

Dynamic Learning in Markets: Pricing, Advertising, and Information Acquisition*

Thomas A. Weber
École Polytechnique Fédérale de Lausanne
thomas.weber@epfl.ch

Abstract

In the face of demand uncertainty, a monopolist can observe sales as a controlled reaction to its price and advertising so as to improve the choice of this marketing mix in the future. Furthermore, to upgrade its knowledge about demand the firm has the option to invest in external market intelligence and thus to directly acquire relevant information. Using a two-period model we determine the firm's profit-maximizing learning strategy using all three of these levers: price, advertising, and information acquisition. This illustrates the firm's tradeoff of actively managing its consumer base through costly marketing, exploiting expected demand through pricing, and increasing the efficiency of its actions by means of costly outside information. An extension of the model to the case with internal budget constraints on information acquisition is provided, and a numerical example is discussed.

1 Introduction

In a dynamic market fraught with uncertainty a firm can acquire information to improve its operational decisions. It can also spend money on advertising, increasing the demand for its products. Here we examine the joint problem of advertising, pricing, and information acquisition. This extends the consideration set of the pricing and advertising decisions in the seminal contribution by Dorfman and Steiner [9] to include information acquisition in a simple and analytically tractable setting. The problem is solved using dynamic programming; informational and operational actions persistently interact, and information can come from an outside source (such as a market-research firm) and from an internal source (through observation of unit sales).

1.1 Literature

In their well-known article, Dorfman and Steiner [9] examine the static problem of jointly optimizing price and advertising expenditures, deriving a system of equations which involve elasticities of demand with respect to price and with respect to advertising. Nerlove and Arrow [18] consider a dynamic version of the joint optimization problem, using simple dynamics with respect to the diminishing effect of past advertising and a functional form that allows for an effective separation of the pricing problem from the system dynamics. Holthausen and Assmus [12] present a model for the allocation of an advertising budget to geographic market segments when the sales response to advertising in each segment is characterized by a probability distribution. It is shown that allocation decisions that are based on the expected sales response may be associated with a relatively large degree of risk and, therefore, non-optimal to a risk-averse manager. Their model produces an “efficient frontier” in terms of the expected profit and its variance resulting from alternative budget allocations. The manager then chooses the optimal allocation based on his/her preference function.

In single-person decision problems, information has a positive benefit at the margin, and more informative information sources are of (weakly) greater value than less informative ones [4].¹ “Garbling” an information source, which amounts to adding noise, decreases its value at least weakly [8]. Naturally, this result persists in a Bayesian setting [13]: for example, the value of a Gaussian information source is increasing in its precision (or “confidence”), defined as the inverse of its variance. Here we use the Gaussian setting because normally distributed prior beliefs update to normally distributed posterior beliefs, by means of observing normally distributed informative signals [19]. Wald [22] introduces the dynamic information-acquisition problem

*The author wishes to thank Ksenia Gattiker for valuable research assistance at the outset of this project.

¹When multiple decision-makers interact strategically, the marginal value of information may well be negative [10, 11, 23].

in a very general framework. Based on elements of that framework, Moscarini and Smith [17] consider information acquisition in continuous time so as to inform a binary decision. A specific binary decision, namely technology adoption, is considered by McCardle [14] in discrete time, with information acquisition. This implies an optimal stopping problem and an optimal policy that depends on belief thresholds. Moore and Whinston [15, 16] also analyze sequential information acquisition, followed by a final action.² By contrast, we are concerned here with problems where (nonbinary) control interventions and information acquisition coexist from period to period, in spirit similar to Weber and Nguyen [25] who examine a linear-quadratic Gaussian optimal control problem over an infinite horizon. In our setting, information acquisition is both internal and external, in the sense that at the beginning of each period there is an external source of information that can decrease the noisiness of the underlying system. Yet, at the end of each period (only relevant for the first period), there is an internal information acquisition by observing a relevant system statistic (in our case, the demand realization) that allows a separate, cost-free Bayesian belief update.

The main contribution of this paper is to examine the joint use of advertising and information acquisition as profitability-enhancing measures the firm can take and to disentangle their effects, in a simple analytically tractable setting.

1.2 Outline

The remainder of this paper is organized as follows. Sec. 2 introduces our model of the firm's dynamic price-advertising and information-acquisition decisions. Sec. 3 provides the quasi-static solution to the firm's problem where the firm acquires outside information and chooses its price and advertising optimally. This serves as the terminal period for the company's dynamic decision problem, which is solved in Sec. 4. Finally, Sec. 5 concludes the paper. The paper is new in that it combines the two types of decisions, operational and informational. In addition, informational actions can be external as well as internal. Indeed, observing the market's reaction in terms of unit sales (demand) is quasi-free, as it is directly related to the company's sales figures which are readily available in a company's accounting department.

²Some applications in healthcare are of this form; see, e.g., Cipriano and Weber [6] who determine the amount of information to collect (using public health screening) to inform a planner's decision to discontinue a public health screening program for hepatitis C.

2 Model

Consider a situation where at each time $t \in \{0, 1\}$ the firm faces a random demand of the form

$$\tilde{D}_t(a_t, p_t) = \tilde{\theta} + a_t - p_t + \tilde{\nu}_t,$$

where a_t is the amount of advertising, p_t the firm's price, and $\tilde{\theta}$ the uncertain market potential whose true value is persistent between periods, and $\tilde{\nu}_t$ a zero-mean unsystematic noise term which is independent and identically distributed (i.i.d.) between periods. For concreteness, we assume that both random variables are normally distributed, so that

$$\tilde{\theta} \sim \mathcal{N}(\mu, \sigma^2) \quad \text{and} \quad \tilde{\nu}_t \sim \mathcal{N}(0, \sigma_\nu^2),$$

