Crowdvoting the Timing of New Product Introduction

Victor Araman[*] René Caldentey[†]

July 2016

Working Paper Olayan School of Business, AUB
Please Dot Not Distribute Without Authors' Permission

Abstract

Launching new products into the marketplace is a complex and risky endeavor that companies must continuously undertake. As a result, it is not uncommon to witness major firms discontinuing a product shortly after its introduction. In this paper, we consider a seller who has the ability to first test the market and gather demand information before deciding whether or not to launch a new product. In particular, we consider the case in which the seller sets up an online voting system that potential customers can use to provide feedback about their willingness to buy the new product. This voting system has the potential of offering a win-win situation whereby a consumer who votes hopes to influence the seller's final assortment, while at the same time these votes and their pace benefit the seller as they provide valuable information to better forecast demand. We investigate the optimal design of such a *crowdvoting* system and its implications on the seller's commercialization strategy.

1 Introduction

Bringing new products into the marketplace offers great opportunities for companies to generate new revenue streams and increase sales. However, such endeavors represent a risky bet, and unsuccessful products are a major liability generating possibly great capital expenditure, early markdowns, serious goodwill cost, and loss of market share. It is not infrequent to witness major brands preferring to discontinue a product, shortly after its introduction, rather than taking more risks and incurring higher draining costs[‡]. For these reasons, companies seek to test the market's reaction to a product before the decision of launching is made. In practice, these market tests can be expensive and difficult to conduct effectively and probably worth doing only for a few radically new products. However, this reality is now changing as companies are beginning to see the potential to *crowdsource* such market testing activities.

Broadly speaking, *crowdsourcing* is a recent phenomenon used by all kind of companies, from small start-ups to multinationals and across industries in the private and public sectors to outsource some difficult tasks, previously impossible, too costly, or inefficient. More specifically,

[*]Olayan School of Business, American University of Beirut, Beirut, Lebanon.
[†]Booth School of Business, The University of Chicago.
[‡]Sell Big or Die Fast, *New York Times*, J. Wortham and V.G. Kopytoff, August 23, 2011

> "Simply defined, *crowdsourcing* represents the act of a company or institution taking a function once performed by employees and outsourcing it to an undefined (and generally large) network of people in the form of an open call." Jeff Howe, 2006.

As an example, a recurrent application of crowdsourcing is when a company publishes (online) specific needs in order to benefit from a greater selection of suppliers (*e.g.*, Unilever has a list of "Wants" available on its website for those interested in suggesting a technical solution and becoming a Unilever partner). However, crowdsourcing is not only a B2B phenomenon but a B2C as well, whereby companies use such platforms to reach out to consumers and channel their feedback; the so-called "wisdom of the crowd". In many cases, crowdsourcing is becoming a primary source of information in the innovation process (*e.g.*, Fiat Mio project was the result of 17,000 participants submitting more than 11,000 ideas that were studied and interpreted by Fiat and resulted in the final concept). One way of "sourcing" information is through the implementation of a *crowdvoting* system, in which the "crowd" is asked to vote on a particular point of interest (*e.g.*, in 2011, 1.8 Million people voted to decide what should be the plural of Prius, Toyota's flagship hybrid car. In this contest, run on Toyota's website, 25% voted for Prii which is now the official term to describe the plural of Prius.)

As the Internet and social network platforms keep growing worldwide, we expect that crowdvoting initiatives, such as the Fiat Mio project or the Toyota Prii example, becomes standard business practices in the design and introduction of new products into the marketplace. As a matter of the fact, the list of companies taking advantage of, or built around crowdvoting systems, is increasing at a fast pace. For example, in order to better understand customers needs and reduce any supply-demand mismatch, the French company myfab.com has witnessed a great success with their new furniture business model. This company offers a set of designs and prototypes on their website. Visitors log in and vote for their favorite product during a given window of time. At the end of this voting phase, the designs that are most liked become available for sales in their online *boutique*. As a result of this model, mismatch costs are significantly reduced allowing myfab.com to offer highly targeted products at very competitive prices. Another example is Threadless.com, a site where anyone can design a T-shirt and submit it to a weekly contest. Viewers vote for their favorite T-shirt and the winning designs are selected for production and their designers get rewarded.

Motivated by this emerging trend, we investigate the question of how a company can approach the task of setting up a crowdvoting system to support the process of launching a new product into the market. Specifically, we focus on two distinguishing features that we believe make these online voting systems particularly valuable. First, a crowdvoting system acts as a demand forecasting mechanism allowing companies to assess the market potential of new products before committing the necessary resources needed for launching them. Second, crowdvoting offers a natural opportunity to pre-sell an item that is being voted. By pre-selling a new product, the seller (i) creates an incentive for customers to vote, (ii) is able to price discriminate between voters and regular buyers and generate more revenues and (iii) is able to build up an "inventory" of pre-orders that hedge the cost and risk of launching the product. In other words, such a system offers an opportunity to *crowdfund* (some of) the costs of bringing new products into the marketplace, which is a feature that is particularly relevant for small firms and entrepreneurs that have limited access to capital markets (*e.g.*, kickstarter.com).

In this paper, we propose a stylized mathematical model to study the design of a crowdvoting system. In particular, we consider a seller who is contemplating launching a new product and is uncertain about its market potential. The seller sets up an online system in which website visitors can review the product description and attributes and, if they like it, can vote (click) on the product without any commitment. To stimulate this process, the seller sometimes offers a price discount (or coupon) for voters valid if the product gets commercialized. Before opening up the voting process, the seller has an initial belief on how successful the product will be. As the votes get casted, the seller learns consumers' preferences and updates his belief about market demand. Eventually, the seller stops the voting phase and either discards the product or launches it and starts a regular selling phase. A critical question in the design of the system is how long the crowdvoting phase should last. On one hand, the longer the duration the more accurate the seller's demand forecast would be. On the other hand, a long voting period can discourage consumers to vote or to purchase the product if it is eventually launched. In addition, the seller incurs an opportunity cost by delaying the start of the selling phase.

In practice, most crowdvoting systems use a fixed pre-determined voting period; those with presales set also a minimum "sales" target to reach by the end of the voting period. However, in an uncertain market environment this is typically suboptimal. In this paper, we propose instead a stopping time formulation to determine the optimal duration of the voting process. Our model allows us to quantify the trade-off between the value of demand information and the financial impact of delaying the product introduction to accumulate pre-orders. Our formulation also sheds some light on how to price the voting phase to stimulate effectively this voting process.

We conclude this introduction with a brief discussion that connects our work with some related literature. First of all, our research contributes to the recent and growing literature on crowdsourcing; see Jeff Howe (2006) and Surowiecki (2004). From a business perspective, this phenomenon seems to encompass at least three main functions: evaluating (crowdvoting), donating or funding (crowdfunding) and developing and designing solutions (crowdcreation). Particularly relevant to us are the first two functions.

Despite the attention that the industry has been giving to crowdsourcing, the academic literature remains limited. On the crowdcreation end, few studies have shed some light on the benefits of crowdsourcing in generating new product ideas (see Terwiesch and Xu, 2008 and Huang et al., 2014). In the field of Operations Research we mention the recent work of Krager et al. (2014) that looks at adaptively allocating small tasks to workers through crowdsourcing while meeting some reliability target. On the crowdvoting end, and more relevant to our work, the paper by Marinesi and Girotra (2013) focuses on measuring the information that is acquired from a customer voting system. Using a two-period game-theoretical model, they prove among other results that by offering a sufficiently high discount during the voting phase, crowdvoting systems - used to decide whether to develop the product or not - represent an effective way to elicit information on customers willingness-to-pay. Our model shares with Marinesi and Girotra (2013) the same sequence of events as well as how the second period depends on the learning that occurred in the first one. As opposed to Marinesi and Girotra (2013) we do not consider strategic consumers; however, we consider a more detailed modelling of the first phase relying on a continuous time setting with dynamic learning and with the possibility of ending the voting phase at anytime.

Finally, crowdfunding has been getting recently a lot of attention not only on the practical level (*e.g.* kickstarter.com or indiegogo.com) but also on the academic level. The latter effort is primarily exploratory. We mention the recent work of Mollick (2014) that looks at the determinant factors of (crowdfunding) project success and at how reliable these projects end up being. Another recent work is that of Crosetto and Regner (2014) that analyzes the correlation between the success of a project and the rate at which the funding has been received during the funding phase. More relevant to our context is the work of Belleflamme et al. (2014), who also consider a two-period model where founders announce the project or the product idea and set a target for the funds requested during the first period. If the target is reached during the first period, investors either get their share in profit or the product itself at a discounted price. The authors also model the asymmetry of information that exists between the seller and the investor with respect to the quality of the product. In the crowdfunding literature, demand learning *per se* is not explicitly modelled and to the best of our knowledge there is no work that has explicitly analyzed how to jointly optimize the duration of the fundraising campaign and the funding target.

Probably, the primary role of a crowdvoting phase is to learn the market's reaction to a potential new product. On the other hand, the primary role of a crowdfunding phase is to generate the necessary funds to launch a new product or develop a new project. By introducing the possibility of voters to pre-sell a product, the voting phase becomes also a funding campaign. Recently, Threadless.com introduced the possibility for voters to pre-buy their favorite T-shirt; interestingly, they called this feature "funding". The main objective of this work is to suggest a model that incorporates some of the main features and dynamics of crowdvoting and analyze their implications on the profitability of the product; in particular, we focus on characterizing the optimal duration of the crowdvoting phase. Implicitly, this duration embodies the trade-off between the value of information acquired from voters and the financial impact generated by the pre-sales phase.

Our work is also related to the advance-purchase literature. Similarly to a crowdvoting phase, the benefits from an advance-purchase tactic are primarily: (i) a better forecast of the selling season demand and, (ii) an opportunity to price discriminate when the consumers have heterogenous valuations. The value of such strategy is often analyzed through a two-period model where the demand of both periods are correlated and consumers are strategic (see the early work of Gale and Holmes, 1992, Gale and Holmes, 1993, Dana, 1998 and more recently, that of Tang et al., 1995 and Prasad et al., 2011). By dropping the strategic consumer piece, Raman and Fisher (1996) and Boyaci and Özer (2010) are able to add additional operational complexities and use the early purchase to have a better hold of their production quantity required for the selling phase. In a multi-product setting, Raman and Fisher (1996) obtain approximations for production quantities when constrained by a finite capacity. On the other hand, Boyaci and Özer (2010) consider a single product with a preselling phase that runs over (finite) multiple periods. Advanced demand information in the context of inventory management is also a relevant stream of literature where customers place orders in advance of their needs due to positive leadtimes. Sellers take this information into account for replenishment purposes (e.g. Gallego and Özer, 2001).

Some of the main differences between our work and the literature around advance-purchase include the fact that we model the crowdvoting phase using a continuous time infinite horizon setting that incorporates dynamically an (active) Bayesian learning of the market size. Our work is also unique in proposing a setting that allows one to analyze the case of pure voting (where no presales occur but only votes are tracked); in this case, we obtain a full characterization of the payoff. As we rely on active demand learning, one of the main decision the seller faces in our case is the duration of the voting phase which is obtained through an optimal stopping problem. Only Boyaci and Özer (2010) allows for such control variable in a finite horizon setting. However, in our case the seller not only decides when to end the crowdvoting phase, but also, whether to discard the product at this point or commercialize it. When presales are allowed, this important caveat makes the crowdvoting phase looks even more like a funding phase; if not enough funds are accrued then the product is discarded, otherwise, these funds are used to support the cash-flows of the selling phase (including when to incur the fixed cost). Similarly, to some of this research we also look in the Extension section how to adapt our model to take pricing of the voting phase into account.

On the demand learning side, our model assumes that voters arrive according to a Poisson process with an unknown rate that the seller estimates using a Bayesian sequential learning approach. There is a long list of papers that have considered such learning framework; in particular we mention the recent work on dynamic pricing and incomplete demand information (see Aviv and Pazgal, 2002, Araman and Caldentey, 2009 and Farias and Van Roy, 2010) who use a similar model for their selling phase, however, with no presales or voting available.

In terms of methodology and solution techniques, our paper builds on the work of Peskir and Shiryaev (2000) on Bayesian sequential testing for a Poisson process. Despite the similarities, there are major differences between Peskir and Shiryaev (2000)'s work and ours. In the former paper, the objective is to minimize a one-dimension total expected cost that includes the running cost and the cost generated from a possible type I or type II error at the time when the test stops. In our case, we maximize an infinite horizon discounted revenues characterized by a piecewise linear function. Moreover, in the pure voting case, we provide an alternative analysis of the problem using a diffusion approximation. In addition, in the case with pre-orders (see Section 5) the solution methods in Peskir and Shiryaev (2000) do not directly apply and we must use a completely different approach.

The rest of the paper is organized as follows. In the next section, we introduce the different components of the general model and discuss the main assumptions. In Section 3, we consider the special case of the full information where the seller knows exactly the demand rate. In particular, we depict the financial value that the voting phase offers to the seller even when demand is fully known. This is then followed in Section 4 by a detailed analysis of the case where voters do not end up buying the product if and when launched; we call them *impulse voters*. We formulate the problem as a set of quasi-variational inequalities which we solve completely and obtain a closed-form solution of the payoff function. These results are further simplified by appropriately scaling the problem and solving for a diffusion approximation. As a result, among others, we obtain a simple relationship for the expected time of the voting phase. In Section 5, we tackle the general case where demand is unknown and is being learned and voters can possibly purchase the product if commercialized, that is, some voters are actual buyers. We first highlight the structural difference between the general case and the special case of impulse voters. We characterize the optimal payoff and suggest an efficient algorithm to solve it. Moreover, we offer approximations of the value function and of the policy to adopt that preserve the main properties of the voting system. We numerically show that the approximations perform well. In Section 6, we discuss some extensions to our model including pricing tactics as well as incorporating the effect of delaying the product introduction on the voters 'patience' to wait and buy it. We conclude in Section 7 and offer some possible directions for future work.

2 Model Description

We consider a firm, "the seller", who is planning to introduce a new product into the marketplace. We assume that the demand for this product is driven by a Poisson process with rate (or market potential) Λ and that the seller is uncertain about its true value. In particular, he believes that Λ is a random variable with values in $\{\lambda_{\text{L}}, \lambda_{\text{H}}\}$ with $0 < \lambda_{\text{L}} \leq \lambda_{\text{H}}$. If $\Lambda = \lambda_{\text{H}}$ then the seller expects the product to be a commercial success and launching it is an optimal decision. On the other hand, if $\Lambda = \lambda_{\text{L}}$ then demand is expected to be low and the seller should discard the product (*e.g.*, the fixed costs of launching it exceed the expected sales profits).

In an attempt to reduce the risks of his decision, the seller sets up an online voting system in which potential customers (those visiting the seller's website) can vote, say by clicking on the product, if they are interested in buying it in the event that it becomes available. For simplicity, we assume that the retailer only tracks the cumulative number of favorable votes, *i.e.*, the system does not allow unfavorable voting. (In practice, we could imagine a more sophisticated interface using a more detailed scoring system, *e.g.*, a 0 to 10 scale, or even allowing for consumer reviews.) This voting phase occurs before the seller decides to launch the product and has the potential of offering a win-win situation whereby a consumer who votes hopes to influence the retailer to commercialize the "right" products; and on the other hand, these votes and their pace provide valuable information that the retailer can use to better forecast the value of Λ.

A key feature of the system is the actual duration of this voting phase, which we denote by τ. On one hand, the seller would like to make τ as large as possible to maximize the quality of his demand forecast. On the other hand, there is the financial cost of extending the voting system and delaying the possible revenues that will be generated during the selling phase. In addition, a long voting phase can discourage customers from voting or reduce their likelihood of purchasing when the product is made available. Hence, the seller's problem can be viewed as one of balancing the trade-off between *exploration* and *exploitation*.

To formalize the seller's problem, we introduce the following mathematical framework. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space endowed with a Poisson process $N(t)$ with rate Λ and let $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ be the usual filtration generated by N. Given this filtration, we define the set $\mathcal{T}$ of stopping times with respect to $\mathbb{F}$. At time $t = 0$, the seller starts a *voting phase* during which he records the number of votes (or clicks) that the product receives. This voting phase continues until a (possibly random) time $\tau \geq 0$ at which a decision to introduce or discard the product is made and we restrict $\tau \in \mathcal{T}$. (A trivial example, which appears to be one of the most common strategies used in the emerging crowdvoting implementations, is to set $\tau = T$ for some fixed time T.) We note that for $t \leq \tau$, N_t counts the number of votes that the product has received up to time t.

In terms of demand forecast, we assume that the seller starts the voting phase with a prior belief $\lambda_0 = \mathbb{E}[\Lambda]$. As time goes by, and the voting process N_t unfolds, the seller updates this prior according to Bayes rule. It follows that, at time t, the seller's belief $\lambda_t := \mathbb{E}[\Lambda|\mathcal{F}_t]$ satisfies:

Lemma 1 (Belief Process) *The process $\{(\lambda_t, \mathcal{F}_t) : t \geq 0\}$ is a martingale that solves the stochastic differential equation*

$$\mathrm{d}\lambda_t = \eta(\lambda_{t-}) \left[\mathrm{d}N_t - \lambda_t \, \mathrm{d}t\right] \qquad \textit{where} \quad \eta(\lambda) := \frac{(\lambda_{\mathrm{H}} - \lambda)\,(\lambda - \lambda_{\mathrm{L}})}{\lambda}.$$

In addition, for a given initial condition $\lambda_0 = \lambda$, the value of λ_t depends on the path of the voting process $\{N_s : 0 \leq s \leq t\}$ only through the value of the likelihood ratio $L(t, n) := (\lambda_{\mathrm{H}}/\lambda_{\mathrm{L}})^n \exp(-(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})\,t)$ according to

$$\lambda_t = \frac{\lambda_{\mathrm{L}}\,(\lambda_{\mathrm{H}} - \lambda) + \lambda_{\mathrm{H}}\,(\lambda - \lambda_{\mathrm{L}})\,L(t, N_t)}{\lambda_{\mathrm{H}} - \lambda + (\lambda - \lambda_{\mathrm{L}})\,L(t, N_t)}.$$

PROOF: All proofs are relegated to the appendix. □

The function $\eta(\lambda)$ is the size of the jump of the belief process at an arrival epoch (*i.e.*, a measure of the 'amount' of learning carried by a vote).

We complete the model description by specifying the objective function. We assume that the seller is a risk-neutral agent who maximizes the expected discounted value of the cash-flows generated by his retail business. Note that all these cash-flows happen after τ, since there are no sales or financial transactions during the voting phase. At time τ, the seller stops the voting process and decides whether or not to launch the product based on the available information at this time. In the case that the product is discarded, we assume the seller receives a fixed reward R (possibly zero) which captures the opportunity cost of his business (see Section 2.1 below for further discussion). On the other hand, if the seller decides to introduce the product then he incurs a fixed cost K and collects the revenues from sales, which come from two sources:

i. REVENUES FROM THE REGULAR SELLING PHASE: If the product is introduced in the market, customers arriving after τ will be able to purchase the product and the seller will collect these sales revenues. We aggregate this future stream of revenues in a function $G_{\mathrm{R}}(\Lambda)$, which we assume depends on the unknown demand rate Λ. In particular, –and for a number of reasons that we discuss in Section 2.1– we focus on the case in which future payoffs can be expressed as a linear function of the demand, that is, $G_{\mathrm{R}}(\Lambda) := \beta + \delta\,\Lambda$, for two constants β and $\delta > 0$.

ii. REVENUES FROM THE VOTING PHASE: If the product is set first onto a voting phase, voters, who by definition have showed interest in the product, would possibly come back and purchase it if commercialized. They represent the second source of revenues. We denote by N_τ the number of voters accumulated by τ of which a fraction (but, not necessarily all) purchases the product; it is useful to think of those revenues as some form of *pre-orders* or *pre-sales* that build-up during the voting phase. We assume that the seller's expected revenues generated during the voting phase are collected at time τ and are aggregated in a function G_{V} which we assume is linear in N_τ, $G_{\mathrm{V}}(N_\tau) := \phi\,N_\tau$. In Section 2.1, we further motivate this point and discuss specifically how this model encompasses for instance the possibility of offering a price discount (e.g., a coupon) to voters that decide to purchase the product; in Section 6 we look at additional variants of this model.

Putting all the pieces together, the seller's payoff at time τ conditional on the value of Λ and the number of voters $N_\tau = n$ is given by

$$G(\Lambda, n) := \max\left\{ \underbrace{R}_{discard}, \underbrace{G_{\mathrm{R}}(\Lambda) + G_{\mathrm{V}}(n) - K}_{commercialize} \right\}. \tag{1}$$

As a result, the seller's problem can be described by the following optimization problem:

$$\Pi(\lambda) := \sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda_\tau, N_\tau)\right] \quad \text{subject to} \quad \mathrm{d}\lambda_t = \eta(\lambda_{t-})\left[\mathrm{d}N_t - \lambda_t\,\mathrm{d}t\right] \quad \text{and} \quad \lambda_0 = \lambda, \tag{2}$$

where r is the seller's discount factor and $\Pi(\lambda)$ is his expected discounted payoff under an optimal strategy, which is parametrized by his initial (prior) belief $\lambda = \mathbb{E}[\Lambda]$. We will tackle the solution of (2) using dynamic programming and so we find convenient to extend $\Pi(\lambda)$ and define the seller's value function $\Pi(\lambda, n)$, which is the seller's expected discounted payoff-to-go if he finds himself in a state where n votes have been already collected and his current belief about the expected value of Λ is λ. In other words,

$$\Pi(\lambda, n) := \sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda_\tau, n + N_\tau)\right] \quad \text{subject to} \quad \mathrm{d}\lambda_t = \eta(\lambda_{t-})\left[\mathrm{d}N_t - \lambda_t\,\mathrm{d}t\right] \quad \text{and} \quad \lambda_0 = \lambda. \tag{3}$$

Despite our stylized mathematical formulation, solving the optimal stopping problem (3) presents some challenges that have limited our ability to provide a simple characterization of $\Pi(\lambda, n)$ and an optimal stopping time $\tau^*(\lambda, n)$. For this reason, we have postponed the analysis of the general case to Section 5. Before, we first discuss two special cases that help us understand the structure of an optimal solution and that we also believe are important in their own right. First, in Section 3, we consider the case of *full demand information* in which the seller knows in advance the true value of Λ. Then, in Section 4, we investigate the case in which $\phi = 0$, which we refer to as a *pure voting system*. We analyze these two special cases separately because they allow us to isolate and measure the impact of the two main features of our crowdvoting system, namely,

(a) Financial Component: Under full information, votes carry no demand information and only impact the firm's cash-flows in the form of pre-orders at the time of introduction, and

(b) Informational Component: In contrast, in a pure voting system, votes have no direct impact on the firm's cash-flows and only affect the seller's forecast of the actual demand.

In addition to isolating these two effects, the payoffs associated with these extreme cases can also be used to derive simple upper and lower bounds for the value of $\Pi(\lambda)$ in (2). In this regard, the following properties of $\Pi(\lambda)$ will prove useful in deriving these bounds as well as in the analysis in the following sections.

