Index Policies for Campaign Promotion Strategies in Reward-based Crowdfunding

Chenguang Wang

c.wang7@outlook.com

Loughborough Business School, Loughborough University, Leicestershire LE11 3TU, United Kingdom

Dong Li

dong.li@lancaster.ac.uk

Department of Management Science, Lancaster University, Lancaster LA1 4YX, United Kingdom

Baibing Li, corresponding author

b.li2@lboro.ac.uk

Loughborough Business School, Loughborough University, Leicestershire LE11 3TU, United Kingdom

Index Policies for Campaign Promotion Strategies in Reward-based Crowdfunding

Abstract

Reward-based crowdfunding plays a crucial role in fundraising for start-up entrepreneurs. Recent studies, however, have shown that the actual success rate of fundraising projects is surprisingly low across multiple crowdfunding platforms. This paper considers crowdfunding platforms' decision-making of selecting projects to highlight on their homepage to boost the chance of success for projects, and investigates promotion strategies aiming at maximizing platforms' revenue over a fixed period. We characterize backers' investment decisions by a discrete choice model with a time-varying coefficient of herding effect, and formulate the problem as a stochastic dynamic program, which is however computationally intractable. To address this issue, we follow the Whittle's Restless Bandit approach to decompose the problem into a collection of single-project problems and prove indexability for each project under some mild conditions. We show that the index values of the proposed index policy can be directly derived from the value-to-go of each project under the non-promotion policy, which is calculated recursively offline with a linear-time complexity. Moreover, to further enhance the scalability we develop a closed-form approximation to calculate the index values online. To the best of our knowledge, this work is the first in the literature to develop index policies for campaign promotions in reward-based crowdfunding. It is also the first attempt to provide indexability analysis of bi-dimensional restless bandits coupled by not only resource but also demand. Extensive numerical experiments show that the proposed index policies outperform the other benchmark heuristics in most of the scenarios considered.

Keywords: Crowdfunding, Dynamic programming, Restless bandits, Stochastic process

1. Introduction

Reward-based crowdfunding, an innovative online financing alternative, has begun to play a crucial role in the fundraising of start-up entrepreneurs. In recent years, many online crowdfunding platforms, e.g., Kickstarter and Indiegogo, have emerged. These platforms bridge the gap between entrepreneurs (fundraisers) who initiate a fundraising campaign for the development of their new products/technologies and the individual investors (backers) who look for investment opportunities. As an emerging and accessible means of entrepreneurial financing and marketing, crowdfunding sits right at the interaction of multiple disciplines and has attracted considerable attention in finance (e.g., Lee and Parlour 2022), marketing (e.g., Zhang and Tian 2021), innovation (e.g., Crosetto and Regner 2018), and operations management (e.g., Du et al. 2020).

A typical crowdfunding campaign or project, terms that are used interchangeably henceforth, includes the following components: the information about the product or technology for crowdfunding, a pre-specified funding goal, a set of pledge/reward options that backers may choose to buy, and the length of the campaign duration. Many platforms use an all-or-nothing (AoN) scheme, where a project is successful only if its funding goal has been reached by the end of the campaign; otherwise, the project fails and all the funds raised are returned to the backers. Platforms profit from successful campaigns by taking commissions as a percentage of the total funds raised. No fees are charged for failed projects.

Despite the popularity of crowdfunding, the success rate of fundraising campaigns is surprisingly low in practice, leading to low revenue for the platforms. For example, Clifford (2016) shows that 69-89% of projects failed to reach their funding goals. The issue of a low success rate has attracted a wide range of researches probing the problem, both empirically and analytically. Most previous researches focus on the factors of individual campaigns, such as fundraisers' experience and expertise (Zhou et al., 2018), funding goals (Kim et al., 2020), duration (Zhang et al., 2023), reward options (Du et al., 2019), and information descriptions such as readability or word, picture, video counts (Liang et al., 2020). Besides these campaign-specific factors, a crowdfunding platform may boost the chance of success of a campaign by highlighting it on the platform's homepage, which is referred to as homepage promotion in the crowdfunding literature (Varma et al., 2021; Song et al., 2022). Due to the limited promotion space, however, an important question faced by the crowdfunding platform is how to choose the right project to promote at the right time in order to achieve the overall maximal expected revenue for the platform.

In this study, we investigate the mathematical modeling of the reward-based crowdfunding process and develop index policies for platforms' campaign promotion strategies, whereby a limited homepage promotion space is dynamically allocated to multiple competing crowdfunding projects such that the total expected revenue is maximized under the AoN scheme. To properly model the crowdfunding process, we propose a random utility function to characterize backers' perceived value of each project which includes the following attributes: (a) the attraction of the projects (Gerber and Hui, 2013); (b) whether a project is currently being promoted on the homepage (Song et al., 2022); (c) the herding effect such as the influence of the amount already pledged on future backers' decisions (Xiao et al., 2021); and (d) the side effect, i.e., market competition or choice overload (Chan et al., 2021; Du et al., 2019). Unlike many in the literature that assume a constant herding effect over time, we allow a more general setting where the herding effect may vary over time. Indeed, as Dai and Zhang (2019) suggest, when backers adhere to the utility-maximization rule and base their backing decisions on project quality and likelihood of success, it is reasonable to anticipate an accelerated fund collection process later in the crowdfunding campaign. To characterize backers' decision-making behavior under promotion, we adopt a Multinomial Logit (MNL) model to describe their decision-making regarding which project to invest. The MNL model is widely used to describe decision-making processes in revenue management (Akçay et al., 2010; Yang et al., 2016; Du et al., 2016), as well as in the recent crowdfunding literature (Nosrat, 2022; Weinmann et al., 2023), where individuals choose a single item from a set of alternatives based on their relative values represented by their utility functions.

We then model the problem as a finite horizon discrete-time Markov decision process and formulate it as a stochastic dynamic program. Mathematically this problem can be solved by backward induction, but only for small instances due to its computational intractability. In this paper, we address this problem by developing an index policy approach. Specifically, we follow the Whittle's Restless Bandit (Whittle, 1988) framework and use Lagrangian relaxation to decompose the original problem into a collection of single-project problems, each of which has two dimensional states (the funds already raised and the remaining campaign time). We characterize the structure of the optimal policies and prove indexability for these problems when the campaign duration is sufficiently long. Moreover, we demonstrate that the index values can directly be calculated from the value-to-go of each project following the non-promotion policy under which the project is never being promoted. This algorithm has a linear-time complexity in the number of states, where the value-to-go for all states is calculated based on a recursive manner of dynamic programming. To further enhance the scalability, we also develop a closed-form approximation to the index values. Two index policies are then developed, with one based on the exact and the other on the approximate index values. Both index policies promote the project with the largest index value at each decision epoch.

1.1. Related Literature

This research follows a couple of existing research strands in the literature which are briefly outlined below.

A dynamic strategy of promoting projects on a platform's homepage is proposed by Varma et al. (2021) so as to enhance the projects' success rate. They focus on two projects with two actions, i.e., either to continue promoting the current project or to switch to promote the other, where the switch of promotion is irreversible. They do not consider the herding effect. The issue of homepage promotions is also considered by Song et al. (2022) who investigate philanthropic (i.e., donor-based) crowdfunding. They take a statistical approach and estimate the probability of awareness of the promoted projects using real-world data. Besides homepage promotions, some studies consider the use of stimulus for enhancing crowdfunding success. For example, Li et al. (2020b) consider dynamic stimulus policies in their study on the crowdfunding pledging process, via promoting a project on social media such as Facebook or Twitter, to attract more potential backers and therefore accelerate the project's funding completion. However, their study focuses on a single crowdfunding project. More recently, stimulus strategies for a single crowdfunding project are also investigated by Du et al. (2022). Following the idea of the cascade effect (i.e., herding effect), they design one-off stimulus strategies, either being reactive (e.g., upgrade a project's features) or proactive (e.g., limited-time bonus offer), to attract more customers to a campaign. For a comprehensive account on crowdfunding problems and the latest development in the operations management literature, readers are referred to a systematic review by Allon and Babich (2020).

On the methodological side, our paper is related to weakly coupled dynamic programs (Adelman and Mersereau, 2008) that involve a collection of sub-problems interconnected only by some constraints on actions, usually resource/capacity limitations. A common approach to solving such a problem is to relax the coupling constraints via Lagrangian relaxation, which allows the problem to be decomposed and a value function approximation to be obtained from the sub-problems. A greedy control policy is then developed via methods such as one-step lookahead. Lagrangian relaxation has been widely studied in the literature, such as Hawkins (2003), Bertsimas and Mišić (2016) and Brown and Smith (2020). The approximate linear programming, another widely studied approximation, is considered by Adelman and Mersereau (2008) who provide a thorough comparison of the approximation errors between linear programming and Lagrangian relaxation, and the performance of the resulting policies. A further analysis on the tightness of both approximations can be found in Brown and Zhang (2023).

Our paper falls into the literature of restless bandits (RB) problems which are essentially weakly coupled dynamic programs with a single linking constraint in resources among the sub-problems (bandits). The RB scheme is a significant extension of the classical Multi-armed Bandit (MAB) problems, with a wide range of applications such as maintenance (Glazebrook et al., 2005) and risk-averse control (Malekipirbazari and Çavuş, 2024). Its core idea is to sequentially allocate a single indivisible resource to multiple bandits that are stateful and the state transition probabilities are Markovian (Maghsudi and Hossain, 2016). With the classical MAB problem, the state of each bandit changes only when it is in receipt of the resource (active); otherwise it remains frozen (passive). The RB scheme relaxes MAB by allowing the bandits to always evolve, even when they are passive. However, such a generalization comes at a cost: the RB scheme is almost certainly intractable as shown by Papadimitriou and Tsitsiklis (1994). In his seminal work, Whittle (1988) develops an index policy to an RB problem by decomposing it into multiple single-bandit sub-problems via Lagrangian relaxation (that breaks the linkage by the resource) and then calibrating state-dependent index values for individual bandits. These index values can be interpreted as fair charges for receiving resources. The benefit of index policies is that the resource allocation decisions can be made by simply sorting the index values among the bandits. Such a heuristic is shown to have strong performance in many problems (e.g., Niño-Mora 2007), and it converges to the optimal under certain conditions (Weber and Weiss, 1991). However, for this method to work, the bandits need to pass an *indexability* test, a structural property that is required to establish the existence of a solution to the Lagrangian relaxation problem (Glazebrook et al., 2014). This is perhaps the biggest hurdle to the development of Whittle's index policies; some researches have contributed to the development of theories to prove indexability under different conditions; see e.g., Glazebrook et al. (2011).

To date, most studies in the RB literature develop Whittle's index policies for infinite time horizon problems, in which case indices are often simpler, as stated by Whittle. Only a few studies have considered finite horizon problems; see for example Li et al. (2020a) and Graczová and Jacko (2014). In addition, almost all the studies concern the bandits with only one-dimensional states (a benefit of infinite horizon where time is not relevant). As far as we know, the only exception is Graczová and Jacko (2014) who consider the allocation of a limited display area to perishable products. They formulate the problem as a knapsack problem with perishable inventories, and develop index policies following the Whittle's RB framework. Each product or bandit is characterized by two-dimensional states, the remaining inventory and the remaining shelf life. To allow decomposition and ensure their problem is indexable, however, they have made the important assumptions that the demands across different products are independent of each other, and the purchasing probabilities are constant regardless of the state. These assumptions are strong as in real life customers normally make choices among multiple potentially substitutable products, and the purchasing probabilities may change with either the remaining inventory, the shelf life, or both.

In this paper, to realistically describe the crowdfunding process, we model backers' investment choice behaviors with the MNL model, where the backing probability of each project takes into account the substitution effect and is dependent on the funding progress of all the other projects. As a result, the backing probabilities are state-dependent (i.e., they rely on the funding progress), which is an important feature to consider in crowdfunding. This, however, brings in some challenging mathematical complications to the analysis for the structure of the optimal policies and the establishment of indexability. Moreover, the projects are interconnected by not only the resource (promotion space), but also the demand. One of the research challenges for our study, therefore, is to explore whether and how the Whittle's RB framework can be applied to such a problem.

1.2. Contributions

This paper contributes to the literature in the following two aspects. First, this is the first study in the crowdfunding literature to develop a mathematical framework to model the crowdfunding process, with a full characterization for backers' purchasing behaviors described via a discrete choice model. Moreover, we show that our analyses and results still hold for more general state-dependent utility functions, provided that the corresponding choice probability satisfies some general conditions (as shown in Lemma 4.1). To the best of our knowledge, this work is also the first on the optimal homepage promotion strategies in reward-based crowdfunding using the Whittle's RB framework.

Second, this paper contributes to the weakly coupled dynamic program and RB literature by an innovative application of the Whittle's index method to the problem where the bandits are coupled by not only resource but also demand (so the coupling is not *weak* anymore), with the incorporation of an action-dependent and state-dependent discrete choice model into the RB scheme. We decompose this finite time horizon problem into a collection of single-bandit (project) problems and prove indexability of each (two-dimensional) bandit when the campaign duration is sufficiently long. We show that the index values can directly be calculated from the value-to-go of each project under the non-promotion policy, which has a complexity of linear time on the number of states. We also develop a closed-form approximation to the index values to further cut the computational time.

1.3. Outline

The remainder of the paper proceeds as follows. Section 2 is devoted to the modeling of crowdfunding processes with homepage promotions where the problem is formulated as a finite horizon stochastic dynamic program. In Section 3, the original problem is relaxed and decomposed into a number of single-project problems following the Whittle's RB approach. In Section 4, we investigate the structural properties of the single-project problems, show they are indexable under some mild condition, and demonstrate how the index values and a closed-form approximation are derived, leading to the development of two index policies. In Section 5, we develop nine heuristic policies (including one based on Lagrangian relaxation and another based on integer programming approximation) and examine their performance against the proposed index policies in a series of numerical experiments including sensitivity analyses. Finally, Section 6 concludes the paper.

2. Modeling of Crowdfunding Processes for Homepage Promotions

This section introduces the formulation of the crowdfunding problem and develops our model.

2.1. Problem Statement and Backing Probabilities

Consider a collection of J substitutable fundraising projects on a reward-based crowdfunding platform that seek financial investment from potential backers. Each project j $(= 1, \ldots, J)$ has a pre-specified funding goal G_j. Following the AoN scheme, a project is successful only if its funding goal has been reached by the end of its campaign. For simplicity, we consider a common duration for all campaigns from time 0 to T. An extension that allows projects to have heterogeneous start and end campaign times is discussed in Section 5.5.

We suppose that backers visit the platform according to a Poisson process. We discretize the time horizon into sufficiently small time intervals, so that the Poisson process may be approximated by Bernoulli arrivals with an arrival probability $\lambda \in (0, 1)$. Upon arrival at time epoch t, each backer will choose to back one project, say j, with probability $p_{j,t}$, or leave the platform without any purchases. For the latter case, the non-purchase scenario is denoted by $j = 0$. Hence, we have $\sum_{j=0}^{J} p_{j,t} = 1$. Having chosen a project j to support, the backer is presented with a collection of pledge options $\mathcal{R}_j = \left\{r : 1 \leq r \leq R_j, r \in \mathbb{Z}^+\right\}$, and they will choose one of them to purchase with a known probability $F_j(r)$, where without loss of generality, it is assumed that each pledge r is an integer between 1 and the maximum pledge R_j; in practice, pledge amounts are usually multiples of a positive integer, and for those not available we can always set their purchasing probabilities to be zero.

At each time epoch, the platform selects one project to highlight on the homepage, making this project clearly visible to backers, though they may browse other projects as well. Hence, the project is expected to have a higher possibility of being invested during the promotion period. Note the projects that have already reached their funding goals stay on the platform till the end of the campaign. The platform charges fees proportional to the total funds raised from every successfully funded project. The objective for the platform manager is to allocate the promotion space to a project at a time epoch, with the aim of maximizing the overall revenue by the end of the campaign.

To model the backing probabilities, we consider a linear random utility function to characterize backers' evaluation of each project, which is widely used in the literature (Akçay et al., 2010; Yang et al., 2016; Du et al., 2016). It is worth noting that our work is applicable to a much wider family of nonlinear utility functions, as long as the utility function satisfies some mild condition as discussed in Remark 4.1 (see Section 4.1). We consider homogeneous backers for whom their utility function includes the following attributes for each project j: (a) the project's attractiveness $m_j > 0$ representing the overall appealing of the project to backers; (b) the promotion power $\beta_1 > 0$ that measures the boost of valuation if the project is being promoted, with a binary variable a_j used to indicate whether the project is being promoted $\left(a_j = 1\right)$ or not $\left(a_j = 0\right)$; (c) the herding effect, i.e., $\beta_{2,t}\left(1 - g_j/G_j\right)$, that measures the influence on backer's valuation after observing the percentages of the funds already raised at time epoch t, where $\beta_{2,t}$ is a time-varying herding parameter and g_j is the amount of shortfall to reach the funding goal, with $g_j = 0$ representing the funding goal has just been reached and $g_j < 0$ representing more money has been raised beyond the funding goal; and finally (d) the side effect $\beta_{3,J} > 0$ that reflects competition intensity across projects, where the presence of more projects campaigning simultaneously amplifies the competition/choice overload faced by potential backers. Mathematically, $\beta_{3,J}$ is a non-decreasing function of J and it takes value 0 when $J = 1$.

Thus, the backers' perceived valuation of a project j, denoted by its utility $U_{j,t}$, is modeled as below

$$U_{j,t} = z_{j,t}\left(g_j, a_j\right) + \epsilon_j = m_j + \beta_1 a_j + \beta_{2,t}\left(1 - \frac{g_j}{G_j}\right) - \beta_{3,J} + \epsilon_j,$$

where $z_{j,t}\left(g_j, a_j\right) = m_j + \beta_1 a_j + \beta_{2,t}\left(1 - \frac{g_j}{G_j}\right) - \beta_{3,J}$ is the systematic component of the utility, and ϵ_j is the stochastic component representing the unobserved attributes of the project. The above utility function shows that a project is valued higher when it is more attractive, being promoted, or has more funds already been raised, whereas it is valued lower when it is faced with more competition from the other projects. Note that when a project is over funded (i.e., $g_j < 0$), it is still appealing to future backers because the funds can be used to enhance the final product, such as adding more songs to an album or including extra elements in a game (Kickstarter, 2024).

Assuming backers are utility maximizers at each time epoch, and the random variables ϵ_j are independently and identically distributed and follow a standard Gumbel distribution with zero mean, we may calculate the backing probability for project j, i.e., $p_{j,t}(\mathbf{g}, \mathbf{a}) = \Pr(U_{j,t} \geq U_{k,t}, \text{for all } k \neq j)$, leading to the following Multinomial Logit model (see, e.g., Li (2011) and references therein):

$$p_{j,t}(\mathbf{g}, \mathbf{a}) = \begin{cases} \frac{\exp\left(z_{j,t}\left(g_j, a_j\right)\right)}{1+\sum_{k=1}^{J} \exp\left(z_{k,t}\left(g_k, a_k\right)\right)}, & 1 \leq j \leq J, \\ \frac{1}{1+\sum_{k=1}^{J} \exp\left(z_{k,t}\left(g_k, a_k\right)\right)}, & j = 0. \end{cases} \tag{1}$$

where $\mathbf{g} = (g_1, \cdots, g_J)$ and $\mathbf{a} = (a_1, \cdots, a_J)$ represent a vector of project shortfalls and a vector of promotion indicators, respectively.

