

DISCUSSION PAPER SERIES

DP11363

LEARNING THROUGH CROWDFUNDING

Gilles Chemla and Katrin Tinn

**FINANCIAL ECONOMICS and
INDUSTRIAL ORGANIZATION**



LEARNING THROUGH CROWDFUNDING

Gilles Chemla and Katrin Tinn

Discussion Paper DP11363

Published 28 June 2016

Submitted 08 December 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS and INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Gilles Chemla and Katrin Tinn

LEARNING THROUGH CROWDFUNDING

Abstract

We develop a model where reward-based crowdfunding enables firms to obtain a reliable proof of concept early in their production cycle. Crowdfunding allows firms to learn about total demand from a limited sample of target consumers pre-ordering a new product. Learning creates a valuable real option as firms invest only if updated in and out of sample demand is sufficiently high. This is particularly valuable for firms facing a high degree of uncertainty about consumer preferences, such as developers of innovative consumer products. The real option value of learning enables these firms to overcome moral hazard, even if diverting funds is costless. The higher the funds raised, the lower the firms' incentives to divert them, provided third-party platforms limit the sample size by restricting campaign length. Expected funds raised are maximized at an intermediate sample size. Our results are consistent with stylized facts and lead to new empirical implications.

JEL Classification: D80, G30, L14, L26, O30

Keywords: Reward-based crowdfunding, moral hazard, real options, learning, uncertainty, Kickstarter

Gilles Chemla - g.chemla@imperial.ac.uk
Imperial College Business School, CNRS, and CEPR and CEPR

Katrin Tinn - k.tinn@imperial.ac.uk
Imperial College Business School

Acknowledgements

We thank Bruno Biais, Matthew Ellman, Sjaak Hurkens, Raj Iyer, Lubos Pastor, Jean-Charles Rochet, Wilfried Sand-Zantmann, Antoinette Schoar, participants at the CEPR 2015 European Summer Symposium in Financial Markets, the ESRC 2015 "Crowdfunding and other alternative forms of entrepreneurial finance" conference, the Imperial College Business School Fintech 2016, the Bank of Estonia's 2015 Winter conference, the Imperial College FinTech Showcase 2016, the Paris-Dauphine 2016 Corporate Finance conference, the Cambridge Centre for Alternative Finance 2017 conference, the Ecole Hoteliere de Lausanne 2017 Private Markets conference, the ESCP 2017 Fintech conference, and seminar participants at Rochester Simon School of Business, Cambridge Judge Business School, SKEMA, Telecom Paristech, and Toulouse for helpful discussions.

Learning through Crowdfunding*

Gilles Chemla

Imperial College Business School, DRM/CNRS, and CEPR.

Katrin Tinn

Imperial College Business School

First Draft: February 2015

This Draft: December 5, 2017

Abstract

We develop a model where reward-based crowdfunding enables firms to obtain a reliable proof of concept early in their production cycle. Crowdfunding allows firms to learn about total demand from a limited sample of target consumers pre-ordering a new product. Learning creates a valuable real option as firms invest only if updated in and out of sample demand is sufficiently high. This is particularly valuable for firms facing a high degree of uncertainty about consumer preferences, such as developers of innovative consumer products. The real option value of learning enables these firms to overcome moral hazard, even if diverting funds is costless. The higher the funds raised, the lower the firms' incentives to divert them, provided third-party platforms limit the sample size by restricting campaign length. Expected funds raised are maximized at an intermediate sample size. Our results are consistent with stylized facts and lead to new empirical implications.

JEL codes: D80, G30, L14, L26, O30

Keywords: Reward-based crowdfunding, moral hazard, real options, learning, uncertainty.

*Correspondance to Imperial College Business School, South Kensington campus, London SW7 2AZ, UK. E-mail: g.chemla@imperial.ac.uk and k.tinn@imperial.ac.uk. We thank Bruno Biais, Matthew Ellman, Sjaak Hurkens, Raj Iyer, Lubos Pastor, Jean-Charles Rochet, Wilfried Sand-Zantmann, Antoinette Schoar, participants at the CEPR 2015 European Summer Symposium in Financial Markets, the ESRC 2015 "Crowd-funding and other alternative forms of entrepreneurial finance" conference, the Imperial College Business School Fintech 2016, the Bank of Estonia's 2015 Winter conference, the Imperial College FinTech Showcase 2016, the Paris-Dauphine 2016 Corporate Finance conference, the Cambridge Centre for Alternative Finance 2017 conference, the Ecole Hoteliere de Lausanne 2017 Private Markets conference, the ESCP 2017 Fintech conference, and seminar participants at Rochester Simon School of Business, Cambridge Judge Business School, SKEMA, Telecom Paristech, and Toulouse for helpful discussions.

1 Introduction

Reward-based crowdfunding platforms enable firms to raise funds directly from future consumers, typically in exchange for a promise of the future delivery of a new product. While they have been seen as means for artists and credit constrained firms to seek support for their projects, primary beneficiaries of these platforms nowadays appear to be firms that develop innovative consumer products.

Indeed, data we have obtained from Kickstarter show that as of November 15, 2017, *i*) 63% of the \$3.0 billion successfully raised on that platform have gone to firms selling technology, design or gaming products, and *ii*) an average successfully funded project in technology, design and games raised \$94K, \$62K and \$55K, respectively, but only \$10K on average across all other categories, and *iii*) 240 out of the 265 (respectively, 3,551 out of the 4,545) Kickstarter projects that have raised over \$1 million (respectively \$100K) were in the three categories aforementioned.

In this paper we develop a theoretical model to understand the prime sources of value creation in reward-based crowdfunding, and why it is particularly attractive for innovative projects. We argue that reward-based crowdfunding platforms play an important role in enabling firms to test out their market at an early stage of product development. Pre-selling a product through these platforms acts as a credible consumer survey where firms learn about target consumer preferences and demand before making their investment decisions. We show that this creates a substantial real option value of learning: observing the decisions of a random sub-sample of consumers (backers) enables the firm to update its beliefs about the preferences of *all* their future consumers. There is value in both success and failure: regardless of whether the project seems profitable ex-ante, firms either benefit from learning that their product appeals to their target consumers or can save on investment costs. We show that this option value of learning is maximized at an intermediate investment cost and is higher for firms that are more uncertain about future demand. This can explain why innovative products with most demand uncertainty are likely to benefit the most.

While there is value in learning, crowdfunding could be hampered by a well-known moral hazard problem (Tirole, 2006): firms may be tempted to divert the funds they raised instead of delivering the products and if anything, innovative firms may be particularly prone to moral hazard and informational frictions (Bussgang 2014, Lerner et. al. 2012). Further, reward-based crowdfunding platforms are not legally responsible for guaranteeing the delivery of rewards and proving that a firm has committed a fraud is difficult.¹ Yet, the vast majority of projects do deliver their rewards,

¹Kickstarter makes it clear that legal protection is limited and that the relationship relies primarily on interactions between the firm and backers "Backers must understand that Kickstarter is not a store. When you back a project, you're helping to create something new — not ordering something that already exists.", see Accountability part at

e.g., Mollick (2014) finds that only 3.6% of successful Kickstarter projects have failed to deliver the promised rewards.

We show that the real option value of learning is a powerful force that enables firms to endogenously overcome moral hazard. After its crowdfunding campaign, a firm that expects there to be high future consumer demand will choose to not divert funds, even if it is costless to do so. Because a firm that faces greater demand uncertainty has greater option value, it can also overcome moral hazard more easily. We allow the firm to set a target of funds to raise and solicit the platform to return the funds to backers if that target is not met. If the firm sets a low or zero target, it can invest and deliver the product in some states and divert funds in other states. We prove that the firm chooses an "All-or-Nothing" (AoN) crowdfunding scheme with a sufficiently high target so that it has incentives to invest after a successful campaign. AoN dominates "Keep-it-All" (KiA) where all funds raised are always passed on to the firm because the discount the firm needs to offer is too high compared to the expected benefit of diverting funds. We also discuss the robustness of our main findings to variable costs and uncertainty about the firm's ability to develop their product.

We further investigate the interplay between learning, moral hazard and the crowdfunding sample size, which can be proxied by the campaign length. We show that in the presence of moral hazard shorter campaigns are more likely to succeed and the expected funds raised and platform fees are maximized at an intermediate sample size. Under moral hazard, it is neither feasible nor optimal to pre-sell the product to all potential consumers. It then benefits the firm to have a third-party platform that ensures that the campaign targets a limited size sample. We derive the optimal sample size.

Our model can explain a number of stylized facts about reward-based crowdfunding discussed in Section 7. An important and potentially surprising stylized fact, consistent with our model, is that successful AoN crowdfunding campaigns systematically raise more funds than the target set at the beginning of the campaign. This pattern is present in all categories, but is most pronounced in the case of Technology projects. Figure 1 illustrates this with all Kickstarter projects in first three quarters on 2015. We can see that around 40% of successful projects raised at least twice the target, around 20% raised at least four times the target, etc. Appendix A.1 confirms this pattern with another sample period for all categories, illustrating that projects involving high demand uncertainty such as technology, design or gaming products are most "overfunded". This pattern cannot be explained if one views crowdfunding as merely a means to help the firm to overcome credit constraints.

<https://www.kickstarter.com/help/faq/kickstarter%20basics>.

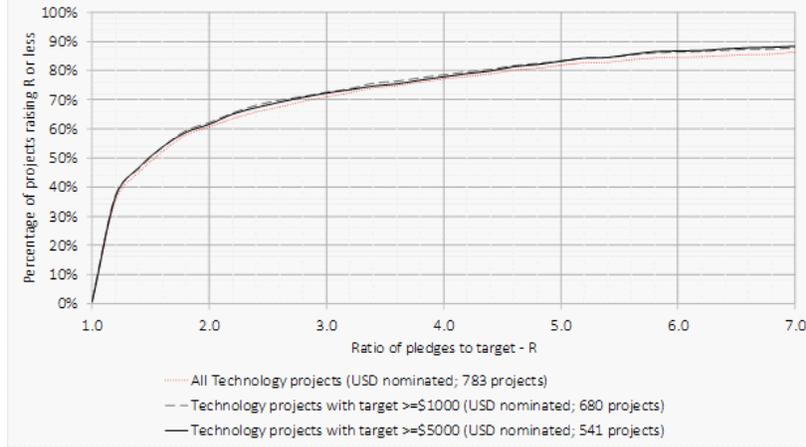


Figure 1: "Oversubscribed" technology projects. The figure includes all successful USD denominated technology projects on Kickstarter between 1 January, 2015 and 17 September, 2015.

Earlier theoretical explanations of reward-based crowdfunding have focused on backer preferences rather than learning and moral hazard. For example, Belleflamme, Lambert, and Schwienbacher (2014) assume that participation in crowdfunding provides backers with an additional utility compared to their valuation for the product, which enables firms to raise funds and to price-discriminate. Varian (2013) endogenizes this additional utility by deriving an equilibrium in which seemingly altruistic backer preferences are due to each of them having a pivotal role in ensuring that the firm has enough funds to invest and to produce the product that the backer values. Yet, these important consumer side effects alone cannot explain some important patterns of successful crowdfunding campaigns, such as products being sold at par or at a discount, and the fact that many products are oversubscribed multiple times over the target.

The argument that pre-selling enables firms to learn before investing also features in a few contemporaneous theoretical papers on reward based crowdfunding. Ellman and Hurkens (2016) focuses on price discrimination in a setting that rules out moral hazard by assuming high reputation costs. Our model focuses on the real option value of learning. It also suggests tests for the presence and magnitude of moral hazard by deriving empirical patterns that should only be observed under moral hazard, some of which have indeed been documented. Additional differences with that paper include our analysis of the effects of the degree of uncertainty about consumer preferences and our comparison of AoN and KiA schemes.

Strausz (2017) does consider moral hazard and argues that it needs to be mitigated by deferred payments and by conditional pledging where backers stop making contributions after the firm meets its target. He further argues that these features are realistic provided "projects that succeed tend to do so by relatively small margins". While this view may not be far from the truth for artistic

categories like "Dance" or "Theatre", it cannot explain the systematic "overpledging" in categories that raise most funds such as technology, games and design (see Figure 1 and Appendix A.1). The key difference between our two models is that Strausz (2017) assumes that the distribution of consumer preferences is known, which implies that learning from a sample of backers does not reveal information about consumer preferences out of sample. In our model, we show that in the limit case where the distribution of consumer preferences is known, the benefit of learning is significantly reduced and firms would choose to not participate in crowdfunding in the absence of financial constraints. Further, in contrast to Strausz (2017) the existing crowdfunding schemes are very close to first best efficient in our setting.

Chang (2016) considers a common value setting in which a firm aims to fund its project via crowdfunding. In his model, backers act as a group that decides whether to invest based on a noisy signal about the value of the project rather than their private valuation for the product. While the common value assumption may be better suited for financial rather than reward-based crowdfunding, his analysis also highlights the benefit of AoN compared to KiA and the benefit of learning in mitigating moral hazard. In contrast to our paper the firm will only use crowdfunding to cover the difference between the project cost and other sources of funding, and his paper does not consider the real option value of learning about consumer valuations.

Our paper differs from these three theoretical contributions in several other respects. We analyze the relationship between the size of the crowdfunding sample and crowdfunding outcomes. Our model explains why firms selling innovative products benefit more from crowdfunding than sellers of other products. We derive empirical predictions consistent with existing findings and that can be used to test our model against alternative theoretical models discussed above.

Our paper also contributes to the literature that points out that investing in entrepreneurial projects enables firms to experiment new technologies (Hellmann 2002, Gromb and Scharfstein 2005, Bettignies and Chemla 2008, Kerr, Nanda, and Rhodes-Kropf 2014). We show that crowdfunding is an efficient mechanism to learn about demand without experimentation costs. Our paper also relates to the industrial organization literature on pre-selling (Tirole, 1988, Rob, 1991, Crawford and Shum, 2005, Chu and Zhang, 2011), which primarily focuses on price-discrimination. While these papers consider the pre-selling of existing products, ours is about pre-selling at an earlier stage of the product cycle.

Our paper can also be related to the strands of corporate finance and monetary economics that view asymmetric information and moral hazard as major sources of financial constraints (Myers 1977, Stiglitz and Weiss 1981, Hart 1995, Tirole 2006, Kiyotaki and Moore 2002). The cam-

campaign itself may either provide actual funding or alleviate the root causes of financial constraints. For example, if financial constraints are driven by asymmetric information about demand, then the crowdfunding campaign that generates public information about demand alleviates these constraints.² In our setting crowdfunding is complementary to outside financing. There are indeed reported cases where firms, after succeeding in reward-based crowdfunding, obtain further funding from angels, venture capitalists or investor-based crowdfunding.³ Relatedly, investor-based crowdfunding highlights the benefits of learning public information from the "crowd" in terms of screening creditworthiness (see Iyer et al 2015) and the signaling of loan quality on lending platforms (Hildendbrand, Puri, and Rocholl 2016).

2 The Model

We consider a three-date model in which a firm has a new product idea, and can learn about demand after observing consumer decisions at date 0. At date 1, the firm updates its beliefs and decides whether or not to invest $I \geq 0$. At date 2 it produces and sets price p_2 at which it sells the product to consumers. For now, the firm's marginal cost of production is zero. All agents are risk neutral and the discount rate is normalized to 1. We do not impose exogenous financial constraints.

