THE JOURNAL OF INDUSTRIAL ECONOMICS 0022-1821
Volume LVIII September 2010 No. 3

THE QUALITY OF INFORMATION AND INCENTIVES FOR EFFORT*

Omer Moav†
Zvika Neeman‡

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 can be nonmonotonic. There exists an efficient plausible equilibrium that induces a threshold beyond which any further improvement in the precision of information weakens the agent's incentive to produce high quality. Accordingly, both very accurate and very inaccurate signals about the agent's performance may destroy its incentive to exert effort. A few applications of this result are discussed.

I. INTRODUCTION

This paper concerns the relationship between the precision of public 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 may be nonmonotonic. There may exist a threshold beyond which any further improvement in the precision of information weakens the agent's incentive to produce high quality. Accordingly, both imprecise and very precise public 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 contractible and is not observable to the consumer at the time of purchase. Such goods are referred to as experience or credence goods in the literature.[1]

*We thank Elchanan Ben-Porath, Asher Wolinsky, and seminar participants at Barcelona (JOCS), Boston University, Duke/UNC, and the Hebrew University for their helpful comments. Financial support from the Falk Institute for Economic Research is gratefully acknowledged.

†Authors' affiliations: Department of Economics, Hebrew University of Jerusalem, Mount Scopus, Jerusalem 91905, Israel; Department of Economics, Royal Holloway, University of London, Egham, Surrey, England; and CEPR.
e-mail: omer.moav100@gmail.com

‡The Eitan Berglas School of Economics, Tel-Aviv University, Tel-Aviv 69978, Israel.
e-mail: zvika@post.tau.ac.il

[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 fully revealed to consumers.

Examples range from food and wine to used cars and expert advice. Importantly, because consumers cannot contract on the quality of the good, the price they are willing to pay depends on their beliefs about the good's expected quality. These beliefs are identified in our model with the reputation of the good's producer. A producer's reputation, in any given period, is determined by consumers' beliefs regarding the producer's type, taking into account the producer's incentives to exert effort to produce high quality. The beliefs regarding each producer's type are updated based on publicly available information about the producer's past performance, and the producers' incentives to exert high effort are determined by the return to reputation that depends on expected future prices.

For any given prior beliefs about the producer's ability, the producer's incentive to exert costly effort in order to produce high quality is increasing with the probability that the true quality of the good will be revealed. Hence, for any given consumer's prior beliefs, an improvement in consumers' ability to detect the quality of a good has an obviously positive effect on the producer's incentives to produce high quality. 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, the consumers' prior beliefs are affected by the precision of their information.

If prior beliefs regarding a producer's ability are very precise, then a contradictory signal about the producer's ability is attributed to either sampling error or a random shock in the production process, and has only a small effect on the consumers' posterior beliefs. It follows that if the prior probability that a producer is competent is sufficiently high, then the producer can exert less effort and 'rest on its laurels' without incurring a significant loss of reputation. Similarly, if the prior probability that a producer is competent is sufficiently low, then it is difficult for the producer to significantly improve its reputation by the production of high quality, because, as before, a signal that is inconsistent with the producer's reputation is heavily discounted. Hence, because precise or concentrated priors are hard to change, precise information 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.

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 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 publicly available information would reduce the incentives to exert effort in order to produce high quality.[2]

The relevant literature can be divided into three parts as follows.

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 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.[3] 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 it has a high ability.

There are two main differences between Holmström's model and ours. First, in our model producers privately know their own types. In contrast, 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.[4] Second, in our model the agent's ability and effort are strategic complements. They are strategic substitutes in Holmström's model. 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

[2] It is well known that more precise *private information* may sometimes lead to a less efficient outcome. Consider for example a standard 'lemons' market in which sellers know the quality of what they sell but buyers do not. Suppose that the adverse selection in this market is so bad that no trade takes place in equilibrium. In this market, if sellers were uninformed about the quality of the goods sold, then the efficient outcome in which buyers and sellers trade at a price that is equal to the average value of the good would prevail. A number of papers (see, e.g., Sakai [1985]; Gal-Or [1988]; Mirman *et al*. [1994]; Harrington [1995]; and Schlee [1996]) describe environments in which *public information* about quality may sometimes have a negative value. However, the reason that public information may have a negative value in these models can be attributed to 'non-convexities' of some kind, which is very different from the reasons discussed in this paper.

[3] Gibbons and Murphy [1992] considered an extension of Holmström's model in which the agent's output is contractible. They showed that the optimal compensation contract in such a setting optimizes over the combination of the implicit incentives from career concerns and the explicit incentives from the compensation contract. As the agent approaches retirement, the explicit incentives induced by the optimal compensation contract become stronger to make up for the weaker career concerns of such an agent.

