

Coordinating Donors on Crowdfunding Websites

Rick Wash
Michigan State University
wash@msu.edu

Jacob Solomon
Michigan State University
solomo93@msu.edu

ABSTRACT

Crowdfunding websites like Kickstarter, Spot.U.s and Donor's Choose seek to fund multiple projects simultaneously by soliciting donations from a large number of donors. Crowdfunding site designers must decide what to do with donations to projects that don't reach their goal by the deadline. Some crowdfunding sites use an all-or-nothing *return rule* in which donations are returned to donors if a project doesn't meet its goal. Other sites use a *direct donation* structure where all donations are kept by the project even if the total is insufficient.

We simulated a crowdfunding site using a threshold public goods game in which a set of donors tries to fund multiple projects that vary in riskiness. We find that the return rule mechanism leads to a marginal improvement in productivity of a site – more money is donated in total – by eliciting more donations. However, the return rule also leads to a potential loss in efficiency (percentage of projects funded) because donations become *spread* across too many projects and are not coordinated to achieve the maximum possible impact. The direct donation model, though, encourages donors to coordinate to create a more efficient but slightly less productive marketplace.

Author Keywords

Crowdfunding; Quantitative; Experiment; Public Goods; Crowdsourcing

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

Raising money to do a project is difficult; projects are inherently risky because they might not succeed in achieving their goals. When a project needs money to succeed, convincing people to contribute money to the project is challenging because potential donors see a lot of risk. Crowdfunding websites attempt to help by publicizing projects that need money and allowing many contributors to each make a small contribution toward the larger project. By aggregating many small donations, crowdfunding websites can fund

large and interesting projects of all kinds. For example, the movie writer Charlie Kaufman and his associates raised over \$400,000 from over 5,000 donors to fund the creation of an animated movie¹. Kickstarter, Spot.U.s, and Donors Choose are all examples of crowdfunding websites targeted at specific types of projects (creative, journalism, and classroom projects respectively).

Crowdfunding websites are two-sided matching marketplaces: they seek to match donors with projects that need money. However, crowdfunding websites introduce a second type of risk to projects: they allow people to contribute less money than the amount needed for the project, with no guarantee that the project will ever receive sufficient funds. For a project to receive its needed total, donors must coordinate. Crowdfunding websites do not use an algorithmic mechanism to achieve coordination, but instead provide an open marketplace for project creators and potential donors to match themselves. However, crowdfunding sites can be designed to enforce rules and mechanisms that help this coordination. Thus, crowdfunding site designers face important design questions that impact how well the site coordinates donation activity to fund projects.

One such design question is what to do when a project receives some donations, but these donations are insufficient to completely fund the project. Many crowdfunding websites have chosen a marketplace-enforced *return rule*: if a project does not receive sufficient funding by a pre-specified deadline, then all donations to that project are returned to the donors. The return rule formalizes a structure already present in donor preferences: donations to projects complement each other. Each donation is itself insufficient to fund a project, and therefore each donor prefers to donate to projects that are receiving donations from others, so that the project receives enough funds to succeed.

Other crowdfunding websites use a more traditional *direct donation* rule: all donated funds (minus the marketplace cut) are delivered to the project, regardless of whether they total up to the pre-specified goal of the project.

We ask two questions. First, how does this return rule affect the donation decisions of individuals who visit a crowdfunding site? Second, given the effects on individual donors, what are the consequences of the return rule for a site as a whole when these individual donations are aggregated? Specifically, does the return rule affect the ability of a site to fund projects and the efficiency with which this is accomplished?

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CSCW'14, February 15–19, 2014, Baltimore, Maryland, USA.
Copyright © 2014 ACM 978-1-4503-2540-0/14/02...\$15.00.
<http://dx.doi.org/10.1145/2531602.2531678>

¹<http://www.kickstarter.com/projects/anomalisa/charlie-kaufmans-anomalisa>

In this paper, we show that the return rule increases donors' willingness to donate. We also show that the return rule leads donors to donate according to their preferences rather than focusing on the projects that are more likely to be funded. These effects on individual donors have two consequences for a crowdfunding site. The increase in donations helps the site produce more projects. However, the return rule simultaneously reduces donor coordination, causing donations to become spread out among more projects, including those with little chance of being completed. This *spread* largely offsets the gains in site productivity obtained from increased donations

We contribute to crowdfunding site design by experimentally demonstrating that the decision to use either a return rule or direct donation model affects both individual and site-level behavior and outcomes on a crowdfunding site. We also demonstrate the importance of designing crowdfunding sites not only to solicit more donations from the crowd, but also designing to facilitate coordination among the crowd.

BACKGROUND

Crowdfunding

Crowdfunding is the act of soliciting, via an open call, resources from a wide variety of contributors in order to realize a new idea. Crowdfunding can be done in many ways – via an open call on one's webpage, through posting a notice in a public place, or through an organized online marketplace called a crowdfunding website. It is these online marketplaces that we are concerned with in this article.

A crowdfunding website is an online marketplace where users – who we call project creators – can post ideas for projects (art projects, businesses, bands, classroom exercises, etc.) and other users – donors – can contribute small amounts of money toward funding those projects. By aggregating large numbers of small-amount contributions, these sites enable project creators to raise funding for a wide variety of projects.

Crowdfunding websites are a relatively new phenomenon; Modern Internet technologies enable crowdfunding because they permit low-cost, centralized advertising of project ideas, secure and trustworthy contribution to those projects, and simultaneous solicitation by a large number of projects. Recent years have seen the rise of a wide variety of crowdfunding websites [8], including Kickstarter (which funds creative projects), IndieGoGo (which funds a wide variety of ideas and new businesses), Spot.U.s (which funds investigative journalism), Sellaband (which funds musicians) and Donors Choose (which funds K–12 classroom projects).

