

Köhler, Wolfgang R.

Working Paper

Optimal Incentive Contracts for Experts

Bonn Econ Discussion Papers, No. 6/2004

Provided in Cooperation with:

Bonn Graduate School of Economics (BGSE), University of Bonn

Suggested Citation: Köhler, Wolfgang R. (2004) : Optimal Incentive Contracts for Experts, Bonn Econ Discussion Papers, No. 6/2004, University of Bonn, Bonn Graduate School of Economics (BGSE), Bonn

This Version is available at:

<https://hdl.handle.net/10419/22884>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

BONN ECON DISCUSSION PAPERS

Discussion Paper 6/2004

Optimal Incentive Contracts for Experts

by

Wolfgang R. Köhler

April 2004



Bonn Graduate School of Economics
Department of Economics
University of Bonn
Adenauerallee 24 - 42
D-53113 Bonn

The Bonn Graduate School of Economics is
sponsored by the

Deutsche Post  World Net

MAIL EXPRESS LOGISTICS FINANCE

Optimal Incentive Contracts for Experts*

Wolfgang R. Köhler[†]

Abstract

This paper analyzes optimal incentive contracts for information acquisition and revelation. A decision maker faces the problem to design a contract that provides an expert with incentives to acquire and reveal information. We show that it is in general **not** optimal to reward the expert if his recommendation is confirmed. The common observation that experts are paid when their recommendation is confirmed can be explained by incomplete information about the expert's cost to increase the precision of his information. We extend the model to analyze contracting with multiple experts, the timing of expertise, and the provision of incentives when the realized state is not verifiable.

*Part of this research was supported by the National Science Foundation (grant SES-0136817). I would like to thank Curtis Taylor and Huseyin Yildirim.

[†]University of Bonn, Institut für Wirtschaftstheorie II, Lennestr. 37, 53113 Bonn, Germany, e-mail: wolfgang.koehler@uni-bonn.de

1. Introduction

It is a common feature of economic problems that actions have to be taken before the exact circumstances that determine the outcome are known. In many situations, there are potentially huge costs of a mismatch between action and the state that is realized after the action has been taken. In these situations, the decision maker would benefit from access to better information. Even when it is impossible to completely resolve the uncertainty about the future, in many situations players can acquire additional information besides what is publicly available. Acquiring information is costly and it is not necessarily the decision maker who has the lowest cost. In many professions exist experts who specialize in forecasting, presumably because they have access to a more efficient technology to acquire information or due to economies of scale. The best known examples of forecasts are probably weather forecasts and predictions of macroeconomic variables. Besides that, there exists a variety of financial forecasts, estimates about the demand for new products or the cost of new technologies, predictions of the outcome of elections, the effect of tax reforms, etc. In these examples, the one who initiates the forecast and whose decision is based on the forecast, is usually not the one who acquires the information. But if information acquisition is delegated to experts, the problem arises how to align the incentives for information acquisition and for truthful information revelation.

We study optimal incentive contracts for information acquisition and revelation. This paper contributes to a growing literature that concentrates on asymmetries in the cost of acquiring information rather than on asymmetric information. Specifically, we consider the following situation: A decision maker has to take some action before the state of the world is realized. The better the action matches the state, the higher is the decision maker's payoff. Instead of basing the choice of the action solely on the common prior, the decision maker can contract with an expert to gather additional information. The decision maker (principal) faces the problem to design a contract that provides the expert (agent) with incentives to increase the precision of his information and

to reveal the information.

When effort is not observable, incentives for information acquisition can only be provided if the agent's compensation depends on some variable that is correlated with the realized state and on the report of the agent. Under complete information, a report is equivalent to a recommendation to implement a certain action. We show that it is in general not optimal to reward the agent when his recommendation is confirmed by the facts. Under the optimal contract, the agent is (at least sometimes) rewarded when his recommendation turns out to be wrong but is not rewarded when the recommendation was correct. The optimal action and, if the agent reveals his information, the recommendation are functions of the agent's posterior. In general, a recommendation is confirmed if the realized state is equal to the agent's posterior mean. When the agent's information has the form of a signal and signal and state are normally distributed, the agent is not rewarded when the recommendation is confirmed and the state is equal to the agent's posterior mean but instead if the state is equal to the signal.

The observation that in reality many experts are rewarded if their recommendations are confirmed can be explained by incomplete information about the agent's cost to exert effort. When the principal knows the agent's cost to exert effort, for a given contract, she can compute the equilibrium precision of the agent's signal. In this case, the revelation of the agent's private information is equivalent to the recommendation of an action. Under incomplete information about the agent's cost to exert effort, a recommendation contains more information than the revelation of the private information of the agent. On the other hand, if the agent recommends an action and is rewarded if the recommendation is confirmed, he has less incentives to increase the precision of his signal. When there is a lot of uncertainty about the cost to exert effort, the first effect (more information) dominates the second (less incentives).

We use the results about optimal incentive contracts to analyze the organization of expertise. The optimal number of agents depends on the cost of information acquisition. Only if the marginal cost of precision increases, it is optimal to contract with as many agents as possible. In many

situations, it is possible to gather information either sequentially or simultaneously. We show that regardless of the cost function, the timing of expertise does not affect the payoffs and, therefore, does not affect the equilibrium amount of information that is acquired. Finally, we analyze the optimal contract when the realized state is not observable or not verifiable. In this case, the principal contracts with two agents and the wage depends on the reports of both agents. Similar to the case where the state is observable, an agent is not rewarded if his report is confirmed by the other agent. Instead, an agent is rewarded if the other agent reports a signal which is more extreme.

There exists an extensive literature that analyzes and tests different theories about the behavior of forecasters¹. While these papers are interested in information revelation, they do not model the decision to acquire information and do not consider the design of incentive contracts.

Ottaviani and Sørensen analyze the strategic behavior of forecasters who are endowed with a signal. Ottaviani and Sørensen (2003) analyze a situation where forecasters compete in a rank-order contest where prizes are exogenously determined. In the presence of prior information, forecasts are excessively differentiated in the sense that the posterior mean lies between the prior mean and the forecast. Ottaviani and Sørensen (2004) model forecasting as a reputational cheap-talk game. Similar to Ottaviani and Sørensen, Krishna and Morgan (2001) do not consider information acquisition or the design of incentive contracts. They analyze strategic information revelation of biased experts who are perfectly informed about the true state.