Proposition 1 *The seller's value function $\Pi(\lambda, n)$ is increasing and convex in both λ and n.*

2.1 Discussion of the Model Assumptions

One aspect of the model that requires further discussion is our assumption that sales revenues are linear in the demand rate Λ, i.e., $G_{\mathrm{R}}(\Lambda) = \beta + \delta\,\Lambda$. On one hand, we can view this assumption as a tractable first-order linear approximation of a more complex dependence of $G_{\mathrm{R}}(\Lambda)$ on Λ. Alternatively, if we denote by p_{R} the fixed price set during the selling phase and suppose that (i) the demand process after introduction is statistically equivalent to the voting process N_t, (ii) the seller is uncapacitated (e.g., has infinite inventory) and (iii) operates over an infinite horizon, then under these conditions, we have that

$$G_{\mathrm{R}}(\Lambda) = \mathbb{E}\left[\int_{\tau}^{\infty} e^{-r\,(t-\tau)} p_{\mathrm{R}}\,\mathrm{d}N_t|\Lambda\right] = \frac{p_{\mathrm{R}}\,\Lambda}{r}$$

(see Chapter II in Brémaud, 1980 for details on how to define the stochastic integral above and compute its expectation). Of the three conditions above, condition (i) is probably the most questionable as one would expect the behavior of voters to be somehow different (although correlated) to that of actual buyers. For instance, since voting is essentially "free", we can expect that people visiting the seller's website will be more prone to vote for the product if they like it than to buy it at the price p_{R}. However, if there is a linear relationship between the demand intensity of buyers after time τ and the arrival rate of voters prior to τ (e.g., $\Lambda_{\text{buyers}} = x\,\Lambda$ for some $x > 0$) then the linearity of $G_{\mathrm{R}}(\Lambda)$ on Λ would still hold.

Another aspect of the model that needs further motivation is the linearity of the expected sales from the voting phase i.e., $G_{\mathrm{V}}(n) = \phi\,n$. Given a number of votes N_τ accumulated at the end of the voting phase, it seems reasonable to assume that not every voter ends up buying the product. For instance, a fraction of those might be *impulse buyers* that would have bought the product immediately at the time of voting (if it was available) but would not return to purchase it at the introduction time τ. To capture this behavior, we assume that with probability $\theta \in [0, 1]$ a vote will be converted into a sale if the product is introduced (i.e., a fraction $1 - \theta$ of voters are impulse buyers). So, for instance, if those voters that decide to purchase the product at introduction are charged the same price p_{R} set during the selling phase then, conditional on N_τ, the seller's expected revenues generated during the voting phase and collected at time τ are given by $G_{\mathrm{V}}(N_\tau) = p_{\mathrm{R}}\,\theta\,N_\tau$ i.e., $\phi = p_{\mathrm{R}}\,\theta$.

The previous discussion touches upon another aspect of our model that is worth commenting on, namely, the effect of prices on the arrival rates of voters and buyers. In order to keep our formulation as simple and parsimonious as possible, we have not explicitly modeled this dependence in our base model. However, we do investigate this connection in Section 6.1; there, we extend our model to the case in which the intensity of the voting process is given by $\Lambda\,\bar{F}_{\mathrm{V}}(p_{\mathrm{V}})$, for some decreasing function $\bar{F}_{\mathrm{V}}$ that captures the voters' reservation price distribution where, p_{V}, is the price that the seller offers to voters if they return to buy the product (if it is commercialized).

Besides affecting the speed of voting, the price p_{V} can also impact the likelihood that a voter will return to purchase the product, i.e., $\theta = \theta(p_{\mathrm{V}})$. More generally, one would expect that the likelihood that a voter returns to buy the product depends not only on the level of discount but also on the amount of time that the voter must wait. Someone that voted a long time ago is probably less likely to buy the product than someone who voted close to the introduction time. We report on this variation of the model in Section 6.2 where we consider the case in which a voter remains interested in buying the product for a randomly distributed amount of time.

Finally, another aspect of the model that deserves some discussion is the parameter R, i.e., the reward that the seller collects if he ends up discarding the product. This parameter should be interpreted as the opportunity cost that the seller incurs as he commits his operations (or a portion of it) to the specific product that he is testing and possibly commercializing. For instance, the retailer can be restricted to test a limited number of products in a given category at any given time. In this case, the opportunity cost should be computed on the basis of the (historical) average profit that a product generates in this category. It is not uncommon, however, to see many online retailers using crowdvoting systems that display a very large list of products (most of them direct substitutes) at any given time. This is a practice that suggests a relatively small value of R. For this reason, as part of our analysis, we will consider the special case of $R = 0$ and discuss the practical implications for the design of a voting system.

2.2 Notation and Conventions

In order to streamline the presentation, and without loss of generality, we find convenient to introduce a couple of changes of variables. From the definition of the payoff function G in equation (1) we have that

$$\frac{G(\Lambda, n)}{\delta} = \max\left\{\frac{R}{\delta}\,,\ \frac{\beta - K + \phi\, n}{\delta} + \Lambda\right\}.$$

By appropriately changing the units that we use to measure payoffs, we can normalize δ to be equal to 1. In addition, we set the change of variable $\beta \leftarrow \beta - K$ so that the payoff function can be rewritten in the following compact form:

$$G(\Lambda, n) = \max\{R\ ,\ \beta + \phi\, n + \Lambda\}. \tag{4}$$

Observe that the parameter β can take negative values. Finally, to avoid trivial cases, we will assume that

$$\beta + \lambda_{\mathrm{L}} < R, \qquad 0 < \lambda_{\mathrm{L}} < \lambda_{\mathrm{H}} \qquad \text{and} \quad r > 0. \tag{5}$$

If the first inequality does not hold then the option of discarding the product can never be optimal. The other two conditions on the values of λ_{L}, λ_{H} and r are self-evident.

3 Full Information

In this section, we analyze the case of *full information* in which the seller knows the true value of the demand rate $\Lambda = \lambda$ at time $t = 0$. (We can think of this as a special case of the general model in which $\lambda_{\mathrm{L}} = \lambda_{\mathrm{H}} = \lambda$.) This is an important benchmark that we will use to quantify the value of a voting system from a demand learning perspective. At the same time, this full information case is helpful to understand the *financial* role played by a voting system. Indeed, as we will see even if there is nothing to learn the seller might still want to use a voting system because of its impact on his cash-flows at the product's introduction time. The full information policy can also be implemented as *certainty equivalent* type of heuristic for the case in which Λ is not known.

To get some intuition as to why a voting system is beneficial in this full information case, note that the seller –by going through a voting phase first and delaying the decision to introduce the product– is essentially postponing the disbursement of the fixed launching cost and is simultaneously accumulating pre-orders. This "inventory" of voters creates at the time of commercialization an instantaneous boost in revenues which could be sufficiently large to make the decision of launching the product an optimal one, even if launching immediately at time $t = 0$ is not.

We will denote by $\Pi^{\mathrm{D}}(\lambda)$ the seller's expected discounted payoff if $\Lambda = \lambda$. More generally, we will denote by $\Pi^{\mathrm{D}}(\lambda, n)$ the seller's value function if he has already collected n votes. It follows that $\Pi^{\mathrm{D}}(\lambda) = \Pi^{\mathrm{D}}(\lambda, 0)$. The seller's optimization problem in this full information case is given by

$$\Pi^{\mathrm{D}}(\lambda, n) := \sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda, n + N_\tau)\right],$$

where N_t is a Poisson process with rate λ. Given the Markovian nature of the problem, we can restrict our attention to hitting times $\tau_m := \inf\{t > 0, N_t = m\}$ for some nonnegative integer m. Furthermore, it is not hard to see that in this full information case it is never optimal for the seller to introduce a voting system and to later decide to discard the product. As a result, we can rewrite the optimization problem as follows:

$$\begin{aligned} \Pi^{\mathrm{D}}(\lambda, n) &= \max\Big\{R\,,\ \max_{m\in\mathbb{N}} \mathbb{E}\left[e^{-r\,\tau_m}\,(\beta + \lambda + \phi\,(n+m))\right]\Big\} \\ &= \max\left\{R\,,\ \max_{m\in\mathbb{N}} \left(\frac{\lambda}{r+\lambda}\right)^m (\beta + \lambda + \phi\,(n+m))\right\}, \end{aligned}$$

where the second equality uses the fact that τ_m has an Erlang distribution with parameters (m, λ). Let us define the auxiliary functions

$$H(\lambda, n, m) := \left(\frac{\lambda}{\lambda + r}\right)^m (\beta + \lambda + \phi\,(m+n)) \qquad \text{and} \qquad m(\lambda, n) := \arg\max_{m\in\mathbb{N}} \Big\{H(\lambda, n, m)\Big\}.$$

For fixed λ and n, the function $H(\lambda, n, m)$ is unimodal in m. As a result, it can be easily proven that

$$m(\lambda, n) := \left\lceil \frac{\lambda\,(\phi - r) - r\,\beta}{\phi\, r} - n \right\rceil^+ \tag{6}$$

where $\lceil x \rceil$ is the *ceiling* function (*i.e.*, the smaller integer greater than or equal to x) and $x^+ := \max\{0, x\}$. The following result follows.

Proposition 2 *Under full information, the seller's value function is equal to*

$$\Pi^{\mathrm{D}}(\lambda, n) = \max\left\{R,\ \left(\frac{\lambda}{r+\lambda}\right)^{m(\lambda,n)} \Big(\beta + \lambda + \phi\,(n + m(\lambda, n))\Big)\right\}.$$

Furthermore, the seller uses the voting system until $m^*(\lambda) := m(\lambda, 0)\, \mathbb{1}(\Pi^{\mathrm{D}}(\lambda) > R)$ *votes are collected and launches the product only if* $\Pi^{\mathrm{D}}(\lambda, 0) > R$.

As we mentioned above, under full information, the value of a voting system is that it gives the seller an opportunity to postpone the costs of introducing the product while she accumulates pre-orders resulting from the voting process. However, if $m^*(\lambda) = 0$, the voting system does not bring any additional value, *i.e.*, immediate stopping is optimal. Intuitively, this happens in the extreme and opposite cases in which the fixed cost of introducing the product is too high or too low compared to the revenues from sales. In the former case the best course of action is to discard the product immediately while in the latter case the best decision is to launch it immediately without going through a voting phase.

The following is a straightforward corollary of Propositions 1 and 2 that provides necessary and sufficient conditions that ensure that a product will never be discarded.

Corollary 1 *The product will never be discarded if and only if* $\Pi^{\mathrm{D}}(\lambda_{\mathrm{L}}) > R$. *For this inequality to hold, the following condition is necessary* $\lambda_{\mathrm{L}}\,(\phi - r) > r\,\beta$.

To prove this result, note that by (i) the monotonicity of $\Pi(\lambda)$ (see Proposition 1) and (ii) the fact that $\lambda = \lambda_{\mathrm{L}}$ is an absorbing state for the dynamics of λ_t (see Lemma 1), we have that $\Pi(\lambda) \geq \Pi(\lambda_{\mathrm{L}}) = \Pi^{\mathrm{D}}(\lambda_{\mathrm{L}}) > R$. As a result discarding is never optimal. Furthermore, for $\Pi^{\mathrm{D}}(\lambda_{\mathrm{L}}) > R$ to hold, we must have $m(\lambda_{\mathrm{L}}, 0) > 0$ which implies the condition $\lambda_{\mathrm{L}}\,(\phi - r) > r\,\beta$.

Figure 1 depicts the form of an optimal policy in the (λ, n) space under full information, where n is the number of votes that the seller has already received. Of course, at time $t = 0$, the seller has received no

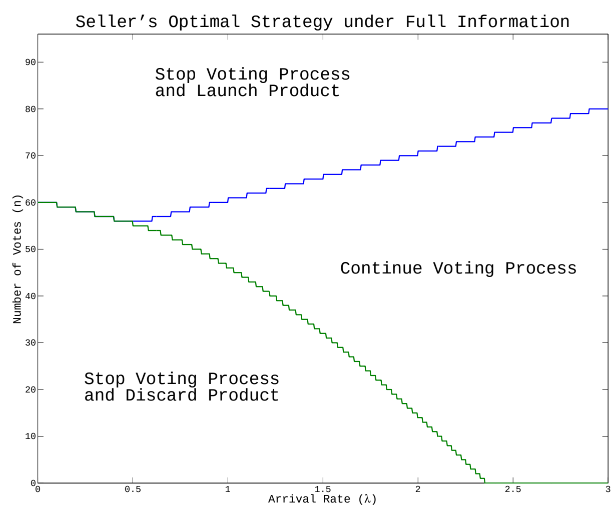

Figure 1: Data: : $\beta = -5$, $R = 1$, $\phi = 0.1$ and $r = 5\%$.

votes and the implementation of this optimal policy depends exclusively on what is optimal at the level $n = 0$ (*i.e.*, the X-axis). For instance, in this example, the voting system will be used only if $\lambda \gtrsim 2.4$. Otherwise, the product should be discarded immediately. It is worth noting that if $R = 0$, then the product will always be commercialized either directly if λ is high enough or after a voting phase that will be long enough to dilute the fixed costs and accumulate the necessary votes to generate positive payoff.

Proposition 2 is also useful in the design of a voting system. For instance, if the seller decides to use the voting system then he can announce at time $t = 0$ that $m^*(\lambda)$ votes will be collected before launching the product; making the system more transparent to the customers; (alternatively, he can estimate the amount of time that the voting system will be open (*e.g.*, in expectation $m^*(\lambda)/\lambda$ units of time) and commit to a fixed duration based on this estimate.). It is interesting to note that the value $m(\lambda, n)$ can be either increasing or decreasing in λ depending on whether $\phi > r$ or $\phi < r$, respectively. Hence, for large values of ϕ the seller will tend to collect more votes before introducing the product as the demand intensity increases. At the same time, the average time $m^*(\lambda)/\lambda$ that the voting system will be open, if it is implemented, is decreasing in λ when $\beta < 0$ (*i.e.*, when there is a positive launching cost). †

Let us conclude this section deriving upper and lower bounds for the seller value function $\Pi(\lambda, n)$ in (3) which are based on the function $H(\lambda, n, m)$ and the full information payoff $\Pi^{\mathrm{D}}(\lambda, n)$.

Proposition 3 *For all $\lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$ and $n \in \mathbb{N}$, the seller's expected discounted payoff $\Pi(\lambda, n)$ satisfies*

$$\underline{\Pi}(\lambda, n) := \max \Big\{ R \ , \ \max_{m \in \mathbb{N}} \mathbb{E}\big[H(\Lambda, n, m)\big] \Big\} \leq \Pi(\lambda, n) \leq \mathbb{E}\Big[\Pi^{\mathrm{D}}(\Lambda, n)\Big] =: \overline{\Pi}(\lambda, n),$$

where the expectations are taken with respect to Λ with prior $\mathbb{E}[\Lambda] = \lambda$.

A few remarks about this result and implications are in order. First of all, it is worth noticing that the bounds $\underline{\Pi}(\lambda, n)$ and $\overline{\Pi}(\lambda, n)$ are tight in the following cases:

(i) When R is large. Indeed, since the function $H(\lambda, n, m)$ is increasing in λ, it follows that if $R \geq \max_{m \in \mathbb{N}} H(\lambda_H, n, m)$ then $\overline{\Pi}(\lambda, n) = R$. In this case, the optimal action is to stop and discard the product.

(ii) When Λ is deterministic (*i.e.* $\lambda_{\mathrm{L}} = \lambda_{\mathrm{H}}$). In this case, we trivially have that $\underline{\Pi}(\lambda, n) = \overline{\Pi}(\lambda, n)$ for all $\lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$ and $n \in \mathbb{N}$.

(iii) When $m(\lambda_{\mathrm{L}}, n) = m(\lambda_{\mathrm{H}}, n) = m^* > 0$ and $\Pi^{\mathrm{D}}(\lambda, n) > R$. Indeed, from equation (6) it follows that if $m(\lambda_{\mathrm{L}}, n) = m(\lambda_{\mathrm{H}}, n) = m^*$ then we must have $m(\lambda, n) = m^*$ for all $\lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$. As a result, $\max_{m \in \mathbb{N}} H(\lambda, n, m) = H(\lambda, n, m^*) \geq H(\lambda_{\mathrm{L}}, n, m^*) > R$ for all $\lambda \in [\lambda_{\mathrm{L}}.\lambda_{\mathrm{H}}]$ (where the strict inequality follows from the assumption $\Pi^{\mathrm{D}}(\lambda, n) > R$). It follows that $R < \max_{m \in \mathbb{N}} \mathbb{E}[H(\Lambda, n, m)] = \mathbb{E}[H(\Lambda, n, m^*)] = \mathbb{E}[\Pi^{\mathrm{D}}(\Lambda, n)]$ and we conclude that $\Pi(\lambda, n) = \overline{\Pi}(\lambda, n)$.

A potential drawback of these bounds is the fact that they are based on a model with full information. As a result, one should not expect that they perform particularly well in capturing the value of demand learning. Take for instance the extreme case in which $\phi = 0$. In this situation, votes generate no incremental revenue and so the voting system only provides an opportunity for demand learning. It is not hard to see that if $\phi = 0$ then $m(\lambda, n) = 0$, that is, the voting system will never be used. To handle this issue, in the following section we explore in detail the case in which $\phi = 0$ and provide a complete characterization of an optimal solution in this *pure voting* scenario.

4 Pure Voting System

In this section we discuss the case in which none of the voters arriving in $[0, \tau]$ will return to buy the product if it is made available. As a result, this voting period becomes a *pure demand learning* phase

†Although theoretically possible, we expect the case $\phi > r$ to be somehow atypical in practice. Indeed, based on our interpretation of the payoff functions $G_{\mathrm{R}}(\Lambda)$ and $G_{\mathrm{V}}(n)$ in Section 2.1 and the scaling of the parameters in Section 2.2, the (normalized) value of ϕ is equal to the product of the retail price p_{R} and the likelihood θ that a voter will return to buy the product divided by the value of $\delta = p_{\mathrm{R}}/r$. Hence, under these assumptions we get $\phi = \theta\, r$ and since $\theta \leq 1$ we expect $\phi \leq r$. Roughly speaking, these calculations reveal that we could expect $\phi > r$ in situations in which the average price paid by voters is substantially larger than the average retail price charged after introduction.

that generates no direct revenues in the form of pre-orders. From a modeling perspective, this scenario corresponds to the case in which $\phi = 0$ and the seller's payoff becomes independent of N_τ.

Despite the fact that this pure demand learning model can be viewed as a special case of the general model with $\phi \geq 0$, the solution method that we use to solve it differs significantly from the case $\phi > 0$. As we discuss in Section 5, the case where $\phi > 0$ is reduced to solving a finite horizon stochastic control problem using backward recursion. On the other hand, the pure voting case, that we tackle in this section, is formulated as an infinite horizon stochastic control problem whose solution relies on an application of the functional Picard fixed point Theorem. The value function satisfies a set of quasi-variational inequalities (QVIs) in the form of a delayed ODE with free boundary conditions. In this regard, our analysis in this section is related to Peskir and Shiryaev (2000).

With a slight abuse of notation, we redefine in this section $G(\lambda) := G(\lambda, 0) = \max\{R,\ \beta + \lambda\}$. In order to avoid trivial solution we assume in this section that $\beta + \lambda_{\mathrm{L}} < R < \beta + \lambda_{\mathrm{H}}$, otherwise immediate stopping would be optimal.

The value function in this pure voting model is given by

$$\Pi(\lambda) = \sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda_\tau)\right] \quad \text{subject to} \quad \mathrm{d}\lambda_t = \eta(\lambda_{t-})\left[\mathrm{d}N_t - \lambda_t\, \mathrm{d}t\right], \quad \lambda_0 = \lambda. \tag{7}$$

In order to solve the stopping problem in (7) we derive optimality conditions in the form of a set of partial differential inequalities that characterize the optimal stopping time. To develop these ideas, we start by introducing the notion of quasi-variational inequalities. To this end, let us define the set $\mathcal{D}$ of functions $Z(\cdot)$ given by

$$\mathcal{D} := \Big\{Z(\lambda) : [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}] \to \mathbb{R} \text{ such that } Z(\cdot) \text{ is continuous and has right derivative for every } \lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}})\Big\}.$$

We also define the operator $\mathcal{H}$ that applies on functions $Z \in \mathcal{D}$ as follows

$$\mathcal{H}Z(\lambda) := \lambda\left[Z(\lambda + \eta(\lambda)) - Z(\lambda) - \eta(\lambda)\, \partial_+ Z(\lambda)\right] - r\, Z(\lambda), \tag{8}$$

where $\partial_+ Z(\lambda) := \lim_{h \downarrow 0}(Z(\lambda + h) - Z(\lambda))/h$.

Definition 1 (QVI) *The function $Z(\lambda) \in \mathcal{D}$ satisfies the quasi-variational inequalities for the seller's problem (7), if for every $\lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}})$,*

$$\begin{aligned} &\mathcal{H}Z(\lambda) \leq 0 \\ &Z(\lambda) - G(\lambda) \geq 0 \\ &\big(Z(\lambda) - G(\lambda)\big)\, \mathcal{H}Z(\lambda) = 0. \qquad \square \end{aligned} \tag{9}$$

As we will show, a solution to these QVI conditions partition the interval $[\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$ into two regions: a *continuation* region in which the seller's optimal strategy is to keep the voting process open and an *intervention* region in which stopping the voting period is optimal.

$$\begin{aligned} \text{Continuation:} \quad & \mathcal{C} := \big\{\lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}] : Z(\lambda) > G(\lambda) \quad \text{and} \quad \mathcal{H}Z(\lambda) = 0\big\} \\ \text{Intervention:} \quad & \mathcal{I} := \big\{\lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}] : Z(\lambda) = G(\lambda) \quad \text{and} \quad \mathcal{H}Z(\lambda) \leq 0\big\} \end{aligned}$$

For every solution of the QVI we can associate a control $\tau \in \mathcal{T}$.

Definition 2 *Let $Z \in \mathcal{D}$ be a solution of the QVI in (9). We define the control τ as follows*

$$\tau = \inf\{t > 0 : Z(\lambda_t) = G(\lambda_t)\}.$$

We call it the QVI-control associated to Z.

We are now ready to connect the QVI conditions to the optimization problem in (7) by mean of the following *verification* theorem, which provides a set of sufficient conditions for the optimal value function of the seller's problem.

Theorem 1 (Verification) *Let $Z(\lambda) \in \mathcal{D}$ be a solution of the QVI in (9). Then,*

$$Z(\lambda) \geq \Pi(\lambda) \quad \textit{for every } \lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}].$$

In addition, if there exists a QVI-control associated to Z then it is optimal and $Z(\lambda) = \Pi(\lambda)$.

According to the previous theorem, we can search for the solution of the seller's optimization problem by solving the system of QVI equations in (9), which we can rewrite in the following more compact form (HJB conditions):

$$0 = \max\Big\{G(\lambda) - Z(\lambda)\ ,\ \lambda\left[Z(\lambda + \eta(\lambda)) - Z(\lambda) - \eta(\lambda)\, Z'(\lambda)\right] - r\, Z(\lambda)\Big\}, \tag{10}$$

with border conditions $Z(\lambda_{\mathrm{H}}) = G(\lambda_{\mathrm{H}})$ and $Z(\lambda_{\mathrm{L}}) = G(\lambda_{\mathrm{L}})$ since both $\lambda = \lambda_{\mathrm{H}}$ and $\lambda = \lambda_{\mathrm{L}}$ are absorbing states.