The firm's potential market consists of N consumers and the firm can sell at most 1 unit of the product to each of them. Each consumer $i \in \{1, \dots, N\}$ has private valuation $v_i = \{0, 1\}$ for one unit of the product and 0 for any additional unit. We refer to a consumer with valuation $v_i = 1$ (resp. $v_i = 0$) as a 1-consumer (resp. a 0-consumer). Private valuations are *conditionally* i.i.d, which implies that consumer i 's valuation is a Bernoulli trial drawn from the true distribution, i.e., $\Pr(v_i = 1|\theta) = \theta$. The probability θ that consumer i is a 1-consumer (and the aggregate share of 1-consumers) is unknown to the firm and it follows a beta distribution, the p.d.f. of which is

$$f_{\theta}(\theta) = \frac{\theta^{\alpha-1} (1-\theta)^{\beta-1}}{B(\alpha, \beta)}, \quad (1)$$

where α, β are positive parameters and $B(\alpha, \beta)$ is the beta function. Beta distributions enable us to capture different distributions of prior beliefs, be they U-shaped, hump-shaped, or uniform.⁴ For

²Further, publicly disclosing information about consumer preferences can alleviate information asymmetry in an unbiased manner. In contrast, information is revealed publicly through the actions of privately informed entrepreneurs may distort both firm value and investment decisions. For example, Myers and Majluf (1984) argue that equity issuance announcements convey negative information to investors, while Tinn (2010) and Angeletos, Lorenzoni and Pavan (2010) show that technology investments can be perceived as a positive public signal.

³Brunstein, Joshua, 2014, "How Kickstarter turned into the Venture Capitalist's Best Friend?", Bloomberg Business, August 11, <http://www.bloomberg.com/bw/articles/2014-08-11/kickstarter-successes-pivot-from-crowdfunding-to-venture-capital>

⁴If parameters $\alpha, \beta \geq 1$ are positive integers, the beta function is given by $B(x, y) = \frac{(x-1)!(y-1)!}{(x+y-1)!}$. More generally,

the sake of clarity, we write $\alpha = \theta_0 \lambda$ and $\beta = (1 - \theta_0) \lambda$, where $\theta_0 \in [0, 1]$ and $\lambda > 0$, such that

$$\mathbb{E}[\theta] = \frac{\alpha}{\alpha + \beta} = \theta_0; \text{Var}[\theta] = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)} = \frac{\theta_0(1 - \theta_0)}{(\lambda + 1)}.$$

That is, θ_0 is the prior mean and a higher λ implies a lower level of uncertainty based on prior beliefs. All agents know the prior distribution. Consumer i knows his own valuation for the product.

It is important to highlight why this information structure captures a realistic feature of learning where observing the preferences of some consumers reveals information about the preferences of other consumers. If the firm observes that one consumer values the product highly, it Bayesian updates its beliefs about the valuation of all other consumers: $\Pr(v_i = 1 | v_j = 1) = \frac{\lambda\theta_0 + 1}{\lambda + 1} > \theta_0$. The less uncertain the firm is about its target consumer preferences the less learning there is. At the limit, when $\lambda \rightarrow \infty$, the consumer preference distribution becomes public information: there is still uncertainty about demand by each consumer, but not about aggregate demand provided that the target market is sufficiently large to ensure that the law of large numbers applies. We will show that the limit case with $\lambda \rightarrow \infty$ is very special and eliminates the real option value created by crowdfunding.

As a benchmark we consider a frictionless consumer survey where $M \leq N$ consumers truthfully and costlessly reveal their preferences at date 0. The firm makes an investment decision based on updated beliefs about consumer preferences at date 1, and provided it invests, it delivers the product to consumers at date 2. Such frictionless consumer survey is unlikely to be implementable in practice because consumers face zero cost when overstating their interest in the product. In contrast, reward-based crowdfunding enables the firm to be sure that consumers who make pledges are not 0-consumers as these consumers would lose money by pre-ordering products that they do not value. In the basic model we keep M fixed and compare the benchmark setting and the crowdfunding model. Section 6 further analyzes the relationship between sample size M and the expected crowdfunding outcomes.

In the main model, the players are the firm and N potential consumers of the firm's product. If the product is produced, then all consumers can buy it at date 2, while a random subset of these consumers are potential backers who could pre-order the product through crowdfunding at date 0. Because pre-ordering takes place before the firm decides whether or not to invest, it involves a potential risk that the product will not be produced and the consumer loses his contribution. Without loss of generality, we index consumers who can participate in crowdfunding as $i = \{1, \dots, M\}$. Let us denote $0 \leq m \leq M$ the number of 1-consumers in set $\{1, \dots, M\}$, and

the beta function can be expressed as $B(x, y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$, where $\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}dt$ is the gamma function. For example, the prior distribution is uniform when $\alpha = \beta = 1$, and hump-shaped if $\alpha, \beta > 1$.

$0 \leq m_B \leq M$ the number of consumers in set $\{1, \dots, M\}$ who choose to pre-order the product.

In the baseline model the crowdfunding platform is a passive interface that, first, enables the firm to post the project in the morning of date 0, second, collects pledges during the campaign which takes place during daytime at date 0, and, third, passes on the funds collected to the firm in the evening of date 0. The platform also enables the firm to commit to a funding target \bar{m}_T , which means that any funds raised during the campaign get automatically returned to the consumers if the firm does not meet the target. We denote the event that the firm meets its target with an indicator function

$$T_{\bar{m}_T} := \begin{cases} 1 & \text{if } m_B \geq \bar{m}_T \\ 0 & \text{if } m_B < \bar{m}_T \end{cases} .$$

The target choice is also, de facto, a choice between AoN and KiA, as the firm can set any finite target \bar{m}_T under AoN, and by construction KiA means that $\bar{m}_T = 0$.⁵

The timing of events and the actions that players take are the following:

Morning of date 0: The firm sets the pre-selling price p_0 and the target \bar{m}_T .

Daytime of date 0: Each consumer $i = \{1, \dots, M\}$ chooses whether or not to pledge/pre-order the product. We denote the pledging decision with an indicator function

$$\mathbf{1}_0^i := \begin{cases} 1 & \text{if consumer } i \text{ chooses to pledge} \\ 0 & \text{otherwise.} \end{cases}$$

Date 1: The firm chooses whether or not to invest. We denote the firm's decision to invest with

$$\mathbf{1}_1^F := \begin{cases} 1 & \text{if the firm invests } I \\ 0 & \text{otherwise.} \end{cases}$$

Date 2: If the firm invested at date 1, it sets a price p_2 for the product. Consumers choose whether to buy the product. Denote the choice of consumer $i = \{1, \dots, N\}$ at date 2 with $\mathbf{1}_2^i = 1$ if the consumer buys the product, and 0 otherwise.

Given these actions the payoff to a firm that chooses to participate in crowdfunding is

$$\pi_F = T_{\bar{m}_T} \cdot p_0 m_B + \mathbf{1}_1^F \cdot \left(p_2 \sum_{i=0}^N \mathbf{1}_2^i - I \right) - \varphi(m_B, T_{\bar{m}_T}) - (1 - T_{\bar{m}_T}) \cdot \varsigma \cdot \mathbf{1}_1^F, \quad (2)$$

where $\varphi(m_B, T_{\bar{m}_T})$ denotes the fees that the platform charges and ς is a cost of investing after a campaign that failed to meet its pre-set target. We assume that $\varsigma > 0$, but it can be arbitrarily small. This assumption guarantees that the firm strictly prefers to meet its crowdfunding target

⁵Even though under real world KiA, there are "targets", they have no economic meaning as all the funds raised get passed on to the firm regardless of whether or not the firm meets the "target".

to failing to meet it.⁶ The payoff function (2) captures the moral hazard problem: if the firm does not invest at date 1, i.e., $\mathbf{1}_1^F = 0$, it diverts $T_{\bar{m}_T} \cdot p_0 m_B \geq 0$ at no cost apart from the platform fees.

The payoff to a consumer who belongs to a set of potential backers, $i = \{1, \dots, M\}$, is

$$u_i = T_{\bar{m}_T} \cdot \mathbf{1}_0^i \cdot (\mathbf{1}_1^F \cdot v_i - p_0) + \mathbf{1}_F^i \cdot \mathbf{1}_2^i \cdot (v_i - p_2), \quad (3)$$

where the term $\mathbf{1}_0^i \cdot (\mathbf{1}_1^F \cdot v_i - p_0)$ means that pledging, $\mathbf{1}_0^i = 1$, comes with the risk of paying p_0 and not getting the product if the firm does not invest after the firm meets its target. The consumer can always choose to not pledge at date 0, i.e., he can set $\mathbf{1}_0^i = 0$, and wait until date 2 to decide whether or not to purchase the product. Whether there is a product to purchase at date 2 depends on $\mathbf{1}_F^i$. However, the consumer faces zero risk of no delivery at date 2. The payoff to consumer $i = \{M + 1, \dots, N\}$ who cannot participate in crowdfunding in date 0 currency units is

$$u_i = \mathbf{1}_F^i \cdot \mathbf{1}_2^i (v_i - p_2) \quad (4)$$

We seek to find a Perfect Bayesian Equilibrium (PBE) of this game. To this end, we need to specify the information sets the players have at different dates. Both the firm and consumers know the parameters of the model and the distribution of v_i . When choosing p_0 and \bar{m}_T in the morning of date 0, the firm only knows the prior distribution (1) and parameters of the model. During the daytime of date 0, consumer i 's information set is $\Omega_0^i = \{p_0, \bar{m}_T, v_i\}$. When deciding whether to invest at date 1 the firm has the information set $\Omega_1^F = \{p_0, \bar{m}_T, m_B\}$. At date 2 uncertainty has resolved and all agents know $\Omega_2 = \{p_0, \bar{m}_T, m_B, \mathbf{1}_1^F, \{v_i\}_{i \in \{1, \dots, N\}}\}$.

Denote $b^i(\Omega_0^i)$ the probability with which a consumer $i \in \{1, \dots, M\}$ expects to receive the product at date 2 if he pledges at date 0. His beliefs depend on \bar{m}_T and his own valuation v_i . On (and off) the equilibrium path, he may also need to form expectations about other consumer's strategies which depends on their valuation v_j . If so, his beliefs must be consistent with Bayes' rule, i.e., the consumer assigns $\Pr(v_j = 1 | v_i) = \frac{\lambda \theta_0 + v_i}{\lambda + 1}$ and takes other consumers' on path strategies as given. We denote $b^F(m | m_B)$ the firm's beliefs about the total number of 1-consumers in set $i \in \{1, \dots, M\}$ conditional on consumers who backed the project.

Denote the equilibrium quantities with "*" and alternative strategies without "*". To shorten the notation, we express the players' expected payoffs as a function of their actions only. In any

⁶The technical reason behind this assumption is to avoid the following potentially profitable deviating strategy: the firm sets an infinite target, perfectly learns consumer preferences and proceeds without the platform ever receiving any fee. Such strategy would make it impossible for crowdfunding platforms to break even, render them irrelevant and make the interaction no different from the frictionless consumer survey that we argued is not to be an incentive compatible mechanism to learn consumer preferences. There are many realistic reasons that could justify this cost: failing to meet the preset target could be a negative signal about the quality of the project, some initially interested consumers may lose their interest, and it may take longer to find alternative sources of external finance.

PBE where the firm chooses to participate in crowdfunding, its equilibrium pricing and target $\{p_0^*, \bar{m}_T^*\}$ choice must satisfy

$$\mathbb{E} [\pi_F (p_0^*, \bar{m}_T^*, \mathbf{1}_1^{F*}, p_2^*)] \geq \mathbb{E} [\pi_F (p_0, \bar{m}_T, \mathbf{1}_1^{F*}, p_2^*)], \quad (5)$$

its equilibrium investment decision, $\mathbf{1}_1^{F*}$, must satisfy

$$\mathbb{E} [\pi_F (p_0^*, \bar{m}_T^*, \mathbf{1}_1^{F*}, p_2^*) | \Omega_1^F] \geq \mathbb{E} [\pi_F (p_0^*, \bar{m}_T^*, \mathbf{1}_1^F, p_2^*) | \Omega_1^F] \quad (6)$$

and its date 2 pricing choice, p_2^* , must satisfy.

$$\mathbb{E} [\pi_F (p_0^*, \bar{m}_T^*, \mathbf{1}_1^{F*}, p_2^*) | \Omega_2] \geq \mathbb{E} [\pi_F (p_0^*, \bar{m}_T^*, \mathbf{1}_1^{F*}, p_2) | \Omega_2] \quad (7)$$

The pledging strategy of consumer $i \in \{1, \dots, M\}$, $\mathbf{1}_0^{i*}$, must satisfy

$$\mathbb{E} [u_i (\mathbf{1}_0^{i*}, \mathbf{1}_2^{i*}) | \Omega_0^i] \geq \mathbb{E} [u_i (\mathbf{1}_0^i, \mathbf{1}_2^{i*}) | \Omega_0^i] \quad (8)$$

and the date 2 purchasing strategy of consumer $i \in \{1, \dots, N\}$, $\mathbf{1}_2^{i*}$, must satisfy

$$\mathbb{E} [u_i (\mathbf{1}_0^{i*}, \mathbf{1}_2^{i*}) | \Omega_2] \geq \mathbb{E} [u_i (\mathbf{1}_0^{i*}, \mathbf{1}_2^i) | \Omega_2] \quad \text{if } i \in \{1, \dots, M\} \quad (9)$$

$$\mathbb{E} [u_i (\mathbf{1}_2^{i*}) | \Omega_2] \geq \mathbb{E} [u_i (\mathbf{1}_2^i) | \Omega_2] \quad \text{if } i \in \{M + 1, \dots, N\}.$$

Because a consumer will not buy the product twice, we must have

$$T_{\bar{m}_T} \cdot \mathbf{1}_0^{i*} + \mathbf{1}_2^{i*} \leq 1 \quad \text{for any } i \in \{1, \dots, M\}.$$

The firm may also choose not to participate in crowdfunding. It is easy to see that all 1-consumers then buy the product at date 2 as long as $0 < p_2 \leq 1$, and it is optimal for the firm to set $p_2 = 1$. This means that the NPV of the project is $-I + N\mathbb{E}[\theta] = -I + N\theta_0$, and the firm invests if, and only if, $I \leq N\theta_0$. We will refer to this benchmark firm as *the reference firm* and the expected value of the reference firm's project is

$$\mathbb{E} [\pi_{ref}^F] = \max [0, (N\theta_0 - I)].$$

This implies that the firm will participate in crowdfunding if, and only if its expected payoff from crowdfunding satisfies

$$U_F = \mathbb{E} [\pi_F (p_0^*, \bar{m}_T^*, \mathbf{1}_1^{F*}, p_2^* | \Omega_0)] - \mathbb{E} [\pi_{ref}^F] \geq 0.$$

We assume that expected fees

$$\mathbb{E} [\varphi (m_B, T_{\bar{m}_T})] = z, \quad (10)$$

where z is a cost of running the project on the platform, which may or may not include a mark-up.