Similar to this paper, Osband (1989) is interested in the optimal contract when the principal observes neither the agent's information nor the effort that determines the precision of the information. His setting differs in that he assumes that the agent is not wealth-constrained. Therefore, it is optimal to sell the project to the agent.² Prendergast (1993) analyzes information acquisition when there is no obvious metric to measure the accuracy of the forecast. The agent receives two

¹ E.g., Ehrbeck and Waldman (1996), Hong et al. (2000), Lamont (2002), Laster et al. (1999).

² When the principal knows the type of agent (i.e., the agent's cost of effort) she sells the project to the agent. Under incomplete information, the principal sells only a fraction of the project to reduce information rents.

signals. The first signal is correlated with the state and the second is correlated with the opinion of the principal. To provide the agent with incentives to exert effort, the wage has to be tied to the principal's opinion. Prendergast shows that the agent's report is biased towards the agent's belief about the opinion of the principal. As a result, the principal cannot infer the signals. If the agent's desire to conform with the principal's opinion is large, the gain from an unbiased report outweighs the loss from lower precision. In this case, the principal does not offer an incentive contract but instead pays a fixed wage to guarantee truth-telling.

In a different context, several authors analyze incentives for information acquisition in procurement settings. Crémer and Khalil (1992), Crémer, Khalil, and Rochet (1998), and Lewis and Sappington (1999) analyze how to deter an agent from acquiring information that generates rents but has no social value. Lewis and Sappington (1997) consider a situation where the principal wants the agent not only to acquire but also to reveal information that is socially valuable. The principal optimally offers a reward schedule with 'super-high-powered' incentives when costs are low and a high degree of cost-sharing when costs are high.

Li (2001) and Szalay (2003) examine information acquisition when players have the same preferences over the implementation of actions and when it is not possible to design monetary incentives. Li assumes that the action space is binary and that players receive a signal about the true state where the precision of the signal is increasing in effort. Li shows that commitment to excessive conservatism in the decision-making of committees can be used to increase the incentives to acquire information. Szalay considers a situation where the action space is continuous and assumes that the probability to learn the state is increasing in effort. In a principal-agent setting, Szalay shows that the principal finds it optimal to let the agent choose the action but excludes actions from the choice set, which are optimal when no information is acquired. In both papers, the commitment to ex-post inefficiencies in the choice of the action is used to mitigate the ex-ante problem to create incentives for the acquisition of information.

Gromb and Martimort (2003) study the implications of optimal incentive contracts for the

organizational design of expertise. There are two possible states: a project is either profitable or not. The agent can draw independent binary signals at a fixed cost per signal. After receiving the signals, the agent recommends to either undertake the project or not. Gromb and Martimort state a 'Principle of Incentives for Expertise': It is optimal to reward the agent if his recommendation is confirmed. We show that this Principle does not necessarily hold when the state and signal space are continuous.

The paper is organized as follows: Section 2 describes the model. Section 3 derives the optimal contract when the principal knows the agent's cost to exert effort. Section 4 discusses the effect of uncertainty about the agent's cost on the optimal contract. Section 5 examines situations where the principal can contract with multiple agents. We analyze the optimal number of agents and show that the equilibrium payoffs and precision do not change when agents acquire information sequentially instead of simultaneously. We also analyze how to provide incentives to acquire information when the agent has to be paid before the state of the world is realized or if the state is not verifiable. Section 6 concludes.

2. The Model

Consider a principal whose expected payoff depends on the precision of her estimate about the realization of some random variable. The principal has to take some action a before the future state of the world is realized. The principal's payoff is decreasing in the distance between the action and the realized state x where $x, a \in R$. Let l denote the loss function that describes the payoff and let $V = -E[l(|a - x|)]$ denote the expected loss with $l' > 0$.

There exists a common prior that the future state x is distributed normally with mean μ and precision $\nu < \infty$, i.e., $x \sim N(\mu, \frac{1}{\nu})$. The agent has access to a technology that generates information about the future state of the world. When the agent exerts effort, he draws a signal s . The signal is distributed normally and the mean of the signal is the future state of the world,

$s \sim N(x, \frac{1}{\tau})$. The precision τ of the signal depends on how much effort the agent exerts. The disutility of the effort that is necessary to generate a signal with precision τ is given by $c(\tau)$ where c is continuously differentiable with $c(0) = 0$, $c'(0) = 0$, $c'(\tau) > 0$ for all $\tau > 0$, and $c'' \geq 0$.

The principal observes neither the signal nor the precision. To improve the precision of her estimate about the future state, the principal can offer a contract to the agent. The agent can costlessly send a message $m \in R$. While signal and precision are not observed by the principal, we assume that the message, the action, and the realized state are verifiable. Let $E[w]$ be the expected wage payment. The expected payoff of the principal is $V - E[w]$. The payoff of the agent depends on the wage w and the disutility of effort. To isolate the problem to create incentives for information acquisition and revelation from the effects of risk-aversion, we assume that the agent is risk-neutral and has zero liability. If the agent accepts the contract, his payoff is $w - c(\tau)$. If the agent rejects the contract, he receives a payoff of zero.

The timing is as follows: The principal offers a contract. The agent accepts or rejects. If he accepts, he exerts effort and draws a signal with the corresponding precision. After observing the signal, the agent sends a message. Given the message, the principal updates her belief about the future state and chooses an action. Finally, the state is realized, the wage is paid, and payoffs are realized.

After the principal receives the message, she updates her beliefs about the distribution of x and chooses an action to minimize the expected loss. Choosing an action is equivalent to computing the appropriate estimator where the estimator minimizes the expected value of the error function l (i.e., of the loss function). The error of the estimator and, therefore, the expected payoff of the principal, depend not only on the precision of the signal but also on how much information the message reveals about the signal. For any non-trivial mapping from signals to messages, the expected payoff of the principal is increasing in the precision of the signal. Similarly, for any precision $\tau > 0$, the expected payoff of the principal increases if the measures of all sets of signals, which induce the same message, decrease. If the contract induces a one-to-one mapping from

signals to messages, the principal learns the signal and her posterior can be described by a normal distribution.

Lemma 1 *When the principal learns the signal and the precision, she chooses the action equal to her posterior mean.*

All proofs are relegated to the appendix.

Given that the message reveals the signal, let $\frac{\partial V}{\partial \tau}$ be the marginal change of the expected loss from a mismatch between action and state with respect to the precision of the signal.

Benchmark: The First-Best Solution

Suppose the principal can observe τ and s or can sell the project to a risk-neutral agent with unlimited liability. In this case, in equilibrium, $a = \frac{\tau^{**}s + \nu\mu}{\tau^{**} + \nu}$, the agent exerts the first-best effort, and the first-best precision τ^{**} is defined by

$$\frac{\partial V}{\partial \tau} = c'(\tau^{**})$$

When the precision of the signal is not observable and the agent is either wealth constrained or risk-averse, implementation of the first best effort is not optimal. In this case, the principal faces the problem to design a contract that provides the agent with incentives to exert effort and to reveal his information.

