Market Contingent Managerial Hierarchies

Kieron Meagher*, Hakan Orbay†, Timothy Van Zandt‡

18 January 2008

(Preliminary and incomplete.)

Abstract

We present a formal model of how the size of a firm's administrative staff depends on the environment in which it operates. We consider a monopolist that introduces a new product in the face of market uncertainty. In a model that incorporates an information-processing view of the managerial hierarchies into a product choice problem, we relate optimal size of organizations to the speed of change in market conditions and to more standard parameters of demand. We show that faster-changing environments lead to smaller management structures, which is consistent with anecdotal evidence of both flattening and decentralization by firms in response to faster-changing consumer tastes. When the environment is changing more quickly, it is better to aggregate less information to reduce the inexorable delay that occurs even when information processing is decentralized. This occurs both for positive managerial wages and in a benchmark model of zero managerial wages. An improvement of information technology, which we take to be an increase in the speed at which each manager can process information, has a similar effect, though it is more subtle. On the one hand, such a speedup is like a decrease in the speed of change of the environment, so that the firm processes more information. On the other hand, any fixed amount of information can be processed with fewer "person-hours" and hence with fewer managers. When the managerial wage is zero, these two effects exactly cancel each other, so that management size does not depend on information technology. However, when the managerial wage is positive, the firm chooses to economize on managers when they become more productive, so that an improvement in information technology leads to a shrinking of the management size. In other results, we show that the change in the optimal size of the organization is non-monotonic in both the dispersion of consumer preferences and the "transport cost" parameter, while the optimal size increases as the consumers' value of the product increases.

*School of Economics, University of New South Wales, Sydney, Australia. Email: k.meagher@unsw.edu.au

†Graduate School of Management, Sabançi University, Istanbul, Turkey. Email: hakan@sabanciuniv.edu

‡INSEAD, Fontainebleau, France. Email:tvz@insead.edu

1. Introduction

Improvements in communication, transportation, and information technology create an environment in which the world changes at an ever faster pace. Keeping up with the fast pace of change in consumer market has lead firms to adopt management structures that are nimble and responsive to market conditions.

We consider this question by shifting the analysis away from the issue of centralization versus decentralization and focusing instead on the size of the management structure. Optimally smaller management structures involve fewer steps in the decision-making process, and fewer steps mean faster decisions. Fewer steps also mean less well-informed decisions, and this trade-off between speed and quality in decision-making is the fundamental intuition underlying our model.

Our model considers the organizational and product-related choices of a monopolist that is developing a new product. The examples that typify the situation being considered are the development of a feature film, a new software package or the design of a new car. For such complex heterogeneous products, producers compete not just on price but also on the characteristics of their products.

Consider the feature-film example: the preferences of the film going public (as in many other markets) are not directly observable and change over time. However, market research and other information regarding public tastes can and does influence the content of films during the production and, especially, the pre-production process. How tolerant are the public of violence? Do people still want a romantic story? Which books are people reading? The answers to such questions reflect the preferences of the consumers and have an impact on the selection of which project to produce, details of the screenplay, which stars to cast, and so on. The preferences change over time and are not directly observable; hence the producers resort to collecting information through market research.

In our model, we distinguish two types of uncertainty. First, any product will attract only a subset of consumers, which is determined by the dispersion of consumer preferences. Secondly, consumer preferences may change over time in an unpredictable manner. We model these uncertainties by assuming that consumers are normally distributed with an unknown mean that changes over time stochastically. Thus, the monopolist has to collect information from the consumers to make a decision on the type of product to launch.

We view the managerial structure as an information processing apparatus that collects and aggregates information. Such information processing models of firms were first introduced by Radner (1993) and later extended by others.[1] The critical assumption in these models is that processing requires time, which is increasing in the amount of information processed and, within limits, decreasing in the size of the organization. Because of our focus on the size of the management rather than on its structure, we take a reduced-form model of this information processing, stating axiomatically a relationship between the amount of information collected, the speed of information processing, the size of the organization, and the management cost.

1. See Van Zandt (1998a) for a survey.

In our model of organizational choice, the monopolist chooses how much information to collect, that is, the sample size of the market research. As the delay in processing increases with sample size, the monopolist faces a trade-off in this choice: While more information improves the description of the historical situation, the associated longer delay reduces the value of information due to the changing environment.

Consider, then, the effect of an increase in the speed at which the environment changes. The delay penalty is larger, leading the organization to process less information, which in turn leads to smaller managerial size.

Improvements in the speed of processing induce the organization to process more information, much like a decrease in the speed at which the environment changes. It is the relative race between the changing environment and the ability of the management to keep up that matters. However, more information does not mean larger hierarchies in this case because any fixed amount of information can be processed with fewer "person-hours". When the managerial wage is zero (as a benchmark), these two effects—the increase in the amount of information and the increase in managerial productivity—exactly cancel each other, so that management size does not depend on information technology. However, when the managerial wage is positive, the firm chooses to economize on managers when they become more productive, so that an improvement in IT leads to a shrinking of the management size.

Related work by Van Zandt and Radner (2001) and Van Zandt (2003) has identified another effect of improvements of information technology: an increase in the scale of the firm. Combining these results, an improvement of IT may lead to mergers but shedding of managerial jobs, with an ambiguous net effect on the size of the resulting managerial structure.

Another contribution of this paper is to relate more standard characteristics of demand to organizational structure. Through the differentiated product model, we show how product quality, transport costs, and dispersion of consumer tastes also affect the optimal choice of organizational structure.

We develop the model in Section 2. First we consider the underlying decision problem of a firm that introduces a product with uncertainty about consumer tastes, providing a characterization of optimal product location, pricing, and the maximized profit (Section 2.1). Then we describe how forecasts would be made conditional on data about consumer tastes (Section 2.2). We add a reduced-form model of information processing that links sample size, lags in information processing, and managerial costs (Section 2.3). Finally, we describe the organization design problem obtained by integrating these components of the model (Section 2.4). In Section 3, we consider a benchmark model in which managers are costless, and in Section 4 we consider the case of positive managerial wage. Then, in Section 5, we relate also output and production employment to the parameters of the model.

2. The Model

A monopolist launches a new product and must chose the product's attributes and price. An administrative staff of boundedly rational agents makes these decisions by aggregating market information collected before the product launch.

2.1. The market and the static decision problem

As a building block in the organizational model of decision making, we consider the underlying decision problem from a static Bayesian perspective.