2.2. The Model

The crowdfunding problem discussed in the previous subsection may be formulated as a discrete-time Markov decision process. The decision epochs are discrete time points $t, 0 \le t \le T-1$. We define the state at each epoch as the vector $\mathbf{g} = (g_1, \cdots, g_J)$, where g_j is the current amount of shortfall to the funding goal for project j. Let Ω_t be the state space at time epoch t given by $\Omega_t = \left\{\mathbf{g} : G_j - tR_j \le g_j \le G_j, 1 \le j \le J\right\}$, where $G_j - tR_j$ represents the maximal reachable state at time epoch t for project j. Thus, $\Omega_0 = \{(G_1, \cdots, G_J)\}$ and $\Omega_t \subset \Omega_{t'}, \forall t < t'$. Note that, as time is an implicit dimension of the state, we explicitly express the state vector $\mathbf{g}$ at time epoch t as $\mathbf{g}(t)$ when necessary; otherwise, we simply use $\mathbf{g}$ for brevity.

At each decision epoch, at most one project is chosen for promotion. The collection of admissible actions is denoted as $\mathcal{A}$ and is given by

$$\mathcal{A} = \left\{(a_1, a_2, ..., a_J) : a_j \in \{0,1\}, \sum_{j=1}^{J} a_j \le 1\right\}.$$

Thus, if an action $\mathbf{a} = (a_1, a_2, ..., a_J)$ is taken at a decision epoch with state $\mathbf{g}$, and a backer chooses to back project j with a purchase of pledge option r, then the state transits to $\tilde{\mathbf{g}} = \mathbf{g} - r\mathbf{e}_j$ with probability $\lambda p_{j,t}(\mathbf{g,a}) F_j(r)$, where $\mathbf{e}_j$ is a J-dimensional vector that takes value 1 on the j-th component and zero elsewhere. If there are no arrivals or no purchases, the system remains unchanged with the probability $(1-\lambda) + \lambda p_{0,t}(\mathbf{g,a})$. We then formulate the problem as a finite-horizon stochastic DP. Define the value function $V_t(\mathbf{g})$ as the maximum expected total funds raised from time epoch t onward, given the system occupies state $\mathbf{g}$ at time t. Then, the value function satisfies the following Bellman equation:

$$V_t(\mathbf{g}) = \max_{\mathbf{a} \in \mathcal{A}} \left\{\lambda \sum_{j=1}^{J} p_{j,t}(\mathbf{g,a}) \sum_{r=1}^{R_j} F_j(r)\,(r + V_{t+1}(\tilde{\mathbf{g}})) + (1 - \lambda + \lambda p_{0,t}(\mathbf{g,a}))\, V_{t+1}(\mathbf{g})\right\}, \forall \mathbf{g} \in \Omega_t, 0 \le t \le T-1, \tag{2a}$$

$$V_T(\mathbf{g}) = \sum_{j=1}^{J} h_T(g_j), \text{ where } h_T(g_j) = \begin{cases} -(G_j - g_j) & if \quad g_j > 0, \\ 0 & if \quad g_j \le 0. \end{cases}, \forall \mathbf{g} \in \Omega_T. \tag{2b}$$

The termination condition (2b) ensures the AoN scheme is followed, where only those projects that have reached their funding goals keep the funds raised; otherwise all the money is returned. It is worth mentioning that the problem may be formulated alternatively such that the money is only collected at the end of campaigns for completed projects. We show that both formulations are equivalent in Appendix A.

A policy π is any non-anticipative rule to choose a project for promotion after observing the system state at each decision epoch. Let Π denote the set of policies $\pi : \Omega_t \to \mathcal{A}, \forall 0 \le t \le T-1$. We aim to find such a policy that maximizes the expected total revenue from the initial state $\mathbf{G} = (G_1, \cdots, G_J)$. In principle, the problem described above can be solved by the standard dynamic programming approach. This, however, is computationally tractable only for small-scale problems due to the *curse of dimensionality*. In this paper, we follow Whittle (1988) and develop index policies that only concern a single project at a time.

Remark 2.1 (Multiple customer segments)**.** *The above model can be readily generalized to accommodate multiple customer segments. Suppose there are K customer segments each denoted by k, and the probability that an arriving customer is of segment k is* q^k*. The utility function can be extended with segment specific parameters, denoted as* $U^k_{j,t}$*, as can the backing probabilities* $p^k_{j,t}(\boldsymbol{g,a})$ *and the pledging probabilities* F^k_j*. Then the Bellman equation (2a) can be modified to*

$$V_t(\boldsymbol{g}) = \max_{\boldsymbol{a} \in \mathcal{A}} \left\{\lambda \sum_{k=1}^{K} q^k \sum_{j=1}^{J} p^k_{j,t}(\boldsymbol{g,a}) \sum_{r=1}^{R_j} F^k_j(r)(r + V_{t+1}(\tilde{\boldsymbol{g}})) + (1 - \lambda + \lambda \sum_{k=1}^{K} q^k p^k_{0,t}(\boldsymbol{g,a}))V_{t+1}(\boldsymbol{g})\right\}.$$

Remark 2.2 (Multiple promotions)**.** *The developed model can be straightforwardly generalized to the scenario of multiple promotions. Suppose there are N promotion slots available at each decision epoch. We modify the action set to be* $\mathcal{A} = \left\{(a_1, a_2, ..., a_J) : a_j \in \{0, 1\}, \sum_{j=1}^{J} a_j \leq N\right\}$*. The Bellman equation remains unchanged. The relaxation and decomposition approach to be developed in the next section will follow in exactly the same way. All the analyses and results on the decomposed single-project problems are not impacted.*

3. Relaxation and Decomposition

In this section, we consider problem relaxation and decompose the problem in the previous section into a collection of single-project problems.

Under any policy $\pi \in \Pi$, we write the expected total funds raised over the entire time horizon as

$$V_0^{\pi}(\mathbf{G}) = \mathbb{E}\left[\sum_{t=0}^{T-1} h_t(\mathbf{g}(t), \pi_t(\mathbf{g}(t))) + \sum_{j=1}^{J} h_T(g_j(T))\right],$$

where the state vector $\mathbf{g}$ at time t is explicitly denoted as $\mathbf{g}(t)$, and $h_t(\mathbf{g}, \mathbf{a}) = \lambda \sum_{j=1}^{J} p_{j,t}(\mathbf{g},\mathbf{a}) \sum_{r=1}^{R_j} rF_j(r)$ is the expected one-period revenue at time t when action $\mathbf{a}$ is taken at state $\mathbf{g}$. Our optimization problem may be expressed as $V_0(\mathbf{G}) = \sup_{\pi \in \Pi} V_0^{\pi}(\mathbf{G})$.

We now consider the problem relaxation with a different set of policies in which more than one project can be promoted at each decision epoch, i.e., $\tilde{\mathcal{A}} = \left\{(a_1, a_2, ..., a_J) : a_j \in \{0, 1\}\right\}$. We denote by $\tilde{\Pi}$ the set of policies $\tilde{\pi} : \Omega_t \to \tilde{\mathcal{A}}, \forall 0 \leq t \leq T - 1$. Clearly we have $\Pi \in \tilde{\Pi}$. Even though we allow multiple projects to be promoted simultaneously, it is reasonable to require the average resource consumed does not exceed the capacity, i.e., $\mathbb{E}\left[\sum_{t=0}^{T-1} \left(1 - \sum_{j=1}^{J} \tilde{\pi}_{j,t}(\mathbf{g}(t))\right)\right] \geq 0$, where $\tilde{\pi}_{j,t}(\mathbf{g}(t))$ is the action for project j under policy $\tilde{\pi}$. Following Whittle (1988), we associate a non-negative Lagrangian multiplier W to the constraint and incorporate it into the objective:

$$\hat{V}_0^{W}(\mathbf{G}) = \sup_{\tilde{\pi} \in \tilde{\Pi}} \mathbb{E}\left[\sum_{t=0}^{T-1} h_t\left(\mathbf{g}(t), \tilde{\pi}_t(\mathbf{g}(t))\right) + \sum_{j=1}^{J} h_T(g_j(T)) + W \sum_{t=0}^{T-1}\left(1 - \sum_{j=1}^{J} \tilde{\pi}_{j,t}(\mathbf{g}(t))\right)\right]. \tag{3}$$

Unlike other RB problems in the literature which would have been decomposed by now, problem (3) is not yet decomposable due to the dependency among the projects via the backing probabilities used to calculate $h_t(\mathbf{g}, \tilde{\pi}_t(\mathbf{g}))$, which are computed from the MNL model involving all the projects.

Recall that our objective is to develop index policies following Whittle's restless bandit approach. As mentioned in the Introduction Section, this approach aims to derive index values and then allocate the resource to the project with the maximum index value. Therefore, unlike the other approaches in the weakly coupled dynamic program literature which revolve around value-function approximation, this approach is concerned with the ranking of the projects. In light of this, we wish to find such a way that decouples the original problem, while at the same time preserves the ranking of the projects as measured by backers' choice probabilities. To this end, we propose to relax the problem by replacing the MNL-based backing probabilities with the following Binomial Logit (BNL) model, one for each project j:

$$p_{j,t}^{a_j}(g_j) = \frac{\exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right)\right)}{1 + \exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right)\right)}, \tag{4}$$

where we use a simplified notation $p_{j,t}^{a_j}(g_j)$ for the backing probability of project j at state g_j under action a_j at time epoch t. Note that here by definition, $\beta_{3,J}$ vanishes from the backing probabilities since competition is no longer an issue for the single-project problems (i.e., $J = 1$).

Remark 3.1 (Rank preservation from the MNL to BNL)**.** *In statistics, for any multinomial distribution having J possible mutually exclusive outcomes with corresponding probabilities $p_1, ..., p_J$, it is well known that the order for the probabilities of occurrence determined among the outcomes is maintained for these outcomes in its marginal and conditional distributions; see, e.g., Balakrishnan and Nevzorov (2004), Ch.25. For the crowdfunding problem investigated in this paper, this rank preservation property also holds. More specifically, if $p_{j,t}(\mathbf{g}, \mathbf{a}) \geq p_{k,t}(\mathbf{g}, \mathbf{a})$ for any two projects j and k in problem (2a), then we have $p^{a_j}_{j,t}(g_j) \geq p^{a_k}_{k,t}(g_k)$, i.e. the order for backers' choice probabilities is maintained when replacing the MNL model in (1) with the BNL model in equation (4). See Lemma B.1 in Appendix B for proof.*

With some algebraic manipulation, we may write our relaxation of the original problem in the following form:

$$V^W_0(\mathbf{G}) = \sup_{\tilde{\pi}\in\tilde{\Pi}} \mathbb{E}\left[\sum_{j=1}^{J}\left(\sum_{t=0}^{T-1}\left(\lambda p^{\tilde{\pi}_{j,t}(\mathbf{g}(t))}_{j,t}(g_j(t))\sum_{r=1}^{R_j} rF_j(r) - W\tilde{\pi}_{j,t}(\mathbf{g}(t))\right) + h_T(g_j(T))\right)\right] + WT. \tag{5}$$

The Lagrangian multiplier W can be viewed as the charge for using the promotion space for one time period. The two terms inside the inner parentheses on the right-hand side of equation (5) are the expected one-period revenue obtained from each project j and the promotion charge in each time period between 0 and $T-1$, respectively. $h_T(g_j(T))$ is the terminal value of each project in time T. The last term WT is a constant and will be discarded in the subsequent discussion. Thus, for the three optimization problems considered in this section, we have the following result.

Proposition 3.1. *$V_0(\boldsymbol{G}) \leq \hat{V}^W_0(\boldsymbol{G}) \leq V^W_0(\boldsymbol{G})$, $\forall W \geq 0$.*

Proof. See Appendix B.1. □

Then, the problem (5) can readily be decomposed into a collection of independent single-project problems ($j = 1, ..., J$):

$$v^W_{j,0}(G_j) = \sup_{\pi_j} \mathbb{E}\left[\left(\sum_{t=0}^{T-1}\left(\lambda p^{\pi_{j,t}(g_j(t))}_{j,t}(g_j(t))\sum_{r=1}^{R_j} rF_j(r) - W\pi_{j,t}(g_j(t))\right) + h_T(g_j(T))\right)\right]. \tag{6}$$

Each single-project problem j in the above equation can be understood as a problem in which project j has a dedicated promotion space, and the action is whether or not to use the space for promotion at each decision epoch in light of the current state. If the action is to promote ($a_j = 1$), the project will be highlighted on the homepage with a cost of W. If the decision is not to promote ($a_j = 0$), the project will not be highlighted and no cost is incurred. Hence, for each single-project problem j, the state space becomes $\Omega_{j,t} = \left\{g_j : G_j - tR_j \leq g_j \leq G_j\right\}$, and the policy becomes $\pi_j : \Omega_{j,t} \to \{0, 1\}, \forall 0 \leq t \leq T-1$.

4. Indexability and Index Policies

In this section, we focus on a single-project problem and investigate the structural properties of the optimal policies. We establish indexability and derive Whittle's index, upon which we propose two index heuristics: one based on the derived index values and the other on an approximation. As we are now concerned with a single-project problem, the subscript j is dropped in all the notation.

4.1. Optimal Policies for a Single-Project Problem with Sufficiently Long Campaign Duration.

We begin with an observation that the optimal policy for problem (6) does not have the monotone structure: the optimal action might switch between 0 and 1 back and forth over time or cross states. If the project is already over-funded or far from the funding goal, there is no motivation to choose promotion when there are just a few time periods left, whereas if the project has almost reached the funding goal the platform would always choose it for promotion to pass the bar in the last minute. This lack of monotone structure complicates the analysis and establishment of indexability (Gittins et al., 2011), which is particularly the case in our problem that has two dimensional states. Nevertheless, it is straightforward to see that when the campaign duration T is sufficiently long (compared to the funding goal G), it is expected that the project will complete, thanks to the strictly positive and bounded purchasing probabilities. We provide a detailed discussion on a mathematical condition for the sufficiently long campaign duration in Appendix C. In such a case the AoN scheme becomes irrelevant and we may set $h_T(g) \equiv 0, \forall g \in \Omega_T$ in (6). As we shall see, with such a simplification the optimal policies become monotonic in both state and time, which greatly simplifies the task of demonstrating indexability and more importantly, allows the development of efficient algorithms to calculate the index values, which is critical for the scalability of the resulting index policies. In the remainder of this section we shall focus on indexability analysis under the condition of sufficiently long campaign duration.

The following results hold for each project. It shows that the value of promotion gradually diminishes with the progression of the campaign if the herding parameter is non-decreasing over time. In the majority of crowdfunding literature the herding parameter is deemed as a constant (Zhang et al., 2023; Xiao et al., 2021), while a few consider it increasing over time (Dai and Zhang, 2019). Therefore, in the remaining of this section we assume that the herding parameter $\beta_{2,t}$ is non-decreasing in t.

Lemma 4.1. *The backing probability $p_t^a(g)$, $a \in \{0, 1\}$, in equation (4) has the following properties for all $g \in (-\infty, G]$:*

(i) $p_t^a(g)$ is decreasing concave in g;

(ii) The difference, $p_t^1(g) - p_t^0(g)$, is positive and monotonically increasing in g;

Moreover, if the herding parameter $\beta_{2,t}$ is non-decreasing in t, we have

(iii) The difference, $p_t^1(g) - p_t^0(g)$, is monotonically non-increasing in t;

(iv) The difference, $p_t^a(g-1) - p_t^a(g)$, is monotonically non-increasing in t.

Proof. See Appendix B.2. □

Remark 4.1 (Indexability for general utility functions)**.** *Lemma 4.1 provides a sufficient condition for passing the indexability test. Although we have discussed a linear utility function so far, the theorems/propositions in this section will hold and the indexability test will pass for a utility function with any form, provided that its corresponding choice probability satisfies the conditions in Lemma 4.1.*

Denote the value function at time t and state g by $v_t^W(g)$ for a given Lagrangian multiplier W. The Bellman equation for the single-project problem can be written as follows:

$$\begin{aligned} v_t^W(g) &= \max_{a\in\{0,1\}} \left\{\lambda p_t^a(g) \textstyle\sum_{r=1}^R F(r)\left(r + v_{t+1}^W(g-r)\right) + (1-\lambda p_t^a(g))\, v_{t+1}^W(g) - Wa\right\} \\ &= \max_{a\in\{0,1\}} \left\{\lambda p_t^a(g)\left(\bar{r} + \triangle v_{t+1}^W(g)\right) + v_{t+1}^W(g) - Wa\right\}, \forall g \in \Omega_t, 0 \le t \le T-1, \\ v_T^W(g) &= h_T(g) = 0, \forall g \in \Omega_T, \end{aligned} \tag{7}$$

where $\bar{r} = \sum_{r=1}^R F(r)r$ and $\triangle v_t^W(g) = \sum_{r=1}^R F(r)v_t^W(g-r) - v_t^W(g)$. Denote by $\mathring{\pi}^W$ the optimal policy under W, for which the action at state g and time t can be specified as

$$\mathring{\pi}_t^W(g) = \begin{cases} 1, & \text{if } \lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\bar{r} + \triangle v_{t+1}^W(g)\right) > W, \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

In the above if-condition, $\lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\bar{r} + \triangle v_{t+1}^W(g)\right)$ is the expected benefit of promotion at state g and time t. To see this, we note that $p_t^1(g) - p_t^0(g)$ is the increment in the backing probability due to promotion, $\bar{r}$ is the immediate (expected) money raised from a single purchase, and $\triangle v_{t+1}^W(g)$ can be understood as the marginal future revenue of this purchase (i.e., the extra future revenue due to an additional purchase at the current time). Hence, the project should be promoted only if the benefit of promotion outweighs the cost W. It is therefore reasonable to expect that the project would prefer to promote if W is small, and prefer not to promote when W is large.

To facilitate the subsequent analysis, we now introduce two special policies π^a ($a \in \{0,1\}$) that always take the same action a throughout the campaign, i.e., always to promote for $a = 1$ (denoted as π^1) and always not to promote for $a = 0$ (denoted as π^0). Define the corresponding value-to-go under W as:

$$v_t^{W,\pi^a}(g) = \lambda p_t^a(g)\left(\bar{r} + \triangle v_{t+1}^{W,\pi^a}(g)\right) + v_{t+1}^{W,\pi^a}(g) - Wa,$$

and let $\triangle v_t^{W,\pi^a}(g) = \sum_{r=1}^R F(r)v_t^{W,\pi^a}(g-r) - v_t^{W,\pi^a}(g)$ denote the marginal future revenue of an additional purchase under π^a. Note that under the non-promotion policy π^0, the value-to-go does not depend on W. We use simplified notations in this case, i.e., $v_t^{\pi^0}$ for v_t^{W,π^0} and $\triangle v_t^{\pi^0}$ for $\triangle v_t^{W,\pi^0}$.

Lemma 4.2. *Under policy π^a, we have:*

(i) $\triangle v_t^{W,\pi^a}(g) > \triangle v_t^{W,\pi^a}(g-1)$*,* $\forall g \in \Omega_t, 0 \le t \le T-1$*;*

(ii) $\triangle v_t^{W,\pi^a}(g) > \triangle v_{t+1}^{W,\pi^a}(g)$*,* $\forall g \in \Omega_t, 0 \le t \le T-1$.

