in G_1$. This proves the first two bullet points for the base case. To show the third point, let $G_2 = (v^H, v]$, for some v. For any $k_2 \leq 1$ and any feasible price $p_2 \leq v^H$, we have $k_2 p_2 \leq v^H$, which is not feasible for any $g \in G_2$, and hence $k_2^* > 1$, but if that is the case then by the same argument as above, $p_2^* < p_1^*$, since if not, the expected utility of the first consumer is again negative and the total expected pledge is equal to zero. The third point follows by selecting $g_1 \in G_1$ and $g_2 \in G_2$ and using the previous argument to conclude that $k_1^* < k_2^*$ and $p_1^* > p_2^*$. This concludes the base case.

[7] We slightly abuse notation for the purposes of this proof and refer to a specific goal value with a lowercase g and to a range of goal values with a capital G.

For the induction hypothesis, consider interval $G_i = [g_i, g_i')$ and let k_i^* be the smallest k such that $\exists p_i \; kp_i \geq g \in G_i$. Denote by k_{i+1}^* and p_{i+1}^* the optimal k_{i+1} and p_{i+1} for $g' \in G_{i+1} = [g_i', g_{i+1})$ and assume that $k_{i+1}^* < k_i^*$, then this implies that k_i^* is not the smallest feasible k for $g \in G_i$, since $k_{i+1}^* p_{i+1}^* \geq g' > g$, contradicting the induction hypothesis and implying that $k_{i+1}^* \geq k_i^*$, which proves the first bullet point. Assume that $k_{i+1}^* = k_i^*$ and denote by p_i^* the optimal price for $g \in G_i$, then $p_{i+1}^* \geq p_i^*$ as in the statement of the theorem. To see this, assume that $p_{i+1}^* < p_i^*$. Because the solution (k_i^*, p_{i+1}^*) is feasible for g', we have $k_i^* p_{i+1}^* \geq g' > g$, and hence (k_i^*, p_{i+1}^*) is feasible for g as well. Similarly, because $p_i^* > p_{i+1}^*$, (k_i^*, p_i^*) is feasible for both g and g'. This means that if the expected pledge under (k_i^*, p_{i+1}^*) is optimal for g', then it must also be optimal for g, contradicting the optimality of p_i^* and hence $p_{i+1}^* \geq p_i^*$, proving the second bullet point. Finally, assume that $k_{i+1}^* > k_i^*$. Pick any $g \in G_i$ and let p_i^* be the optimal price corresponding to that goal. Let g_{i+1} be such that $k_{i+1}^* p_i^* \geq g_{i+1}$ (such g_{i+1} exists by continuity, by taking $g_{i+1} = g^m + \epsilon$ where $g^m = \max g \in G_i$ and $\epsilon > 0$). Now assume that $p_{i+1}^* > p_i^*$ for some $g' \in G_{i+1}$. This means that the expected pledge under the solution (k_{i+1}^*, p_{i+1}^*) is higher than the expected pledge under the solution (k_{i+1}^*, p_i^*), but if that is the case, then the solution (k_i^*, p_{i+1}^*) also has higher expected pledge for $g \in G_i$, since the decrease in k can only lead to higher expected pledge, contradicting the optimality of p_i^* and implying that $p_{i+1}^* \leq p_i^*$.

■

B Recursion for Expected Pledge

Consider the solution to the dynamic program in Equation (4). This solution provides the success probabilities s_i^j for all arrivals i and pledges j . The probability threshold for pledging is given by $\beta = \frac{c}{v-c-p}$. Define ρ_i^j to be the probability that consumer i pledges when there has been j pledges made up to but *not including* i and let γ_i^j be the probability that there have been j pledges made in i arrivals. Using the output of Equation (4), ρ_i^j is given by

$$\rho_i^j = \begin{cases} \pi & \text{if } s_i^j \geq \beta \\ 0 & \text{else.} \end{cases}$$

We can then compute the probabilities γ_i^j going forward as follows

$$\gamma_i^j = \begin{cases} \gamma_{i-1}^{j-1} \rho_i^{j-1} + \gamma_{i-1}^j (1 - \rho_i^j) & \text{if } i, j > 1 \\ \rho_1^0 & i = j = 1 \\ 1 - \rho_1^0 & i = 1, j = 0. \end{cases}$$

and the total expected pledge is then given by $\sum_{j=k}^n j \gamma_n^j$, which is the sum of the products of the probability of getting k or more pledges with the corresponding number of pledges.

References

Acemoglu, Daron, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar (2011), "Bayesian learning in social networks." *The Review of Economic Studies*, 78, 1201–1236.

Acemoglu, Daron and Asuman Ozdaglar (2011), "Opinion dynamics and learning in social networks." *Dynamic Games and Applications*, 1, 3–49.

