The Quality of Information and Incentives for Effort*

Omer Moav† and Zvika Neeman‡

September 20, 2004

Abstract

We study the relationship between the precision of information about the performance of an agent in a market, and the incentives this agent has for exerting effort to produce high quality. We show that this relationship is not monotonic. There exists a threshold beyond which any further improvement in the precision of information weakens the agent's incentive to produce high quality. Accordingly, both very precise and very imprecise information about the agent's performance may destroy its incentive to exert effort. Applications to real markets are discussed.

JEL CLASSIFICATION NUMBERS:
KEYWORDS:

* Acknowledgements to be added.
† Department of Economics, the Hebrew University of Jerusalem, Jerusalem, Israel 91905. Email `msmoav@mscc.huji.ac.il`, URL http://economics.huji.ac.il/facultye/moav/moav.html
‡ Department of Economics, Boston University, 270 Bay State Road, Boston MA 02215, Center for Rationality and Department of Economics, the Hebrew University of Jerusalem, Jerusalem, Israel 91904. Email `zvika@BU.edu`, URL http://people.bu.edu/zvika/.

1. Introduction

This paper concerns the relationship between the precision of information about the performance of an agent in a market, and the incentives this agent has for exerting effort to produce high quality. We show that this relationship is not monotonic. There exists a threshold beyond which any further improvement in the precision of information weakens the agent's incentive to produce high quality. Accordingly, both very imprecise but also very precise information about the agent's performance may destroy its incentive to exert effort.

We consider a dynamic model of a market for a good whose quality is not observable to the consumer at the time of purchase and is not contractible. Such goods are referred to as experience or credence goods in the literature.[1] Examples range from food and wine, to used cars, to expert advice. Importantly, because consumers cannot contract on the quality of the good, the price they are willing to pay for it depends on their beliefs about the expected quality of the good, which are synonymous with the reputation of the producer who produces that good. Consumers form their beliefs about the expected quality of the good in any given period (equivalently, the producer's reputation in any given period is obtained) by updating their beliefs about the producer's ability (type) based on the producers' past performance, taking into account the producer's incentives to exert effort to produce high quality given future prices.

For any given prior beliefs about the producer's ability, the producer's incentive to exert a costly effort in order to produce high quality is increasing with the probability that the true quality of the good would be revealed. Hence, for any given consumers' prior beliefs, an improvement in consumers' ability to detect high and low quality goods has an obvious positive effect on the producer's incentives to produce high quality in the current period. However, as we show here, such an improvement may have an overall negative effect on the incentives to produce high quality because, in a dynamic setting, consumers' prior beliefs are affected by the precision of their information.

If prior beliefs regarding a producer's ability are very accurate, then a contradictory signal about the producer's ability would be attributed to either sampling error or a random shock in the production process, and would only have a small effect on the consumers' posterior beliefs. It follows that if the prior probability that a producer is competent is sufficiently high, then this producer could avoid effort and 'rest on its laurels' without incurring a significant loss of reputation. Hence, precise information that induces accurate priors might generate perverse incentives. Consequently, an equilibrium in which competent producers exert costly effort in order to produce high quality goods and maintain their reputation could unravel as consumers' information becomes more precise.

Consider for example the case of a restaurant that serves sushi. An important aspect of the quality of sushi, which consumers have no way of verifying at the spot, is whether

[1] An experience or a credence good is distinguished by the fact that its quality cannot be determined by consumers at the time of purchase. The true quality of an experience good is revealed later when consumers experience its consumption. The true quality of a credence good is never revealed to consumers.

or not the sushi contains "germs" that cause food disease. Sushi that contains such germs often causes stomach upset but not always. And not every stomach upset is related to the consumption of contaminated sushi. Therefore, a consumer of sushi can never tell for sure whether the sushi she ate is of high or low quality.

Sushi may contain germs either because the sushi chef is incompetent and simply does not know how to pick fresh fish, or because the sushi chef saves money and effort and buys inferior fish, or fails to handle the fish properly after he buys them. Suppose that consumers are aware of whether or not other consumers experienced a stomach upset in the previous couple of periods after they had ate at the restaurant, and that furthermore, the price of sushi in the restaurant depends on whether or not consumers suffered a stomach upset after eating there. The results reported in this paper suggest that as the consumers' information about the restaurant becomes more precise, for example, if the pool of other consumers with whom the consumer shares its information about the quality of the sushi increases in size, then the incentives of competent sushi masters to exert effort may become weaker. Intuitively, suppose that consumers believe that the restaurant chef is competent with a very high probability, as would be the case if consumers are familiar with a large number of other satisfied consumers who did not suffer a stomach upset after eating in the restaurant. Under such circumstances, news about a few cases of stomach upset would most likely not be attributed to the poor quality of sushi but to some other reason, and this might give the restaurant an incentive to "rest on its laurels" for a while, and cut cost by buying inferior quality fish.

This paper combines several ideas about the relationship between reputation and incentives, which have already been discussed in the literature, which we present below.

Career Concerns

Holmström (1999) considered a model in which an agent's future career concerns influence its incentives to exert effort. The output produced by the agent is not contractible, so it is impossible to directly reward or penalize the agent based on its past performance. Rather, in each period, the agent's wage is determined based on the belief about its ability (the agent's reputation) and its expected future effort. Initially, an agent in Holmström's model may exert some effort, but with time, as information about the agent's true ability becomes more and more precise, the agent's incentive to exert effort weakens, and the agent's level of effort decreases to zero. Holmström shows that if the agent's ability changes stochastically over time, then an incentive to exert effort can be sustained, because in every period, the agent still has an incentive to prove anew that its ability is high.