Proof. See Appendix B.4. □

The above lemma shows that the marginal future revenue of an additional purchase under policy π^a always decreases when there are more funds already raised or less time remaining in the campaign. Therefore, under π^a, the marginal future revenue of an additional purchase gradually tails off with the progress of the campaign. Note that the result of Lemma 4.2 (i) can be straightforwardly extended to $\triangle v_t^{W,\pi^a}(g) > \triangle v_t^{W,\pi^a}(g-r)$ for any $r \in \mathcal{R}$.

Theorem 4.1 (Two Critical Values of the Promotion Cost)**.**

(i) There exists a positive and finite value $\overline{W}$, such that the optimal policy takes the form of π^0 if and only if the promotion cost satisfies $W > \overline{W}$.

(ii) There exists a positive and finite value, given by $\underline{W} = \lambda\bar{r}\left(p_{T-1}^1(G-(T-1)R) - p_{T-1}^0(G-(T-1)R)\right)$, such that the optimal policy takes the form of π^1 if and only if the promotion cost satisfies $W \le \underline{W}$.

Proof. See Appendix B.5. □

These two critical values jointly define a range for W such that beyond this range, the optimal policy takes the form of π^a. The following theorem shows how the optimal action changes within this range.

Theorem 4.2 (The Change Pattern of the Optimal Action)**.** *Starting from $\overline{W}$, as the cost W gradually decreases to $\underline{W}$, the optimal action for each state at each time epoch switches from $a = 0$ to $a = 1$ in the following manner:*

(i) The first to switch takes place at state G and time 0.

(ii) The last to switch takes place at state $G - (T - 1)R$ and time $T - 1$.

(iii) For the same state g, the switch takes place at time t before $t + 1$.

(iv) For the same time t, the switch takes place at state g before $g - 1$.

Proof. See Appendix B.6. □

We are now ready to present our main results on the single-project problem (7).

Proposition 4.1 (Monotonicity of the Optimal Policy)**.** *For any $W \geq 0$, the optimal policy $\overset{*}{\pi}{}_t^W$ satisfies:*

(i) $\overset{*}{\pi}{}_t^W(g) \geq \overset{*}{\pi}{}_t^W(g-1)$, $\forall g \in \Omega_t, 0 \leq t \leq T-1$;

(ii) $\overset{*}{\pi}{}_t^W(g) \geq \overset{*}{\pi}{}_{t+1}^W(g)$, $\forall g \in \Omega_t, 0 \leq t \leq T-1$.

Proof. See Appendix B.7. □

Proposition 4.1 shows that the optimal policy is monotonic in both g and t. At each time epoch, if the optimal action is not to promote at a certain shortfall level, this action will remain optimal as the shortfall will either reduce or stay at the same level. On the other hand, for a given shortfall g, once the optimal action is not to promote at a time epoch, this action will remain optimal in the remainder of the campaign. Therefore, the optimal policy chooses to promote the project earlier rather than later. Once it switches to non-promotion, it will remain so until the end of the campaign. We now use a small running example to illustrate the results.

Running Example. Consider a project with the funding goal $G = 2$, attraction $m = 0.1$, and two pledge options with the same purchasing probability $F(r = 1) = F(r = 2) = 0.5$. The campaign duration is $T = 12$, which is sufficiently long compared to its funding goal. The other parameters are chosen as follows: arrival probability $\lambda = 0.7$, promotion power $\beta_1 = 1$, and the market side effect $\beta_{3,1} = 0$. Moreover, we assume a time-increasing herding parameter in the form of $\beta_{2,t} = 0.01 + \gamma \cdot t/T$, where $\gamma = 0.01$.

For any given W, we solve the Bellman equation (7) for this project and obtain the optimal policy. Fig. 1 depicts the optimal policy for some selected W values. For example, Fig. 1 (a) corresponds to the scenario where the promotion cost $W > \overline{W}$ (the method of deriving $\overline{W}$ is discussed in Section 4.2) and it shows that the optimal action for the project is always not to promote, as indicated by the red area. As the promotion cost reduces from $\overline{W}$ to W^1 and further to W^2 ($W^1 > W^2$), the collection of states in which it is optimal to promote gradually increases, as shown by the blue area in Fig. 1 (b) and (c) respectively. Finally, Fig. 1 (d) shows the scenario of $W \leq \underline{W}$ where the optimal action becomes to promote the project all the time.

Figure 1: An illustration of the optimal actions of the project for selected W values in the running example.

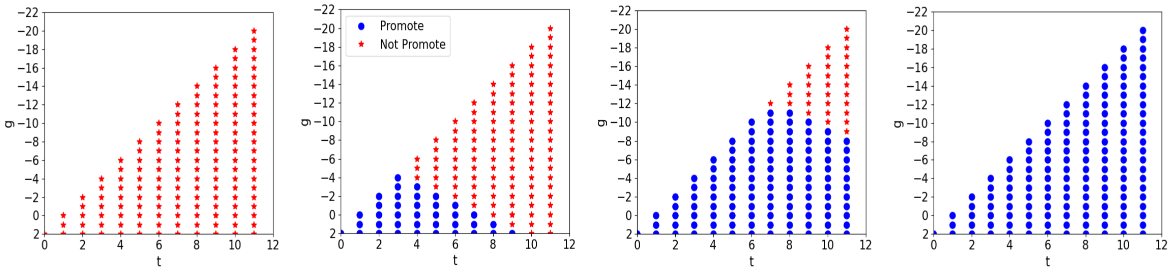

(a) $W > \overline{W} = 0.240$ (b) $W = W^1 = 0.237$ (c) $W = W^2 = 0.233$ (d) $W \le \underline{W} = 0.228$

4.2. Indexability for Projects with Sufficiently Long Campaign Duration.

For any fixed promotion cost W and time t, define the optimal promotion set as $B_t(W) = \left\{g : \overset{*}{\pi}{}_t^W(g) = 1, g \in \Omega_t\right\}$. Note that for some W values and time t the optimal action might be indifferent between promotion and no promotion at some states. Without causing any ambiguity we include these states into the corresponding optimal promotion sets. Hence, it is optimal to promote the project at time t if and only if the state belongs to $B_t(W)$. An important implication of Proposition 4.1 is that for each time epoch t and a given W, there exists a critical state, denoted as $g^*(W, t)$, such that

$$B_t(W) = \{G, G-1, \cdots, g^*(W, t)\}. \tag{9}$$

We write $g^*(W, t) = \infty$ if $B_t(W) = \varnothing$ and $g^*(W, t) = -\infty$ if $B_t(W) = \Omega_t$. We now formally define indexability of a project as follows.

Definition 4.1 (Indexability)**.** *A project is indexable if $B_t(W) \subseteq B_t(W')$ for any $W \ge W'$, $\forall 0 \le t \le T-1$.*

We have the following result.

Proposition 4.2 (Indexable Projects)**.** *The projects with sufficiently long campaign duration are indexable in our problem.*

Proof. See Appendix B.9. □

Therefore, as $W \to 0$, the optimal promotion set grows until $B_t(W) = \Omega_t$, whereas $B_t(W) \to \varnothing$ when $W \to \infty$. We are now ready to define the Whittle's index.

Definition 4.2 (Whittle's Index)**.** *For an indexable project, the Whittle's index is defined as*

$$w(g, t) = \max\{W : g \in B_t(W)\}.$$

Thus, $w(g, t)$ is the maximum promotion cost under which it is optimal to promote at state g and time t. In the RB literature this is usually termed the fair charge, at which the project is indifferent between being promoted or not. The index value $w(g, t)$ changes with both state and time, and it is monotonic in either dimension, as shown in the following lemma.

Lemma 4.3. *The index value $w(g, t)$ (i) increases in state g; and (ii) decreases in time t.*

Proof. See Appendix B.8. □

From Lemma 4.3, we know that $\overline{W}$ is the index value at state G and time 0, i.e., $\overline{W} = w(G, 0)$. The following proposition shows how to calculate the index values.

Proposition 4.3 (Whittle's Index Values)**.** *The Whittle index is evaluated as follows:*

$$w(g,t) = \lambda \left(p_t^1(g) - p_t^0(g)\right) \left(\bar{r} + \triangle v_{t+1}^{\pi^0}(g)\right), \tag{10}$$

where $\triangle v_t^{\pi^0}(g) = \sum_{r=1}^{R} F(r) v_t^{\pi^0}(g-r) - v_t^{\pi^0}(g)$ *is the marginal future revenue of an additional purchase under policy* π^0.

This is the direct result of the definition of index values and Theorem 4.2. In equation (8), when the equality holds in the if-condition for any g and t, the benefit and cost of promotion are just the same and the corresponding W is the indifference cost, i.e., the index value. By Theorem 4.2, the optimal action for all the future states in the remaining campaign has not yet switched from $a = 0$ to $a = 1$, so the optimal policy (under this W) for them is still π^0, and thus equation (10). The value-to-go $v_t^{\pi^0}(g)$ can be conveniently obtained via solving the single-project DP (7) for $a = 0$. As the number of pledging options R is usually so small (e.g., 3-10 in most projects (Kickstarter, 2023)) that it is negligible compared to T, this algorithm has the complexity of linear-time (in number of states $0.5RT(T-1)+1$), significantly faster than the general algorithm that takes cubic time (Niño-Mora, 2007).

Nevertheless, we note that due to the recursive format of equation (7), the $v_t^{\pi^0}(g)$ values must be calculated off-line, i.e., in advance for all possible states g. To further reduce the computational time, we propose a closed-form approximation of the Whittle's index values below. Specifically, for each state g at time t, we assume that the herding effect from $t+1$ onwards and for all the future states remains the same as it is evaluated at time t for state g. Under this assumption, the purchasing probabilities become a constant for the remaining time horizon. Hence, an approximation to the value-to-go, denoted by $\hat{v}_t^{\pi^0}(g)$, can be derived as follows,

$$\hat{v}_t^{\pi^0}(g) = \lambda p_t^0(g) \sum_{t}^{T-1} \bar{r} = (T-t)\lambda p_t^0(g)\bar{r}.$$

Substituting $\hat{v}_t^{\pi^0}(g)$ into (10), the Whittle's index value can be approximated as

$$\hat{w}(g,t) = \lambda(p_t^1(g) - p_t^0(g)) \left[\bar{r} + \lambda\bar{r}(T-t-1)\left(\sum_{r=1}^{R} F(r)(p_t^0(g-r) - p_t^0(g))\right)\right]. \tag{11}$$

It is straightforward to see that the approximate index values via equation (11) preserve the monotonicity property as stated in Lemma 4.3. More precisely, they decrease with time t; by Lemma 4.1, both $p_t^1(g) - p_t^0(g)$ and $p_t^0(g-r) - p_t^0(g)$ increase in state g, so do the index values.

4.3. Index Policies for the Original Problem

Thus far, we have shown that the projects with sufficiently long campaign duration are indexable and how to calculate their index values. This allows us to develop two efficient index policies for the original problem:

- *Index Policy (IP)*: At each decision epoch t when state $\mathbf{g} = (g_1, \cdots, g_J)$ is occupied, the index policy IP always chooses to promote the project j^* with the largest index value, i.e., $j^* = \arg\max_{1\le j\le J} w_j(g_j, t)$, where for each project j, the index value $w_j(g_j, t)$ is calculated from Proposition 4.3.

- *Index Policy Approximation (IPx)*: The IPx policy follows the same idea as IP but uses the approximate Whittle's index values $\hat{w}_j(g_j, t)$, which are calculated via equation (11).

As mentioned previously, for projects with short campaign duration, the optimal policy is not monotonic due to the AoN scheme. In such cases indexability cannot be easily established, and the calculation of index values is much more involved; we may calculate the Whittle index values under the AoN scheme using a generic algorithm such as the method suggested by Niño-Mora (2007) or the Whittle's index approach in Brown and Smith (2020), but these algorithms are computationally expensive as they need to solve the single-project DPs (7) once for each state g and time t. In contrast, our proposed algorithm only needs to solve a single instance of DP (7) under the non-promotion policy. Hence, for practical purposes we propose to use the above IP and IPx polices, even if the campaign duration is not long. Their performance is examined in detail in the next section.

5. Numerical Experiments

In this section, we evaluate the performance of the proposed index policies against alternative heuristics in a number of problem instances. Section 5.1 introduces these alternative heuristic policies and Section 5.2 describes the experiment settings. Section 5.3 reports all the results: (i) Section 5.3.1 shows the sub-optimality gap of each policy in the running example, (ii) Section 5.3.2 explores each policy's performance in larger problem instances for which it is no longer possible to find the exact DP solution, and (iii) Section 5.3.3 discusses the robustness of the index policies. Section 5.4 provides a detailed discussion on the behavior of the index values and the nuances of the two index policies. Finally, in Section 5.5, the experiments are extended to the problems where projects have different start and end times in their campaigns.

5.1. Alternative Policies

As we are not aware of any readily applicable policies in the literature for our crowdfunding problem, we develop two dynamic policies as benchmarks for comparison purposes. One is based on the Lagrangian relaxation, while the other is a deterministic approximation to the problem. In addition, we also consider several simple heuristics in the experiments.

5.1.1. Lagrangian Relaxation (LR)

In Section 3 we decomposed the Lagrangian relaxed problem (5) into a collection of independent single-project problems (6). It is straightforward to see that the value function for the Lagrangian relaxed problem can be calculated as $V_0^W(\mathbf{G}) = \sum_{j=1}^{J} v_{j,0}^W(G_j) + WT$. We have the following result.

Proposition 5.1 (Lagrangian Convexity)**.** *For all $\boldsymbol{G}$, $V_0^W(\boldsymbol{G})$ is piecewise linear and convex in W.*

Proof. See Appendix B.10. □

From Proposition 3.1, the Lagrangian value function $V_0^W(\mathbf{G})$ can be used as an upper bound to the original problem. We want to find value W that produces the best upper bound, which can be identified by solving the following Lagrangian dual problem

$$\min_{W\geq 0} V_0^W(\mathbf{G}). \tag{12}$$

Based on Proposition 5.1, the Lagrangian dual problem is a convex optimization problem, which can be solved by a sub-gradient algorithm. We first define the following notation for each project j. Let $\mathring{\pi}_j^W$ be the optimal policy that solves the single-project problem (6) for a given W. Denote by $\mathbf{U}_{j,t}$ the vector of actions under $\mathring{\pi}_j^W$ in all states at time t, and by $\mathbf{Q}_{j,t}$ the corresponding one-step transition probability matrix at time t.

Proposition 5.2 (Sub-gradient)**.** *A sub-gradient of $V_0^W(\boldsymbol{G})$ at W, denoted by $\nabla V(W)$, is given by*

$$\nabla V(W) = T - \sum_{j=1}^{J} \left[\mathbf{U}_{j,0} + \sum_{s=1}^{T-1} \left(\prod_{\tau=0}^{s-1} \mathbf{Q}_{j,\tau} \right) \times \mathbf{U}_{j,s} \right] (G_j),$$

where $[\cdot](g)$ is the component corresponding to state g of a vector. Specifically, $\mathbf{U}_{j,t}$ denotes a (TR_j+1)-dimensional vector for the optimal actions at time t, $\mathbf{P}_{j,t}$ denotes a (TR_j+1)-dimensional vector for the corresponding backing probabilities at time t, and $\mathbf{Q}_{j,t}$ denotes a $(TR_j+1)\times(TR_j+1)$ matrix for the one-step transition probabilities at time t under the optimal policy.

Proof. See Appendix B.11 □

Then the Lagrangian dual problem (12) can be solved by the sub-gradient optimization methods (Brown and Smith, 2020). Denote the solution by W^*. The *Lagrangian Relaxation (LR) policy* then chooses the action that achieves the following at epoch t when state $\mathbf{g}$ is occupied.

$$\arg\max_{\mathbf{a}\in\mathcal{A}} \left\{ \lambda \sum_{j=1}^{J} p_{j,t}(\mathbf{g,a}) \sum_{r=1}^{R_j} F_j(r) \left(r + V_{t+1}^{W^*}(\tilde{\mathbf{g}}) \right) + (1 - \lambda + \lambda p_{0,t}(\mathbf{g,a}))\, V_{t+1}^{W^*}(\mathbf{g}) \right\},$$

where $V_t^{W^*}(\mathbf{g}) = \sum_{j=1}^{J} v_{j,t}^{W^*}(g_j) + W^*(T-t)$.

5.1.2. Deterministic Integer Program (DIP)

At each time epoch t when state $\mathbf{g}$ is occupied, assuming the backing probabilities for each project remain unchanged in the rest of the time horizon, we can formulate a deterministic integer program to approximate the original problem. More specifically, let x_j^s be a binary decision variable which takes value 1 if project j is promoted at time s, and 0 otherwise. Denote by y_j the total fund raised for project j by the end of the campaign and Y_j the fund raised under AoN. Let $\delta_j = \{0, 1\}$ be an indicator variable which takes value 1 if the project is successful and 0 otherwise. The integer program takes the following form.

$$\begin{aligned}
\text{DIP}(\mathbf{g}, t) : \max \quad & \sum_{j=1}^{J} Y_j \\
s.t. \quad & \sum_{j=1}^{J} x_j^s = 1, \forall t \le s \le T-1 \\
& y_j = \left(G_j - g_j\right) + \sum_{s=t}^{T-1} \lambda \left(p_j x_j^s + \sum_{k\neq j}^{J} q_{j,k} x_k^s \right) \bar{r}_j, \forall j \\
& G_j - M\left(1-\delta_j\right) \le y_j < G_j + M\delta_j, \forall j \\
& y_j - M\left(1-\delta_j\right) \le Y_j \le y_j + M\left(1-\delta_j\right), \forall j \\
& -M\delta_j \le Y_j \le M\delta_j, \forall j,
\end{aligned} \qquad (13)$$

where p_j is the backing probability of project j if it is promoted in state $\mathbf{g}$ at time t, and $q_{j,k}$ is the backing probability of project j when another project k is being promoted. Note again both probabilities remain unchanged over time, always the same as they were evaluated at time t. M is a big number. The objective function maximizes the total funds raised under the AoN scheme. The first constraint ensures only one project is promoted each time. The second constraint calculates the total money raised before considering AoN, while the third constraint determines if or not a project is successful. The last two constraints calculate the total funds raised under AoN.

Ideally such an integer program is solved at every time epoch to produce a dynamic policy; only the solution of the current epoch (i.e., x_j^t) is implemented while the rest are all discarded. Such an approach, however, is time consuming especially for large problems. Instead, a more practical approach is that one may re-optimize at specific (regular) intervals, where within each interval the solutions are implemented deterministically. By the end of each interval another integer program is solved for the state then occupied, and the same implementation repeats until the end of time horizon. This policy is termed *DIP* hereafter.

5.1.3. Simple Heuristics

We also consider a few simple heuristics, including:

- *Smallest/largest shortfall first (SSF/LSF)*: These two policies always promote the project with the smallest/largest percentage shortfall g_j/G_j at every decision epoch.

- *Smallest/largest utility first (SUF/LUF)*: These two policies always promote the project with the smallest/largest expected utility $z_j\left(g_j, 1\right)$ at every decision epoch. As far as we know, LUF is a policy that resembles the algorithm used in some crowdfunding platforms, e.g., Kickstarter (2015).