3. The Optimal Contract

Since signal and effort are not observed by the principal, the wage can only be conditioned on the message, the action, and the realized state. If the wage depends only on the message or only on the state, the agent has no incentive to exert effort. Since $\nu < \infty$, additional information about the future state is valuable for principal. The assumptions on c imply that at $\tau = 0$, the

marginal cost to provide the agent with incentives to exert effort is zero. Therefore, it cannot be optimal to offer a contract where the wage depends only on the message or only on the state.

On the other hand, contracts where the wage depends on the action are feasible, but not optimal. Since the agent's payoff does not directly depend on the action, it is not possible to provide additional incentives to exert effort by making the wage dependent on the action. Additionally, if the wage depends in a non-trivial way on the action, the choice of the action is distorted, because the principal takes into account how the action affects the wage. Therefore, under the optimal contract, the wage is a function of message and state only.

We derive the optimal contract in two steps. Initially, we suppress the question whether signals are revealed and characterize the contract that maximizes the incentive to exert effort for an arbitrary expected wage payment. Using the result that under these contracts the message reveals the signal, we compute the optimal precision that the principal implements.

Given a signal s with precision τ , the agent's conditional expectation is $E[x|s, \tau] = \frac{\tau s + \nu \mu}{\tau + \nu}$. From the perspective of the agent, the future state is distributed normally with $x|s, \tau \sim N\left(\frac{\tau s + \nu \mu}{\tau + \nu}, \frac{1}{\tau + \nu}\right)$. Let $q(x|s, \tau)$ denote the corresponding conditional density with $q(x|s, \tau) = \frac{\sqrt{\tau + \nu}}{\sqrt{2\pi}} \exp -\frac{1}{2} \{(x - E[x])^2(\tau + \nu)\}$. Let $w(m, x)$ denote the agent's wage if he sends message m and the realized state is x . Given w , let $p(m, s, \tau)$ denote the expected utility from wages given message m and signal s observed with precision τ ,

$$p(m, s, \tau) = \int w(m, x) q(x|s, \tau) dx \quad (1)$$

Let $m^*(s, \tau)$ be the message that maximizes $p(m, s, \tau)$. We suspend the proof that m^* is unique until Proposition 1.

Before the agent observes a signal, he chooses effort to maximize his payoff

$$\int \int w(m^*, x) q(x|s, \tau) dx f(s, \tau) ds - c(\tau) \quad (2)$$

where $f(s, \tau)$ denotes the unconditional density of the signal as a function of the effort. Since the conditional distribution of s is normal with $s \sim N\left(x, \frac{1}{\tau}\right)$ and since $x \sim N\left(\mu, \frac{1}{\nu}\right)$, the un-

conditional distribution of s is normal with $s \sim N\left(\mu, \frac{\tau+\nu}{\tau\nu}\right)$ and $f(s, \tau)$ is given by $f(s, \tau) = \frac{\sqrt{\tau\nu}}{\sqrt{2\pi}\sqrt{\tau+\nu}} \exp -\frac{1}{2} \left\{ \frac{\tau\nu(s-\mu)^2}{\tau+\nu} \right\}$.

Let τ^* be the precision that maximizes the agent's payoff. The first order condition that corresponds to the agent's problem to maximize eqn.(2) is given by:

$$\begin{aligned} & \int \left[\int \left(\frac{\partial q}{\partial \tau} + \frac{\partial q}{\partial E[x]} \frac{\partial E[x]}{\partial \tau} \right) w(m^*, x) dx + \frac{\partial p(m, s, \tau^*)}{\partial m} \frac{\partial m}{\partial \tau} \right] f ds \\ & + \int \int w(m^*, x) q dx \frac{\partial f}{\partial \tau} ds = c'(\tau^*) \end{aligned} \quad (3)$$

where w is evaluated at m^* . Since m^* is the message that maximizes the agent's payoff,

$\frac{\partial p(m, s, \tau)}{\partial m} = 0$ at m^* . After substituting for q , f , and $E[x]$, the first-order condition in eqn.(3) can be written as

$$\int \int \left(\frac{1}{2\tau^*} - \frac{1}{2}(x-s)^2 \right) w(m^*, x) q(x|s, \tau^*) dx f(s, \tau^*) ds = c'(\tau^*) \quad (4)$$

where the LHS of eqn.(4) is the marginal increase of the expected wage if the precision of his signal increases.

The principal designs a contract to maximize her payoff

$$V - \int \int w(m^*, x) q(x|s, \tau^*) dx f(s, \tau^*) ds$$

where second term is the expected payment to the agent. The contract specifies a non-negative wage for every message and for every realized state.

If the optimal contract induces the agent to reveal his signal, it follows from the revelation principle that it is sufficient to restrict attention to contracts where $m^*(s) = s$ and the agent announces his signal. If signals are revealed, maximizing the payoff of the principal is equivalent to maximizing the incentive to increase the precision subject to the constraint that wages are non-negative, that the agent reveals the signal, and that the expected wage payment is less or equal to some optimally chosen constant.

The incentive to increase the precision is the marginal effect of precision on the expected wage and is equal to the LHS of eqn.(4). Since $m^*(s) = s$, the expected wage $E[w]$ is equal

to $\int \int w(s, x)q(x|s, \tau^*)dx f(s, \tau^*)ds$. For any $E[w] > 0$, a necessary condition for the LHS of eqn.(4) to be maximal is that $w(s, x) = 0$ for $x \neq s$ and that $w(s, s) > 0$ for some s . Hence, the principal pays a positive wage only if the realized state is equal to the signal. Using the result that $w(s, x) = 0$ if $x \neq s$, eqn.(4) can be written as

$$\frac{1}{2\tau^*} \int \int w(s, s)q(s|s, \tau^*)dx f(s, \tau^*)ds = c'(\tau^*) \quad (5)$$

Eqn.(5) is derived from the maximization of $E[w] - c(\tau)$. The LHS of eqn.(5) is equal to $\frac{\partial E[w]}{\partial \tau}$. Since $E[w] = \int \int w(s, x)q(x|s, \tau)dx f(s, \tau)ds$, the expected wage can be written as a differential equation with $\frac{\partial E[w]}{\partial \tau} = \frac{1}{2\tau}E[w]$. Then

$$E[w] = K\sqrt{\tau} \quad (6)$$

where K is a constant that is chosen by the principal.