where μ is the positive expected demand, σ is the positive standard deviation of the market potential, and σ_ν is the nonnegative standard deviation of the unsystematic noise. At the beginning of each period, the firm has the option to acquire an informative signal which—conditional on the realized market potential θ —is of the form:

$$\tilde{s}_t | \theta = \theta + \tilde{\varepsilon}_t,$$

where $\tilde{\varepsilon}_t \sim \mathcal{N}(0, 1/\kappa_t)$ is a zero-mean normally distributed observational noise, with precision $\kappa_t \geq 0$. If the signal's precision vanishes ($\kappa_t = 0$), then the firm acquires no information at all. The company's cost realization at time t depends on the demand realization D_t , its advertising activity a_t , and the chosen precision κ_t of the acquired information:

$$C_t(a_t, p_t, \kappa_t) = cD_t + (\alpha/2) a_t^2 + \beta \kappa_t^2.$$

The nonnegative marginal production cost $c \geq 0$ is constant, the cost of advertising is quadratic (scaled by the positive constant α), and the cost of information acquisition is quadratic in the precision (scaled by the nonnegative constant β). To avoid the degenerate outcome that the firm can increase its demand through advertising at a marginal cost that is below its marginal benefit, one needs to assume that marketing expenditures are “significant” in the sense that

$$\alpha > 1/2;$$

as shown in the next section, this restriction guarantees that the firm's absolute spending on advertising stays finite.

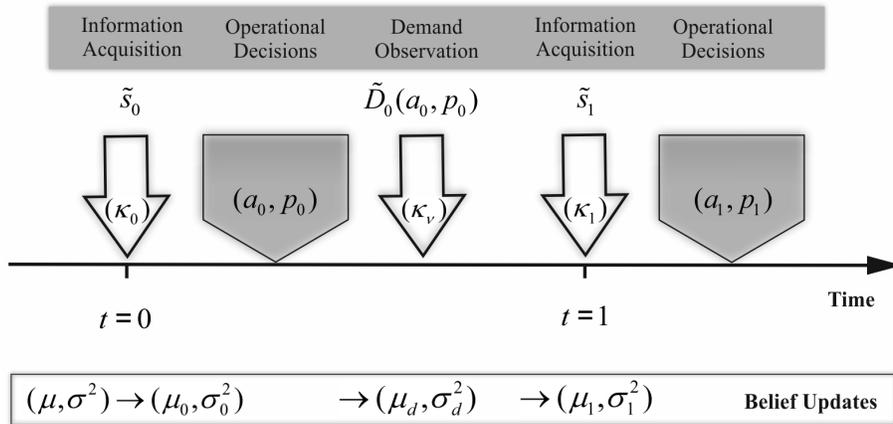


Figure 1: Timeline.

Remark 1 (Special Cases). When marginal cost is zero ($c = 0$), the firm effectively focuses on revenue maximization. In practice, this situation may occur for information products, such as computer software and digital content (e.g., music or ebooks)—at least approximately.³ The case where advertising becomes prohibitively expensive ($\alpha \rightarrow \infty$) reduces the firm's operational decision variables to price only, which may serve as an interesting base case. Finally, if the cost of information acquisition vanishes ($\beta \rightarrow 0^+$), then the firm will simply maximize expected profit (subject to the remaining unsystematic risk) and if outside information becomes prohibitively expensive ($\beta \rightarrow \infty$), then the firm refrains from acquiring outside market intelligence and its learning is restricted to what it can gather by observing demand realizations as a reaction to its operational price-advertising decisions. As shown in the next two sections, there exists a positive but finite threshold for the marginal information-acquisition cost β , above which the firm will not actively acquire information and thus is indifferent to any further increase of β .

Let \mathcal{I}_t denote the information available to the firm at time t . Then the firm's decision problem consists in finding a policy $\pi_t = \pi_t(\mathcal{I}_t)$, with $\pi_t = (a_t, p_t, \kappa_t)$, which maximizes its expected discounted profits,

$$\max_{(\pi_0, \pi_1)} \sum_{t=0}^1 \mathbb{E} \left[\delta^t \tilde{\Pi}_t(a_t, p_t, \kappa_t) \middle| \mathcal{I}_t \right],$$

where the firm's uncertain time- t profit $\tilde{\Pi}_t(a_t, p_t, \kappa_t)$ has the realization

$$\tilde{\Pi}_t(a_t, p_t, \kappa_t) = p_t D_t(a_t, p_t) - C_t(a_t, p_t, \kappa_t),$$

³Upon closer examination, the provision of most information products does in fact have a positive marginal cost, albeit possibly very small.

⁴The interest r can be viewed as the firm's opportunity cost of capital, corresponding to the investors' next-best investment return at the same risk. Interestingly, the firm's risk in this model is endogenously determined by its actions, and consequently the company's internal risk-return tradeoff may not be the same as the risk-return tradeoff in the market (as exemplified by the local slope of the efficient frontier in a CAPM-setting). The latter means that in fact the applicable discount rate may vary (at least in the medium-run) as a function of the firm's actions, an interesting second-order effect we are neglecting here.

and where $\delta = 1/(1+r) \in (0, 1]$ is the firm's per-period discount rate based on the per-period interest rate $r \geq 0$.⁴

The timing and mechanism of the firm's actions and Bayesian belief updates are as follows. At the beginning of period $t = 0$, the firm starts out with a prior belief about $\tilde{\theta}$ which is reflected in the vector of first and second moments (μ, σ^2) , which provide a complete description of the underlying normal distribution $\mathcal{N}(\mu, \sigma^2)$. Then the firm chooses κ_0 , effectively deciding about how much external market intelligence to acquire, which will allow it to observe the signal realization s_0 , based on which its belief is updated to (μ_0, σ_0^2) . The firm then selects its time-0 operational variables, advertising (a_0) and price (p_0), and observes the demand realization $d = D_0(a_0, p_0)$ which allows for a further Bayesian update of its belief about $\tilde{\theta}$, to (μ_d, σ_d^2) . At the beginning of period $t = 1$, the firm decides about how much information to acquire in terms of the signal precision κ_1 , and then updates its belief accordingly. Finally, the firm chooses its time-1 operational variables a_1 and p_1 so as to maximize this period's profits. Fig. 1 summarizes the timing of the information updates and the firm's operational decisions.