[4] Although the market is not able to observe the agent's effort directly in Holmström's model, it can infer it by solving the agent's optimization problem (Holmström [1999], p. 171).

above, in Holmström's model it unambiguously leads to stronger incentives to exert effort.

Dewatripont, Jewitt and Tirole [1999] build on Holmström's model to characterize information structures in terms of their effect on the agent's incentives.[5] 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 monotone likelihood ratio 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 high probability that any producer is incompetent to provide this 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

[5] See also Bar-Isaac and Ganuza [2008].

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.

In a related paper, Hörner [2002] has showed that competition among producers can also serve as a disciplining device as consumers would rationally abandon producers who fail them and switch to the competition.

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 develop when producers of a certain good occasionally exit the market and sell their reputations to new entrants to the market. They have shown 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 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 discuss a few extensions of the basic model. We conclude in Section 5 with a discussion of the possible implications of our analysis.

II. MODEL

We describe a simple model in which we can elucidate our main argument. A few extensions of the basic model are presented in Section 4. We consider a dynamic model of a market for an experience good. Time is discrete, and periods are indexed by $t \in \{\ldots, 0, 1, \ldots\}$. There is a continuum of measure one of producers. There are two types of producers, competent and incompetent. The measure of incompetent producers is given by $\eta \in (0, 1)$. Producers discount future payoffs at the rate $\delta < 1$.

In every period, each competent producer may either exert a costly effort, at cost $c \in (0, 1)$, to produce one unit of a high quality good, or it may costlessly produce one unit of a low quality good. Incompetent producers are incapable of producing high quality goods. However, they may each costlessly produce one unit of the low quality good in every period.[6]

[6] Assuming instead that the produced quality is stochastic so that a competent producer who incurs the cost of producing high quality sometimes produces low quality, and vice-versa, does not change our results.

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. In 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, where $0 < \pi^L < \pi^H \leq 1$. Whether or not each producer passes or fails inspection is public information.[7] This public information is forgotten after $n \geq 2$ periods. It should be noted that if $n = 1$ then there is a monotonic effect of precision on effort, because in this case producers cannot free ride on their past reputation. The rational for the limited memory of the market is that over time producers' types may change and so the memory of long past events is increasingly irrelevant. We discuss this issue further in Section 4.2.

It follows that in every period, producers are sorted into 2^n submarkets, depending on whether they passed or failed inspection in the previous n periods. Let h_n denote an n-dimensional vector of passes and fails, and let H_n denote the set of all such n-dimensional vectors. 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 some n-dimensional profile of passes and fails $h_n \in H_n$.

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, let $q_t^{h_n} \in \{High, Low\}$ denote the quality of the good produced by competent producers in submarket $h_n \in H_n$ at t, and let $p_t^{h_n}$ denote the price in submarket $h_n \in H_n$ at t. Thus, the information about the producers' last n passes and fails are encoded into market prices. Our notion of market-equilibrium is defined as follows.

Definition. A sequence of qualities and prices $\left\{ \left(q_t^{h_n}, p_t^{h_n} \right)_{h_n \in H_n} \right\}_t$ is a market-equilibrium if:

1. In every period t, producers produce the quality that maximizes the discounted value of their expected profits given the sequence of prices $\left\{ \left(p_t^{h_n} \right)_{h_n \in H_n} \right\}_t$. We assume that in case of indifference, producers produce high quality.
2. In every period t, the price in each submarket is equal to the expected quality of the good for consumers in that submarket given competent producers' levels of effort.

[7] The conclusions of the model remain qualitatively unchanged if a producer who has failed inspection is subject to a fine, provided, of course, that this fine is not so large as to cause producers to stay out of the market.

Remark 1. Existence of a Market-Equilibrium. The model admits the existence of at least one market-equilibrium. In particular, the sequence of qualities and prices $\left\{\left(q_t^{h_n}, p_t^{h_n}\right)_{h_n \in H_n}\right\}_t$ where $q_t^{h_n} = Low$ and $p_t^{h_n} = 0$ for every $h_n \in H_n$ and every t, 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.

Remark 2. Interpretation as a Model of Career Concerns. The description above has emphasized an interpretation of the model as that of a market for an experience good. However, the same assumptions also admit an interpretation of the model as that 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.

III. THE PRECISION OF INFORMATION

For a given market-equilibrium, *ceteris paribus*, 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 the threshold cost is above which competent producers may sometimes choose not to produce high quality although they are supposed to. The higher this threshold, the stronger is the incentive 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 not well defined because competent producers are not supposed to produce high quality, but in any other equilibrium this threshold may be positive.