In what follows, we will solve the HJB equation above by taking advantage of the fact that the belief process λ_t has only forward jumps, i.e., $\eta(\lambda) > 0$ for all $\lambda \in (\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}})$. Our method works in three steps. First, we postulate that the stopping time τ that solves problem (7) is of the threshold type. That is, we assume that there exist two thresholds $\underline{\lambda}, \bar{\lambda} \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$ with $\underline{\lambda} \leq \bar{\lambda}$ such that the *continuation region* is the interval $(\underline{\lambda}, \bar{\lambda})$. Figure 2 depicts an example. We note that without lost of generality we can assume that $G(\bar{\lambda}) \geq R$. Otherwise, for any λ in the continuation region $G(\lambda) = R$ and an optimal strategy would be to stop immediately. The condition $G(\bar{\lambda}) \geq R$ is equivalent to $\bar{\lambda} \geq \tilde{\lambda}$, where $\tilde{\lambda} := R - \beta$. In passing, we note that the convexity of $\Pi(\lambda)$ (see Proposition 1) implies that the difference $\Pi(\lambda) - G(\lambda)$ is maximized at $\lambda = \tilde{\lambda}$. In other words, the value of using the voting system to learn demand is maximized when the seller's belief is equal to $\tilde{\lambda}$.
Second, in this continuation region, the HJB condition above leads to the ODE

$$\lambda\left[Z(\lambda + \eta(\lambda)) - Z(\lambda) - \eta(\lambda)\, Z'(\lambda)\right] - r\, Z(\lambda) = 0, \quad \text{for all } \lambda \in (\underline{\lambda}, \bar{\lambda}). \tag{11}$$

We solve this ODE and impose value-matching and smooth-pasting conditions to derive the values of $\underline{\lambda}$ and $\bar{\lambda}$. Finally, we use the verification Theorem 1 to prove that the candidate solution $Z(\lambda)$ that we have derived is indeed an optimal solution, that is, $Z(\lambda) = \Pi(\lambda)$ for all $\lambda \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$. To ease the exposition, we have relegated most of the technical steps of this program to the appendix and have concentrated on providing an intuitive description of the methodology.

Given a value of $\bar{\lambda}$ in $(\tilde{\lambda}, \lambda_{\mathrm{H}})$, we first solve the ODE in (11) in $(\lambda_{\mathrm{L}}, \bar{\lambda})$. To this end, we use the method of *integrating factors* to rewrite it as follows

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{(\lambda - \lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}} - \lambda)^{\alpha_{\mathrm{H}}}}\, Z(\lambda)\right] = \frac{(\lambda - \lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}} - \lambda)^{\alpha_{\mathrm{H}}}}\, \frac{Z(\lambda + \eta(\lambda))}{\eta(\lambda)}, \tag{12}$$

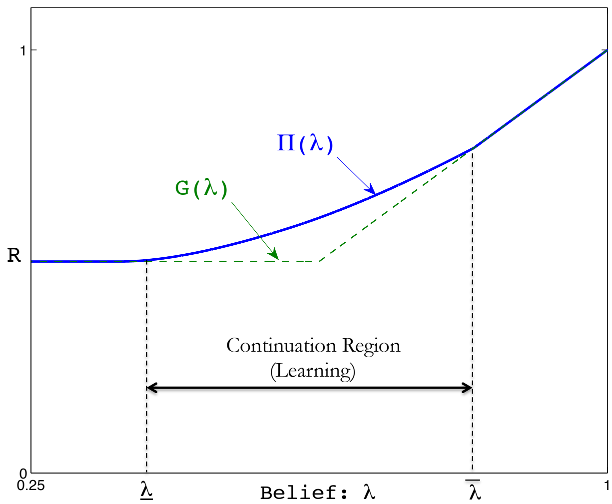

Figure 2: Value function $\Pi(\lambda)$ for $\beta = 0$, $R = 0.5$, $\lambda_{\rm L} = 0.25$, $\lambda_{\rm H} = 1$ and $r = 5\%$. The dashed line represents $G(\lambda)$.

where

$$\alpha_{\rm H} := \frac{r + \lambda_{\rm H}}{\lambda_{\rm H} - \lambda_{\rm L}} \quad \text{and} \quad \alpha_{\rm L} := \frac{r + \lambda_{\rm L}}{\lambda_{\rm H} - \lambda_{\rm L}}.$$

Note that (12) is a delayed differential equation since in its right-hand side the function Z is evaluated at $\lambda + \eta(\lambda)$. We take advantage of this feature and of the fact that $\eta(\lambda) > 0$ to develop a simple iterative procedure that solves this ODE in $(\lambda_{\rm L}, \bar{\lambda})$. The following definition will prove useful in our description of the algorithm.

Definition 3 (Lag Operator) *For any $\lambda \in [\lambda_{\rm L}, \lambda_{\rm H}]$ we define the lag operator $\mathcal{L}(\lambda)$ that satisfies*

$$\mathcal{L}(\lambda) + \eta(\mathcal{L}(\lambda)) = \lambda.$$

It is straightforward to show that $\mathcal{L}(\lambda)$ is increasing in λ and satisfies

$$\mathcal{L}(\lambda) = \frac{\lambda_{\rm L}\,\lambda_{\rm H}}{\lambda_{\rm L} + \lambda_{\rm H} - \lambda}.$$

Furthermore, for any $\lambda_0 \in [\lambda_{\rm L}, \lambda_{\rm H})$, the sequence $\lambda_n = \mathcal{L}(\lambda_{n-1})$ converges to $\lambda_{\rm L}$ as $n \to \infty$.

Iterative Procedure:

1. Input: Choose $\bar{\lambda} \in (\tilde{\lambda}, \lambda_{\rm H})$.

2. Initialization: Set $i = 0$, $\lambda_i = \bar{\lambda}$, $\mathcal{I}_i = [\lambda_i, \lambda_{\rm H}]$ and $Z_i(\lambda) = G(\lambda) = \beta + \lambda$ for all $\lambda \in \mathcal{I}_i$. Set $i = 1$.

3. Computation: Set $\lambda_i = \mathcal{L}(\lambda_{i-1})$ and $\mathcal{I}_i = [\lambda_i, \lambda_{i-1})$. For all $\lambda \in \mathcal{I}_i$, compute $Z_i(\lambda)$ solving the ODE

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}\, Z(\lambda)\right] = \frac{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}\, \frac{Z(\lambda + \eta(\lambda))}{\eta(\lambda)},$$

   with border condition $Z_{i-1}(\lambda_{i-1}) = Z_i(\lambda_{i-1})$.

4. Iteration: Set $i \leftarrow i + 1$ and go to Step 3 and iterate.

5. Output: The algorithm produces a function $\tilde{Z}(\lambda;\bar{\lambda})$ in $(\lambda_{\rm L},\lambda_{\rm H})$ parametrized by the value of $\bar{\lambda}$ in Step 1 given by $Z(\lambda;\bar{\lambda}) = Z_i(\lambda)$ for $\lambda \in \mathcal{I}_i,\ i=0,1,\ldots$ □

A few remarks about this procedure are in order. First of all, we note that the ODE in Step 3 is a conventional (i.e., non-delayed) differential equation in $Z_i(\lambda)$. This follows from the fact that the right-hand depends exclusively on the function $Z_{i-1}(\lambda)$ in the interval $\mathcal{I}_{i-1}$, which has been computed in the previous iteration of the algorithm. One direct implication of this feature is that the algorithm can be easily implemented using standard numerical methods for ODE. Another important property of the algorithm is that it produces a function $Z(\lambda;\bar{\lambda})$ that is continuous in $(\lambda_{\rm L},\lambda_{\rm H}]$. This follows from the border condition in Step 3. Furthermore, because the different pieces $Z_i(\lambda)$, $i=1,2,\ldots$ solve the same ODE in Step 3 and the function $\eta(\lambda)$ is continuous, one can show that $Z(\lambda,\bar{\lambda})$ is also differentiable in the interval $(\lambda_{\rm L},\lambda_{\rm H}]$ except (possibly) at $\lambda=\bar{\lambda}$.

The following proposition provides an explicit characterization of the function $Z(\lambda;\bar{\lambda})$ produced by the previous algorithm.

Proposition 4 *Let $\bar{\lambda} \in (\tilde{\lambda},\lambda_{\rm H})$ and define a decreasing sequence $\{\lambda_i\}_{i\geq -1}$ such that $\lambda_{-1}=\lambda_{\rm H}$, $\lambda_0=\bar{\lambda}$ and $\lambda_{i+1}=\mathcal{L}(\lambda_i)$ for $i\geq 1$. The ODE (12) admits a solution $Z(\lambda;\bar{\lambda})$, on $(\lambda_{\rm L},\bar{\lambda}]$ parameterized by $\bar{\lambda}$ and given by*

$$Z(\lambda;\bar{\lambda}) = Z_i(\lambda) \quad \textit{for all} \quad \lambda \in [\lambda_i,\lambda_{i-1}], \qquad i\geq 0$$

where,

$$Z_i(\lambda) = A_i\,(\lambda_{\rm H}-\lambda) + B_i\,(\lambda-\lambda_{\rm L}) + \frac{(\lambda_{\rm H}-\lambda)^{\alpha_{\rm H}}}{(\lambda-\lambda_{\rm L})^{\alpha_{\rm L}}}\sum_{n=0}^{i-1} C_{i,n}\,\ln^n\left(\frac{\lambda_{\rm H}\,(\lambda-\lambda_{\rm L})}{\lambda_{\rm L}\,(\lambda_{\rm H}-\lambda)}\right), \qquad \lambda\in[\lambda_i,\lambda_{i-1}].$$

The coefficients A_i, B_i and $\{C_{i,n} : n=0,\ldots,i-1\}$ satisfy the recursions $A_0=(\beta+\lambda_{\rm L})/(\lambda_{\rm H}-\lambda_{\rm L})$, $B_0=(\beta+\lambda_{\rm H})/(\lambda_{\rm H}-\lambda_{\rm L})$, $C_{0,0}=0$ and

$$A_i := \frac{A_{i-1}\,\lambda_{\rm L}}{(\lambda_{\rm H}-\lambda_{\rm L})\,\alpha_{\rm L}}, \qquad B_i := \frac{B_{i-1}\,\lambda_{\rm H}}{(\lambda_{\rm H}-\lambda_{\rm L})\,\alpha_{\rm H}}, \qquad C_{i,n} = \frac{\lambda_{\rm L}^{\alpha_{\rm H}}}{(\lambda_{\rm H}-\lambda_{\rm L})\,\lambda_{\rm H}^{\alpha_{\rm L}}}\,\frac{C_{i-1,n-1}}{n}, \quad n=1,2,\ldots,i-1,$$

and $C_{i,0}=K_i$, where K_i is a constant of integration that is computed imposing the value-matching condition $Z_i(\lambda_{i-1})=Z_i(\lambda_{i-1})$, for all $i\geq 1$.The function $Z(\lambda;\bar{\lambda})$ is convex in $(\lambda_{\rm L},\bar{\lambda})$ and satisfies $Z(\lambda;\bar{\lambda})\to\infty$ as $\lambda\downarrow\lambda_{\rm L}$.

The next step in our characterization of the value function $Z(\lambda)$ is to impose the optimality conditions in the HJB equation (10) to our candidate solution $Z(\lambda,\bar{\lambda})$ derived in Proposition 4. Using these conditions, we will determine the value of $\bar{\lambda}$ as well as $\underline{\lambda}$ that define the continuation region. To get some intuition about how we will use these optimality conditions, let us consider Figure 3 which depicts the function $Z(\lambda;\bar{\lambda})$ produced by the iterative procedure for three different values of $\bar{\lambda}\in\{0.8,0.826,0.85\}$. The dashed line represents $G(\lambda)$. We note that for $\bar{\lambda}=0.826$, the function $Z(\lambda;0.826)$ achieves its minimum at a value $\hat{\lambda}$ such that $Z(\hat{\lambda};0.826)=G(\hat{\lambda})=R$. In what follows, we argue that this property implies that $Z(\lambda,0.826)$ is indeed the value function $\Pi(\lambda)$ and that $\underline{\lambda}=\hat{\lambda}$.

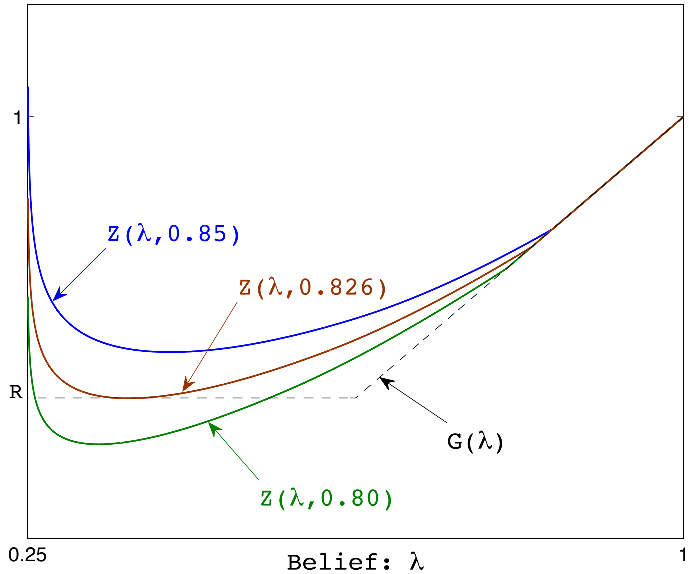

Figure 3: Value of $Z(\lambda;\bar\lambda)$ produced by the iterative procedure for four values of $\bar\lambda \in \{0.8, 0.826, 0.85\}$. For $\bar\lambda = 0.826$, the function $Z(\lambda;\bar\lambda)$ satisfies the smooth-pasting condition at $\underline{\lambda}$. The dashed line represents $G(\lambda)$. Data: $\beta = 0$, $R = 0.5$, $\lambda_{\rm H} = 1$, $\lambda_{\rm L} = 0.25$ and $r = 5\%$.

To this end, let us consider first a value of $\bar\lambda$ such that the function $Z(\lambda;\bar\lambda) > G(\lambda)$ in the interval $(\lambda_{\rm L}, \bar\lambda)$. This is the case of $\bar\lambda = 0.85$ in Figure 3. This strict inequality and the HJB equation (10) imply that $Z(\lambda) = Z(\lambda;\bar\lambda)$ in $(\lambda_{\rm L}, \bar\lambda)$. But from Proposition 4 we know that $Z(\lambda;\bar\lambda)$ increases to ∞ as $\lambda \downarrow \lambda_{\rm L}$ and from Proposition 1 we know that the value function $\Pi(\lambda)$ is non-decreasing and converges to $G(\lambda_{\rm L}) = R$ as λ approaches $\lambda_{\rm L}$. This contradiction implies that the value of $\bar\lambda$ must satisfy the condition: $Z(\lambda;\bar\lambda) \le G(\lambda)$ for some $\lambda \in (\lambda_{\rm L}, \bar\lambda)$, which rules out the value $\bar\lambda = 0.85$ in our example.
Consider now a value of $\bar\lambda$ such that $Z(\lambda;\bar\lambda)$ intersects $G(\lambda)$ at some value $\hat\lambda \le \bar\lambda$. (This is the case of $\bar\lambda = 0.826$ and $\bar\lambda = 0.80$ in the figure.) Let $\underline{\lambda} := \max\{\lambda \le \bar\lambda : Z(\lambda;\bar\lambda) = G(\lambda)\}$. Because of the convexity of $Z(\lambda;\bar\lambda)$, one can show that $\underline{\lambda} \le \tilde\lambda$, that is, at $\underline{\lambda}$ we have $Z(\underline{\lambda};\bar\lambda) = G(\underline{\lambda}) = R$. The monotonicity of the value function and the HJB equation suggest that the only candidate solution in this case is

$$Z(\lambda) = \begin{cases} Z(\lambda;\bar\lambda) & \text{if } \lambda \in (\underline{\lambda}, \lambda_{\rm H}) \\ R & \text{if } \lambda \in [\lambda_{\rm L}, \underline{\lambda}]. \end{cases} \tag{13}$$

It is not hard to see that this function $Z(\lambda)$ satisfies the HJB condition (10) only if the two pieces match smoothly at $\lambda = \underline{\lambda}$[†]. In other words, the derivative of $Z(\lambda,\bar\lambda)$ at $\lambda = \underline{\lambda}$ has to be equal to 0. Note that in Figure 3 $Z(\lambda, 0.826)$ satisfies this condition (by construction) while $Z(\lambda, 0.80)$ violates it. To get some additional intuition about why this *smooth-pasting* condition is satisfied at optimality, note that the dynamics of the belief process in equation (7) are such that λ_t decreases smoothly in between consecutive arrival epochs of N_t and jumps forward when a voter arrives. This suggests that the value function $Z(\lambda)$ should be smooth at the lower threshold $\underline{\lambda}$ as this value is reached smoothly from the interior of the continuation region. Using this smooth-pasting condition, we can pinpoint the value of $\bar\lambda$, which in the

[†]To see this, suppose that the derivate of $Z(\lambda;\bar\lambda)$ is strictly positive at $\lambda = \underline{\lambda}$ (e.g., for $\bar\lambda = 0.80$ in Figure 3). Then, by continuity one can show that for the proposed solution $Z(\lambda)$ in (13), $\lambda, \big[Z(\lambda + \eta(\lambda)) - Z(\lambda) - \eta(\lambda)\, Z'(\lambda)\big] - r\, Z(\lambda) > 0$ in the interval $(\underline{\lambda} - \epsilon, \underline{\lambda})$ for some $\epsilon > 0$.

case of the example in Figure 3 leads to $\bar{\lambda} = 0.3784$.

The following proposition summarizes our previous discussion and provides some additional features of an optimal solution, *e.g.*, specific conditions on the parameters of the problem that guarantee the existence of $\bar{\lambda}$ satisfiying the required smooth-pasting condition.

Proposition 5
a) Let $Z(\lambda)$ be the function defined in (13), where $Z(\lambda;\bar{\lambda})$ is the function computed in Proposition 4 for some $\bar{\lambda} \in (\tilde{\lambda}, \lambda_{\rm H})$. Furthermore, supposed that $\bar{\lambda}$ is such that $Z(\lambda;\bar{\lambda})$ satisfies the value-matching and smooth-pasting conditions:

$$Z(\underline{\lambda};\bar{\lambda}) = R \qquad \textit{and} \qquad Z'(\underline{\lambda};\bar{\lambda}) = 0,$$

for some $\underline{\lambda} < \tilde{\lambda}$. Then, $\Pi(\lambda) = Z(\lambda)$ for all $\lambda \in (\lambda_{\rm L}, \lambda_{\rm H})$.
b) The existence of $\bar{\lambda} \in (\tilde{\lambda}, 1)$ satisfying the conditions above is guaranteed if and only if

$$(\lambda_{\rm H} - \tilde{\lambda})\,(\tilde{\lambda} - \lambda_{\rm L}) > R\,r. \tag{14}$$

The following corollary is an important byproduct of the previous proposition. For notational convenience, let us define

$$X := \frac{4\,R\,r}{(\lambda_{\rm H} - \lambda_{\rm L})^2} \qquad \text{and} \qquad \tilde{q} := \frac{\lambda_{\rm H} - \tilde{\lambda}}{\lambda_{\rm H} - \lambda_{\rm L}}.$$

Corollary 2 *For any prior $\lambda \in [\lambda_{\rm L}, \lambda_{\rm H}]$, immediate stopping (i.e., no learning) is an optimal strategy if condition (14) is not satisfied, that is, if (i) $X \geq 1$ or (ii)*

$$\tilde{q} \leq \frac{1}{2}\left(1 - \sqrt{1-X}\right) \quad \textit{or} \quad \tilde{q} \geq \frac{1}{2}\left(1 + \sqrt{1-X}\right).$$

In this case, $\Pi(\lambda) = G(\lambda)$ for all $\lambda \in [\lambda_{\rm L}, \lambda_{\rm H}]$.

A few words about this result are in order. First of all, to get some intuition about condition (i), we can view rR as a measure of the opportunity cost of waiting. This cost increase with the discount rate r and the opportunity cost R. On the other hand, the term $\lambda_{\rm H} - \lambda_{\rm L}$ is a measure of the *speed* of learning[‡]. Hence, condition (i) formalizes the intuition that if the opportunity cost is high and/or the speed of learning is low, then learning is not profitable. To understand condition (ii), let us recall first that we have assumed that $\beta + \lambda_{\rm L} < R < \beta + \lambda_{\rm H}$ to avoid trivial cases. As a result, we have that $\tilde{q} = (G(\lambda_{\rm H}) - G(\lambda_{\rm L}))/(\lambda_{\rm H} - \lambda_{\rm L})$. It follows then that condition (ii) formalizes the idea that if the maximum possible value of learning (measured by the difference of the payoff function under perfect information $G(\lambda_{\rm H}) - G(\lambda_{\rm L})$) is too small or too large then immediate stopping is optimal.
Table 1 summarizes the *value of learning* as a function of $\tilde{q}$ and X measured as follows

$$\mathcal{V} := \sup_{\lambda \in [\lambda_{\rm L}, \lambda_{\rm H}]} \left\{ \frac{\Pi(\lambda) - G(\lambda)}{G(\lambda)} \right\}.$$

Note that $\tilde{q} \in (0,1)$. Also, from condition (i) in Corollary 2, learning is not valuable if $X \geq 1$, i.e., $\mathcal{V} = 0$.

[‡]To see this, let $\mathbb{P}_{\rm H}$ and $\mathbb{P}_{\rm L}$ be two probability measures under which N_t is a Poisson process with rate $\lambda_{\rm H}$ and $\lambda_{\rm L}$, respectively. Then, the likelihood ratio $L = {\rm d}\mathbb{P}_{\rm H}/{\rm d}\mathbb{P}_{\rm L}$ satisfies

$$L_t = \left(\frac{\lambda_{\rm H}}{\lambda_{\rm L}}\right)^{N_t} \exp(-(\lambda_{\rm H} - \lambda_{\rm L})\,t),$$

and so the speed of learning is determined by the difference $\lambda_{\rm H} - \lambda_{\rm L}$.

		X				
		0.1	0.3	0.5	0.7	0.9
	0.1	15.43%	0.44%	0	0	0
	0.3	45.18%	4.45%	1.00%	0.32%	0
$\tilde{q}$	0.5	53.44%	6.32%	1.80%	0.46%	0.29%
	0.7	40.71%	4.40%	1.14%	0.38%	0
	0.9	10.37%	0.24%	0	0	0

Table 1: "Value of Learning" $\mathcal{V}$ as a function of $\tilde{q}$ and X. Data: $R = 0.3$ and $r = 8\%$.

Consistent with our previous discussion, the value of learning is maximized when X is small and the value of $\tilde{q}$ is close to 0.5. One particular scenario when this holds is the case of $R = 0$. In this numerical example, the possibility of learning can increase the seller's profits by more than 50% in some cases.