Proposition 1 summarizes the equilibrium.³

Proposition 1 *The principal offers the following contract:*

$$w(m, x) = k \exp \left\{ \frac{1}{2}\nu m^2 - \nu m\mu \right\} \quad \text{if } m = x$$

$$w(m, x) = 0 \quad \text{otherwise}$$

(i) *The agent accepts and sends message $m(s) = s$.*

(ii) *The principal believes that $x \sim N\left(\frac{\tau^*m + \nu\mu}{\tau^* + \nu}, \frac{1}{\tau^* + \nu}\right)$ and chooses action $a = \frac{\tau^*m + \nu\mu}{\tau^* + \nu}$.*

(iii) *The equilibrium precision τ^* is defined by $\frac{\partial V}{\partial \tau} = 2c'(\tau^*) + 2\tau^*c''(\tau^*)$.*

(iv) *In equilibrium, $E[w] = 2\tau^*c'(\tau^*)$. The expected equilibrium payoff of the agent is $2\tau^*c'(\tau^*) - c(\tau^*)$ and the expected equilibrium payoff of the principal is $V(\tau^*) - 2\tau^*c'(\tau^*)$.*

³ The assumption that state and signal are normally distributed simplifies the analysis and allows to compute explicit solutions but complicates the interpretation of the results. The probability that the state is equal to the signal is zero. Since the support of the normal distribution is unbounded, the wage is unbounded. Since $w(s, s)q(s|s, \tau)f(s, \tau)$ is constant, for $0 < k < \infty$, the expected wage is positive and finite although the wage is paid with probability zero. The technical reason for this zero-probability result is that at the optimum, the principal rewards the agent if the state lies in the smallest interval that is centered around the signal. Therefore, equilibrium existence requires that the agent is rewarded only if $x = s$, which occurs with probability zero. On the other hand, it is impossible to measure the realization of a continuous variable (e.g., GDP) with infinite accuracy. Proposition 1 should be interpreted as the wage being positive if the state lies in some small interval that is centered around the message.

One advantage of the model is that it allows to compute explicit expressions for the equilibrium precision and the equilibrium payoffs.

Example: Suppose that the loss is quadratic with $l(|x - a|) = (x - a)^2$, that $c(\tau) = \tau$, and that $\nu < \sqrt{\frac{1}{2}}$. (Since $c(\tau) = \tau$ does not satisfy $c'(0) = 0$, the assumption that $\nu < \sqrt{\frac{1}{2}}$ is necessary to guarantee an interior solution.)

At the equilibrium, $\frac{\partial V}{\partial \tau} = 2c'(\tau^*) + 2\tau^*c''(\tau^*)$. From $\frac{\partial V}{\partial \tau} = \frac{1}{(\tau+\nu)^2}$ and $c'(\tau) = 1$ follows that $\tau^* = \sqrt{\frac{1}{2}} - \nu$. From eqn.(6) and the equilibrium condition that $\frac{\partial E[w]}{\partial \tau} = c'(\tau^*)$ at τ^* follows that the expected wage is $E[w] = \sqrt{2} - 2\nu$. The expected payoff of the principal is $-2\sqrt{2} + 2\nu$, and the expected payoff of the agent is $\sqrt{\frac{1}{2}} - \nu$.

Gromb and Martimort (2003) consider an agent who can pay a fixed cost to draw a binary signal. There are two states: a project is either profitable or not. The agent recommends to either undertake the project or not. Gromb and Martimort state a 'Principle of Incentives for Expertise': It is optimal to reward the agent if his recommendation is confirmed (either by the facts or by the recommendation of another expert).

Similar to the setting of Gromb and Martimort, in the model above, the agent is indifferent towards which action is chosen. The contract in Proposition 1 is equivalent to a contract where the agent recommends to choose the action $\frac{\tau s + \nu \mu}{\tau + \nu}$ and the principal uses the recommendation to compute the agent's signal and pays the same wages as in Proposition 1. Contrary to the binary setting of Gromb and Martimort, their 'Principle of Incentives for Expertise' does not hold in the continuous environment that is considered here. Contrary to common intuition, this paper shows that it is in general not optimal to reward the agent when his recommendation is confirmed. Specifically, under the optimal contract, the agent is rewarded (at least sometimes) when his recommendation is wrong but not rewarded when his recommendation is confirmed.

The principal faces two problems: To design a contract and to choose the action that maximizes her expected payoff. The objective in the first problem is to maximize the agent's incentive to

exert effort and to ensure that the agent reveals his information. The objective in the second problem is to maximize the expected payoff after the agent has revealed his private information. To provide the agent with incentives to exert effort, the wage has to depend on the realized state and the message (i.e., the recommendation). However, there is no reason why the wage should be tied to the question how well the recommendation matches the state.

In the model above, the recommendation of an action is equivalent to the announcement of the posterior mean and the recommendation is confirmed if it is equal to the realized state.⁴

To see why it is not optimal to reward the agent if his recommendation is confirmed, fix an arbitrary signal s . Given s , from the perspective of the agent, an increase of the precision implies that the likelihood that some states are realized increases ($\frac{dq(x|s,\tau)}{d\tau} > 0$) and decreases for others ($\frac{dq(x|s,\tau)}{d\tau} < 0$). The principal chooses when to reward the agent. Hence, the principal maximizes the RHS of eqn.(3) with respect to x . It is optimal to reward the agent when the relative effect of an increase of the precision on the likelihood that a certain state is realized, is maximal. Therefore, if the agent observes (and reveals) s , it is optimal to reward him when the realized state is equal to the maximizer of $\frac{dq(x|s,\tau)}{d\tau} \frac{1}{q(x|s,\tau)}$. Note that eqn.(3) makes no use of the assumption that state and signal are normally distributed. In the special case where state and signal are normally distributed, $x = s$ is the maximizer of $\frac{dq(x|s,\tau)}{d\tau} \frac{1}{q(x|s,\tau)}$ and it is optimal to reward the agent when the state is equal to the signal.

In the following paragraphs, we summarize a few results which follow from Proposition 1. Not only are signals revealed under the optimal contract, but the equilibrium precision and expected wage payment are the same as when the signal is verifiable.

Result 1 *The principal implements the same effort level regardless whether the signal is verifiable or not.*

⁴ More general, the optimal action and, therefore, the recommendation, are functions of the posterior mean and a recommendation is confirmed if the realized state is equal to the agent's posterior mean.

The optimal contract in Proposition 1 is derived for the case where the signal is not observable. Result 1 shows that the analysis applies not only to the case where information is soft (i.e., where the signal is not observed or is not verifiable) but also to the case where information is hard and the agent's signal can be verified. In both cases, the contract described in Proposition 1 is optimal.