Definition. The strength of incentives that is induced by a sequence of qualities $\left\{\left(q_t^{h_n}\right)_{h_n \in H_n}\right\}_t$ is equal to the threshold cost of producing high quality c below which the sequence of qualities $\left\{\left(q_t^{h_n}\right)_{h_n \in H_n}\right\}_t$ can be part of a market-equilibrium $\left\{\left(q_t^{h_n}, p_t^{h_n}\right)_{h_n \in H_n}\right\}_t$. When there is no ambiguity with respect to the market-equilibrium, this threshold cost is denoted $c_n^*(\pi^H, \pi^L, \eta, \delta)$.

We are interested in the magnitude of this threshold in sequences of qualities that are part of *stationary* equilibria.

Definition. A market-equilibrium $\left\{\left(q_t^{h_n}, p_t^{h_n}\right)_{h_n \in H_n}\right\}_t$ is stationary if the quality produced by competent producers and the price in every submarket h_n, $q_t^{h_n}$ and $p_t^{h_n}$, respectively, remain constant over time, or in other words if $q_t^{h_n}$ and $p_t^{h_n}$ are independent of t. A stationary market-equilibrium may therefore by described by just a vector of qualities and prices, $\left\{\left(q^{h_n}, p^{h_n}\right)_{h_n \in H_n}\right\}$

We pay special attention to the sequence of qualities and market-equilibrium where competent producers produce high quality in every submarket and after every history. Because production of high quality costs $c < 1$ but generates a value of one to consumers this is the most efficient sequence of qualities and market-equilibria out of all possible sequences of actions and market-equilibria, including non stationary ones.

Given a stationary market-equilibrium $\left\{\left(q^{h_n}, p^{h_n}\right)_{h_n \in H_n}\right\}$, let U^{h_n} denote the expected discounted payoff of a competent producer with a history $h_n \in H_n$ in any period t who proceeds to behave optimally in period t and afterwards.

The next lemma shows that the behavior of competent producers in a stationary market-equilibrium depends only on their history of passes and fails in the last $n-1$ periods.[8] It thus implies that the number of stationary market equilibria is equal to $2^{2^{n-1}}$ when the length of memory is given by n.

Lemma. In a stationary market-equilibrium $\left\{\left(q^{h_n}, p^{h_n}\right)_{h_n \in H_n}\right\}$, in every period t, a competent producer with a history h_n produces high quality if and only if

$$(1) \qquad c \leq \delta\left(\pi^H - \pi^L\right)\left(U^{h_n P} - U^{h_n F}\right)$$

where $h_n P$, $h_n F \in H_n$ denote histories whose $n-1$ first coordinates coincide with the last $n-1$ coordinates of h_n and that have a pass and fail, respectively, in the n-th coordinate.

[8] Inspection of the proof of the lemma reveals that it holds for all market equilibria, not just stationary market equilibria.

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

$$\delta\Big(\pi^H U_{t+1}^{h_nP} + \big(1-\pi^H\big)U_{t+1}^{h_nF}\Big) - c \geq \delta\Big(\pi^L U_{t+1}^{h_nP} + \big(1-\pi^L\big)U_{t+1}^{h_nF}\Big)$$

if and only if (1). ■

Because it is possible to express the main insights of this paper in a model where the length of consumers' memory is $n = 2$ periods, for simplicity, we restrict attention to this case for the rest of this section.

In the case where $n = 2$ there are four stationary market-equilibria: one where competent producers never produce high quality ((1) in Table I), one where competent producers always produce high quality after every history ((2) in Table I), one where competent producers produce high quality after they pass inspection but low quality after they fail inspection ((3) in Table I), and one where competent producers produce high quality after they fail inspection but low quality after they pass ((4) in Table I).

The stationary market-equilibrium (1) where producers never produce high quality is very inefficient. The stationary market-equilibrium (4) is also inefficient because if signals are sufficiently accurate, then competent producers produce low quality approximately every other period. The stationary market-equilibrium (2) is the most efficient equilibrium among all market-equilibria as explained above. The stationary market-equilibrium (3) might seem to be pretty efficient because a competent producer who produces high quality and passes inspection continues to produce high quality, but as we shall see below, this is not the case when inspection becomes more precise. The 'unforgiving nature' of this equilibrium (producers who failed inspection are expected to produce low quality) implies that competent producers who failed inspection are locked into producing low quality, and as inspection becomes more precise, the fraction of these producers does not decrease to zero.