Project creators sometimes solicit purely altruistic donations, but often creators offer something in exchange for donation. These exchanges fall into three categories: rewards, product pre-orders, and equity. Rewards are usually small tokens of appreciation that are attached to different levels of contribution, such as an acknowledgement at the end of a crowd-funded movie, or thank you letters from students who used crowd-funded supplies in their classroom. Product pre-orders are contributions that effectively "pre-order" a product, to be delivered at a later date. Product pre-orders are an effective

way to raise initial capital to create a new product. Equity is currently relatively rare as a crowdfunding exchange in the US, though the recent JOBS act (H.R. 3606, 2012) explicitly legalized using crowdfunding for equity exchanges. The various crowdfunding websites have differing levels of support for each of the types of exchange.

In addition to the explicit exchange, crowdfunding almost always includes some amount of public good being funded. For product pre-orders, the public good being funded is the ability to purchase a product on the open market, as the product won't exist without startup capital. For example, the TikTok iPod Nano wristwatch that was famously funded on Kickstarter needed a minimum order size for the factory to be able to produce the watches. In more altruistic marketplaces like Donors Choose or Spot.U.s, the public good is a more traditional public good such as news or education. Indeed, Belleflamme et al.[5] argue that in the absence of any public good aspect, crowdfunding theoretically "yields exactly the same outcome as seeking money from a bank or equity investor." In other words, the primary reason that someone would prefer crowdfunding as a method of raising capital is because of the public good aspects of what is being produced.

Project creators choose to use crowdfunding websites to raise capital for a number of reasons. Gerber et al. [7] identify six reasons beyond the obvious reason of needing capital: expanding awareness of one's work, maintaining creative control over the outcome, measuring supporter interest in the project, making longer-term connections with customers, gaining approval and confidence in the project, and to learn more about business and fundraising. Belleflamme et al. [5] find that, at least when offering product pre-orders via crowdfunding, it is the high-value customers that are most likely to contribute to a crowdfunding project, thus suggesting that the connections made with customers through crowdfunding are likely to be high-value.

Many crowd-funded projects are discrete goods, for which there is no or little value if insufficient funds are raised. Products cannot be manufactured without enough capital to fund a production run. Films cannot be made without necessary equipment. Classroom projects cannot be completed without supplies for all students. The design of many crowdfunding sites may be appealing to these types of projects because they are focused around achieving target amounts. There has been some speculation that these type of discrete goods that can't use partial funds are better for sites that use a return rule, while more continuously funded goods like operating capital for a charity are better funded on sites that use a direct donation model.

Projects posted on crowdfunding websites have a surprisingly high likelihood of being funded. 43–47% of projects on Kickstarter are fully funded [12, 10]; 43.5% of projects on Spot.U.s are fully funded [9], and almost 70% of projects on Donors Choose are fully funded [20]. Mollick [12] uses data from Kickstarter to find that projects are more likely to be successful if they ask for smaller amounts of money, they are created by someone with a large social network, are higher quality projects to start with (as indicated by the presence of

a video), and fit in with the culture of the city the creator lives in (such as music projects in Nashville, TN or film projects in Los Angeles, CA).

People donate to crowdfunding projects irrespective of geography [1], though friends and family are often the initial donors that provide the first contributions [18]. Using qualitative interviews, Gerber et al. [7] find five motivations for people who contribute to crowdfunding projects: collecting rewards, helping others achieve their goals, supporting like-minded people, being part of a community, and supporting a cause. Confirming this with empirical data, Shin and Jian [18] find that people who find contributing fun, or who like contributing to friends and family contribute the most money to crowdfunded projects.

Complementarities and Coordination

Coordinating the decisions and actions of a large group of people is an important challenge in the design of crowdfunding, as well as many other forms of socio-technical systems. One of the things that complicates coordination in socio-technical systems is the presence of *complementarities* in the preferences or interests of individuals within the crowd. A complementarity occurs when an individual has a preference for some decision or outcome that depends on the preferences or decisions of others. For example, a visitor to a crowdfunding site may prefer a project and want to donate only if there are enough others that donate that the project can be completed. If the project cannot be completed, the visitor may prefer not to donate.

Complementarities can make it difficult for a group to successfully coordinate their actions to produce the best possible outcome. Economists studying Market Design have worked to develop algorithms that effectively coordinate the preferences of many individuals [15]. For example, this work has been used to design an algorithm that matches medical interns to hospital positions, taking into account both the preferences of the hospitals and the interns [16].

Complementarities can disrupt coordination by such algorithms. For example, when couples began graduating from medical school together, the individuals in these couples preferred positions near wherever their partner received a position. This created a complementarity in the preferences of the interns that led to an unraveling of the medical intern market [14]. The algorithm became unable to match interns to positions such that nobody wanted to ignore the decision of the algorithm. That is to say, the algorithm was no longer able to produce *stable* matches of interns to hospitals. However, Roth and Sotomayor [17] demonstrated a solution to this problem that works as long as complementarities are the exception rather than the norm in the market.

Socio-technical systems, including crowdfunding systems, often involve complementarities. For example, a person may only be interested in participating in an online community if she thinks others will also participate. Or a person may not want to provide an answer to a crowdsourced question if he thinks that there are many others who can provide the answer, but if nobody else will answer, he would prefer to answer the

question. Additionally, many socio-technical systems are designed for self-coordination, where users make the best decisions they can about what to do or how to contribute based on the information they have. This is logical given that the number of complementarities in such systems likely makes the algorithmic approach to coordination excessively difficult.

We are interested here in whether the design of the rules of a socio-technical system, in this case a crowdfunding site, can influence the ability of users to self-coordinate. The nature of most crowdfunding projects as threshold goods creates complementarities in the preferences of users of that site. For example, a visitor to the site who sees a project for a new product she is interested in would only want to donate as long as there are enough others who donate that the product can be produced. In order to achieve the goal of producing the product, all donors with similar preferences must coordinate their actions.