The dynamic programming principle [2, 3] implies that a solution to this finite-horizon discrete-time dynamic optimization problem can be obtained via backward induction, starting with the second period. Accordingly we obtain the quasi-static policy as the terminal policy of the two-period problem, which is discussed next in Sec. 3. The first-period problem is then solved in Sec. 4.

3 Quasi-Static Policy

At the end of the last period ($t = 1$), the firm's belief about the market potential $\tilde{\theta}$ is given by (μ_1, σ_1^2) . The firm now solves a profit-maximization problem with respect to its operational variables,

$$\max_{a_1, p_1} \{ (p_1 - c)(\mu_1 + a_1 - p_1) - (\alpha/2) a_1^2 \},$$

which yields the unique solution

$$\begin{aligned} a_1^* &= \hat{\alpha} (\mu_1 - c), \\ p_1^* &= (1 + \hat{\alpha}) \frac{\mu_1 + c}{2}, \end{aligned}$$

where we denote by

$$\hat{\alpha} \triangleq \frac{1}{2\alpha - 1} \in [0, \infty)$$

the advertising effectiveness which is decreasing in $\alpha > 1/2$ and tends to zero as α goes to infinity (i.e., when advertising becomes prohibitively expensive). The firm's maximized terminal expected profit is obtained by substituting (a_1^*, p_1^*) into the preceding objective function:

$$\Pi_1^* = (1 + \hat{\alpha}) \frac{(\mu_1 - c)^2}{4}.$$

Given the firm's belief (μ_d, σ_d^2) , information acquisition (via κ_1) therefore results in expected profits of

$$\hat{\Pi}_1(\kappa_1) \triangleq \left(\frac{1 + \hat{\alpha}}{4} \right) \mathbb{E} [(\mu_1(\tilde{s}_1) - c)^2] - \beta \kappa_1^2.$$

The value for the firm derives from the fact that its terminal operational decisions can be adjusted. Indeed, after having observed the informative signal realization s_1 the firm can update its market-potential estimator from μ_d to the weighted average

$$\mu_1(s_1) = \lambda_1 s_1 + (1 - \lambda_1) \mu_d,$$

where $\lambda_1 \triangleq \kappa_1 \sigma_d^2 / (1 + \kappa_1 \sigma_d^2) \in [0, 1]$ is increasing in κ_1 .⁵ As a result, the terminal profit (prior to information acquisition) becomes

$$\begin{aligned} \hat{\Pi}_1(\kappa_1) &= \frac{1 + \hat{\alpha}}{4} [(\mu_d - c)^2 + \lambda_1^2 \sigma_d^2] - \beta \kappa_1^2 \\ &= \frac{1 + \hat{\alpha}}{4} \left[(\mu_d - c)^2 + \frac{\kappa_1^2 \sigma_d^6}{(1 + \kappa_1 \sigma_d^2)^2} \right] - \beta \kappa_1^2. \end{aligned}$$

The first term of the firm's payoff is independent of the informational precision and corresponds to the firm's profit without information acquisition. The sum of the following two terms is concave in κ_1 and maximization

⁵The firm also updates the variance of its belief, from σ_d^2 to $\sigma_d^2 / (1 + \kappa_1 \sigma_d^2)$, independent of s_1 .

⁶Here and in what follows we use the abbreviation $[x]_+ \triangleq \max\{0, x\}$ for all $x \in \mathbb{R}$.

yields the firm's optimal terminal information acquisition,⁶

$$\kappa_1^* = \left[\left(\frac{1 + \hat{\alpha}}{4\beta} \right)^{1/3} - \frac{1}{\sigma_d^2} \right]_+ = \left[\frac{1}{\underline{\sigma}_d^2} - \frac{1}{\sigma_d^2} \right]_+,$$

where $\underline{\sigma}_d^2 \triangleq (4\beta / (1 + \hat{\alpha}))^{1/3}$ denotes the variance threshold. For belief variances σ_d^2 below the threshold $\underline{\sigma}_d^2$ at the beginning of the second period it is in the firm's best interest not to acquire any information prior to taking its terminal operational decisions; see Fig. 2.

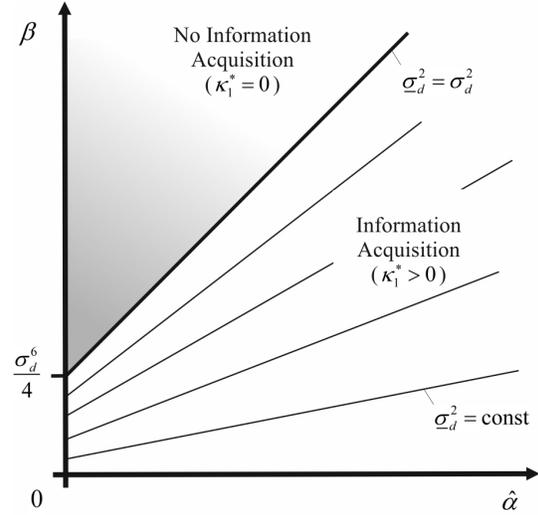


Figure 2: Optimal information acquisition at time $t = 1$.

Remark 2 (Information-Cost Threshold). The expression for κ_1^* implies the cost threshold $\bar{\beta}_d \triangleq (1 + \hat{\alpha}) \sigma_d^6 / 4$, so that nontrivial market intelligence is economical if and only if the firm's information-acquisition cost is such that $\beta \in (0, \bar{\beta}_d)$.