There are three important differences between Holmström's model and ours. First, while in our model memory is finite, in Holmström's model, the agent's entire history of past performance is observable. Second, while in our model producers know their types, in Holmström's model the information about the agent's true ability is symmetric. Namely, the agent and the market are equally well informed about the agent's true ability.[2] Finally,

[2] Although the market is not able to directly observe the agent's effort in Holmström's model, it can infer

third, while in our model producers' types are fixed, in Holmström's model they are subject to random shocks. Accordingly, Holmström's conclusions about the effect of a change in the precision of the information about the agent's performance is very different from ours. While in our model, more precise information may weaken the agent's incentives to exert effort as explained above, in Holmström's model it unambiguously leads to stronger incentives to exert effort. We believe that this difference is due to Holmström's assumption that the agent is not better informed than the market about its own ability. While in our model, an agent who knows it is of high ability has an incentive to rest on its laurels for a while after it has established a good reputation, an agent in Holmström's model would be more reluctant to cash in on its reputation because it cannot count on being able to rehabilitate its tarnished reputation in the same way that an agent who knows it is highly capable can.

Dewatripont, Jewitt, and Tirole (1999) build on Holmström's model to characterize information structures in terms of their effect on the agent's incentives. They identify information structures where more precise information may weaken the agent's incentives. They describe a number of examples that are all based on the following insight. Consider the agent's incentives to exert effort when information about its performance is given by the more informative signal (y, z) compared to the less informative signal y. Conditional on the realization of the signal y, suppose that the market's expectation of the agent's talent is increased when higher values of statistic z are observed. If higher effort by the agent tends to increase z (which follows from the MLRP property), then having z in the market information set enhances effort. However, if more effort on part of the agent tends to decrease z, then having z in the market information set would reduce the incentive for effort (pp. 193-4). Thus, the reason that better information may weaken incentives in Dewatripont et al.'s model is different from the reason that is described in this paper. Because, unlike in Holmström's model and in this paper, Dewatripont et al. only consider a two-period model, the informativeness of the signal has no effect on the market's prior beliefs at the beginning of each period as in our model. Furthermore, unlike the results obtained here, they show that under a number of "regularity" conditions, more precise information about the agent's performance unambiguously improves the agent's incentives to exert effort.

Reputation as Separation from Less Competent Types

In our model, as well as in all similar models, competent producers exert effort to produce high quality in order to maintain a 'reputation for competence.' The existence of incompetent producers is thus crucial for our results. For if all producers were known to be equally competent, then producers would not be capable of distinguishing themselves as 'more' competent than others, and would thus lose the incentive to exert costly effort. See for example, Mailath and Samuelson (2001), and in the different context of the enforcement of cooperation in repeated community prisoner's dilemma like games, Ghosh and Ray (1996). For obvious reasons, the mere existence of incompetent producers, by itself, is insufficient to provide competent producers with sufficient incentives. It must be that consumers assign a sufficiently

it by solving the agent's optimization problem (Holmström, 1999, p. 171).

high likelihood that any producer is incompetent to provide any particular producer with the incentive to exert the costly effort associated with the production of high quality so as to distinguish itself from less competent producers. Thus, another intuitive explanation for our main result is that as information becomes more precise, a competent producer finds it easier to distinguish itself from less competent types. The fact that separation becomes easier might imply that the incentive to exert costly effort in order to distinguish oneself is weakened.

The Market for Names

The relationship between reputation and incentives has also been explored in the context of the "market for names" where names serve as repositories for reputations (see Mailath and Samuelson, 2001; Tadelis 1999, 2002, and 2003; and the references therein). These authors studied the market for names that develops when producers of a certain good occasionally exit the market and sell their reputations to new entrants to the market. They have showed that such a 'market for names' provides an incentive to exert effort to produce high quality so as to build a 'name' or a reputation that can later be sold.

The contribution of this paper is that it provides an intuitive account of the relationship between the precision of consumers' beliefs, and hence producers' reputation, and producers' incentives. We build on the familiar idea that a producer with a good reputation might 'rest on its laurels' and produce low quality to argue that, by strengthening the inference that a producer with a good reputation is indeed very likely to be truly competent, an improvement in the precision of information may have a perverse effect on incentives. We identify a threshold beyond which any further improvement in the precision of information would reduce the incentives to exert effort in order to produce high quality.

The rest of the paper proceeds as follows. In the next section, we describe the model. In Section 3, we show the existence of a threshold, beyond which any further improvement in the precision of information would weaken the incentives to produce high quality. In section 4, we show that increasing the length of memory, which was held fixed in Section 3, and which provides yet another way in which information about past performance can be made more precise, may have a similar negative effect on incentives. We conclude in Section 5 with a discussion of the possible implications of our analysis.

2. Model

We describe the simplest possible model in which we can elucidate our main argument. For this reason, we describe a model in which there are two types of producers (competent and incompetent), who might exert two levels of effort (exert effort or not), to each produce one unit of a good that may either have a high or low quality. Consumers, or the "market," have an n period memory. That is, they observe whether or not each producer has passed or failed inspection in the previous $n \geq 1$ periods.

Specifically, we consider a dynamic model of a market economy. Time is discrete, and periods are indexed by $t \in \mathbb{I} \equiv \{..., 0, 1, ...\}$. In every period, each producer produces one unit of the good. He has to decide what quality to produce, high or low.[3] High and low quality goods cannot be distinguished by consumers at the time of purchase. A high quality good has value 1 and a low quality good has value 0 for consumers.

There is a continuum of each of the two types of producers. The measures of competent and incompetent producers are normalized to $1 - \eta$ and η, respectively, where $\eta \in (0, 1)$. Producers discount future payoffs at the rate $\delta < 1$. Competent producers can produce high quality at a cost $c > 0$. Incompetent producers are incapable of producing high quality. The cost of producing a low quality good is zero for both types of producers, in every period.