Threshold Public Goods

Many crowdfunding projects are examples of threshold public goods as described by Bagnoli and McKee [4]. A public good is anything that benefits everyone (no one can be excluded from benefits) and also does not get completely used up when it is used (i.e. its use is non-rivalrous). Most public goods grow in size and benefit with additional funds; additional funding to the army helps everyone be better protected. Threshold public goods, though, have a minimum threshold of funding necessary to receive any benefit. For example, if a bridge project has insufficient funding, then no bridge is built.

Threshold public goods are frequently studied using experimental economics and game theory. Threshold public goods games ask subjects to make donations to a group account, and if donations exceed a given threshold, all members of the group earn an additional payout. Bagnoli and Lipman [3] prove that in a threshold public goods game, any outcome where the sum of all contributions is exactly equal to the threshold required for producing the good is both a Nash Equilibrium and a Pareto Efficient outcome of the game, suggesting that it can be rational for individuals to contribute to threshold public goods. Unfortunately, other equilibria also exist which are not efficient. Particularly, there exists a strong free-riding equilibrium in threshold public goods games wherein nobody contributes to the good and therefore the good is not created. This multiplicity of equilibria means it is difficult to predict from theory alone what outcomes to expect when a group of individuals has an opportunity to contribute to a threshold public good.

Cadsby and Maynes [6] conducted a series of threshold public goods experiments which varied the size of the threshold for the good, the size of the reward for funding the good, the type of contribution allowable (continuous vs. discrete), and whether or not contributions would be refunded if the threshold was not met. In their experiments, higher rewards and allowing continuous contributions led to higher donations and more frequent funding of the public good. They also found that goods with a higher threshold received less contribution. However, they found that offering refunds when the threshold

was not met increased the likelihood of reaching the threshold, particularly when the threshold was high.

This refund mechanism is analogous to the return rule in crowdfunding sites, so we expect that crowdfunding websites using the return rule will be more successful at generating contributions to projects. But crowdfunding websites are markets for multiple public goods that are simultaneously soliciting contributions from potential donors who visit the site. The presence of multiple goods simultaneously introduces a necessity for coordination that has not been examined in the existing literature on threshold public goods.

A BRIEF THEORY OF CROWDFUNDING SITE DESIGN

Crowdfunding markets are not guaranteed to have a stable match. Consider a market with 3 donors A , B , and C , and three projects 1, 2, 3. Each project has the the same preference: It needs donations from any two donors to succeed, but weakly prefers zero donors to having only one contribute. This is a complementary preference; one donor is not valued unless accompanied by another and therefore donors are complements. Now, assume the donors have the following preferences over projects:

- A: $1 > 2 > 3$
- B: $2 > 3 > 1$
- C: $3 > 1 > 2$

Given the preference against having only one donor contribute, only one project can possibly be matched to donors. However, no matter which project is matched to the donors, the two donors who don't list that project as their top choice have an incentive to defect and fund a different project. This means that no matter which project is matched to donors, the match is not stable. Thus, **crowdfunding websites do not always have a stable outcome**, and certainly there does not exist any mechanism that always produces a stable matching.

The presence of multiple simultaneous threshold public goods also creates new inefficiencies not seen in previous threshold public goods. In the example above, each donor might choose to contribute only to his or her top choice, hoping someone else will follow their lead. In that case, even though 3 people contributed, no project is funded because the donors spread out their donations among too many projects. Therefore, **crowdfunding websites can fail to fund projects even when there are sufficient donations and donor interest**. Designing a crowdfunding website to increase contributions isn't enough; donations need to be *coordinated* to fund as many projects as possible.

EXPERIMENT SETUP

The return rule is an innovation from industry that may or may not help crowdfunding websites overcome these theoretical challenges. To understand how the return rule affects contributions and funding of projects on crowdfunding sites, we simulated a crowdfunding site using a threshold public goods game similar to the games reported by Bagnoli and Lipman [4] and Cadsby and Maynes [6].

The public goods in our game were presented as projects on a crowdfunding site which were each requesting a minimum

Table 1. The payout structure for each participant for each project

Subject	Low Risk	Medium Risk	High Risk	Unfundable
1	200	150	100	50
2	50	200	150	100
3	100	50	200	150
4	150	100	50	200
5	200	100	150	50
6	50	150	100	200

of 400 credits. Unlike previous threshold public goods experiments, our simulation involved multiple projects which simultaneously requested contributions from the public.

The experiment was designed to simulate three important tensions that can exist in a crowdfunding market. First, projects generally need more than one donor to achieve their funding goal. Second, not all projects can be funded and each donor has separate preferences; this simulates the coordination problem faced by the site as a whole. Third, projects differ in the risk of not being funded; donors who must choose between donating according to their preferences and donating according to the funding risk. These three tensions are all sources of potential inefficiency in a crowdfunding market.

Crowdfunding game

The crowdfunding simulation involved six players simultaneously making contributions to four fictional projects. Each player was given an endowment of 150 credits which could be allocated in any way among the four projects. Players could also keep some or all of the credits.

Subjects receive the payout described in Table 1 for each of the projects if that project receives its goal of 400 credits donated. Subjects receive their payout regardless of whether they contributed to the project; this reflects the public goods nature of the projects, and also provides an incentive to free-ride. At the end of the experiment, subjects exchange credits for cash.

Donors on real crowdfunding sites can have many reasons to prefer some projects over others, including an explicit reward, personal relevance, or a feeling of warm glow. In this experiment, we use the payouts to represent and induce preferences over the projects. Vernon Smith showed that in an experimental setting, paying subjects different amounts of real money for acquiring fake goods can induce those subjects to prefer one good over another, with the higher-payout good being more highly preferred [19]. His Induced Value theory [19] suggests that payouts such as ours can induce preferences in public goods experiments, even when the differences in payouts are very small. Since we are only interested how donors behave once they have formed their preferences, and not in the process of forming preferences for crowdfunding projects, we chose our payouts based on Induced Value Theory such that subjects will prefer the project they are supposed to prefer by being offered a higher payout for its completion.