- *Greedy policy (GP)*: At each decision epoch, this policy promotes the unfinished project with the highest funding goal, i.e., $j^* = \arg\max_{j:g_j>0} \left\{G_j\right\}$. Once all the projects have met their funding goals, the policy randomly chooses one project to promote.

- *Conservative policy (CP)*: At each decision epoch, this policy promotes the unfinished project with the lowest funding goal, i.e., $j^* = \arg\min_{j:g_j>0} \left\{G_j\right\}$. Once all the projects have met their funding goals, the policy randomly chooses one project to promote.

- *Myopic policy (MP)*: Denote the total immediate reward when project j is promoted at state $\mathbf{g}$ by $L_j(\mathbf{g}) = \lambda \sum_{k=1}^{J} p_{k,t}\left(\mathbf{g}, \mathbf{e}_j\right) \bar{r}_k$, where $\mathbf{e}_j$ and $p_{k,t}(\cdot,\cdot)$ were defined in Section 2. At each decision epoch, this policy chooses to promote project $j^* = \arg\max_{1\le j\le J} \left\{L_j(\mathbf{g})\right\}$.

5.2. Experiment Settings

We consider two scenarios, $J = 3$ and $J = 5$. There are two types of parameters in our problem. The global parameters include the promotion power β_1, herding effect $\beta_{2,t}$, market side effect $\beta_{3,J}$, and arrival rate λ. We follow the ideas of the existing crowdfunding literature to set these parameters. In Varma et al. (2021), the visibility of a project (which directly affects the investment) is increased by 30% when being promoted. Since the promotion in our model influences the amount of investment indirectly via customers' backing probabilities, we slightly increase the promotion power. The market side effect is set at a similar magnitude to the promotion power. The herding effect is the same as specified in the running example. Specifically, we set $\beta_1 = 1$, $\beta_{2,t} = 0.01 + \gamma \cdot (t/T)$ with $\gamma = 0.01$, $\beta_{3,3} = 2.5$, and $\beta_{3,5} = 3$. The arrival rate is set as $\lambda = 0.7$ in all instances. The local parameters for each project, including the attraction m_j, funding goal G_j, and pledge probability $F_j(r)$ for $r \in \mathcal{R}_j$, are shown in Table 1. For each problem, the more attractive projects also have larger funding goals, and the higher pledging probabilities are associated with more expensive reward options. Such settings avoid trivial cases where the optimal promotion actions are obvious.

Besides the above parameters, the campaign duration T plays a pivotal role in promotion strategies under the AoN scheme. If T was too short, only those projects with the smallest funding goals could complete and thus they would always be promoted first. If T was too long, most projects would complete so the ones with the largest funding goals would be promoted first. The problem is really of interest when the campaign duration is moderate, in which case the optimal promotion policy is far from clear. The choice of campaign duration is an interesting

Table 1: Parameter setting for the local parameters.

Project j	$J=3$			$J=5$				
	1	2	3	1	2	3	4	5
m	0.04	0.08	0.16	0.02	0.04	0.08	0.12	0.16
G	40	80	160	20	40	80	120	160
$F(r=1)$	0.55	0.5	0.48	0.8	0.6	0.5	0.46	0.45
$F(r=2)$	0.45	0.5	0.52	0.2	0.4	0.5	0.54	0.55
$\bar{r}$	1.45	1.5	1.52	1.2	1.4	1.5	1.54	1.55

problem in its own right. Readers may refer to Zhang et al. (2023) for a detailed discussion on the influence of campaign duration in crowdfunding. In our experiments, we consider a reference duration $\overline{T}$, which is the expected time needed to *just* complete all the projects without promotion and herding effect. Specifically, we note that the total purchasing probability reduces to $P = 1 - \frac{1}{1+\sum_{j=1}^{J} e^{m_j - \beta_3}}$ in this case, and $\overline{T} = \frac{\sum_{j=1}^{J} G_j/\bar{r}_j}{\lambda \cdot P}$, where the numerator is the number of purchases required to complete all the projects, and the denominator can be understood as the mean number of purchases per time epoch. To evaluate the performance of the index policies for a variety of campaign durations, we consider a range of T as different proportions of the reference duration. For $J = 3$, we consider $T \in \left\{40\%\overline{T}, 50\%\overline{T}, 60\%\overline{T}, 70\%\overline{T}\right\}$. Due to more projects competing for funding when $J = 5$, we consider slightly higher proportions with $T \in \left\{50\%\overline{T}, 60\%\overline{T}, 70\%\overline{T}, 80\%\overline{T}\right\}$.

The revenue of the heuristic policies for each problem instance is obtained by simulation as detailed below: in each replication, a stream of arrivals are randomly generated and then at each time epoch a policy is implemented to select one project for promotion. The backing probabilities are then updated, upon which one project that the backer is to support is sampled, followed by drawing a pledge from the pledging probabilities that the backer is to purchase. This process repeats until the end of the time horizon at which the total funds raised for each project are tallied. For those projects that are not completed, all the money is returned to the backers. To reduce the sampling variation, the revenue of each policy is averaged across 2,000 replications. Besides revenue, two other important performance metrics are also considered: a) the completion rate for each project defined to be the proportion of all the replications where the project is completed, and b) the average percentage of the raised funds compared to the funding goal of each project, regardless of being completed or not. All the experiments are executed with high performance computing clusters consisting of multiple nodes, each with Intel Xeon X5650 CPUs and 24GB RAM.

5.3. *Results*

5.3.1. *The Running Example*

First, we return to the running example and add two more projects, a smaller one and a larger one. All the parameters are shown in Table 2, where the global parameters β_1, $\beta_{2,t}$, $\beta_{3,3}$ and λ are as specified in Section 5.2. For the project-specific parameters, we set $m_1 < m_2 < m_3$, $G_1 < G_2 < G_3$ and $\bar{r}_1 < \bar{r}_2 < \bar{r}_3$. Recall that the campaign duration is $T = 12$ and the problem is small enough to obtain the optimal policy by solving the original DP in equation (2).

Table 2: The parameter setting in the running example.

Project	β_1	$\beta_{2,t}$	$\beta_{3,3}$	T	λ	G	m	$F(r=1)$	$F(r=2)$
1						1	0.05	0.8	0.2
2	1	$0.01 + 0.01t/T$	2.5	12	0.7	2	0.1	0.5	0.5
3						4	0.2	0.48	0.52

For each policy, we calculate its revenue using the simulation method described in Section 5.2. The percentage gaps to the optimal revenue (i.e., sub-optimalities) are shown in Table 3. In addition, Table 3 also reports the CPU time when solving the original DP and when developing the three dynamic policies. For the IP policy, it is the time to calculate all the index values for a single project; for LR it is the time to solve the Lagrangian-dual problem (12); while for DIP it is the total time to solve problem (13) at every time epoch. The computational time for IPx is negligible, thanks to the closed-form approximation. Also, the computational time for all the simple heuristics is negligible. Note that DIP policy reoptimizes at each time epoch for this small problem.

Table 3: The revenue performance and the CPU time in the running example.

Solutions	Optimal(DP)	DIP	LR	IP	IPx	SSF	LSF	SUF	LUF	GP	CP	MP
Sub-optimality in %	0	31.569	10.236	2.882	5.816	8.104	16.294	17.820	15.599	14.310	11.665	11.682
CPU time (seconds)	4.232	0.536	0.083	0.003	Negligible							

We report sub-optimality in percentage (%), i.e., $(V_{DP} - V_\pi)/V_{DP} \cdot 100\%$, where V_π is the revenue of policy π obtained by the simulation.

It is shown in Table 3 that the two index policies comfortably beat the other alternatives and the sub-optimality of IP/IPx is only 2.882%/5.816%, significantly lower than the others. Neither LR nor DIP shows strong performance compared to the simple heuristics. In fact, DIP yields the lowest revenue among all the policies, which is not surprising as the approximation error of DIP increases with the campaign duration. Note that $T = 12$ is relatively long compared to the funding goals. Moreover, IP is significantly faster than the exact DP algorithm; it is also faster than LR and DIP. The scalability of IP and especially IPx is clearly better than the latter two alternatives.

5.3.2. Larger Problem Instances

It is no longer possible to obtain the optimal policy for the larger problem instances specified in Section 5.2. To measure the performance of each alternative policy π against the index policy IP, we report its revenue gap relative to that of IP, i.e., $(V_{IP} - V_\pi)/V_{IP} \cdot 100\%$, where V_π is the revenue of policy π obtained by simulation. For DIP, we explored different re-optimization intervals. Although more frequent re-optimization may lead to some improvement, it is accompanied by a significant increase in running time. Consequently, we set a re-optimization interval of 25 time epochs in all the subsequent numerical experiments.

The results for the $J = 3$ problem with different choices of campaign duration are presented in Table 4A. Overall, the two proposed index policies demonstrate superior and robust performance compared to the alternative heuristics in most of the scenarios considered. They are outperformed only in the scenario of minimal campaign duration (i.e., 40% of $\overline{T}$), where SSF and SUF yield the highest revenue. In such a case the campaign duration is so short that larger projects would never be able to complete, and thus the promotion should focus on the smaller ones, as SSF and SUF do. The other two dynamic policies, DIP and LR, show respectable performance as well, beating the simple heuristics in all but the first scenario. As the campaign duration increases, the performance (relative to IP) of LSF, LUF, GP, CP, and MP gradually improves, while the performance of SSF, SUF, DIP and LR shows a downwards trend. It is also shown that with longer campaign durations (e.g., 70% of $\overline{T}$), the revenue gap among the policies narrows, as all the projects have a better chance to reach completion. Similar results and performance patterns are observed in the results for $J = 5$, as shown in Table 4B. The main difference is that in the shorter duration scenarios, the index policies are also outperformed by DIP and LR, both of which produce strong performance in these cases.

Table 4A: Extra revenue in % of the IP over the alternatives for the $J = 3$ problem with different campaign durations.

Percentile of $\overline{T}$	Duration	SSF	LSF	SUF	LUF	GP	CP	MP	IPx	DIP	LR
40%	500	-5.132	91.727	-5.174	95.361	93.773	23.243	94.182	-1.312	1.038	0.440
50%	625	25.534	84.101	29.040	87.369	86.532	16.062	88.002	9.879	0.707	8.536
60%	750	26.141	75.249	32.935	63.543	61.838	10.313	58.425	3.497	2.761	2.890
70%	875	24.533	52.119	17.888	12.046	10.139	14.640	8.931	-0.667	6.398	9.531

We report revenue gap in percentage (%), i.e., $(V_{IP} - V_{\pi})/V_{IP} \cdot 100\%$, where where V_{π} is the revenue of policy π obtained by the simulation.

Table 4B: Extra revenue in % of the IP over the alternatives for the $J = 5$ problem with different campaign durations.

Percentile of $\overline{T}$	Duration	SSF	LSF	SUF	LUF	GP	CP	MP	IPx	DIP	LR
50%	975	-3.397	52.513	-23.026	55.286	54.379	-19.517	55.027	4.734	-55.912	-36.144
60%	1170	15.926	54.966	8.473	35.039	37.043	-14.061	36.012	-22.322	-29.493	-18.970
70%	1365	36.178	64.588	34.832	15.990	17.485	18.535	15.764	7.297	1.636	19.488
80%	1560	39.305	56.369	37.319	10.502	15.612	21.059	10.083	11.612	5.450	11.468

We report revenue gap in percentage (%), i.e., $(V_{IP} - V_{\pi})/V_{IP} \cdot 100\%$, where where V_{π} is the revenue of policy π obtained by the simulation.

Fig. 2 provides an illustration of the completion rate (left bars) and the average percentage of funds raised (right bars) for each project with different choices of campaign duration. When the duration is short, as shown in Fig. 2 (a), at most 1/3 of the projects are successful under the simple policies, which focus too much on either the small project 1 or the large project 3 (that fails anyway). In contrast, the four dynamic policies (i.e., IP/IPx/DIP/LR) spend more effort in promoting the moderate project 2, resulting higher completion rates of this project and thus higher revenue overall.

With longer duration, there is an increase in either the completion rate, the percentage of raised funds, or both. In the longest duration scenario, the small project would complete regardless of the policies, as shown in Fig. 2 (c). In such a case, IP and IPx achieve 100% completions for both project 1 and project 2. Moreover, the percentage of funds raised for project 2 reaches 180% under IPx, leading to the highest revenue among all the heuristics. CP also achieves 100% completions for projects 1 and 2, but it spends too much time on promoting the small project and thus collects less funds for project 2. In contrast, MP, GP and LUF prioritize the large project and they do manage to achieve more than 60% completion rate, but this is at the expense of no completion of project 2 at all, resulting in less revenue. Similarly, LR and DIP also try to boost project 3, which however fails to complete anyway in most cases. The other two policies SUF and SSF prioritize project 1 too much. Even though the percentage of funds raised is over 300%, there is no revenue from the other two projects. The worst one is LSF, which spends a lot of effort to promote the largest project but before its completion switches to the others, leading to low completion rates for all the projects.

5.3.3. Robustness of Policies

To further investigate the robustness of the index policies, we now focus on the $J = 3$ scenario with $T = 875$ (the baseline) and perturb some key problem parameters. For the global parameters, we vary the promotion power and the herding effect by certain percentages, while for the local parameters we consider both smaller and larger differences of the attraction and pledging probability between the projects, as shown in Table 5.

The results are presented in Table 6, where the baseline scenario is highlighted in boldface. The perturbation of promotion power and herding effect clearly impacts on the performance of all the policies. With an increased promotion power, the performance (relative to IP) of most simple policies (except CP) improves, and those priori-

Figure 2: The completion rate and the percentage of funds raised for the $J = 3$ problem with different campaign duration.

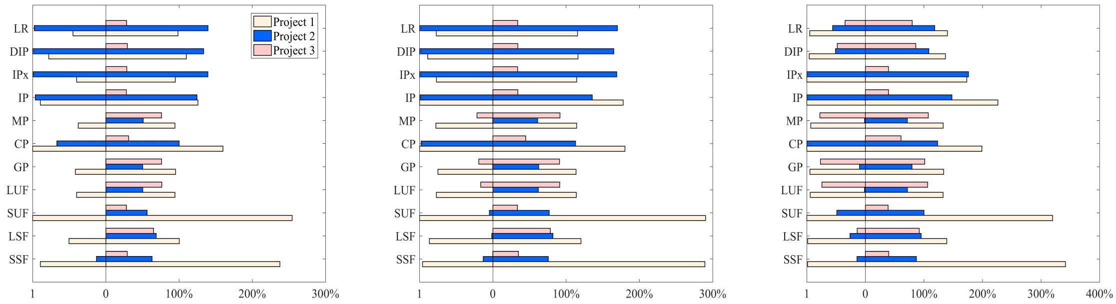

(a) $T = 625$ (b) $T = 750$ (c) $T = 875$

Note that in each graph the left bars refer to the completion rates while the right bars the average percentages of funds raised under different policies.

Table 5: Local parameter perturbation for $J = 3$ scenario with $T = 875$.

		Pledging probability $F(r)$			Attraction m		
Projects		1	2	3	1	2	3
Smaller differences	$r = 1$	0.54	0.5	0.49	0.05	0.08	0.15
	$r = 2$	0.46	0.5	0.51			
Larger differences	$r = 1$	0.56	0.5	0.47	0.03	0.08	0.17
	$r = 2$	0.44	0.5	0.53			

tizing the promotion of larger projects, such as LUF, GP, and MP, demonstrate significant improvement. Note that the performance of LR and DIP improves with a lower promotion power as well. Meanwhile, the increase in the herding effect also positively influences the performance of all the policies except IP and IPx, whose performance slightly reduces (more discussion on this is provided in the next section), leading to narrower revenue gaps among the policies. For most of them the change in performance is relatively small. The only exception is SUF, which is the second weakest for $\gamma = 0$ but the second strongest for $\gamma = 0.03$. The perturbation of the local parameters also influences the performance of the policies to some degree, though not significantly.

Table 6: Extra revenue in % of the IP over the alternatives for the $J = 3$ problem with parameter perturbation.

Scenarios	Perturbation	SSF	LSF	SUF	LUF	GP	CP	MP	IPx	DIP	LR
Promotion power β_1											
1.1	+10%	23.449	44.108	14.359	1.010	0.035	17.450	3.638	-0.861	1.065	6.458
0.9	-10%	26.826	65.703	30.437	43.090	47.065	11.236	45.065	-0.082	0.013	1.167
Herding parameter $\beta_{2,t} = 0.01 + \gamma \cdot (t/T)$											
γ											
0.00	-100%	26.549	59.020	29.539	23.235	22.419	14.967	19.742	0.157	11.915	13.093
0.01	**0%**	**24.533**	**52.119**	**17.888**	**12.046**	**10.139**	**14.640**	**8.931**	**-0.667**	**6.398**	**9.531**
0.03	+200%	20.853	54.630	-1.049	15.872	15.088	11.612	16.710	-2.011	2.409	3.910
Attraction m											
Smaller Differences		24.570	60.991	16.874	21.958	23.675	13.824	25.126	1.794	4.631	11.710
Larger Differences		25.988	56.717	30.772	15.127	16.879	14.273	17.519	0.053	8.490	0.632
Pledging probability $F(r)$											
Smaller Differences		22.959	59.450	19.342	18.718	21.729	12.137	19.477	-2.041	8.499	9.233
Larger Differences		21.936	55.643	19.367	12.840	11.822	9.329	11.389	2.154	7.838	7.745

We report revenue gap in percentage (%), i.e., $(V_{IP} - V_{\pi})/V_{IP} \cdot 100\%$, where where V_{π} is the revenue of policy π obtained by the simulation.

Next, we turn our attention to the impact of the side-effect parameter $\beta_{3,J}$ on the performance of the two proposed policies, i.e. IP and IPx. From the discussion in Section 4, we see that for single-project sub-problems, the two proposed policies do not involve parameter $\beta_{3,J}$. This may potentially be a weakness of the proposed policies. To assess the optimality gap between the proposed policies and the optimal DP solution across different values of $\beta_{3,J}$, we consider several small problem instances, using the baseline parameter setting shown in Table 7. Table 8 displays the sub-optimality of the two proposed policies for $\beta_{3,3} = 2.5$ (the baseline setting) with several levels of perturbation where $\beta_{3,3}$ is increased by a percentage, ranging from 5% to 20%. From Table 8, we can see that the two proposed policies exhibit weaker performance for increased $\beta_{3,3}$. This is primarily because a higher $\beta_{3,J}$ reduces the customer's backing probability for all the projects (see equation (1)), making it more difficult for the policies that favor large projects to reach their funding goals.

Table 7: Scenarios of the baseline parameter setting for the $J = 3$ problems.

Instance	Project	β_1	$\beta_{2,0}$	$\beta_{3,3}$	T	λ	G	m	$F(r=1)$	$F(r=2)$
	1						1	0.05	0.8	0.2
1	2	1	0.01	2.5	12	0.7	2	0.1	0.5	0.5
	3						4	0.2	0.48	0.52
	1						1	0.05	0.99	0.01
2	2	1	0.01	2.5	13	0.7	2	0.1	0.92	0.08
	3						4	0.2	0.9	0.1

From a practical perspective, since the duration of a campaign is dependent on customers' purchasing probability which in turn is affected (negatively) by the side effect, one solution to this problem is to adjust the campaign duration to compensate for a strong side effect, whenever it exists for the current projects.