In any PBE, the strategies and beliefs must satisfy both sequential rationality and the consistency of beliefs, i.e., at date 0 the firm anticipates that each consumer $i \in \{1, \dots, M\}$ will follow the equilibrium strategy given his type v_i , and each consumer $i \in \{1, \dots, M\}$ anticipates the firm's optimal investment decision given m_B . In our main setting we focus on deriving a PBE in pure strategies. Because all consumers of the same type are identical, we expect there to be a symmetric equilibrium. If the symmetric PBE is fully revealing, then consistency of beliefs requires that the firm's beliefs on the equilibrium path are *skeptical*, which means that observing m^B , the firms considers that all backers that did not pledge were 0-consumers, i.e., $b^F(m|m_B) = \Pr(m = m^B) = 1$. We consider the off-equilibrium beliefs to be the same.

3 Benchmark and the value of learning

3.1 Updated beliefs

Before analyzing the crowdfunding game described in Section 2, we consider the benchmark of the frictionless survey, where at date 0 the firm directly learns the preferences of consumer $i \in \{0, 1, \dots, M\}$. After the firm observes number of 1-consumers in the sample, m , it updates its expectations about the share of 1-consumers in the entire population, N . Since $v_i|\theta$ is a Bernoulli trial, $m|\theta$ follows the binomial distribution. Bayes' rule implies that the posterior distribution is also Beta, with $\theta|m \sim Be(\lambda\theta_0 + m, \lambda(1 - \theta_0) + M - m)$. Therefore, the posterior expectations are

$$\mathbb{E}[\theta|m] = \frac{\lambda\theta_0 + m}{\lambda + M}. \quad (11)$$

Updated expectations of the share of 1-consumers in the full population is monotonically increasing in m , e.g. when $\frac{m}{M} > \theta_0$ the firm learns that there are more 1-consumers in its target market than it expected based on its prior beliefs. We derive the distribution of m

$$\begin{aligned} q_m &\equiv \Pr(m) = \mathbb{E}[\Pr(m|\theta)] = \int_0^1 \Pr(m|\theta) f(\theta) d\theta \\ &= \binom{M}{m} \frac{B(\lambda\theta_0 + m, \lambda(1 - \theta_0) + M - m)}{B(\lambda\theta_0, \lambda(1 - \theta_0))}, \end{aligned} \quad (12)$$

which means that m is a beta-binomial random variable. The shape of the beta-binomial distribution replicates the shape of the underlying prior beta distribution.⁷ Unconditional central moments

⁷For example, when the prior is uniform, i.e., $\lambda = 2$ and $\theta_0 = \frac{1}{2}$, $\Pr(m) = \frac{1}{M+1}$, i.e. the distribution of m is discrete uniform. If the firm's prior beliefs are U-shaped (hump-shaped), then the distribution of m , is also U-shaped (hump-shaped).

of the beta-binomial variable can be written

$$\mathbb{E}[m] = M\theta_0 \text{ and } Var[m] = M(\lambda + M) Var[\theta] = \frac{M(\lambda + M)\theta_0(1 - \theta_0)}{(\lambda + 1)}. \quad (13)$$

The expected m is proportional to the prior mean of the share of 1-consumers, θ_0 , and for any $M > 1$, an increase in λ implies a decrease in the degree of uncertainty about both θ and m . In the limit case where the distribution of consumer preferences is known, i.e. $\lambda \rightarrow \infty$, the firm does not update its beliefs about consumer preferences after observing m : $\lim_{\lambda \rightarrow \infty} E[\theta|m] = \theta_0 = \theta$ and q_m converges to a binomial distribution with parameters M and θ_0 , which is precisely the distribution assumed in Strausz (2015) and Ellman and Hurkens' (2016) baseline model.⁸

3.2 Benchmark investment decision

The profit of a firm is given by $\pi_F = \mathbf{1}_1^F \left(p_2 \sum_{i=0}^N 1_2^i - I \right)$. At date 2, the firm can only sell its product to 1-consumers at price $p_2 > 0$. Each of these consumers then obtains a surplus $1 - p_2$, and at date 1. If the firm invests at date 1, its expected profit is $p_2(m + (N - M)\mathbb{E}[\theta|m]) - I$, where m consumers are known to value the product and $(N - M)\mathbb{E}[\theta|m]$ other future consumers are expected to be 1-consumers. It follows immediately that both the firm value and the joint surplus, i.e. the sum of the payoff to the firm and of the consumer surplus, are maximized when $p_2 = 1$. The firm extracts all consumer surplus and invests in all projects that are non-negative NPV based on updated beliefs.

Denoting $\lceil x \rceil$ the ceiling function, i.e. the nearest integer rounded up, we obtain the following proposition:

Proposition 1 *At date 1, the first best investment decision is as follows: If $I < I_0 \equiv \frac{\lambda\theta_0(N-M)}{\lambda+M}$, then the firm invests regardless of the realization of m . If $I \geq I_0$, the firm invests if, and only if,*

$$m \geq \tilde{m} \equiv \left\lceil \frac{\lambda + M}{\lambda + N} (I - I_0) \right\rceil. \quad (14)$$

When the firm invests, its NPV at date 1 is

$$D(m) = m + (N - M)\mathbb{E}[\theta|m] - I = \frac{\lambda + N}{\lambda + M}m + I_0 - I \quad (15)$$

Proof. Follows from (11), the NPV being non-negative at date 1 and m being an integer. ■

The first part of Proposition 1 highlights that if $N > M$, then the firm with an investment cost lower than I_0 does not benefit from learning about consumer preferences, as it would invest even if $m = 0$. Furthermore, if $N > M$, then $\lim_{\lambda \rightarrow \infty} I_0 \rightarrow \infty$, which implies that the consumer survey brings

⁸Additionally, Section 5.2 in Ellman and Hurkens (2016) considers θ , which can take two values.

no value to a firm that is not uncertain about the distribution of its target consumer preferences. The second part of Proposition 1 shows that as long as there is some uncertainty about consumer preferences, i.e., λ is finite, then learning is valuable.

3.3 The real option value of learning under the benchmark

At date 0, the value of learning is the expected value of a firm that has the opportunity to learn about demand minus the expected value of the reference firm. Recall that at date 1 the expected firm profit $D(m)$ is given by (15). Given (13) the unconditional expectation of the date 1 profit is $\mathbb{E}[D(m)] = N\theta_0 - I$. The expected value of learning about demand is

$$U_B^F = \begin{cases} U_F^{B,I} & \text{if } I < N\theta_0 \\ U_F^{B,NI} & \text{if } I \geq N\theta_0 \end{cases}, \text{ where} \quad (16)$$

$$U_F^{B,I} \equiv \mathbb{E}[D(m) | m \geq \tilde{m}] \Pr(m \geq \tilde{m}) - (N\theta_0 - I) = -\mathbb{E}[D(m) | m < \tilde{m}] \Pr(m < \tilde{m}) \quad (17)$$

$$U_F^{B,NI} \equiv \mathbb{E}[D(m) | m \geq \tilde{m}] \Pr(m \geq \tilde{m}), \quad (18)$$

and the superscripts "I" and "NI" denote whether or not the reference firm invests. Since $D(m) < (>) 0$ for any $m < (>) \tilde{m}$ and $D(m) \geq 0$ if $m = \tilde{m}$, both $U_F^{B,I}$ and $U_F^{B,NI}$ are positive. Further, $U_F^{B,I}$ and $U_F^{B,NI}$ represent both the value of learning for the firm and the joint surplus from learning for consumers and the firm. From (14), the value of learning is positive if, and only if, the firm's investment cost satisfies

$$I_0 < I < I_0 + M \frac{\lambda + N}{\lambda + M}.$$

This guarantees that the investment threshold, \tilde{m} , is above 0 and can be met with a positive probability, i.e., $\tilde{m} \in \{1, 2, \dots, M\}$.

Proposition 2 *The value of learning is maximized when $I = N\theta_0$. Further, $U_F^{B,I}$ is increasing in I and $U_F^{B,NI}$ is decreasing in I .*

Proof. See Online Appendix B.2. ■

Proposition 2 shows that a firm that expects to break even based on prior beliefs has most to gain from learning, while the overall relationship between I and the value of learning is hump-shaped. If $I < N\theta_0$, the firm benefits from avoiding a sub-optimal investment. This benefit increases with the investment cost that it expects to save. If $I > N\theta_0$, the firm benefits mostly because it can learn that investment is worth undertaking, and thus the higher I the lower the returns from the investment. The opportunity to learn about demand provides the firm with a real option.

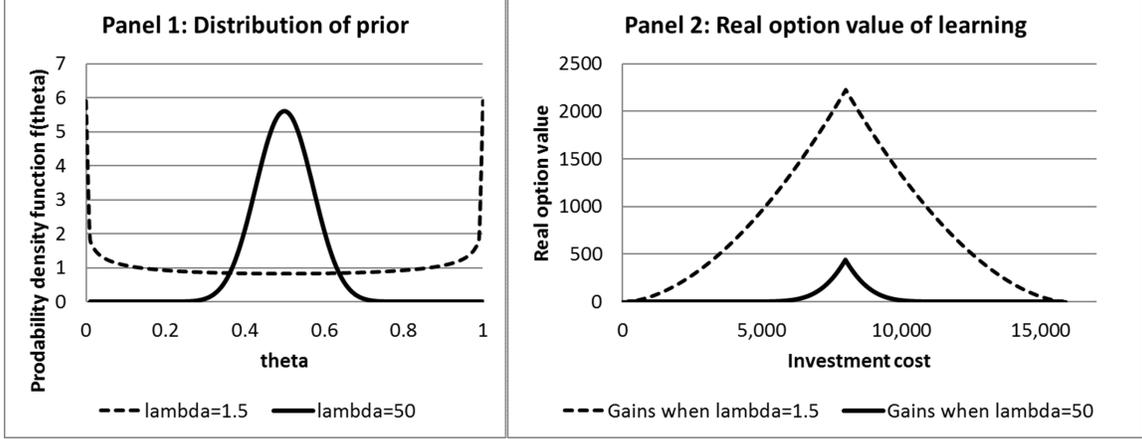


Figure 2: First-best value of learning

Proposition 3 *Ceteris paribus, the value of learning increases with the degree of prior uncertainty about demand, i.e., it decreases with λ .*

Proof. See Online Appendix B.3. ■

Proposition 3 shows that a firm that is more uncertain about the preferences of its target consumers, has more to gain. For example, a firm that expects its product to be valued either by many or by few consumers, has most to gain from learning. Since novel, creative consumer products (e.g., new technology gadgets) are more likely to be characterized by such belief structures than other types of products, we expect innovative firms to benefit most from learning.

Note that at date 0, λ affects the firm value through three channels. First, an increase in uncertainty increases the difference between the firm prior and updated beliefs about the share of 1-consumers, i.e., $|\mathbb{E}[\theta|m] - \mathbb{E}[\theta]| = \frac{|m - M\theta_0|}{M + \lambda}$. Second, from equation (14), higher uncertainty may affect the threshold, \tilde{m} , at which the firm finds it optimal to invest. Third, the distribution of m with a higher λ second order stochastically dominates the one with a lower λ . Overall, the second order stochastic dominance, which implies that the probability of extreme realizations of m decreases with λ , ensures that the effect of an increase in uncertainty is positive. The effect of λ on the first best investment threshold further enhances the benefit of learning when there is higher uncertainty.

Figure 2 illustrates these comparative statics by plotting the value of learning about demand as a function of λ and I . We consider two possible prior distributions with the same mean ($\theta_0 = \frac{1}{2}$) and different values of λ . We chose $\lambda_1 = 1.5 < \lambda_2 = 50$ so that the distribution with λ_1 is U-shaped and the one with λ_2 is hump-shaped, as illustrated on Panel 1. The figure assumes $\theta_0 = \frac{1}{2}$, $M = 1600$, and $M/N = 0.1$. Firms break even if $I = 8000$ for both $\lambda = \lambda_1$ and $\lambda = \lambda_2$. The value

of learning is positive if $I \in (6.7, 15993)$ for $\lambda = \lambda_1$, and if $I \in (218.2, 15570)$ for $\lambda = \lambda_2$. Panel 2 plots the value of learning about demand as a function of I for both λ_1 and λ_2 .

4 Reward-based crowdfunding

While a survey is unlikely to induce consumers to truthfully reveal their valuation, pre-ordering decisions could generate credible information as only 1-consumers would pre-order the product at $p_0 > 0$. This section explores the existence of a fully revealing PBE of the crowdfunding game specified in Section 2, and we proceed by backward induction.

4.1 Firm's investment decision.

It clearly remains optimal for the firm to extract all consumer surplus at date 2 and to set $p_2 = 1$. From this, beliefs $b(m|m_B) = \Pr(m = m_B) = 1$ and (6), the firm chooses to invest at date 1 if, and only if,

$$\begin{aligned} m_B &\geq \bar{m}^T \ \& \ p_0 m_B + (N - M) E[\theta|m_B] - \varphi(m_B, \bar{m}^T) - I \geq p_0 m_B - \varphi(m_B, \bar{m}^T) \quad \parallel \quad (19) \\ m_B &< \bar{m}^T \ \& \ m_B + (N - M) E[\theta|m_B] - \varphi(m_B, \bar{m}^T) - I - \varsigma \geq -\varphi(m_B, \bar{m}^T) \end{aligned}$$

We can see that the fee structure $\varphi(m_B, \bar{m}^T)$ does not affect the firm's investment incentives. If $m_B \geq \bar{m}^T$ there is a moral hazard problem because the firm has already received an amount $p_0 m^B$, which it can divert at no extra cost. The firm's incentives to invest in these states are preserved when expected future demand by the $(N - M)$ consumers that did not have the chance to participate in crowdfunding is sufficiently high, i.e., (19) implies that if $m_B \geq \bar{m}^T$, then the incentive compatibility constraint can be written

$$(N - M) E[\theta|m_B] - I = (N - M) \frac{\lambda \theta_0 + m_B}{\lambda + M} - I \geq 0. \quad (20)$$

Learning about future consumer preferences is essential here as the firm's expected future demand is proportional to $E[\theta|m^1] = E[\theta|m^B] = \frac{\lambda \theta_0 + m_B}{\lambda + M}$, which itself is increasing in m_B as long as λ is finite.

Whether or not the firm always collects funds is the crucial difference between AoN and KiA. Under KiA ($\bar{m}^T = 0$) the firm's investment decision is always distorted by the moral hazard problem. In contrast, under AoN, the firm may still choose to invest if its campaign fails, as $m_B < \bar{m}^T$ means that it will invest if, and only if

$$m_B + (N - M) E[\theta|m_B] - I - \varsigma \geq 0, \quad (21)$$

which is a less restrictive condition than (20). The date 1 investment decision is summarized in the following Lemma.