Agrawal, Ajay K, Christian Catalini, and Avi Goldfarb (2011), "The geography of crowdfunding." Technical report, National bureau of economic research.

Agrawal, Ajay K, Christian Catalini, and Avi Goldfarb (2013), "Some simple economics of crowdfunding." Technical report, National Bureau of Economic Research.

Anand, Krishnan S and Ravi Aron (2003), "Group buying on the web: A comparison of price-discovery mechanisms." *Management Science*, 49, 1546–1562.

Andreoni, James (2006), "Leadership giving in charitable fund-raising." *Journal of Public Economic Theory*, 8, 1–22.

Araman, Victor F and Rene Caldentey (2016), "Crowdvoting the timing of new product introduction." *Available at SSRN 2723515*.

Bagnoli, Mark and Barton L Lipman (1989), "Provision of public goods: Fully implementing the core through private contributions." *The Review of Economic Studies*, 56, 583–601.

Belleflamme, Paul, Thomas Lambert, and Armin Schwienbacher (2013), "Crowdfunding: Tapping the right crowd." *Journal of Business Venturing*.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy*, 992–1026.

Cachon, Gérard and Christian Terwiesch (2009), *Matching supply with demand*, volume 2. McGraw-Hill Singapore.

Cachon, Gérard P (2004), "The allocation of inventory risk in a supply chain: Push, pull, and advance-purchase discount contracts." *Management Science*, 50, 222–238.

Cachon, Gérard P and Martin A Lariviere (2001), "Contracting to assure supply: How to share demand forecasts in a supply chain." *Management science*, 47, 629–646.

Cachon, Gérard P and Martin A Lariviere (2005), "Supply chain coordination with revenue-sharing contracts: strengths and limitations." *Management science*, 51, 30–44.

Cellan-Jones, Rory (2015), "Kickstarter's zano drone fails to fly." URL `http://www.bbc.com/news/34787404`.

Drakopoulos, Kimon, Asuman Ozdaglar, and John N Tsitsiklis (2013), "On learning with finite memory." *Information Theory, IEEE Transactions on*, 59, 6859–6872.

Hu, Ming, Xi Li, and Mengze Shi (2015), "Product and pricing decisions in crowdfunding." *Marketing Science*.

Hu, Ming, Mengze Shi, and Jiahua Wu (2013), "Simultaneous vs. sequential group-buying mechanisms." *Management Science*, 59, 2805–2822.

Huang, Yan, Param Vir Singh, and Kannan Srinivasan (2014), "Crowdsourcing new product ideas under consumer learning." *Management Science*, 60, 2138–2159.

Jing, Xiaoqing and Jinhong Xie (2011), "Group buying: A new mechanism for selling through social interactions." *Management science*, 57, 1354–1372.

Kuppuswamy, Venkat and Barry L Bayus (2013), "Crowdfunding creative ideas: the dynamics of projects backers in kickstarter." *SSRN Electronic Journal*.

Lewis-Kraus, Gideon (2015), "Zpm espresso and the rage of the jilted crowdfunder." URL http://www.nytimes.com/2015/05/03/magazine/zpm-espresso-and-the-rage-of-the-jilted-crowdfunder.html.

List, John A and David Lucking-Reiley (2002), "The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign." *Journal of Political Economy*, 110, 215–233.

Lobel, Ilan, Evan D Sadler, and Lav R Varshney (2015), "Customer referral incentives and social media." *Forthcoming, Management Science*.

Marinesi, Simone and Karan Girotra (2013), "Information acquisition through customer voting systems."

Mollick, Ethan (2014), "The dynamics of crowdfunding: An exploratory study." *Journal of business venturing*, 29, 1–16.

Mollick, Ethan and Venkat Kuppuswamy (2014), "After the campaign: Outcomes of crowdfunding." *Available at SSRN*.

Mutz, Diana C (1995), "Effects of horse-race coverage on campaign coffers: Strategic contributing in presidential primaries." *The Journal of Politics*, 57, 1015–1042.

Schifferle, Lisa Weintraub (2015), "Dont let crowdfunding be your doom." URL https://www.ftc.gov/news-events/blogs/business-blog/2015/06/dont-let-crowdfunding-be-your-doom.

Shontell, Alyson (2015), "Kickstarter's zano drone fails to fly." URL `http://www.businessinsider.com/how-one-stupid-mistake-and-35000-from-kickstarter-made-an-ave`

Varian, Hal R (1994), "Sequential contributions to public goods." *Journal of Public Economics*, 53, 165–186.

Vesterlund, Lise (2003), "The informational value of sequential fundraising." *Journal of Public Economics*, 87, 627–657.

Wu, Jiahua, Mengze Shi, and Ming Hu (2014), "Threshold effects in online group buying." *Management Science*, 61, 2025–2040.