Substituting the firm's terminal information-acquisition policy into its expected payoffs yields

$$\hat{\Pi}_1(\kappa_1^*) = (1 + \hat{\alpha}) \left(\frac{(\mu_d - c)^2}{4} + \frac{\sigma_d^2}{4} \left[1 - \frac{\sigma_d^2}{\sigma_d^2} \right]_+^3 \right).$$

The last expression neatly illustrates the separate effects of information acquisition and of advertising in the company's quasi-static decision problem that it faces in the last period. Information acquisition improves the effectiveness of both pricing and advertising, resulting in an additive improvement $\Delta_1 \geq 0$ of the no-advertising-no-information profit $(\mu_d - c)^2 / 4$. This augmented profit is scaled by the factor $(1 + \hat{\alpha})$ which contains the advertising effectiveness $\hat{\alpha} \geq 0$. In other words,

$$\hat{\Pi}_1^* \triangleq \hat{\Pi}_1(\kappa_1^*) = (1 + \hat{\alpha}) (\Pi_1(0, p_1^m, 0) + \Delta_1),$$

where $p_1^m \triangleq (\mu_d + c)/2$ is the standard solution of the monopoly pricing problem [21] with linear demand curve without advertising and without information acquisition, while

$$\Delta_1 \triangleq \frac{\sigma_d^2}{4} \left[1 - \frac{\sigma_d^2}{\sigma_d^2} \right]_+^3$$

is the information-related (pre-advertising) profit increment. The firm's expected monopoly profit without advertising and information acquisition is

$$\Pi_1(0, p_1^m, 0) = \frac{(\mu_d - c)^2}{4}.$$

The expected value of sample information (EVSI) about the demand at time $t = 1$ is the difference between the firm's optimal time-1 profit and its profit without information acquisition:

$$\text{EVSI}_1 = \hat{\Pi}_1^* - \hat{\Pi}_1(0) = (1 + \hat{\alpha})\Delta_1.$$

This expression implies a complementarity between advertising and information acquisition, in the sense that a better advertising effectiveness also increases the value of information, as noted in the following remark.

Remark 3 (Parameter Sensitivity). The firm's optimal time-1 profit $\hat{\Pi}_1^*$ is increasing in the advertising effectiveness $\hat{\alpha}$ and (weakly) decreasing in the marginal information-acquisition cost β , with corresponding sensitivities

$$\frac{\partial \hat{\Pi}_1^*}{\partial \hat{\alpha}} = \Pi_1(0, p_1^m, 0) + \frac{\sigma_d^2}{4} \left[1 - \frac{\sigma_d^2}{\sigma_d^2} \right]_+^2 > 0,$$

and

$$\frac{\partial \hat{\Pi}_1^*}{\partial \beta} = \frac{\partial(\text{EVSI}_1)}{\partial \beta} = - \left[\frac{1}{\sigma_d^2} - \frac{1}{\sigma_d^2} \right]_+^2 \leq 0,$$

respectively.⁷

Remark 4 (Perfect Information). In the fortunate case where $\beta \rightarrow 0^+$, information-acquisition costs vanish and the firm can obtain perfect demand information at no cost. This implies an expected value of perfect information (EVPI),

$$\text{EVPI}_1 = \frac{(1 + \hat{\alpha})\sigma_d^2}{4}.$$

The latter is an upper bound for the expected value of sample information, EVSI_1 , determined earlier.

⁷Moreover, the firm's optimal time-1 profit is (weakly) convex in $\hat{\alpha}$ and β , respectively. However, the cross-effect of the two parameters is (weakly) negative, $\partial^2 \hat{\Pi}_1^* / (\partial \hat{\alpha} \partial \beta) \leq 0$, so an increase in the advertising effectiveness increases the firm's sensitivity with respect to *decreases* in the marginal information-acquisition cost.

4 Dynamic Policy

We are now ready to consider the firm's dynamic two-period decision problem, as formulated in Sec. 2. The firm's second-period problem amounts to the quasi-static profit-maximization problem discussed in the last section. Pursuing the dynamic-programming strategy of backward induction, in what follows we first look at the firm's belief update from the observed demand realization. The latter allows learning even if no information is acquired from the outside. We then turn our attention to the operational decisions, before we solve the initial information-acquisition problem.

4.1 Learning from Demand Observation

In addition to externally furnished costly market intelligence, the firm usually has access to the information revealed by the market's response to its product offering, which in our case is modulated by the product price p_0 and the company's advertising activity a_0 . Indeed, the firm's realized demand at the end of the first period is

$$D_0 = \theta + a_0 - p_0 + \nu_0,$$

where θ and ν_0 are the realizations of the unknown (and therefore random) market potential $\tilde{\theta}$ and the unsystematic noise $\tilde{\nu}_0$, respectively. Similar to the belief update in the terminal period, the belief update about $\tilde{\theta}$ is informed by the realization

$$s_d = D_0 - (a_0 - p_0)$$

of an informative signal, with conditional distribution

$$\tilde{s}_d | \theta \sim \mathcal{N}(\theta, \sigma_\nu^2).$$

That is, the confidence related to the firm's demand observation is given by

$$\kappa_\nu \triangleq 1/\sigma_\nu^2.$$

As before, we obtain the firm's updated estimate for the market potential as weighted average,

$$\mu_d(s_d) = \lambda_d s_d + (1 - \lambda_d)\mu_0,$$

where $\lambda_d \triangleq \kappa_\nu \sigma_0^2 / (1 + \kappa_\nu \sigma_0^2)$. Meanwhile, the updated variance is given by (see footnote 5)

$$\sigma_d^2 = \frac{\sigma_0^2}{1 + \kappa_\nu \sigma_0^2} = \frac{\sigma_0^2 \sigma_\nu^2}{\sigma_0^2 + \sigma_\nu^2}.$$