Lemma 4 *The firm invests if, and only if*

$$m^B \geq \begin{cases} \tilde{m}_Y & \text{if } m^B \geq \tilde{m}^T \\ \tilde{m}_N & \text{if } m^B < \tilde{m}^T \end{cases}, \quad (22)$$

where

$$\tilde{m}_Y \equiv \left\lceil \frac{\lambda + M}{(N - M)} (I - I_0) \right\rceil = \left\lceil \frac{\lambda + M}{(N - M)} \cdot I - \lambda \theta_0 \right\rceil. \quad (23)$$

and

$$\tilde{m}_N = \left\lceil \frac{\lambda + M}{\lambda + N} (I + \varsigma - I_0) \right\rceil \quad (24)$$

Provided ς is small⁹, $\tilde{m}_Y \geq \tilde{m}_N \geq \tilde{m}$

Proof. Follows from (20) and (21), $E[\theta|m_B] = \frac{\lambda\theta_0 + m_B}{\lambda + M}$ and the constraint that any threshold needs to be a integer. The claim $\tilde{m}_Y \geq \tilde{m}_N \geq \tilde{m}$ follows from the assumption that ς is arbitrarily small and from (14), (23), (24). ■

Lemma 4 shows that if the firm meets its target \tilde{m}^T it invests if $m \geq \tilde{m}_Y$. The threshold \tilde{m}_Y is (weakly) higher than the optimal investment threshold. Yet, this does not necessarily mean suboptimal investment decisions. A firm that fails to meet its target will invest if the NPV based on posterior beliefs is positive. That firm received no funding and does not face moral hazard, its investment threshold is (weakly) lower than \tilde{m}_Y and converges to the first best when $\varsigma \rightarrow 0$. Furthermore, the threshold \tilde{m}_Y is significantly lower than it would be in a setting where the crowdfunding sample does not reveal information about consumer preferences out of sample (as, e.g., in Strausz, 2017): notice that when $\lambda \rightarrow \infty$, then the threshold $\tilde{m}_Y \rightarrow \infty$ with $I > I_0$, while an unconstrained firm with an investment cost $I \leq I_0$ would not participate in crowdfunding as it does not benefit from the option value of learning (Section 3.2).

Lemma 4 also shows that if $M \rightarrow N$ raising funds via crowdfunding becomes impossible, which suggests that limiting the campaign length is essential to ensure that crowdfunding is feasible.

4.2 Backer pledging decisions

From 3, the expected utility to consumer $i = \{1, \dots, M\}$ is

$$\mathbb{E}[u_i|\Omega_0^i] = \mathbb{E}[T_{\tilde{m}_T} \cdot \mathbf{1}_0^i \cdot (\mathbf{1}_1^F \cdot v_i - p_0) | \Omega_0^i] + \mathbb{E}[(1 - \mathbf{1}_0^i) \cdot \mathbf{1}_F^i \cdot \mathbf{1}_2^i \cdot (v_i - p_2) | \Omega_0^i].$$

⁹In particular, assuming ς arbitrarily small ensures here that $\varsigma \leq \frac{\lambda + M}{N - M} (I - I_0)$.

Since $p_2 = 1$ and $\mathbf{1}_2^i = 0$ for any 0-consumer, it is clear that the second term is zero and we obtain that consumer $i = \{1, \dots, M\}$ sets $\mathbf{1}_0^i = 1$ and this condition becomes $\mathbb{E} [T_{\bar{m}_T} (\mathbf{1}_1^F \cdot v_i - p_0) | \Omega_0^i] \geq 0$. Using the law of iterated expectations, it then follows that consumer i pledges if, and only if,

$$\Pr (T_{\bar{m}_T} = 1 \& \mathbf{1}_1^F = 1 | \Omega_0^i) v_i - p_0 \Pr (T_{\bar{m}_T} = 1 | \Omega_0^i) \geq 0 \quad (25)$$

As consumer i must have consistent beliefs, he correctly anticipates that if the firm meets its target, i.e., $T_{\bar{m}_T} = 1$, then it sets $\mathbf{1}_1^F = 1$ whenever $m^B \geq \tilde{m}_Y$ as shown in Lemma 4. As the target \bar{m}_T is known to consumer i , the consistency of beliefs and Bayes's rule require that consumer i 's beliefs satisfy

$$b^i (\Omega_0^i) = \Pr (T_{\bar{m}_T} = 1 \& \mathbf{1}_1^F = 1 | \Omega_0^i) = \begin{cases} \Pr (T_{\bar{m}_T} = 1 | \Omega_0^i) & \text{if } \bar{m}_T \geq \tilde{m}_Y \\ \Pr (T_{\bar{m}_T} = 1 | \Omega_0^i) \Pr (\mathbf{1}_1^F = 1 | \Omega_0^i, T_{\bar{m}_T} = 1) & \text{if } \bar{m}_T < \tilde{m}_Y \end{cases}$$

In a fully revealing equilibrium, the consumer expects all other 1-consumers to pledge and all 0-consumers not to pledge. We denote the total pledges of all other consumers except i , with $m_{-i} = \sum_{j=1, j \neq i}^M v_j$. Consumer i expects the total pledges to be $m = m^B = m_{-i} + \mathbf{1}_0^i$. Bayes' rule implies that conditional on his own type, i 's beliefs about θ are $\theta | v_i \sim Be(\lambda\theta_0 + v_i, \lambda(1 - \theta_0) + 1 - v_i)$ and his beliefs about other consumer pledges, $m_{-i} | v_i$, follow a beta-binomial distribution with parameters $M - 1$, $\lambda\theta_0 + v_i$ and $\lambda(1 - \theta_0) + 1 - v_i$.

Lemma 5 *Each consumer $i \in \{1, \dots, M\}$ with $v_i = 0$ chooses $\mathbf{1}_0^i = 0$*

Each consumer $i \in \{1, \dots, M\}$ with $v_i = 1$ sets $\mathbf{1}_0^i = 1$ either if

$$\bar{m}_T \geq \tilde{m}_Y \text{ and } p_0 \leq 1 \quad (26)$$

or if

$$\bar{m}_T < \tilde{m}_Y \text{ and } p_0 \leq \frac{\Pr (m_{-i} \geq \tilde{m}_Y - 1 | v_i = 1)}{\Pr (m_{-i} \geq \bar{m}_T - 1 | v_i = 1)} \leq 1 \quad (27)$$

and sets $\mathbf{1}_0^i = 0$ otherwise.

Proof. From (25) it is immediate that when $v_i = 0$, pledging leads to a strictly negative expected payoff, which proves the first sentence of Lemma 5. When $\bar{m}_T \geq \tilde{m}_Y$, the consumer anticipates the firm will invest with probability 1 if it meets its target, which proves that $\mathbf{1}_0^i = 1$ whenever (26) holds. When $\bar{m}_T < \tilde{m}_Y$, then $\Pr (T_{\bar{m}_T} = 1 | \Omega_0^i) = \Pr (m^B \geq \bar{m}_T | v_i = 1) = \Pr (m_{-i} + 1 \geq \bar{m}_T | v_i = 1)$ and $\Pr (T_{\bar{m}_T} = 1 | \Omega_0^i) \Pr (\mathbf{1}_1^F = 1 | \Omega_0^i, T_{\bar{m}_T} = 1) = \Pr (m^B \geq \tilde{m}_Y | v_i = 1) = \Pr (m_{-i} \geq \tilde{m}_Y - 1 | v_i = 1)$, using (25), we then find 27. The last inequality follows from the fact that $\tilde{m}_Y - 1 > \bar{m}_T - 1 \iff \Pr (m_{-i} \geq \tilde{m}_Y - 1 | v_i = 1) \leq \Pr (m_{-i} \geq \bar{m}_T - 1 | v_i = 1)$. ■

Lemma 5 shows that the full revelation of preferences is indeed plausible in equilibrium. If the firm has set a target that is sufficiently high to ensure that the firm has incentives to invest at date 1, then each consumer is insured against the risk of no investment. This is why he is willing to make a pledge as long as the pre-selling price is not above his private valuation for the product. This case requires the firm to set a positive target as $\bar{m}_T \geq \tilde{m}_Y > 0$, which is only possible with an AoN crowdfunding campaign. If the firm has set a target that is lower than \tilde{m}_Y , the consumer is still willing to participate as long as he can pre-order the product at a large enough discount that compensates for the risk of the firm failing to invest. Notice that under a KiA scheme, we would always have $\bar{m}_T = 0 < \tilde{m}_Y$. For most parameters of the model, the last inequality of condition 27 is strict.

4.3 The fully revealing equilibrium

The firm faces a trade-off between target and discount. On the one hand, the firm can set a target $\bar{m}_T \geq \tilde{m}_Y$ and pre-sell the product at a minimal discount, $p_0 = 1$. On the other hand, it can set a lower target, sell the product at a larger discount and divert some funds raised following some outcomes of its crowdfunding campaign. It turns out that the discount that the firm needs to offer is too large to outweigh the benefits of diverting funds.

Lemma 6 *The firm sets the price to $p_0 = 1$ and the target to $\bar{m}_T = \tilde{m}_Y$. The date 0 expected value of participation in AoN crowdfunding is*

$$U_F = U_F^B - z - \mathbb{E}[D(m) | \tilde{m} < m \leq \tilde{m}_N] \Pr(\tilde{m} < m \leq \tilde{m}_N) - \Pr(\tilde{m}_N \leq m < \bar{m}_Y) \varsigma \quad (28)$$

where U_F^B is the first best utility defined in (16) and \tilde{m}_N and \bar{m}_Y are defined in (23) and (24). When $\varsigma \rightarrow 0$ the investment target after a failed campaign is the same as the optimal target \tilde{m} defined in (14) and the gains from crowdfunding are

$$\lim_{\varsigma \rightarrow 0} U_F = U_F^B - z \quad (29)$$

Proof. See Appendix A.2. ■

Lemma 6 shows that the firm's utility from crowdfunding is maximized when it chooses AoN and sets a target high enough not to have incentives to divert funds after a successful campaign. Furthermore, the firm's utility from crowdfunding is (almost) as high as under the benchmark case. The crowdfunding target $\bar{m}_T = \tilde{m}_Y$, has no direct effect on the firm's utility: the firm also benefits from learning when it fails to meet the target of its crowdfunding campaign. The firm's utility under crowdfunding only differs from that under the benchmark of a frictionless consumer survey because of the platform fees and a possible cost of investment under a failed campaign, ς .

Proposition 7 *There is a fully revealing Perfect Bayesian Equilibrium of the crowdfunding game, where*

- 1) *all 1-consumers in set $\{1, \dots, M\}$ pledge at date 0 and purchase the product at date 2 if the firm fails to meet its crowdfunding target and invests. All 1-consumers in set $\{M + 1, \dots, N\}$ purchase the product at date 2 if the firm invests. 0-consumers neither pledge nor purchase the product;*
- 2) *The firm with $U_F \geq 0$ participates in crowdfunding. A firm that participates chooses an AoN campaign, sets a crowdfunding target $\bar{m}_T = \tilde{m}_Y$, prices $\{p_0, p_2\} = \{1, 1\}$ and invests as long as $m^B \geq \tilde{m}_N$.*

Proof. This follows from Lemmas 4, 5 and 6. ■

The investment costs that satisfy $U_F \geq 0$ must be a proper subset of $(I_0, I_0 - \varsigma + M \frac{\lambda+N}{\lambda+M})$: a firm that participates must have a positive value of learning and must have an investment threshold that it can meet with a positive probability.

5 Robustness and extensions of the main model

Other equilibria While Section 4 showed that there is a fully revealing PBE in pure strategies, there are other PBE of the game described in Section 2 under different player beliefs. An obvious one is a non-informative equilibrium where the firm expects none of the 1-consumers to pledge, and because consumers are indifferent between pledging or not, they may as well not pledge at date 0. The PBE under such beliefs is one where the firm chooses not to participate in crowdfunding: without learning, the firm would make exactly the same investment choice as the *reference firm* and the cost of participating in crowdfunding, $z > 0$, makes the firm strictly better off not participating. Because consumers are indifferent about pledging, there are also mixed strategy equilibria, e.g., one where each 1-consumer in sample M pledges with probability $0 < \rho < 1$. Consistency of beliefs then requires that the firm believes that $m = \frac{m^B}{\rho}$, which implies that it still learns the total number of 1-consumers in the sample and there is a qualitatively similar fully revealing PBE to the one in Section 4. Such equilibria are however not robust to a small modification in consumer preferences. Namely, if each consumer preferences are lexicographic, so that his primary preference is for the product itself as in the basic model (25), and his secondary preference is to "tell" the firm he likes the product, then consumers strictly prefer to pledge and only the fully revealing equilibrium in pure strategies described in Proposition 7 survives. Such preferences would also justify the assumption that a consumer survey as described in the benchmark would not lead to a truthful revelation of consumer preferences.

Communication In the main model, consumers only know their own preferences and have

beliefs about other consumer preferences and strategies when deciding whether or not to pledge. In reality, platforms such as Kickstarter and Indiegogo make it possible for consumers to see pledges that have been made on a running basis. However, each backer can make and withdraw his pledges at any time during the campaign, and he becomes committed to his pledge only at the end of the campaign when funds are taken from his account. One may view the interaction during the campaign as "cheap talk", viewing pledges during the campaign as "messages" followed by a pledging decision at the last moment of the campaign (see Farrell and Rabin, 1996). The equilibrium derived in our main model then corresponds to a "babbling" equilibrium of the messaging game where all potential backers disregard everyone else's messages. Clearly, the PBE described in Proposition 7 is also consistent with a fully revealing messaging game as consumers do not benefit from hiding their type. What can change as a result of communication are the backer strategies on an off-equilibrium path where the firm has set a target $\bar{m}_T < \tilde{m}_Y$ (see Lemma 5 in Section 4.2). Revealing communication enables consumers to coordinate and not back a project with $m < \tilde{m}_N$.¹⁰

Reputation costs. So far we have assumed that a firm can costlessly divert all funds raised if it chooses not to invest after meeting its crowdfunding target. In reality, failing to invest and to deliver the product to backers may come with a reputation cost or some litigation risk for the firm. Such reputation cost may easily be incorporated in our model, e.g., assume that if the firm fails to invest and deliver the product, it pays a cost $\chi(m_B) = \chi_Y + \kappa m_B$, where $\chi_Y, \kappa \geq 0$. The firm's incentive compatibility constraint (20) after a successful campaign, $m^B \geq \bar{m}_T$, then becomes

$$(N - M) \frac{\lambda \theta_0 + m_B}{\lambda + M} - I \geq \chi_Y + \kappa m_B.$$

Hence, the new investment threshold after a successful campaign becomes

$$\tilde{m}_Y = \left\lceil \frac{\lambda + M}{(N - M) + (\lambda + M) \kappa} (I_0 - \chi_Y - I) \right\rceil.$$

Notice that reputation costs reduce the threshold at which the firm invests. If χ_Y (or κ) is low enough, such that $\tilde{m}_Y > \tilde{m}$ as defined in (14), the equilibrium as in described in Proposition 7. When reputation costs are sufficiently high, such that $\tilde{m}_Y \leq \tilde{m}$, then the firm sets a target $\bar{m}_T = \tilde{m}$, makes a first best investment decision, and achieves the same utility as in the benchmark model (minus platform fees z). This extension is primarily relevant to empirical implications, as it allows us to propose a way to test for the presence and magnitude of the moral hazard problem (Section

¹⁰Another way to see the communication during the campaign is in the spirit of Okuno-Fujiwara et. al. (1990) where the platform can certify some messages (e.g., "I am a type 1-consumer, because the platform can take funds from the backers' account and 0-consumers would not pledge) and not others (e.g., "I am a type 0-consumer", which is not distinguishable from a consumer who chooses not to reveal his type). Again we can expect that full information revelation is an equilibrium and it leads to the same outcome as Proposition 7.

7). This extension also shows that prominent firms that have high reputation costs can commit to return funds to their backers directly and hence pre-sell new products via their own website, while less well-known firms benefit from a third-party platform.