Table 8: Sub-optimality in % of the IP/IPx for the $J = 3$ problem with different $\beta_{3,3}$.

		$\beta_{3,3} = 2.5$ and its perturbation in %				
Instance	Policies	2.5 (0%)	2.625 (+5%)	2.75 (+10%)	2.875 (+15%)	3 (+20%)
1	IP	2.832	5.505	4.904	5.365	7.810
	IPx	8.401	4.136	5.133	5.473	9.133
2	IP	4.396	4.523	5.104	8.835	11.593
	IPx	4.967	6.170	6.859	7.835	11.892

We report sub-optimality in percentage (%) relative to the exact solution obtained by DP , i.e., $(V_{DP} - V_\pi)/V_{DP} \cdot 100\%$, where V_{DP} is the revenue of the optimal solution obtained by DP and V_π is the revenue of policy π obtained by the simulation.

5.4. Further Discussion on IP and IPx

So far we have discussed the performance of the two index policies, IP and IPx. Next, to gain deeper insights into the nuances of them, we closely examine the index policies for the $J = 3$ problem. Our analysis begins by showing how their index values change over time and state. Fig. 3 (a) - (c) and (d) - (f) depict the index values of IP and IPx, respectively, for all the projects in the baseline scenario at $t = 300, 500,$ and 700. It is evident that, for both IP and IPx, the index values of all the projects (at each time point t) decrease with lower shortfalls (or equivalently, increase with state g), as indicated by Lemma 4.3 (i). Additionally, the rate of descent for project 1 (the smallest project) is higher than that of project 2 (the moderate project), which, in turn, is higher than that of project 3 (the largest project). This is due to the difference in funding goals, where the value of promoting a smaller project decreases more rapidly than that of a larger project as more funds are raised. In these cases, the herding effect gradually accumulates for both projects, but more prominently for the smaller one. Furthermore, for the same state g, the index values of all the projects decrease with time t, as indicated by Lemma 4.3 (ii). The rate of descent over time is also higher for the smaller projects, signifying a faster drop in the value of promotion for these projects as their herding effect becomes more prominent when shortfalls gradually diminish over time.

Figure 3: The index values of IP and IPx at $t = 300$, 500 and 700 in the $J = 3$ baseline scenario.

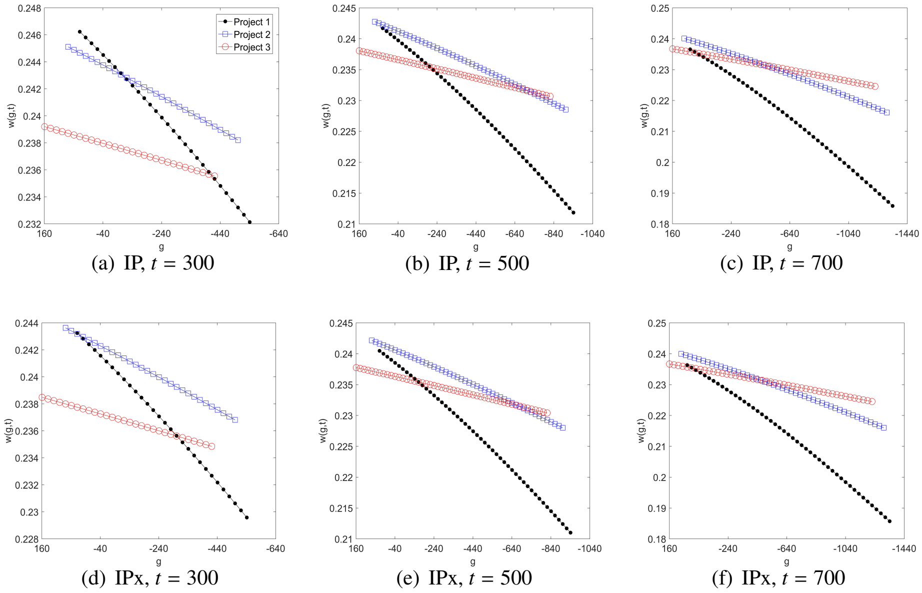

(a) IP, $t = 300$ (b) IP, $t = 500$ (c) IP, $t = 700$

(d) IPx, $t = 300$ (e) IPx, $t = 500$ (f) IPx, $t = 700$

Between the two index policies, it is evident that although IPx is an approximation of IP, its index values exhibit a similar trend to those obtained by IP. Moreover, the two sets of index values gradually converge as time increases. For example, even though the index values by IP and IPx show different rates of descent over state at $t = 300$, their values become almost the same at $t = 700$. The above observation indicates that IP and IPx would converge to each other over time.

Figure 4: The actions of IP and IPx at time $t = 300$, 500 and 700 when $g_1 = 10$ in the $J = 3$ baseline scenario.

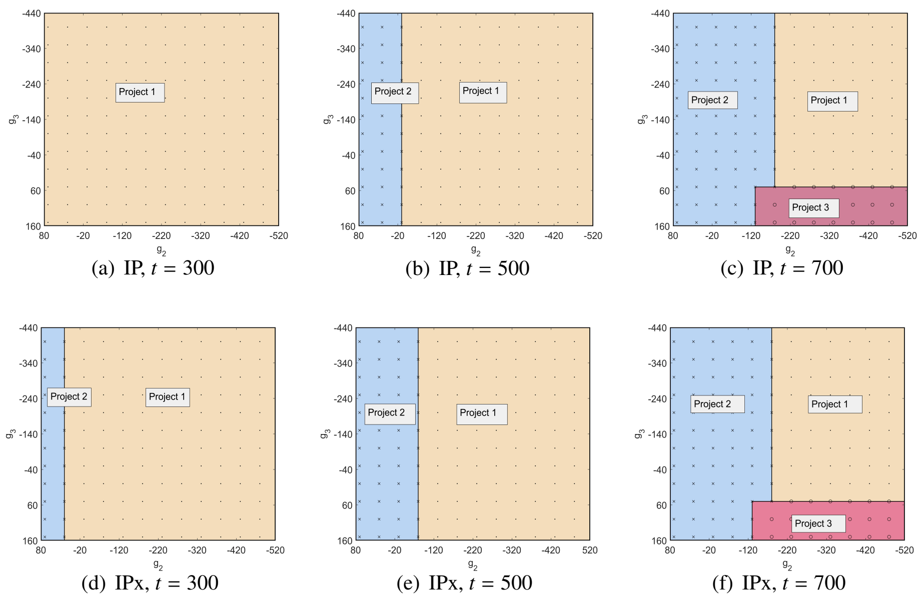

(a) IP, $t = 300$ (b) IP, $t = 500$ (c) IP, $t = 700$

(d) IPx, $t = 300$ (e) IPx, $t = 500$ (f) IPx, $t = 700$

Building on the insights gained from Fig. 3, we further analyze the allocation of promotion under each policy at different states and time, as illustrated in Fig. 4. Specifically, we plot the actions by IP and IPx in the baseline scenario at $t = 300, 500$ and 700 for a selected state of project 1, i.e., $g_1 = 10$. Each graph is two-dimensional, representing the states of projects 2 and 3. It is clear that both IP and IPx gradually shift their promotion priority towards larger projects over time, evidenced by the shrinking area of project 1 and growing areas of project 2 and 3. Moreover, the switches occur earlier under IPx than IP, and thus the former sees bigger areas for larger projects at each time point (see Fig. 4 (a)-(b), (d)-(e)). Towards the end of the campaign the actions become the same between these two policies, as shown in Fig. 4 (c) and (f). We also plot the average actions for the project being promoted by IP and IPx, respectively, over time across all the simulation replications for the baseline scenario in Fig. 5; it clearly shows this pattern.

Figure 5: Average IP/IPx actions for the project being promoted over time across all simulation replications for the baseline scenario.

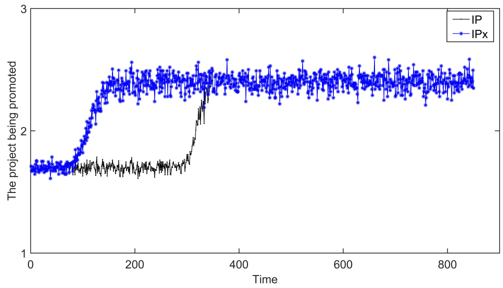

The earlier switching behavior, as exhibited in Fig. 5, helps IPx achieve higher revenue when the herding effect increases quickly over time. In such cases, the impact of promotion is dominated by the herding effect after a short time for smaller projects. Therefore, switching earlier to larger projects that need more boost would increase their chance of completion, leading to higher revenue overall. This also explains why the performance of the index policies reduces with faster growing herding effect (see again Table 6), as in such cases even IPx switches to larger projects later than preferred. On the other hand, when the herding effect is small and does not increase much over time, switching too early would lead to detrimental outcomes as the smaller projects may not be able to complete themselves, which explains (again) why the performance of the index policies (and especially IP) increases in such cases.

5.5. Different Start and End Times Across Projects

In this section we relax the assumption of common duration for all the campaigns and allow them to be flexible, i.e., they may start and end at different times. Both index policies are readily applicable to such problems, thanks to the decomposition, and hence the index values for each project are still calculated by Proposition 4.3 or equation (11). DIP is modified by forcing the backing probabilities to zeros before and after the campaign of each project. Accordingly, LR is modified by summing over each project's campaign duration when calculating the sub-gradient, and by considering only the live projects when determining promotion actions. The extension of all the simple heuristics is straightforward. We investigate the numerical performance of the alternative heuristic policies against the index policies. The revenue of all the policies are computed via the same simulation method as described in Section 5.2. The only difference is that at each decision epoch, only live projects are considered for promotion.

In our experiments, we consider a problem with $J = 7$ projects and their project-specific parameters are displayed in Table 9. For the global parameters, the promotion power and arrival rate remain as $\beta_1 = 1$ and

$\lambda = 0.7$. The herding parameter still takes the form $\beta_{2,t} = 0.01 + 0.01t/T$. The side effect is set to be a piecewise function of J, i.e. (a) $\beta_{3,J} = 0$ if $J = 1$, (b) $\beta_{3,J} = 2$ if $J = 2, 3$, and (c) $\beta_{3,J} = 3$ if $J \geq 4$.

Table 9: Local parameters for the $J = 7$ problem.

Project j	1	2	3	4	5	6	7
m	0.02	0.04	0.06	0.08	0.1	0.12	0.16
G	20	40	60	80	100	120	160
$F(r = 1)$	0.7	0.6	0.55	0.5	0.48	0.46	0.45
$F(r = 2)$	0.3	0.4	0.45	0.5	0.52	0.54	0.55
$\bar{r}$	1.3	1.4	1.45	1.5	1.52	1.54	1.55

The duration of each project is determined by $T_j = G_j / \left(\bar{r}_j \lambda \theta_j\right)$, where θ_j is a control parameter dependent on the average backing probability of each arriving customer. Therefore, T_j is the duration required to complete project j under the given θ_j; the smaller θ_j, the longer T_j, and vice versa. For simplicity we choose $\theta_j = 0.1$ for all projects in our experiments, which means that 1 out of 10 customers will back a project on average. We consider three sequences of how projects appear over time, as shown in Fig. 6, which represent different situations that a platform manager might encounter in practice. The numbers in the brackets represent the start and end time for each project. Table 10 reports the percentage revenue of IP over the other policies. The results indicate that the

Figure 6: Three sequences of how projects appear over time.

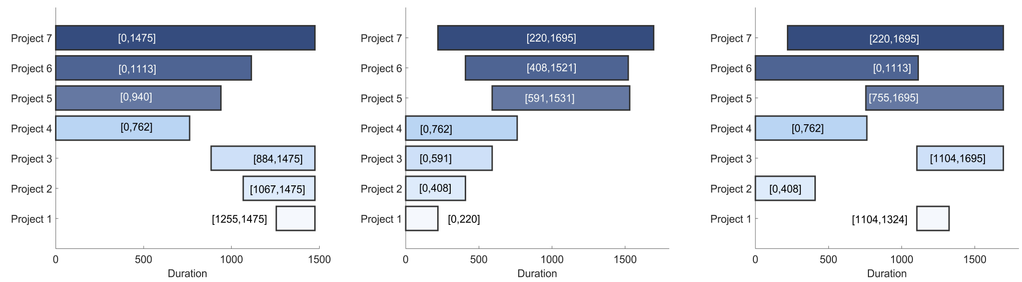

(a) Scenario 1: larger-projects-first (b) Scenario 2: smaller-projects-first (c) Scenario 3: random

index policies comfortably outperform all the other alternatives in most scenarios considered. The only exception is the smaller-projects-first scenario, where DIP produces higher revenue by a small margin of 3.78%. Between the two index policies, IPx is clearly the better one in the larger-projects-first scenario, where the reverse is the case in the smaller-projects-first scenario. There is no clear difference between them in the random scenario.

Table 10: Extra revenue in % of the IP over the alternatives for the $J = 7$ problem with different campaign timelines.

Scenario	SSF	LSF	SUF	LUF	GP	CP	MP	IPx	DIP	LR
1	29.103	56.769	30.292	5.405	3.439	16.888	5.096	-9.357	17.023	25.362
2	18.418	25.102	23.817	10.501	0.932	11.170	9.969	10.770	-3.780	15.325
3	13.234	26.496	8.669	15.362	18.670	11.980	15.256	-0.540	27.987	35.657

We report revenue gap in percentage (%), i.e., $(V_{IP} - V_{\pi})/V_{IP} \cdot 100\%$, where where V_{π} is the revenue of policy π obtained by the simulation.

In conclusion, the numerical analysis in this section shows that the two proposed index policies outperform the other alternatives, leading to higher revenue and more completions of crowdfunding projects in most scenarios considered. DIP and LR perform reasonably well but they incur a much higher computational cost. In addition, although there are a few simple heuristic policies exhibiting superior performance in one scenario or another, in general their performance is weak and sensitive to the problem parameters. In contrast, IP/IPx exhibits robustness with strong performance in all scenarios considered.

6. Concluding Remarks

In recent years, online-platform based market analytics (e.g., algorithms for web advertising and display ranking) that concerns the management of demand and revenue under limited resources has become increasingly important for revenue management research. In this paper, we aim to maximize the platform's expected total profit while addressing the challenge of low success rates in crowdfunding projects. In practice, one useful lever for platforms is to use homepage promotions to highlight projects on their homepage to increase their visibility and boost their chance of success. Given the limited homepage space, a key decision faced by a platform is which project to promote at each time epoch throughout a campaign. To address this research challenge, we have developed a mathematical framework to model the crowdfunding process, with an MNL model to describe backers' choice behaviors among multiple competing projects and a Markov decision process to capture the dynamics of the process. To address the computational complexity of the resulting dynamic program, we have developed two index policies following the celebrated Whittle's RB approach. Importantly, the developed approach is applicable to a wide range of backers' utility functions. In fact, no particular specification for the functional form of backers' utility is required, provided that its corresponding choice probability satisfies some general conditions. Our numerical experiments show that the proposed index policies outperform a number of benchmark heuristics in most of the scenarios considered.

From a methodological point of view, our work is an innovative application of Whittle's RB approach to a class of problems where the bandits are not only coupled by resources, but also by customer demand. Indeed, in our proposed mathematical model for crowdfunding processes, the backing probability of one project depends on the features and the funding progress of all the other projects via the MNL model. The incorporation of the MNL into the mathematical model for crowdfunding processes, however, brings about some challenges in the problem decomposition of the RB approach. To address this complexity, we have proposed to replace the MNL with a collection of BNL models, one for each project. Such a relaxation not only allows the original problem to be decomposed into a number of single-project problems, but also preserves the ranking among the projects (as measured by backers' choice probabilities). Mathematically, each decomposed problem has a state-dependent backing probability due to the herding effect which is an important characteristic to be considered in crowdfunding, and this state-dependency has introduced additional complexities in the indexability analysis. In this paper, we have proved that a project is indexable when the campaign duration is sufficiently long and shown that the index values, which are bi-dimensional, can be conveniently derived from the value-to-go under the non-promotion policy for each project. To further reduce the computational time we have developed a close-form approximation to the index values. Our results show that the approximated values are close to the exact index values and the resulting index policy yields comparable performance to that of the true index values.

From a managerial perspective, our study yields some useful insights on homepage promotions for crowdfunding platforms. We have shown that the length of campaigns is an important factor to consider when deciding which project to promote. If the campaign duration is short, small projects should be promoted first until their completions; otherwise, large projects should be promoted first until their completions. For a moderate campaign duration, it is preferable to promote the smallest projects first, and once the funds are accumulated to a certain level (but not yet completed) and the herding effect takes over, the promotion should switch to larger projects. This paper has also demonstrated that the index policies can be of practical use to provide tailored recommendations regarding the project selection and the timing of switches for promotion purposes. Once the index values are evaluated, the platform managers can simply pick up the project with the largest index value to promote at each time epoch; the index value, which may be interpreted as the fair charge for promotion, is also easy to understand by the managers. More importantly, the index policies, and in particular the one based on the closed-form approximation, are scalable and the index values can be quickly re-evaluated online as needed, which holds particular significance in practical scenarios where there are usually many projects competing against each other for investment/resources at the same time.

Finally, we outline a few potential future research directions. First, we note that in Section 5, we investigated the performance of an LR method with a single Lagrangian multiplier for all time epochs. The main takeaway is that the LR is outperformed by the proposed IP policies in most of the instances considered. The LR's performance might be improved by using the LR method developed in Brown and Smith (2020), where more granular, time-dependent Lagrangian multipliers are used. Brown and Smith (2020) show that their policies are asymptotically optimal under mild conditions, although the run time increases rapidly with time horizon and hence these polices may not be scalable for problems with a long time horizon, such as the ones considered in this paper. Nevertheless, it would be an interesting future research topic to formally investigate the performance and runtime of the LR policies and tie-breaking rules proposed in Brown and Smith (2020) in crowdfunding applications, in which the decoupling of demand by replacement of MNL with BNL might hinder the performance of LR policies, even with more granular Lagrangian multipliers.

Next, we note that Zhang et al. (2023) have recently investigated the optimal campaign duration, although they consider a single-project problem only. They suggest that its length should be tailored to the composition of the backer population, i.e., longer campaigns are better suited to herding backers (namely backers who are affected by herding effect), while shorter campaigns cater to independent backers. It is of interest to extend their research from single-project problems to multiple projects, and explore to how much to adjust the campaign duration to compensate for a large side effect for the case of multiple projects.

One limitation of this work is the assumption of MNL choice model for investors' decision-making. Although widely used in practice, the MNL model imposes some restrictive conditions, e.g., investors' utility follows the Gumbel distribution. Such conditions may not always hold in practice. In future research, one can address this issue and hence extend our analysis by incorporating a generalized choice model, e.g., Li (2011).

Another interesting direction of future research is to consider personalized promotion strategies. In practice, many backers are returning customers, and platforms may recommend personalized projects to individual backers. It is also of interest to consider homepage promotion strategies in other types of crowdfunding, such as donor-based crowdfunding, where the objective is to raise enough money for the most urgent projects first. Besides homepage promotions, our modeling approach to crowdfunding processes may also be employed to study other problems such as the optimal design of the reward schemes. Beyond crowdfunding, it would be interesting to extend the proposed methodology to other applications of Whittle's RB where bandits are coupled by both resource and demand.