All projects requested 400 credits. However, in order to simulate variation in risk between different crowdfunding projects,

Table 2. Details of the four different projects

	Amount Needed	Seed Funding	Probability of Being Funded
Low Risk	400	300	88%
Medium Risk	400	200	72%
High Risk	400	100	13%
Unfundable	400	0	0%

The different seed levels induced different levels of risk. In the experiment, the projects were given uninformative names. The probability of being funded is the frequency that each project was funded in the experiment.

we seeded varying amounts into the projects. This created four types of projects; a low-risk project, a medium risk project, a high risk project, and an unfundable project (Table 2). The seed funding did induce different levels of risk, as represented by the actual probability of being funded in Table 2.

Projects vary in riskiness, and donors vary in which projects they prefer. This creates a tension for donors, and the best decision isn't obvious. Even though the collective potential payouts for each project are equal (750 credits), some donors prefer high risk projects and must choose between donating according to their preferences, risking failure, and donating to low risk projects for a lower payout.

Experimental Conditions

The crowdfunding game is played under two conditions. In the direct contribution condition, all donations to projects are final and participants lose their contribution regardless of whether the project to which it was made reaches the funding threshold. In the return rule condition, any contributions made to a project that does not reach its funding threshold are refunded to participants.

These two conditions result in two different payoff functions of the game. P is the set of all contributions (including contributions of zero) made by a player to each project. F is the set of contributions made by the player to projects which were funded. R is the set of rewards earned by the player from funded projects. ϵ is the initial endowment of 150 credits given to each player. Equation 1 gives the function for the payoff (π) in the direct donation condition, in which donations to all projects are subtracted from the total payout. Equation 2 gives the function under the the return rule, where only contributions to funded projects are subtracted.

$$\pi = \epsilon + \sum_{r \in R} r - \sum_{p \in P} p \quad (1)$$

$$\pi = \epsilon + \sum_{r \in R} r - \sum_{f \in F} f \quad (2)$$

Participants

We recruited a random sample of undergraduate students from our university through an email from the registrar to participate in the study.

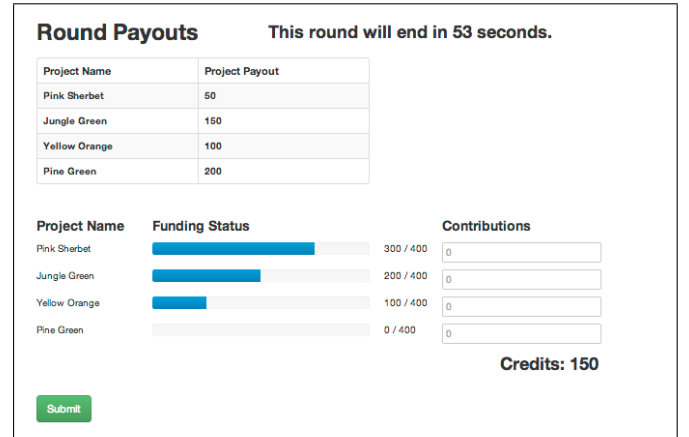


Figure 1. Crowdfunding simulation donation screen

The email promised \$5 for signing up and showing up to the study, and an additional reward based on their performance and the performance of others in the experiment, with an expected average of \$20 for participating. 4% of students contacted responded to the email, and a total of 168 participants were scheduled to participate in 14 experimental sessions (8 sessions under the return rule and 6 under direct donation).

Participants were 48% female with a median age of 20. This is equivalent to the demographics of the undergraduate population at our university from which we recruited participants, and so we do not suspect that our recruiting methods were biased.

Only 18% of participants had ever visited a crowdfunding website before, and only one participant had ever created a project on a crowdfunding site before. Because participants were generally unfamiliar with crowdfunding, we do not suspect that previous experience with crowdfunding sites would have influenced participants' behavior in the game.

Experiment Procedure

Each experimental session involved 12 players playing under a single experimental condition (return rule or direct donation). Each session consisted of 18 rounds of the game. In each round, the 12 players were randomly divided into two groups of six, and players were never aware of which other players were in their group in a given round. This follows the procedure used by [2] to prevent players from learning their partners' strategies in a public goods game.

In each round, participants were given a different set of preferences (see Table 2). Over the 18 rounds, each participant was assigned each of the six unique preferences three times.

Figure 1 shows the interface for the crowdfunding simulation. The interface provided participants with all necessary information and enforced the rules and timing of the game. All donors make their donations simultaneously. To help players learn the game, a detailed summary was provided at the end of each round that displayed the outcome of the round and a detailed calculation of the player's payout (see Figure 2).

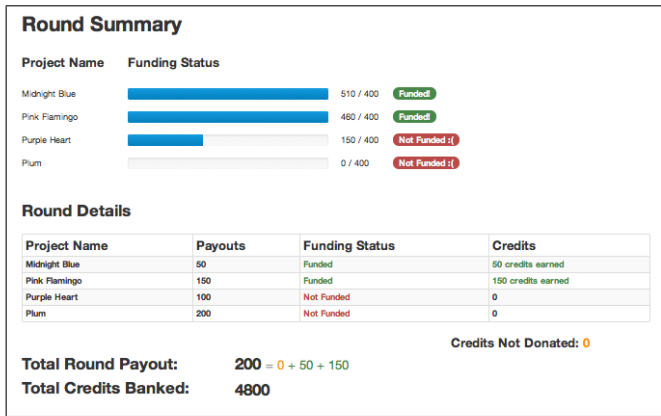


Figure 2. Crowdfunding simulation round summary screen

Table 3. Perceptions of risk for each project

Project	Return Rule		Direct Donation	
	Mean	(SD)	Mean	(SD)
Low risk	1.97	(1.36)	1.64	(0.83)
Medium risk	3.43	(1.00)	3.33	(0.80)
High risk	4.82	(1.08)	5.40	(0.80)
Unfundable	5.89	(1.99)	6.82	(0.78)

Manipulation Checks

After the game, we verified participants' understanding of the rules of the game in a short questionnaire.