TABLE I

STATIONARY MARKET-EQUILIBRIA FOR THE CASE WHERE $n = 2$ PP STANDS FOR A HISTORY WITH PASSES IN THE TWO LAST PERIODS, PF STANDS FOR A HISTORY WITH A PASS FOLLOWED BY A FAILURE IN THE TWO LAST PERIODS, AND SO ON.

submarket:	*PP*	*PF*	*FP*	*FF*
(1)	*Low*	*Low*	*Low*	*Low*
(2)	*High*	*High*	*High*	*High*
(3)	*High*	*Low*	*High*	*Low*
(4)	*Low*	*High*	*Low*	*High*

In the efficient stationary market-equilibrium (2) where competent producers exert high effort after every history, in every period t,

$$
(6)\qquad
\begin{aligned}
U^{PP} &= p^{PP} - c + \delta\big(\pi^H U^{PP} + \big(1-\pi^H\big)U^{PF}\big)\\
U^{PF} &= p^{PF} - c + \delta\big(\pi^H U^{FP} + \big(1-\pi^H\big)U^{FF}\big)\\
U^{FP} &= p^{FP} - c + \delta\big(\pi^H U^{PP} + \big(1-\pi^H\big)U^{PF}\big)\\
U^{FF} &= p^{FF} - c + \delta\big(\pi^H U^{FP} + \big(1-\pi^H\big)U^{FF}\big)
\end{aligned}
$$

It is possible explicitly to solve for $(U^{PP}, U^{PF}, U^{FP}, U^{FF})$ in the four linear equations in (6) in terms of the market prices $(p^{PP}, p^{PF}, p^{FP}, p^{FF})$. The solution is long and because it is not necessary for the rest of the analysis, is not presented here.

In the efficient stationary market-equilibrium, the measure of competent and incompetent producers in any period in submarket PP is given by $(1-\eta)\,(\pi^H)^2$ and $\eta\,(\pi^L)^2$, respectively. The price in submarket PP is therefore equal to

$$
(7)\qquad p^{PP} = \frac{(1-\eta)\,(\pi^H)^2}{(1-\eta)\,(\pi^H)^2 + \eta\,(\pi^L)^2}
$$

in every period. Similarly, the prices in submarkets PF, FP, and FF, are equal to

$$
(8)\qquad p^{PF} = p^{FP} = \frac{(1-\eta)\pi^H(1-\pi^H)}{(1-\eta)\pi^H(1-\pi^H) + \eta\pi^L(1-\pi^L)},
$$

and

$$
(9)\qquad p^{FF} = \frac{(1-\eta)\,(1-\pi^H)^2}{(1-\eta)\,(1-\pi^H)^2 + \eta\,(1-\pi^L)^2},
$$

respectively, in every period.

Lemma 1 implies that the threshold cost $c_2^*(\pi^H, \pi^L, \eta, \delta)$ in this equilibrium is equal to:

$$
(10)\qquad
\begin{aligned}
c_2^*\big(\pi^H,\pi^L,\eta,\delta\big) &= \delta\big(\pi^H-\pi^L\big)\min\big\{U^{PP}-U^{PF}, U^{FP}-U^{FF}\big\}\\
&= \delta\big(\pi^H-\pi^L\big)\min\left\{\begin{array}{l}(1+\delta\pi^H)p^{PP}-(1-\delta+2\delta\pi^H)p^{PF}-\delta(1-\pi^H)p^{FF}\\ \delta\pi^H p^{PP}+(1+\delta-2\delta\pi^H)p^{PF}-(1+\delta(1-\pi^H))p^{FF}\end{array}\right\}.
\end{aligned}
$$

This equation follows from (1), (6), and the observation that $p^{PF} = p^{FP}$, and that (6) also implies that $U^{PP} - U^{FP} = p^{PP} - p^{FP}$, and $U^{PF} - U^{FF} = p^{PF} - p^{FF}$.

For a fixed π^L,

$$
(11)\quad
\begin{aligned}
&\lim_{\pi^H \nearrow 1} p^{PP} = \frac{1-\eta}{1-\eta+\eta\,(\pi^L)^2} &&\text{and}\quad \lim_{\pi^H \nearrow 1} \tfrac{dp^{PP}}{d\pi^H} = \frac{2(1-\eta)\eta\,(\pi^L)^2}{\left(1-\eta+\eta\,(\pi^L)^2\right)^2};\\
&\lim_{\pi^H \nearrow 1} p^{PF} = 0 &&\text{and}\quad \lim_{\pi^H \nearrow 1} \tfrac{dp^{PF}}{d\pi^H} = -\frac{1-\eta}{\eta\pi^L(1-\pi^L)};\\
&\lim_{\pi^H \nearrow 1} p^{FF} = 0 &&\text{and}\quad \lim_{\pi^H \nearrow 1} \tfrac{dp^{FF}}{d\pi^H} = 0.
\end{aligned}
$$

Therefore for a fixed π^L, as π^H tends to 1, (10) implies that $c_2^*(\pi^H, \pi^L, \eta, \delta)$ tends to $\delta^2(1-\pi^L)p^{PP}$ *from above* because when π^H tends to 1 the minimum in (10) is obtained on the second term, and as implied by (11) the limit of $\frac{dp^{PF}}{d\pi^L}$ becomes arbitrarily larger than $\frac{dp^{FF}}{d\pi^L}$ in absolute value as π^L becomes closer to zero.