The expressions for the demand-induced updated belief (μ_d, σ_d^2) can now be substituted into the terminal profit $\hat{\Pi}_1^*$ obtained in the preceding section:

$$\bar{\Pi}_1^* = (1 + \hat{\alpha}) \left[\frac{(\mu_0 - c)^2}{4} + \Delta_0 \right],$$

where

$$\Delta_0 \triangleq \frac{\sigma_0^6/4}{(\sigma_0^2 + \sigma_\nu^2)^2} + \frac{\sigma_0^2 \sigma_\nu^2/4}{\sigma_0^2 + \sigma_\nu^2} \left[1 - \frac{\sigma_0^2 + \sigma_\nu^2}{\sigma_0^2} \frac{\sigma_d^2}{\sigma_\nu^2} \right]_+^3.$$

The first term on the right-hand side represents the firm's EVSI for the demand observation, which is increasing in its belief uncertainty σ_0^2 . Viewed from the end of the first period, the second term on the right-hand side captures the fact that information is anticipated to be acquired in the following terminal period if and only if

$$\frac{1}{\sigma_0^2} + \frac{1}{\sigma_\nu^2} < \frac{1}{\sigma_d^2}.$$

In particular, the “target variance” σ_d^2 needs to be strictly smaller than $\min\{\sigma_0^2, \sigma_\nu^2\}$, that is, smaller than the variance of the unsystematic noise *and* smaller than the residual variance after external information acquisition at the beginning of the first period.

Remark 5 (Information-Value Consistency). To examine the consistency of the expected information values Δ_0 and Δ_1 before and immediately after the firm's observing demand, consider the extreme cases of perfect inference (where $\sigma_\nu \rightarrow 0^+$) and uninformative inference (where $\sigma_\nu \rightarrow \infty$), respectively. (i) With noiseless demand information one obtains $\lim_{\sigma_\nu \rightarrow 0^+} \Delta_0 = \sigma_0^2/4$ and $\Delta_1 = 0$. Moreover, the expected value of perfect information after the demand update vanishes (i.e., $\text{EVPI}_1 = 0$); earlier, the expected value of perfect demand information is $\text{EVPI}_0 = (1 + \hat{\alpha})\sigma_0^2/4$, analogous to the earlier expression for EVPI_1 with acquisition of perfect information (for $\kappa_1 \rightarrow \infty$) in period 1 (with σ_d^2 replaced by σ_0^2). (ii) With uninformative demand information it is $\lim_{\sigma_\nu \rightarrow \infty} \Delta_0 = \Delta_1|_{\sigma_d=\sigma_0}$: the firm's beliefs are transmitted “unfiltered” from period 0 to period 1, in the sense that $(\mu_d, \sigma_d^2) = (\mu_0, \sigma_0^2)$.

The expected time-0 value of future information acquisition Δ_0 is increasing in the firm's belief uncertainty $\sigma_0^2 \in [0, \infty)$, and

$$\Delta_1|_{\sigma_d=\sigma_0} = \frac{\sigma_0^2}{4} \left[1 - \frac{\sigma_d^2}{\sigma_0^2} \right]_+^3 \leq \Delta_0 \leq \frac{\sigma_0^2}{4},$$

with the bounds obtained in Remark 5.

⁸The internal information may be provided by the accounting department or a market-intelligence group tasked with evaluating demand.

4.2 Operational Decisions

Based on the firm's “free” belief update, by actively using its own internal market intelligence⁸ the firm's profit-maximization problem, after initial information acquisition (κ_0), is similar to the problem in the terminal period (our “quasi-static setting” in Sec. 3):

$$\max_{a_0, p_0} \{ (p_0 - c)(\mu_0 + a_0 - p_0) - (\alpha/2) a_0^2 \},$$

with the unique solution

$$\begin{aligned} a_0^* &= \hat{\alpha}(\mu_0 - c), \\ p_0^* &= (1 + \hat{\alpha}) \frac{\mu_0 + c}{2}, \end{aligned}$$

as before. As a result, the optimal expected discounted profit (after initial information acquisition) becomes

$$\hat{\Pi}_0(\kappa_0) = \mathbb{E}_{\mu_0} \left[\frac{(1 + \hat{\alpha})(\mu_0 - c)^2}{4} + \delta \bar{\Pi}_1^* \right] - \beta \kappa_0^2,$$

given the firm's initial belief (μ, σ^2) .

4.3 Information-Acquisition Decisions

Consider now the problem of external information acquisition at the outset. As before, Bayesian updating of the firm's initial belief (μ, σ^2) using the informative external signal realization s_0 at the chosen confidence level κ_0 yields

$$\mu_0(s_0) = \lambda_0 s_0 + (1 - \lambda_0) \mu,$$

where $\lambda_0 = \kappa_0 \sigma^2 / (1 + \kappa_0 \sigma^2)$, and

$$\sigma_0^2 = \frac{\sigma^2}{1 + \kappa_0 \sigma^2}.$$

Substituting this into the last formula of the previous subsection leads to an expected discounted profit,

$$\begin{aligned} \hat{\Pi}_0(\kappa_0) &= \frac{(1 + \hat{\alpha})(1 + \delta)}{4} \left[(\mu - c)^2 + \frac{\kappa_0^2 \sigma^6}{(1 + \kappa_0 \sigma^2)^2} \right] \\ &\quad + (1 + \hat{\alpha}) \delta \hat{\Delta}_0(\kappa_0) - \beta \kappa_0^2, \end{aligned}$$

where the expected future informational gain,

$$\hat{\Delta}_0(\kappa_0) \triangleq \Delta_0|_{\sigma_0=\sigma^2/(1+\kappa_0\sigma^2)},$$

is now evaluated with respect to the firm's initial belief (μ, σ^2) . As discussed earlier, the latter is tightly bounded:

$$\hat{\Delta}_{0,\min} \leq \hat{\Delta}_0 \leq \hat{\Delta}_{0,\max},$$

where $\hat{\Delta}_{0,\max} \triangleq \sigma^2 / (4(1 + \kappa_0 \sigma^2))$ and

$$\hat{\Delta}_{0,\min} \triangleq \frac{\sigma^2/4}{1 + \kappa_0 \sigma^2} \left[1 - (1 + \kappa_0 \sigma^2) \frac{\sigma_d^2}{\sigma^2} \right]_+^3.$$