We use a standard "location cum transportation costs" parameterization of product attributes and consumer preferences. The valuation of a consumer of type $y \in \mathbb{R}$ for a product with location (attributes) $x \in \mathbb{R}$ is $a - |x - y|$, where $a > 0$. When the price is p, such a consumer purchases the good if and only if her valuation is at least p, that is, if and only if

$$x - (a - p) \le y \le x + (a - p).$$

Each consumer's type is distributed normally with mean μ and variance σ^2. Thus, when the product location is x and the price is p, the probability that she purchases the product is $F(x + (a - p)) - F(x - (a - p))$, where F is the cumulative distribution function (c.d.f.) of $N(\mu, \sigma^2)$. Let $\varphi(\cdot)$ and $\Phi(\cdot)$ denote, respectively, the density and the c.d.f. of the standard normal distribution. Let $\tau = 1/\sigma$ be the precision of y.[2] Then $F(y) = \Phi((y - \mu)\tau)$ and the probability that the consumer purchases the product equals

$$\Phi((x - \mu + (a - p))\tau) - \Phi((x - \mu - (a - p))\tau). \tag{1}$$

Whether there is one such consumer or many and whether the consumer types are correlated or independent, the expected per-capita demand when the monopolist chooses location x and price p is given by equation (1). We assume that the monopolist maximizes expected profit and has no production cost, so that it maximizes expected revenue. Given whatever price it charges, the location should be chosen to maximize the expected demand or probability that each customer purchases the product. As shown in Proposition 1, that implies setting the location to μ.

PROPOSITION 1. *Given any price $p < a$, the unique location that maximizes expected demand is equal to μ.*

(All proofs are in the appendix.)

Taking as given that the location is set optimally to μ and normalizing the total market size to 1, we can write the revenue as a function of p and τ as follows.

$$\begin{aligned} R(p; \tau) &:= p\,(\Phi((a - p)\tau) - \Phi(-(a - p)\tau)) \\ &= p\,(2\Phi((a - p)\tau) - 1) \end{aligned}$$

The second equality follows from the fact that Φ is the c.d.f. of a symmetric distribution with mean 0.

2. We follow the custom in the sciences of defining the precision of a random variable to be the inverse of the standard deviation. In some economics papers, it has been defined as the inverse of the variance. No confusion should arise since the interpretation of our ordinal comparative-statics statements would be the same for either definition.

We explicitly include the parameter τ in the revenue function because this parameter will vary when the precision depends on how much information about consumers is processed. We are interested in how the revenue-maximizing price and the maximum revenue depend on τ.

PROPOSITION 2. *For all $\tau > 0$, there is a unique price $p^*(\tau)$ that maximizes $R(p;\tau)$. It is the solution to*

$$\Phi((a-p)\tau) - p\tau\varphi((a-p)\tau) - 1/2 \;=\; 0 \tag{2}$$

and it lies between $a/2$ and a. Furthermore, $p^*(\tau)$ is strictly increasing, with $\lim_{\tau\downarrow 0} p^*(\tau) = a/2$ *and* $\lim_{\tau\uparrow\infty} p^*(\tau) = a$.

The maximized revenue, $R^*(\tau) = R(p^*(\tau),\tau)$, depends on τ but not on μ. Proposition 3 characterizes this dependence, as follows. Since τ measures the lack of dispersion of consumers or the information about a consumer, $R^*(\tau)$ is increasing. As $\tau \downarrow 0$, consumers become infinitely dispersed or there is no information about the consumer, and hence $R^*(\tau)$ converges to 0. As $\tau \uparrow \infty$, the firm can perfectly adapt the product to the consumer(s) and can perfectly price discriminate, and hence its revenue equals the highest possible valuation a.

PROPOSITION 3. *$R^*(\tau)$ is strictly increasing, with* $\lim_{\tau\downarrow 0} R^*(\tau) = 0$ *and* $\lim_{\tau\uparrow\infty} R^*(\tau) = a$.

The increase in revenue when τ rises is a combination of the increase in price shown in Proposition 2 and an increase in output, shown in Proposition 4.

PROPOSITION 4. *Expected output, given by $Q = 2\Phi((a-p^*)\tau) - 1$, is increasing in τ.*

2.2. The sampling and forecasting problem

The product launch takes place at a specific time which we denote by 0. The monopolist is uncertain about the market conditions at the time of the product launch. It forecasts the market conditions by gathering data beforehand. The monopolist ends up with incomplete information both because the data provide an incomplete snapshot of the market conditions at the time they are gathered and because the market conditions change between when the data are gathered and when the product is launched.

At time $t \in \mathbb{R}$, there is a continuum of heterogeneous consumers whose types are normally distributed with mean $\mu_y(t)$ and variance σ_y^2. The variance σ_y^2 does not change but the mean evolves according to

$$d\mu_y(t) := \sigma_\varepsilon \, d\varepsilon(t)\,, \tag{3}$$

where $\varepsilon(t)$ is the standard Brownian motion and σ_ε is the measure of environmental volatility.

The monopolist's problem is to estimate $\mu_y(0)$. For this purpose, it randomly samples n consumers at time $-L$, observing the type of each consumer sampled. (For example, the monopolist uses focus groups or a marketing survey; consumers are assumed to report their locations truthfully.) In statistical terms, the type of consumer i in the sample corresponds to the observation of the random variable

$$y_i = \mu_y(-L) + \zeta_i\,, \tag{4}$$

where $\zeta_i \sim N(0, \sigma_y^2)$.

We consider a limiting case in which the monopolist has a "diffuse prior". Formally, we could either (a) specify a time t_0 at which $\mu(t_0)$ is known and allow $t_0 \to -\infty$, (b) specify a prior belief on $\mu(-L)$ that is normally distributed and allow the variance of this distribution to increase to ∞, or (c) replace the Brownian motion by a stationary mean-reverting process such as $d\mu_y(t) = -\gamma\mu(t) + \sigma_\varepsilon d\varepsilon(t)$ and let $\gamma \downarrow 0$. In each case, the posterior distribution of $\mu_y(-L)$ conditional on the sample is normal and, in the limit, its mean and variance are $\bar{y}$ and σ_y^2/n, resp., where $\bar{y} = (1/n)\sum_{i=1}^n y_i$ is the sample mean.

Since $\mu_y(0)$ can be written $\mu_y(0) = \mu_y(-L) + e_L$, where e_L is uncorrelated with $\mu_y(-L)$ and has distribution $N(0, L\sigma_\varepsilon^2)$, the posterior distribution of $\mu_y(0)$ is $N(\bar{y}, \sigma_y^2/n + L\sigma_\varepsilon^2)$. We can now use our analysis of the static model in Section 2.1 to characterize the optimal location and pricing decisions. The mean μ in the static model is the posterior $\bar{y}$. The variance σ^2 in the static model is the variance $\sigma_y^2/n+L\sigma_\varepsilon^2$ of the posterior distribution of $\mu(0)$ plus the variance σ_y^2 of the true distribution of types, so that

$$\sigma^2 = (1 + 1/n)\sigma_y^2 + L\sigma_\varepsilon^2. \tag{5}$$

The monopolist thus chooses location $\bar{y}$ and price $p^*(\tau)$, where $\tau = 1/\sigma$ and σ is defined by equation (5).

2.3. Organizational decision processes

We have presented a stylized model of the sampling, estimation, and decision making of a firm, which is of course much simpler than real-life decision problems. Such simplification is not only standard in economic theory but it is also an objective thereof. The model stands merely as a proxy for the real-life decision processes of firms, which are complicated, use soft data, and cannot be fully represented in a mathematical model.