Every period, produced goods are subject to inspection. We assume that high and low quality goods pass the inspection with probabilities π^H, and π^L, respectively, and that $0 < \pi^L < \pi^H \leq 1$. Whether or not each producer passes or fails inspection is public information.[4] This public information is forgotten after $n \geq 1$ periods. Accordingly, in every period, producers are sorted into 2^n sub-markets, depending on whether or not they passed or failed inspection in the previous n periods.

Denote an infinite history of passes and fails that ends in period $t - 1$, that a producer might have in period t, by h_t. Let H_t denote the set of all such histories. Let $h_t P$ denote an infinite history of passes and fails that coincides with the history h_t up to period $t - 1$, and that is followed by a pass in period t. The history $h_t F$ is similarly defined. Recursively applying this definition implies that $h_t PP$ denotes the infinite history of passes and fails that coincides with the history h_t up to period $t - 1$, and that is followed by two additional passes in periods t and $t + 1$, respectively, and so on.

Let $h_t(\tau) \in \{P, F\}$ denote whether a producer with history h_t at t passed or failed inspection in period $\tau \leq t - 1$. Because the market is assumed to have an n period memory, all the producers with histories $h_t \in H_t$ that have the same profile of realizations $(h_t(t-n), ..., h_t(t-1))$ at t are sorted into the same submarket in period t, which with a slight abuse of notation, we also denote by either h_t or more straightforwardly by the profile of pass/fail realizations in the last n periods, $\widehat{h}_{n,t} = (h_t(t-n), ..., h_t(t-1))$.

We assume that in every period, demand in each submarket is infinitely elastic at the expected value of the good to consumers in that submarket.

For every period $t \in \mathbb{I}$, let $p_t^{h_t}$ denote the price in submarket $h_t \in H_t$ at t. Our notion of market-equilibrium is defined as follows.

Definition 1. *A sequence of prices* $\left\{ \left(p_t^{h_t} \right)_{h_t \in H_t} \right\}_{t \in \mathbb{I}}$ *is a market-equilibrium if:*

---

[3]The conclusions of the model remain qualitatively the same if the produced quality is stochastic so that a producer who incurs the cost of producing high quality produces low quality, and vice-versa.

[4]The conclusions of the model remain qualitatively the same if a producer who has failed inspection is subject to a fine, provided, of course, that this fine is not so large so as to cause to producer to not want to enter the market in the first place.

1. *In every period $t \in \mathbb{I}$, producers produce the quality that maximizes the discounted value of their expected profits given the sequence of prices $\left\{ \left(p_t^{h_t}\right)_{h_t \in H_t} \right\}_{t \in \mathbb{I}}$. We assume, for simplicity, that in case of indifference, producers produce high quality.*

2. *In every period $t \in \mathbb{I}$, the prices in each submarket is equal to the expected quality of the good for consumers in that submarket.*

Remark 1. Interpretation as a Model of Career Concerns

The description above has emphasized an interpretation of the model as a model of the market for an experience good. However, the same assumptions also admit an interpretation of the model as a model of an agent whose future career concerns influence its incentives to exert effort as in the "career concern" literature mentioned in the introduction. Under this alternative interpretation, instead of a continuum of producers, there is only one agent, who is initially believed to be competent with probability $1 - \eta$. The output of the agent is assumed to be non contractible, so that in every period, the agent is paid a wage that is based on the belief about its competence and its expected effort in that period. Under this alternative interpretation, the 2^n different submarkets may be thought of as the 2^n reputations that an agent might have in an environment where the market only obtains noisy signals about the agent's performance in the last n periods.

Remark 2. Existence of a Market-Equilibrium

The model described above admits the existence of at least one market-equilibrium. In particular, the sequence of prices $\left\{ \left(p_t^{h_t}\right)_{h_t \in H_t} \right\}_{t \in \mathbb{I}}$ where $p_t^{h_t} = 0$ for every $h_t \in H_t$ and $t \in \mathbb{I}$ and where no producer ever produces high quality is a market-equilibrium. To see this, observe that if prices are zero in every period, then no producer has any incentive to incur the cost required to produce high quality. Under this equilibrium, any passing of inspection would be attributed to inspection error.

The next lemma characterizes the behavior of competent producers in a market equilibrium. Denote $\pi \equiv \pi^H - \pi^L$. Given a market equilibrium $\left\{ \left(p_t^{h_t}\right)_{h_t \in H_t} \right\}_{t \in \mathbb{I}}$, let $U_t^{h_t}$ denote the expected discounted payoff of a competent producer with a history $h_t \in H_t$ in period t who proceeds to behave optimally in period t and onwards.

Lemma 1. *Fix a market-equilibrium $\left\{ \left(p_t^{h_t}\right)_{h_t \in H_t} \right\}_{t \in \mathbb{I}}$. In every period $t \in \mathbb{I}$, a competent producer with a history h_t produces high quality if and only if*

$$c \leq \delta \pi \left(U_{t+1}^{h_t P} - U_{t+1}^{h_t F} \right). \tag{1}$$

Proof. In every period t, a competent producer with a history h_t produces high quality if and only if,

$$\delta \left(\pi^H U_{t+1}^{h_t P} + \left(1 - \pi^H\right) U_{t+1}^{h_t F} \right) - c \geq \delta \left(\pi^L U_{t+1}^{h_t P} + \left(1 - \pi^L\right) U_{t+1}^{h_t F} \right)$$

if and only if (1). ■

The fact that the market has an n period memory implies that whatever has happened more than n periods ago has no effect on the producers' incentives. It therefore follows that two competent producers whose histories of passes and fails coincide in the last $n-1$ periods face identical incentives because they would be treated in exactly the same way as of the next period. This implies that if two histories of passes and fails $h_t, h_t' \in H_t$ coincide in the last $n-1$ periods, then

$$U_t^{h_t} - U_t^{h_t'} = p_t^{h_t} - p_t^{h_t'}.$$

Namely, the difference between the discounted expected payoffs of a competent producer in period t after histories h_t and h_t' is due only to the different price they would fetch in period t itself. It therefore follows that a competent producer with history h_t produces high quality if and only if a competent producer with history h_t' does.