While the time-0 information-acquisition problem does not have an explicit solution, closed-form expressions for “boundary solutions” are available when $\hat{\Delta}_0$ is replaced by one of its bounds.⁹ The optimal time-0 information acquisition κ_0^* satisfies

$$\kappa_0^* \in \arg \max_{\kappa_0 \geq 0} \left\{ \frac{(1 + \delta)\kappa_0^2 \sigma^6}{4(1 + \kappa_0 \sigma^2)^2} + \delta \hat{\Delta}_0(\kappa_0) - \frac{\beta \kappa_0^2}{1 + \hat{\alpha}} \right\},$$

and can be determined numerically using standard methods.¹⁰

The belief update in the future depends on the firm’s choice of κ_0 . Indeed, the variance after the demand observation is

$$\sigma_d^2 = \frac{\sigma_0^2}{1 + \kappa_\nu \sigma_0^2} \Big|_{\sigma_0^2 = \frac{\sigma^2}{1 + \kappa_0 \sigma^2}} = \frac{\sigma^2}{1 + (\kappa_0 + \kappa_\nu) \sigma^2},$$

which illustrates the fact that the confidence levels of subsequent belief updates are additive.¹¹ In this manner, the firm’s optimal time-1 information-acquisition decision can be written as a function of its time-0 information acquisition:

$$\kappa_1^*(\kappa_0) = \left[\frac{1}{\frac{\sigma^2}{\sigma_d^2}} - \frac{1}{\sigma^2} - (\kappa_0 + \kappa_\nu) \right]_+.$$

This means that the firm acquires information in the second period if and only if its accumulated confidence does not exceed the target confidence level $\kappa_d \triangleq 1/\sigma_d^2$:

$$\kappa_1^*(\kappa_0) > 0 \Leftrightarrow \kappa_0 + \kappa_\nu + \frac{1}{\sigma^2} > \kappa_d.$$

More specifically, the firm’s confidence at the end of the horizon (see footnote 11), after having observed all available information, will be equal to

$$\frac{1}{\sigma_1^2} = \max \left\{ \kappa_d, \kappa_0 + \kappa_\nu + \frac{1}{\sigma^2} \right\},$$

where $\kappa_d = ((1 + \hat{\alpha})/(4\beta))^{1/3}$. In the second period, provided there is an incentive to acquire a positive amount of information (so $\kappa_1^*(\kappa_0) > 0$), the firm considers κ_1 and κ_0 to be perfect substitutes:

$$\frac{d\kappa_1^*(\kappa_0)}{d\kappa_0} = -\mathbf{1}_{\{\kappa_0 < \kappa_d - \kappa_\nu - (1/\sigma^2)\}}.$$

⁹As roots of fourth-degree polynomials these expressions are, however, not particularly informative, and have therefore been omitted.

¹⁰Here we use the fact that the solution to a maximization problem is homogeneous of degree zero. A solution to this nonconvex problem exists by the Weierstrass theorem ([3], p. 540); in this context, one can restrict attention to a compact interval (i.e., to $\kappa_0 \in [0, \bar{\kappa}_0]$ for some $\bar{\kappa}_0 \in (0, \infty)$) because the objective function is coercive for large values of κ_0 ([3], p. 539). For global optimality conditions, see [24].

¹¹The firm’s belief variance, after having observed the signal \tilde{s}_1 with confidence κ_1 , is $\sigma_1^2 = \sigma^2/(1 + (\kappa_0 + \kappa_1 + \kappa_\nu) \sigma^2)$.

¹²A sufficient condition for submodularity is that $\partial^2 \hat{\Delta}_0(\kappa_0 | \kappa_1^* = 0)/(\partial \kappa_0 \partial \kappa_\nu) = -(3/2)\kappa_\nu \sigma^8/(1 + (\kappa_0 + \kappa_\nu) \sigma^2)^4 < 0$.

¹³One possible threshold is $\kappa_\nu^* \triangleq \max\{0, \kappa_d - (1/\sigma^2) - \kappa_0^*\}$, where $\kappa_0^* \triangleq \lim_{\kappa_\nu \rightarrow \infty} \kappa_0^*$ is the asymptotic confidence of time-0 information collection as $\kappa_\nu \rightarrow \infty$. Note that even for $\kappa_\nu \rightarrow \infty$, information acquisition in the first period may be optimal ($\kappa_0^* > 0$) because it carries an immediate benefit for the firm’s operational decisions before demand is observed.

¹⁴In that case, one would obtain $\kappa_0^* = \left[\frac{1+\delta}{\sigma_d^2} - \frac{1}{\sigma^2} \right]_+$ which would imply that $\kappa_1^*(\kappa_0^*) = 0$ because $\frac{1}{\sigma_1^2} = \max \left\{ \frac{1}{\sigma^2}, \frac{1+\delta}{\sigma_d^2} \right\} \geq \frac{1}{\sigma_d^2}$.