20% of participants gave at least one incorrect answer. 10% did not understand the return rule (or lack of return rule). 12% did not know how many other donors were in their group. 7% of participants did not understand that others had different preferences for projects. Overall, participants appeared to understand the rules of structure of the game at the end, but there was not perfect comprehension of the game and it is possible that some results are skewed as a result.

We further verified in the questionnaire that participants interpreted the variation in seed credits among the four projects as variation in risk as was intended. Participants clearly interpreted projects with fewer seed credits as being more risky in both conditions (see Table 3). Thus, we are satisfied that the seed credits created the desired perception of risk for the projects.

RESULTS

Return Rule Increases Donations to High-Risk Projects

Table 4. Donations Decisions

	Direct Donation	Return Rule
Donation Amount	91	122
Donated Budget	27%	59%
# Projects Donated To	1.87	2.37

Donation amount is the average number of credits donated in a round by a subject. Donated budget is the proportion of donations decisions where the subject donated all 150 credits they were allocated.

Table 6. Donations Correlations

	Direct Donation	Return Rule
Payouts and Donations	0.34	0.50
Risk and Donations	0.41	0.16

Donor's average correlation between donations to projects and either payouts or riskiness of projects. Wilcoxon rank sum tests found that both differences between conditions were statistically significant ($p < .001$).

Logically, the return rule should reduce the risk to donors involved in contributing to crowdfunding projects. If a project doesn't receive enough funding, then your money is returned and you are no worse off than if you hadn't donated. Taken to its logical extreme, this argument predicts that subjects will donate all of their 150 credits in the return rule condition.

On average, people donated approximately 30 credits more under the return rule (Table 4). However, subjects in the return rule condition only donated an average of 121 out of their 150 credit budget, and only 60% of the time did they donate their full budget.

Even under the return rule, there is some risk; credits donated above what was requested are effectively lost. Also, because the return rule reduces risk, high-risk projects will seem less risky. For these two reasons, we expected more and larger donations to higher-risk projects.

The return rule had very little effect on the low-risk and medium-risk projects. However, it over doubled the total contributions made to the high-risk and unfundable projects, and over doubled the number of people donating to those projects (Table 5). When people donate, however, they donate similar amounts in both the direct donation and return rule conditions.

We also found that each donor on average donated to about 0.5 more projects under the return rule. The return rule caused donors to spread their donations out among a larger percentage of projects on the site. A Wilcoxon rank sum test confirmed that this difference was statistically significant.

The Return Rule Favors Individual's Preferences

The experiment created a tension for each donor between their personal preferences for projects (as induced by the payouts) and the amount of risk involved in donating (as manipulated by the seed credits). This tension was particularly evident when subjects had higher preferences for riskier projects.

We calculated the average correlation coefficient for each donor between his or her payouts for each project and donations to each project. We wanted to see if donors would generally donate more to projects with higher payouts and less to projects with lower payouts, and whether this would differ for return rule donors from direct donation donors. Table 6 describes these results. We found that under the return rule, there was a higher correlation between payouts and donations than under direct donation.

Similarly, for each donor we calculated the average correlation coefficient between donations and the amount of risk in

Table 5. Contributions to Each Project

	Total Contributions		Number of Donors		Average Donation	
	Direct Donation	Return Rule	Direct Donation	Return Rule	Direct Donation	Return Rule
Low Risk	166	161	4.07	4.13	41.5	39.5
Medium Risk	245	237	4.39	4.55	56.3	* 52.9
High Risk	102	* 226	1.75	* 3.69	62.7	62.2
Unfundable	33	* 108	0.71	* 1.84	48.1	* 60.3

Statistically significant differences ($p < 0.01$) from a Wilcoxon rank sum test between experimental conditions are marked with a *.

projects, as operationalized by the number of seed credits. Higher donations to projects with more seed credits would indicate stronger risk-based donating behavior. There was a higher correlation between risk and donations under direct donation than under the return rule.

Together, these findings show that direct donation leads crowdfunding donors to coordinate; they base their donation choices more on the relative risk between projects than on their actual preferences for projects. On the other hand, the return rule leads users to donate more in line with their own preferences with less adherence to the riskiness of projects. The riskiness of projects is a shared characteristic— all donors see the same risk in projects. Preferences are an individual characteristic— each donor has unique payoffs. Thus, by emphasizing the importance of the projects shared characteristic, direct donation leads users to coordinate around that characteristic.

Site-Level Outcomes are Mixed

Crowdfunding by definition aggregates the behavior and money of many individuals to fund multiple projects. As we illustrated above, simply receiving more donations doesn't necessarily lead to more projects being funded. The aggregate effects of the return rule on the marketplace are complicated, but important for site designers. In Table 7, we evaluate four market-level effects of the return rule.

Donors funded an average of 0.17 more projects per round under the return rule than under direct donation. While this difference is statistically significant, it is nonetheless a fairly modest increase given that contributions to the higher risk projects doubled under the return rule. As a result, there was a loss in efficiency under the return rule: Donations provided an 83% return on investment to donors for each credit donated under the return rule, compared to a 126% return on investment for direct donation.

This lack of efficiency can be characterized in two other ways. As noted above, donors individually *spread* their donations to more projects under the return rule. We operationalize spread using the Gini coefficient. The Gini coefficient is normally used to measure income inequality in a population, but here we use it to measure donation inequality across projects. A market is more spread out if the contributions to projects are more equal. Collectively, we found that the aggregated donations were more evenly spread out among projects under the return rule; the Gini coefficient is much smaller (0.16) under the return rule than under direct donation.

Table 7. Site-level outcomes

	Direct Donation	Return Rule
# Projects Funded	1.63	1.80
Gini Coefficient	0.39	0.23
Return on Investment	126%	83%
Uncoordinated Rounds	30%	60%

For the Gini coefficient, 0 means equal contributions to all projects, and 1 means all of the money is concentrated on a single project. Uncoordinated rounds are the percentage of rounds where the total donations to unfunded projects were enough to fund at least one more project. All comparisons are statistically significant ($p < 0.02$) from a Wilcoxon rank sum test.