Variable costs In the main model we considered a zero variable cost. Suppose instead that a fixed cost I_F is paid at date 1 and a variable cost $I_V \in (0, 1)$ is paid at date 2. As in the main setting, the highest price at which the firm can pre-sell the product is $p_0 = 1$. Writing $I \equiv \frac{I_F}{(1-I_V)}$, the date 1 expected profit is $(1 - I_V) D(m)$, where $D(m)$ is as in (15). The first best investment threshold in terms of I is the same as the one derived in Section 3 and related comparative statics remain valid. Since I is increasing in I_V , an increase in the variable cost increases the optimal investment threshold at a given I_F .

When diverting funds, the firm can now also avoid the variable cost. The firm now finds it optimal to invest after a successful campaign, if

$$m \geq \tilde{m}_Y = \left[\left(\frac{N - M}{\lambda + M} - \frac{I_V}{1 - I_V} \right)^{-1} \left(I - \frac{I_0}{(1 - I_V)} \right) \right],$$

which is higher than in the main setting. Intuitively, the variable cost increases the severity of moral hazard. This implies that firms with high variable costs find it more difficult to succeed in crowdfunding. All the other results in the preceding sections remain.

Uncertainty So far we have assumed that whenever the firm invests, it will successfully produce the product. Suppose instead that the investment itself only succeeds with probability $\gamma \in (0, 1]$. Hence, the firm's date 1 profit is

$$\frac{\lambda + N}{\lambda + M} p_0 m + \gamma I_0 - I.$$

In addition, the consumers who pre-order the product know that they will receive it with probability γ only. Hence, they are willing to buy if, and only, if $p_0 \leq \gamma$. The firm has to offer a discount in order to convince consumers to participate and sets $p_0 = \gamma$. All our qualitative results remain valid and the targets are scaled up with γ as follows:

$$\tilde{m}_Y = \left[\frac{\lambda + M}{N - M} \left(\frac{I}{\gamma} - I_0 \right) \right] \text{ and } \tilde{m}_N = \left[\frac{\lambda + M}{\lambda + N} \left(\frac{I + \varsigma}{\gamma} - I_0 \right) \right]$$

Overall, both uncertainty about the firm's ability to deliver the product and variable costs imply higher investment thresholds, and make successful crowdfunding more difficult. The intuition for these effects is simple, as both imply lower expected payoffs to the firm. In addition, uncertainty about the firm's ability to deliver the product implies that the firm must offer consumers larger discounts at the pre-selling stage, and variable costs worsen the moral hazard problem.

6 Role of platforms and campaign length

Analysis in Sections 2 to 5 indicates why third party platforms are needed for firms with no (or a small) reputation cost: these platforms are beneficial for the firm as they facilitate its commitment to target \bar{m}^T and can ensure that $M < N$. If a firm crowdfunds a project via its own website, it may be tempted to keep the campaign running with the purpose of diverting funds should the product fail to attract enough consumer interest, and may be unable to credibly commit to return funds to backers if the target is not met. In this Section we first explore how the size of the crowdfunding sample, M , affects the probability of crowdfunding success, expected fee income, and the firm's expected utility taking as given the current AoN crowdfunding rules. We then discuss the efficiency of current schemes.

6.1 Sample size

Even though our main setting allows for any fee structure, the real world AoN crowdfunding sites such as Kickstarter charge fees proportional to the funds raised only if the campaign is successful.

We examine the impact of sample size on our results in an environment with zero reputation cost and $\varsigma \approx 0$. Then the crowdfunding target is $\bar{m}^T = \tilde{m}_Y \geq \tilde{m}$ and the investment target is $\tilde{m}_N \approx \tilde{m}$. Our variables of interest are therefore the probability of success of a crowdfunding campaign, $\Pr(m \geq \tilde{m}_Y)$, the platform's fee income, which is proportional to $\mathbb{E}[m|m \geq \tilde{m}_Y] \Pr(m \geq \tilde{m}_Y)$, and the firm's utility, which depends on \tilde{m} and is given by equations (29), (16) and (15).

We start by plotting an example of these variables by considering the example of a firm with a break-even investment cost $I = N\theta_0$. Recall that such a firm has most to gain from learning about demand and therefore represents a typical crowdfunding participant. We consider the same parameters as on Figure 2 in Section 3 and we additionally consider a uniform prior (with $\theta_0 = 0.5$ and $\lambda = 2$). We also assume that the fees represent 5% of the funds raised as on Kickstarter.

Panel 1 of Figure 3 shows that the probability of crowdfunding success indeed tends to be higher as the sample size gets smaller, provided the sample is not too small. Panel 2 of Figure 3 shows that the expected revenue from fees is highest at an intermediate sample size. Intuitively, an increase in sample size increases the expected fee income conditional on the firm meeting the target, but reduces the probability that the firm meets its target. The figure also shows that an increase in uncertainty (a decrease in λ) leads to a slower decrease in the probability of success, an increase in expected fee income, and an increase in the sample size that maximizes fee income. Panel 3 of Figure 3 shows that the firm's utility is increasing in M , but the marginal benefit of a higher M quickly becomes negligible. This reflects the fact that a firm does not need a very large sample size

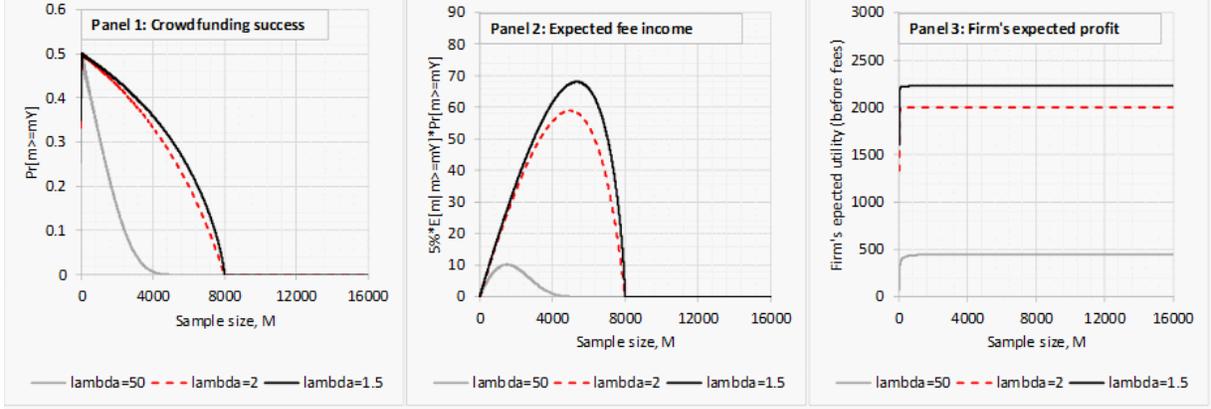


Figure 3: Probability of crowdfunding success, platform's expected fee income and firm's expected profit.

to make reasonably accurate inferences about the value of θ . Overall, Figure 3 suggest that if the crowdfunding platform charges fees that are proportional to the funds raised, it has an incentive to set an upper limit on M , and the firm does not benefit from shortening the campaign further.

We can highlight the main forces that drive these effects. Let us denote $q(m; M)$ and $q(m; M + 1)$ the probability mass functions (12) and $Q(m; M) = \sum_{k=0}^m q(k; M)$ and $Q(m; M + 1) = \sum_{k=0}^m q(k; M + 1)$ the cumulative distribution functions, under M and $M + 1$, respectively.

Lemma 8 For any $m = 0, \dots, M - 1$, we have $Q(m; M) > Q(m; M + 1)$.

Proof. See Online Appendix B.4. ■

Lemma (8) shows that the distribution $Q(m; M)$ first-order stochastically dominates $Q(m; M + 1)$. At the same time, notice that the term inside the ceiling function $\lceil \cdot \rceil$ that defines the target threshold $\tilde{m}_T = \tilde{m}_Y$ is increasing and convex in M . This implies that there are two opposite forces that affect the probability of crowdfunding success $\Pr(m \geq \tilde{m}_Y) = 1 - Q(\tilde{m}_Y - 1; M)$. First, as the sample size increases it becomes easier to meet a fixed target. Second, as the sample size increases the crowdfunding target \tilde{m}_Y itself increases and becomes harder to meet. A higher M requires the firm to set an increasingly ambitious target to guarantee that the firm has incentives to invest rather than to divert funds. Convexity ensures that the second effect becomes the dominating force when M increases. It is worth noticing that the decrease in the probability of success with M is a symptom of a moral hazard problem. If reputation costs were sufficiently high to eliminate moral hazard, then the crowdfunding target would be \tilde{m} (which is roughly linear in M), and the probability of crowdfunding success for a break-even firm would be roughly constant.

In order to identify the forces that drive the fee income, we can use Abel's Lemma to decompose

$$\mathbb{E}(m|m \geq \tilde{m}_Y) \Pr(m \geq \tilde{m}_Y) = M\theta_0 + \sum_{m=0}^{\tilde{m}_Y-1} Q(m) - \tilde{m}_Y Q(\tilde{m}_Y - 1).$$

We can see that increasing M affects fee income via three channels: 1) a direct positive effect via unconditional expectations, $M\theta_0$; 2) indirect effects via threshold \tilde{m}_Y , which tend to increase the term $\sum_{m=0}^{\tilde{m}_Y-1} Q(m)$ and decrease the term $-\tilde{m}_Y Q(\tilde{m}_Y - 1)$; 3) an indirect effect via a decrease in $Q(\cdot)$ as shown in Lemma 8.

We gain further intuition and comparative statics while we consider a continuous approximation of the model that assumes that M and N are large enough (due to the integer problem, the lines on Figure 3 have many small kinks). Instead of the number of backers, M , it is now better to focus on the share of backers among target consumers, $\mu \equiv \frac{M}{N}$. In a sufficiently large sample, the firm obtains an accurate and unbiased estimate of θ , i.e., $E[\theta|\mu M] \approx \theta$, and we can redefine the target and investment thresholds in terms of θ . The firm's date 1 expected profit is now $\pi_F \approx \mu\theta N + N(1 - \mu)\theta - I$, which implies that the firm's optimal investment decision is given by $\theta \geq \tilde{\theta} = \frac{I}{N}$, and that the firm sets a target $\tilde{\theta}_Y = \frac{I}{N(1-\mu)}$. The prior distribution of θ is beta with density (1).

We confirm our earlier results and intuition: the real option value of learning is positive¹¹, while the probability of meeting the crowdfunding target is decreasing in μ .¹² By construction, with a large sample, the firm's expected utility and probability of meeting the optimal investment target $\tilde{\theta}$ do not depend on μ .

Proposition 9 *Under the continuous approximation of the model, there is a unique $0 < \mu^* < 1$ that maximizes the expected funds raised through crowdfunding, and hence the platform fees proportional to them. Further, μ^* is decreasing in I and increasing in N .*

Proof. See Appendix A.3. ■

Proposition 9 confirms the intuition of the numerical example above that the expected funds raised via crowdfunding are highest when the sample size is limited and at an intermediate value. This proposition also highlights that *ceteris paribus*, firms that have more profitable projects ex ante, i.e. a larger target market and smaller investment costs, maximize the funds they raise with a longer campaign relative to firms that have less profitable projects.

¹¹If $I \geq N\theta_0$, then the real option value of learning is $\int_{\frac{I}{N}}^1 (N\theta - I) f_\theta(\theta) d\theta > 0$ and if $I < N\theta_0$, then the value of learning is $\int_{\frac{I}{N}}^1 (N\theta - I) f_\theta(\theta) d\theta - (N\theta_0 - I) = \int_0^{\frac{I}{N}} (N\theta_0 - I) f_\theta(\theta) d\theta$.

¹²Note that $\Pr(\theta \geq \tilde{\theta}_Y) = \int_{\tilde{\theta}_Y}^1 f_\theta(\theta) d\theta$, and $\frac{\partial \Pr(\theta \geq \tilde{\theta}_Y)}{\partial \mu} = -f_\theta(\tilde{\theta}_Y) \cdot \frac{\partial \tilde{\theta}_Y}{\partial \mu} < 0$, because $\frac{\partial \tilde{\theta}_Y}{\partial \mu} = \frac{I}{N(1-\mu)^2} > 0$.

In the special case where prior beliefs about θ are uniform, i.e., $\theta_0 = 0.5$ and $\lambda = 2$, we can derive that the optimal μ^* solves a cubic equation $(\mu^* - 1)^3 + (1 + \mu^*) \left(\frac{I}{N}\right)^2 = 0$. The discriminant of this is $-4 \left(\frac{I}{N}\right)^4 \left(\left(\frac{I}{N}\right)^2 + 27\right) < 0$, which implies that the cubic equation has only one real root. Assuming the same parameter values as in the numerical example on Figure 3, we obtain that $\mu^* = 0.3106$. This is almost the same value as in our discrete numerical example, where the optimal sample size is $M^* = 4966$, which implies $\mu^* = M^*/16000 = 0.3104$.

While we assume that the firm is not credit constrained at the time of investment, it is worth highlighting that a credit constrained firm would share the objective to maximize the funds raised and also benefit from setting the same intermediate limit for M^* .

6.2 Efficiency of the existing crowdfunding schemes and deferred payments.

An important finding of our paper is that AoN crowdfunding, as it takes place in the current environment, leads to outcomes very close to the first best despite the fact that we assumed an extreme moral hazard problem in which the firm can costlessly divert all funds raised. This suggests that the scope for improvement on current AoN crowdfunding schemes is somewhat limited.

Strausz (2017) suggests that to mitigate moral hazard, the platform could defer passing the funds on to the firm until the product is actually produced and delivered. Delaying passing the funds on to the firm would indeed enable the firm to set the crowdfunding target to the optimal investment target \tilde{m} instead of the higher target \tilde{m}_Y in our setting. The platform could, in principle, use the funds raised during the campaign as a collateral, and make a loan to the entrepreneur. Real world reward-based crowdfunding platforms do not engage in such practices. One likely reason is that holding backer funds as collateral and lending funds to a firm would come with additional costs and regulatory burden for the platform. Holding backer funds as a collateral may become particularly complicated in an environment where projects may fail due to exogenous reasons (as analyzed in Section 5). In such case the backers need to trust the platform to truthfully report the uncertain outcome of the investment, which could lead to a whole new range of agency problems between the platform and the backers. Another reason why platforms may choose to remain an information intermediary rather than become an investor is highlighted in Condorelli et. al, 2015. They consider platforms that have private information about consumer preferences and choose between referrals (passing on this information) or intermediating the trade. They show that referrals are optimal due to information unravelling. The current crowdfunding platform business model fits with this framework as a campaign allows the platform to obtain private information about consumer preferences, which it indeed currently transmits to the firm instead of becoming

an active investor. Platforms can truthfully intermediate information as their income depends on attracting firms to participate rather than on any individual firm's success.

A more simple but so far hypothetical mechanism that would ensure implementation of the optimal target \tilde{m} would enable firms and backers to use *conditional payments* where firms describe the investment cost breakdown: backer contributions go first to an escrow account, conditional on the proof of pre-agreed investment, and funds are only subsequently passed on to the firm. Such conditional payments are not yet possible on a large scale, but could become feasible in the future with, e.g., evolving blockchain technology. The Ethereum Project discusses the possible application of smart contracts to crowdfunding in a similar spirit.

7 Empirical implications

As we discuss below, our model of crowdfunding is consistent with existing empirical evidence and our findings suggests some new avenues for further empirical research.