The Precision of the Prior

The wage is convex in the message with a minimum at $m = \mu$. This is the consequence of the result that it is optimal to reward the agent only if the realized state is equal to the signal. Since $q(s|s, \tau)$ decreases in $|s - \mu|$, the larger the distance between signal and prior mean, the smaller is the likelihood that the state is equal to the signal. To ensure that the agent reveals the signal, the wage has to be convex. The difference between posterior mean and prior mean depends on signal and precision. If the prior is very precise, the agent expects (before he draws the signal) that his posterior mean is fairly close to the prior mean. Therefore, the probability that the state is equal to the signal decreases faster in the distance between signal and prior mean, if the prior is more precise. This is the reason that the convexity of the wage increases when the precision of the prior increases.

Result 2 *The convexity of the wage is increasing in the precision of the prior: $\frac{w''}{w}$ increases (decreases) in ν for $m > \mu$ ($m < \mu$).*

When the precision of the prior is high, the agent has a fairly good idea about the future state even if he exerts no effort. Intuitively, one would expect that this makes it more costly to induce the agent to exert effort when the precision of the prior is high.⁵ It turns out that this intuition is wrong. Under the optimal contract, the incentive to exert effort depends only on the expected wage $E[w]$. If ν increases, the optimal contract changes as the wage becomes more convex, while the equilibrium effort of the agent does not change as long as $E[w]$ does not change.

⁵ For the simple (suboptimal) contract where the agent is rewarded if his recommendation is correct, it is straightforward to show that it is more costly to induce effort for higher values of ν . In this case, the marginal incentive to exert effort is $\frac{1}{2(\tau+\nu)}E[w]$ which is decreasing in ν .

Result 3 *The expected wage that is necessary to implement a given effort does not depend on the precision of the prior.*

Of course, if the precision of the prior increases, the marginal value of the precision of the agent's signal decreases. Therefore, the principal implements a lower effort level if the precision of the prior is high.

Compared to a signal that is close to the prior mean, extreme signals contain more information. More extreme signals trigger a larger change in the action that the principal chooses compared to the action that would have been chosen if no signal were observed. In this sense, extreme signals are more valuable for the principal. On the other hand, the observation of an extreme signal is good news for the agent since the expected wage payment is increasing in the distance between signal and prior mean. However, the fact that the expected wage for extreme signals is higher has nothing to do with the fact that extreme signals are more valuable for the principal. The only reason for the higher wage is to ensure that the agent reveals the signal.

4. Incomplete Information

Proposition 1 shows that it is optimal to reward the agent if the state is equal to the signal. This seems to contradict the observation that many experts are rewarded if their recommendation is confirmed by the facts. We show in this section that incomplete information about the agent's cost to exert effort offers an explanation for this observation.⁶

In the model above, the optimal action coincides with the posterior mean. This coincidence is a result of the assumptions on the loss function but does not restrict the generality of model. The

⁶ There are other explanations. If the agent is risk-averse, it is optimal to reward the agent if the state lies in an interval that is centered around a value between $\frac{\tau s + \nu \mu}{\tau + \nu}$ and s (and might include $\frac{\tau s + \nu \mu}{\tau + \nu}$). Since $q(s|s, \tau)$ is decreasing in $|s - \mu|$, it gets increasingly costly to reward the agent when the interval is centered around s .

recommendation of an optimal action is equivalent to the revelation of the agent's posterior mean. A recommendation of the agent is confirmed by the facts if the realized state coincides with the recommended action.

The optimal contract in Proposition 1 is derived under the assumption that the principal knows the agent's cost to exert effort. The contract in Proposition 1 is equivalent to a contract where $w(\tilde{s}, x) = k \exp \left\{ \frac{1}{2} \nu \tilde{s}^2 - \nu \tilde{s} \mu \right\}$ if $\tilde{s} = x$ and $w(\tilde{s}, x) = 0$ otherwise where $\tilde{s} = \frac{m(\tau + \nu) - \nu \mu}{\tau}$. The agent 'recommends an action' by sending a message that is equal to his posterior: $m(s) = \frac{\tau s + \nu \mu}{\tau + \nu}$. The principal follows the recommendation and implements $a = m$.

When the principal does not know the agent's cost to exert effort, the revelation of the signal is no longer equivalent to the revelation of the posterior. Agents with different costs, exert different effort levels and receive signals with different precision. In this case, the principal does not know the precision of the signal which underlies the agent's report. If the principal learns the signal but does not know the precision of the signal, she cannot compute the efficient action.

On other hand, the principal can design a contract such that the agent recommends the efficient action regardless of the precision of his signal. The optimal contract that ensures that the agent always recommends the efficient action can be written as: $w(m, x) = k$ if $m = x$ and $w(m, x) = 0$ if $m \neq x$. Under this contract, the agent sends the message $m(s) = \frac{\tau s + \nu \mu}{\tau + \nu}$. The agent is rewarded if and only if his recommendation is confirmed by the facts and the principal follows the recommendation. To avoid confusion, from now on, the term 'contract where agent recommends an action' refers to a contract where the agent recommends the efficient action regardless of the precision of his signal.

The incentive to exert effort is maximal if the agent is rewarded only if $x = s$. But if the principal does not know the cost to exert effort and the agent recommends an action, the principal does not learn the signal and it is not possible to reward the agent only if $x = s$.

Under some conditions, this problem has a simple solution: The principal asks the agent to make two reports. The first report reveals the signal, the second reveals the precision of the

signal, and wages are the same as in Proposition 1. Since the wage is independent of the reported precision, the agent has no incentive to misreport the precision and there is no efficiency loss through incomplete information.

The separation of types can be costly and is not always possible. Consider the case where the cost to exert effort is either high or low. The simple solution where the agent sends two reports does not work if the benefit for the high cost type if he is seen as the low cost type is larger than the loss for the low type if he is seen as the high type.⁷ In this case, the principal faces a trade-off: A contract where the agent is rewarded if the state is equal to the signal maximizes the incentive to increase the precision of the signal. But for a given precision, it is optimal if the agent recommends an action because only this guarantees that the principal learns all relevant information.

For a given contract, let τ_l and τ_h be the lowest and the highest precision that different types of the agent choose. If the contract asks the agent to recommend an action, a type i agent recommends $E[x|s, \tau_i]$, the principal follows the recommendation and, therefore, implements the optimal action. If the principal offers a contract under which she learns the signal but not the precision, the action that the principal chooses lies between $E[x|s, \tau_l]$ and $E[x|s, \tau_h]$.⁸ Except for possibly one type j , for who $E[x|s, \tau_j]$ equals the action, the principal does not implement the optimal action. Contrary to the posterior mean $E[x|s, \tau_i]$, the signal s does not contain all information that is relevant for the principal. Therefore, for a given precision, the payoff of the principal is maximal when the agent is asked to recommend an action (Information Effect).