4.1 Asymptotic Analysis

One of the main difficulties in solving the firm's optimization problem in (7) is the fact that the belief process λ_t has discontinuous paths and the corresponding HJB equation is a delayed ODE. One possible way to go around this nuisance is to consider an approximation in which the magnitude of the jumps are small so that the paths of λ_t can be approximated by continuous ones. In order to derive this approximation but at the same time retain as much as possible the original structure of the problem, we construct a sequence of instances of the problem indexed by a non-negative integer k in such a way that as k grows large, the size of the jumps of λ_t converges to zero (i.e., λ_t converges to a continuous process), while the instantaneous volatility of λ_t is preserved.

For each $k \geq 1$, we consider an instance of the problem in which the demand rates $\lambda_{\rm H}$ and $\lambda_{\rm L}$, are given by

$$\lambda^k_{\rm H} := k + \varphi \sqrt{k} \quad \text{and} \qquad \lambda^k_{\rm L} := k \tag{15}$$

and let λ^k_t be the corresponding belief process for the $k^{\rm th}$ system. The *raison d'être* of the transformation in (15) is that the magnitude $\eta(\lambda^k_t)$ of the jumps of λ^k satisfies

$$\eta(\lambda^k_t) = \frac{(\lambda^k_{\rm H} - \lambda^k_t)(\lambda^k_t - \lambda^k_{\rm L})}{\lambda^k_t} \leq \frac{(\lambda^k_{\rm H} - \lambda^k_{\rm L})^2}{\lambda^k_t} = \frac{\varphi^2\, k}{\lambda^k_t} \leq \frac{\varphi^2\, k}{\lambda^k_{\rm L}} = \varphi^2.$$

Hence, as k grows large the magnitude of the jumps remains bounded by φ^2 while the value of λ^k_t grows linearly with k. As we will see, this property implies that the quality of the inference that we can make about whether $\Lambda = \lambda^k_{\rm H}$ or $\Lambda = \lambda^k_{\rm L}$ at any jump epoch of N_t is decreasing in k, or under an appropriate scaling, the jumps of λ^k_t converge to zero as $k \to \infty$. It is also worth noticing that this type of asymptotic regime that we are considering resembles the typical conditions imposed in the heavy traffic literature to prove weak convergence of discrete Markov processes into fluid and diffusion processes (e.g., Kurtz, 1978).

In order to formalize the ideas in the previous paragraph, we first need to change the state (belief) space since λ^k_t grows linearly with k and hence cannot have a well-defined limit. Instead, let us define the auxiliary belief process $q_t^k := \mathbb{P}(\Lambda = \lambda_{\mathrm{L}}^k | \mathcal{F}_t) = (\lambda_{\mathrm{H}}^k - \lambda_t^k)/(\lambda_{\mathrm{H}}^k - \lambda_{\mathrm{L}}^k)$ for the k^{th} system. From Lemma 1, it follows that q_t^k satisfies the SDE

$$\mathrm{d}q_t^k = q_t^k\,(1-q_t^k)(\lambda_{\mathrm{H}}^k - \lambda_{\mathrm{L}}^k)\,\left[\mathrm{d}t - \frac{\mathrm{d}N_t^k}{\bar{\Lambda}^k(q_t^k)}\right],$$

where N_t^k is a point process with intensity $\bar{\Lambda}^k(q_t^k) := \lambda_{\mathrm{L}}^k\,q_t^k + \lambda_{\mathrm{H}}^k\,(1-q_t^k)$.

Proposition 6 *Consider a sequence of problems indexed by k for which the arrival rates λ_{H}^k and λ_{L}^k are given by (15). Let $q_t^k = \mathbb{P}(\Lambda = \lambda_{\mathrm{L}}^k|\mathcal{F}_t)$ be the seller's belief process that $\Lambda = \lambda_{\mathrm{L}}^k$ with $q_0^k = q_0$, independent of k. Then, q_t^k converges weakly to q_t, solution of the SDE*

$$\mathrm{d}q_t = \varphi\,q_t\,(1-q_t)\,\mathrm{d}W_t,$$

where W_t is a Wiener process.

According to Proposition 6, the scaling in (15) leads to a diffusion approximation for the dynamics of the auxiliary belief process $q_t = \mathbb{P}(\Lambda = \lambda_{\mathrm{L}}|\mathcal{F}_t)$.

Let us turn now to the seller's objective function. Given our asymptotic result on q_t, we find convenient to rewrite the seller's expected payoff in terms of q_t rather than λ_t. Since $\lambda_t = \lambda_{\mathrm{L}}\,q_t + \lambda_{\mathrm{H}}\,(1-q_t)$, we set $G(q) := \max\{R, \beta + \lambda_{\mathrm{L}}\,q + \lambda_{\mathrm{H}}\,(1-q)\}$. As we scale the system with k, we would like the seller's optimization problem to remain well-posed, that is, we would like to ensure that the profit function remains bounded. For this, we must use an appropriate scaling for how the parameters R and β change with k. We define $R^k := \tilde{R}\,\sqrt{k}$ and $\beta^k = \tilde{\beta}\,\sqrt{k} - k$ for some fixed constants $\tilde{R}$ and $\tilde{\beta}$[§]. Under this scaling, we get that

$$\begin{aligned} G^k(q) \;&:=\; \max\left\{R^k\,,\,\beta^k + \lambda_{\mathrm{L}}^k\,q + \lambda_{\mathrm{H}}^k\,(1-q)\right\} \\ &=\; \sqrt{k}\,\max\left\{\tilde{R}\,,\,\tilde{\beta} + \varphi(1-q)\right\}. \end{aligned}$$

We define

$$\widetilde{G}(q) := \frac{G^k(q)}{\sqrt{k}} = \max\left\{\tilde{R}\,,\,\tilde{\beta} + \varphi(1-q)\right\}. \tag{16}$$

Based on the asymptotic results in Propositions 6 and the asymptotic scaled objective, we can approximate the firm's optimization problem as follows

$$\widetilde{\Pi}(q) = \sup_{\tau\in\mathcal{T}} \mathbb{E}_q\left[e^{-r\,\tau}\,\widetilde{G}(q_\tau)\right] \tag{17}$$

$$\text{subject to}\qquad \mathrm{d}q_t = \varphi\,q_t\,(1-q_t)\,\mathrm{d}W_t, \qquad q_0 = q. \tag{18}$$

As before the solution is of the threshold type, that is, there exist two constants $\underline{q}$ and $\bar{q}$ such that

$$\frac{(\varphi\,q\,(1-q))^2}{2}\,\Pi''(q) - r\,\Pi(q) = 0, \qquad \underline{q} < q < \bar{q} \tag{19}$$

$$\Pi(q) = G(q), \qquad q \in [0,1] - (\underline{q}, \bar{q}) \tag{20}$$

$$\Pi'(\underline{q}) = -1 \qquad \text{and} \qquad \Pi'(\bar{q}) = 0. \tag{21}$$

[§]Note that the condition $\beta^k + \lambda_{\mathrm{L}}^k < R^k < \beta^k + \lambda_{\mathrm{H}}^k$ that we need to avoid trivial cases boils down to $\tilde{\beta} < \tilde{R} < \tilde{\beta} + \varphi$.

After some manipulations, one can show that a general solution for the ODE in (19) is given by

$$\Pi(q) = K_0 \, \frac{(1-q)^A}{q^{A-1}} + K_1 \frac{q^A}{(1-q)^{A-1}}, \quad \text{where } A := \frac{1+\sqrt{1+8r/\varphi^2}}{2} \tag{22}$$

and K_0 and K_1 are two constants of integration. The next result follows.

Proposition 7 *Let $A := (1+\sqrt{1+8r/\varphi^2})/2$. There exist constants $\underline{q}, \bar{q} \in (0,1)$, K_0 and K_1 such that*

$$\widetilde{\Pi}(q) = \begin{cases} \tilde{\beta} + \varphi\,(1-q) & \text{if} \quad 0 \le q \le \underline{q} \\ K_0\,(1-q)^A\, q^{1-A} + K_1\,(1-q)^{1-A}\, q^A & \text{if} \quad \underline{q} \le q \le \bar{q} \\ \tilde{R} & \text{if} \quad \bar{q} \le q \le 1. \end{cases}$$

The constants are determined imposing smooth-pasting *conditions in* (20)-(21) *at $\underline{q}$ and $\bar{q}$. The function $\widetilde{\Pi}(q)$ is twice-continuously differentiable and convex in $(0,1)$.*

To complete our diffusion approximation, we now interpret the results in Proposition 7 in terms of the original (unscaled) primitive model. For this, we reverse the scaling using the fact that $k = \lambda_{\rm L}$ to get

$$\varphi = \frac{\lambda_{\rm H}-\lambda_{\rm L}}{\sqrt{\lambda_{\rm L}}}, \qquad \tilde{R} = \frac{R}{\sqrt{\lambda_{\rm L}}}, \qquad \tilde{\beta} = \frac{\beta+\lambda_{\rm L}}{\sqrt{\lambda_{\rm L}}}, \qquad \text{and} \qquad \widetilde{\Pi}(\lambda) = \sqrt{\lambda_{\rm L}}\,\widetilde{\Pi}(q).$$

Note that we have used a slight abuse of notation writing $\widetilde{\Pi}(\lambda)$ for the unscaled approximated value function. In addition, we must also write the threshold $\underline{q}$ and $\bar{q}$ in terms of threshold on the values of λ, specifically we have

$$\bar{\lambda} = \underline{q}\,\lambda_{\rm L} + (1-\underline{q})\,\lambda_{\rm H} \qquad \text{and} \qquad \underline{\lambda} = \bar{q}\,\lambda_{\rm L} + (1-\bar{q})\,\lambda_{\rm H}.$$

Putting all the pieces together, we get the following diffusion approximation:

$$\widetilde{\Pi}(\lambda) = \begin{cases} \beta + \lambda & \text{if} \quad \bar{\lambda} \le \lambda \le \lambda_{\rm H} \\ \mathcal{K}_0\,(\lambda-\lambda_{\rm L})^A\,(\lambda_{\rm H}-\lambda)^{1-A} + \mathcal{K}_1\,(\lambda-\lambda_{\rm L})^{1-A}\,(\lambda_{\rm H}-\lambda)^A & \text{if} \quad \underline{\lambda} \le \lambda \le \bar{\lambda} \\ R & \text{if} \quad \lambda_{\rm L} \le \lambda \le \underline{\lambda}, \end{cases} \tag{23}$$

where the constants $\underline{\lambda}$, $\bar{\lambda}$, $\mathcal{K}_0$ and $\mathcal{K}_1$ are determined imposing value-matching and smooth-pasting conditions at $\lambda = \bar{\lambda}$ and $\lambda = \underline{\lambda}$.

Figure 4 depicts an example comparing the value function $\Pi(\lambda)$ and the diffusion approximation $\widetilde{\Pi}(\lambda)$. We note that $\widetilde{\Pi}(\lambda)$ satisfies value-matching and smooth-pasting conditions at $\underline{\lambda}$ and $\bar{\lambda}$. This is in contrast to $\Pi(\lambda)$ that satisfies these conditions only at $\bar{\lambda}$.
Table 2 summarizes a set of computations experiments that numerically evaluate the quality of the diffusion approximation. In particular, the values reported in the table correspond to the *relative error*

$$\mathcal{E} := \sup_{\lambda\in[\lambda_{\rm L},\lambda_{\rm H}]} \left\{ \frac{|\Pi(\lambda) - \widetilde{\Pi}(\lambda)|}{\Pi(\lambda)} \right\}$$

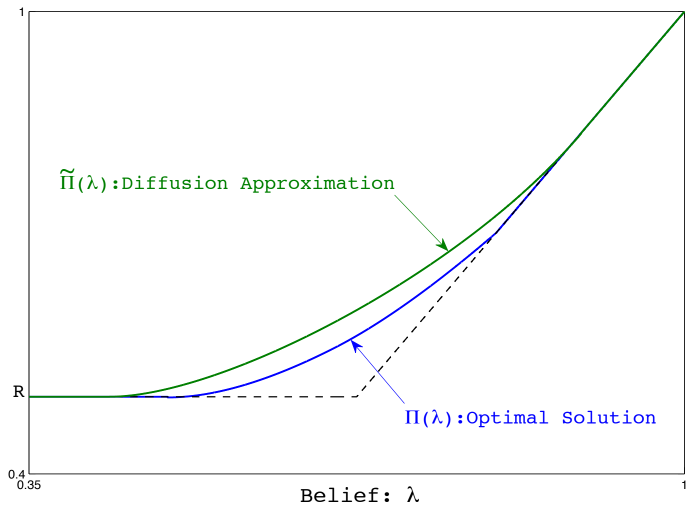

Figure 4: Diffusion approximation: Data: : $\beta = 0$, $\lambda_{\rm H} = 1$, $\lambda_{\rm L} = 0.35$, $R = 0.5$ and $r = 5\%$.

		$\varphi = (\lambda_{\rm H} - \lambda_{\rm L})/\sqrt{\lambda_{\rm L}}$				
		0.1	0.3	0.5	0.7	0.9
	0.1	0.85%	0.86%	2.73%	5.19%	6.61%
	0.3	0.36%	0.90%	2.21%	4.11%	5.94%
R	0.5	0.36%	0.69%	0.14%	2.91%	4.43%
	0.7	0.12%	0.68%	0.99%	1.54%	2.31%
	0.9	0.08%	0.12%	0.36%	0.62%	0.93%

Table 2: Relative error $\mathcal{E}$ as a function of R and φ. Data: $\beta = 0$, $\lambda_{\rm H} = 1$ and $r = 5\%$.

as a function of R and φ. Overall, the results show that the diffusion approximation is very accurate across a wide range of parameters. As expected, the approximation is particularly good for small values of φ.

Besides providing a simple representation of the value function, the diffusion approximation is also useful as it allows us to use standard results for one-dimensional diffusion processes (e.g., Section 5.5 in Karatzas and Shreve, 1991) to analyze its optimal solution. For instance, the following corollary characterizes the expected duration of the voting phase and the likelihood that a product will be eventually introduced into the market.

Corollary 3 *Let $(\underline{\lambda}, \bar{\lambda})$ be the optimal continuation region where $\underline{\lambda}$ and $\bar{\lambda}$ are given in equation (23). For any initial condition $\lambda \in (\underline{\lambda}, \bar{\lambda})$, let $\tau^* := \inf\{t > 0 : \lambda_t \in (\underline{\lambda}, \bar{\lambda})\}$ and $\gamma^* := \mathbb{P}(\lambda_{\tau^*} = \bar{\lambda} \,|\, \lambda_0 = \lambda)$. Then,*

$$\gamma^* = \frac{\lambda - \underline{\lambda}}{\bar{\lambda} - \underline{\lambda}} \qquad \text{and} \qquad \mathbb{E}[\tau^*] = \gamma^* \, \mathcal{T}(\bar{\lambda}) + (1 - \gamma^*) \, \mathcal{T}(\underline{\lambda}) - \mathcal{T}(\lambda),$$

where $\mathcal{T}(\lambda)$ is the function

$$\mathcal{T}(\lambda) := \frac{2}{\varphi^2} \frac{\lambda_H + \lambda_L - 2\lambda}{\lambda_H - \lambda_L} \ln\left(\frac{\lambda_H - \lambda}{\lambda - \lambda_L}\right).$$

5 Voting System with Preselling

In this section we study the seller's optimization problem defined in equation (2) for the case in which $\phi > 0$; that is, in expectation, every vote generates a positive revenue if the product is eventually launched. Recall that $\Pi(\lambda, n)$ is the seller's value function which is given by

$$\Pi(\lambda,n) = \sup_{\tau\in\mathcal{T}} \Big\{ \mathbb{E}\Big[e^{-r\,\tau} G(\lambda_\tau, N_\tau) \,|\, \lambda_0 = \lambda, N_0 = n \Big] \Big\}, \quad \text{for all } (\lambda,n) \in (\lambda_{\rm L}, \lambda_{\rm H}) \times \mathbb{N}. \tag{24}$$

Our solution method takes full advantage of the fact that the number of votes N_t is a monotonically increasing process, which is a property that allows us to interpret the two dimensional dynamic programming formulation as a one dimensional finite-horizon stochastic control problem. As a result we are able to use a recursive method that computes the seller's value function and characterizes the continuation and intervention regions. The method is obtained through a backward induction on the value of n; for this, we first argue that if the number of votes n reaches a sufficiently large value, then it is in the seller's best interest to introduce immediately the product in the market and capitalize on the instantaneous expected revenues $\phi\, n$ that these voters generate. We formalize this intuition in the following lemma, which uses the mapping $m(\lambda) = m(\lambda, 0)$ defined in equation (6). (Recall that $m(\lambda)$ is the optimal number of votes needed to launch the product in the full information case, *i.e.*, when $\Lambda = \lambda$ is known at time 0.)

Lemma 2 *Let $\bar n$ be given by*

$$\bar n := \max\left\{ m\left(\frac{r(R-\lambda_{\rm L})}{\phi - r}\right), m(\lambda_{\rm L}), m(\lambda_{\rm H}) \right\}.$$

Then, for all states (λ, n) with $\lambda_{\rm L} \le \lambda \le \lambda_{\rm H}$ and $n \ge \bar n$ it is optimal to stop immediately and to launch the product in the market. In this case, the seller's value function is equal to

$$\Pi(\lambda,n) = \beta + \lambda + \phi\, n, \qquad \text{for all } n \ge \bar n.$$

It is worth noticing that we can use $\bar n$ to (upper) bound the time of the voting campaign. Indeed, the quantity $\tau_{\bar n}$ which is the time it takes the voting process to collect $\bar n$ votes, is clearly an upper bound on the actual voting duration; and for example, the seller can easily compute its expected value,

$$\mathbb{E}[\tau_{\bar n}] = \left(\frac{\lambda_{\rm H} - \lambda}{\lambda_{\rm H} - \lambda_{\rm L}}\right)\frac{\bar n}{\lambda_{\rm L}} + \left(\frac{\lambda - \lambda_{\rm L}}{\lambda_{\rm H} - \lambda_{\rm L}}\right)\frac{\bar n}{\lambda_{\rm H}}.$$

Equipped with the upper bound $\bar n$, we can compute the value of $\Pi(\lambda, n)$ using backward induction in n.

Proposition 8 *The value function $\Pi(\lambda, n)$ satisfies recursively the optimality equation*

$$\Pi(\lambda,n) = \frac{(\lambda_{\rm H}-\lambda)^{\alpha_{\rm H}}}{(\lambda-\lambda_{\rm L})^{\alpha_{\rm L}}} \max_{\lambda_1\in(\lambda_{\rm L},\lambda]} \left[\int_{\lambda_1}^{\lambda} \frac{(x-\lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H}-x)^{\alpha_{\rm H}}} \frac{\Pi(x+\eta(x), n+1)}{\eta(x)}\, {\rm d}x + \frac{(\lambda_1-\lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H}-\lambda_1)^{\alpha_{\rm H}}} G(\lambda_1, n) \right], \tag{25}$$

with boundary condition $\Pi(\lambda, n) = \beta + \lambda + \phi\, n$, for all $\lambda \in [\lambda_{\rm L}, \lambda_{\rm H}]$ and $n \ge \bar n$.

Using arguments similar to those used in the previous section, one can show that the value function $\Pi(\lambda, n)$ satisfies the HJB equation:

$$0 = \max\Big\{ G(\lambda,n) - \Pi(\lambda,n)\ ,\ \lambda\left[\Pi(\lambda+\eta(\lambda), n+1) - \Pi(\lambda,n) - \eta(\lambda)\,\partial_+\Pi(\lambda,n)\right] - r\,\Pi(\lambda,n) \Big\}, \tag{26}$$

with border conditions $\Pi(\lambda_{\rm H}, n) = \Pi^{\rm D}(\lambda_{\rm H}, n)$, $\Pi(\lambda_{\rm L}, n) = \Pi^{\rm D}(\lambda_{\rm L}, n)$, for all $n \geq 0$ and $\Pi(\lambda, n) = \beta + \lambda + \phi\, n$ for all $\lambda \in [\lambda_{\rm L}, \lambda_{\rm H}]$ and $n \geq \bar{n}$ (recall the definition of the full information value function $\Pi^{\rm D}(\lambda, n)$ in Proposition 2).

In contrast to the pure voting case discussed in the previous section, we do not have a closed-form solution for the value function $\Pi(\lambda, n)$ and we must rely on equations (25) or (26) to compute an optimal solution numerically. As with most dynamic programming equations, part of the challenge in solving (25) is dealing with inner optimization over λ_1. In this regard, the following result greatly simplifies these numerical calculations.

Proposition 9 *For any initial belief $\lambda \in (\lambda_{\rm L}, \Lambda_{\rm H})$ and $n \geq 0$, let*

$$\mathcal{L}^*(\lambda, n) = \underset{\lambda_1 \in (\lambda_{\rm L}, \lambda]}{\operatorname{argmax}} \left[\int_{\lambda_1}^{\lambda} \frac{(x - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - x)^{\alpha_{\rm H}}} \frac{\Pi(x + \eta(x), n+1)}{\eta(x)} \, {\rm d}x + \frac{(\lambda_1 - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - \lambda_1)^{\alpha_{\rm H}}} G(\lambda_1, n) \right], \tag{27}$$

be the optimal hitting threshold in equation (25). Suppose that for some λ_0 and n, $\mathcal{L}^*(\lambda_0, n) = \hat{\lambda}_0 < \lambda_0$. Then, $\mathcal{L}^*(\lambda, n) = \hat{\lambda}_0$ for all $\lambda \in [\hat{\lambda}_0, \lambda_0]$.

This result essentially formalizes the intuition that the seller's only source of information is the one coming from the jumps of the voting process N_t. Indeed, if no jumps are occuring then the threshold to stop is fixed independent of the current value of λ_t. (We omit its proof as it follows directly from the characterization of $\Pi(\lambda, n)$ in Proposition 8.) We note that the mapping $\mathcal{L}^*(\lambda, n)$ fully characterizes the seller's optimal strategy. Indeed, the state (λ, n) belongs to the continuation or intervention regions depending on whether $\mathcal{L}^*(\lambda, n) < \lambda$ or $\mathcal{L}^*(\lambda, n) = \lambda$, respectively.

Figure 5 depicts the seller's optimal strategy in the (λ, n) space for two instances of the problem that differs only on the relationship between ϕ and r. In each panel, the shaded area corresponds to the *intervention region*, which is further divided into two subregions: in the top region it is optimal to stop the voting and launch the product while in the lower region it is optimal to stop and discard the product. Note that while the lower boundary of the *continuation region* is decreasing in λ the upper boundary can be either increasing (left panel) or decreasing (right panel). This seems consistent with the fact that $m(\lambda)$ is increasing or decreasing in λ which in turns depends on whether $\phi > r$ or $\phi < r$, respectively (see equation (6)).

5.1 Computational Experiments

In this section we conduct two types of numerical experiments to further investigate the properties and performance of our proposed voting system. First, we do a sensitivity analysis with respect to the main parameters of the model. Second, we consider two simple policies which are commonly used in practices and use them as benchmarks to evaluate the benefits of our proposed solution.