Acknowledgment

The authors thank the three anonymous reviewers for their insightful and valuable comments that have substantially improved the quality of this paper.

References

Adelman, D., Mersereau, A.J., 2008. Relaxations of weakly coupled stochastic dynamic programs. Operations Research 56, 712–727.

Akçay, Y., Natarajan, H.P., Xu, S.H., 2010. Joint dynamic pricing of multiple perishable products under consumer choice. Management Science 56, 1345–1361.

Allon, G., Babich, V., 2020. Crowdsourcing and crowdfunding in the manufacturing and services sectors. Manufacturing & Service Operations Management 22, 102–112.

Balakrishnan, N., Nevzorov, V., 2004. A primer on statistical distributions. John Wiley & Sons, Inc.

Bertsimas, D., Mišić, V.V., 2016. Decomposable markov decision processes: A fluid optimization approach. Operations Research 64, 1537–1555.

Brown, D.B., Smith, J.E., 2020. Index policies and performance bounds for dynamic selection problems. Management Science 66, 3029–3050.

Brown, D.B., Zhang, J., 2023. On the strength of relaxations of weakly coupled stochastic dynamic programs. Operations Research 71, 2374–2389.

Chan, H.F., Moy, N., Schaffner, M., Torgler, B., 2021. The effects of money saliency and sustainability orientation on reward based crowdfunding success. Journal of Business Research 125, 443–455.

Clifford, C., 2016. Less than a third of crowdfunding campaigns reach their goals. `https://www.entrepreneur.com/article/269663`. Accessed: 08/05/2024.

Crosetto, P., Regner, T., 2018. It's never too late: Funding dynamics and self pledges in reward-based crowdfunding. Research Policy 47, 1463–1477.

Dai, H., Zhang, D.J., 2019. Prosocial goal pursuit in crowdfunding: Evidence from kickstarter. Journal of Marketing Research 56, 498–517.

Du, C., Cooper, W.L., Wang, Z., 2016. Optimal pricing for a multinomial logit choice model with network effects. Operations Research 64, 441–455.

Du, L., Hu, M., Wu, J., 2022. Contingent stimulus in crowdfunding. Production and Operations Management 31, 3543–3558.

Du, S., Peng, J., Nie, T., Yu, Y., 2020. Pricing strategies and mechanism choice in reward-based crowdfunding. European Journal of Operational Research 284, 951–966.

Du, Z., Li, M., Wang, K., 2019. the more options, the better? investigating the impact of the number of options on backers decisions in reward-based crowdfunding projects. Information & Management 56, 429–444.

Gerber, E.M., Hui, J., 2013. Crowdfunding: Motivations and deterrents for participation. ACM Transactions on Computer-Human Interaction (TOCHI) 20, 1–32.

Gittins, J., Glazebrook, K., Weber, R., 2011. Multi-armed bandit allocation indices. John Wiley & Sons.

Glazebrook, K.D., Hodge, D.J., Kirkbride, C., 2011. General notions of indexability for queueing control and asset management. The Annals of Applied Probability 21, 876–907.

Glazebrook, K.D., Hodge, D.J., Kirkbride, C., Minty, R., 2014. Stochastic scheduling: A short history of index policies and new approaches to index generation for dynamic resource allocation. Journal of Scheduling 17, 407–425.

Glazebrook, K.D., Mitchell, H., Ansell, P., 2005. Index policies for the maintenance of a collection of machines by a set of repairmen. European Journal of Operational Research 165, 267–284.

Graczová, D., Jacko, P., 2014. Generalized restless bandits and the knapsack problem for perishable inventories. Operations Research 62, 696–711.

Hawkins, J.T., 2003. A Langrangian decomposition approach to weakly coupled dynamic optimization problems and its applications. Ph.D. thesis. Massachusetts Institute of Technology.

Kickstarter, 2015. How to get featured on kickstarter. `https://www.kickstarter.com/blog/how-to-get-featured-on-kickstarter`. Accessed: 08/05/2024.

Kickstarter, 2023. Creator handbook - building rewards. `https://www.kickstarter.com/help/handbook/rewards`. Accessed: 08/05/2024.

Kickstarter, 2024. What happens when a project is overfunded? https://help.kickstarter.com/hc/en-us/articles/115005132473-What-happens-when-a-project-is-overfunded. Accessed: 08/05/2024.

Kim, C., Kannan, P., Trusov, M., Ordanini, A., 2020. Modeling dynamics in crowdfunding. Marketing Science 39, 339–365.

Lee, J., Parlour, C.A., 2022. Consumers as financiers: Consumer surplus, crowdfunding, and initial coin offerings. The Review of Financial Studies 35, 1105–1140.

Li, B., 2011. The multinomial logit model revisited: A semi-parametric approach in discrete choice analysis. Transportation Research Part B: Methodological 45, 461–473.

Li, D., Ding, L., Connor, S., 2020a. When to switch? index policies for resource scheduling in emergency response. Production and Operations Management 29, 241–262.

Li, Z., Duan, J.A., Ransbotham, S., 2020b. Coordination and dynamic promotion strategies in crowdfunding with network externalities. Production and Operations Management 29, 1032–1049.

Liang, X., Hu, X., Jiang, J., 2020. Research on the effects of information description on crowdfunding success within a sustainable economythe perspective of information communication. Sustainability 12, 650.

Maghsudi, S., Hossain, E., 2016. Multi-armed bandits with application to 5G small cells. IEEE Wireless Communications 23, 64–73.

Malekipirbazari, M., Çavuş, Ö., 2024. Index policy for multiarmed bandit problem with dynamic risk measures. European Journal of Operational Research 312, 627–640.

Niño-Mora, J., 2007. Dynamic priority allocation via restless bandit marginal productivity indices. Top 15, 161–198.

Nosrat, F., 2022. Assortment optimization with customer choice modeling in a crowdfunding setting. arXiv preprint arXiv:2207.07222 .

Papadimitriou, C.H., Tsitsiklis, J.N., 1994. The complexity of optimal queueing network control, in: Proceedings of IEEE 9th Annual Conference on Structure in Complexity Theory, IEEE. pp. 318–322.

Song, Y., Li, Z., Sahoo, N., 2022. Matching returning donors to projects on philanthropic crowdfunding platforms. Management Science 68, 355–375.

Varma, V., Salahaldin, L., Elayoubi, S.E.E., 2021. A real options model for crowdfunding platforms: optimal project promotion strategies, in: 24th Annual International Real Options Conference.

Weber, R.R., Weiss, G., 1991. Addendum to 'On an index policy for restless bandits'. Advances in Applied Probability 23, 429–430.

Weinmann, M., Mishra, A.N., Kaiser, L.F., vom Brocke, J., 2023. The attraction effect in crowdfunding. Information Systems Research 34, 1276–1295.

Whittle, P., 1988. Restless bandits: Activity allocation in a changing world. Journal of Applied Probability 25, 287–298.

Xiao, S., Ho, Y.C., Che, H., 2021. Building the momentum: Information disclosure and herding in online crowdfunding. Production and Operations Management 30, 3213–3230.

Yang, X., Strauss, A.K., Currie, C.S., Eglese, R., 2016. Choice-based demand management and vehicle routing in e-fulfillment. Transportation Science 50, 473–488.

Zhang, J., Savin, S., Veeraraghavan, S., 2023. Revenue management in crowdfunding. Manufacturing & Service Operations Management 25, 200–217.

Zhang, Y., Tian, Y., 2021. Choice of pricing and marketing strategies in reward-based crowdfunding campaigns. Decision Support Systems 144, 113520.

Zhou, M.J., Lu, B., Fan, W.P., Wang, G.A., 2018. Project description and crowdfunding success: an exploratory study. Information Systems Frontiers 20, 259–274.

Appendix A. The Alternative Dynamic Program Formulation

Define by $V'_t(\mathbf{g})$ the value function under the alternative formulation. We have:

$$V'_t(\mathbf{g}) = \max_{\mathbf{a}\in\mathcal{A}} \left\{ \lambda \sum_{j=1}^{J} p_{j,t}(\mathbf{g},\mathbf{a}) \sum_{r=1}^{R_j} F_j(r) V'_{t+1}(\tilde{\mathbf{g}}) + (1 - \lambda + \lambda p_{0,t}(\mathbf{g},\mathbf{a})) V'_{t+1}(\mathbf{g}) \right\}, \forall \mathbf{g} \in \Omega_t, 0 \le t \le T-1,$$

$$V'_T(\mathbf{g}) = \sum_{j=1}^{J} h'_T(g_j), \text{ where } h'_T(g_j) = \begin{cases} 0 & if \quad g_j > 0, \\ G_j - g_j & if \quad g_j \le 0. \end{cases}, \forall \mathbf{g} \in \Omega_T.$$

We have the following results.

Proposition A.1. *For all* $\boldsymbol{g}$ *and* t, $V_t(\boldsymbol{g}) = V_t'(\boldsymbol{g}) - \sum_{j=1}^{J}(G_j - g_j)$.

Proof. We prove it by induction. Clearly it is true for $t = T$ and any state $\mathbf{g}$. Suppose it is true for some epoch t. We then have $V_t(\mathbf{g}) = V_t'(\mathbf{g}) - \sum_{k=1}^{J}(G_k - g_k)$. Substituting it into Equation (2a), we have

$$\begin{aligned}
V_{t-1}(\mathbf{g}) &= \max_{\mathbf{a}\in\mathcal{A}}\left\{\lambda\sum_{j=1}^{J}p_{j,t}(\mathbf{g},\mathbf{a})\sum_{r=1}^{R_j}F_j(r)\left(r + V_t'(\tilde{\mathbf{g}}) - \sum_{k=1}^{J}(G_k-\tilde{g}_k)\right) + (1-\lambda+\lambda p_{0,t}(\mathbf{g},\mathbf{a}))\left(V_t'(\mathbf{g}) - \sum_{k=1}^{J}(G_k-g_k)\right)\right\}\\
&= \max_{\mathbf{a}\in\mathcal{A}}\left\{\lambda\sum_{j=1}^{J}p_{j,t}(\mathbf{g},\mathbf{a})\sum_{r=1}^{R_j}F_j(r)\left(V_t'(\tilde{\mathbf{g}}) - \sum_{k=1}^{J}(G_k-g_k)\right) + (1-\lambda+\lambda p_{0,t}(\mathbf{g},\mathbf{a}))\left(V_t'(\mathbf{g}) - \sum_{k=1}^{J}(G_k-g_k)\right)\right\}\\
&= \max_{\mathbf{a}\in\mathcal{A}}\left\{\lambda\sum_{j=1}^{J}p_{j,t}(\mathbf{g},\mathbf{a})\sum_{r=1}^{R_j}F_j(r)V_t'(\tilde{\mathbf{g}}) + (1-\lambda+\lambda p_{0,t}(\mathbf{g},\mathbf{a}))V_t'(\mathbf{g}) - \sum_{k=1}^{J}(G_k-g_k)\right\}\\
&= V_{t-1}'(\mathbf{g}) - \sum_{k=1}^{J}(G_k-g_k),
\end{aligned}$$

where the second equality is by incorporating r into $\tilde{g}_k$ and reversing the state transition, the third equality is by the definition of $p_{0,t}(\mathbf{g},\mathbf{a})$, and the last equality is by the definition of $V_t'(\mathbf{g})$. □

Appendix B. Proofs of the Theorems/Propositions/Lemmas.

To start with, we first define $q_t^W(g,a), \forall g \in \Omega_t, 0 \le t \le T-1$, as the value-to-go when action a is taken at state g and time t, and the optimal policy (under cost W) is followed thereafter, i.e.,

$$q_t^W(g,a) = \lambda p_t^a(g)\left(\bar{r} + \triangle v_{t+1}^W(g)\right) + v_{t+1}^W(g) - Wa,$$

which is termed the optimal q-factor in the literature. Therefore, equation (7) can be rewritten as

$$v_t^W(g) = \max_{a\in\{0,1\}}\left\{q_t^W(g,a)\right\}.$$

Appendix B.1. Proof of Proposition 3.1

We first introduce the following lemma.

Lemma B.1. *For any system state* $\mathbf{g}$ *and action* $\mathbf{a}$ *at time epoch* t*, we have:*

(i) if $p_{j,t}(\mathbf{g},\mathbf{a}) \ge p_{k,t}(\mathbf{g},\mathbf{a})$*, then* $p_{j,t}^{a_j}(g_j) \ge p_{k,t}^{a_k}(g_k)$ *for any projects* j *and* k*;*

(ii) $p_{j,t}(\mathbf{g},\mathbf{a}) \le p_{j,t}^{a_j}(g_j)$ *for any project* j*.*

Proof. For part (i), we obtain from $p_{j,t}(\mathbf{g},\mathbf{a}) \ge p_{k,t}(\mathbf{g},\mathbf{a})$ that

$$\exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right) - \beta_{3,J}\right) \ge \exp\left(m_k + \beta_1 a_k + \beta_{2,t}\left(1 - g_k/G_k\right) - \beta_{3,J}\right),$$

which is equivalent to

$$\frac{\exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right) - \beta_{3,J}\right)}{1 + \exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right) - \beta_{3,J}\right)} \ge \frac{\exp\left(m_k + \beta_1 a_k + \beta_{2,t}\left(1 - g_k/G_k\right) - \beta_{3,J}\right)}{1 + \exp\left(m_k + \beta_1 a_k + \beta_{2,t}\left(1 - g_k/G_k\right) - \beta_{3,J}\right)}.$$

Then, it follows immediately by noting that the above inequality leads to the following inequality:

$$\frac{\exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right)\right)}{1 + \exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right)\right)} \ge \frac{\exp\left(m_k + \beta_1 a_k + \beta_{2,t}\left(1 - g_k/G_k\right)\right)}{1 + \exp\left(m_k + \beta_1 a_k + \beta_{2,t}\left(1 - g_k/G_k\right)\right)}.$$

Next, for part (ii), consider the following two probabilities under the same action a_j:

$$p_{j,t}(\mathbf{g},\mathbf{a}) = \frac{\exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right) - \beta_{3,J}\right)}{1 + \exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right) - \beta_{3,J}\right) + \sum\limits_{k=1,k\neq j} \exp\left(z_{k,t}(g_k, a_k)\right)},$$

$$\tilde{p}^{a_j}_{j,t}(g_j) = \frac{\exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right) - \beta_{3,J}\right)}{1 + \exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right) - \beta_{3,J}\right)}.$$

It is straightforward to verify that $p_{j,t}(\mathbf{g},\mathbf{a}) < \tilde{p}^{a_j}_{j,t}(g_j)$, as $\exp\left(z_{k,t}(g_k, a_k)\right) > 0$. Next, consider the BNL-based backing probability $p^{a_j}_{j,t}(g_j)$ in equation (4):

$$p^{a_j}_{j,t}(g_j) = \frac{\exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right)\right)}{1 + \exp\left(m_j + \beta_1 a_j + \beta_{2,t}\left(1 - g_j/G_j\right)\right)}.$$

Since $\beta_{3,J} \geq 0$, it follows immediately that $\tilde{p}^{a_j}_{j,t}(g_j) \leq p^{a_j}_{j,t}(g_j)$. □

Proof of Proposition 3.1. For the first inequality $V_0(\mathbf{G}) \leq \hat{V}^W_0(\mathbf{G})$, we note that with $W \geq 0$ the last term on the right-hand-side of $\hat{V}^W_0(\mathbf{G})$, i.e. $W\sum_{t=0}^{T-1}\left(1 - \sum_{j=1}^{J}\tilde{\pi}_{j,t}(\mathbf{g}(t))\right)$, is always non-negative for all policies satisfying the original resource constraint. Since we have now relaxed the resource constraint to allow more than one projects to be promoted each time, the optimal value function will only increase further so the Lagrangian value function provides an upper bound on the true value function.

For the second inequality $\hat{V}^W_0(\mathbf{G}) \leq V^W_0(\mathbf{G})$, we note that the difference between the two value functions arises from their definitions in (3) and (5), where the probabilities calculated using the MNL model (1) for $\hat{V}^W_0(\mathbf{G})$ are replaced with those defined in (4) for $V^W_0(\mathbf{G})$. First, recall that equation (3) is given by:

$$\hat{V}^W_0(\mathbf{G}) = \sup_{\tilde{\pi}\in\tilde{\Pi}} \hat{V}^{\tilde{\pi},W}_0(\mathbf{G}) = \sup_{\tilde{\pi}\in\tilde{\Pi}} \mathbb{E}\left[\sum_{t=0}^{T-1} h_t\left(\mathbf{g}(t), \tilde{\pi}_t(\mathbf{g}(t))\right) + \sum_{j=1}^{J} h_T(g_j(T)) + W\sum_{t=0}^{T-1}\left(1 - \sum_{j=1}^{J}\tilde{\pi}_{j,t}(\mathbf{g}(t))\right)\right],$$

where $h_t(\mathbf{g},\mathbf{a}) = \lambda\sum_{j=1}^{J} p_{j,t}(\mathbf{g},\mathbf{a})\sum_{r=1}^{R_j} rF_j(r)$. Substituting it into the above equation and re-arranging the order of summation, we have

$$\hat{V}^W_0(\mathbf{G}) = \sup_{\tilde{\pi}\in\tilde{\Pi}} \mathbb{E}\left[\sum_{j=1}^{J}\left(\sum_{t=0}^{T-1}\left(\lambda p_{j,t}(\mathbf{g}(t), \tilde{\pi}_t(\mathbf{g}(t)))\sum_{r=1}^{R_j} rF_j(r) - W\tilde{\pi}_{j,t}(\mathbf{g}(t))\right) + h_T(g_j(T))\right)\right] + WT.$$

Next, recall that equation (5) is given by:

$$V^W_0(\mathbf{G}) = \sup_{\tilde{\pi}\in\tilde{\Pi}} V^{\tilde{\pi},W}_0(\mathbf{G}) = \sup_{\tilde{\pi}\in\tilde{\Pi}} \mathbb{E}\left[\sum_{j=1}^{J}\left(\sum_{t=0}^{T-1}\left(\lambda p^{\tilde{\pi}_{j,t}(\mathbf{g}(t))}_{j,t}(g_j(t))\sum_{r=1}^{R_j} rF_j(r) - W\tilde{\pi}_{j,t}(\mathbf{g}(t))\right) + h_T(g_j(T))\right)\right] + WT.$$

For each policy $\tilde{\pi}$, the only difference between the right-hand-side of the above two equations is the customer backing probabilities, i.e., MNL-based backing probability $p_{j,t}(\mathbf{g}(t), \tilde{\pi}(\mathbf{g}(t)))$ and BNL-based backing probability $p^{\tilde{\pi}_{j,t}(\mathbf{g}(t))}_{j,t}(g_j(t))$. Finally, from Lemma B.1, we have $p_{j,t}(\mathbf{g}(t), \tilde{\pi}(\mathbf{g}(t)) \leq p^{\tilde{\pi}_{j,t}(\mathbf{g}(t))}_{j,t}(g_j(t)))$ for each stage $\mathbf{g}$ at epoch t. Therefore we have $\hat{V}^{\tilde{\pi},W}_0(\mathbf{G}) \leq V^{\tilde{\pi},W}_0(\mathbf{G})$ for all $\tilde{\pi} \in \tilde{\Pi}$. Hence, by definition of the value functions, we have $\hat{V}^W_0(\mathbf{G}) \leq V^W_0(\mathbf{G})$. □

Appendix B.2. Proof of Lemma 4.1

Proof. **(i)** For any $a \in \{0,1\}$ and time epoch t, it is easy to verify that the first derivative of $p^a_t(g)$ is negative and hence $p^a_t(g)$ monotonically decreases in g, $\forall g \in (-\infty, G]$. The concavity of $p^0_t(g)$ follows by noting that it has a negative second derivative:

$$\frac{d^2}{dg^2}p^0_t(g) = -\frac{\beta^2_{2,t}\left(e^{\beta_{2,t}(1-g/G)+m} - 1\right)e^{\beta_{2,t}(1-g/G)+m}}{G^2\left(e^{\beta_{2,t}(1-g/G)+m} + 1\right)^3} < 0,\ \forall g \in (-\infty, G].$$

The concavity of $p_t^1(g)$ can be shown similarly.