On the other hand, if the agent only recommends an action (i.e., reports his posterior mean), the principal cannot infer the signal and it is not possible to reward the agent when the state is equal to the signal. Instead, the agent is rewarded if the recommendation is confirmed, i.e., if the

⁷ There exists an extensive literature that discusses career concerns and reputation building as reasons for the benefit or loss from being perceived as a certain type (e.g., Chevalier and Ellison (1999) and Hong et al. (2000)).

⁸ The density function f that describes the unconditional distribution of s is a function of τ . When the agent reports a signal, the principal updates her beliefs over the type of the agent. The action that the principal chooses depends on the loss function and on the updated belief over the type of the agent.

state is equal to $E[x|s, \tau_i]$. In this case, at the equilibrium precision τ^* holds $\frac{1}{2(\tau^* + \nu)} E[w] = c'(\tau^*)$. Comparison with eqn.(5) reveals that the equilibrium precision is lower when the agent is rewarded if his recommendation is correct (Incentive Effect).

Suppose that the principal does not know the agent's cost to exert effort and that it is not possible to separate the types. Under incomplete information, neither the contract in Proposition 1 where the agent reveals the signal regardless of his precision nor the contract where the agent reveals the posterior regardless of his precision is optimal. Instead, under the optimal contract, the agent is rewarded if the realized state is equal to some value which lies between the agent's posterior and the signal.

The first step to find the optimal contract is to determine for which realized state the agent gets rewarded. Wlog. let the message be such that the agent gets rewarded if and only if the message is equal to the realized state. If $m = E[x|s, \tau]$, the principal chooses $a = m$. If the message is marginally larger or smaller than $E[x|s, \tau]$, the principal does not learn which action is efficient. However, since $\frac{\partial V}{\partial a} = 0$ at $a = E[x|s, \tau]$, the principal incurs only a second-order loss if the message differs marginally from $E[x|s, \tau]$.⁹

On the other hand, from eqn.(4) follows that at $m = E[x|s, \tau]$ we have $\frac{\partial(\frac{\partial E[w]}{\partial \tau})}{\partial m} > 0$ for $m > \mu$ (and $\frac{\partial(\frac{\partial E[w]}{\partial \tau})}{\partial m} < 0$ for $m < \mu$). Recall that $\frac{\partial E[w]}{\partial \tau}$ is the incentive to exert effort. Hence, at $m = E[x|s, \tau]$, changing the contract such that the message is marginally closer to s results in a first-order increase in the incentive to exert effort.

Similar, at $m = s$, changing the contract such that m is marginally closer to $E[x|s, \tau]$ results in a second-order decrease in the incentive to exert effort but in a first-order gain from a more efficient action.

⁹ If the agent is rewarded when the message is equal to the state, the agent sends the message where $\frac{\partial q(m|s, \tau)}{\partial m} \frac{1}{q(m|s, \tau)} = -\frac{\partial w(m, m)}{\partial m} \frac{1}{w(m, m)}$. Hence, agent i sends the message $m_i^*(s) = \frac{\partial w(m, m)}{\partial m} \frac{1}{w(m, m)} \frac{1}{\tau_i + \nu} + E[x|s, \tau_i]$. From the perspective of the principal, the message is equal to the posterior plus noise (because the principal does not know the agent's type, i.e., does not know τ_i). The smaller $\frac{\partial w(m, m)}{\partial m} \frac{1}{w(m, m)}$, the smaller is the noise term. Therefore, the smaller $\frac{\partial w(m, m)}{\partial m} \frac{1}{w(m, m)}$, the smaller is the expected difference between the action that the principal chooses and the efficient action.

Therefore, under the optimal contract, the agent sends a message that lies between $E[x|s, \tau_i]$ and s and the agent is rewarded if the realized state is equal to the message.

Suppose that the agent's cost to exert effort is either high or low. If the difference between high and low costs is large, the difference between equilibrium precisions τ_l^* and τ_h^* is large. Consider a contract as in Proposition 1 where the principal learns only the signal. If the difference between τ_l^* and τ_h^* is large, the expected difference between $\frac{\tau_l^* s + \nu \mu}{\tau_l^* + \nu}$ and $\frac{\tau_h^* s + \nu \mu}{\tau_h^* + \nu}$ is large and the information loss when the principal learns the signal instead of the posterior mean $\frac{\tau_i^* s + \nu \mu}{\tau_i^* + \nu}$ is large. If the difference between τ_l^* and τ_h^* is large enough, the Information Effect dominates the Incentive Effect. In this case, under the optimal contract, the agent is rewarded if the realized state is close to the posterior mean $E[x|s, \tau]$. If the difference between the equilibrium precisions is small, the Incentive Effect dominates. In this case, the agent is rewarded if the realized state is close to his signal.

The probability distribution over the types of the agent also affects the design of the optimal contract. If there is little uncertainty over the cost of the agent, the Incentive Effect dominates the Information Effect. Suppose that the probability that the agent has high costs is close to one. If the principal learns only the signal, she chooses an action close to $\frac{\tau_h^* s + \nu \mu}{\tau_h^* + \nu}$. With high probability, the principal's action is close to the efficient action and only with small probability, there is a large difference between the principal's action and the efficient action. Therefore, if the probability that the agent has high costs is either large enough or small enough, the Incentive Effect dominates the Information Effect. In this case, the agent is rewarded if the realized state is close to his signal.

5. Multiple Experts, the Timing of Expertise, and Non-Verifiable States

Decision makers routinely seek the advice of more than one expert. Especially in the public policy area, it is common that several experts comment on projects. Duplication of expertise

is usually justified with benefits from including different points of view in the decision making process. Several authors examine situations where a principal can contract with several experts to acquire information. Pesendorfer and Wolinsky (2003) analyze the optimal strategy of a principal who has to rely on experts to identify the correct type of service that is needed. The principal samples experts until he receives two matching recommendations. Krishna and Morgan (2001) find that a principal, who faces biased experts, benefits from consulting two experts if the biases are opposed. In a binary model, Gromb and Martimort (2003) show that agency costs exhibit diseconomies of scale and that contracting with multiple experts reduces costs. Contrary to Gromb and Martimort, we show that in the continuous model that is discussed above, agency costs do not exhibit diseconomies of scale (at least if the marginal cost of precision is constant as assumed by Gromb and Martimort).