In the real-life decision problem, the analysis of the marketing data (hard and soft) is typically more costly and time consuming than the acquisition of this data. The task may be shared among the members of a marketing research department and outside consultants because it is too large for a single person to do quickly. Because of the complexity, very thorough market research based on extensive data involves a trade-off: the process will yield an accurate characterization of the market conditions that existed when the data were gathered but, because of the time the data analysis takes, the data will be less relevant to the actual market conditions when the product is launched.

To model this trade-off in our formal model, we treat the calculations that must be performed, although they only entail arithmetic, as complicated and time consuming. The organizational choice problem is then one of choosing the sample size, a time frame for processing the sample, and an organizational structure with which to process the sample, based on a model of decision-making technology. Such modeling of information processing has been used, for example, by Radner (1993), Bolton and Dewatripont

(1994), and Van Zandt (1998b) to study the internal structure of organizations. We will instead focus on the more easily measured size of the administrative staff as the endogenous variable we wish to explain. As a consequence, we need not fully specify the way in which information is processed but rather merely a relationship between sample size, processing delay, and managerial size.

The main task in processing the data is to aggregate them. In this model, this simply means summing the n observations. An essential feature of managerial decision making is that it involves distributed computation, meaning that the tasks are shared among multiple administrators. When n data are aggregated via an associative operation and the computation is distributed, the delay has roughly the form $d(n) = \delta(\alpha + \log n)$ and the amount of information processing that must be performed—i.e., the number of hours managers are busy processing information—is proportional to δn. The parameter δ is a measure of how long it takes individual managers to process information. Since the size of an administrative staff is roughly proportional to the the number of hours worked, we use δn as our measure of managerial size. The cost of information processing is $w\delta n$, where w is a measure of the managerial wage. We summarize this specification with the following axiom.

Assumption 1. There are parameters $\delta > 0$, $\alpha \geq 0$, and $w > 0$ such that, if the organization aggregates n data, then the delay is $d(n) = \delta(\alpha + \log n)$, the managerial size is proportional to δn, and the managerial cost is $w\delta n$.

For example, for the computation model of Radner (1993), the minimum delay is approximately $\delta(1 + \log_2 n)$. Consider a simper model, called the PRAM (parallel random access machine). The unit of analysis is operations rather than managers, because of the simplifying assumption that all managers have equal and instantaneous access to all data and partial results at any time. Suppose it takes δ units of time, called a cycle, to aggregate two data or partial results via an associative operation. A PRAM can sum n numbers in $\lceil \log_2 n \rceil$ cycles as follows: The data are divided into $n/2$ pairs (one per manager), which are concurrently summed in one cycle. The $n/2$ partial results are then paired and summed in the next period, and so on. The number of inputs are cut in half in each period, hence it takes $\lceil \log_2 n \rceil$ cycles to have a single partial result left, which is the sum of the n numbers. This process requires $n-1$ operations; hence the total managerial cost is $w\delta(n-1)$ if w is the managerial wage.

2.4. Organizational Choice

The organizational choice problem is to select the number n of inputs processed, which, in turn, determines the delay of the organization and the size of the administrative staff. As the sample size increases, the firm gains more precise information regarding the distribution of the consumers at the time of sampling. However, as the larger sample size also takes longer to process, the value of this precision is reduced due to the stochastic evolution of the mean.

Let $V(n)$ denote the conditional variance of the consumers' types as a function of the sample size n. Following the discussion in Section 2.3, this is given by substituting

$L = d(n) = \delta(\alpha + \log n)$ into equation (5), so that

$$V(n) = (1 + 1/n)\sigma_y^2 + \delta(\alpha + \log n)\sigma_\varepsilon^2.$$

Since the managerial cost is $w\delta n$, the profit—expected revenue minus managerial overhead—as a function of n is

$$\Pi(n) := R^*(V(n)^{-1/2}) - w\delta n.$$

In setting up this model, we have already ignored the fact that information processing involves operations whose discreteness may be reflected in the formula. For example, the delay might more accurately be $\delta\lceil\log_2 n\rceil$. In solving for the optimal sample size, we will also ignore the integer constraint on n. It is by now well known that such approximations do not entail large qualitative errors.

3. Organizations with Costless Managers

In this section, we consider the benchmark case where the processing is costless, meaning that $w = 0$. Therefore, the monopolist chooses n to maximize $R^*(V(n)^{-1/2})$. Since $R^*(\tau)$ is increasing in τ (more information is better), n is chosen to minimize $V(n)$. The first-order condition is

$$V'(n) = -\sigma_y^2/n^2 + \delta\sigma_\varepsilon^2/n = 0, \tag{6}$$

and the solution is $n = \sigma_y^2/(\delta\sigma_\varepsilon^2)$. One can show that $V''(n) = (\delta\sigma_\varepsilon^2)^3/(\sigma_y^2)^2 > 0$ when evaluated at $n = \sigma_y^2/(\delta\sigma_\varepsilon^2)$, and hence this solution is a local minimum. Because there are no other stationary points, the solution is a global minimum. Thus, we have the following proposition.

PROPOSITION 5. *If $w = 0$, then the optimal sample size is $n = \sigma_y^2/(\delta\sigma_\varepsilon^2)$. Organization size is proportional to $\delta n = \sigma_y^2/\sigma_\varepsilon^2$. Thus, organization size is increasing in the dispersion of consumer preferences and decreasing in the speed of change of the environment. It is not affected by information processing speed.*

Observe the effect of an increase in information processing speed. Although sample size increases, this is exactly offset by the fact that each manager can process more information per unit time. *Thus, the overall effect on organization size is neutral.*

Proposition 5 suggests that a decrease in organizational size can be attributed to either an increase in environmental volatility (faster-changing market conditions) or a decrease in dispersion of consumer preferences. The next proposition indicates how these two factors can be distinguished.

PROPOSITION 6. *When the environmental volatility σ_ε^2 increases, both organization size and monopoly rents decrease. When instead the dispersion σ_y^2 of consumer preferences decreases, organization size decreases but monopoly rents increase.*

4. Organizations with Costly Decision Making

Let us now consider the case where decision making is costly, that is, $w > 0$. Our first observation is that organization size is decreasing in the managerial wage.

PROPOSITION 7. *Organization size is decreasing in the managerial wage. That is, if n_1 and n_2 are optimal organization sizes for wages w_1 and w_2, resp., with $w_2 > w_1$, then $n_2 \leq n_1$.*

COROLLARY 1. *Optimal organization size is at most $\sigma_y^2/\sigma_\varepsilon^2$.*

We now turn to the comparative statics presented in Proposition 5 for the case of costless processing.

PROPOSITION 8. *Optimal organization size is decreasing in environmental volatility σ_ε^2.*

The effect of a decrease in δ, i.e., an improvement in information technology, is interesting and subtle. As identified already in the benchmark of zero managerial wage, there are two countervailing effects: (a) the sample size goes up, just as when the environment changes more slowly; (b) the person-hours required to process any fixed amount of information decreases. When the managerial wage is zero, these two effects cancel each other (Proposition 5). When $w > 0$, there is an additional effect: as δ decreases, the firm chooses to economize on the number of managers. The net effect is therefore a decrease in the management size. We state this as a conjecture because, as of this version of the paper, we have only some numerical tests.