Before moving into the details of these computations, let us summarize some general guidelines that we have used in the design of our experiments. First of all, we impose the normalization $\lambda_{\rm H} = 1$. This is without loss of generality as we can always scale the units of time accordingly. Also, to avoid trivial cases, we assume that the condition $\lambda_{\rm L} + \beta \leq R$ is satisfied (otherwise, discarding the products is never an optimal decision). We also assume that $\beta \leq 0$, $R \geq 0$ and $R \leq \beta + \lambda_{\rm H}$, that is, launching the product

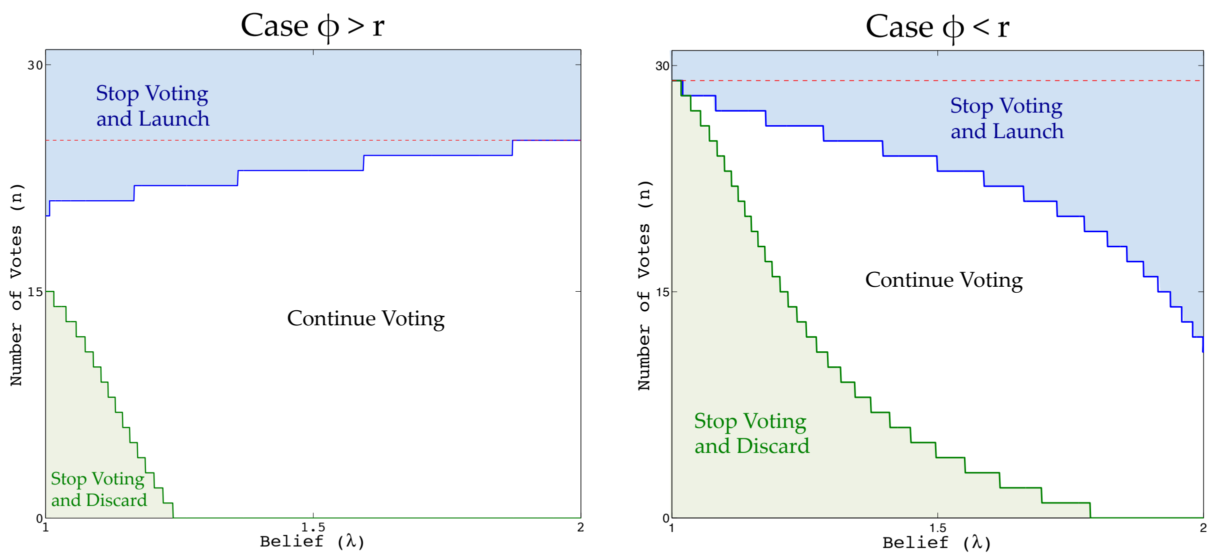

Figure 5: Optimal strategy for the instances $\lambda_{\mathrm{L}} = 1$, $\lambda_{\mathrm{H}} = 2$, $\beta = -1$, $R = 1$, $r = 5\%$ and $\phi = 0.065$ (left panel) and $\phi = 0.035$ (right panel). The top dashed line is at the level $\bar{n} = 25$ (left panel) and $\bar{n} = 29$ (right panel).

is costly, the opportunity cost of discarding the product is nonnegative and a high demand product will always be launched (even in a pure voting scenario with $\phi = 0$). It follows that $R \in [0, 1]$ and $\beta \in [-1, 0]$. Finally, in an effort to standardize our result across different instances, we measure payoffs relative to the case without a voting system. Specifically, we focus on the following performance measures:

$$\mathcal{G} := \max_{\lambda_{\mathrm{L}} \leq \lambda \leq \lambda_{\mathrm{H}}} \left\{ \frac{\Pi(\lambda, 0) - G(\lambda, 0)}{G(\lambda, 0)} \right\} \qquad \text{and} \qquad \overline{\mathcal{G}} := \frac{1}{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}} \int_{\lambda_{\mathrm{L}}}^{\lambda_{\mathrm{H}}} \frac{\Pi(\lambda, 0) - G(\lambda, 0)}{G(\lambda, 0)} \, \mathrm{d}\lambda,$$

which are the maximum and average[†] relative gains of using a voting system, respectively.

I) Sensitivity Analysis: In what follows, we investigate the sensitivity of $\mathcal{G}$ and $\overline{\mathcal{G}}$ with respect to the two main components embedded in our voting system: (1) informational (demand learning) and (2) financial (pre-selling revenues). To this end, we define the following quantities:

$$\text{Opportunity for Learning :} \qquad \mathcal{L} := \frac{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}} \quad \text{and}$$

$$\text{Opportunity for Pre-Selling :} \qquad \mathcal{PS} := \frac{\phi\,(\lambda_{\mathrm{H}} + \lambda_{\mathrm{L}})}{2\,r\,R}.$$

Our definition of the *opportunity for learning* $\mathcal{L}$ is motivated by two characteristics of the learning process. On one hand, $\mathcal{L}$ is a normalized –in $[0, 1]$– measure of the "amount" of uncertainty that the seller faces. If $\mathcal{L}$ is close to zero then there is not much to learn while if $\mathcal{L}$ is close to one then there is a significant difference between the intensities of high and low demand products and the opportunity for learning is high. On the other hand, $\mathcal{L}$ also measures the *speed of learning* (see foonote (‡) for mode details). Higher values of $\mathcal{L}$ mean that the voting process is more informative and the seller can learn quicker the value of Λ.

[†] In our numerical experiments we have approximated the integrals by summation over a partition with mesh size $\Delta\lambda = 0.001$.

Let us turn to our second measure, the *opportunity for pre-selling value* $\mathcal{PS}$. To understand its definition, note that our goal is to isolate the financial value of pre-selling. For this, we consider a deterministic model in which the voting process N_t is replaced by a fluid process $N_t^{\mathrm{F}} = \Lambda\, t$ and so the seller's value function in (24) becomes

$$\Pi^{\mathrm{F}}(\lambda, n) = \max\left\{R\ ,\ \max_{t\geq 0}\left\{e^{-r\,t}\big(\beta + \lambda + \phi\,(n + \lambda\,t)\big)\right\}\right\}.$$

Letting t^* be the optimal solution in the inner maximization, one can show that the option of pre selling is valuable (*i.e.*, generates enough cash-flows in the form of pre-orders to support launching the product) if

$$\frac{\phi\,\lambda}{r\,R} \geq e^{r\,t^*}.$$

Our definition of $\mathcal{PS}$ is motivated by the quantity on the left-hand side of this inequality, replacing the value of λ by the mid-value $(\lambda_{\mathrm{H}} + \lambda_{\mathrm{L}})/2$. Hence, we expect that the voting system to be more (financially) beneficial to the seller for large values of $\mathcal{PS}$.

Figure 8 plots the values of $\mathcal{G}$ (left panel) and $\overline{\mathcal{G}}$ (right panel) as function of the *opportunity for learning* $\mathcal{L}$ for two scenarios with $\phi > 0$ and $\phi = 0$ (pure voting). Interestingly, the value of the voting system increases monotonically with $\mathcal{L}$. On one hand, this is expected since a large value of $\mathcal{L}$ means that there is more room for learning. However, on the other hand, a large value of $\mathcal{L}$ implies that the seller faces larger demand risk and the likelihood of discarding the product is higher. This example also reveals that the gains of using a voting system can be quite significant, as much as 15% even in a pure voting system with $\phi = 0$.

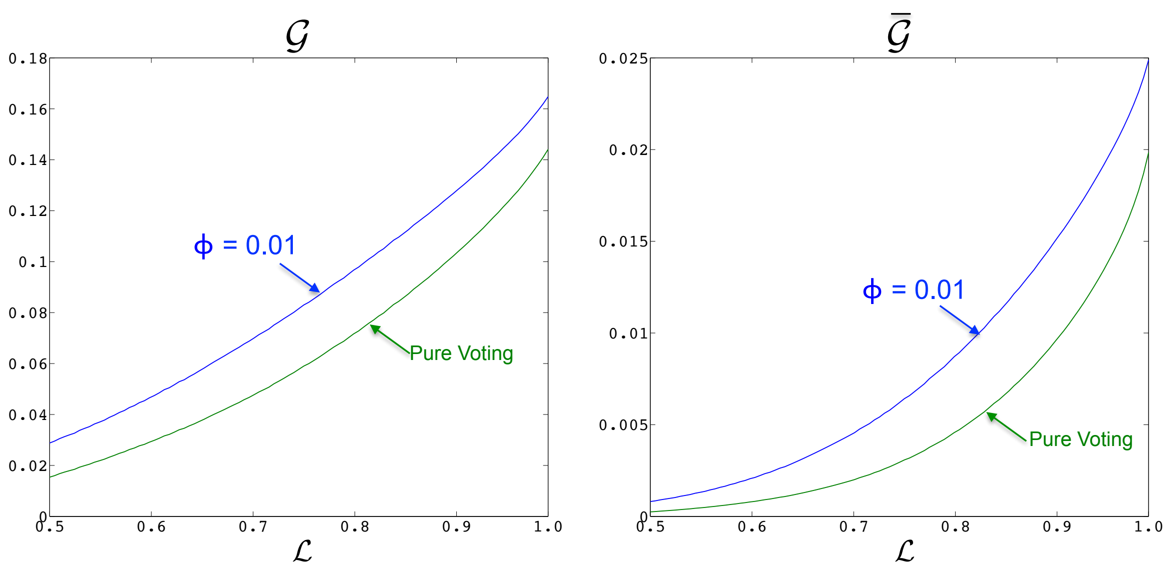

Figure 6: Maximum ($\mathcal{G}$) and average ($\overline{\mathcal{G}}$) relative gains as function of the *opportunity for learning* $\mathcal{L}$. Data: $R = 0.4$, $\beta = -0.5$ and $r = 0.05$.

Turning to the sensitivity of the seller's payoffs with respect to the opportunity to pre-sell, Figure 7 depicts the behavior of maximum and average relative gains as a function of $\mathcal{PS}$. In these computations, we have varied the value of the opportunity cost R (in the range $[0.1, 0.4]$) keeping everything else fixed.

It is interesting to note that the relative gains associated with the voting system grow rapidly with $\mathcal{PS}$ if $\phi > 0$. However, these gains are limited and remain bounded in a pure voting context. In other words, the seller should pay special attention to identify the type of mechanisms that one can use to induce voters to become buyers, for instance, by offering price discounts and other benefits.

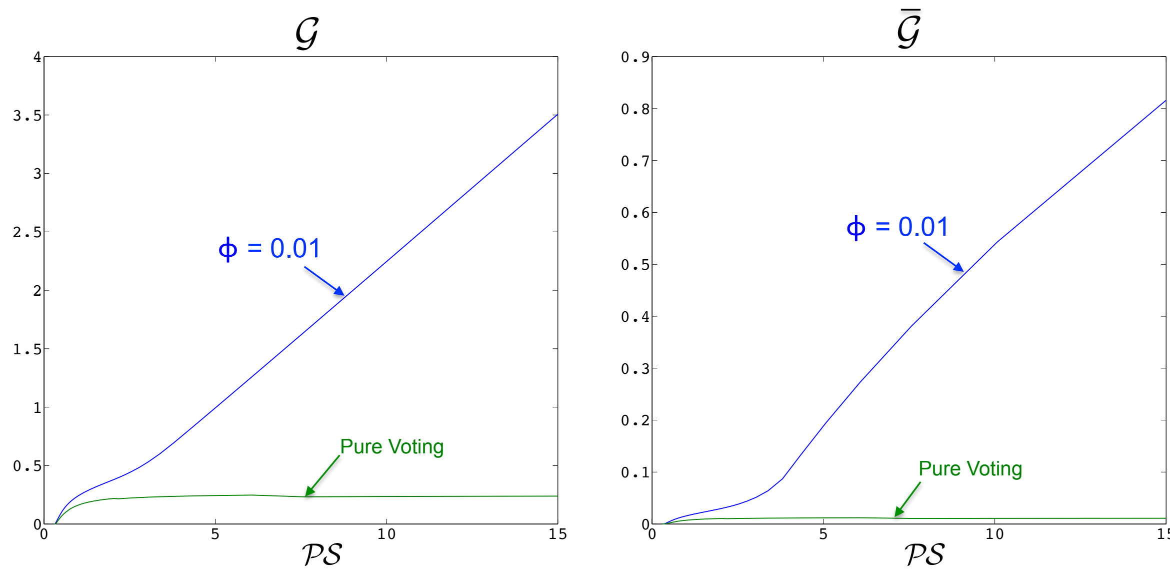

Figure 7: Maximum ($\mathcal{G}$) and average ($\overline{\mathcal{G}}$) relative gains as function of the *opportunity for pre-selll* $\mathcal{PS}$. Data: $\lambda_{\mathrm{L}} = 0.5$, $\beta = -0.5$ and $r = 0.05$.

Benchmarks: Let us now compare the performance of our voting system to two alternative policies in which the seller announces, at time $t = 0$, either (i) a fixed duration $T^* \geq 0$ for the voting phase or (ii) a fixed number of votes $M^* \geq 0$ that will be collected before a launching decision is made. We have chosen these simple (open-loop) policies as benchmarks because they are popular and commonly used in practice as they are easy to communicate and implement.

For a given prior λ on the value of Λ, we let $\Pi^{\mathrm{T}}(\lambda, n)$ and $\Pi^{\mathrm{M}}(\lambda, n)$ be the seller's expected discounted payoffs under the two static policies mentioned. From the representation of λ_t in terms of the likelihood ratio $L(t, N_t)$ in Lemma 1, we can compute the values of $\Pi^{\mathrm{T}}(\lambda, n)$ and $\Pi^{\mathrm{M}}(\lambda, n)$ as follows:

$$\Pi^{\mathrm{T}}(\lambda, n) := \sup_{T \geq 0} \mathbb{E}\left[e^{-r\,T} \max\left\{ R\ ,\ \beta + \frac{\lambda_{\mathrm{L}}\,(\lambda_{\mathrm{H}} - \lambda) + \lambda_{\mathrm{H}}\,(\lambda - \lambda_{\mathrm{L}})\,L(t, N_T)}{\lambda_{\mathrm{H}} - \lambda + (\lambda - \lambda_{\mathrm{L}})\,L(t, N_T)} + \phi\,(n + N_T) \right\}\right]$$

and

$$\Pi^{\mathrm{M}}(\lambda, n) := \sup_{m \in \mathbb{N}} \mathbb{E}\left[e^{-r\,\tau_m} \max\left\{ R\ ,\ \beta + \frac{\lambda_{\mathrm{L}}\,(\lambda_{\mathrm{H}} - \lambda) + \lambda_{\mathrm{H}}\,(\lambda - \lambda_{\mathrm{L}})\,L(\tau_m, m)}{\lambda_{\mathrm{H}} - \lambda + (\lambda - \lambda_{\mathrm{L}})\,L(\tau_m, m)} + \phi\,(n + m)) \right\}\right],$$

where τ_m is the m^{th} epoch of the voting process N_t. Obviously, $\Pi^{\mathrm{T}}(\lambda, n)$ and $\Pi^{\mathrm{M}}(\lambda, n)$ are lower bounds for the value function $\Pi(\lambda, n)$. To assess the performance of our proposed voting system with respect to these approximations we compute the worst-case relative errors:

$$\mathcal{E}^{\mathrm{T}} := \max_{\lambda_{\mathrm{L}} \leq \lambda \leq \lambda_{\mathrm{H}}} \left\{ \frac{\Pi(\lambda, 0) - \Pi^{\mathrm{T}}(\lambda, 0)}{\Pi(\lambda, 0)} \right\} \qquad \text{and} \qquad \mathcal{E}^{\mathrm{M}} := \max_{\lambda_{\mathrm{L}} \leq \lambda \leq \lambda_{\mathrm{H}}} \left\{ \frac{\Pi(\lambda, 0) - \Pi^{\mathrm{M}}(\lambda, 0)}{\Pi(\lambda, 0)} \right\},$$

as well as the average relative errors:

$$\bar{\mathcal{E}}^{\mathrm{T}} := \frac{1}{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}} \int_{\lambda_{\mathrm{L}}}^{\lambda_{\mathrm{H}}} \frac{\Pi(\lambda,0) - \Pi^{\mathrm{T}}(\lambda,0)}{\Pi(\lambda,0)}\,\mathrm{d}\lambda \qquad \text{and} \qquad \bar{\mathcal{E}}^{\mathrm{M}} := \frac{1}{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}} \int_{\lambda_{\mathrm{L}}}^{\lambda_{\mathrm{H}}} \frac{\Pi(\lambda,0) - \Pi^{\mathrm{M}}(\lambda,0)}{\Pi(\lambda,0)}\,\mathrm{d}\lambda.$$

Fixed Duration

Worst-Case Error: $\mathcal{E}^{\mathrm{T}}$

	λ_{L}			
ϕ	0.2	0.4	0.6	0.8
0.005	3.2%	1.6%	0.5%	0.0%
0.01	3.3%	1.8%	0.7%	0.0%
0.05	8.0%	7.1%	5.1%	1.9%
0.1	13.0%	10.5%	3.5%	2.4%

Average Error: $\bar{\mathcal{E}}^{\mathrm{T}}$

	λ_{L}			
ϕ	0.2	0.4	0.6	0.8
0.005	0.4%	0.1%	0.0%	0.0%
0.01	0.6%	0.1%	0.0%	0.0%
0.05	2.0%	1.4%	1.0%	1.4%
0.1	3.3%	2.3%	2.1%	1.8%

Fixed Number of Votes

Worst-Case Error: $\mathcal{E}^{\mathrm{M}}$

	λ_{L}			
ϕ	0.2	0.4	0.6	0.8
0.005	4.5%	2.4%	1.0%	0.0%
0.01	5.5%	2.8%	1.5%	0.0%
0.05	17.9%	12.1%	4.1%	0.0%
0.1	26.9%	9.2%	0.2%	0.0%

Average Error: $\bar{\mathcal{E}}^{\mathrm{M}}$

	λ_{L}			
ϕ	0.2	0.4	0.6	0.8
0.005	0.5%	0.1%	0.0%	0.0%
0.01	0.7%	0.2%	0.0%	0.0%
0.05	3.4%	1.7%	0.1%	0.0%
0.1	5.0%	1.4%	0.0%	0.0%

Table 3: *Percentage relative errors $\mathcal{E}^{\mathrm{T}}$ and $\mathcal{E}^{\mathrm{M}}$. Data: $R = 0.5$, $\beta = -1$, $\lambda_{\mathrm{H}} = 1$ and $r = 5\%$.*

As we can see from the values of the worst-case errors $\mathcal{E}^{\mathrm{T}}$ and $\mathcal{E}^{\mathrm{M}}$ in Table 3, using these static approximations can lead to substantial suboptimality, in some cases as large as to 27%. (As expected, the average errors $\bar{\mathcal{E}}^{\mathrm{T}}$ and $\bar{\mathcal{E}}^{\mathrm{M}}$ are smaller but still significant.) If we compare the two approximations, we also note that using a fixed duration for the voting process leads to better performance when λ_{L} is small but as λ_{L} increases setting a fixed number of votes is a better strategy. In other words, under high demand uncertainty a voting campaign with fixed duration is preferred to one with a fixed target of votes, and vice-versa. The fact that a fixed number of votes is better when demand uncertainty is low is somehow expected since choosing a fixed number of votes is an optimal strategy under full information (see Section 3). Another important observation that emerges from these numerics is that the value of using an optimal stopping rule to determine the duration of the voting phase (instead of a fixed static policy) increases in ϕ, that is, when the opportunities for pre-selling are higher. It follows that when the conversion of voters into buyers is low a static voting campaign could lead to good performance. However, as the conversion rate increases the seller has more incentives to use a dynamic stopping rule.

6 Extensions

6.1 Pricing Tactic

Up to this point, we have focused on determining the optimal duration of the voting phase. However, there is another factor that can have a significant impact on the overall success of a crowdvoting system and that is pricing. In this section, we extend our model to investigate this issue. As we have been discussing throughout, the performance of a crowdvoting system depends on its capacity to (a) improve the seller's demand forecast and (b) generate pre-orders that will translate into sales if and when the product is launched. These two objectives, however, are not necessarily aligned when it comes to determining a pricing policy. On one hand, in an effort to speed up the voting and demand learning processes, the seller would like to offer voters deep price discounts to encourage them to vote. On the other hand, price discounts will reduce the revenues that these pre-orders will generate. Hence, an optimal pricing tactic should try to appropriately balance these two conflicting objectives.

To put the previous discussion on a concrete mathematical footing, let us consider the following variation of our base model. As before, we divide the planning horizon into two phases –voting and selling phases– and we assume that consumers arrive to the system according to a Poisson process with rate Λ during the entire planning horizon. In this case, however, we also assume that arriving consumers have a maximum willingness-to-pay for the product and that their voting and/or purchasing behavior depends on whether or not the price they see exceeds this threshold. Since it is reasonable to assume that the reservation price of individuals that arrive during the voting phase is different than the one of those arriving during the selling phase, we let $F_{\mathrm{V}}(p)$ and $F_{\mathrm{R}}(p)$ be the cumulative probability distribution of the willingness-to-pay of voters and buyers, respectively. On the seller side, we assume that he sets a price p_{V} during the crowdvoting phase (this will be the price that voters will pay at the time the product is launched) and a price p_{R} during the regular selling phase if this latter phase eventually materializes. As a result, the voting process is driven by a Poisson process with rate $\Lambda\,\bar{F}_{\mathrm{V}}(p_{\mathrm{V}})$, where $\bar{F}_{\mathrm{V}}(p_{\mathrm{V}}) := 1 - F_{\mathrm{V}}(p_{\mathrm{V}})$ is the probability that an arriving voter clicks at the product at price p_{V}. Similarly, the demand for the product during the regular selling phase follows a Poisson process with intensity $\Lambda\,\bar{F}_{\mathrm{R}}(p_{\mathrm{R}})$, where $\bar{F}_{\mathrm{R}}(p_{\mathrm{R}}) := 1 - F_{\mathrm{R}}(p_{\mathrm{R}})$. From a practical standpoint, one would expect that $p_{\mathrm{V}} \leq p_{\mathrm{R}}$, that is, that the seller is willing to offer voters a discount to incentivize them to vote as well as to compensate them for the risk and delay that they face in getting the product. We will not impose this constraint directly but rather we will investigate how the optimal prices p_{V} and p_{R} depend on the values of the different parameters of the model, and in particular on F_{V} and F_{R}.