(ii) It follows by noting that the first derivative of $p_t^1(g) - p_t^0(g)$ is positive:

$$\frac{d}{dg}\left(p_t^1(g) - p_t^0(g)\right) = \frac{\left(e^{\beta_{2,t}} - 1\right)\beta_{2,t} e^{\beta_{2,t}(1-g/G)+m}\left(e^{2\beta_{2,t}(1-g/G)+2m+\beta_1} - 1\right)}{G\left(e^{\beta_{2,t}(1-g/G)+m} + 1\right)^2\left(e^{\beta_{2,t}(1-g/G)+m+\beta_1} + 1\right)^2} > 0, \ \forall g \in (-\infty, G].$$

(iii) For any state g at time epoch t, we have

$$p_t^1(g) - p_t^0(g) = \frac{\exp(m + \beta_1 + \beta_{2,t}(1 - g/G))}{1 + \exp(m + \beta_1 + \beta_{2,t}(1 - g/G))} - \frac{\exp(m + \beta_{2,t}(1 - g/G))}{1 + \exp(m + \beta_{2,t}(1 - g/G))}.$$

Then, for the same state g at time epoch $t + 1$, we have

$$\begin{aligned} p_{t+1}^1(g) - p_{t+1}^0(g) &= \frac{\exp(m+\beta_1+\beta_{2,t+1}(1-g/G))}{1+\exp(m+\beta_1+\beta_{2,t+1}(1-g/G))} - \frac{\exp(m+\beta_{2,t+1}(1-g/G))}{1+\exp(m+\beta_{2,t+1}(1-g/G))} \\ &= \frac{\exp(m+\beta_1+\beta_{2,t}(1-g/G)+d)}{1+\exp(m+\beta_1+\beta_{2,t}(1-g/G)+d)} - \frac{\exp(m+\beta_{2,t}(1-g/G)+d)}{1+\exp(m+\beta_{2,t}(1-g/G)+d)} \\ &\leq \frac{\exp(m+\beta_1+\beta_{2,t}(1-g/G))}{1+\exp(m+\beta_1+\beta_{2,t}(1-g/G))} - \frac{\exp(m+\beta_{2,t}(1-g/G))}{1+\exp(m+\beta_{2,t}(1-g/G))} \\ &= p_t^1(g) - p_t^0(g) \end{aligned}$$

where $d = (\beta_{2,t+1} - \beta_{2,t})(1 - g/G) \geq 0$ as $\beta_{2,t+1} \geq \beta_{2,t}$. The inequality holds by **(ii)** since, given a fixed state g, $\beta_{2,t}(1-g/G)+d$ is equivalent to $\beta_{2,t}(1-g'/G)$ where $g' < g$. Similarly, we can also show that **(iv)** holds as well. □

Appendix B.3. Proof of Lemma B.2

Next, we show that the value function $v_t^W(g)$ is monotonic in both time t and state g.

Lemma B.2. *(Monotonicity of the Value Function) For any $W \geq 0$, the value function satisfies:*

(i) $v_t^W(g) < v_t^W(g - 1)$, $0 \leq t \leq T - 1$, $v_T^W(g) = 0$, $\forall g \in \Omega_t$;

(ii) $v_t^W(g) > v_{t+1}^W(g)$, $0 \leq t \leq T - 1$, $v_T^W(g) = 0$, $\forall g \in \Omega_t$.

Proof. **(i)** The proof is by induction. First, the inequality holds at time $T - 1$ by the boundary condition (7), i.e., $v_T^W(g) = 0$. We assume that it is true at time $t + 1$. Let $v_{t+1}^W(g)$ denote the optimal solution at time $t + 1$ for the problem in (7). We will show the inequality also holds for time t. There are two actions (i.e., $a \in \{0, 1\}$) available to be taken at time t. We first focus on action $a = 0$. We have

$$\begin{aligned} q_t^W(g-1,0) - q_t^W(g,0) &= \lambda\bar{r}\left(p_t^0(g-1) - p_t^0(g)\right) + v_{t+1}^W(g-1) - v_{t+1}^W(g) \\ &\quad + \lambda p_t^0(g-1)\left(\textstyle\sum_{r=1}^R F(r)v_{t+1}^W(g-r-1) - v_{t+1}^W(g-1)\right) \\ &\quad - \lambda p_t^0(g)\left(\textstyle\sum_{r=1}^R F(r)v_{t+1}^W(g-r) - v_{t+1}^W(g)\right). \end{aligned}$$

Since $p_t^0(g-1) > p_t^0(g)$ from Lemma 4.1 (i) and $v_{t+1}^W(g-1) > v_{t+1}^W(g)$ by induction, we obtain

$$\begin{aligned} q_t^W(g-1,0) - q_t^W(g,0) &> \lambda\bar{r}\left(p_t^0(g-1) - p_t^0(g)\right) + v_{t+1}^W(g-1) - v_{t+1}^W(g) \\ &\quad + \lambda p_t^0(g)\left(\textstyle\sum_{r=1}^R F(r)v_{t+1}^W(g-r-1) - v_{t+1}^W(g-1)\right) \\ &\quad - \lambda p_t^0(g)\left(\textstyle\sum_{r=1}^R F(r)v_{t+1}^W(g-r) - v_{t+1}^W(g)\right). \end{aligned}$$

Rearranging terms yields

$$\begin{aligned} q_t^W(g-1,0) - q_t^W(g,0) &> \lambda\bar{r}\left(p_t^0(g-1) - p_t^0(g)\right) + \left(v_{t+1}^W(g-1) - v_{t+1}^W(g)\right)(1 - \lambda p_t^0(g)) \\ &\quad + \lambda p_t^0(g)\left(\textstyle\sum_{r=1}^R F(r)v_{t+1}^W(g-r-1) - v_{t+1}^W(g-r)\right) \\ &> 0. \end{aligned}$$

Hence, the result holds by induction for $a = 0$. Similarly, we can show that the result also holds for $a = 1$. Based on above results, we can straightforwardly conclude that $\max_{a\in\{0,1\}} \left\{q_t^W(g,a)\right\}$ monotonically decreases in g at any time epoch, i.e., $\max_{a\in\{0,1\}} \left\{q_t^W(g-1,a)\right\} > \max_{a\in\{0,1\}} \left\{q_t^W(g,a)\right\}$. By induction, this completes the proof of **(i)** for $a \in \{0,1\}$.

(ii) From the result of Part (**i**), we could easily deduce that

$$\triangle v_t^W(g) = \sum_{r=1}^{R} F(r) v_t^W(g-r) - v_t^W(g) > 0, 0 \le t \le T-1 \text{ and } \triangle v_T^W(g) = 0, \forall g \in \Omega_t. \tag{B.1}$$

thus, we have

$$v_t^W(g) = \max_{a\in\{0,1\}} \left\{q_t^W(g,a)\right\} \ge q_t^W(g,0) = \lambda p_t^0(g)\left(\bar{r} + \triangle v_{t+1}^W(g)\right) + v_{t+1}^W(g) > v_{t+1}^W(g).$$

Therefore, $v_t^W(g)$ monotonically decreases in t. □

Appendix B.4. Proof of Lemma 4.2

Proof. For both parts **(i)** and **(ii)**, we focus on the case of $a = 0$. The proofs for $a = 1$ are similar.

(i) The proof is by induction on t. First the statement is trivially true at time $T-1$ by the boundary condition in (7), i.e., $v_T^W(g) = 0$. Assume that it is true at time $t+1$ and let $v_{t+1}^{\pi^0}(g)$ denote the value function at time $t+1$, $\forall g \in \Omega_{t+1}$. Then

$$\begin{aligned} \triangle v_t^{\pi^0}(g) - \triangle v_t^{\pi^0}(g-1) \;=\; & \triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g-1) \\ & + \lambda \bar{r}\left(\textstyle\sum_{r=1}^{R} F(r)(p_t^0(g-r) - p_t^0(g-r-1)) + p_t^0(g-1) - p_t^0(g)\right) \\ & + \textstyle\sum_{r=1}^{R} F(r)\left(\lambda p_t^0(g-r) \triangle v_{t+1}^{\pi^0}(g-r) - \lambda p_t^0(g-r-1) \triangle v_{t+1}^{\pi^0}(g-r-1)\right) \\ & + \lambda p_t^0(g-1) \triangle v_{t+1}^{\pi^0}(g-1) - \lambda p_t^0(g) \triangle v_{t+1}^{\pi^0}(g) \end{aligned} \tag{B.2}$$

Since $\triangle v_{t+1}^{\pi^0}(g) > \triangle v_{t+1}^{\pi^0}(g-1), \forall g \in \Omega_{t+1}$ by induction, the following inequality holds:

$$\triangle v_{t+1}^{\pi^0}(g-1) - \triangle v_{t+1}^{\pi^0}(g-r) \quad < \quad \triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g-r-1)$$

Furthermore, by Lemma 4.1 (i) that $p_t^0(g-1) > p_t^0(g)$, $\forall g \in \Omega_{t+1}$, we obtain

$$\begin{aligned} p_t^0(g-r)\left(\triangle v_{t+1}^{\pi^0}(g-1) - \triangle v_{t+1}^{\pi^0}(g-r)\right) \quad &< \quad p_t^0(g-r-1)\left(\triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g-r-1)\right) \\ p_t^0(g-r) \triangle v_{t+1}^{\pi^0}(g-1) - p_t^0(g-r-1) \triangle v_{t+1}^{\pi^0}(g) \quad &< \quad p_t^0(g-r) \triangle v_{t+1}^{\pi^0}(g-r) - p_t^0(g-r-1) \triangle v_{t+1}^{\pi^0}(g-r-1) \end{aligned}$$

In addition, from Lemma 4.1 (i), we have $p_t^0(g-1) + p_t^0(g-r) > p_t^0(g) + p_t^0(g-r-1)$ for any $r > 0$. Substituting the above inequalities into (B.2), we obtain

$$\begin{aligned} \triangle v_t^{\pi^0}(g) - \triangle v_t^{\pi^0}(g-1) \quad > \quad & \triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g-1) \\ & + \lambda \bar{r}\left(\textstyle\sum_{r=1}^{R} F(r)\left(p_t^0(g-r) - p_t^0(g-r-1)\right) + p_t^0(g-1) - p_t^0(g)\right) \\ & + \lambda \textstyle\sum_{r=1}^{R} F(r)\left(p_t^0(g-1) + p_t^0(g-r)\right) \triangle v_{t+1}^{\pi^0}(g-1) - \left(p_t^0(g-1) + p_t^0(g-r)\right) \triangle v_{t+1}^{\pi^0}(g) \end{aligned}$$

Let $P_t(g,r) = p_t^0(g-1) + p_t^0(g-r) \in [0,2]$, then

$$\begin{aligned} \triangle v_t^{\pi^0}(g) - \triangle v_t^{\pi^0}(g-1) \quad = \quad & \lambda \bar{r}\left(\textstyle\sum_{r=1}^{R} F(r)(p_t^0(g-r) - p_t^0(g-r-1)) + p_t^0(g-1) - p_t^0(g)\right) \\ & + \textstyle\sum_{r=1}^{R} F(r)\left(1 - \lambda P_t(g,r)\right)\left(\triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g-1)\right) \end{aligned}$$

Further, let $FP_t(g,r) = \sum_{r=1}^{R} F(r)(1 - \lambda P_t(g,r)) \in [-1,1]$ and rearrange terms, by Lemma 4.1 (iv) and (B.1) in

Lemma B.2, we obtain

$$
\begin{aligned}
& \triangle v_t^{\pi^0}(g) - \triangle v_t^{\pi^0}(g-1) \\
> \; & \lambda \bar{r} \left(\textstyle\sum_{r=1}^R F(r) \left(p_{t+1}^0(g-r) - p_{t+1}^0(g-r-1) \right) + p_{t+1}^0(g-1) - p_{t+1}^0(g) \right) (1 + FP_{t+1}(g,r)) \\
& + FP_{t+1}(g,r) \left(\triangle v_{t+2}^{\pi^0}(g) - \triangle v_{t+2}^{\pi^0}(g-1) \right) \\
> \; & \min_{s \in [t,T-1]} \left\{ \lambda \bar{r} \left(\textstyle\sum_{r=1}^R F(r) \left(p_s^0(g-r) - p_s^0(g-r-1) \right) + p_s^0(g-1) - p_s^0(g) \right) \left(1 + \cdots + FP_T(g,r)^{T-t-1} \right) \right\} \\
& + FP_T(g,r)^{T-t} \left(\triangle v_T^{\pi^0}(g) - \triangle v_T^{\pi^0}(g-1) \right) \\
= \; & \min_{s \in [t,T-1]} \left\{ \lambda \bar{r} \left(\textstyle\sum_{r=1}^R F(r) \left(p_s^0(g-r) - p_s^0(g-r-1) \right) + p_s^0(g-1) - p_s^0(g) \right) \frac{1 - FP_s(g,r)^{T-t}}{1 - FP_s(g,r)} \right\}, \\
> \; & 0
\end{aligned}
$$

The result holds since

$$
\min_{s \in [t,T-1]} \left\{ \lambda \bar{r} \left(\sum_{r=1}^R F(r) \left(p_s^0(g-r) - p_s^0(g-r-1) \right) + p_s^0(g-1) - p_s^0(g) \right) \right\} > 0,
$$

$$
\min_{s \in [t,T-1]} \left\{ \frac{1 - FP_s(g,r)^{T-t}}{1 - FP_s(g,r)} \right\} = \min_{s \in [t,T-1]} \left\{ \frac{1 - FP_s(g,r)^{T-t}}{\lambda \sum_{r=1}^R F(r) P_s(g,r)} \right\} > 0
$$

Hence the proof for **(i)** is completed.

(ii) We note that by the definition of $\triangle v_t^{\pi^0}(g)$, we have

$$
\triangle v_t^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g) = \lambda \sum_{r=1}^R F(r) p_t^0(g-r) \left(\bar{r} + \triangle v_{t+1}^{\pi^0}(g-r) \right) - \lambda p_t^0(g) \left(\bar{r} + \triangle v_{t+1}^{\pi^0}(g) \right).
$$

The proof for **(ii)** is done by induction. First, it holds for $t = T-1$ by the boundary condition in (7). Assume it is true for time $t+1$ that $\triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+2}^{\pi^0}(g) > 0$, $\forall g \in \Omega_{t+1}$ and let $v_{t+1}^{\pi^0}(g)$ denote the value function at time $t+1$. After some algebraic manipulation, we may obtain

$$
\begin{aligned}
\triangle v_t^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g) = \; & \lambda \textstyle\sum_{r=1}^R F(r) p_t^0(g-r) \left(\sum_{r_1=1}^R F(r_1) \left(\lambda p_{t+1}^0(g-r-r_1) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g-r-r_1) \right) \right) \right) \\
& - \lambda \textstyle\sum_{r=1}^R F(r) p_t^0(g-r) \left(\lambda p_{t+1}^0(g-r) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g-r) \right) \right) \\
& - \lambda p_t^0(g) \left(\textstyle\sum_{r_1=1}^R F(r_1) \left(\lambda p_{t+1}^0(g-r_1) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g-r_1) \right) \right) \right) + \lambda p_t^0(g) \left(\lambda p_{t+1}^0(g) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g) \right) \right) \\
& + \lambda \bar{r} \textstyle\sum_{r=1}^R F(r) p_t^0(g-r) - \lambda \bar{r} p_t^0(g) \\
& + \lambda \textstyle\sum_{r=1}^R F(r) p_t^0(g-r) \left(\sum_{r_1=1}^R F(r_1) v_{t+2}^{\pi^0}(g-r-r_1) \right) - \lambda \sum_{r=1}^R F(r) p_t^0(g-r) v_{t+2}^{\pi^0}(g-r) \\
& - \lambda p_t^0(g) \textstyle\sum_{r_1=1}^R F(r_1) v_{t+2}^{\pi^0}(g-r_1) + \lambda p_t^0(g) v_{t+2}^{\pi^0}(g)
\end{aligned} \tag{B.3}
$$

Rearrange all the terms, the first two lines give $\lambda \sum_{r=1}^R F(r) p_t^0(g-r) \left(\triangle v_{t+1}^{\pi^0}(g-r) - \triangle v_{t+2}^{\pi^0}(g-r) \right)$, the third line gives $-\lambda p_t^0(g) \left(\triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+2}^{\pi^0}(g) \right)$, and the rest of lines gives $\lambda \sum_{r=1}^R F(r) p_t^0(g-r) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g-r) \right) - \lambda p_t^0(g) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g) \right)$. By induction and Lemma 4.1 (iv), we obtain

$$
\begin{aligned}
& \lambda \textstyle\sum_{r=1}^R F(r) p_t^0(g-r) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g-r) \right) - \lambda p_t^0(g) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g) \right) \\
\geq \; & \lambda \textstyle\sum_{r=1}^R F(r) p_{t+1}^0(g-r) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g-r) \right) - \lambda p_{t+1}^0(g) \left(\bar{r} + \triangle v_{t+2}^{\pi^0}(g) \right) \\
= \; & \triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+2}^{\pi^0}(g)
\end{aligned}
$$

Finally we have the final result by arranging the above terms, as

$$
\triangle v_t^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g) \geq \lambda \sum_{r=1}^R F(r) p_t^0(g-r) \left(\triangle v_{t+1}^{\pi^0}(g-r) - \triangle v_{t+2}^{\pi^0}(g-r) \right) + \left(1 - \lambda p_t^0(g) \right) \left(\triangle v_{t+1}^{\pi^0}(g) - \triangle v_{t+2}^{\pi^0}(g) \right) > 0.
$$

Hence, the inequality in **(ii)** holds by induction. □

Appendix B.5. Proof of Theorem 4.1

(i) Note that policy π^0 is optimal if and only if $q_t^W(g,1) < q_t^W(g,0), \forall t, g \in \Omega_t$. Consider

$$\triangle q_t^W(g) = q_t^W(g,1) - q_t^W(g,0) = \lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\bar{r} + \triangle v_{t+1}^W(g)\right) - W, \forall t, g \in \Omega_t. \tag{B.4}$$

It is clear from equation (B.4) that when W is sufficiently large, we have $\triangle q_t^W(g) = q_t^W(g,1) - q_t^W(g,0) < 0, \forall t, g \in \Omega_t$, and hence the policy π^0 is optimal for sufficiently large W.