Similarly, for a fixed π^H,

$$
(12)\quad
\begin{aligned}
&\lim_{\pi^L \searrow 0} p^{PP} = 1 &&\text{and}\quad \lim_{\pi^L \searrow 0} \tfrac{dp^{PP}}{d\pi^L} = 0;\\
&\lim_{\pi^L \searrow 0} p^{PF} = 1 &&\text{and}\quad \lim_{\pi^L \searrow 0} \tfrac{dp^{PF}}{d\pi^L} = -\tfrac{\eta}{(1-\eta)\pi^H(1-\pi^H)};\\
&\lim_{\pi^L \searrow 0} p^{FF} = \tfrac{(1-\eta)\left(1-\pi^H\right)^2}{(1-\eta)\,(1-\pi^H)^2+\eta} &&\text{and}\quad \lim_{\pi^L \searrow 0} \tfrac{dp^{FF}}{d\pi^L} = \tfrac{2\eta(1-\eta)\left(1-\pi^H\right)^2}{\left((1-\eta)\,(1-\pi^H)^2+\eta\right)^2}.
\end{aligned}
$$

Therefore for a fixed π^H, as π^L tends to 0, $c_2^*(\pi^H, \pi^L, \eta, \delta)$ tends to $\delta^2\pi^H(1-\pi^H)(1-p^{FF})$ from above because when π^L tends to 0 the minimum in (10) is obtained on the first term, and as implied by (12) the limit of $\frac{dp^{PF}}{d\pi^L}$ becomes arbitrarily larger than $dp^{FF}d\pi^L$ in absolute value as π^H becomes closer to one.

We summarize these results in the following proposition.

Proposition 1. Consider the efficient stationary market-equilibrium in the model with a 2-period memory, where competent producers produce high quality after every history. For every fixed values of $\eta \in (0, 1)$ 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 the threshold $c_2^*(\pi^H, \pi^L, \eta, \delta)$.

Intuitively, what is going on is that as π^L decreases to zero, so few incompetent producers pass inspection that both p^{PF} and p^{PP} converge to 1. A competent producer who has passed inspection in the previous period (and is therefore 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 decreases 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 in the next period it can then 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}, in the period after that. This implies that unless the cost of producing high quality c is very low, such a producer would indeed produce low quality. It follows that in order to sustain the efficient market equilibrium, the cost of producing high quality c has to decrease as π^L decreases to zero. More formally, the binding incentive constraint in this case is inequality (1) with $h_n \in \{PP, FP\}$, 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 similar. In this case, it is inequality (1) with $h_n \in \{PF, FF\}$ 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. 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. It follows that unless the cost of producing high quality c is very low, a competent producer who has failed the last inspection would produce low quality, which, in turn, implies that in order to to sustain the efficient market equilibrium, the cost of producing high quality c has to decrease as π^H increases to one.

Remark 3. Calculation of $c_n^*(\pi^H, \pi^L, \eta, \delta)$ *for* $n > 2$. In principle, it is possible to explicitly calculate $c_n^*(\pi^H, \pi^L, \eta, \delta)$ for every n in the same way it was calculated above for the case where $n = 2$. However, since the dimensionality of the calculation increases with n, such a calculation becomes rather tedious already with $n = 3$. Nevertheless, it is still possible to generalize Proposition 1 for any $n \geq 3$ in a way that does not involve an explicit calculation of $c_n^*(\pi^H, \pi^L, \eta, \delta)$. The argument can be obtained from the authors upon request.

Proposition 1 shows that the efficient stationary market-equilibrium (2) suffers from a weakness: it induces weaker incentives to produce high quality as information becomes more precise. This observation raises the question of what the relationship is between the precision of information and the strength of incentives in other stationary market-equilibria, and especially in stationary market-equilibrium (3), which, as explained above, seems to be rather efficient.