In this modified setting, the voting process N_t has intensity $\Lambda\,\bar{F}_{\mathrm{V}}(p_{\mathrm{V}})$. Hence, the belief process $\lambda_t = \mathbb{E}[\Lambda|\mathcal{F}_t]$ satisfies

$$\mathrm{d}\lambda_t = \eta(\lambda_{t-})\left[\mathrm{d}N_t - \bar{F}_{\mathrm{V}}(p_{\mathrm{V}})\,\lambda_t\,\mathrm{d}t\right] \qquad \text{where} \quad \eta(\lambda) = \frac{(\lambda_{\mathrm{H}} - \lambda)\,(\lambda - \lambda_{\mathrm{L}})}{\lambda}.$$

Based on our discussion in Section 2.1, the seller's payoff at time τ conditional on the value of $\lambda_\tau = \lambda$ and the number of voters $N_\tau = n$ is given by

$$G(\lambda, n) = \max\{R\,,\, G_{\mathrm{R}}(\lambda) + G_{\mathrm{V}}(n) - K\} = \max\left\{R\,,\; \frac{p_{\mathrm{R}}\,\bar{F}_{\mathrm{R}}(p_{\mathrm{R}})}{r}\,\lambda + \theta\,p_{\mathrm{V}}\,n - K\right\}.$$

Recall that $K > 0$ is the fixed cost of launching the product and $\theta \in [0, 1]$ is the probability that a voter will return to buy the product if it is eventually put on the marketplace. As a result, the seller's optimization problem is now given by

$$\begin{aligned}\Pi(\lambda) &= \sup_{p_{\rm V}, p_{\rm R}, \tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\,\max\left\{R\,,\ \frac{p_{\rm R}\,\bar F_{\rm R}(p_{\rm R})}{r}\,\lambda_\tau + \theta\,p_{\rm V}\,N_\tau - K\right\}\right]\\ \text{subject to} \qquad & \mathrm{d}\lambda_t = \eta(\lambda_{t-})\left[\mathrm{d}N_t - \bar F_{\rm V}(p_{\rm V})\,\lambda_t\,\mathrm{d}t\right] \quad \text{and} \quad \lambda_0 = \lambda.\end{aligned}$$

Since the dynamics of λ_t are independent of $p_{\rm R}$, it follows that the optimal regular price is given by $p_{\rm R}^* := \arg\max\{p\,\bar F_{\rm R}(p)\}$. This result should be intuitively clear, since once the product is introduced on the market, the seller focuses exclusively on maximizing the instantaneous revenue rate. In other words, demand learning provides no additional value at this point. Let us set $\delta^* := p_{\rm R}^*\,\bar F_{\rm R}(p_{\rm R}^*)/r$.

On the other hand, the dependence of $G(\Lambda, n)$ on $p_{\rm V}$ is more complex since this price affects both the revenues at launching as well as the dynamics of λ_t and N_t. In order to isolate the effect of $p_{\rm V}$ on the seller's expected payoff, let us introduce a time change to remove the dependence of the state variable on $p_{\rm V}$. Specifically, let us define $\hat\lambda(t) := \lambda(t/\bar F_{\rm V}(p_{\rm V}))$ and $\hat N(t) := N(t/\bar F_{\rm V}(p_{\rm V}))$. (Note that $\hat N_t$ is a Poisson process with rate Λ.) Also, by appropriately scaling R, K, θ and $\Pi(\lambda)$ by δ^*, we can assume that $\delta^* = 1$. With these changes, the seller's optimization becomes

$$\begin{aligned}\Pi(\lambda) &= \sup_{p_{\rm V}, \tau\in\mathcal{T}} \mathbb{E}\left[\exp\left(-\frac{r}{\bar F_{\rm V}(p_{\rm V})}\,\tau\right)\,\max\left\{R\,, \hat\lambda_\tau + \theta\,p_{\rm V}\,\hat N_\tau - K\right\}\right] \qquad (28)\\ \text{subject to} \qquad & \mathrm{d}\hat\lambda_t = \eta(\hat\lambda_{t-})\left[\mathrm{d}\hat N_t - \,\hat\lambda_t\,\mathrm{d}t\right] \quad \text{and} \quad \hat\lambda_0 = \lambda.\end{aligned}$$

We note that this formulation is equivalent to our base model with $\beta = -K$, $\phi = \theta\,p_{\rm V}$ and a modified discount fact $r/\bar F_{\rm V}(p_{\rm V})$. So, for a fixed value of $p_{\rm V}$, with $\bar F_{\rm V}(p_{\rm V}) > 0$, we can use the solution techniques that we developed in the previous sections to derive an optimal stopping rule. The optimization in (28) provides an alternative perspective on the trade-off that the seller faces when choosing $p_{\rm V}$. On one hand, a large $p_{\rm V}$ will increase the revenues at introduction $\theta\,p_{\rm V}\,\hat N_\tau$. On the other hand, increasing $p_{\rm V}$ will decrease $\bar F_{\rm V}(p_{\rm V})$, effectively increasing the (modified) discount factor $r/\bar F_{\rm V}(p_{\rm V})$. The following result follows directly from (28).

Proposition 10 *In a* pure voting *system, that is when $\theta = 0$, it is optimal to set $p_V = 0$.*

In other words, when the voting system is a pure demand learning mechanism and has no direct impact on the seller's cash-flows, it is in the seller best interest to stimulate the pace of voting as much as possible by offering voters deep discounts, which in our model corresponds to setting $p_{\rm V} = 0$.

Figure 8 plots the optimal value of $p_{\rm V}$ as a function of the conversion rate θ. As we can see, the result in Proposition 10 extends to a range of values θ. So even if a voter can generate a positive revenue at introduction, it is better for the seller to prioritize learning speed by giving voters a deep discount. On the other hand, as θ grows large, the optimal price $p_{\rm V}$ also increases and rapidly converges to the myopic value $p_{\rm V}^* = \arg\max\{p\,\bar F_{\rm V}(p)\}$

6.2 Impatient Voters

Another aspect of our model that deserves further discussion is voters' purchasing behavior. Up to this point, we have assumed that there is a fixed probability θ that a voter returns to buy the product at

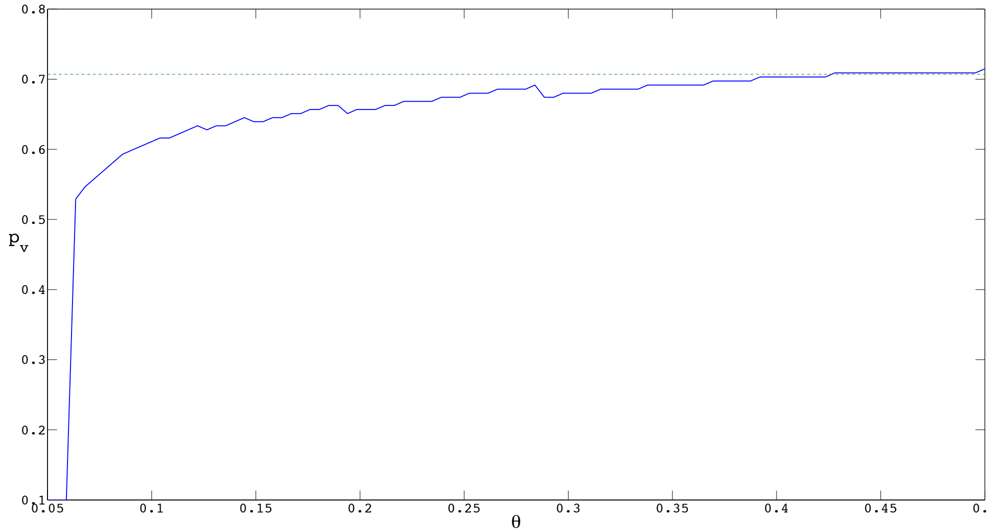

Figure 8: Optimal crowdvoting prive p_V as a function of the conversation θ of voters into buyers. The dashed line is at the myopic price level $p^*_V = \arg\max\{p\,\bar{F}_V(p)\}$. Data: $R = 0.4$, $\beta = -0.5$ and $r = 0.05$, $\lambda_L = 0.2$, $\lambda_H = 1$ and $\lambda = 0.5$.

its introduction time. Our motivation behind this formulation is the idea that a fraction $1-\theta$ of the voters are *impulse buyers* that would have bought the product if it would have been available at the time they voted but they will not return to buy it at a later time. In addition, we have also assumed that this probability θ is independent of how long *non-impulse* voters have to wait. It seems reasonable to expect, however, that these *non-impulse buyers* have a maximum tolerance on how long they are willing to wait. To capture this *impatience*, we assume that each non-impulse customer is willing to wait for an exponential amount of time with rate α. If the product is not available after this time, the customer will no longer be interested (*i.e.*, her preferences have changed or she has bought a substitute product somewhere else).

Let us denote by S_τ the expected number of sales that would be realized if the product is launched at time τ.

Lemma 3 (Sales from Voters) *If the seller introduces the product at time $\tau \in \mathcal{T}$, then the expected number of sales at introduction, S_τ satisfies the following stochastic differential equation*

$$\mathrm{d}S_t = -\alpha\, S_t\,\mathrm{d}t + \theta\,\mathrm{d}N_t \quad S_0 = 0, \quad 0 \le t \le \tau.$$

Using this lemma, we can rewrite the seller's optimization problem in this case with impatient voters as the following optimal stopping problem:

$$\begin{aligned} & \Pi(\lambda) = \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda_\tau, S_\tau)\right] \\ \text{subject to} \quad & \mathrm{d}\lambda_t = \eta(\lambda_{t-})\left[\mathrm{d}N_t - \lambda_t\,\mathrm{d}t\right], \quad \lambda_0 = \lambda \\ & \mathrm{d}S_t = -\alpha\, S_t\,\mathrm{d}t + \theta\,\mathrm{d}N_t, \qquad S_0 = 0, \end{aligned}$$

where $G(\lambda, S) = \max\{R, \beta + \lambda + p_V\, S\}$ with p_V the price that a voter will pay to purchase the product.

As before, we can tackle the solution to this problem using dynamic programming methods. However, a key difference between the optimization problem above and the one discussed in Section 5 is that the corresponding value function $\Pi(\lambda, S)$ depends on the state variable S (*i.e.*, expected number of sales at introduction) as opposed to the number of votes. Because the process S_t is no longer discrete and monotonic in time, we can no longer use backward induction as we did in the previous section to solve the dynamic program. A detailed analysis of this model is beyond the scope of this paper. Instead, we provide a simple approximation that builds upon our previous results.

To this end, let us denote by X_t the number of non-impulse voters that have arrived in $[0, t]$ and who are still willing to buy the product (see the proof of Lemma 3 for a precise definition). It follows that $S_t = \mathbb{E}[X_t|\mathcal{F}_t]$. Note that from the seller's perspective, X_t evolves as a hidden birth-death process. Actually, the arrivals of new voters (births) are partially visible since the seller can see the voting process but cannot differentiate between impulse and non-impulse buyers. On the other hand, the seller cannot see deaths which correspond to the time epochs at which voters "renege" because their patient limit has been exceeded. Births occur at a rate $\theta\,\Lambda$ while deaths occur at a rate $\alpha\, X_t$. From standard birth-death theory it follows that, conditional on $X_0 = 0$, X_t has a Poisson distribution and

$$\mathbb{E}[X_t|X_0 = 0, \Lambda] = \frac{\theta\,\Lambda}{\alpha}\,(1 - \exp(-\alpha\, t)).$$

Based on the above, we could approximate S_τ by the constant $\bar{S} := \theta\,\Lambda/\alpha$, *i.e.* by the long-term average value of X_t. We expect this approximation to be accurate in those situation in which the voting process is open for a relatively long period of time relative to the rate at which voters renege. Under this approximation, the seller's payoff becomes a function of Λ exclusively and is given by

$$G(\Lambda) = \max\Big\{R\ ,\ \beta + \Lambda + \frac{p_{\mathrm{V}}\,\theta\,\Lambda}{\alpha}\Big\}.$$

As before, if we denote by $\phi = p_{\mathrm{V}}\,\theta$ we can use the change of variables $\bar{R} := \alpha\, R/(\alpha+\phi)$ and $\bar{\beta} := \alpha\,\beta/(\alpha+\phi)$, we can rewrite this payoff as

$$G(R) = \frac{\alpha+\phi}{\alpha}\,\bar{G}(\Lambda), \qquad \text{where } \bar{G}(\lambda) := \max\{\bar{R}\ ,\ \bar{\beta}+\Lambda\}.$$

As a result, and modulo the factor $(\alpha+\phi)/\alpha$, the seller's optimization problem is given

$$\bar{\Pi}(\lambda) := \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\max\{\bar{R}\ ,\ \bar{\beta}+\Lambda\}\right]$$
$$\text{subject to} \qquad \mathrm{d}\lambda_t = \eta(\lambda_{t-})\,\Big[\mathrm{d}N_t - \lambda_t\,\mathrm{d}t\Big], \quad \lambda_0 = \lambda.$$

The resulting optimization problem is exactly the same as the one we solved under pure voting but replacing β and R by $\bar{\beta}$ and $\bar{R}$ and so we can directly export all the results in Section 4 to solve this problem explicitly.

7 Concluding Remarks

In this paper we analyze the value of crowdsourcing in the context of new product introduction. More specifically, we consider a seller contemplating the possibility of launching a new product into the market.

The seller is uncertain on whether the product will be successful or not- if commercialized - and rely on a voting platform (crowdvoting) to make a more informed decision. The seller makes available the design online for potential buyers to vote on. As the voting progresses, the seller has to decide when to end the voting phase and, as a result of the information gathered, whether to discard the product or to commercialize it. In the case it is commercialized, the seller incurs a fixed cost and collects a stream of revenues from sales which depend on the market size. If it is discarded, the seller collects a payoff R that captures the product's opportunity cost. This voting phase can also be used to generate pre-orders and hence some positive cash flow at introduction. The voting system plays then two main functions. First and foremost, it represents a market learning mechanism, whereby the seller would continuously update his estimate of the market size as votes get casted. Secondly, it represents a funding mechanism whereby, through the pre-orders generated, the voting phase allows the seller to reduce, at introduction, the impact of the fixed costs at the expense of a delay in launching the product and thus in the materialization of the total profits.

In order to dissect the impact of each of these two functions, we study each separately. In Section 3, we assume that the market demand rate is known, and hence the voting phase is only used to better manage the product's cash flow. In this full information case, the optimal policy is, from the start, to either discard the product or to announce a '*funding target*' which is given by the number of votes that need to be collected before the product is launched. This full information policy also represents a myopic approximation to the general case that we use to generate a lower and an upper bounds that are particularly tight when the uncertainty is low or the opportunity cost, R is high. Next, in Section 4, we focus on the pure voting case where visitors do not have the opportunity to pre-order the item. In this case, the crowdvoting platform is primarily used to learn the market size, which is done through Bayesian updating. In the pure voting, we obtain an explicit expression of the payoff in closed-form using QVI machinery. The optimal policy, characterized by the following stopping time

$$\tau = \inf\{t \geq 0 : \Pi(\lambda_t) = G(\lambda_t)\},$$

is then fully defined and easy to implement. The closed form expression brings also additional advantages; in particular, an exact measure of the value of learning, $\Pi(\lambda) - G(\lambda)$. Maximum learning is shown to be obtained when the opportunity cost from running the voting phase is small or when the speed of learning measured by the difference $\lambda_{\text{H}} - \lambda_{\text{L}}$ is large. We also determine simple conditions on the problem parameters' to assess whether a crowdvoting phase is worthwhile or not. Furthermore, we conduct in this pure voting setting, an asymptotic analysis by scaling some of the parameters, in particular λ_{H} and λ_{L} by n, while keeping their difference in $\sqrt{n}$ and the magnitude of the jumps η, bounded. We prove in this case that the belief process $\mathbb{P}(\Lambda = \lambda_{\text{L}}|\mathcal{F}_t)$ converges to a diffusion process. As a result, we convert the scaled optimization problem at the limit, into a simple second order ODE with smooth border conditions. The solution to the ODE reveals a simple formulation of the payoff function and an explicit characterization of the distribution of the optimal stopping time (i.e. the duration of the crowdvoting session), in particular, of its expected value.

In Section 5, we investigate the general case in which the voting platform is used to presell as well as to conduct a Bayesian learning on the market size Λ. The seller needs to track in this case both the updated value of λ and the number of votes over time. Given the complexity of the problem in this general case, we cannot obtain closed form solutions. Nevertheless, we were able to fully characterize the optimal payoff as a solution of a recursive relationship with given boundary conditions that can be solved rather easily numerically. We were also able to fully characterize the continuation region, which is particularly helpful in implementing the optimal policy in practice. Finally, we conducted intensive numerical experiments and our conclusions were of two folds. First, we identified the parameters to which the payoff is sensitive to and showed that the value of the voting phase increase with respect to these parameters. Despite the many parameters defining the model, we reduced our analysis to how sensitive the payoff is with respect to the opportunity of learning and to that of pre-selling, each embodied by one well identified quantity. As expected, the value of a crowdvoting platform increases (and becomes quite sizeable) as the uncertainty on the market size and the opportunity of pre-selling increase. Secondly, we showed numerically that offering a voting phase for which the duration is set dynamically (as a result of a sophisticated stopping rule that accounts for both the market learning level and the funds generated through pre-sales) generates a clear edge compared to those simple (myopic and open loop) policies used in practice; whether those relying on a fixed deterministic duration or set based on a funding target (number of votes). Again this edge is sizeable when both the opportunities of learning and preselling are high. From this analysis we also observe that when the opportunity of learning is small while the opportunity of preselling is high, a preset number of votes performs well. On the other hand, when the opportunity of learning is large (i.e. the market uncertainty is high) then a preset number of votes is too risky and the seller would be better off fixing the duration of the voting phase. The performance of the latter policy deteriorates as the opportunity of preselling increases.

Finally in Section 6, we discuss two extensions of our model to take into account pricing strategies and voters patience. In the first one, we introduce the price of the voting phase as a static decision variable that needs to be set by the seller on top of the duration. For that we show that the earlier model can easily accommodate this important addition and can leverage on the previous solution techniques. We numerically explore the impact of the conversion rate (a proxy for the opportunity of preselling) on the optimal price and observe that the latter is either set to be zero when the conversion rate is very small in greater support of the informational component; or converging quickly to the value that maximizes the preselling rate and hence in support of the financial component. This conclusion argues (surprisingly) for a simple pricing rule despite the complexity of the general model. Next, we consider the case where the likelihood for every vote to be converted into sales decreases with time and hence, the later the product is introduced the more revenues from early voters are lost. In many contexts, this is a more natural setting, whereby the accumulated funds follows a stochastic process that is not necessarily increasing with time. We again formulate the problem in the general case and suggest to approximate the sales process by its long-term average average. This approximation reduces the problem to a formulation similar to the pure voting setting which can be solved in closed form. Finally, this section also shows that the model we are suggesting in this paper to analyze crowdvoting platforms is flexible enough and can be adapted to include other relevant factors and additional layers of complexities while keeping a good level of tractability.

This work also identifies many interesting avenues to tackle in future research. One major avenue is that of incorporating strategic customers that could use voting as a way to increase their overall surplus. In such game theoretic setting, one important question is whether the seller is better off announcing the number of votes accumulated at this point or not. Indeed, strategic customers might try to “manipulate” the seller by not voting if the number of votes is already high. But, at the same time a customer uncertain about her true valuation can see her expected valuation increasing with a high number of votes. Another avenue is that of multi products, where customers visit the platform and are offered multiple choices (e.g. T-shirts designs). Their decision to vote or not would not only take their own valuation into account but also the probability that their favorite choice would be commercialized or not. Other avenues related to dynamic pricing and inventory management could also be worth exploring as well.

References

Araman, V., R. Caldentey. 2009. Dynamic pricing for nonperishable products with demand learning. *Oper. Res.* **57**(5) 1169–1188. 5

Aviv, Y., A. Pazgal. 2002. Pricing of short life-cycle products through active learning. Working Paper, Olin School of Business, Washington University, St. Louis. 5

Belleflamme, P., T. Lambert, A. Schwienbacher. 2014. Crowdfunding: Tapping the right crowd. *Journal of Business Venturing* **29**(5) 585–609. 4

Billingsley, P. 1968. *Convergence of probability measures*. Wiley series in probability and mathematical statistics, Wiley, New York. 43

Boyaci, T., Ö. Özer. 2010. Information acquisition for capacity planning via pricing and advance selling: When to stop and act? *Oper. Res.* **58**(5) 1328–1349. 4, 5

Brémaud, P. 1980. *Point Processes and Queues, Martingale Dynamics*. Springer-Verlag, New York, NY. 9

Crosetto, P., T. Regner. 2014. Crowdfunding: Determinants of success and funding dynamics. *Jena Economic Research Papers* **8**. 4

Dana, J. 1998. Advanced purchase discounts and price discrimination in competitive markets. *J. Political Econom.* **106**(2) 395–422. 4

Farias, V.F., B. Van Roy. 2010. Dynamic pricing with a prior on market response. *Operations Research* **58** 16–29. 5

Gale, I. L., T.J. Holmes. 1992. The efficiency of advanced-purchase discounts in the presence of aggregate demand uncertainty. *International Journal of Industrial Organization* **10**(4) 413–437. 4

Gale, I. L., T.J. Holmes. 1993. Advance-purchase discounts and monopoly allocation of capacity. *American Economoc Review* **83**(1) 135–146. 4

Gallego, G., Ö. Özer. 2001. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **47**(10) 1344–1360. 4

Huang, Y., P.V. Singh, Y. Xu. 2014. Crowdsourcing new product ideas under consumer learning. *Management Sci.* **60**(9) 2138–2159. 3

Jeff Howe. 2006. Crowdsourcing: A definition. URL http://crowdsourcing.typepad.com/cs/2006/06/crowdsourcing_a.html. 2, 3

Karatzas, I., S.E. Shreve. 1991. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York, NY. 23

Krager, D.R., S. Oh, D. Shah. 2014. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research* **62**(1) 1–24. 3

Kurtz, T. G. 1978. Strong approximation theorems for density dependent markov chains. *Stochastic Processes and their Applications* **6** 223–240. 20

Kurtz, T.G., P. Protter. 1991. Weak limit theorems for stochastic integrals and stochastic differential equations. *The Annals of Probability* **19**(3) 1035–1070. 44

Marinesi, S., K. Girotra. 2013. Information acquisition through customer voting systems. Tech. rep., Working Paper, INSEAD. 3

Mollick, E. 2014. The dynamics of crowdfunding: An exporatory study. *Journal of Business Venturing* **29**(1) 1–16. 4

Peskir, G., A.N. Shiryaev. 2000. Sequential testing problems for poisson processes. *Ann. Stat.* **28** 837–859. 5, 14

Prasad, A., K.E Stecke, X. Zao. 2011. Advance selling by a newsvendor retailer **20**(1) 1059–1478. 4

Protter, P.E. 2004. *Stochastic Integration and Differential Equations*. Springer, Germany. 38, 40, 44

Raman, A., M. Fisher. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Oper. Res.* **44**(4) 87–99. 4

Surowiecki, J. 2004. *The Wisdom of Crowds*. Random House, Inc., New York, NY. 3

Tang, C.S., K. Rajaram, A. Alptekinoglu, J. Ou. 1995. The benefits of advanced booking discount programs: Model and analysis. *Management Sci.* **50**(4) 465–478. 4

Terwiesch, C., Y. Xu. 2008. Innovation contests, open innovation and multi-agent problem solving. *Management Sci.* **54**(9) 1529–1543. 3

Whitt, W. 1980. Some useful functions for functional limit theorems. *Mathematics of Operations Research* **5**(1) 67–85. 43

Appendix A: Proofs

Proof of Lemma 1: Let us define the auxiliary state variable $q_t := \mathbb{P}(\Lambda = \lambda_{\mathrm{L}} | \mathcal{F}_t)$ with initial condition $q_0 = q$. Note that $\lambda_t = \lambda_{\mathrm{L}}\, q_t + \lambda_{\mathrm{H}}\,(1 - q_t)$ and $q_t = (\lambda_{\mathrm{H}} - \lambda_t)/(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})$. Now, using the Poisson distribution of cumulative votes in $[0, t]$ and Bayes' rule we get that

$$\begin{aligned} q_t &= \frac{q\,(\lambda_{\mathrm{L}}\, t)^{N_t} \exp(-\lambda_{\mathrm{L}}\, t)/N_t!}{q\,(\lambda_{\mathrm{L}}\, t)^{N_t} \exp(-\lambda_{\mathrm{L}}\, t)/N_t! + (1-q)\,(\lambda_{\mathrm{H}}\, t)^{N_t} \exp(-\lambda_{\mathrm{H}}\, t)/N_t!} \\ &= \frac{q}{q + (1-q)(\lambda_{\mathrm{H}}/\lambda_{\mathrm{L}})^{N_t} \exp(-(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})\, t)} = \frac{q}{q + (1-q)\, L(t, N_t)}, \end{aligned} \tag{29}$$

where the first equality follows from the Markov property of the voting process. This proves the second part of the lemma.