CONJECTURE 1. *Management size is increasing in δ.*

This is illustrated in Figure 1.

Let us now consider how the optimal size changes as σ_y^2 increases. The main idea is that organization size is initially increasing and then decreasing in σ_y^2. For small σ_y^2, a very accurate estimate is possible using just one datum, and hence it is not worth gathering more data and thereby increasing the delay. For large σ_y^2, the value of information processing is small, and therefore again a large organization cannot be optimal. For intermediate values of σ_y^2, the optimal organization size may be larger.

PROPOSITION 9. *Fix $w > 0$. There are $\ell, h > 0$ such that* (a) *the optimal sample size is at most* 1 *if $\sigma_y^2 < \ell$ and* (b) *the optimal sample size is* 0 *if $\sigma_y^2 > h$.*

The inverted-U shape relationship between n and σ_y^2 is illustrated in Figure 2.

Finally, because higher intrinsic valuation makes information processing more valuable, optimal size is increasing in a.

PROPOSITION 10. *Optimal size is increasing in the customer valuation parameter a.*

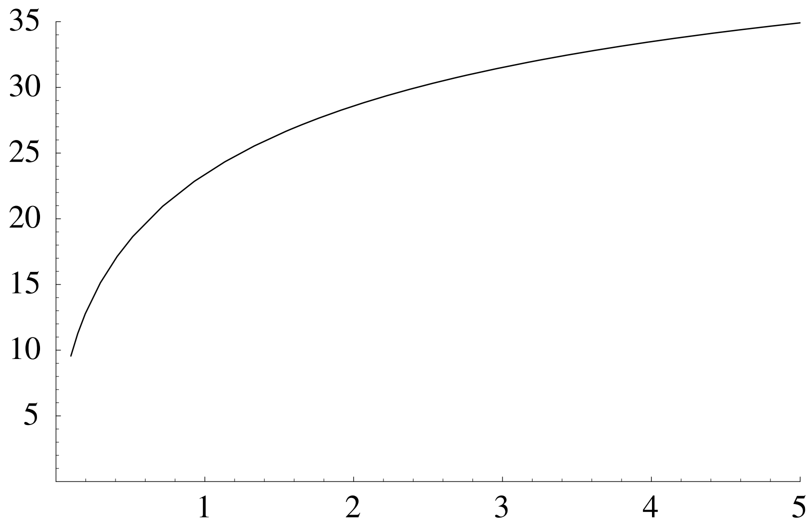

FIGURE 1. Management size as a function of processing delay δ, for $a = 8$, $\sigma_y^2 = 5$, $\sigma_\varepsilon^2 = 0.1$, $\alpha = 1$, and $w = .001$.

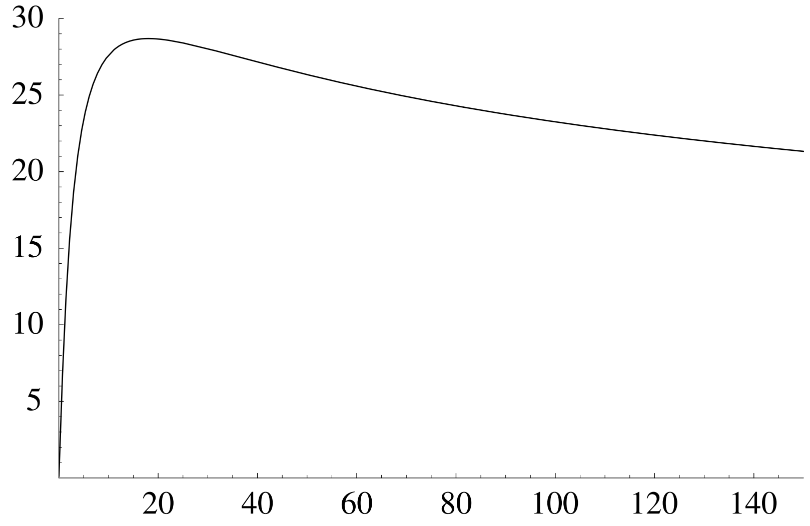

FIGURE 2. Sample size as a function of consumer dispersion σ_y^2, for $a = 8$, $\sigma_\varepsilon^2 = 0.1$, $\delta = 1$, $\alpha = 1$, and $w = .001$.

We can introduce a transport cost T into the model by letting the valuation of a consumer of type y be $a - T|x - y|$. If we insert this parameter in the initial development of the model in Section 2.1, we find that revenue given p, τ, and T is $R(p;\tau/T) = p(2\Phi((a-p)\tau/T) - 1)$. Thus, in the static decision problem, an increase in T has the same effect as an increase in the standard deviation $1/\tau$ of the distribution of consumers. When we consider how n depends on T, we obtain a result analogous to the relationship between n and σ_y^2 shown in Proposition 9, although T and σ_y^2 do not enter into the model in the exact same way.

PROPOSITION 11. *Fix $w > 0$. There are $\ell, h > 0$ such that* (a) *the optimal sample size is at most* 1 *if $T < \ell$ and* (b) *the optimal sample size is* 0 *if $T > h$.*

5. Complementarity between management and factors of production

In Section 4, we examined how managerial employment is affected by changes in the environment in which the monopolist operates. We now turn our attention to two more classical issues, namely, how the monopolist's environment affects output and production employment. By production employees, we mean all those people involved directly in the production of the monopolist's output, rather than those people involved in the managerial decision-making process, who have been our focus up to now.

Our explicit inclusion of the managerial team who formulate the firm's strategy allows us to compare the effects of external environmental factors on both managerial and production employment. For simplicity, production employees are assumed to be homogeneous and a normal input into the production process, so that production employment and output move in the same direction as a result of external factors.

Consider the effect of an increase in the managerial wage. We know from Proposition 7 that this results in a decrease in the sample size. Because the optimal sample size is always in a range where $V'(n) < 0$, this causes a decrease in the precision of the distribution of consumers. That is, an increase in the managerial wage rate leads to the employment of fewer managers and as a result poorer quality decisions. From Proposition 4, poor quality decisions, in the form of a larger forecast error, results in a decrease in output and hence less demand for production labor.

PROPOSITION 12. *An increase in managerial wages causes output and hence production employment to fall.*

More generally, any change in an exogenous parameter that leads to an increase in the total variance will lead to a decrease in both output and employment in the production division. The most obvious example of this comparative statics occurs when the pace of change in the monopolist's environment increases.

PROPOSITION 13. *A faster-changing environment (increase in σ_ε^2) leads to a decrease in the precision of the distribution of consumers, to lower output, and to lower managerial and production employment.*

An improvement in IT causes a reduction in the size of management. Yet it also causes an increase in the total effectiveness of management (i.e., a decrease in total variance) and hence an increase in output and production employment.

PROPOSITION 14. *Smaller δ leads to lower management size and higher production employment.*

6. Ongoing decision making

[*This paper is an ongoing project. This section outlines an extension that has not yet been fully analyzed.*]

We have studied a one-shot decision problem. Suppose instead the product is sold repeatedly. The firm modifies the product characteristics in order to adapt to changing consumer tastes. Decisions must again be computed from past information.