Existing empirical evidence: Mollick (2014) shows that very few successful Kickstarter projects (3.6%) fail to deliver their promised rewards. Our analysis highlights the reasons why the existing crowdfunding mechanisms are efficient enough to endogenously overcome (even an extreme form of) moral hazard. We also argue that the real option value of learning, rather than credit constraints, is the main value driver of crowdfunding. Mollick and Kuppuswamy (2014) present evidence that the number one reason why both successful and unsuccessful firms in technology, project design, and video games category sought crowdfunding was "to see if there is demand for the project" (68% and 60% of successful and unsuccessful projects agreeing, respectively). Financing the project was only the fourth reason cited, while demand-related marketing and connecting with community were ranked second and third, respectively. Mollick and Kuppuswamy (2014) also find that 30% of firms continue to pursue their projects after failing to meet their target. Xu (2017) finds further evidence of Bayesian learning among Kickstarter participants. Regarding more specific findings of our paper Cumming et. al. (2015) finds evidence that AoN crowdfunding dominates KiA. Mollick (2014) find that shorter campaigns succeed with a higher probability than longer campaigns. In fact, Kickstarter itself shortened its maximum campaign duration from 90 to 60 days in 2011 referring to the observation that shorter campaigns are more likely to succeed.¹³

New empirical avenues: Uncertainty about target consumer preferences plays a central role in our model, and would warrant further empirical investigation. We predict that more uncertain project should be relatively more "overfunded" compared to less uncertain projects as the wedge

¹³<https://www.kickstarter.com/blog/shortening-the-maximum-project-length>

$\mathbb{E} \left[\frac{p_0 m}{p_0 \tilde{m}_Y} | p_0 m \geq p_0 \tilde{m}_Y \right] - p_0 \tilde{m}_Y$ in increasing in the degree of uncertainty. This is consistent with stylized facts presented in our Introduction. We also analyzed all successfully funded Kickstarter projects between January 1, 2015 and September 17, 2015 in extreme opposite categories: Technology and Theatre, and constructed unconditional distributions. Online Appendix B.5. confirms that Technology projects are more uncertain in the sense of second order stochastic dominance than Theatre ones, and our data indicates that the average successful US-based technology project raised 5.8 times its target, while an average successful US-based theatre project raised 1.3 times its target. Our model also predicts that firms with uncertain projects should participate more often, fail to meet their crowdfunding target with a higher probability, and complete their projects with a higher probability after an unsuccessful campaign.

Our model also provides a structure to test the severity of moral hazard and/or the importance of reputation costs and legal enforcement. We have shown that firms that benefit the most from crowdfunding are those with a break-even investment cost without crowdfunding. Hence these firms are likely to represent the typical crowdfunding participant. In the absence of moral hazard, these firms set the optimal target such that $\mathbb{E} \left[\frac{p_0 m}{p_0 \tilde{m}} \right] = 1$, i.e., the average completion ratio, $\frac{pledges}{goal}$, across all projects should be one as well. Instead, under moral hazard such typical firm would set a target $p_0 \tilde{m}_Y > p_0 \tilde{m}$, which implies a completion ratio below one. Indeed, Cumming et. al. (2015) analyze AoN and KiA projects at Indiegogo, and report an average completion ratio, $\frac{pledges}{goal} = \frac{pm}{p\tilde{m}}$, of 0.403, ranging from 0.337 to 0.617 across innovative, creative, and social categories and with both KiA and AoN schemes.

In addition to distinguishing projects according to preference uncertainty, it would also be interesting to distinguish projects according to variable costs. Our model suggests that firms with higher variable costs face a higher degree of moral hazard, and are thus less likely to participate, to set a higher target if they participate, and fail to meet their target with a higher probability.

Our off-equilibrium path results further suggest that projects that do not set ambitious targets are more likely fraudulent, unless they offer a bigger discount, and non-fraudulent KiA projects should offer larger discounts than comparable AoN projects.

Empirical patterns that help to distinguish our paper from other theoretical models: Ellman and Hurkens (2016) argue that reputation costs are high enough to overcome moral hazard. We show that both a low completion ratio and shorter campaigns being associated with higher failure rate are symptoms of moral hazard, which appear to be present in data. Strausz's (2017) and Varian's (2013) models predict that successful campaigns should not systematically raise funds in large excess compared to the target. The reasons are different: Strausz (2017) argues that backers

pursue a conditional pledging strategy to overcome moral hazard, Varian (2013) allows for some over-pledging, but only to the degree that each pivotal individual wants to obtain the full gift for his contribution. The magnitude of over-pledging observed, especially for technology projects, suggests that backers do not become less willing to participate after the firm has met its target. Quite the opposite, our model suggests that backers face less risk and are more willing to participate if a firm raises more funds during a limited length campaign. Our model also predict that the products are sold at par or at a discount at the crowdfunding stage compared to the ultimate retail price. While we do not have enough data to test this, anecdotal evidence based on firms' announced retail prices or on examples such as Pebble Watch suggests that this is the case. With more systematic data, our predictions could be tested as alternatives to those in Belleflamme et. al. (2013) who explain crowdfunding as a means to price-discriminate individuals who enjoy the crowdfunding experience. In contrast to our model, their results imply that the prices would be set at the crowdfunding stage at a premium over the retail price.

8 Concluding remarks

In our model, the only parameter the firm can learn about is total demand. We show that firms that face highly uncertain demand benefit the most from reward-based crowdfunding. We have shown that these firms can also overcome moral hazard most easily. In reality, the benefits can extend to learning about consumer preferences about the specifications of the product, e.g., the color of new widget or the features of a new game.¹⁴ We argue that to fully understand the success of reward-based crowdfunding, it is important to consider its role as a learning device, rather than focusing on a mere funding scheme.

Our model can explain the following stylized facts: 1) a noticeable proportion of firms that fail to meet their target in a crowdfunding campaign still complete the project, 2) many successful projects receive amounts significantly higher than the target amount, 3) rich and famous companies have crowdfunded projects, 4) firms that develop riskier products that involve sufficiently high investment costs and have high demand uncertainty, e.g. technology gadgets, appear to seek reward-based crowdfunding most often. We also predict that firms offer products at par or at a discount rather than premium during crowdfunding campaigns. These features are specific to our model and would enable to test our predictions relative to alternative papers.

Some established firms may incur high reputation costs. Indeed Sony and Apple routinely pre-

¹⁴Interestingly, new projects can then lead to the development of new products that can be either spawned or retained (Habib, Hege, and Mella-Barral 2013).

sell products via their own websites and may not need third-party platforms. However, most firms are widely unknown, do not have much reputation to lose, and cannot pre-commit to money-back guarantees. We have shown that third-party crowdfunding platforms are valuable to those firms because they facilitate commitment to features, in particular the crowdfunding targets and limited campaign length, that help them to overcome moral hazard.

References

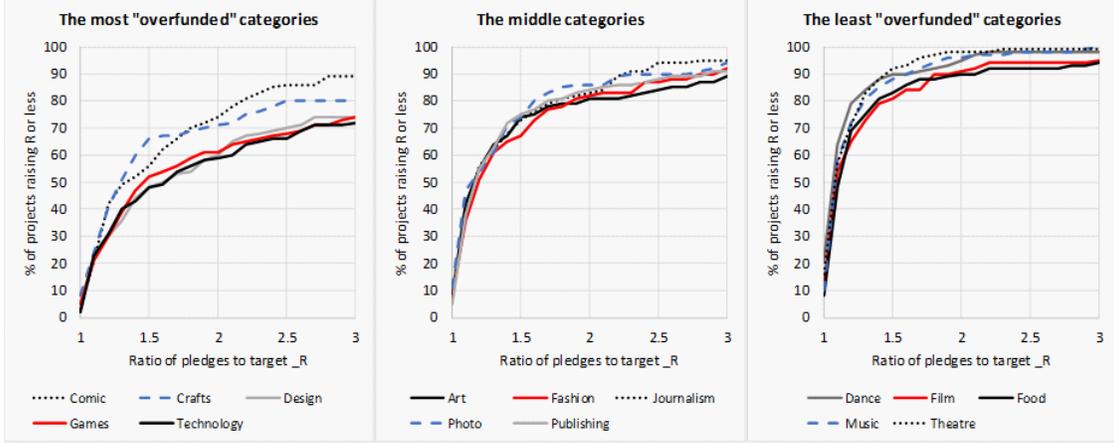
- [1] Angeletos G., G. Lorenzoni and A. Pavan, 2010, "Beauty Contests and Irrational Exuberance: A Neoclassical Approach", MIT.
- [2] Bagnoli, Mark, and Ted Bergstrom, 2005, "Log-concave probability and its application", *Economic Theory*, 26, 445-469.
- [3] Belleflamme, Paul, Lambert, Thomas, and Armin Schwiendbacher, 2014, "Crowdfunding: Tapping the Right Crowd", *Journal of Business Venturing*, 29, 585–609.
- [4] Bettignies, J. and G. Chemla, 2008, "Corporate Venture Capital, Allocation of Talent, and Competition for Star Managers", *Management Science*, 54, 505-521.
- [5] Condorelli, Daniele & Galeotti, Andrea, and Skreta, Vasiliki, 2017. "Selling Through Referrals," CEPR Discussion Paper number 12048.
- [6] Crawford, G., and M. Shum, 2005, "Uncertainty and Learning in Pharmaceutical Demand", *Econometrica*, 73(4), 1137–1173.
- [7] Chu, L. Y., and H. Zhang, 2011, "Optimal Pre-order Strategy with Endogenous Information Control", *Management Science*, 57, 6, 1055-77.
- [8] Dewatripont M. (1988), "Commitment through Renegotiation-Proof Contracts with Third Parties", *Review of Economic Studies*, 55, 377-90.
- [9] Edlin, Aaron and Chris Shannon, 1998, "Strict Monotonicity in Comparative Statics", *Journal of Economic Theory*, 81, 201-219.
- [10] Ellman, R., and S. Hurkens, 2016, "Optimal Crowdfunding Design", mimeo IESE.
- [11] Farrell, Joseph and Matthew Rabin, 1996, "Cheap Talk", *Journal of Economic Perspectives*, 10(3), 103-118.
- [12] Gromb, D., and D. Scharfstein, 2005, "Entrepreneurship in Equilibrium", mimeo Harvard Business School and HEC.
- [13] Habib, M., Hege, U., and P. Mella-Barral, 2013, "Entrepreneurial Spawning and Firm Characteristics", *Management Science*, 59, 12, .2790-2804.
- [14] Hart, O., 1995, "Firms, Contracts and Financial Structure", *Clarendon Lectures in Economics*, Oxford University Press.

- [15] Hellmann, T., 2002, "A Theory of Strategic Venture Investing", *Journal of Financial Economics*, 64, 285-314.
- [16] Hildenbrand, T., Puri, M., and J. Rocholl, 2016, "Averse Incentives in Crowdfunding", *Management Science*, 63, 3, 587-608.
- [17] InfoDev, 2013, *Crowdfunding's Potential for the Developing World*, Finance and Private Sector Development Department, The World Bank | 1818 H Street, NW | Washington DC, 20433.
- [18] Iyer, R., Khwaja, A., Luttmer, E., and K. Shue, 2015, "Screening Peers Softly: Inferring the Quality of Small Borrowers", *Management Science*, 62, 6, 1554-1577.
- [19] Kerr, W, Nanda, R., and M. Rhodes-Kropf, "Entrepreneurship as experimentation", *Journal of Economic Perspectives*, 28 (3), 25-48.
- [20] Kiyotaki, N., and J. Moore, 2002, "Evil Is the Root of All Money", *American Economic Review, Papers and Proceedings*, 92(2) 46-50.
- [21] Mollick, E., 2014, "The dynamics of crowdfunding: An exploratory study", *Journal of Private Equity*, 29, 1-16.
- [22] Mollick, E., and V. Kuppuswamy, 2014, "After the Campaign: Outcomes of Crowdfunding.", mimeo Wharton.
- [23] S. Myers, 1977, "Determinants of corporate borrowing", *Journal of Financial Economics*, 147-175.
- [24] Myers S. and Majluf (1984), "Corporate Financing and Investment Decisions when Firms have Information that Investors do not Have", *Journal of Financial Economics*, 13, 187-221.
- [25] Okuno-Fujiwara, Masahiro, Andrew Postlewaite and Kotaro Suzumura, 1990. "Strategic Information Revelation", *Review of Economic Studies*, 57(1), 25-47.
- [26] Rob R., 1991, "Learning and Capacity Expansion under Demand Uncertainty", *Review of Economic Studies*, 58, 4, 655-75.
- [27] Stiglitz, J., and A. Weiss, 1981, "Credit Rationing in Markets with Imperfect Information", *American Economic Review*, 71(3), 393-410.
- [28] Strausz, Roland, 2017, "A Theory of Crowdfunding - a mechanism design approach with demand uncertainty and moral hazard.", *American Economic Review*, 107, 6, 1430-1476.
- [29] Tirole, J., 1988, "The Theory of Industrial Organization", MIT Press.
- [30] Tirole, J., 2006, "The Theory of Corporate Finance", Princeton University Press.
- [31] Tinn, K., 2010, "Technology Adoption with Exit in Imperfectly Informed Equity Markets". *American Economic Review*, 100(3), 925-57.
- [32] Varian, H., 2013, "Public goods and private gifts", mimeo UC Berkeley.
- [33] Xu, T., 2017, "Learning from the Crowd: The Feedback Value of Crowdfunding", mimeo UBC.

A Appendix

A.1 "Oversubscription" across Kickstarter categories

The figures use a sample of 100 successful projects in each category completed by October 30, 2014. The classification to "most", "middle" and "least" oversubscribed projects is based on median values conditional on success.



A.2 Proof of Lemma 6

On the equilibrium path the firm anticipates that all 1-consumers in set $\{1, \dots, M\}$ choose to pledge and we obtain

$$\mathbb{E} [\pi^F | \Omega_1^F] = \mathbf{1}_1^F \cdot D(m) + T_{\bar{m}^T} (p_0 - \mathbf{1}_1^F) m - (1 - T_{\bar{m}^T}) \varsigma \cdot \mathbf{1}_1^F - \varphi(m, T_{\bar{m}^T})$$

By the law of total expectations and (10), the expected value of this is

$$\begin{aligned} \mathbb{E} [\pi^F] &= \Pr(\mathbf{1}_1^F = 1 \& T_{\bar{m}^T} = 1) \mathbb{E} [D(m) + (p_0 - 1) m | \mathbf{1}_1^F = 1 \& T_{\bar{m}^T} = 1] \\ &+ \Pr(\mathbf{1}_1^F = 1 \& T_{\bar{m}^T} = 0) \mathbb{E} [D(m) - \varsigma | \mathbf{1}_1^F = 1 \& T_{\bar{m}^T} = 0] \\ &+ \Pr(\mathbf{1}_1^F = 0 \& T_{\bar{m}^T} = 1) \mathbb{E} [p_0 m | \mathbf{1}_1^F = 0 \& T_{\bar{m}^T} = 1] - z \end{aligned} \quad (30)$$

We need to separately consider the possible targets $\bar{m}^T \in [\tilde{m}_Y, \infty)$, $\bar{m}^T \in [\tilde{m}_N, \tilde{m}_Y)$ and $\bar{m}^T \in [0, \tilde{m}_N)$.