3. The Precision of Information

The higher the cost of producing a high quality good, c, the weaker the incentive to produce it. The strength of incentives can therefore be measured by how high is the threshold cost above which competent producers may sometimes refuse to produce high quality. This threshold obviously depends on the particular market equilibrium that is considered. In the equilibrium in which no producer ever produces high quality, this threshold is zero. We are interested in the behavior of this threshold in the equilibrium where competent producers always produce high quality, regardless of the particular submarket to which they have access in any given period.

For simplicity, we assume in this section that the market has a 2-period memory. As should become clear below, our results should hold for any memory of finite length $n \geq 2$.

In a market equilibrium in which competent producers produce always produce high quality regardless of the submarket in which they find themselves in any given period, the measure of competent and incompetent producers in any period in submarket PP is given by $(1-\eta)\left(\pi^H\right)^2$ and $\eta\left(\pi^L\right)^2$, respectively. The price in submarket PP is therefore given by

$$p^{PP} = \frac{(1-\eta)\left(\pi^H\right)^2}{(1-\eta)\left(\pi^H\right)^2 + \eta\left(\pi^L\right)^2} \tag{3}$$

in every period $t \in \mathbb{I}$. Similarly, the prices in submarkets PF, FP, and FF, are given by

$$p^{PF} = p^{FP} = \frac{(1-\eta)\,\pi^H\left(1-\pi^H\right)}{(1-\eta)\,\pi^H\left(1-\pi^H\right) + \eta\pi^L\left(1-\pi^L\right)}, \tag{4}$$

and

$$p^{FF} = \frac{(1-\eta)\left(1-\pi^H\right)^2}{(1-\eta)\left(1-\pi^H\right)^2 + \eta\left(1-\pi^L\right)^2}, \tag{5}$$

respectively, in every period $t \in \mathbb{I}$.[5]

By Lemma 1, a situation in which competent producers always produce high quality regardless of the submarket in which they find themselves in any given period is a market equilibrium if and only if (**??**) and (**??**) hold for every period $t \in \mathbb{I}$. In this case, the definition of the U_t^{ij}'s implies that

$$\begin{aligned} U_t^{PP} &= p^{PP} - c + \delta\left(\pi^H U_{t+1}^{PP} + \left(1-\pi^H\right) U_{t+1}^{PF}\right) \\ U_t^{PF} &= p^{PF} - c + \delta\left(\pi^H U_{t+1}^{FP} + \left(1-\pi^H\right) U_{t+1}^{FF}\right) \\ U_t^{FP} &= p^{FP} - c + \delta\left(\pi^H U_{t+1}^{PP} + \left(1-\pi^H\right) U_{t+1}^{PF}\right) \\ U_t^{FF} &= p^{FF} - c + \delta\left(\pi^H U_{t+1}^{FP} + \left(1-\pi^H\right) U_{t+1}^{FF}\right) \end{aligned}$$

for every $t \in \mathbb{I}$. The fact that in this case also $U_t^{ij} = U_{t+1}^{ij} \equiv U^{ij}$ for every $ij \in \{PP, PF, FP, FF\}$ and $t \in \mathbb{I}$, implies that (**??**) and (**??**) are equivalent to

$$c \leq \delta\pi\left(p^{PP} - p^{PF} + \delta\left(\pi^H\left(U^{PP} - U^{FP}\right) + \left(1-\pi^H\right)\left(U^{PF} - U^{FF}\right)\right)\right) \tag{6}$$

and

$$c \leq \delta\pi\left(p^{FP} - p^{FF} + \delta\left(\pi^H\left(U^{PP} - U^{FP}\right) + \delta\left(1-\pi^H\right)\left(U^{PF} - U^{FF}\right)\right)\right) \tag{7}$$

respectively. If competent producers always produce high quality regardless of the submarket in which they find themselves in any given period, then because memory is only two periods long, both $U^{PP} - U^{FP} = p^{PP} - p^{FP}$ and $U^{PP} - U^{FP} = p^{PP} - p^{FP}$. It is therefore possible to express (6) and (7) directly in terms of the prices in the different submarkets as follows

$$c \leq \delta\pi\left(\left(1+\delta\pi^H\right)p^{PP} - \left(1-\delta+2\delta\pi^H\right)p^{PF} - \delta\left(1-\pi^H\right)p^{FF}\right) \tag{8}$$

and

$$c \leq \delta\pi\left(\delta\pi^H p^{PP} + \left(1+\delta-2\delta\pi^H\right)p^{FP} - \left(1+\delta-\delta\pi^H\right)p^{FF}\right), \tag{9}$$

where (8) is equivalent to (3), and (9) is equivalent to (5). By plugging (3)-(5) into (8) and (9) it is then possible to determine the threshold cost above which it is no longer optimal for competent producers to always produce high quality regardless of the submarket in which they find themselves in any given period as a function of the parameters of the model π^H, π^L, η, and δ, alone. Denote this threshold by $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$.

The next proposition establishes the existence of a threshold on the precision of information, beyond which any further improvement in the precision of information would weaken the incentives of competent producers to produce high quality.

[5]Notice that

$$p^{PP} \geq p^{PF} = p^{FP} \geq p^{FF}$$

if and only if

$$\pi^H \geq \pi^L.$$

Proposition 1. *For every fixed values of* $\eta > 0$ *and* $\delta < 1$, *there exists threshold values of* π^L *and* π^H *such that any further improvement in the precision of information, namely either an increase in the value of* π^H *or a decrease in the value of* π^L *results in a lower value of* $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$.