This introduces a modeling dilemma concerning the information available to the firm for making its decisions. In the real world, the firm faces a complicated many-dimensional decision problem that involves gathering information directly about the customers, as in the one-shot problem. However, in our economic model, a firm can magically become fully informed about the one-dimensional state of the market just by observing the actual level of sales.

Rather than make the model much more complicated so that the sales level does not fully reveal the state, we simply restrict the firm from using this information. That is, the only source of information that the firm can use are samples of consumer tastes, as in the one-shot model.

We consider decision procedures of the following form: Every k periods, the firm gathers a new sample of data of size n and aggregates it, as in the one-shot problem. The firm then combines this information with previously aggregated information using an updating rule, thereby incorporating the cumulative information into its decision about the product location.

First consider which kind of decision rules are statistically optimal, in the sense that they maximize the profit conditional on the information they use. Owing to our statistical assumptions, this requires only a simple updating rule (an example of a Kalman filter).

Suppose first that the aggregation of each sample has no computational delay. Let t be a period in which a sample is taken, let $Y(t) = \{y_{1t}, \ldots, y_{nt}\}$ be the sample gathered that period, and let $Y^p(t)$ denote all previous samples. Let $\hat{\mu}^n_y(t) = E[\mu_y(t)|Y(t)]$ and $\hat{\mu}^p_y(t) = E[\mu_y(t)|Y^p(t)]$; let Σ^n and Σ^p be the respective mean-squared errors of these estimates. The errors of these estimates are uncorrelated. (The error for $\hat{\mu}^n_y(t)$ is the average sample error, i.e., deviation of sampled consumer types from the mean, and the error for $\hat{\mu}^p_y(t)$ is a function of previous sample errors and innovations.) Therefore, from the projection formulas,

$$\hat{\mu}_y(t) := E[\mu_y(t)|Y(t), Y^p(t)] = \frac{\Sigma^n}{\Sigma^p + \Sigma^n}\hat{\mu}^p_y(t) + \frac{\Sigma^p}{\Sigma^p + \Sigma^n}\hat{\mu}^n_y(t)\,, \tag{7}$$

and the mean-squared error, i.e., the conditional variance of $\mu_y(t)$, is $(\Sigma^p\Sigma^n)/(\Sigma^p + \Sigma^n)$.

For period s as far back as $t-k$ (when the previous sample was gathered), $\hat{\mu}_y^p(t)$ is also the estimate of $\mu_y(s)$. Furthermore, as explained in Section 2.2, $\hat{\mu}_y^n(t)$ is the sample average of $Y(t)$. Hence, $\hat{\mu}_y(t)$ is calculated by summing $\bar{y}(t) = (1/n)\sum_{i=1}^n y_{it}$ and then averaging

$$\hat{\mu}_y(t) = (1-\alpha)\hat{\mu}_y(t-k) + \alpha\bar{y}(t)\,, \tag{8}$$

where $\alpha = \Sigma^p/(\Sigma^p + \Sigma^n)$.

We consider the steady state in which the mean-squared error is the same in each period in which a sample is gathered; let Σ^* be its steady-state value. Then $\Sigma^* = (\Sigma^p\Sigma^n)/(\Sigma^p + \Sigma^n)$. Since the last sample was gathered in period $t-k$, we have $\hat{\mu}_y^p(t) = E[\mu_y(t-k)\,|\,Y_t^p]$. Thus, the mean-squared error of $\hat{\mu}_y^p(t)$ as an estimate of $\mu_y(t-k)$ is also Σ^*; as an estimate of $\mu_y(t)$, the error of $\hat{\mu}_y^p(t)$ is augmented by the variance of the intervening innovations, so $\Sigma^p = \Sigma^* + k\sigma_\varepsilon^2$. We thus have the identity $\Sigma^* = (\Sigma^* + k\sigma_\varepsilon^2)\Sigma^n/(\Sigma^* + k\sigma_\varepsilon^2 + \Sigma^n)$, which simplifies to $(\Sigma^*)^2 + k\sigma_\varepsilon^2\Sigma^* - k\sigma_\varepsilon^2\Sigma^n = 0$. The positive solution to this quadratic equation, with the substitution $\Sigma^n = \sigma_y^2/n$ from Section 2.2, is

$$\Sigma^* = \frac{1}{2}\left(-k\sigma_\varepsilon^2 + \sqrt{(k\sigma_\varepsilon^2)^2 + 4k\sigma_\varepsilon^2\sigma_y^2/n}\right). \tag{9}$$

The previous three paragraphs and hence equation (9) presumed that there was no computational delay when updating the estimate with new information. Furthermore, they only considered the estimate of $\mu(t)$ and its MSE on dates in which new information was incorporate. With delay d, the information $\{Y(t), Y^p(t)\}$ is first used to estimate $\mu(t+d)$ rather than $\mu(t)$ and then, over the planning cycle, it is also used to estimate $\mu(t+d+1), \ldots, \mu(t+d+k-1)$. Because $E[\mu_y(t+s)\,|\,Y(t), Y^p(t)] = E[\mu_y(t)\,|\,Y(t), Y^p(t)]$, the updating rule is the same and the estimate does not change over the planning cycle. However, for each extra period of age that the information has, the MSE increases by the per-period variance $\sigma^2\varepsilon$ of the innovation. Hence, the MSE period j into the planning cycle is

$$\Sigma^*(j;n,d,k) = \frac{1}{2}\left(-k\sigma_\varepsilon^2 + \sqrt{(k\sigma_\varepsilon^2)^2 + 4k\sigma_\varepsilon^2\sigma_y^2/n}\right) + (d+j)\sigma_\varepsilon^2\,.$$

The average expected revenue over the planning cycle is

$$\bar{R}(n,d,k) = \sum_{j=0}^{k-1} R^*\left(\Sigma^*(j;n,d,k)^{-1/2}\right).$$

Compared to the one-shot model, the computational delay is increased by the additional time δm required to combine the two aggregate statistics for old and new information. However, for within-model analysis of the case of dynamic updating (e.g., analysis of optimal n and k and how these depend on the parameters), this constant m can be combined into the constant α from the earlier formula for delay. Thus, setting $d = \delta(\alpha + \log n)$ and using the fact that the average managerial wage is $w\delta n/k$ over the planning cycle, we have that the average profit as a function of n and k is

$$\bar{\Pi}(n,k) = \bar{R}(n, \delta(\alpha + \log n), k) - w\delta n/k\,.$$

Because the estimation problem starts in period 0, the average profit over the planning cycle is not at its steady-state value but rather converges to it as $t \to \infty$; nevertheless, $\bar{\Pi}(n, k)$ is the long-run average profit and is the measure of firm performance that we use.