With respect to the organization of expertise, we are interested in two questions: What is the optimal number of agents with whom the principal contracts and what is the optimal timing of expertise. We consider two different timings. Under 'simultaneous expertise', agents acquire information simultaneously and agents cannot communicate with each other before they send a message to the principal. The assumption that communication is impossible rules out collusion. Under 'sequential expertise', agents acquire information sequentially and agents observe earlier messages and contracts.

Suppose there is a pool of $N < \infty$ identical agents. Let i denote an arbitrary agent with whom the principal contracts and let τ_i^* be the equilibrium precision of agent i 's signal and let $\tau^* = \sum \tau_i^*$. Let $\frac{\partial V}{\partial \tau}$ be the derivative with respect to τ^* . Let n^* be the equilibrium number of agents who contract with the principal.¹⁰

Proposition 2 (i) *If signals are conditionally independent, $n^* = 1$ if $c'' < 0$, $n^* = N$ if $c'' > 0$, and $n^* = 1, 2, \dots, N$ if $c'' = 0$. The equilibrium precision τ^* is defined by $\frac{\partial V}{\partial \tau} =$*

¹⁰ In the model above, we assumed that $c'' \geq 0$ to ensure that there exists a unique precision that maximizes the principal's payoff. If $c'' < 0$, Proposition 2 refers to the case when τ^* is unique.

$2c' \left(\frac{\tau^*}{n^*} \right) + 2\tau_i^* c'' \left(\frac{\tau^*}{n^*} \right)$. If $c'' \neq 0$, $\tau_i^* = \frac{\tau^*}{n^*}$. If $c'' = 0$, $\sum \tau_i^* = \tau^*$ and $\tau_i^* \geq 0$ for all i . In equilibrium, $E[w_i] = 2\tau_i^* c'(\tau_i^*)$. The expected equilibrium payoff of the agent is $2\tau_i^* c'(\tau_i^*) - c(\tau_i^*)$ and the expected equilibrium payoff of the principal is $V(\tau^*) - 2\tau^* c'(\tau_i^*)$.

(ii) For any vector $\tau_i^*, \tau_j^*, \dots$, payoffs and expected wages are the same under simultaneous and sequential expertise.

Proposition 2 shows that the costs of expertise are independent of the timing of expertise. Besides a possible preference to receive information early, there exist other factors that influence the timing of expertise. When the loss of the principal is a function of the action¹¹, sequential expertise can be optimal even when $c'' < 0$. Consider a firm which has to decide whether or not to invest in a project. The project is profitable only if the state is above some threshold. In this case, it can be optimal that the agent initially exerts only little effort and acquires only general information. Only if the first assessment is not too negative, it is worthwhile to ask an agent to acquire more detailed information.¹²

A different reason for sequential expertise or contracting with multiple agents is that it is not always possible to write contracts where transfers are contingent on the realized state. If the state is realized at some distant point in the future, it might be necessary to pay the agent before the state is realized. Additionally, in some situations, the realized state is not verifiable while messages are. Consider an expert who is asked to estimate how many people will be killed by a disease. While the actual number of deaths might not be verifiable, an estimate is a number that the expert announces and there is no reason why it should not be possible to contract on this

¹¹ In the model above, we assume that the loss function is symmetric and that it depends only on the distance between action and state. This assumption simplifies the proof of uniqueness because it guarantees that $\frac{\partial^2 V}{\partial z^2} < 0$. Except to ensure uniqueness, there are no other reasons for this assumption.

¹² Whenever $\tau^* > 0$, the agent earns some rent. If there is only one agent, the agent takes into account that only a message that induces an optimistic belief guarantees that the principal wants to acquire additional information which generates additional rents for the agent. To ensure that the agent reveals the signal, the wage has to be adjusted. Specifically, for low signals (where the principal does not ask the agent to acquire additional information) the wage is higher compared to the contract in Proposition 1. The agent's concern about future rents makes initial information acquisition more costly and, therefore, leads to inefficiencies. If there are many agents, the source of the inefficiencies disappears. After every message, the principal decides whether she should ask another agent to acquire additional information.

announcement.

When the state is either not verifiable or is not observed before wages are paid, the contract cannot specify wages as a function of the realized state. In order to provide the agent with incentives to acquire information, the principal has to contract with a second agent.¹³ Since signals are correlated with the true state, information about the signal of a second agent can be used to provide the first agent with incentives to exert effort. The cost to provide an agent with incentives to exert effort depends on the precision of the signal of the other agent. The less precise the signal of the other agent, the more expensive it is to provide an agent with incentives to exert effort. In the limit as the precision of the signal of the other agent goes to infinity, the principal is in the same situation as in Proposition 1 where state is observable.

We do not consider collusion. Collusion is a less severe problem than one would expect. Under the equilibrium contract in Proposition 3, at most one agent receives a positive transfer (except if both send the message $m = \mu$). If side payments between agents are not possible, the principal can prevent collusion by setting the wage for $m = \mu$ to zero.

Suppose that the principal can contract with two agents and let i, j denote the agents. Additionally, we assume that $c''' \geq 0$.¹⁴

Proposition 3 (i) *If transfers cannot be conditioned on the state, the principal contracts with both agents. Each agent is offered the following contract:*

$$w_i(m_i, m_j) = k \exp \left\{ \left(1 + \frac{\nu}{\tau_j^*} \right) \left(\frac{1}{2} \nu m_i^2 - \nu m_i \mu \right) \right\} \quad \text{if } m_j = m_i + \frac{\nu(m_i - \mu)}{\tau_j^*}$$

$$w_i(m_i, m_j) = 0 \quad \text{if } m_j \neq m_i + \frac{\nu(m_i - \mu)}{\tau_j^*}$$

The agents accept and send the message $m(s) = s$.

(ii) *The principal believes that $x \sim N \left(\frac{\tau_i^* m_i + \tau_j^* m_j + \nu \mu}{\tau_i^* + \tau_j^* + \nu}, \frac{1}{\tau_i^* + \tau_j^* + \nu} \right)$ and chooses action $a =$*

¹³ In our model, the principal has no private information. Prendergast (1993) analyzes a situation where the state is not observed and where the principal receives a signal about the true state that is her private information.