Let us turn now to the first part. Using (29) we can derive the dynamics of the process $(q_t : t \geq 0)$. To this end, we write $q_t = f(Y_t)$, where $Y_t := \ln(\lambda_{\mathrm{H}}/\lambda_{\mathrm{L}}) N_t - (\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})\, t$ is an $\mathcal{F}_t$-semimartingale and f is a twice differentiable and bounded function given by $f(n) := \frac{q}{q+(1-q)\exp(n)}$. From Itô's lemma (*e.g.*, Protter, 2004) and the fact that Y_t is a finite variation process (which follows from the fact that $N(t)$ is a pure-jump process), we get

$$\mathrm{d}q_t = f'(Y_{t-})\, \mathrm{d}Y_t + f(Y_t) - f(Y_{t-}) - f'(Y_{t-})\, \Delta Y_t.$$

Taking advantage of the pure-jump nature of N_t, we have $\mathrm{d}N_t = \Delta N_t$, $\mathrm{d}Y_t = \Delta Y_t - (\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})\, \mathrm{d}t$, and $f(Y_t) - f(Y_{t-}) = [f(Y_{t-} + \ln(\lambda_{\mathrm{L}}/\lambda_{\mathrm{H}})) - f(Y_{t-})]\mathrm{d}N_t$, so that

$$\begin{aligned} \mathrm{d}q_t &= -f'(Y_{t-})(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})\, \mathrm{d}t + [f(Y_{t-} + \ln(\lambda_{\mathrm{L}}/\lambda_{\mathrm{H}})) - f(Y_{t-})]\, \mathrm{d}N_t \\ &= (\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}) \frac{q(1-q)\exp(Y_{t-})}{(q + (1-q)\exp(Y_t-))^2}\, \mathrm{d}t + \left[\frac{q}{q + (1-q)\, \exp(Y_{t-}) \frac{\lambda_{\mathrm{H}}}{\lambda_{\mathrm{L}}}} - \frac{q}{q + (1-q)\, \exp(Y_{t-})} \right] \mathrm{d}N_t \\ &= -\eta(q_{t-}) \Big[\mathrm{d}N_t - (\lambda_{\mathrm{L}} q_{t-} + \lambda_{\mathrm{H}}(1 - q_{t-}))\mathrm{d}t \Big], \quad \text{where } \eta(q_t) := \frac{q_t(1-q_t)(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})}{\lambda_{\mathrm{L}} q_t + \lambda_{\mathrm{H}}(1-q_t)}. \end{aligned} \tag{30}$$

Finally, since $\lambda_t = \lambda_{\mathrm{L}}\, q_t + \lambda_{\mathrm{H}}\,(1 - q_t)$, it follows that $\mathrm{d}\lambda_t = -(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})\, \mathrm{d}q_t$. Plugging these identities in (30), the result follows. $\square$

Proof of Proposition 1: First of all, to prove the monotonicity of $\Pi(\lambda, n)$ on λ and n it suffices to note that $G(\lambda, n)$ is increasing in both λ and n and that the belief process λ_t is pathwise increasing in the initial value λ_0 (see Lemma 1).

To prove the convexity of $\Pi(\lambda, n)$ on λ, let us denote by $\mathbb{E}_{\mathrm{L}}$ and $\mathbb{E}_{\mathrm{H}}$ the expectation operators conditional on $\Lambda = \lambda_{\mathrm{L}}$ and $\Lambda = \lambda_{\mathrm{H}}$, respectively. It follows that

$$\begin{aligned} \Pi(\lambda, n) &= \sup_{\tau \in \mathcal{T}} \mathbb{E}[e^{-r\,\tau}\, G(\lambda_\tau, n + N_\tau)] \\ &= \sup_{\tau \in \mathcal{T}} \left\{ \mathbb{E}_{\mathrm{L}}[e^{-r\,\tau}\, G(\lambda_\tau, n + N_\tau)]\, \mathbb{P}(\Lambda = \lambda_{\mathrm{L}}) + \mathbb{E}_{\mathrm{H}}[e^{-r\,\tau}\, G(\lambda_\tau, n + N_\tau)]\, \mathbb{P}(\Lambda = \lambda_{\mathrm{H}}) \right\} \\ &= \sup_{\tau \in \mathcal{T}} \left\{ \mathbb{E}_{\mathrm{L}}[e^{-r\,\tau}\, G(\lambda_\tau, n + N_\tau)] \left(\frac{\lambda_{\mathrm{H}} - \lambda}{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}} \right) + \mathbb{E}_{\mathrm{H}}[e^{-r\,\tau}\, G(\lambda_\tau, n + N_\tau)] \left(\frac{\lambda - \lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}} \right) \right\}. \end{aligned}$$

Consider a pair of beliefs $\lambda_1, \lambda_2 \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$ and set $\lambda = \alpha\,\lambda_1 + (1-\alpha)\,\lambda_2$ for some $\alpha \in [0,1]$. Then, convexity of $\Pi(\lambda)$ follows from

$$\begin{aligned}
\Pi(\lambda, n) &= \sup_{\tau\in\mathcal{T}} \left\{ \mathbb{E}_{\mathrm{L}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_{\mathrm{H}}-\lambda}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) + \mathbb{E}_{\mathrm{H}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) \right\} \\
&= \sup_{\tau\in\mathcal{T}} \left\{ \alpha \left[\mathbb{E}_{\mathrm{L}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_{\mathrm{H}}-\lambda_1}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) + \mathbb{E}_{\mathrm{H}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_1-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) \right] + \right. \\
&\qquad \left. (1-\alpha) \left[\mathbb{E}_{\mathrm{L}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_{\mathrm{H}}-\lambda_2}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) + \mathbb{E}_{\mathrm{H}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_2-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) \right] \right\} \\
&\le \alpha \sup_{\tau\in\mathcal{T}} \left\{ \left[\mathbb{E}_{\mathrm{L}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_{\mathrm{H}}-\lambda_1}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) + \mathbb{E}_{\mathrm{H}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_1-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) \right] \right\} + \\
&\qquad (1-\alpha) \sup_{\tau\in\mathcal{T}} \left\{ \left[\mathbb{E}_{\mathrm{L}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_{\mathrm{H}}-\lambda_2}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) + \mathbb{E}_{\mathrm{H}}[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)] \left(\frac{\lambda_2-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}}}\right) \right] \right\} \\
&= \alpha\,\Pi(\lambda_1) + (1-\alpha)\,\Pi(\lambda_2).
\end{aligned}$$

Finally, to prove the convexity of $\Pi(\lambda, n)$ on n note that

$$\begin{aligned}
\Pi(\lambda, n+1) + \Pi(\lambda, n-1) &= \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda_\tau, n+1+N_\tau)\right] + \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda_\tau, n-1+N_\tau)\right] \\
&\ge \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\,\tau} \Big(G(\lambda_\tau, n+1+N_\tau) + G(\lambda_\tau, n-1+N_\tau) \Big)\right] \\
&\ge 2 \sup_{\tau\in\mathcal{T}} \mathbb{E}\left[e^{-r\,\tau}\, G(\lambda_\tau, n+N_\tau)\right] \\
&= 2\,\Pi(\lambda, n),
\end{aligned}$$

where the second inequality follows from the convexity of $G(\lambda, n)$ on n. □

Proof of Proposition 3: Let us first derive the lower bound. Suppose the seller uses the strategy of running the voting system until time τ_m at which the m^{th} vote is collected and then at this time he launches the product (*i.e.*, the option of discarding is not considered). Since this strategy is feasible and $\tau_m \in \mathcal{T}$ we have that

$$\Pi(\lambda, n) \ge \mathbb{E}\Big[e^{-r\,\tau_m}\, G(\lambda_{\tau_m}, N_{\tau_m}) | N_0 = n\Big] \ge \mathbb{E}\Big[e^{-r\,\tau_m} \big(\beta + \lambda_{\tau_m} + \phi\,(n+m)\big)\Big].$$

In addition, $\lambda_{\tau_m} = \mathbb{E}[\Lambda | \mathcal{F}_{\tau_m}]$ and so we also have that

$$\begin{aligned}
\Pi(\lambda, n) &\ge \mathbb{E}\Big[e^{-r\,\tau_m} \big(\beta + \mathbb{E}[\Lambda|\mathcal{F}_{\tau_m}] + \phi\,(n+m)\big)\Big] = \mathbb{E}\Big[\mathbb{E}\Big[e^{-r\,\tau_m} \big(\beta + \Lambda + \phi\,(n+m)\big) \Big| \mathcal{F}_{\tau_m}\Big]\Big] \\
&= \mathbb{E}\left[e^{-r\,\tau_m} \big(\beta + \Lambda + \phi\,(n+m)\big)\right] = \mathbb{E}\Big[\mathbb{E}\left[e^{-r\,\tau_m} \big(\beta + \Lambda + \phi\,(n+m)\big) \big| \Lambda\right]\Big] \\
&= \mathbb{E}\left[\left(\frac{\Lambda}{\Lambda + r}\right)^m \Big(\beta + \Lambda + \phi\,(n+m)\Big)\right] = \mathbb{E}[H(\Lambda, n, m)].
\end{aligned}$$

Taking 'max' over m and noticing that $\Pi(\lambda, n)$ is trivially greater than R, we conclude that

$$\Pi(\lambda, n) \ge \max\Big\{R\ ,\ \max_{m\ge 0} \mathbb{E}\big[H(\Lambda, n, m)\big]\Big\}.$$

To derive the upper bound, we use the convexity $\Pi(\lambda, n)$ on λ derived in Proposition 1. It follows by Jensen's inequality that

$$
\begin{aligned}
\Pi(\lambda, n) &= \Pi(\mathbb{E}[\Lambda], n) \le \mathbb{E}[\Pi(\Lambda, n)] = \mathbb{P}(\Lambda = \lambda_{\mathrm{L}})\, \Pi(\lambda_{\mathrm{L}}, n) + \mathbb{P}(\Lambda = \lambda_{\mathrm{H}})\, \Pi(\lambda_{\mathrm{H}}, n) \\
&= \mathbb{P}(\Lambda = \lambda_{\mathrm{L}})\, \Pi^{\mathrm{D}}(\lambda_{\mathrm{L}}, n) + \mathbb{P}(\Lambda = \lambda_{\mathrm{H}})\, \Pi^{\mathrm{D}}(\lambda_{\mathrm{H}}, n) \\
&= \mathbb{E}\left[\Pi^{\mathrm{D}}(\Lambda, n)\right].
\end{aligned}
$$

Note that we have used the identities $\Pi(\lambda_{\mathrm{L}}, n) = \Pi^{\mathrm{D}}(\lambda_{\mathrm{L}}, n)$ and $\Pi(\lambda_{\mathrm{H}}, n) = \Pi^{\mathrm{D}}(\lambda_{\mathrm{H}}, n)$, which follow from the fact that the states $\lambda = \lambda_{\mathrm{L}}$ and $\lambda = \lambda_{\mathrm{H}}$ are absorbing for the dynamics of λ_t. $\square$

Proof of Theorem 1: Let $Z(\lambda) \in \mathcal{D}$ be a function that satisfies the QVI conditions and consider an arbitrary admissible policy $\tau \in \mathcal{T}$ and let λ_t be the corresponding trajectory of the belief process starting at $\lambda_0 = \lambda$.

Given the assumptions on $Z(\cdot)$, we can apply integration by part followed by Itô's lemma (see Protter, 2004) to get that

$$
\begin{aligned}
e^{-r\tau} Z(\lambda_\tau) &= Z(\lambda) + \int_0^\tau e^{-r\,t} \Big(\lambda_t \left[Z(\lambda_t + \eta(\lambda_t)) - Z(\lambda_t) - \eta(\lambda_t)\, \partial_+ Z(\lambda_t)\right] - r\, Z(\lambda_t)\Big)\, \mathrm{d}t \\
&\quad + \int_0^\tau e^{-r\,t} [Z(\lambda_{t-} + \eta(\lambda_{t-})) - Z(\lambda_{t-})]\, (\mathrm{d}N_t - \lambda_t\, \mathrm{d}t) \\
&= Z(\lambda) + \int_0^\tau e^{-r\,t}\, \mathcal{H}Z(\lambda_t)\, \mathrm{d}t + \int_0^\tau e^{-r\,t} \left[Z(\lambda_{t-} + \eta(\lambda_{t-})) - Z(\lambda_{t-})\right] \mathrm{d}M_t,
\end{aligned}
$$

where $\mathcal{H}$ is the operator defined in (8) and M_t is an $\mathcal{F}_t$-martingale given by

$$
M_t := N_t - \int_0^t \lambda_s\, \mathrm{d}s.
$$

Since Z satisfies the QVI in (9), it follows that $\mathcal{H}Z(\lambda_t) \le 0$ which implies

$$
e^{-r\,\tau} Z(\lambda_\tau) \le Z(\lambda_0) + \int_0^\tau e^{-r\,t} \left[Z(\lambda_{t-} + \eta(\lambda_{t-})) - Z(\lambda_{t-})\right] \mathrm{d}M_t.
$$

Now, note that the term $e^{-r\,t}[Z(\lambda_{t-} + \eta(\lambda_{t-})) - Z(\lambda_{t-})]$ is a bounded previsible process while M_t is a martingale and thus the expectation of this last integral is equal to zero. As a result, we get that

$$
\mathbb{E}[e^{-r\tau} Z(\lambda_\tau)] \le Z(\lambda).
$$

In addition, from the second QVI condition in Definition 1 we have that $Z(\lambda) \ge G(\lambda)$ and get

$$
Z(\lambda) \ge \mathbb{E}[e^{-r\tau} G(\lambda_\tau)].
$$

Since this inequality follows for any admissible $\tau \in \mathcal{T}$, we conclude that

$$
Z(\lambda) \ge \Pi(\lambda).
$$

To complete the proof, note that all the inequalities above become equalities for the QVI-control associated to $Z(\lambda)$. $\square$

Proof of Proposition 4: For $i = 1$, we can solve directly (12) since $Z_0(\lambda) = \beta + \lambda$. After some straightforward calculation we get that

$$Z_1(\lambda) = K_1 \frac{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}} + A_1 (\lambda_{\rm H} - \lambda) + B_1 (\lambda - \lambda_{\rm L}) \qquad \lambda \in [\lambda_1, \lambda_0], \tag{31}$$

where K_1 is a constant of integration to be determined and

$$A_1 := \frac{\lambda_{\rm L} (\beta + \lambda_{\rm L})}{\alpha_{\rm L} (\lambda_{\rm H} - \lambda_{\rm L})^2} \quad \text{and} \quad B_1 := \frac{\lambda_{\rm H} (\beta + \lambda_{\rm H})}{\alpha_{\rm H} (\lambda_{\rm H} - \lambda_{\rm L})^2}.$$

We can determine the value of K_1 as to guarantee continuity of the function Z and so by using the value matching condition $Z_0(\bar\lambda) = Z_1(\bar\lambda)$

$$K_1 = \frac{(\bar\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - \bar\lambda)^{\alpha_{\rm H}}} \left[\beta + \bar\lambda - A_1(\lambda_{\rm H} - \bar\lambda) - B_1(\bar\lambda - \lambda_{\rm L})\right].$$

Given the solution for $Z_1(\lambda)$ in (31), we can solves for $Z_2(\lambda)$ in $\lambda \in [\lambda_2, \lambda_1]$ integrating the ODE

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}} Z_2(\lambda)\right] = \frac{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}} \frac{Z_1(\lambda + \eta(\lambda))}{\eta(\lambda)}, \qquad \lambda \in [\lambda_2, \lambda_1].$$

After some tedious but straightforward calculations we get

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}}{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}} Z_2(\lambda)\right] = \frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{K_1 \lambda_{\rm L}^{\alpha_{\rm H}}}{(\lambda_{\rm H} - \lambda_{\rm L})\lambda_{\rm H}^{\alpha_{\rm L}}} \ln\left(\frac{\lambda - \lambda_{\rm L}}{\lambda_{\rm H} - \lambda}\right) + \frac{A_1 \lambda_{\rm L}}{(\lambda_{\rm H} - \lambda_{\rm L})\alpha_{\rm L}}\left(\frac{\lambda - \lambda_{\rm L}}{\lambda_{\rm H} - \lambda}\right)^{\alpha_{\rm L}} + \frac{B_1 \lambda_{\rm H}}{(\lambda_{\rm H} - \lambda_{\rm L})\alpha_{\rm H}}\left(\frac{\lambda - \lambda_{\rm L}}{\lambda_{\rm H} - \lambda}\right)^{\alpha_{\rm H}}\right]$$

which leads to

$$Z_2(\lambda) = K_2 \frac{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}} + A_2 (\lambda_{\rm H} - \lambda) + B_2 (\lambda - \lambda_{\rm L}) + C_2 \frac{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}} \ln\left(\frac{\lambda - \lambda_{\rm L}}{\lambda_{\rm H} - \lambda}\right) \qquad \lambda \in [\lambda_2, \lambda_1], \tag{32}$$

where

$$A_2 := \frac{A_1 \lambda_{\rm L}}{(\lambda_{\rm H} - \lambda_{\rm L})\alpha_{\rm L}}, \quad B_2 := \frac{B_1 \lambda_{\rm H}}{(\lambda_{\rm H} - \lambda_{\rm L})\alpha_{\rm H}}, \quad C_2 := \frac{K_1 \lambda_{\rm L}^{\alpha_{\rm H}}}{(\lambda_{\rm H} - \lambda_{\rm L})\lambda_{\rm H}^{\alpha_{\rm L}}}$$

and K_2 is a constant of integration that we find imposing the value matching condition $Z_1(\lambda_1) = Z_2(\lambda_1)$. We can see that the first three terms in the representation of $Z_2(\lambda)$ in (32) have the same functional form as those in the representation of $Z_1(\lambda)$ in (31). Hence, as we iterate this sequential resolution of the ODE, we will expect a similar pattern repeating on each iteration. That is, we expect $Z_i(\lambda)$ to satisfy

$$Z_i(\lambda) = K_i \frac{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}} + A_i (\lambda_{\rm H} - \lambda) + B_i (\lambda - \lambda_{\rm L}) + C_2 \frac{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}} \ln\left(\frac{\lambda - \lambda_{\rm L}}{\lambda_{\rm H} - \lambda}\right) + \text{other terms},$$

where

$$A_i := \frac{A_{i-1} \lambda_{\rm L}}{(\lambda_{\rm H} - \lambda_{\rm L})\alpha_{\rm L}}, \quad B_i := \frac{B_{i-1} \lambda_{\rm H}}{(\lambda_{\rm H} - \lambda_{\rm L})\alpha_{\rm H}}, \quad C_i := \frac{K_{i-1} \lambda_{\rm L}^{\alpha_{\rm H}}}{(\lambda_{\rm H} - \lambda_{\rm L})\lambda_{\rm H}^{\alpha_{\rm L}}}$$

and K_i is a constant of integration that we find imposing the value matching condition $Z_{i-1}(\lambda_{i-1}) = Z_i(\lambda_{i-1})$.

In order to find the '*other terms*' in the representation of $Z_i(\lambda)$ we need to identify how the logarithmic term in (32) will expand as we keep integrating the ODE for $i = 3, 4, \ldots$. Rather than iterating the ODE to derive this new term, we will use an "educated guess" to postulate a generic solution for $Z_i(\lambda)$ and then use induction to verify the prosed solution is indeed the right one. To this end, let us suppose that $Z_i(\lambda)$ admits the following representation

$$Z_i(\lambda) = A_i (\lambda_{\rm H} - \lambda) + B_i (\lambda - \lambda_{\rm L}) + \frac{(\lambda_{\rm H} - \lambda)^{\alpha_{\rm H}}}{(\lambda - \lambda_{\rm L})^{\alpha_{\rm L}}} \sum_{n=0}^{i-1} C_{i,n} \ln^n\left(\frac{\lambda_{\rm H}(\lambda - \lambda_{\rm L})}{\lambda_{\rm L}(\lambda_{\rm H} - \lambda)}\right), \qquad \lambda \in [\lambda_i, \lambda_{i-1}], \tag{33}$$

where A_i, B_i and $C_{i,n}$ are constant to be determined. We use the convention $\sum_{k=0}^{-1} x_k = 0$ so that condition (33) holds for $i = 0$ with $A_0 = (\beta + \lambda_{\mathrm{L}})/(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})$ and $B_0 = (\beta + \lambda_{\mathrm{H}})/(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})$.
We will now use the ODE (12) to compute $Z_{i+1}(\lambda)$ in the domain $\lambda \in [\lambda_{i+1}, \lambda_i]$ and we will verify that it has the desired form given by (33).