More generally, the question is whether an 'unforgiving' market equilibrium in which producers are 'penalized' for failing inspection by the expectation that they would produce low quality for at least one period or more as in market-equilibrium (3), would not induce weaker incentives as information becomes more precise, and not be that much more inefficient than the 'forgiving' efficient market-equilibrium in which competent producers are expected to continue to produce high quality after every history[9].

We answer this question below for the case where $n = 2$ by showing that market-equilibrium (3) is in fact very inefficient, as inefficient as market equilibrium (4).

Denote the measure of competent producers in the stationary market-equilibrium (3) in submarkets PP, FP, PF, and FF by μ^{PP}, μ^{FP}, μ^{PF}, and μ^{FF}, respectively. The behavior of competent producers in this market equilibrium implies that

$$\begin{aligned}
\mu^{PP} &= \pi^{H}\left(\mu^{PP} + \mu^{FP}\right) \\
\mu^{FP} &= \pi^{L}\left(\mu^{PF} + \mu^{FF}\right) \\
\mu^{PF} &= \left(1 - \pi^{H}\right)\left(\mu^{PP} + \mu^{FP}\right) \\
\mu^{FF} &= \left(1 - \pi^{L}\right)\left(\mu^{PF} + \mu^{FF}\right).
\end{aligned}$$

Solving this system of linear equations for the values of μ^{PP}, μ^{FP}, μ^{PF}, and μ^{FF} we obtain:

$$\begin{aligned}
\mu^{PP} &= \frac{(1-\eta)\pi^{H}\pi^{L}}{1 - \pi^{H} + \pi^{L}} \\
\mu^{FP} = \mu^{PF} &= \frac{(1-\eta)\pi^{L}(1-\pi^{H})}{1 - \pi^{H} + \pi^{L}} \\
\mu^{FF} &= \frac{(1-\eta)(1 - \pi^{H} - \pi^{L} + \pi^{H}\pi^{L})}{1 - \pi^{H} + \pi^{L}}
\end{aligned}$$

As information becomes more precise, or as $\pi^{L} \searrow 0$ and $\pi^{H} \nearrow 1$, the measure of competent producers in submarkets FP and PF decreases to zero, while the measures of competent producers in submarkets PP and FF are of the same order of magnitude. For example, if we let $\pi^{L} = 1/k$ and $\pi^{H} = 1 - 1/k$ and let k increase to infinity, then both μ^{PP} and μ^{FF} would tend to $(1-\eta)/2$, which implies that competent producers would spend approximately half their time inefficiently producing low quality goods, as in market-equilibrium (4). This limit inefficiency result depends on the relative speed of convergence of π^{H} and π^{L} to their respective limits. If π^{L} decreases to zero faster or at the same speed as π^{H}, then equilibrium (3) is inefficient in the limit. But if π^{L} decreases to zero at a slower rate than π^{H}, then equilibrium (3) is efficient in the limit. Thus, for 'most' sequences (more than half of the sequences), equilibrium (3) is indeed inefficient in the limit. Of course,

[9] We are grateful to a referee for raising this point.

equilibrium (3) is less efficient than equilibrium (2), which is the focus of our analysis, for any values of π^H and π^L.

This observation for the case where $n = 2$ leads to the following conjecture. For any $n \geq 3$, the set of stationary market equilibria can be divided into two subsets: one subset includes all the 'forgiving' stationary market-equilibria where producers are 'forgiven' for failing inspection if they passed many inspections in the past and are still expected to produce high quality goods in the future. Such forgiving equilibria are efficient, but for the same reason as explained above, induce weaker incentives to produce high quality goods as information becomes more precise because they induce competent producers who have passed many inspections to rest on their laurels. The other subset includes all the 'unforgiving' stationary market equilibria where producers are penalized for failing inspection by the belief that they would continue to produce low quality goods for at least a number of periods. Such equilibria induce strong incentives for producing high quality goods, but are possibly very inefficient in a world in which inspection is subject to errors. There may also exist 'quasi-forgiving' or 'quasi-unforgiving' equilibria where producers who fail inspection after passing many inspections would be expected to produce high quality in a large fraction but not all future periods. Such equilibria could possibly combine the efficiency of the forgiving equilibria with the good incentives induced by the unforgiving equilibria. However, investigation of this class of equilibria and the general question of what is the 'best market equilibrium' is beyond the scope of this paper.

IV. EXTENSIONS

IV(i). *The Length of Memory*

The precision of information and the length of memory are substitutes: both provide more precise 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 show that as the length of memory increases beyond a certain threshold, then the incentive to produce high quality in the efficient market equilibrium is undermined, and as the length of memory tends to infinity, the incentive to produce high quality is completely eliminated.