In case the firm does not want to acquire any information in the second period, so $\kappa_1^*(\kappa_0) = 0$, the expected future informational gain,

$$\hat{\Delta}_0(\kappa_0 | \kappa_1^* = 0) = \frac{1}{4} \left(\frac{\kappa_\nu \sigma^2}{1 + (\kappa_0 + \kappa_\nu) \sigma^2} \right)^2 \frac{\sigma^2}{1 + \kappa_0 \sigma^2},$$

stems only from the upcoming costless observation of demand. As κ_ν increases, the firm’s incentive to acquire information in the first period decreases. Indeed, since $\hat{\Delta}_0(\kappa_0 | \kappa_1^* = 0)$ is submodular in (κ_0, κ_ν) ,¹² κ_0^* is decreasing in the confidence κ_ν with respect to demand observations, provided that $\kappa_1^* = 0$. On the other hand, if κ_ν becomes sufficiently large, it will be best for the firm to stop acquiring information in the second period: there is a threshold confidence level κ_ν^* , so that¹³

$$\kappa_\nu \geq \kappa_\nu^* \Rightarrow \kappa_1^* = 0.$$

On the other hand, if the demand observation is uninformative (i.e., if $\kappa_\nu = 0$), then

$$\hat{\Delta}_0(\kappa_0) = \frac{\sigma^2/4}{1 + \kappa_0 \sigma^2},$$

which leads to an optimal time-0 confidence level κ_0^* that is less than the solution one would obtain conditional on $\kappa_1^* = 0$ when the firm’s information-acquisition problem would be isomorphic to its information-acquisition problem at time 1, up to the factor $(1 + \delta)$.¹⁴ The following main result largely summarizes the preceding discussions.

Theorem (Optimal Information Acquisition). *Consider variations of the firm’s (nonnegative) confidence κ_ν in demand observations.*

(i) *If the precision κ_0 of the information collected in the first period (weakly) exceeds the optimal amount κ_0^* , then information acquisition in the second period is never optimal (i.e., $\kappa_1^*(\kappa_0) = 0$); in particular: $\kappa_1^*(\kappa_0^*) = 0$.*

(ii) *If, for some suboptimal κ_0 implemented in the first period, information collection is optimal in the second period (i.e., $\kappa_1^*(\kappa_0) > 0$), then κ_ν is a direct substitute for κ_1^* (in the sense that $\partial \kappa_1^*/\partial \kappa_\nu = -1$), as long as κ_ν stays below the threshold $\kappa_\nu^*(\kappa_0) \triangleq \kappa_d - \kappa_0$ beyond which $\kappa_1^*(\kappa_0)$ vanishes;*

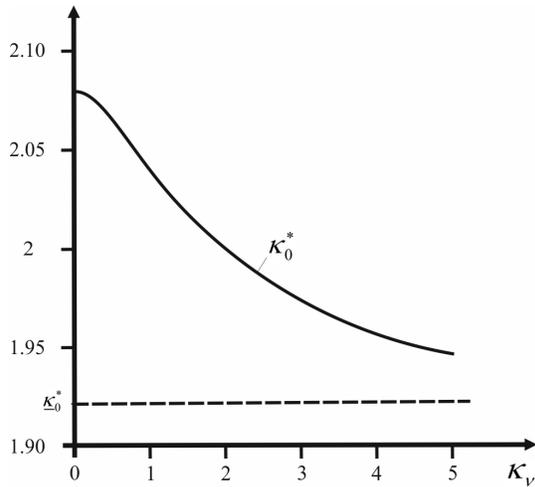


Figure 4: Unconstrained first-period information acquisition κ_0^* as a function of κ_ν in the numerical example.

Assuming that the firm’s initial belief about demand is given by (μ, σ^2) with expectation $\mu = 100$ and standard deviation $\sigma = 30$ (so the coefficient of variation is 30%), the initial confidence level is $1/\sigma^2 \approx 0.0011$. This is three orders of magnitude below the target confidence κ_d , which suggests that active information acquisition may be desirable—even if it is costly. Fig. 4 shows the optimal first-period information acquisition policy as a function of κ_ν in the absence of any info-budget constraint. As shown earlier, the second-period information acquisition vanishes in that case, so $\kappa_1^* = 0$ for all $\kappa_\nu \geq 0$.

Budget-Constrained Information Acquisition. Suppose now that the firm has a positive info-budget B so that $\kappa_B < \kappa_0^*$ at least for small values of κ_ν . Then qualitatively the optimal-information acquisition policy is as in Fig. 3. Specifically, for our numerical values we obtain a second-period information acquisition of the form

$$\kappa_1^*(\hat{\kappa}_1^*) \approx \min \left\{ [1.7089 - \hat{\kappa}_1^* - \kappa_\nu]_+, \sqrt{10B} \right\},$$

for all $\kappa_\nu \geq 0$, where $\hat{\kappa}_0^*$ is the truncated unconstrained first-period information acquisition, as noted in Sec. 4.4.

5 Conclusion

The ability to advertise allows the firm to scale up its profitability. It also increases the returns on additional

information about its uncertain demand. In the absence of spending constraints and given a stationary random demand, we obtain that information should be acquired upfront so as to be used multiple times. In that sense, information can behave like an investment good that pays a periodic rent through the operational improvements it enables.

The firm’s short-term interest in information exhibits a threshold; its goal in the quasi-static setting (at the end of the finite time horizon) is to achieve a target confidence level, as if it was following an (s, S) -inventory policy [1, 20]. When actively considering the future, such as at the initial time ($t = 0$), the firm rationally anticipates future benefits of acquired information in addition to the expected effects of any costless complementary signals observed later, for example, as a result of intermittent demand observations via the data stream from its accounting department. It is interesting that threshold effects arise despite the fact that information-acquisition cost is quadratic in the signal precision.¹⁶ To achieve a very high precision the firm can expect a substantial information-acquisition expense which may well violate internal budget constraints.¹⁷ The latter may lead to the (constrained) smoothing of information consumption over time instead of full upfront information collection in the unconstrained case.

Our work fits in with an increasing interest in the interaction of costly information and operational decisions, for example in inventory control problems [7] or in pricing problems [5]. The fundamental difference between advertising and information as drivers of profitability is that advertising tends to enlarge the pie, whereas information renders the size of the pie and the effect of the operational decisions less uncertain. Information is therefore not an end in itself,¹⁸ but rather a complement to managerial actions.