Proof. Fix some values of $\eta > 0$ and $\delta < 1$. We say that inequality (8) is binding if and only if the RHS of inequality (8) is smaller than the RHS of inequality (9). Note that inequality (8) is binding if and only if the threshold c^* is determined by inequality (8). Algebraic manipulation shows that inequality (8) is binding if and only if

$$p^{PP} - p^{PF} < p^{FP} - p^{FF}.$$

The fact that $\lim_{\pi^L \searrow 0} p^{PP} = 1$, $\lim_{\pi^L \searrow 0} p^{PF} = \lim_{\pi^L \searrow 0} p^{FP} = 1$, and $\lim_{\pi^L \searrow 0} p^{FF} \in (0,1)$, implies that (8) is binding as π^L decreases to zero; and the fact that $\lim_{\pi^H \nearrow 1} p^{PP} \in (0,1)$, $\lim_{\pi^H \nearrow 1} p^{PF} = \lim_{\pi^H \nearrow 1} p^{FP} = 0$, and $\lim_{\pi^H \nearrow 1} p^{FF} = 0$, implies that inequality (9) is binding as π^H increases to 1.

We now show that for values of π^H that are close to 1, the RHS of inequality (8) is increasing in π^L whenever π^L is small enough. This follows from the fact that $\frac{dp^{FF}}{d\pi^L} > 0$ so that p^{FF} is increasing in π^L, and from the fact that while $\frac{dp^{PP}}{d\pi^L}$ tends to zero as π^L tends to zero, $\frac{dp^{PF}}{d\pi^L} = \frac{dp^{FP}}{d\pi^L}$ tends to $-\frac{\eta}{(1-\eta)\pi^H(1-\pi^H)}$ which becomes arbitrarily large as π^H tends to 1.

Finally, we show that for values of π^L that are close to 0, the RHS of inequality (9) is decreasing in π^H whenever π^H close enough to 1. This follows from the fact that $\frac{dp^{FF}}{d\pi^L}$ decreases to zero as π^H tends to 1, and from the fact that while $\frac{dp^{PP}}{d\pi^H}$ tends to $\frac{2\eta(1-\eta)\left(\pi^L\right)^2}{\left((1-\eta)+\eta(\pi^L)^2\right)^2} < \infty$ as π^H tends to one, $\frac{dp^{PF}}{d\pi^L} = \frac{dp^{FP}}{d\pi^L}$ tends to $-\frac{1-\eta}{\eta\pi^L(1-\pi^L)}$ which is negative and becomes arbitrarily large in absolute value as π^L tends to 0. ■

Remark 3. Numerical Analysis

Numerical analysis reveals that the sense in which more precise information can undermine incentives is in fact a little stronger than stated in Proposition 1. It reveals, specifically, that for every fixed parameters π^H, η, and δ, there exists a threshold value of π^L (smaller than π^H) below which $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$ is increasing in π^L, and that for every fixed parameters π^L, η, and δ, there exists a threshold value of π^H (larger than π^L) above which $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$ is decreasing in π^H. As an illustration, in Figure 1a below we plot $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$ as a function of π^L holding the values of π^H, η, and δ constant, and in Figure 1b below we plot $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$ as a function of π^H holding the values of π^L, η, and δ constant.

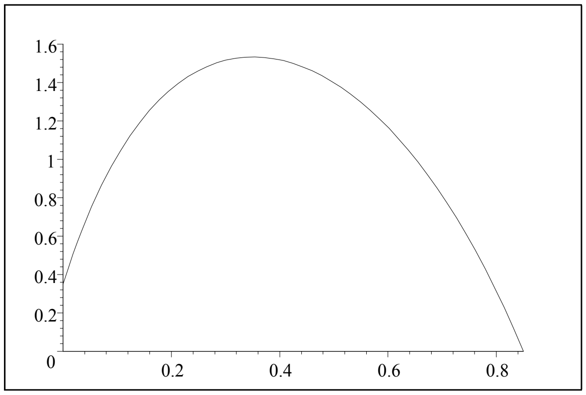

Figure 1a. $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$ as a function of π^L (for $\pi^H = .85$, $\eta = .2$, and $\delta = .9$).

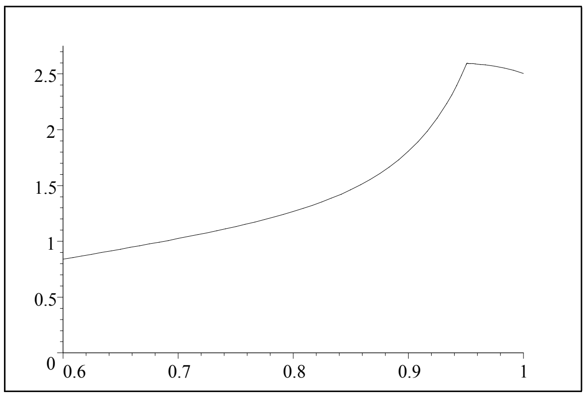

Figure 1b. $c^*\left(\pi^H, \pi^L, \eta, \delta\right)$ as a function of π^H (for $\pi^L = .25$, $\eta = .2$, and $\delta = .9$).

The intuition for the result is the following. As π^L decreases to zero, so few incompetent producers pass inspection that both p^{PF} and p^{PP} converge to 1. Therefore, a competent producer who has passed inspection in the previous period (either in submarket FP or PP) realizes that even if it fails inspection in the current period, it would still get a very good price, namely p^{PF}, in the next period, which significantly reduces its incentive to produce high quality. Such a producer need not fear the stigma associated with having produced low quality because by producing high quality it can pass inspection with a high probability and so gain access to submarket FP where the price p^{FP} is almost equal to the highest possible price it could get, p^{PP}. More formally, as shown in the proof of Proposition 1, the binding incentive constraint in this case is inequality (8), which implies that producers who have passed inspection in the previous period may want to rest on their laurels and produce low quality in the current period.