We show that the threshold cost in the efficient stationary market equilibrium, $c_n^*(\pi^H, \pi^L, \eta, \delta)$, beyond which the efficient market equilibrium cannot be sustained decreases to zero as n increases.

A version of the 'one-stage-deviation principle' (Fudenberg and Tirole [1991] pp. 108–110) implies that the efficient market-equilibrium can be sustained if and only if no producer in any submarket can benefit from a single deviation in which it produces low quality once and then high quality thereafter.

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

$$(13) \qquad p^{h_n} = \frac{(1-\eta)\left(\pi^H\right)^{h_n(P)}\left(1-\pi^H\right)^{n-h_n(P)}}{(1-\eta)\left(\pi^H\right)^{h_n(P)}\left(1-\pi^H\right)^{n-h_n(P)} + \eta\left(\pi^L\right)^{h_n(P)}\left(1-\pi^L\right)^{n-h_n(P)}}$$

for every submarket $h_n \in H_n$.

In the efficient stationary market 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 put a relatively bigger weight on submarkets with a larger number of fails and a smaller number of passes. It therefore follows that in order for producers always to produce high quality, regardless of the submarket in which they happen to find themselves, to be an equilibrium, it must be that such deviations are not profitable, or that the cost of a deviation

$$(14) \quad \delta \max_{h_n \in H_n, k \in \{1,\ldots,n\}} \left\{p^{h_n} - p^{(h_n:k\to F)}\right\} + \cdots + \delta^n \max_{h_n \in H_n, k \in \{1,\ldots,n\}} \left\{p^{h_n} - p^{(h_n:k\to F)}\right\}$$

where $(h_n\colon k \to F)$ denotes a vector that is identical to h_n except that it has a fail in the k-th place is larger than the benefit of a deviation, c.

The next proposition shows that such deviations become more and more attractive as n increases.

Proposition 2. The cost of a one-time deviation (14) converges to zero as n increases.

The proof of the proposition is available on the *Journal's* editorial website, and from the authors on request. Intuitively, the reason that the efficient market-equilibrium becomes impossible to sustain as the length of memory increases is that as n increases, it becomes clearer whether any producer is competent or not. Prices in submarkets with histories that suggest that the producers there are competent converge to one, and prices in submarkets with histories that suggest that the producers there are incompetent converge to zero. It follows that as n increases the market makes increasingly similar inferences about the competence of producers whose record differs by only one failure. This implies that competent producers with a good record of passes may rest on their laurels and produce low quality without seriously damaging their reputations. This weakens the incentive to produce high quality in every period and undermines the efficient market-equilibrium.

Remark 4. Speed of Convergence. Inspection of the proof of Proposition 2 reveals that both p^{h_n} and $p^{(h_n:k\to F)}$ converge to their limits exponentially fast in n. This implies that c_n^* decreases to zero exponentially fast with n.

Remark 5. Other Equilibria. The same intuition suggests that any equilibrium in which high quality is produced often would also be 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 n tends to infinity remains an open problem.

IV(ii). *Endogenizing the Length of Memory, n*

Many share the intuition that whatever happened in the distant past is of little relevance in the present. In the context of the model presented in Section 2, such an intuition implies that there is no need to consider the case where n is very large. In the context of this paper, an explanation of why such an intuition may be justified may proceed along the following lines: producers' abilities are subject to random shocks. It therefore follows that there is little reason to believe that there is any relationship between the types of a producer in periods t and t' if t and t' are very far from each other.

More formally, consider a model that is identical to the one presented in Section 2, except that each producer draws a new type (competent with probability $1-\eta$ and incompetent η) every k periods, where k is uniformly distributed over the set $\{1,\ldots,K\}$ for some $K\geq 2$, independently across different producers and over each producer's personal history. Such an assumption implies:

1. that whatever happened more than K periods earlier is of no relevance for a producer's reputation, and
2. that what has happened k periods earlier is more relevant for a producer's reputation than what has happened $k+1$ periods earlier, for every $k\in\{1,\ldots,K-1\}$. This observation is a consequence of the fact that whatever happened in the previous period is relevant for the reputation of all the producers who did not draw a new type in the current period; that whatever happened two periods earlier is relevant for the reputation of all the producers who did not draw a new type either in the current period or in the previous period; and so on.

The effect of this assumption on the prices in the different submarkets is that prices would put relatively more weight on what has happened in the recent past relative to what has happened in the more distant past. However, the properties of these prices, and the way they would respond to changes in the values of π^H and π^L remains qualitatively unchanged.[10] This implies, in particular, that our result about the negative effect of more precise information on incentives reported in Proposition 1 would hold in this case as well.