[*Analysis to be continued.*]

7. Conclusion

We demonstrate that due to the inevitable delay caused by information processing, a firm must trade off how fast it can react to market conditions against how well it can be informed about market conditions. The optimal organizational solution to this problem is contingent on how fast market conditions change and the nature of the product the monopolist is selling. Thus, like Chandler (1962), we view the internal structure as being influenced by the information flows coming from the market. In our model there is only one kind of information and the issue is speed of response. In Chandler, there are different kinds of information, for example from marketing and production, which have to be combined. Thus we are considering a different information channel than Chandler and hence find results on the size of the decision-making unit while his results are on the grouping of tasks.

A number of directions for future research suggest themselves. On the organizational side there are a number of well discussed issues already in the literature—such as heterogeneous interrelated tasks, differing abilities and all kinds of incentive problems—the integration of which present enormous technical difficulties in information processing model. On the market side, the effect of competition, or more generally an endogenous environment, would be an interesting extension but beyond the scope of this paper.

Appendix

Proof of Proposition 1. From equation (1), the expected demand is equal to the probability that a particular random variable takes values in an interval of length $2(a - p)$. Varying x changes the centerpoint but not the size of this interval. For any unimodal distribution that is symmetric around the mean, such as the normal distribution, the above-mentioned probability is maximized by centering the interval at the mean. If also the distribution has a density that is strictly positive everywhere, then this is the unique solution. (The first-order condition is that the density takes on the same value at the two endpoints of the interval.) In this location-choice problem, the interval is centered on the mean by setting $x = \mu$. □

Proof of Proposition 2. Fix $\tau > 0$ and consider R as a function of p.

The set of p that maximize $R(p; \tau)$ is a non-empty subset of $(0, a)$ because (i) R is continuous in p, (ii) $R(p; \tau) \leq 0$ for $p \in (-\infty, 0] \cup [a, \infty)$, and (iii) $R(p; \tau) > 0$ for $p \in (0, a)$.

R is differentiable in p on $(0, a)$ and

$$\frac{\partial R(p;\tau)}{\partial p} = 2\Phi((a-p)\tau) - 2p\tau\varphi((a-p)\tau) - 1.$$

Any maximizer is a solution to the first-order condition $\partial R/\partial p = 0$, which is shown in equation (2).

Let $g(p,\tau) = \Phi((a-p)\tau) - p\tau\varphi((a-p)\tau)$, so that the first-order condition can be written $g(p,\tau) = 1/2$. This equation has only one solution if $g(p,\tau)$ is strictly decreasing in p on $(0, a)$. This condition holds because (i) $a-p > 0$ and hence $\varphi((a-p)\tau)$ increases toward its peak as p increases towards a and (ii) higher p means that $(a-p)\tau$ is smaller and hence so is $\Phi((a-p)\tau)$.

The solution p^* must be greater than $a/2$, as follows. Since

$$\Phi((a-p)\tau) - 1/2 = \Phi((a-p)\tau) - \Phi(0) = \int_0^{(a-p)\tau} \varphi(z)\,dz,$$

we can rewrite equation (2), evaluated at p^*, as

$$\int_0^{(a-p^*)\tau} \varphi(z)\,dz = p^*\tau\varphi((a-p^*)\tau). \tag{10}$$

Since φ is decreasing on $[0,\infty)$, the following holds for any $p \in (0, a)$:

$$\int_0^{(a-p)\tau} \varphi(z)\,dz \;>\; (a-p)\tau\,\varphi((a-p)\tau). \tag{11}$$

Combining equation (10) with equation (11) for $p = p^*$, we have $p^*\tau\varphi((a-p^*)\tau) > (a-p^*)\tau\,\varphi((a-p^*)\tau)$ and hence $(a-p^*) < p^*$, or $p^* > a/2$.

We next show that p^* is a strictly increasing function of τ. One can show that R does not have increasing differences in (p,τ) because, for any τ, $\partial^2 R(p;\tau)/\partial p\,\partial\tau < 0$ for p close enough to a. Therefore, we resort to local techniques. Specifically, we show that, when evaluated at $(p^*(\tau),\tau)$, we have $\partial g/\partial\tau > 0$ and hence $\partial^2 R(p;\tau)/\partial p\,\partial\tau > 0$. It then follows from the implicit function theorem that p^* is strictly decreasing.

Observe that

$$\frac{\partial g(p,\tau)}{\partial\tau} = (a-p)\varphi((a-p)\tau) - p\varphi((a-p)\tau) - p\tau(a-p)\varphi'((a-p)\tau). \tag{12}$$

The density φ of $N(0,1)$ satisfies $\varphi'(z) = -z\varphi(z)$. Making the substitution $\varphi'((a-p)\tau) = -(a-p)\tau\,\varphi((a-p)\tau)$ in equation (12) and collecting terms yields

$$\frac{\partial g(q,\tau)}{\partial\tau} = \underbrace{\left(a - 2p + p(a-p)^2\tau^2\right)}_{(a)}\varphi((a-p)\tau).$$

Thus, $\partial g/\partial\tau > 0$ if term (a) is positive. This condition is equivalent to

$$(2 - a/p)(a-p)^{-2} < \tau^2. \tag{13}$$

We need to show that it holds when $p = p^*(\tau)$.

Let $h(p)$ denote the left-hand side of inequality (13). On the range $(a/2, a)$, $h(p)$ is strictly increasing (since $(2 - a/p)$ is increasing in p and $(a - p)$ is decreasing in p) and it ranges from 0 when $p \approx a/2$ to ∞ when $p \approx a$. Therefore, there is a unique solution $\hat{p}(\tau) \in (a/2, a)$ to $h(p) = \tau^2$.

The final step is to show that $p^*(\tau) < \hat{p}(\tau)$; this yields the desired conclusion $h(p^*(\tau)) < \tau^2$ because $h(\hat{p}(\tau)) = \tau^2$ and h is increasing. With this objective, we show that $g(\hat{p}(\tau), \tau) < 1/2$; since $g(p^*(\tau), \tau) = 1/2$ and since g is decreasing in p, we have $p^*(\tau) < \hat{p}(\tau)$.

Fix τ, denote $\hat{p}(\tau)$ merely by $\hat{p}$, and let $\hat{\beta} = (a - \hat{p})\tau$. Since $\hat{p}$ is defined by $(2 - a/\hat{p})(a - \hat{p})^{-2} = \tau^2$, we have $\hat{\beta}^2 = (2 - a/\hat{p}) \in (0, 1)$. Hence, $1 - \hat{\beta}^2 = (a - \hat{p})/\hat{p}$ and $\hat{\beta}/(1 - \hat{\beta}^2) = \hat{p}\tau$. Therefore, we can write

$$g(\hat{p}, \tau) = \Phi(\hat{\beta}) - \frac{\hat{\beta}}{1 - \hat{\beta}^2}\varphi(\hat{\beta}).$$

Define the function f by $f(\beta) = \Phi(\beta) - (\beta/(1 - \beta^2))\varphi(\beta)$ for $\beta \in \mathbb{R}$, so that $g(\hat{p}, \tau) = f(\hat{\beta})$. One can verify that

$$\frac{d}{d\beta}\left(\frac{\beta}{1 - \beta^2}\right) = \frac{1 + \beta^2}{(1 - \beta^2)^2}.$$

Using again the identify $\varphi'(z) = -z\varphi(z)$, we have

$$\begin{aligned} f'(\beta) &= \varphi(\beta) - \frac{1 + \beta^2}{(1 - \beta^2)^2}\varphi(\beta) + \frac{\beta^2}{1 - \beta^2}\varphi(\beta) \\ &= \frac{1}{(1 - \beta^2)^2}\left((1 - 2\beta^2 + \beta^4) - (1 + \beta^2) + (\beta^2 - \beta^4)\right)\varphi(\beta) \\ &= \frac{-2\beta^2}{(1 - \beta^2)^2}\varphi(\beta) \\ &< 0. \end{aligned}$$

Since also $f(0) = 1/2$ and $\hat{\beta} > 0$, $f(\hat{\beta}) < 1/2$ as claimed.