The intuition for what happens to the incentive to produce high quality as π^H increases to one is different. In this case, it is inequality (9) that is binding, which implies that producers who have failed inspection in the previous period become discouraged and stop producing high quality. The reason for this is that when π^H is very close to 1, then a failure to pass inspection indicates that the producer is incompetent. This should be contrasted with what happens when π^L is very close to 0, where passing inspection indicates that the producer is competent. Therefore, when π^H is very close to 1, the price in submarket p^{FP} is very close to zero, which undermines the incentives of producers who have failed inspection in the previous period to produce high quality.

4. The Length of Memory

The precision of information and the length of memory are substitutes: both provide more accurate information with respect to producers' competence. Therefore, in much the same way that too precise information can undermine the incentive to produce high quality as shown in the previous section, a longer memory can also undermine the incentive to produce high quality. We describe two results that illustrate this effect. We first show that as the length of memory is increased from one to two periods, the incentive to produce high quality becomes stronger when information is not very precise, but is undermined when it is more precise. Next, we show that as the length of memory increases beyond a certain threshold, then the incentive to produce high quality is undermined, and as the length of memory increases to infinity, the incentive to produce high quality is completely eliminated.

Consider a model that is identical to that described in Section 2, except that memory is only one period long. Consequently, instead of four submarkets, there are only two, one for those who passed inspection in the previous period, denoted P, and one for those who failed, denoted F. In a market equilibrium in which competent producers produce always produce high quality regardless of the submarket in which they find themselves in any given period, the measure of competent and incompetent producers in any period in submarket P is given by $(1-\eta)\,\pi^H$ and $\eta\pi^L$, respectively. The price in submarket P is therefore given by

$$p^P = \frac{(1-\eta)\,\pi^H}{(1-\eta)\,\pi^H + \eta\pi^L} \tag{10}$$

in every period $t \in \mathbb{I}$. Similarly, the price in submarket F is given by

$$p^F = \frac{(1-\eta)\left(1-\pi^H\right)}{(1-\eta)\left(1-\pi^H\right) + \eta\left(1-\pi^L\right)}. \tag{11}$$

By the analog of Lemma 1 for a 1-period memory model, a situation in which competent producers always produce high quality regardless of the submarket in which they find themselves in any given period is a market equilibrium if and only if

$$c \le \delta\pi\left(p^P - p^F\right), \tag{12}$$

and upon plugging (10) and (11) into (12), if and only if

$$c \le \frac{\delta\eta\left(\pi^H - \pi^L\right)^2}{\left(\pi^H + \eta\pi^L\right)\left(1 + \eta - \pi^H - \eta\pi^L\right)}. \tag{13}$$

Examination of inequality (13) reveals that if inspection is already precise, that is, if π^H and π^L are close to 1 and 0, respectively, then the RHS of inequality (13) is very large, which implies that the incentive to produce high quality is very strong. Because, as π^H and π^L converge to 1 and 0, respectively, the RHS of inequality (13) becomes arbitrarily large, increasing the length of memory from one to two periods would only weaken the incentive to produce high quality because, as can be readily verified, the RHS of inequalities (8) and (9), are bounded from above by $\delta(1+\delta)$ and $\delta(2-\delta)$, respectively.

On the other hand, if π^H and π^L are not close to 1 and 0, respectively, then it can be shown, numerically, that the RHS of inequality (13) is smaller than the RHS of inequalities (8) and (9), which implies that increasing the length of memory from one to two periods improves incentives in this case.

Consider now a model that is identical to the one described in Section 2 with an $n > 2$ periods long memory. In such a model, instead of 4, there are 2^n submarkets into which producers are sorted depending on whether or not they passed or failed inspection in the previous n periods. Consider the threshold cost $c_n^*\left(\pi^H, \pi^L, \eta, \delta\right)$ beyond which competent producers would not produce high quality in any submarket they happen to find themselves in any period in the equilibrium where they are always supposed to produce high quality. We show that this threshold c_n^* decreases to zero with n.

Recall that $\widehat{h}_{n,t}$ denotes an n-dimensional vector of passes and fails that ends in period t. Since, for the purpose of this section, the fact that $\widehat{h}_{n,t}$ ends in period t is unimportant, we drop the index t, and denote an n-dimensional vector of passes and fails by $\widehat{h}_n$. Let $\widehat{H}_n$ denote the set of all n-dimensional vectors of passes and fails. Recall that because market participants are assumed to have an n period memory, in any period t, all the producers who have the same realized profile of passes and fails in the last n periods are sorted into the same submarket at t. We can thus identify every submarket with an n-dimensional profile of passes and fails $\widehat{h}_n \in \widehat{H}_n$.

If all competent producers always produce high quality, regardless of the submarket in which they happen to find themselves in any given period, then the prices in every submarket remain constant, and do not change over time. We can therefore denote the price in submarket $\widehat{h}_n$ by $p^{\widehat{h}_n}$, independently of the period. If we let $\widehat{h}_n(P)$ denote the number of passes in the vector $\widehat{h}_n$, then Bayesian updating implies that

$$p^{\widehat{h}_n} = \frac{(1-\eta)\left(\pi^H\right)^{\widehat{h}_n(P)}\left(1-\pi^H\right)^{n-\widehat{h}_n(P)}}{(1-\eta)\left(\pi^H\right)^{\widehat{h}_n(P)}\left(1-\pi^H\right)^{n-\widehat{h}_n(P)} + \eta\left(\pi^L\right)^{\widehat{h}_n(P)}\left(1-\pi^L\right)^{n-\widehat{h}_n(P)}}$$

for every submarket $\widehat{h}_n \in \widehat{H}_n$.