First of all, we note that under condition (33) and for $\lambda \in [\lambda_{i+1}, \lambda_i]$ we have

$$\begin{aligned}\frac{(\lambda-\lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}}-\lambda)^{\alpha_{\mathrm{H}}}} \frac{Z_i(\lambda+\eta(\lambda))}{\eta(\lambda)} \quad &= \quad A_i\,\lambda_{\mathrm{L}}\,\frac{(\lambda-\lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}-1}}{(\lambda_{\mathrm{H}}-\lambda)^{\alpha_{\mathrm{H}}}} + B_i\,\lambda_{\mathrm{H}}\,\frac{(\lambda-\lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}}-\lambda)^{\alpha_{\mathrm{H}}+1}} \\ &+ \quad \frac{\lambda_{\mathrm{L}}^{\alpha_{\mathrm{H}}}}{\lambda_{\mathrm{H}}^{\alpha_{\mathrm{L}}}}\,\frac{1}{(\lambda_{\mathrm{H}}-\lambda)\,(\lambda-\lambda_{\mathrm{L}})} \sum_{n=0}^{i-1} C_{i,n}\,\ln^n\left(\frac{\lambda_{\mathrm{H}}\,(\lambda-\lambda_{\mathrm{L}})}{\lambda_{\mathrm{L}}\,(\lambda_{\mathrm{H}}-\lambda)}\right).\end{aligned}$$

Now, combing the fact that $\alpha_{\mathrm{H}} = \alpha_{\mathrm{L}} + 1$ and the identities

$$\frac{(\lambda-\lambda_{\mathrm{L}})^{\alpha-1}}{(\lambda_{\mathrm{H}}-\lambda)^{\alpha+1}} = \frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{1}{\alpha\,(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})}\left(\frac{\lambda-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda}\right)^{\alpha}\right] \quad \text{and}$$
$$\frac{1}{(\lambda_{\mathrm{H}}-\lambda)\,(\lambda-\lambda_{\mathrm{L}})}\ln^n\left(\frac{\lambda_{\mathrm{H}}\,(\lambda-\lambda_{\mathrm{L}})}{\lambda_{\mathrm{L}}\,(\lambda_{\mathrm{H}}-\lambda)}\right) = \frac{1}{(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})}\frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{1}{n+1}\ln^{n+1}\left(\frac{\lambda_{\mathrm{H}}\,(\lambda-\lambda_{\mathrm{L}})}{\lambda_{\mathrm{L}}\,(\lambda_{\mathrm{H}}-\lambda)}\right)\right]$$

we conclude that

$$\begin{aligned}\frac{(\lambda-\lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}}-\lambda)^{\alpha_{\mathrm{H}}}} \frac{Z_i(\lambda+\eta(\lambda))}{\eta(\lambda)} \quad &= \quad \frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{A_i\,\lambda_{\mathrm{L}}}{(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})\,\alpha_{\mathrm{L}}}\left(\frac{\lambda-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda}\right)^{\alpha_{\mathrm{L}}} + \frac{B_i\,\lambda_{\mathrm{H}}}{(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})\,\alpha_{\mathrm{H}}}\left(\frac{\lambda-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda}\right)^{\alpha_{\mathrm{H}}}\right. \\ &\left. + \quad \frac{\lambda_{\mathrm{L}}^{\alpha_{\mathrm{H}}}}{(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})\,\lambda_{\mathrm{H}}^{\alpha_{\mathrm{L}}}} \sum_{n=0}^{i-1} \frac{C_{i,n}}{n+1}\ln^{n+1}\left(\frac{\lambda_{\mathrm{H}}\,(\lambda-\lambda_{\mathrm{L}})}{\lambda_{\mathrm{L}}\,(\lambda_{\mathrm{H}}-\lambda)}\right)\right].\end{aligned}$$

From this observation, it is not hard to see that the ODE

$$\frac{\mathrm{d}}{\mathrm{d}\lambda}\left[\frac{(\lambda-\lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}}-\lambda)^{\alpha_{\mathrm{H}}}} Z_{i+1}(\lambda)\right] = \frac{(\lambda-\lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}}-\lambda)^{\alpha_{\mathrm{H}}}} \frac{Z_i(\lambda+\eta(\lambda))}{\eta(\lambda)},.$$

has a solution of that is similar in form to that in equation (33), that is,

$$Z_{i+1}(\lambda) = A_{i+1}\,(\lambda_{\mathrm{H}}-\lambda) + B_{i+1}\,(\lambda-\lambda_{\mathrm{L}}) + \frac{(\lambda_{\mathrm{H}}-\lambda)^{\alpha_{\mathrm{H}}}}{(\lambda-\lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}} \sum_{n=0}^{i} C_{i+1,n}\,\ln^n\left(\frac{\lambda-\lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}}-\lambda}\right), \qquad \lambda \in [\lambda_{i+1}, \lambda_i],$$

where the coefficients A_{i+1}, B_{i+1} and $\{C_{i+1,n} : n = 0, \ldots, i\}$ satisfy the recursions

$$A_{i+1} := \frac{A_i\,\lambda_{\mathrm{L}}}{(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})\,\alpha_{\mathrm{L}}}, \qquad B_{i+1} := \frac{B_i\,\lambda_{\mathrm{H}}}{(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})\,\alpha_{\mathrm{H}}}, \qquad C_{i+1,n} = \frac{\lambda_{\mathrm{L}}^{\alpha_{\mathrm{H}}}}{(\lambda_{\mathrm{H}}-\lambda_{\mathrm{L}})\,\lambda_{\mathrm{H}}^{\alpha_{\mathrm{L}}}}\,\frac{C_{i,n-1}}{n}, \quad n = 1, 2, \ldots, i,$$

and $C_{i+1,0} = K_{i+1}$, where K_{i+1} is a constant of integration that is computed imposing the value-matching $Z_{i+1}(\lambda_i) = Z_i(\lambda_i)$. □

Proof of Proposition 5: To prove part (a), it is not hard to show that the function $Z(\lambda)$ defined in the proposition satisfies the QVI conditions given in (9) for all $\lambda \in (\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}})$ and the border conditions $Z(\lambda_{\mathrm{L}}) = G(\lambda_{\mathrm{L}})$ and $Z(\lambda_{\mathrm{H}}) = G(\lambda_{\mathrm{H}})$. As a result, a direct application of the *Verification Theorem* 1 implies that $\Pi(\lambda) = Z(\lambda)$.

The proof of part (b) is based on the convexity of $Z(\lambda;\bar\lambda)$ on λ and its continuity on $\bar\lambda$. The argument in this proof can be easily explained using again the three curves $Z(\lambda,\bar\lambda)$ depicted in Figure 3. First, note that if $\bar\lambda$ is sufficiently high (i.e., $\bar\lambda = 0.85$ in the figure) then $Z(\lambda;\bar\lambda)$ will not intersect $G(\lambda)$ in $(\lambda_{\text{L}},\bar\lambda)$ and the value-matching condition in part (a) cannot be satisfied. To see this, note that from the ODE (11) that defines $Z(\lambda;\bar\lambda)$ we have that

$$Z'(\lambda;\bar\lambda) = \frac{Z(\lambda+\eta(\lambda);\bar\lambda) - Z(\lambda;\bar\lambda)}{\eta(\lambda)} - \frac{r\,Z(\lambda;\bar\lambda)}{(\lambda_{\text{H}}-\lambda)(\lambda-\lambda_{\text{L}})}, \quad \lambda \in (\lambda_{\text{L}},\bar\lambda).$$

For $\bar\lambda$ sufficiently high and $\lambda \uparrow \bar\lambda$, it is not hard to see that $Z'(\lambda;\bar\lambda) < 0$. Hence, by the convexity of $Z(\lambda;\bar\lambda)$ on λ, the function $Z(\lambda;\bar\lambda)$ is decreasing in $(\lambda_{\text{L}},\bar\lambda)$ and therefore does not intersect $G(\lambda)$.
Consider now the case in which $\bar\lambda$ approaches $\tilde\lambda$ from above. We want to show that $Z(\lambda;\bar\lambda)$ will intersect $G(\lambda)$, this is the same as requiring that $\lim_{\lambda\uparrow\bar\lambda} Z'(\lambda;\bar\lambda) > 0$ as $\bar\lambda \downarrow \tilde\lambda$. Using again the ODE (11), this condition is equivalent to

$$(\lambda_{\text{H}}-\tilde\lambda)\,(\tilde\lambda-\lambda_{\text{L}}) > R\,r.$$

Recall that $\tilde\lambda = R-\beta$. To complete the proof, note that $Z(\lambda;\bar\lambda)$ is a solution to the first-order ODE (11), with border condition $Z(\bar\lambda;\bar\lambda) = \beta+\bar\lambda$, which is a continues function of $\bar\lambda$. Hence, by the previous argument it follows that if the previous inequality holds then there must exist a unique $\lambda \in (\tilde\lambda,\lambda_{\text{H}})$ such that the value-matching and smooth-pasting conditions are satisfied. $\square$

Proof of Proposition 6: For the sake of the proof, we denote by $\bar Q^k_t = \int_0^t (1-q^k_s)\mathrm{d}s$ and $N_1(t)$ the homogenous Poisson process with unit rate. We write $\bar\Lambda^k(q) = k + \varphi(1-q)\sqrt{k}$. We denote by $N^k(t)$ the Poisson process with rate $\int_0^t \bar\Lambda^k(q^k_s)\mathrm{d}s = kt + \sqrt{k}\,\varphi\,\bar Q^k_t$ and we observe that $N^k(t) \overset{\text{d}}{=} N_1(kt) + N_1(\sqrt{k}\,\varphi\,\bar Q^k_t)$. We wish to prove a FCLT for $N^k(t)$ whereby,

$$\frac{N^k(t) - \int_0^t \bar\Lambda^k(q^k_s)\mathrm{d}s}{\sqrt{k}} \overset{\text{d}}{=} \frac{N_1(kt)-kt}{\sqrt{k}} + \left(\frac{N_1(\sqrt{k}\,\varphi\,\bar Q^k_t)}{\sqrt{k}} - \varphi\,\bar Q^k_t\right) \Rightarrow W_t$$

as $k \to \infty$, in the Skorohod topology on $D[0,\infty)$ (see Billingsley, 1968). This result follows first by appealing to the FCLT for a homogenous Poisson process, secondly by recalling that $\bar Q^k_t$ is stochastically bounded and arguing that a variant of the SLLN applies forcing the second term in the previous sum to converge almost surely to zero, and finally by noticing that the limit of the sum converges to the sum of the limits and that is true in this case because of the continuity (a.s.) of both limits (see Whitt, 1980). What remains is a careful argument of why the term

$$\left(\frac{N_1(\sqrt{k}\,\varphi\,\bar Q^k_t)}{\sqrt{k}} - \varphi\,\bar Q^k_t\right) \to 0 \;\; a.s.$$

We recall that $\bar Q^k_t$ lives in a compact set. Hence, any subsequence of $\bar Q^k_t$ contains a sub-subsequence $(m_n : n \geq 0)$ that converges to some limit $l \geq 0$. Thus, for any $\epsilon > 0$, and for n large enough,

$$\Big|\frac{N_1(\sqrt{m_n}\,\varphi\,\bar Q^{m_n}_t)}{\sqrt{m_n}} - \varphi\,\bar Q^{m_n}_t\Big| \leq \Big|\frac{N_1(\sqrt{m_n}\,\varphi\,(l+\epsilon))}{\sqrt{m_n}} - l\,\varphi\Big| + \big|\varphi\,(l-\bar Q^{m_n}_t)\big| \leq 2\,\epsilon.$$

It is not hard to see that this forces any converging subsequence of $\frac{N_1(\sqrt{k}\,\varphi\,\bar{Q}^k_t)}{\sqrt{k}} - \varphi\,\bar{Q}^k_t$ to converge to zero almost surely as $k\to\infty$. But, the sequence itself is clearly almost surely bounded, hence we conclude that it must converge to zero almost surely as $k\to\infty$. By noticing that $\big|\frac{N_1(\sqrt{k}\,\varphi\,\bar{Q}^k_t)}{\sqrt{k}} - \varphi\,\bar{Q}^k_t\big|$ is uniformly bounded in t by a converging sequence $\frac{N_1(\sqrt{k}\,\varphi)}{\sqrt{k}} + \varphi$ shows that the convergence of $\big|\frac{N_1(\sqrt{k}\,\varphi\,\bar{Q}^k_t)}{\sqrt{k}} - \varphi\,\bar{Q}^k_t\big|$ to zero is uniform in t.

Going back to the dynamics of the belief process, we have that

$$\mathrm{d}q^k_t = -\sqrt{k}(\lambda^k_{\mathrm{H}} - \lambda^k_{\mathrm{L}})\,\frac{q^k_t(1-q^k_t)}{\bar{\Lambda}^k(q^k_t)}\,\mathrm{d}Y^k_t,$$

where $Y^k_t = \frac{1}{\sqrt{k}}\left[dN^k_t - \bar{\Lambda}^k(q^k_t)\mathrm{d}t\right]$. It is also easily verified that if (q_k, Y^k) converges to (q, Y) in the Skorohod topology, then $\big(q_k, -\sqrt{k}(\lambda^k_{\mathrm{H}} - \lambda^k_{\mathrm{L}})\,\frac{q_k(1-q_k)}{\bar{\Lambda}^k(q_k)}, Y_k\big) \to \big(q, \varphi\, q(1-q), Y\big)$, as $k\to\infty$. Recalling the FCLT type convergence obtained above and the fact that $N^k(t) - \int_0^t \bar{\Lambda}^k(q^k_s)\mathrm{d}s$ is a well behaving martingale, we conclude from Proposition 5.1 of Kurtz and Protter (1991) on convergence of sequences of solutions of stochastic differential equations, that every limit point of the sequence $q_k(\cdot)$, satisfies the following SDE

$$\mathrm{d}q_t = \varphi\, q_t\,(1-q_t)\,\mathrm{d}W_t.$$

Hence, any converging subsequence of q^k converges to the same limit that is solution to the previous SDE. Given that the sequence q^k is uniformly bounded and take values in the compact $[0,1]$, we conclude that the sequence itself q^k, converges to q - solution to the SDE.

Proof of Corollary 3: To compute the probability $\gamma^* = \mathbb{P}(\tau^* = \underline{q})$ that the product will be introduced, we use Dynkin's formula where for some stopping time τ_q

$$\mathbb{E}[f(q_{\tau_q})] = f(q) + \mathbb{E}\left[\int_0^{\tau_q} \mathcal{G}f(q_t)\,\mathrm{d}t\right], \qquad \text{where} \quad \mathcal{G}f(q) := \frac{(\varphi\, q\,(1-q))^2}{2}\,\frac{\mathrm{d}^2 f(q)}{\mathrm{d}q^2},$$

and $\mathcal{G}$ is the infinitesimal generator of the diffusion process q_t in equation (18). Consider the identity function $f(q) = q$. It follows that $\mathcal{G}f(q) = 0$ and by Dynkin's formula $\mathbb{E}[q_{\tau_q}] = q$. But since τ_q is the first exit time of the process q_t from the continuation region $(\underline{q}, \bar{q})$ we have that $\mathbb{E}[q_{\tau_q}] = \gamma^*\,\underline{q} + (1-\gamma^*)\,\bar{q}$ and the result follows by replacing $\underline{q}$ and $\bar{q}$ by their values as a function respectively of $\bar{\lambda}$ and $\underline{\lambda}$.

To compute the expectation $\mathbb{E}[\tau^*]$, we consider a function $\tilde{\mathcal{T}}(q)$ such that $\mathcal{G}(\tilde{\mathcal{T}}) = 1$. One can verify that the function $\tilde{\mathcal{T}}(q) = \frac{2}{\varphi^2}\,(2q-1)\,\ln\left(\frac{q}{1-q}\right)$ satisfies this condition. It follows from Dynkin's formula that $\mathbb{E}[\tilde{\mathcal{T}}(q_{\tau^*})] = \mathcal{T}(q) + \mathbb{E}[\tau^*]$. The result follows directly from this equality when we replace $\underline{q}$ and $\bar{q}$ by their values as a function respectively of $\bar{\lambda}$ and $\underline{\lambda}$. □

Proof of Lemma 2: First of all, it is easy to see that for all $n \geq \bar{n}$ we have $G(\lambda, N) = \beta + \lambda + \phi\, n \geq R$. That is, the number of voters is so large that the seller will never discard the product. Hence, the value function $\Pi(\lambda, n)$ satisfies

$$\Pi(\lambda, n) = \sup_{\tau\in\mathcal{T}}\left\{\mathbb{E}\Big[e^{-r\,\tau}\,\big(\beta + \lambda_\tau + \phi(n + N_\tau)\big)\,|\,\lambda_0 = \lambda, N_0 = n\Big]\right\}.$$

Now, from Itô's lemma (see Protter, 2004) we have that

$$e^{-r\,\tau}\,(\beta+\lambda_\tau+\phi(n+N_\tau)) = \beta+\lambda+\phi\,n+\int_0^\tau e^{-r\,t}\,\Big(\lambda_t\,\phi - r\,(\beta+\lambda_t+\phi\,(n+N_t))\Big)\,\mathrm{d}t + \int_0^\tau e^{-r\,t}(\eta(\lambda_{t-})+\phi)\,\mathrm{d}M_t,$$

where $M_t = N_t - \int_0^t \lambda_s \, \mathrm{d}s$ is an $\mathcal{F}_t$-martingale. Also, since $n \geq \bar{n}$, one can show that the argument inside the '$\mathrm{d}t$' integral is non-positive for all $\lambda_t \in [\lambda_{\mathrm{L}}, \lambda_{\mathrm{H}}]$. We conclue that

$$\mathbb{E}[e^{-r\,\tau}\,(\beta + \lambda_\tau + \phi(n + N_\tau)] \leq \beta + \lambda + \phi\, n$$

and therefore stopping immediately (i.e., $\tau = 0$) is optimal. $\square$

Proof of Proposition 8: Suppose the system is currently in state (λ_0, n) and let τ_0 be the time until the next epoch of N_t (i.e., the time until the next vote) assuming the seller keeps the voting system open. Now, starting from the state (λ_0, n), the belief process λ_t evolves in the time interval $[0, \tau_0)$ according to the dynamics

$$\mathrm{d}\lambda_t = -\,(\lambda_{\mathrm{H}} - \lambda_t)\,(\lambda_t - \lambda_{\mathrm{L}})\,\mathrm{d}t. \tag{34}$$

Because λ_t evolves deterministically before τ_0, it should be clear that if the seller decides to stop at a time $\tau_1 < \tau_0$ (i.e., before a new vote) then this choice of τ_1 could have been made with the information available at time $t = 0$. In other words, at time $t = 0$, the seller can select a deterministic time $\tau_1 \geq 0$ (possibly infinite) at which to stop if a new voter has not arrived before this time. It follows that

$$\Pi(\lambda_0, n) = \sup_{\tau_1 \geq 0} \mathbb{E}\left[e^{-r\,(\tau_1 \wedge \tau_0)}\left(G(\lambda_{\tau_1}, n)\,\mathbb{1}(\tau_1 < \tau_0) + \Pi(\lambda_{\tau_0} + \eta(\lambda_{\tau_0}), n+1)\,\mathbb{1}(\tau_1 \geq \tau_0)\right)\right],$$

where the expectation is taken over τ_0, the epoch of the next vote. It is not hard to see that the probability distribution of τ_0 satisfies

$$\mathbb{P}(\tau_0 > t) = \exp\left(-\int_0^t \lambda_s \mathrm{d}s\right) = \left(\frac{\lambda_{\mathrm{H}} - \lambda_0}{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}}\right) e^{-\lambda_{\mathrm{L}} t} + \left(\frac{\lambda_0 - \lambda_{\mathrm{L}}}{\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}}}\right) e^{-\lambda_{\mathrm{H}} t}.$$

As a result, we can rewrite the value of $\Pi(\lambda_0, n)$ as follows

$$\Pi(\lambda_0, n) = \sup_{\tau_1 \geq 0}\left[\int_0^{\tau_1} e^{-r\,t}\,\Pi(\lambda_t + \eta(\lambda_t), n+1)\,\lambda_t\,\mathbb{P}(\tau_0 > t)\,\mathrm{d}t + e^{-r\,\tau_1}\,G(\lambda_{\tau_1}, n)\,\mathbb{P}(\tau_0 > \tau_1)\right].$$

Furthermore, because of the one-to-one correspondence between t and λ_t, we can replace the integration with respect to t by an integration with respect to λ. For this, note that equation (34) implies the following change of variables

$$t(\lambda) = \frac{1}{(\lambda_{\mathrm{H}} - \lambda_{\mathrm{L}})}\,\ln\left(\frac{(\lambda_{\mathrm{H}} - \lambda)\,(\lambda_0 - \lambda_{\mathrm{L}})}{(\lambda - \lambda_{\mathrm{L}})\,(\lambda_{\mathrm{H}} - \lambda_0)}\right).$$

With a slight abuse of notation, let us denote $\lambda_1 = \lambda_{\tau_1}$ and rewrite the seller optimization problem using the decision variable λ_1 instead of τ_1 as follows

$$\Pi(\lambda_0, n) = \max_{\lambda_1 \in (\lambda_{\mathrm{L}}, \lambda_0]}\left[\int_{\lambda_1}^{\lambda_0} e^{-r\,t(\lambda)}\,\frac{\Pi(\lambda + \eta(\lambda), n+1)}{\eta(\lambda)}\,\mathbb{P}(\tau_0 > t(\lambda))\,\mathrm{d}\lambda + e^{-r\,t(\lambda_1)}\,G(\lambda_1, n)\,\mathbb{P}(\tau_0 > t(\lambda_1))\right].$$

After some manipulations, we get that (recall the definitions of α_{L} and α_{H} in (12))

$$\Pi(\lambda_0, n) = \frac{(\lambda_{\mathrm{H}} - \lambda_0)^{\alpha_{\mathrm{H}}}}{(\lambda_0 - \lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}\,\max_{\lambda_1 \in (\lambda_{\mathrm{L}}, \lambda_0]}\left[\int_{\lambda_1}^{\lambda_0} \frac{(\lambda - \lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}} - \lambda)^{\alpha_{\mathrm{H}}}}\,\frac{\Pi(\lambda + \eta(\lambda), n+1)}{\eta(\lambda)}\,\mathrm{d}\lambda + \frac{(\lambda_1 - \lambda_{\mathrm{L}})^{\alpha_{\mathrm{L}}}}{(\lambda_{\mathrm{H}} - \lambda_1)^{\alpha_{\mathrm{H}}}}\,G(\lambda_1, n)\right]. \;\square$$

Proof of Lemma 3: Let $\{\tau_1, \tau_2, \ldots, \tau_T\}$ be the arrival epochs of $\{N_t,\ 0 \leq t \leq \tau\}$ with T the total number of voters (possibly infinite). For voter n with $1 \leq n \leq T$, we define two random variables ν_n and ζ_n such that ν_n is a Bernoulli random variable with probability θ of success and ζ_n is an exponential random variable with rate α. If $\nu_n = 0$ then voter n is an impulse buyer which will not be interested in buying the product later on. On the other hand, conditional on $\nu_n = 1$, i.e. voter n is non-impulse, ζ_n measures the amount of time he will be willing to wait for the product. We assume that $\{\nu_n, \zeta_n,\ 1 \leq n \leq T\}$ are independent random variables.

It follows that the number of sales generated at an introduction time $t \leq \tau$ is equal to

$$X_t := \sum_{n=1}^{T} \mathbb{1}\big(\tau_n \leq t \quad \text{and} \quad \nu_n = 1 \quad \text{and} \quad \zeta_n \geq t - \tau_n\big).$$

As a result, the expected number of sales given $\mathcal{F}_t$ is equal to

$$N_t = \mathbb{E}[X_t|\mathcal{F}_t] = \sum_{n=1}^{T} \mathbb{1}(\tau_n \leq t)\, \mathbb{P}(\nu_n = 1 \quad \text{and} \quad \zeta_n \geq t - \tau_n) = \sum_{n=1}^{T} \mathbb{1}(t_n \leq t)\, \theta\, e^{-\alpha\,(t-\tau_n)}.$$

From this expression, it is easy to see that N_t satisfies the SDE in the Lemma. We end this proof by getting a closed form solution by writing

$$\begin{aligned}
\int_0^t \mathrm{d}(e^{\alpha\, u} N_u) = e^{\alpha\, t}\, N_t &= N_0 + \int_0^t e^{\alpha\, u} \mathrm{d}N_u + \int_0^t \alpha N_u e^{\alpha\, u} \mathrm{d}u \\
&= N_0 + \int_0^t e^{\alpha\, u}(-\alpha N_u \mathrm{d}u + \theta \mathrm{d}N_u) + \alpha \int_0^t S_u e^{\alpha\, u} \mathrm{d}u \\
&= N_0 + \int_0^t \theta\, e^{\alpha\, u}\, \mathrm{d}N_u. \quad \square
\end{aligned}$$