An interesting direction for future research may be an extension of this framework to situations where the quality of the costly information depends on the operational actions (e.g., the demand information gathered from an advertising campaign or by means of a price discount), giving rise to “monopoly experimentation.” Then the firm would need to trade off the possible lack in optimality of its operational experiments in the short-run against the increased confidence in its demand information in the medium-run together with the improved operational decisions the firm will be able to take.

¹⁶The cost of external information is naturally convex because it becomes more and more difficult to increase the confidence in demand information, e.g., by enlarging the sample further, beyond an established panel. The threshold nature of information acquisition also obtains in a linear-quadratic infinite-horizon setting with Gaussian noise and linear cost of information acquisition [25].

¹⁷Even if the external market intelligence provides a noise-free signal (of infinite confidence), the firm continues to experience uncertainty from period to period because of the persistent unsystematic noise in the system. In particular, the perfect adaptation of price and advertising to actually realized demand conditions in any given period remains a zero-probability event.

¹⁸This is notwithstanding investors’ having a vested interest to reduce profit uncertainty, all else equal; see also footnote 4.

References

- [1] ARROW, K.J., HARRIS, T., MARSCHAK, J. (1951) "Optimal Inventory Policy," *Econometrica*, Vol. 19, No. 3, pp. 250–272.
- [2] BELLMAN, R.E. (1954) "The Theory of Dynamic Programming," *Bulletin of the American Mathematical Society*, Vol. 60, No. 6, pp. 503–516.
- [3] BERTSEKAS, D.P. (1995) *Nonlinear Programming*, Athena Scientific, Belmont, MA.
- [4] BLACKWELL, D. (1951) "Comparison of Experiments," *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, University of California Press, Berkeley, CA, pp. 93–102.
- [5] BONATTI, A. (2011) "Menu Pricing and Learning," *American Economic Journal: Microeconomics*, Vol. 3, No. 3, pp. 124–163.
- [6] CIPRIANO, L.E., WEBER, T.A. (2018) "Population-Level Intervention and Information Collection in Dynamic Healthcare Policy," *Health Care Management Science*, Vol. 21, No. 4, Forthcoming. [DOI 10.1007/s10729-017-9415-5]
- [7] DECROIX, G.A., MOOKERJEE, V.S. (1997) "Purchasing Demand Information in a Stochastic-Demand Inventory System," *European Journal of Operational Research*, Vol. 102, No. 1, pp. 36–57.
- [8] DEGROOT, M. (1962) "Uncertainty, Information, and Sequential Experiments," *Annals of Mathematical Statistics*, Vol. 33, No. 2, pp. 404–419.
- [9] DORFMAN, R., STEINER, P.O. (1954) "Optimal Advertising and Optimal Quality," *American Economic Review*, Vol. 44, No. 5, pp. 826–836.
- [10] GAL-OR, E. (1988) "The Advantages of Imprecise Information," *Rand Journal of Economics*, Vol. 19, No. 2, pp. 266–275.
- [11] HIRSHLEIFER, J. (1971) "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, Vol. 61, No. 4, pp. 561–574.
- [12] HOLTHAUSEN, D.M., ASSMUS, G. (1982) "Advertising Budget Allocation under Uncertainty," *Management Science*, Vol. 28, No. 5, pp. 487–499.
- [13] KIHILSTROM, R.E. (1984) "A Bayesian Exposition of Blackwell's Theorem on the Comparison of Experiments," in: Boyer, M., Kihlstrom, R.E. (Eds.) *Bayesian Models in Economic Theory*. North-Holland, Amsterdam, NL, pp. 13–31.
- [14] MCCARDLE, K.F. (1984) "Information Acquisition and the Adoption of New Technology," *Management Science*, Vol. 31, No. 11, pp. 1372–1389.
- [15] MOORE, J.C., WHINSTON, A.B. (1986) "A Model of Decision-Making with Sequential Information Acquisition (Part 1)," *Decision Support Systems*, Vol. 2, No. 4, pp. 285–307.
- [16] MOORE, J.C., WHINSTON, A.B. (1987) "A Model of Decision-Making with Sequential Information Acquisition (Part 2)," *Decision Support Systems*, Vol. 3, No. 1, pp. 47–72.
- [17] MOSCARINI, G., SMITH, L. (2001) "The Optimal Level of Experimentation," *Econometrica*, Vol. 69, No. 6, pp. 1629–1644.
- [18] NERLOVE, M., ARROW, K.J. (1962) "Optimal Advertising Policy under Dynamic Conditions," *Economica*, Vol. 29, No. 114, pp. 129–142.
- [19] RAIFFA, H., SCHLAIFER, R. (1961) *Applied Statistical Decision Theory*, Division of Research, Graduate School of Business Administration, Harvard University, Boston, MA.
- [20] SCARF, H.E. (1960) "The Optimality of (s, S) Policies in the Dynamic Inventory Problem," in: Arrow, K.J., Karlin, S., Suppes, P. (Eds.) *Mathematical Methods in the Social Sciences*, Stanford University Press, Stanford, CA, Chapter 11.
- [21] TIROLE, J. (1988) *The Theory of Industrial Organization*, MIT Press, Cambridge, MA.
- [22] WALD, A. (1947) "Foundations of a General Theory of Sequential Decision Functions," *Econometrica*, Vol. 15, No. 4, pp. 279–313.
- [23] WEBER, T.A., CROSON, D.C. (2003) "Selling Less Information for More: Garbling with Benefits," *Economics Letters*, Vol. 83, No. 2, pp. 165–171.
- [24] WEBER, T.A. (2017) "Global Optimization on an Interval," *Journal of Optimization Theory and Applications*, Vol. 182, No. 2, pp. 684–705.
- [25] WEBER, T.A., NGUYEN, V.A. (2018) "A Linear-Quadratic Gaussian Approach to Dynamic Information Acquisition," *European Journal of Operational Research*, Vol. 270, No. 1, pp. 260–281.