Next, we note that $\triangle v_0^{\pi^0}(G) > \triangle v_t^{\pi^0}(g), \forall g, t$ from Lemma 4.2 and $p_0^1(G) - p_0^0(G) > p_t^1(g) - p_t^0(g), \forall t, g \in \Omega_t$ from Lemma 4.1 (ii) and (iii). Hence, we have

$$\triangle q_0^W(G) = \max_{\forall g,t}\left\{\triangle q_t^W(g)\right\} = \max_{\forall g,t}\left\{\lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\bar{r} + \triangle v_{t+1}^W(g)\right) - W\right\}.$$

Now we consider the scenario where W decreases to a cost $\overline{W}$ such that the optimal action is indifferent between $a = 1$ and $a = 0$ at state G and time 0, whereas in the rest of states it takes action $a = 0$, i.e.,

$$\begin{aligned}
\triangle q_0^{\overline{W}}(G) &= \max_{\forall g,t}\left\{\triangle q_t^{\overline{W}}(g)\right\} = 0,\\
\triangle q_t^{\overline{W}}(g) &< 0, \text{ otherwise.}
\end{aligned}$$

This indicates that $\triangle q_0^W(G) \geq 0$ for any $W \leq \overline{W}$, and hence the policy π^0 is not optimal anymore. We thus conclude that the policy π^0 is optimal if and only if $W > \overline{W}$.

(ii) Note $p_t^1(g) - p_t^0(g) > 0$ from Lemma 4.1 (ii) and (iii) and $\triangle v_{t+1}^W(g) > 0$ from (B.1) in Lemma B.2. Hence, from equation (B.4), we conclude that $\triangle q_t^W(g) = q_t^W(g,1) - q_t^W(g,0) > 0, \forall t, g \in \Omega_t$, when $W \geq 0$ is sufficiently small. This means that the policy π^1 is optimal for sufficiently small $W \geq 0$. Next, we consider the scenario where W gradually increases from a very small $W \geq 0$. From (B.1) in Lemma B.2, we have $\triangle v_{t+1}^W(g) > 0$ so that for any t and $g \in \Omega_t$:

$$\begin{aligned}
\triangle q_t^W(g) &= \lambda\bar{r}p_t^1(g) - W - \lambda\bar{r}p_t^0(g) + \lambda\left(p_t^1(g) - p_t^0(g)\right)\triangle v_{t+1}^W(g)\\
&> \lambda\bar{r}p_{T-1}^1(g) - W - \lambda\bar{r}p_{T-1}^0(g)\\
&= q_{T-1}^W(g,1) - q_{T-1}^W(g,0) = \triangle q_{T-1}^W(g).
\end{aligned}$$

In addition, from Lemma 4.1 (ii) we have

$$\begin{aligned}
\triangle q_{T-1}^W(g) &= \lambda\bar{r}p_{T-1}^1(g) - W - \lambda\bar{r}p_{T-1}^0(g)\\
&\geq \lambda\bar{r}(p_{T-1}^1(G-(T-1)R) - W - p_{T-1}^0(G-(T-1)R))\\
&= q_{T-1}^W(G-(T-1)R, 1) - q_{T-1}^W(G-(T-1)R, 0)\\
&= \triangle q_{T-1}^W(G-(T-1)R).
\end{aligned}$$

Thus, by setting $\underline{W} = \lambda\bar{r}\left(p_{T-1}^1(G-(T-1)R) - p_{T-1}^0(G-(T-1)R)\right)$ such that $\triangle q_{T-1}^{\underline{W}}(G-(T-1)R) = 0$, we obtain

$$\begin{aligned}
&\triangle q_{T-1}^{\underline{W}}(G-(T-1)R) = 0,\\
&\triangle q_t^{\underline{W}}(g) > 0, \text{ otherwise.}
\end{aligned}$$

This indicates that $\triangle q_{T-1}^W(G-(T-1)R) < 0$ for any $W > \underline{W}$, and hence the policy π^1 is not optimal anymore. We thus conclude that the policy π^1 is optimal if and only if $W \leq \underline{W}$.

Appendix B.6. Proof of Theorem 4.2

First, we note that from the proof of Theorem 4.1, parts (i) and (ii) are trivially true. We now focus on the proofs for parts (iii) and (iv).

Let B_t be the set of states at time t for which the optimal action is to promote. Without causing any ambiguity we include into B_t the states where the optimal action is indifferent between promotion or not. From the proof of Theorem 4.1, for any $W > \overline{W}$, we have $B_t = \varnothing, \forall t$. Consequently, its complement (non-promotion set at t) is $B'_t = \Omega_t$. Also note that for any $W > \overline{W}$, we have

$$\triangle q_t^W(g) = q_t^W(g,1) - q_t^W(g,0) = \lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\bar{r} + \triangle v_{t+1}^{\pi^0}(g)\right) - W < 0, \forall g \in B'_t, \tag{B.5}$$

and, along with Lemmas 4.1 (ii) and (iii) and Lemma 4.2,

$$\triangle q_t^W(g-1) < \triangle q_t^W(g) < 0 \text{ and } \triangle q_{t+1}^W(g) < \triangle q_t^W(g) < 0, \forall g \in B'_t. \tag{B.6}$$

When W reduces to $\overline{W}$, according to part (i), we have $B_0 = \{G\}$ and $B_t = \varnothing, \forall 1 \leq t \leq T-1$. Both (B.5) and (B.6) still hold for the states in the non-promotion sets B'_t. When W further decreases from $\overline{W}$ to some smaller value W^1, there will be the second state that becomes indifferent between the two actions, and this must be the state G at time $t = 1$, since from (B.6) it has the largest value of $\triangle q_t^{W^1}(g)$ among all the states in the non-promotion sets. The promotion set at $t = 1$ is then updated to $B_1 = \{G\}$, while those for all the other time epochs remain unchanged.

Following the same vein, as W continues to decrease to some even smaller value W^2, there will be the next state that becomes indifferent between the two actions. From (B.6), this can take place either at state $G-1$ in time $t = 1$ or at state G in time $t = 2$. If the former switches first, the promotion set at $t = 1$ grows to $B_1 = \{G, G-1\}$; otherwise the promotion set at $t = 2$ grows to $B_2 = \{G\}$. As W continues to decrease, this process continues, and there will be more states switching from the non-promotion sets to the promotion sets, one state at a time. These switching states either have the largest state g, the smallest time t, or both, among the remaining states in the non-promotion sets. Since the optimal q-factor at state g and time t only depends on the value functions for states $g' \leq g$ at time $t' > t$, we conclude that (B.5) and (B.6) are always true for those states in the non-promotion sets under any non-negative W. This completes the proofs for parts (iii) and (iv).

The process of the reasoning is illustrated in Fig. B.1.

Figure B.1: An illustration of the optimal actions as W decreases

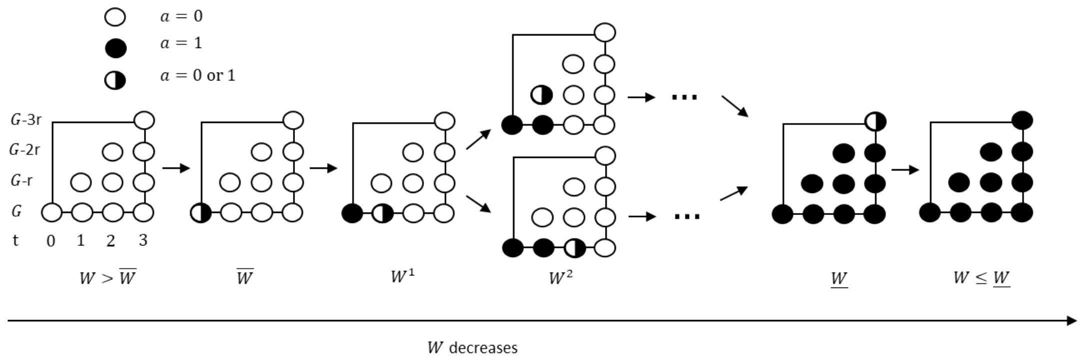

Appendix B.7. Proof of Proposition 4.1

For part (i), Theorem 4.2 (iv) states that the change will take place at state g before $g-1$ for the same time t as the cost W decreases. In other words, for the same time t, the change of the optimal action in state $g-1$ must happen later than the change of the optimal action in state g under the same cost. Therefore, the result in part (i) holds.

Similarly for part (ii), Theorem 4.2 (iii) states that the change will take place at time t before $t+1$ at the same state g. This indicates that for the same state g, the change of the optimal action in time t must happen earlier than the change of the optimal action in time $t+1$ under the same cost. Hence, the result in part (ii) holds.

Appendix B.8. Proof of Lemma 4.3

(i) To show $w(g,t)$ decreases in state g, we first note that from (B.1) in Lemma B.2, we have $\triangle v_{t+1}^{\pi^0}(g) > 0$. Hence, we obtain

$$\begin{aligned} w(g,t) - w(g-1,t) &= \lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\bar{r} + \triangle v_{t+1}^{\pi^0}(g)\right) - \lambda\left(p_t^1(g-1) - p_t^0(g-1)\right)\left(\bar{r} + \triangle v_{t+1}^{\pi^0}(g-1)\right) \\ &> \lambda\left(\left(p_t^1(g) - p_t^0(g)\right) - \left(p_t^1(g-1) - p_t^0(g-1)\right)\right)(\bar{r} + \triangle v_{t+1}^{\pi^0}(g)) \\ &> 0, \end{aligned}$$

where the second inequality is held by Lemma 4.1 (ii).

(ii) To show $w(g,t)$ decreases in time t, we have:

$$\begin{aligned} w(g,t) - w(g,t+1) &= \lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\bar{r} + \triangle v_t^{\pi^0}(g)\right) - \lambda\left(p_{t+1}^1(g) - p_{t+1}^0(g)\right)\left(\bar{r} + \triangle v_{t+1}^{\pi^0}(g)\right) \\ &\geq \lambda\left(p_t^1(g) - p_t^0(g)\right)\left(\triangle v_t^{\pi^0}(g) - \triangle v_{t+1}^{\pi^0}(g)\right) \\ &> 0 \end{aligned}$$

where the inequalities hold by Lemma 4.1 (iii) and Lemma 4.2 (ii).

Appendix B.9. Proof of Proposition 4.2

Proof. From the proof of Theorem 4.2, we observe that for any $W \geq W'$, the collection of states that switch from $a = 0$ to $a = 1$ under W' always include those states which have already switched from $a = 0$ to $a = 1$ under W. This means that the optimal promotion set under W' is larger than that under W. This indicates that, as W decreases, the critical state becomes smaller, i.e., $g^*(W,t) \geq g^*(W',t)$. Therefore, by the definition of indexability (Definition 4.1), we conclude that the project with sufficiently long duration is indexable. □

Appendix B.10. Proof of Proposition 5.1

For a fixed single project policy π_j, the value function $v_{j,0}^W(G_j)$ is linear in W. The pointwise maximum over these linear functions produces a piecewise linear convex function. The sum of these functions of $v_{j,0}^W(G_j)$ over all $1 \leq j \leq J$ with an additional linear term WT, which gives $V_0^W(\mathbf{G})$, is also piecewise linear and convex.

Appendix B.11. Proof of Proposition 5.2

We first find a sub-gradient of the value function for the single-project problem (6). The Bellman equation takes the same form as equation (7) but with the AoN termination condition, written as follows. For any $g \in \Omega_{j,t}, 0 \leq t \leq T-1$,

$$v_{j,t}^W(g) = \max_{a \in \{0,1\}} \left\{ \lambda p_{j,t}^a(g) \sum_{r=1}^{R_j} F_j(r)\left(r + v_{j,t+1}^W(g-r)\right) + \left(1 - \lambda p_{j,t}^a(g)\right) v_{j,t+1}^W(g) - Wa \right\}, \forall g \in \Omega_{j,t}, 0 \leq t \leq T-1 \tag{B.7}$$

with termination for any $g \in \Omega_{j,T}$,

$$v_{j,T}^W(g) = \begin{cases} -(G_j - g) & if \quad g > 0, \\ 0 & if \quad g \leq 0. \end{cases}$$

where the backing probability $p^a_{j,t}(g)$ is calculated below (note that $\beta_{3,J}$ is now included):

$$p^a_{j,t}(g) = \frac{\exp(m_j + \beta_1 a + \beta_{2,t}(1 - g/G_j) - \beta_{3,J})}{1 + \exp(m_j + \beta_1 a + \beta_{2,t}(1 - g/G_j) - \beta_{3,J})}.$$

For convenience we introduce the matrix form for (B.7). For each project j, let $\overset{*}{\pi}{}^W_j$ be the optimal policy under the Lagrangian multiplier W. Denote by $\mathbf{U}_{j,t}$ a $(TR_j + 1)$ dimensional vector for the optimal actions at time t, with each element $u_{j,t}(g)$ the optimal action at state g and time t. Similarly, denote by $\mathbf{P}_{j,t}$ a $(TR_j + 1)$ dimensional vector for the corresponding backing probabilities at time t, with each element $p_{j,t}(g)$ the backing probability at state g and time t under the optimal action $u_{j,t}(g)$. In addition, denote by $\mathbf{Q}_{j,t}$ a $(TR_j + 1) \times (TR_j + 1)$ dimensional matrix for the one-step transition probabilities at time t under the optimal policy. Its $(g, \tilde{g})$ element is given by:

$$q_{j,t}(g, g) = 1 - \lambda p_{j,t}(g); q_{j,t}(g, g - r) = \lambda p_{j,t}(g) F_j(r), \forall 1 \le r \le R_j; \text{ otherwise } q_{j,t}(g, \tilde{g}) = 0.$$

Let $\mathbf{v}^W_{j,t}$ denote the vector $\{v^W_{j,t}(g) : G_j - TR_j \le g \le G_j\}$. We can write the Bellman equation (B.7) under policy $\overset{*}{\pi}{}^W_j$ in the following matrix form

$$\mathbf{v}^W_{j,t} = \lambda \bar{r} \mathbf{P}_{j,t} + \mathbf{Q}_{j,t} \times \mathbf{v}^W_{j,t+1} - W \mathbf{U}_{j,t}. \tag{B.8}$$

For some other $W' = W + \delta$, we have

$$\mathbf{v}^{W'}_{j,t} \ge \lambda \bar{r}_j \mathbf{P}_{j,t} + \mathbf{Q}_{j,t} \times \mathbf{v}^{W'}_{j,t+1} - W' \mathbf{U}_{j,t}, \tag{B.9}$$

which holds as $\mathbf{Q}_{j,t}$, $\mathbf{P}_{j,t}$ and $\mathbf{U}_{j,t}$ are all evaluated under policy $\overset{*}{\pi}{}^W_j$, which may be sub-optimal with respect to W'. Subtracting both sides of equation (B.8) from that of (B.9) yields

$$\mathbf{v}^{W'}_{j,t} - \mathbf{v}^W_{j,t} \ge \mathbf{Q}_{j,t} \times \left(\mathbf{v}^{W'}_{j,t+1} - \mathbf{v}^W_{j,t+1} \right) - \delta \mathbf{U}_{j,t}. \tag{B.10}$$

Expanding the above inequality forward recursively over time, we have

$$\begin{aligned} \mathbf{v}^{W'}_{j,t} - \mathbf{v}^W_{j,t} &\ge \mathbf{Q}_{j,t} \times \left(\mathbf{v}^{W'}_{j,t+1} - \mathbf{v}^W_{j,t+1} \right) - \delta \mathbf{U}_{j,t} \ge \cdots \ge \prod_{\tau=t}^{T-1} \mathbf{Q}_{j,\tau} \times \left(\mathbf{v}^{W'}_{j,T} - \mathbf{v}^W_{j,T} \right) - \delta \left(\mathbf{U}_{j,t} + \sum_{s=t+1}^{T-1} \left(\prod_{\tau=t}^{s-1} \mathbf{Q}_{j,\tau} \right) \times \mathbf{U}_{j,s} \right) \\ &= -\delta \left(\mathbf{U}_{j,t} + \sum_{s=t+1}^{T-1} \left(\prod_{\tau=t}^{s-1} \mathbf{Q}_{j,\tau} \right) \times \mathbf{U}_{j,s} \right) \end{aligned}$$

where the last equality holds as $v^W_{j,T}(g) \equiv v^{W'}_{j,T}(g)$ for all g by the definition of the termination condition. In particular, for the value function at G_j and $t = 0$ we have

$$v^{W'}_{j,0}(G_j) - v^W_{j,0}(G_j) \ge -\delta \left[\mathbf{U}_{j,0} + \sum_{s=1}^{T-1} \left(\prod_{\tau=0}^{s-1} \mathbf{Q}_{j,\tau} \right) \times \mathbf{U}_{j,s} \right] (G_j).$$

Hence, for the value function of the original problem at $\mathbf{G}$ and $t = 0$, we have

$$V^{W'}_0(\mathbf{G}) - V^W_0(\mathbf{G}) = \sum_{j=1}^{J} \left(v^{W'}_{j,0}(G_j) - v^W_{j,0}(G_j) \right) + (W' - W)T \ge \delta \left(T - \sum_{j=1}^{J} \left[\mathbf{U}_{j,0} + \sum_{s=1}^{T-1} \left(\prod_{\tau=0}^{s-1} \mathbf{Q}_{j,\tau} \right) \times \mathbf{U}_{j,s} \right] (G_j) \right).$$

Appendix C. A Condition for Sufficiently Long Duration

In this appendix, we provide a mathematical condition under which the AoN scheme becomes irrelevant.

First, we note that, by definition, the total expected revenue from a single project over the entire campaign duration T under any policy π is given by:

$$\begin{aligned} v^\pi &= \mathbb{E}\left[\left(\sum_{t=0}^{T-1} \left(\lambda p_t^{\pi_t(g(t))}(g(t)) \sum_{r=1}^{R} rF(r) \right) + h_T(g(T)) \right) \right] = \lambda \sum_{r=1}^{R_j} rF(r) \sum_{t=0}^{T-1} \mathbb{E}\left[p_t^{\pi_t(g(t))}(g(t)) \right] + \mathbb{E}[h_T(g(T))] \\ &> \lambda \sum_{r=1}^{R} rF(r) \sum_{t=0}^{T-1} \mathbb{E}\left[p_t^{\pi_t(g(t))}(g(t)) \right] - G, \end{aligned}$$

where the inequality is obtained by the definition of $h_T(g(T))$ in equation (2b). It can easily be shown that the choice probabilities in the above equation are strictly positive and bounded (refer to equation (4) in the paper), i.e., $0 < p^0(G) < p_t^{\pi_t(g(t))}(g(t))$ where $p^0(G)$ represents the purchasing probability without promotion at the initial time epoch. This leads to $v^\pi > T\left[\lambda p^0(G)\sum_{r=1}^{R} rF(r)\right] - G$. It indicates that when the campaign duration T is sufficiently long, the total expected revenue from the project will exceed the pre-set funding goal G. Therefore, a sufficiently long campaign duration for a single project may be specified as

$$T \geq T^* = \frac{G}{\lambda p^0(G)\sum_{r=1}^{R} rF(r)}.$$

For any fixed attractiveness parameter m, we have $T^* < +\infty$. It is therefore expected that the project will complete within T^* and thus the AoN scheme effectively becomes irrelevant.