Consider the interval $\bar{m}^T \in [\tilde{m}_Y, \infty)$. The firm always invests if it meets the target \tilde{m}_Y and if $m \geq \tilde{m}_N$. From (30) the firm must set $\{p_0, \bar{m}^T\}$ such that it maximizes

$$\Pr(m \geq \tilde{m}_N) \mathbb{E} [D(m) | m \geq \tilde{m}_N] + (p_0 - 1) \Pr(m \geq \bar{m}^T) \mathbb{E} [m | m \geq \bar{m}^T] - \varsigma \Pr(\tilde{m}_N \leq m < \bar{m}^T) - z$$

subject to the consumer participation constraint $p_0 \leq 1$. The first and the last term of this expression do not depend on p_0 and \bar{m}^T , the middle terms are maximized when $p_0 = 1$, and $\bar{m}^T = \tilde{m}_Y$. The firm's payoff is then

$$\mathbb{E}[\pi^F] |_{\bar{m}^T = \tilde{m}_Y} = \Pr(m \geq \tilde{m}_N) \mathbb{E}[D(m) | m \geq \tilde{m}_N] - \varsigma \Pr(\tilde{m}_N \leq m < \tilde{m}_Y) - z \quad (31)$$

In the other two cases, the pre-selling price must satisfy the constraint $p_0 \leq \frac{\Pr(m_{-i} \geq \tilde{m}_Y - 1 | v_i = 1)}{\Pr(m_{-i} \geq \bar{m}_T - 1 | v_i = 1)}$. From Bayes' rule and the law of iterated expectations the consumers' participation constraint is

$$p_0 \leq \frac{\lambda \theta_0 + \mathbb{E}[m | m \geq \tilde{m}_Y] \Pr(m \geq \tilde{m}_Y)}{\lambda \theta_0 + \mathbb{E}[m | m \geq \bar{m}_T] \Pr(m \geq \bar{m}_T)}. \quad (32)$$

Since now $\tilde{m}_Y \geq \bar{m}_T$, we know from Lemma 10 in Online Appendix B.1. that $\mathbb{E}[m | m \geq \tilde{m}_Y] \geq \mathbb{E}[m | m \geq \bar{m}_T]$, which implies

$$p_0 \leq \frac{\lambda \theta_0 + \mathbb{E}[m | m \geq \tilde{m}_Y] \Pr(m \geq \tilde{m}_Y)}{\lambda \theta_0 + \mathbb{E}[m | m \geq \bar{m}_T] \Pr(m \geq \bar{m}_T)} \leq \frac{\mathbb{E}[m | m \geq \tilde{m}_Y] \Pr(m \geq \tilde{m}_Y)}{\mathbb{E}[m | m \geq \bar{m}_T] \Pr(m \geq \bar{m}_T)}. \quad (33)$$

Consider the interval $\bar{m}^T \in [\tilde{m}_N, \tilde{m}_Y)$. Expanding (30) and using (33) imply that the firm's expected profit is lower than (31) and the firm will not set a target $\bar{m}^T \in [\tilde{m}_N, \tilde{m}_Y)$. A similar argument holds for $\bar{m}^T \in [0, \tilde{m}_N)$. This proves that the firm's optimal strategy at date 0 is to set $\{p_0, \bar{m}^T\} = \{1, \tilde{m}_Y\}$.

A.3 Proof of Proposition 9

Under the approximation described in the text, the expected funds raised via crowdfunding are given by

$$\mathbb{E}[funds] = N\mu \int_0^1 \theta \frac{\theta^{\lambda \theta_0 - 1} (1 - \theta)^{\lambda \theta_0 - 1}}{B(\lambda \theta_0, \lambda(1 - \theta_0))} d\theta = N\mu \frac{B(\lambda \theta_0 + 1, \lambda(1 - \theta_0))}{B(\lambda \theta_0, \lambda(1 - \theta_0))} \int_{\tilde{\theta}_Y}^1 \frac{\theta^{\lambda \theta_0} (1 - \theta)^{\lambda \theta_0 - 1}}{B(\lambda \theta_0 + 1, \lambda(1 - \theta_0))} d\theta.$$

The term $\frac{\theta^{\lambda \theta_0} (1 - \theta)^{\lambda \theta_0 - 1}}{B(\lambda \theta_0 + 1, \lambda(1 - \theta_0))}$ is the density of a random variable drawn from beta distribution $Be(\lambda \theta_0 + 1, \lambda(1 - \theta_0))$. Denote x this random variable such that $f_x(x)$ is the probability distribution function and $F_x(\tilde{\theta}_Y) = \int_0^{\tilde{\theta}_Y} f_x(x) dx$ is the cumulative density of $x \sim Be(\lambda \theta_0 + 1, \lambda(1 - \theta_0))$.

The term $N \frac{B(\lambda \theta_0 + 1, \lambda(1 - \theta_0))}{B(\lambda \theta_0, \lambda(1 - \theta_0))}$ does not depend on μ and thus $\mathbb{E}[funds] \propto \mu (1 - F_x(\tilde{\theta}_Y))$. A corner solution $\mu = 0$ or $\mu = 1$ cannot be a maximum because $\mu (1 - F_x(\tilde{\theta}_Y)) > 0$ for intermediate values of μ , but is 0 at extremes. To prove that $\mu^* = \arg \max \mu (1 - F_x(\tilde{\theta}_Y))$ is a unique interior solution it is sufficient to show that $\mathbb{E}[funds]$ is log-concave. We find that

$$\frac{\partial^2 \ln \left(\mu \left(1 - F_x(\tilde{\theta}_Y) \right) \right)}{\partial \mu \partial \mu} = -\frac{1}{\mu^2} - \frac{\partial \left(\frac{f_x(\tilde{\theta}_Y)}{1 - F_x(\tilde{\theta}_Y)} \right)}{\partial \tilde{\theta}_Y} \frac{\partial \tilde{\theta}_Y}{\partial \mu} - \frac{f_x(\tilde{\theta}_Y)}{1 - F_x(\tilde{\theta}_Y)} \frac{\partial^2 \tilde{\theta}_Y}{\partial \mu \partial \mu} < 0$$

It is straightforward that $\frac{1}{\mu^2} > 0$, $\frac{\partial \tilde{\theta}_Y}{\partial \mu} = \frac{I}{N(1-\mu)^2} > 0$ and $\frac{\partial^2 \tilde{\theta}_Y}{\partial \mu \partial \mu} = 2 \frac{I}{N(1-\mu)^3} > 0$. Furthermore, as proved in Bagnoli and Bergstrom (2005) $\frac{f_x(\tilde{\theta}_Y)}{1-F_x(\tilde{\theta}_Y)}$ is increasing $\tilde{\theta}_Y$ if the density function $f_x(\cdot)$ is log-concave (Corollary 2 of Theorem 3 there) or if it is monotone increasing (Corollary 3 of Theorem 3 there). Indeed, the former holds if $\lambda(1-\theta_0) \geq 1$ and the latter holds if $\lambda(1-\theta_0) < 1$. This proves that the problem has a unique interior maximum.

In order to derive comparative statics, notice that parameters I and N have an impact on $\mathbb{E}[funds]$ only via the term $\frac{I}{N} = \tilde{\theta}$, and $\tilde{\theta}_Y = \frac{\tilde{\theta}}{(1-\mu)}$. We can express

$$\mathbb{E}[funds] \propto \Psi\left(\mu, \tilde{\theta}\right) = \mu \int_{\frac{\tilde{\theta}}{(1-\mu)}}^1 x^{\lambda\theta_0} (1-x)^{\lambda\theta_0-1} dx$$

as the rest of the terms do not depend on μ and $\tilde{\theta}$. By strict comparative statics results proved by Edlin and Shannon (1998) $\mu^*(\tilde{\theta})$ is increasing (decreasing) in $\tilde{\theta}$ if $\frac{\partial \Psi(\mu, \tilde{\theta})}{\partial \mu}$ is strictly increasing (decreasing) in $\tilde{\theta}$. We find

$$\frac{\partial^2 \Psi(\mu, \tilde{\theta})}{\partial \mu \partial \tilde{\theta}} = - \left(\frac{\tilde{\theta}}{(1-\mu)} \right)^{\lambda\theta_0} \left(1 - \frac{\tilde{\theta}}{(1-\mu)} \right)^{\lambda(1-\theta_0)-1} \frac{1}{(1-\mu)^2} T(\mu)$$

The terms $\frac{\tilde{\theta}}{(1-\mu)}$, $\left(1 - \frac{\tilde{\theta}}{(1-\mu)}\right)$, $\frac{1}{(1-\mu)^2} > 0$. What we need to determine is the sign of $T(\mu) \equiv 1 + \mu\lambda\theta_0 - \mu(\lambda(1-\theta_0) - 1) \frac{\tilde{\theta}}{1-\mu-\tilde{\theta}}$. Note that $T(\mu=0) = 1 > 0$ and $T(\mu=1) = \lambda > 0$. We find $\frac{\partial T(\mu)}{\partial \mu} = \lambda\theta_0 - (\lambda(1-\theta_0) - 1) \tilde{\theta} \frac{1-\theta}{(1-\mu-\tilde{\theta})^2}$. If $\lambda(1-\theta_0) \geq 1$, then $T(\mu)$ is increasing, which implies that $T(\mu) \geq T(\mu=0) = 1 > 0$ for any $\mu \in [0, 1]$. If $\lambda(1-\theta_0) < 1$, then we find that $T(\mu)$ is concave (as $\frac{\partial^2 T(\mu)}{\partial \mu \partial \mu} \propto -2 \frac{1}{(1-\mu-\tilde{\theta})^3} < 0$), and thus $T(\mu) \geq \min[1, \lambda] > 0$ for any $\mu \in [0, 1]$. Therefore, $\frac{\partial^2 \Psi(\mu, \tilde{\theta})}{\partial \mu \partial \tilde{\theta}} < 0$. This together with $\tilde{\theta} = \frac{I}{N}$ being increasing in I and decreasing in N proves the comparative statics claim.

B Online appendices

B.1 Useful general results

Lemma 10 For any $c \in \{0, 1, \dots, M\}$, $\mathbb{E}[m|m \geq c]$ is increasing in c .

Proof. Since c is an integer, it is sufficient to show that for any $c \in \{0, 1, \dots, M\}$ we have

$$\mathbb{E}[m|m \geq c+1] > \mathbb{E}[m|m \geq c],$$

which can be expressed as

$$\frac{\sum_{m=c+1}^M m q_m}{\sum_{m=c+1}^M q_m} > \frac{\sum_{m=c}^M m q_m}{\sum_{m=c}^M q_m} = \frac{c q_c + \sum_{m=c+1}^M m q_m}{q_c + \sum_{m=c+1}^M q_m}$$

Simplifying, we obtain

$$q_c \sum_{m=c+1}^M m q_m > c q_c \sum_{m=c+1}^M q_m \Leftrightarrow q_c \sum_{m=c+1}^M (m-c) q_m > 0,$$

where the inequality holds because $m-c > 0$ for any $c+1 \leq m \leq M$ and with a beta-binomial distribution, $q_m > 0$ for any m . ■

Lemma 11 For any $c \in \{0, 1, \dots, M\}$, $\Pr(m \geq c) \mathbb{E}[D(m)|m \geq c]$ is monotonically increasing (decreasing) in c for any $c < (>) \tilde{m}$; $\Pr(m \geq c) \mathbb{E}[D(m)|m \geq c]$ is maximized when $c = \tilde{m}$, where \tilde{m} is given by (14).

Proof. As c is an integer, consider the difference

$$\begin{aligned} & \Pr(m \geq c+1) \mathbb{E}[D(m)|m \geq c+1] - \Pr(m \geq c) \mathbb{E}[D(m)|m \geq c] \\ &= \sum_{m=c+1}^M D(m) q_m - \sum_{m=c}^M D(m) q_m = -D(c) q_c. \end{aligned}$$

Since $q_c > 0$ for any c , the difference has the same sign as $-D(c) q_c$. From (14) and (15), we obtain

$$\begin{aligned} D(c) &< 0 \text{ if } c < \tilde{m} \\ D(c) &\geq 0 \text{ if } c = \tilde{m}, \\ D(c) &> 0 \text{ if } c > \tilde{m} \end{aligned}$$

which implies that $\Pr(m \geq c) \mathbb{E}[D(m)|m \geq c]$ is indeed monotonically increasing (decreasing) in c for any $c < (>) \tilde{m}$. This also implies that $\Pr(m \geq c) \mathbb{E}[D(m)|m \geq c]$ is maximized when $c = \tilde{m}$.

■

B.2 Proof of Proposition 2

We only need to consider $I_0 < I < I_0 + M \frac{\lambda+N}{\lambda+M}$, as the value of learning is zero otherwise. From (14), it is clear that \tilde{m} is weakly increasing in I .

Consider $I < N\theta_0$. From (15) and (17), the expected value of learning is

$$U_F^{B,I} = -\mathbb{E}[D(m) | m < \tilde{m}] \Pr(m < \tilde{m}) = -\sum_{m=0}^{\tilde{m}-1} D(m) q_m.$$

Since $-D(m) = \left(I - \frac{\lambda+N}{\lambda+M}m - I_0\right)$ is positive and increasing in I for any $m < \tilde{m}$, and since \tilde{m} is weakly increasing in I , we obtain that $U_F^{B,I}$ is increasing in I .

Consider $I > N\theta_0$. From (15) and (18), the expected value of learning is

$$U_F^{B,NI} = \mathbb{E}[D(m) | m \geq \tilde{m}] \Pr(m \geq \tilde{m}) = \sum_{m=\tilde{m}}^M D(m) q_m.$$

Since $D(m) = \left(\frac{\lambda+N}{\lambda+M}m + I_0 - I\right)$ is positive and decreasing in I for any $m \geq \tilde{m}$ and since \tilde{m} is weakly increasing in I , $U_F^{B,NI}$ is decreasing in I . This implies that $I = N\theta_0$ maximizes the expected value of learning.

B.3 Proof of Proposition 3

For the sake of exposition, we have limited indexing the variables of interest in the main text and only highlighted the dependence on m . For this proof it is necessary to consider the dependence on λ . We denote

$$\begin{aligned} q_m(\lambda) &= \binom{M}{m} \frac{B(\lambda\theta_0 + m, \lambda(1-\theta_0) + M - m)}{B(\lambda\theta_0, \lambda(1-\theta_0))} \\ \tilde{m}(\lambda) &= \left\lceil \frac{\lambda + M}{\lambda + N} (I - I_0(\lambda)) \right\rceil \\ D(m, \lambda) &= \frac{\lambda + N}{\lambda + M} m + I_0(\lambda) - I, \end{aligned} \tag{34}$$

where $I_0(\lambda) = \frac{\lambda\theta_0(N-M)}{\lambda+M}$. The equations in (34) are identical to (12), (14) and (15) respectively.