The fact that $p^{\widehat{h}_n}$ is increasing in the number of passes implies that for producers to always produce high quality regardless of the submarket in which they happen to find themselves in any given period, is indeed a market-equilibrium, provided the cost of producing high quality, c, is sufficiently small.

In this equilibrium, if a competent producer were to produce low quality in some period t, and then to continue producing high quality thereafter, then the distribution of submarkets to which this producer would have access to in the following n periods, after which the effect of this single deviation would disappear, would be the same as the distribution of submarkets to which a producer who always produces high quality has access to, with the difference that whereas the producer who always produces high quality is likely to pass inspection at t, the producer who produces low quality at t is likely to fail inspection at t. It therefore follows that in order for producers to always produce high quality regardless of the submarket in which they happen to find themselves in any given period to be a market-equilibrium, it must be that such deviations are not profitable, or that

$$\frac{\delta}{1-\delta} \max_{h_n \in H_n, k \in \{1,...,n\}} \left\{ p^{\widehat{h}_n} - p^{\left(\widehat{h}_n : k \to F\right)} \right\} < c,$$

where $\left(\widehat{h}_n : k \to F\right)$ denotes a vector that is identical to $\widehat{h}_n$ except that it has a fail in the k-th place.

Lemma 1. *The maximum difference*

$$\max_{h_n \in H_n, k \in \{1,...,n\}} \left\{ p^{\widehat{h}_n} - p^{\left(\widehat{h}_n : k \to F\right)} \right\}$$

converges to zero as n increases.

Proof. Rewrite $p^{\widehat{h}_n}$ as

$$p^{\widehat{h}_n} = \frac{1}{1 + \frac{\eta}{(1-\eta)} \left(\frac{\pi^L}{\pi^H}\right)^{\widehat{h}_n(P)} \left(\frac{1-\pi^L}{1-\pi^H}\right)^{n-\widehat{h}_n(P)}}.$$

Suppose that $\widehat{h}_n(P) \in \{1, ..., n\}$ is such that $\frac{\widehat{h}_n(P)}{n}$ converges to some constant $\kappa \in [0, 1]$. Observe that

$$\begin{aligned} \lim_{n \nearrow \infty} \left(\frac{\pi^L}{\pi^H}\right)^{\widehat{h}_n(P)} \left(\frac{1-\pi^L}{1-\pi^H}\right)^{n-\widehat{h}_n(P)} &= \lim_{n \nearrow \infty} \left(\left(\frac{\pi^L}{\pi^H}\right)^{\kappa} \left(\frac{1-\pi^L}{1-\pi^H}\right)^{1-\kappa} \right)^n \\ &= \begin{cases} \infty & \text{if } \kappa < \kappa^* \\ 1 & \text{if } \kappa = \kappa^* \\ 0 & \text{if } \kappa > \kappa^* \end{cases} \end{aligned}$$

where $\kappa^* \equiv \frac{\log\left(\frac{1-\pi^L}{1-\pi^H}\right)}{\log\left(\frac{1-\pi^L}{1-\pi^H}\right) - \log\left(\frac{\pi^L}{\pi^H}\right)}$ is such that $\left(\frac{\pi^L}{\pi^H}\right)^{\kappa^*} \left(\frac{1-\pi^L}{1-\pi^H}\right)^{1-\kappa^*} = 1$. It therefore follows that

$$\lim_{n \nearrow \infty} p^{\widehat{h}_n} = \begin{cases} 0 & \text{if } \lim_{n \nearrow \infty} \frac{\widehat{h}_n(P)}{n} < \kappa^* \\ 1-\eta & \text{if } \lim_{n \nearrow \infty} \frac{\widehat{h}_n(P)}{n} = \kappa^* \\ 1 & \text{if } \lim_{n \nearrow \infty} \frac{\widehat{h}_n(P)}{n} > \kappa^* \end{cases}$$

Since the number of passes in the vector $\left(\widehat{h}_n : k \to F\right)$, denoted $\left(\widehat{h}_n : k \to F\right)(P)$, is between $\widehat{h}_n(P) - 1$ and $\widehat{h}_n(P)$, depending on whether $\widehat{h}_n$ has a pass or fail in the k-th place, $\lim_{n \nearrow \infty} \frac{\widehat{h}_n(P)}{n} = \kappa$ if and only if $\lim_{n \nearrow \infty} \frac{\left(\widehat{h}_n : k \to F\right)(P)}{n} = \kappa$, and so

$$\lim_{n \nearrow \infty} p^{\left(\widehat{h}_n : k \to F\right)} = \lim_{n \nearrow \infty} p^{\widehat{h}_n}.$$

■

Intuitively, the reason that the market-equilibrium where competent producers always produce high quality regardless of the submarket in which they happen to find themselves in any given period becomes impossible to sustain as the length of memory increases is that the market makes increasingly similar inferences about the competence of producers whose record differs by only one failure. This weakens the incentive to produce high quality in every period and undermines the equilibrium.

Remark 4. Speed of Convergence

Inspection of the proof of Lemma 2 reveals that both $p^{\widehat{h}_n}$ and $p^{\left(\widehat{h}_n : k \to F\right)}$ converge to their limits exponentially fast in n. This implies that implies that c_n^* decreases to zero exponentially fast with n.

The same intuition suggests that any equilibrium in which high quality is often produced would also become destabilized as the length of memory increases. However, we have not been able to formally establish such a result, and the question of what is the highest possible quality that can be sustained in a market-equilibrium as the length of memory tends to infinity is left as an interesting open problem.