We conclude by demonstrating the limits stated at the end of the proposition. Consider first the limit as $\tau \uparrow \infty$. Suppose $p^*(\tau)$ converges to some $p_\infty < a$ rather than to a. Then $\Phi((a - p^*(\tau))\tau)$ converges to 1. However, $p^*(\tau)\tau\,\varphi((a - p^*(\tau))\tau)$, which is proportional to $\tau(\exp\{(a - p^*(\tau))/2\})^{-\tau^2}$, converges to zero. Hence, the first-order condition in equation (2) is eventually violated.

Next consider the limit as $\tau \downarrow 0$. Let $p_0 = \lim_{\tau\downarrow 0} p^*(\tau)$. As $\tau \downarrow 0$ the left-hand side of equation (10) converges to $(a - p_0)\varphi(0)$ and the right-hand side converges to $p_0\varphi(0)$. Hence, $a - p_0 = p_0$, so that $p_0 = a/2$. □

Proof of Proposition 3. $R^*(\tau)$ is increasing in τ since $R(p;\tau)$ is increasing in τ for fixed p. Even at a price of 0 the expected demand $2\Phi(a\tau) - 1$ converges to 0 as $\tau \downarrow 0$; hence $\lim_{\tau\to 0} R^*(\tau) = 0$. At any price $a - \varepsilon$, the expected demand $2\Phi(\varepsilon\tau) - 1$ converges to the maximum possible value of 1 as $\tau \uparrow \infty$ and the revenue converges to $a - \varepsilon$. Hence $\lim_{\tau\to\infty} R^*(\tau) = a$. □

Proof of Proposition 4. Since Φ is an increasing function, the sign of $\partial Q/\partial\tau$ at the optimal price equals the sign of $\partial((a-p^*)\tau)/\partial\tau$.

$$\frac{\partial((a-p^*)\tau)}{\partial\tau} = a - p^* - \tau\frac{\partial p^*}{\partial\tau}. \tag{14}$$

From the implicit function theorem, $\partial p^*/\partial\tau = -(\partial g/\partial\tau)/(\partial g/\partial p)|_{p=p^*}$, where g is the function defined in the proof of Proposition 2. From that proof, $\partial g/\partial\tau = (a-2p+pz^2)\varphi(z)$, where $z=(a-p)\tau$. Using the fact that $\varphi'(z) = z\varphi(z)$, we have

$$\frac{\partial g}{\partial p} = -\tau\varphi(z) - \tau\varphi(z) + p\tau^2\varphi'(z) = -(2\tau + p\tau^2 z)\varphi(z).$$

Therefore, dropping the * from p^* to simplify notation,

$$\begin{aligned}\frac{\partial((a-p)\tau)}{\partial\tau} &= (a-p) - \tau\frac{(a-2p+pz^2)}{2\tau+p\tau^2 z}\\ &= \frac{(a-p)(2+p\tau z)-(a-2p+pz^2)}{2+p\tau z}\\ &= \frac{2a-2p+pz^2-a+2p-pz^2}{2+p\tau z}\\ &= \frac{a}{2+p\tau^z} > 0.\end{aligned}$$

□

Proof of Proposition 5. See the paragraph preceding the proposition. □

Proof of Proposition 6. The effect of σ_ε^2 and σ_y^2 on n follows from the fact that optimal sample size is $n = \sigma_y^2/(\delta\sigma_\varepsilon^2)$. The effect on monopoly profit follows from the fact that, for fixed n, $V(n)$ is increasing—and hence $\Pi(n)$ is decreasing—in both σ_ε^2 and σ_y^2. □

Proof of Proposition 7. The interaction between n and w in $\Pi(n)$ is only through the term $-w\delta n$. Therefore, Π has strictly decreasing differences in (n,w). □

Proof of Corollary 1. According to Proposition 5, optimal size equals $\sigma_y^2/\sigma_\varepsilon^2$ when $w=0$. According to Proposition 7, optimal size is smaller if $w>0$. □

Lemma 1 is used in the proof of Proposition 8.

LEMMA 1. *$V(n)$ is strictly convex in n for $n \le \sigma_y^2/(\delta\sigma_\varepsilon^2)$.*

Proof. Observe that

$$V''(n) = \frac{2\sigma_y^2}{n^3} - \frac{\delta\sigma_\varepsilon^2}{n^2}.$$

Then $\operatorname{sign}(V''(n)) = \operatorname{sign}(2\sigma_y^2 - n\delta\sigma_\varepsilon^2)$. Hence, $V''(n) > 0$ if $n < 2\sigma_y^2/(\delta\sigma_\varepsilon^2)$. □

Proof of Proposition 8. We show that Π has strictly decreasing differences in $(n, \sigma_\varepsilon^2)$ for $n \leq \sigma_y^2/(\delta\sigma_\varepsilon^2)$ (the range in which we know that any optimal sample size must lie).

Let $f(v) = R^*(v^{-1/2})$. Then $f'(v) < 0$ since R^* is increasing. We show numerically that f is strictly convex and hence f' is increasing.[3] Since $\Pi(n) = f(V(n)) - w\delta n$,

$$\frac{\partial \Pi}{\partial n} = f'(V(n))V'(n) - \delta w.$$

$V(n)$ is increasing in σ_ε^2. On the range $n < \sigma_y^2/(\delta\sigma_\varepsilon^2)$, an increase in σ_ε^2 has the following effects on $\partial\Pi/\partial n$:

1. $f'(V(n))$ is negative as noted above, and it becomes smaller in magnitude because $V(n)$ goes up and $f'(v)$ is increasing.
2. $V'(n)$ is negative for $n \leq \sigma_y^2/(\delta\sigma_\varepsilon^2)$, and since $V'(n) = -\sigma_y^2/n^2 + \delta\sigma_\varepsilon^2/n$, it becomes smaller in magnitude when σ_ε^2 increases.

Therefore, the term $f'(V(n))V'(n)$ is positive and becomes smaller as σ_ε^2 increases. Thus $\partial\Pi/\partial n$ is decreasing in σ_ε^2 and Π has strictly decreasing differences in $(n, \sigma_\varepsilon^2)$ □

Proof of Proposition 9. We take into account the constraint that n cannot lie between 0 and 1. Recall from Proposition 5 that optimal sample size when $w = 0$ is equal to $n = \sigma_y^2/(\delta\sigma_\varepsilon^2)$ and recall from Proposition 7 that optimal sample size is decreasing in w. Hence, optimal sample size is at most 1 if $\sigma_y^2 \leq \delta\sigma_\varepsilon^2$.