We also denote the cumulative distribution of m

$$Q_m(\lambda) = \sum_{k=0}^m q_k(\lambda).$$

Since the distribution of m is beta-binomial, for any pair $\lambda \in \{\lambda_1, \lambda_2\}$ such that $\lambda_2 > \lambda_1$, the distribution with higher λ_2 second order stochastically dominates (SOSD) the one with λ_1 , i.e.,

$$\sum_{k=0}^m Q_k(\lambda_2) \leq \sum_{k=0}^m Q_k(\lambda_1) \tag{35}$$

for all $m = \{0, 1, \dots, M\}$; the inequality is strict for $m < M - 1$.¹⁵

From (17) and (18), we obtain

$$\begin{aligned} U_F^{B,I}(\lambda) &= - \sum_{m=0}^{\tilde{m}(\lambda)-1} D(m, \lambda) q_m(\lambda) \text{ for } I < N\theta_0 \\ U_F^{B,NI}(\lambda) &= (N\theta_0 - I) - \sum_{m=0}^{\tilde{m}(\lambda)-1} D(m, \lambda) q_m(\lambda) \text{ for } I \geq N\theta_0 \end{aligned} \quad (36)$$

From Abel's Lemma¹⁶, and $D(m+1, \lambda) - D(m, \lambda) = \frac{\lambda+N}{\lambda+M}$,

$$\sum_{m=0}^{\tilde{m}(\lambda)-1} D(m, \lambda) q_m(\lambda) = D(\tilde{m}(\lambda), \lambda) Q_{\tilde{m}(\lambda)-1}(\lambda) - \frac{\lambda+N}{\lambda+M} \sum_{m=0}^{\tilde{m}(\lambda)-1} Q_m(\lambda) \quad (37)$$

We can express the investment threshold as

$$\tilde{m}(\lambda) = \frac{\lambda+M}{\lambda+N} (I - I_0(\lambda)) + \varepsilon(\lambda), \quad (38)$$

where $0 \leq \varepsilon(\lambda) < 1$ is the rounding error. Using (34) and (38) we obtain

$$D(\tilde{m}(\lambda), \lambda) = \frac{\lambda+N}{\lambda+M} \tilde{m}(\lambda) + I_0(\lambda) - I = \frac{\lambda+N}{\lambda+M} \varepsilon(\lambda) \quad (39)$$

Replacing (37) and (39) in (36), we express the firm's expected value of learning as

$$U_F^{B,I}(\lambda) = \frac{\lambda+N}{\lambda+M} \left(\varepsilon(\lambda) \sum_{m=0}^{\tilde{m}(\lambda)-2} Q_m(\lambda) + (1 - \varepsilon(\lambda)) \sum_{m=0}^{\tilde{m}(\lambda)-1} Q_m(\lambda) \right), \quad (40)$$

$$U_F^{B,NI}(\lambda) = (N\theta_0 - I) + \frac{\lambda+N}{\lambda+M} \left(\varepsilon(\lambda) \sum_{m=0}^{\tilde{m}(\lambda)-2} Q_m(\lambda) + (1 - \varepsilon(\lambda)) \sum_{m=0}^{\tilde{m}(\lambda)-1} Q_m(\lambda) \right). \quad (41)$$

From (34), we obtain

$$\begin{aligned} \tilde{m}(\lambda_2) &\leq \tilde{m}(\lambda_1) \text{ if } I < N\theta_0 \\ \tilde{m}(\lambda_2) &\geq \tilde{m}(\lambda_1) \text{ if } I > N\theta_0, \end{aligned}$$

If $I = N\theta_0$, then $\tilde{m}(\lambda_2) = \tilde{m}(\lambda_1) = \lceil I \rceil$ is independent of λ . Since the effect of λ on the firms that would break even without learning is different from its effect on the firms that would not, we analyze these two cases separately.

Case 1: firms with $I < N\theta_0$

In this case, we have $\tilde{m}(\lambda_2) \leq \tilde{m}(\lambda_1)$. We denote $\Delta\tilde{m} = -(\tilde{m}(\lambda_2) - \tilde{m}(\lambda_1))$, where $\Delta\tilde{m} = \{0, 1, \dots\}$. From (38), we obtain

$$\varepsilon(\lambda_2) = \varepsilon(\lambda_1) - \Delta\tilde{m} + \frac{(\lambda_2 - \lambda_1)(N - M)}{(\lambda_2 + N)(\lambda_1 + N)} (N\theta_0 - I)$$

¹⁵This can be proved analytically from the fact that beta-binomial distributions have a unimodal likelihood ratio, which implies second order stochastic dominance. See Hopkins, Kornienko, 2003 for the proof of continuous distributions, the proof for discrete distributions is available upon request.

¹⁶ $\sum_{i=0}^N a_i b_i = A_N b_N - \sum_{i=0}^{N-1} A_i (b_{i+1} - b_i)$, where $A_n = \sum_{i=0}^n a_i$.

This and (40) enable us to decompose the effect of an increase in λ as

$$U_F^{B,I}(\lambda_2) - U_F^{B,I}(\lambda_1) = -\frac{\lambda_1 + N}{\lambda_1 + M} (G_1^I + G_2^I + G_3^I) \quad (42)$$

$$-\frac{(\lambda_2 - \lambda_1)(N - M)}{(\lambda_2 + N)(\lambda_1 + M)} \left(U_F^{B,I}(\lambda_2) + Q_{\tilde{m}(\lambda_2)-1}(\lambda_2)(N\theta_0 - I) \right),$$

where

$$G_1^I \equiv \varepsilon(\lambda_1) \left(\sum_{m=0}^{\tilde{m}(\lambda_1)-2} Q_m(\lambda_1) - \sum_{m=0}^{\tilde{m}(\lambda_2)-2} Q_m(\lambda_2) \right),$$

$$G_2^I \equiv (1 - \varepsilon(\lambda_1)) \left(\sum_{m=0}^{\tilde{m}(\lambda_1)-1} Q_m(\lambda_1) - \sum_{m=0}^{\tilde{m}(\lambda_2)-1} Q_m(\lambda_2) \right),$$

$$G_3^I \equiv -\Delta\tilde{m}Q_{\tilde{m}(\lambda_2)-1}(\lambda_2).$$

The second term in (42) is non-positive and it is strictly negative if $N < M$, because $U_F^{B,I}(\lambda_2) > 0$ and $I < N\theta_0$. The sign of the first term in (42) is determined by $G_1^I + G_2^I + G_3^I$.

If $\Delta\tilde{m} = 0$, i.e., $\tilde{m}(\lambda_1) = \tilde{m}(\lambda_2)$, then $G_1^I + G_2^I + G_3^I \geq 0$ follows from $G_3 = 0$, and $G_1^I, G_2^I \geq 0$ due to second order stochastic dominance (35).

If $\Delta\tilde{m} > 0$, then we can further decompose

$$G_1^I + G_2^I + G_3^I = \varepsilon(\lambda_1) \left(\sum_{m=0}^{\tilde{m}(\lambda_1)-2} Q_m(\lambda_1) - \sum_{m=0}^{\tilde{m}(\lambda_1)-2} Q_m(\lambda_2) \right)$$

$$+ (1 - \varepsilon(\lambda_1)) \left(\sum_{m=0}^{\tilde{m}(\lambda_1)-1} Q_m(\lambda_1) - \sum_{m=0}^{\tilde{m}(\lambda_1)-1} Q_m(\lambda_2) \right)$$

$$+ \sum_{m=\tilde{m}(\lambda_2)-1}^{\tilde{m}(\lambda_2) + \Delta\tilde{m}-2} (Q_m(\lambda_2) - Q_{\tilde{m}(\lambda_2)-1}(\lambda_2))$$

$$+ (1 - \varepsilon(\lambda_1)) (Q_{\tilde{m}(\lambda_2) + \Delta\tilde{m}-1}(\lambda_2) - Q_{\tilde{m}(\lambda_2)-1}(\lambda_2)),$$

which is non-negative for any $\Delta\tilde{m} > 0$, because of second order stochastic dominance (35) and because the cumulative probability is increasing in m . Hence the value of learning decreases with λ for all firms with $I < N\theta_0$.

Case 2: firms with $I \geq N\theta_0$

In this case, we have $\tilde{m}(\lambda_2) \geq \tilde{m}(\lambda_1)$. We again denote $\Delta\tilde{m} = \tilde{m}(\lambda_2) - \tilde{m}(\lambda_1)$, where $\Delta\tilde{m} = \{0, 1, \dots\}$ and from (38) we obtain the following relationship between $\varepsilon(\lambda_1)$ and $\varepsilon(\lambda_2)$:

$$\varepsilon(\lambda_1) = \varepsilon(\lambda_2) - \Delta\tilde{m} + \frac{(\lambda_2 - \lambda_1)(N - M)(I - N\theta_0)}{(\lambda_2 + N)(\lambda_1 + N)}. \quad (43)$$

The above and (40) enable us to decompose again the effect of an increase in λ as

$$\begin{aligned}
U_F^{B,NI}(\lambda_2) - U_F^{B,NI}(\lambda_1) &= -\frac{\lambda_1 + N}{\lambda_1 + M} (G_1^{NI} + G_2^{NI} + G_3^{NI}) \\
&\quad - \frac{(\lambda_2 - \lambda_1)(N - M)}{(\lambda_2 + N)(\lambda_1 + M)} \left(U_F^{B,NI}(\lambda_2) + (I - N\theta_0)(1 - Q_{\tilde{m}(\lambda_1)-1}(\lambda_1)) \right),
\end{aligned} \tag{44}$$

where

$$\begin{aligned}
G_1^{NI} &\equiv \varepsilon(\lambda_2) \left(\sum_{m=0}^{\tilde{m}(\lambda_2)-2} Q_m(\lambda_2) - \sum_{m=0}^{\tilde{m}(\lambda_1)-2} Q_m(\lambda_1) \right), \\
G_2^{NI} &\equiv (1 - \varepsilon(\lambda_2)) \left(\sum_{m=0}^{\tilde{m}(\lambda_2)-1} Q_m(\lambda_2) - \sum_{m=0}^{\tilde{m}(\lambda_1)-1} Q_m(\lambda_1) \right), \\
G_3^{NI} &\equiv -\Delta\tilde{m}Q_{\tilde{m}(\lambda_1)-1}(\lambda_1).
\end{aligned}$$

It is clear that the second term of (44) is non-positive and it is strictly negative if $N < M$, because $\mathbb{E}[\pi_F]^{B,NI}(\lambda_2) > 0$, and $I \geq N\theta_0$. As in case 1, it is easy to see that if $\Delta\tilde{m} = 0$, then $\tilde{m}(\lambda_1) = \tilde{m}(\lambda_2)$ and $G_1^{NI} + G_2^{NI} + G_3^{NI} \geq 0$. If $\Delta\tilde{m} > 0$, then we can follow the same derivation as before to prove that $G_1^{NI} + G_2^{NI} + G_3^{NI} \geq 0$ (keeping in mind that now $\tilde{m}(\lambda_2) > \tilde{m}(\lambda_1)$).

This proves that value of learning increases with the level of uncertainty, i.e., decreases with λ .

B.4 Proof of Lemma 8

From (12), the definition of the beta function $B(x, y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$, and the fact that the gamma function $\Gamma(x)$ satisfies the *recurrence relation* $\Gamma(x+1) = x\Gamma(x)$, we find that the likelihood ratio satisfies

$$\frac{q(m; M+1)}{q(m; M)} = \frac{M+1}{M+1-m} \cdot \frac{\lambda(1-\theta_0) + M - m}{\lambda + M}$$

Differentiating this with respect to m , we get $\frac{\partial \frac{q(m; M+1)}{q(m; M)}}{\partial m} = \frac{(M+1)(\lambda(1-\theta_0)-1)}{(M-m+1)^2(M+\lambda)}$. Depending on the parameters of the model, the likelihood ratio can be monotonically decreasing ($\lambda(1-\theta_0) < 1$), constant ($\lambda(1-\theta_0) = 1$) or monotonically increasing ($\lambda(1-\theta_0) > 1$).

Consider first the high uncertainty case where $\lambda(1-\theta_0) < 1$. We find that

$$Q(m; M+1) - Q(m; M) = \sum_{k=0}^m \left(\frac{q(k; M+1)}{q(k; M)} - 1 \right) q(k; M)$$

Since $\frac{q(0; M+1)}{q(0; M)} = \frac{\lambda(1-\theta_0)+M}{\lambda+M} < 1$ and the likelihood ratio is monotonically decreasing, FOSD holds for M .

Consider then the intermediate case where $\lambda(1-\theta_0) = 1 \implies \lambda = \frac{1}{(1-\theta_0)}$. We obtain

$$Q(m; M+1) = \sum_{k=0}^m \frac{q(k; M+1)}{q(k; M)} q(k; M) = \frac{M+1}{M+\lambda} Q(m; M)$$

Since $\frac{M+1}{M+\lambda} = \frac{M+1}{\frac{1}{1-\theta_0}+M} < 1$, FOSD holds for M .

Finally, consider the low uncertainty case where $\lambda(1-\theta_0) > 1$. Because the likelihood ratio is monotonically increasing, we must have

$$Q(m; M+1) = \sum_{k=0}^m \frac{q(k; M+1)}{q(k; M)} q(k; M) < \frac{q(m; M+1)}{q(m; M)} \sum_{k=0}^m q(k; M)$$

where the last inequality holds because the likelihood ratio is monotonically increasing. At the same time

$$\begin{aligned} 1 - Q(m; M+1) &= q(M+1; M+1) + \sum_{k=m+1}^M q(k; M+1) \\ &= q(M+1; M+1) + \sum_{k=m+1}^M \frac{q(k; M+1)}{q(k; M)} q(k; M) > \\ &> q(M+1; M+1) + \frac{q(m+1; M+1)}{q(m+1; M)} (1 - Q(m; M)), \end{aligned}$$

where again the inequality holds because the likelihood ratio is monotonically increasing. From these two inequalities we obtain that for any $m = \{0, \dots, M-1\}$

$$\frac{Q(m; M+1)}{Q(m; M)} < \frac{q(m; M+1)}{q(m; M)} < \frac{q(M+1; M+1)}{(1 - Q(m; M))} + \frac{q(m+1; M+1)}{q(m+1; M)} < \frac{1 - Q(m; M+1)}{1 - Q(m; M)},$$

where the middle inequality holds because $\frac{q(m; M+1)}{q(m; M)} < \frac{q(m+1; M+1)}{q(m+1; M)}$ and $\frac{q(M+1; M+1)}{(1 - Q(m; M))} > 0$. From here,

$$\frac{Q(m; M+1)}{Q(m; M)} < \frac{1 - Q(m; M+1)}{1 - Q(m; M)} \Leftrightarrow Q(m; M+1) < Q(m; M)$$

for any $m = \{0, \dots, M-1\}$ and $Q(M; M+1) < Q(M; M) = 1$ holds because $Q(M; M+1) = 1 - q(M+1; M+1) < 1$.

B.5 Distribution of the ratio of pledges to target

The figures below are based on data from Kickstarter. The upper figure, which plots $\Pr\left(\frac{pm}{pm} \leq x \mid \frac{pm}{pm} \geq 1\right)$, is constructed using all 2015 Technology and Theatre projects completed from January 1, 2015 to September 17, 2015. We then constructed the unconditional distribution, $\Pr\left(\frac{pm}{pm} \leq x\right)$, thanks to the summary statistics available on Kickstarter on September 17, 2015. This statistics cover the full history of Kickstarter, so we assume that the 2015 data is statistically similar to earlier data. We use $\Pr\left(\frac{pm}{pm} \geq 1\right)$ reported on the site, which is 0.2018 for Technology category and 0.6089 for Theatre, and the frequency of unsuccessful projects raising 0%, 1 – 20%, ..., 81 – 99% in these categories. We then use the law of total expectations to find $\Pr\left(\frac{pm}{pm} \leq x\right)$.