This loss in efficiency when a site uses the return rule is due to a lack of coordination among donors. Under the return rule, in 60% of all rounds, the donations made to unfunded projects were sufficient to fund at least one more project had the donors coordinated their donations on a single project rather than spreading them among multiple projects. Under direct donation, only 30% of rounds could have funded additional projects. This lack of coordination offsets much of the gain made by the return rule through increased total donations.

The overall effect of the return rule is difficult to determine. First, subjects might be donating anew to higher risk projects, but not changing their low-risk donations. This basically means that the return rule only affects the likelihood of donation, but not the amount of donation. Second, subjects may spread all their money out across more projects, but choose to increase the total donated, offsetting the changes in donation amounts to existing projects. Both effects of the return rule — increased total donation amount and decreased donations due to spread — affect donation amounts. Our experiment is not able to distinguish between these two explanations.

Without the Return Rule, People Learn to Only Fund Low Risk Projects

Our experiment asked subjects to play this simulated crowdfunding game repeatedly. This allows us to examine how subjects learned what strategies were effective. Subjects' ability to learn over repeated rounds and collectively improve performance is another indicator of coordination on a crowdfunding site.

The return rule condition did not exhibit much learning; the number of donors to each project and the average amount

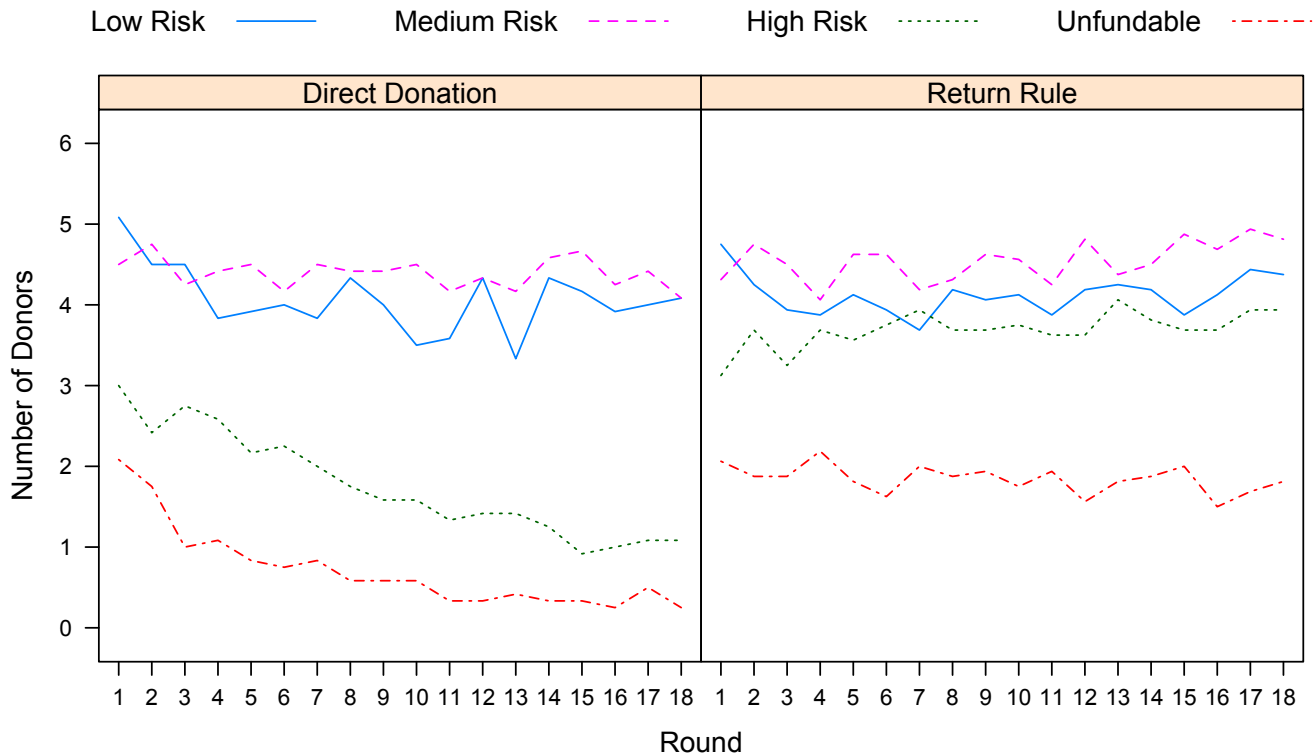


Figure 3. Number of donors who contribute to each project, by round. The left panel is for the Direct Donation condition. The right panel is for the Return Rule condition.

contributed to each project remained fairly stable over the course of the 18 rounds. The right sides of Figures 3 and 4 show little trend over time. Subjects did appear to learn to contribute more to the high-risk project, with the average contribution increasing from 25.8 in the first round to 40.6 in the last.

However, the direct donation condition did exhibit noticeable learning. Subjects appear to be learning not to donate to the high risk and unfunderable projects. The number of subjects donating to the high risk project averages 3.0 in the first round, and decreases to 1.08 in the last round. Likewise, the number of subjects donating to the unfunderable project averages 2.08 in the first round, and decreases to 0.25 in the last round. Average contribution amount shows a similar decrease.

We conclude that the direct donation condition helped donors learn over time to coordinate their donations on only those projects with a high likelihood of being funded. Under the return rule, donors did not learn to coordinate as the experiment went on and generally maintained a consistent behavior throughout.

DISCUSSION

The main goal of this study was to examine the effect that the all-or-nothing return rule has on a crowdfunding marketplace, relative to a direct donation model. We found that this rule has two distinct effects on the way individual donors make

contributions to projects. These two individual-level effects lead to two important market-level outcomes that influence the efficiency of the market in opposite directions.

First, under the return rule, individual donors are more willing to contribute money to projects, and in particular high-risk projects, leading to an overall increase in donations. Second, the available donations are more spread out over more projects. Increased spread in the market makes it more difficult for each project to achieve its goal, and can potentially lead to fewer projects being funded. These two effects appear to approximately offset in our experimental setting, leading to a very small efficiency gain when using the return rule.