Recall that $V(n) \geq \sigma_y^2$ for any n and that, from Proposition 3, $\lim_{\tau\to 0} R^*(\tau) = 0$. Therefore, $R^*(1/V(n)^2)$ converges to zero uniformly in n as $\sigma_y^2 \to \infty$. Hence, there is $h > 0$ such that, if $\sigma_y^2 > h$, then $R^*(1/V(n)^2) < w\delta n$ for all $n \geq 1$. That is, there is no positive sample size in which the revenue covers the information processing cost. It is better to have a sample size of zero, which leads to zero profit. □

Proof of Proposition 10. We show that Π has increasing differences in (n, a) when we restrict n to the convex domain of values for which $V'(n) < 0$, which we know must hold for optimal n (this domain does not depend on a).

We start with

$$\frac{\partial \Pi}{\partial n} = \frac{\partial R^*}{\partial \tau}\frac{\partial \tau}{\partial n} - w\delta.$$

The term $w\delta$ does not depend on a; the term $\partial\tau/\partial n$ is positive for n such that $V'(n) < 0$ and it does not depend on a either. Therefore, $\partial\Pi/\partial n$ is increasing in a if $\partial R^*/\partial\tau$ is increasing in a.

Recall that $R^*(\tau) = R(p^*(\tau), \tau)$. By the envelope theorem, $\partial R^*/\partial\tau = \partial R/\partial\tau|_{p=p^*}$. Since $R(p, \tau) = p(2\Phi((a - p)\tau) - 1)$,

$$\frac{\partial R}{\partial \tau} = 2p\varphi((a - p)\tau)(a - p). \tag{15}$$

3. The only parameter that R depends on is a. We tested the strict convexity condition

$$f(\lambda v_1 + (1 - \lambda v_2); a) > \lambda f(v_1; a) + (1 - \lambda) f(v_2; a)$$

for 2.5 million trials. In each trial, a, v_1, and v_2 are drawn from log-normal distributions ($a \overset{d}{=} 8^{\tilde{x}}$, $v_1 \overset{d}{=} 1.5^{\tilde{x}}$, and $v_2 \overset{d}{=} 1.5^{\tilde{x}}$, where $\tilde{x} \sim N(0, 1)$) and λ is drawn from a uniform distribution on $[0, 1]$.

When showing that $\partial R/\partial\tau|_{p=p^*}$ is increasing in a, we have to take into account that p^* depends on a. From the first-order condition $\partial R(p,\tau)/\partial p = 0$ that p^* must satisfy, which is shown in equation (2), we have

$$2p^*\varphi((a-p^*)\tau) \;=\; \frac{1}{\tau}(2\Phi((a-p^*)\tau)-1). \tag{16}$$

Evaluating equation (15) at p^* and making the substitution from equation (16) yields

$$\left.\frac{\partial R}{\partial \tau}\right|_{p=p^*} \;=\; \frac{1}{\tau}(2\Phi((a-p^*)\tau)-1)(a-p^*).$$

Since $\Phi(\cdot)$ is increasing, $\partial R/\partial\tau|_{p=p^*}$ is increasing in a if $a-p^*$ is increasing in a. This holds if $\partial p^*/\partial a < 1$, which we now show by use of the implicit function theorem.

As in the proof of Proposition 2, we let $g = \Phi((a-p)\tau) - p\tau\varphi((a-p)\tau)$, so that the first-order condition that determines p^* is $g = 0$. Then $\partial p^*/\partial a = -(\partial g/\partial a)/(\partial g/\partial p)$. Letting $z = (a-p)\tau$ and using the fact that $\varphi'(z) = z\varphi(z)$, we have

$$\frac{\partial g}{\partial a} \;=\; \varphi(z)\tau - p\tau^2\varphi'(z) \;=\; (\tau + p\tau^2 z)\varphi(z),$$

$$\frac{\partial g}{\partial p} \;=\; -\tau\varphi(z) - \tau\varphi(z) + p\tau^2\varphi'(z) \;=\; -(2\tau + p\tau^2 z)\varphi(z).$$

Hence,

$$\frac{\partial p^*}{\partial a} \;=\; \frac{\tau + p\tau^2 z}{2\tau + p\tau^2 z} \;=\; \frac{1+p\tau z}{2+p\tau z} \;<\; 1,$$

as claimed. □

Proof of Proposition 11. Observe that Proposition 5, Proposition 7, and Corollary 1 are unaffected by the introduction of the parameter T. Therefore, as in the first step of the proof of Proposition 9, the optimal sample size is at most 1 if $\sigma_y^2 \leq \delta\sigma_\varepsilon^2$.

To mimic the second step of the proof of Proposition 9, we note that maximized revenue is $R^*(\tau/T)$, where R^* has the same functional form as before. Since $V(n) \geq \sigma_y^2$ for any n and since $\lim_{\tau\to 0} R^*(\tau) = 0$, $R^*((1/T)V(n)^{-1/2})$ converges to zero, uniformly in n, as $T \to \infty$. We can then proceed as in the proof of Proposition 9. □

Proof of Proposition 12. See the paragraph preceding the proposition. □

Proof of Proposition 13. Holding fixed n, an increase in σ_ε^2 causes $V(n)$ to increase. From Proposition 8, an increase in σ_ε^2 causes n to fall, and we know that optimal n lies in a range in which $V'(n) < 0$. Therefore, the effect of σ_ε^2 on n also causes the variance to rise. □

Proof of Proposition 14. See the paragraph preceding the proposition. □

References

Bolton, Patrick and Matias Dewatripont (1994). The Firm as a Communication Network. *Quarterly Journal of Economics*, 109, 809–839.

Chandler, Alfred D. (1962). *Strategy and Structure*. New York: Doubleday.

Meagher, Kieron, Hakan Orbay, and Timothy Van Zandt (2003). Hierarchy Size and Environmental Uncertainty. In *Advances in Economic Design.*, pp. 439–457. Dordrecht: Kluwer. (Proceedings of the Conference on Economics Design, 2000.).

Radner, Roy (1993). The Organization of Decentralized Information Processing. *Econometrica*, 62, 1109–1146.

Van Zandt, Timothy (1998a). Organizations with an Endogenous Number of Information Processing Agents. In M. Majumdar (Ed.), *Organizations with Incomplete Information*. Cambridge: Cambridge University Press.

Van Zandt, Timothy (1998b). The Scheduling and Organization of Periodic Associative Computation: Efficient Networks. *Economic Design*, 3, 93–127.

Van Zandt, Timothy (2003). Structure and Returns to Scale of Real-Time Hierarchical Resource Allocation. INSEAD.

Van Zandt, Timothy and Roy Radner (2001). Real-Time Decentralized Information Processing and Returns to Scale. *Economic Theory*, 17, 497–544.