5. Discussion

Recently, testing and general dissemination of test result have become very popular for students, teachers, caregivers, doctors, schools, nursing homes, and for other professions and for other institutions. The results reported in this study suggest that increased reliance on testing to improve incentives may fall short of expectations, and may even weaken incentives.

There are very few empirical studies of the benefit of testing. Presumably, the subject received little attention because it is difficult to establish objective measures for quality. Jin and Leslie (2003) showed that a Los Angeles county requirement that restaurants post hygiene quality grade cards on their windows led to an increase in restaurants health inspection scores and to a decrease in the number of foodborne illness hospitalizations, which suggests that food quality has improved. Dranove et al. (2001) showed that doctors who are required to post their health care report cards tend to decline to treat more difficult, severely ill, patients. Consequently, health report cards may lead to a decrease in healthcare quality. Chipty (1995) exploited the cross-state variation in the choice of day-care regulations to identify the effect of regulation on the performance of the day-care market. She found that an increase in mandated annual inspections decreased equilibrium quality (as measured by staff/child ratios) for family day-care. Rosenthal (2004) has examined the effect of school inspections on the observed exam performance of the state secondary schools in the UK and concluded that inspection had a small but well-determined adverse effect on inspected schools. Finally, Clark and Tomlinson (2001) reported that the extent of monitoring does not seem to affect workers' effort levels based on employees' self-reported effort levels from the 1992 Survey of Employment in Britain. It thus appears that the evidence is consistent with the notion that the effect of improved inspection on outcomes is ambiguous.

We conclude with the following anecdotal evidence about the effect of safety regulation. In the U.S. the Occupational Safety and Health Administration (OSHA) requires employers to comply with a large number of regulations whose purpose is to ensure the safety and health of employees. OSHA routinely monitors violations of its regulations through surprise inspections, and fines those employers that are found to be violating its regulation. In order to relate our theory to the data provided by OSHA, suppose that an employer who has been found to violate OSHA regulation is perceived as riskier by employees, and holding every thing else fixed, it has to pay higher wages to its employees. That is, we assume that two employers with the same safety record would pay their employees different wages depending on whether they have been fined by OSHA or not.

Suppose that the probability of detection of a safety violation is increasing in the number of annual inspections, and that "average safety" can be measured by the average number of violations per inspection. We are interested in relationship between the number of inspections and the number of average violations per inspection over time. Is it the case that as the number of inspections increases, the number of average violations per inspection increases too, as would be the case if increasing the number of inspections weakens the incentives to maintain safety?

This question can be answered using the data provided by OSHA on its homepage. OSHA's executive summary of its 20th century enforcement data reports that the history of OSHA regulation can be roughly divided into five period.

> "The first period, from the formation of OSHA in 1971 to 1976, was characterized by rapid growth in staff, inspections, violations, and penalties. During the second period, from 1977 to 1980, the agency revised its enforcement program to focus on inspection quality rather than quantity, instituted a new complaint policy and revised the penalty assessment methodology. Total inspections dropped dramatically during this period. During the 1981-1985/1986 period OSHA focused its efforts on 'cooperation rather than confrontation.' The agency reduced its enforcement staff, implemented 'records-review-only' inspections, and shifted towards more construction industry and small establishment inspections. As a result, inspection numbers increased while violations and penalties decreased. During the 1985/1986-1991 period, the agency instituted the egregious case pol-

icy, eliminated the 'records-review-only' inspections, and reemphasized quality inspections. These changes resulted in a large decline in inspections but a sharp increase in violations and penalties."

Because for the first four periods, from 1971 to 1991, the nature of OSHA inspections changed from one period to the next, the resulting changes in the numbers of violations per inspection cannot be interpreted as indicative of changes in average safety. However, the nature of inspections remained more or less the same over the 1990s. OSHA's report concludes by mentioning that "For the year 1992 through 2000, the number of inspections conducted by OSHA declined by about 14 percent and the number violations dropped by about 48 percent (compared with the year 1991)." Thus over the 1990s, both the number of inspections and the number of violations per inspection, which is inversely related to average safety, have decreased. Although the number of inspections has gone down, average safety seems to have improved, as would be the case if the precision of information about employers' safety records was already past the threshold beyond which any further improvement would hurt the incentives to improve safety.

References

Clark, K. and M. Tomlinson (2001) "The determinants of work effort: evidence from the Employment in Britain Survey," University of Manchester School of Economics Discussion Paper No. 0113.

Dewatripont, M., I. Jewitt, and J. Tirole (1999) "The Economics of Career Concerns, Part I: Information Structures," *Review of Economic Studies* 66, 183-198.

Dranove, D., D. Kessler, M. McClellan, and M. Satterthwaite (2001) "Is More Information Better? The Effects of Health Care Quality Report Cards" *Journal of Political Economy* 111, 555-88.

Ghosh, P. and D. Ray (1996) "Cooperation in Community Interaction without Information Flows," *Review of Economic Studies* 63, 491-519.

Holmström, B. (1999) "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies* 66, 169-182.

Mailath, G. J. and L. Samuelson (2001) "Who Wants a Good Reputation?" *Review of Economic Studies* 68, 415-441.

Rosenthal, L. (2004) "Do School Inspections Improve School Quality? Ofsted Inspections and School Examination Results in the UK," *Economics of Education Review* 23, 143-151.

Tadelis, S. (1999) "What's in a Name? Reputation as a Tradable Asset," *American Economic Review* 89, 548-563.

Tadelis, S. (2002) "The Market for Reputations as an Incentive Mechanism," *Journal of Political Economy* 92, 854-882.

Tadelis, S. (2003) "Firm Reputations with Hidden information," *Economic Theory* 21, 635-651