When donors use a direct donation system, they seem to learn to coordinate on low-risk projects that are highly likely to be funded. This implicit coordination makes it relatively easy to fund low-risk projects and also leads to a higher return (in terms of the value of the newly produced public good) on each donation. In this experiment, the risk inherent in the projects, largely due to complementarities inherent in projects that need a certain amount of funding to succeed, appears to have served as a coordination device.

On the other hand, in a return rule system, there was much less coordination among donors. Donors seem to feel like it is safe to donate to any project, since they get their money back if the project doesn't achieve its funding goal. The return rule, therefore, eliminated much of the incentive to coordinate with other donors about which projects are being funded. This is

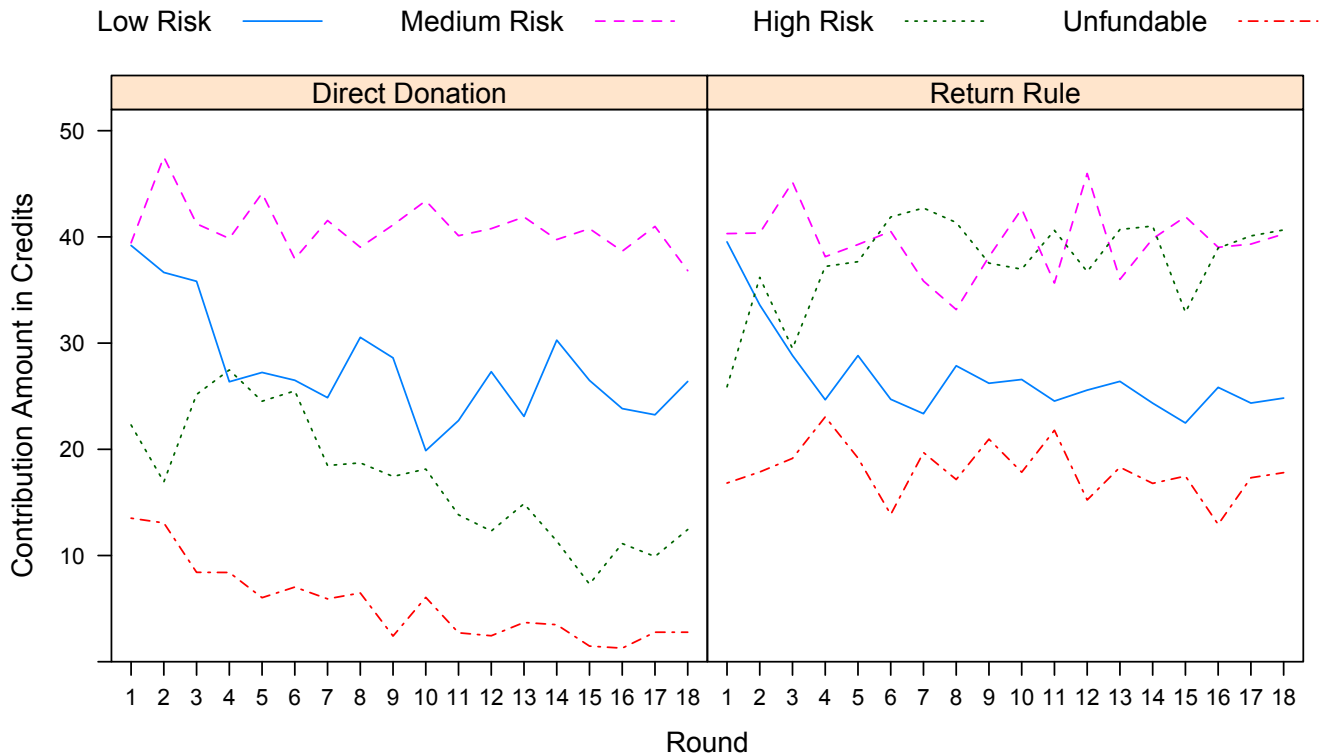


Figure 4. Average contribution to a project by round. The left panel is for the Direct Donation condition. The right panel is for the Return Rule condition.

why we saw an increased spread in donations; donors were focused more on the projects they preferred and less on solving the coordination problem.

These findings have implications for market designers working to design better crowdfunding systems. First, they suggest that there isn't a clear answer whether using an all-or-nothing return rule or using a direct donation rule is better. If the marketplace has a large number of projects simultaneously seeking support, the return rule has the risk that donors will spread out their donations, leading to many donations but low funding rates for projects. On the other hand, if projects are largely undifferentiated (so there is nothing to coordinate on) or if many projects are seen as high risk, then the increased donations that result from the return rule might enable the market to fund more projects.

Additionally, if the goal of the marketplace is to fund high-risk high-reward projects, the return rule is more likely to succeed. If the goal of the marketplace is to weed out high-risk projects and focus donor attention on projects with a high likelihood of success, using a direct donation mechanism might work better.

On a crowdfunding website, whether a project is able to successfully raise the funds it needs is a complex proposition. Mollick [13] suggests that project success is largely a function of the team of people working on the project, much like how venture capitalists fund traditional startup companies. Both Shin and Jian [18] and Gerber et al. [7] suggest that how

hard a project's creator works to fund the project also plays an important role. Our results suggest that the marketplace also has a surprisingly strong influence; other projects available on the crowdfunding website can draw money away, and the structure of the website (return rule vs. direct donation) can also affect the likelihood of being funded.

Increasingly we are seeing crowdfunding websites that offer project creators a choice — on the same website, they can either offer their project as a direct donation project or as a return rule project. For example, IndieGoGo recently changed from a direct donation marketplace to one where each project gets to choose. Our results would suggest that higher risk projects should choose to be offered under a return rule. On the other hand, projects that can benefit from coordination and a strong word-of-mouth should probably choose direct donation, as that provides a stronger incentive for donors to coordinate and spread the word.