IV(iii). *Nonstationary Equilibria*

We believe that in a model with bounded memory, it makes no sense to consider non-stationary equilibria. By considering stationary equilibria we analyze situations that have been going on forever and that are expected to continue forever. For example, a nonstationary trigger-strategy equilibrium where consumers stop purchasing from a producer who failed inspection would require that we specify a first period of the model, and we don't think that such a first period exists in the type of situations that are analyzed in this paper. It is true that if such a first period is fixed, then the expected discounted efficiency losses in the trigger-strategy equilibrium would be small. But such a trigger-strategy equilibrium makes no sense if one is reluctant to specify any specific period as the 'first-period' because then the 'first time that a producer fails inspection,' which is necessary for such an equilibrium, is not well defined.

V. DISCUSSION

Recently, testing and the general dissemination of the results of such testing 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 benefits of testing. 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 had improved. Dranove *et al.* [2001] showed that doctors who were required to post their health care report cards tended 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

[10] Observe that the properties of prices that were used to prove Proposition 1 involve the convergence of prices to 1 and 0 as π^L tends to zero and as π^H tends to one, respectively, and the limit values of the derivatives of prices with respect to π^L and π^L, respectively.

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 U.K. 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 did 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 everything else fixed, it has to pay higher wages to its employees. That is, we assume that two employers 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 the 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 nature of OSHA's inspections remained more or less the same over the 1990's. OSHA's report concludes by mentioning that 'For the years 1992 through 2000, the number of inspections conducted by OSHA declined by about 14 per cent and the number violations dropped by about 48 per cent (compared with the year 1991).' Thus over the 1990's, 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

Bar-Isaac, H. and Ganuza, J.-J., 2008, 'Recruitment, Training and Career Concerns,' *Journal of Economics & Management Strategy*, 17, pp. 839–864.

Clark, K. and Tomlinson, M., 2001 'The determinants of Work Effort: Evidence from the Employment in Britain Survey,' University of Manchester School of Economics discussion paper No. 0113, Manchester, England.

Dewatripont, M.; Jewitt, I. and Tirole, J., 1999, 'The Economics of Career Concerns, Part I: Information Structures,' *Review of Economic Studies*, 66, pp. 183–198.

Dranove, D.; Kessler, D.; McClellan, M. and Satterthwaite, M., 2001, 'Is More Information Better? The Effects of Health Care Quality Report Cards,' *Journal of Political Economy*, 111, pp. 555–588.

Fudenberg, D. and Tirole, J., 1991, *Game Theory* (MIT Press, Cambridge, Massachusetts, U.S.A.).

Gal-Or, E., 1988, 'The Advantages of Imprecise Information,' *Rand Journal of Economics*, 19, pp. 266–275.

Ghosh, P. and Ray, D., 1996, 'Cooperation in Community Interaction without Information Flows,' *Review of Economic Studies*, 63, pp. 491–519.

Gibbons, R. and Murphy, K. J., 1992, 'Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence,' *Journal of Political Economy*, 100, pp. 468–505.

Harrington, J., 1995, 'Experimentation and Learning in a Differentiated Products Duopoly,' *Journal of Economic Theory*, 66, pp. 275–288.

Holmström, B., 1999, 'Managerial Incentive Problems: A Dynamic Perspective,' *Review of Economic Studies*, 66, pp. 169–182.

Hörner, J., 2002, 'Reputation and Competition,' *American Economic Review*, 92, pp. 644–663.

Mailath, G. J. and Samuelson, L., 2001, 'Who Wants a Good Reputation?,' *Review of Economic Studies*, 68, pp. 415–441.

Mirman, L.; Samuelson, L. and Schlee, E., 1994, 'Strategic Information Manipulation Duopolies,' *Journal of Economic Theory*, 62, pp. 363–384.

Rosenthal, L., 2004, 'Do School Inspections Improve School Quality? Ofsted Inspections and School Examination Results in the U.K.,' *Economics of Education Review*, 23, pp. 143–151.

Sakai, Y., 1985, 'The Value of Information in a Simple Duopoly Model,' *Journal of Economic Theory*, 36, pp. 36–54.

Schlee, E., 1996, 'The Value of Information about Product Quality,' *Rand Journal of Economics*, 27, pp. 803–815.

Tadelis, S., 1999, 'What's in a Name? Reputation as a Tradable Asset,' *American Economic Review*, 89, pp. 548–563.

Tadelis, S., 2002, 'The Market for Reputations as an Incentive Mechanism,' *Journal of Political Economy*, 92, pp. 854–882.

Tadelis, S., 2003, 'Firm Reputations with Hidden Information,' *Economic Theory*, 21, pp. 635–651.