Our work also highlights an interesting tension that appears in the design of many groupware systems. There are design options like the return rule that can help people work more independently; these designs are a double-edged sword. By making the work easier, it encourages people to do more work overall. But they also remove the incentive to coordinate their work with others, which may lead to groups failing because the work didn't fit together well enough. This is sometimes seen on Wikipedia, for example. Wikipedia makes independent editing of articles easy, which encourages editors to work. However, sometimes that means that the articles feel

disjointed and are written in different styles or tones, or are self-contradictory [11].

Limitations

There are a few important limitations to this study that should be considered. First, not all crowdfunded projects are necessarily threshold goods. Some projects could still benefit from even insufficient donations, and may be better classified as continuous goods. However, we believe such projects are generally the minority of crowdfunded projects. If continuous goods represented a strong majority of crowdfunded projects, we would expect that there would be little reason to even consider a return rule mechanism.

A second limitation is that the simultaneous nature of donations in our experiment differs from real crowdfunding sites, where donations are made sequentially. We did this to adhere to the public goods game framework and to avoid an additional variable that would have added complexity to the experiment. However, we believe this is an important variable that demands further research. The risk involved in donating to a crowdfunding project may be different depending on how much time is left for the project to meet its goal. It may be riskier, for instance, to donate on the first day than the last because the likelihood of the target being successful is less clear. Thus, the return rule could have the effect of soliciting more donations early in a project's timeline. We hope to explore this question in future extensions of this research.

CONCLUSION

Crowdfunding websites like Kickstarter, IndieGoGo, Donors Choose, and Spot.Us pose an interesting challenge for system designers. They are matching marketplaces, matching people who have ideas for projects with donors willing to contribute money to those projects. But they have a strong inherent complementarity in the market structure: multiple donors are needed to fund a project, and the project often cannot succeed without full funding. Complementaries such as this cause theoretical problems for market designers; matching markets with complementarities frequently do not have stable equilibria, and regularly have problems with market failure or market inefficiency.

There are two major strategies that existing crowdfunding websites have used to deal with this problem: allow direct donations to projects (used by IndieGoGo), or use an all-or-nothing return rule that returns donations if a project doesn't achieve its goal (used by Kickstarter). We conducted an experimental study of these two mechanisms to better understand how these strategies enable matching markets with complementarities to function.

We find that each of these two mechanisms has distinct strengths and weaknesses. Direct donations provide an incentive for donors to coordinate to solve the complementarity problem, and are more likely to fully fund projects even when total donations in the marketplace are small. The return rule encourages people to donate to more projects and to higher risk projects, but has the potential of spreading out contributions among too many projects, causing few of them to reach their goal. Both mechanisms appear to be viable means for

partially solving the complementarity problem in matching marketplaces.

ACKNOWLEDGEMENTS

The authors would like to thank Kyle Safran for his help designing and running this experiment. The authors would also like to thank Jan Boehmer and the whole BITLab Lab Group for feedback. This material is based upon work supported by the National Science Foundation under Grant No. CCF-1101266.

REFERENCES

1. Agrawal, A., Catalini, C., and Goldfarb, A. The geography of crowdfunding. NBER Working Paper #16820, February 2011.
2. Andreoni, J. Why free ride?: Strategies and learning in public goods experiments. *Journal of Public Economics* 37, 3 (1988), 291 – 304.
3. Bagnoli, M., and Lipman, B. L. Provision of public goods: Fully implementing the core through private contributions. *The Review of Economic Studies* 56, 4 (1989), pp. 583–601.
4. Bagnoli, M., and McKee, M. Voluntary contribution games: Efficient private provision of public goods. *Economic Inquiry* 29, 2 (April 1991), 351—366.
5. Belleflamme, P., Lambert, T., and Schwienbacher, A. Crowdfunding: Tapping the right crowd. Tech. Rep. 1578175, Social Science Research Network (SSRN), April 2012.
6. Cadsby, C. B., and Maynes, E. Voluntary provision of threshold public goods with continuous contributions: experimental evidence. *Journal of Public Economics* 71, 1 (1999), 53 – 73.
7. Gerber, E., Hui, J., and Kuo, P.-Y. Crowdfunding: Why people are motivated to participate. Tech. Rep. 12-02, Northwestern University Segal Design Institute, 2012.
8. Greenberg, M., and Gerber, E. Crowdfunding: A survey and taxonomy. Tech. Rep. 12-03, Northwestern University Segal Design Institute, 2012.
9. Jian, L., and Usher, N. Crowd-funded journalism. *Journal of Computer Mediated Communication* (2013). (Forthcoming).
10. Kickstarter. Statistics. <http://www.kickstarter.com/help/stats>, December 2013.
11. Kittur, A., and Kraut, R. Harnessing the wisdom of crowds in wikipedia: Quality through coordination. In *Proceedings of Computer Supported Cooperative Work (CSCW)*, ACM Press (New York, New York, USA, 2008).
12. Mollick, E. The dynamics of crowdfunding: Determinants of success and failure. Tech. Rep. 2088298, Social Science Research Network (SSRN), July 2012.

13. Mollick, E. Swept away by the crowd? crowdfunding, venture capital, and the selection of entrepreneurs. Tech. rep., Social Science Research Network (SSRN), March 2013.
14. Roth, A., and Xing, X. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *American Economic Review* 84 (1994), 992–1044.
15. Roth, A. E. The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. *Econometrica* 70, 4 (2002), 1341–1378.
16. Roth, A. E., and Peranson, E. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review* 89, 4 (September 1999), 748–780.
17. Roth, A. E., and Sotomayor, M. A. O. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1990.
18. Shin, J., and Jian, L. Driving forces behind readers' donation to crowd-funded journalism: The case of spot.us. Working paper, University of Southern California, 2012.
19. Smith, V. L. Experimental economics: Induced value theory. *American Economic Review* 66, 2 (May 1976), 274–279.
20. Wash, R. The value of completing crowdfunding projects. In *International Conference on Weblogs and Social Media (ICWSM)*, AAAI Press (July 2013).