¹⁴ This assumption is only necessary to derive an explicit expression for the equilibrium precision. Proposition 3(i)-(ii) do not change if $c''' < 0$. However, if $c''' < 0$, it is possible that the principal offers different contracts to agent i and j and, therefore, that the equilibrium precisions τ_i^* and τ_j^* differ.

$$\frac{\tau_i^* m_i + \tau_j^* m_j + \nu \mu}{\tau_i^* + \tau_j^* + \nu}.$$

- (iii) In equilibrium, $\tau_i^* = \tau_j^*$. The equilibrium precision τ^* is defined by $\frac{\partial V}{\partial \tau} = \left(4 - \frac{2\nu^2}{(\tau_i^* + \nu)^2}\right) c' \left(\frac{\tau^*}{2}\right) + \left(2 + \frac{2\tau_i^{*2}}{(\tau_i^* + \nu)}\right) c'' \left(\frac{\tau^*}{2}\right)$ with $\tau^* = \tau_i^* + \tau_j^*$.
- (iv) In equilibrium, $E[w_i] = E[w_j] = \frac{2\tau^{*2} + 2\tau^* \nu}{\tau^* + 2\nu} c' \left(\frac{\tau^*}{2}\right)$. The expected equilibrium payoff of an agent is $\frac{2\tau^{*2} + 2\tau^* \nu}{\tau^* + 2\nu} c' \left(\frac{\tau^*}{2}\right) - c \left(\frac{\tau^*}{2}\right)$ and the expected equilibrium payoff of the principal is $V(\tau^*) - \frac{4\tau^{*2} + 4\tau^* \nu}{\tau^* + 2\nu} c' \left(\frac{\tau^*}{2}\right)$.

Contrary to common intuition, agents are not rewarded if they confirm the recommendation of each other. If both agents receive the same signal, they send the same message. Except if $s = \mu$, agents receive a wage of zero. An agent is rewarded only if the other agent received a signal which is more extreme. The reason for this result is similar to the reasoning behind Proposition 1. Under the optimal contract, the agent reveals his signal. For each signal s_i , the principal decides for which realizations of the other signal s_j agent i is rewarded. It is optimal to reward the agent when the relative effect of an increase of the precision on the likelihood that the other agent receives a certain signal, is maximal. If agent i observes s_i , it is optimal to reward him when s_j is the maximizer of $\frac{dq(s_j|s_i, \tau_i, \tau_j)}{d\tau_i} \frac{1}{q(s_j|s_i, \tau_i, \tau_j)}$.

For a given precision, agents earn a higher rent when the state is not observable or not verifiable. In this case, wages depend on the signal of the other agent. Since agents are protected by zero liability, the expected wage that is necessary to implement a certain precision, increases in the variance of the signal of the other agent.

6. Conclusion

It is common that decision makers ask experts to acquire additional information and to recommend an action. When information acquisition is delegated, the problem arises how to align the

incentives for information acquisition and for truthful information revelation. This paper analyzes the optimal design of contracts that provide incentives to acquire and reveal information.

We show that it is in general not optimal to reward the expert when his recommendation is confirmed by the facts. Under the optimal contract, the expert is (at least sometimes) rewarded when his recommendation turns out to be wrong but is not rewarded when the recommendation was correct. When the signal of the expert and the state are normally distributed, the expert is rewarded when the state is equal to his signal. The reason for this surprising result is that it is optimal to reward the expert when the relative effect of an increase of the precision on the likelihood that a certain state is realized, is maximal.

While it runs against common intuition to reward an expert when his recommendation was incorrect but not when it was correct, it also seems to contradict the observation that many experts are rewarded when their recommendation is confirmed. We show that incomplete information about the expert's cost to exert effort can explain why many experts are rewarded if their recommendations are confirmed. When the principal knows the expert's cost to exert effort, she also knows the equilibrium precision of the expert's signal. In this case, the revelation of the signal is equivalent to the recommendation of an action. When there is incomplete information about the expert's cost to exert effort, a recommendation contains more information than the revelation of the private information. On the other hand, if the expert recommends an action and is rewarded if the recommendation is confirmed, he has less incentives to increase the precision of his signal. When there is a lot of uncertainty about the cost to exert effort, the first effect (more precise information) dominates the second (less incentives).

One advantage of the model is that it generates explicit solutions. This allows to analyze the optimal organization of expertise. We show that the cost of expertise is exclusively determined by the cost to acquire information. The optimal number of agents depends on the cost of information acquisition. Only if the marginal cost of precision increases, the costs of expertise exhibit diseconomies of scale and it is optimal to contract with as many agents as possible. In many situations,

it is possible to gather information either sequentially or simultaneously. We show that regardless of the cost function, the timing of expertise does not affect the payoffs and, therefore, does not affect the equilibrium amount of information that is acquired. Finally, we analyze the optimal contract when the realized state is not observable or not verifiable. In this case, the principal contracts with two agents and the wage depends on the reports of both agents. Similar to the case where the state is observable, an agent is not rewarded if his report is confirmed by the other agent. Instead, an agent is rewarded if the other agent reports a signal which is more extreme.

Appendix

Proof of Lemma 1

If the principal learns the signal s and precision τ , her posterior belief is that x is distributed normally with $x \sim N\left(\frac{\tau s + \nu \mu}{\tau + \nu}, \frac{1}{\tau + \nu}\right)$ and her posterior mean is $E[x] = \frac{\tau s + \nu \mu}{\tau + \nu}$. Let q_p denote the associated density function. The principal solves

$$\min_a \int l(|a - x|) q_p(x|s, \tau) dx$$

Write the first-order-condition as

$$\int_{-\infty}^a l'(|a - x|) q_p(x|s, \tau) dx - \int_a^{\infty} l'(|a - x|) q_p(x|s, \tau) dx = 0 \quad (\text{A1})$$

Since q_p is normal, $q_p(E[x] - \Delta|s, \tau) = q_p(E[x] + \Delta|s, \tau)$. Hence, eqn.(A1) holds at $a = E[x]$. To see that $a = E[x]$ is the absolute minimum of the expected loss, let $\Delta = a - E[x]$. Rewrite the first-order-condition in eqn.(A1) as

$$\int_{-\infty}^{E[x] + \Delta} l'(|E[x] + \Delta - x|) q_p(x|s, \tau) dx - \int_{E[x] + \Delta}^{\infty} l'(|E[x] + \Delta - x|) q_p(x|s, \tau) dx = 0 \quad (\text{A2})$$

For $x_1 \leq E[x] + \Delta$, define $x_2 = 2E[x] + 2\Delta - x_1$. Then $l(|E[x] + \Delta - x_1|) = l(|E[x] + \Delta - x_2|)$.

Rewrite eqn.(A2) as

$$\int_{-\infty}^{E[x]+\Delta} l'(|E[x] + \Delta - x_1|) \cdot (q_p(x_1|s, \tau) - q_p(x_2|s, \tau)) dx_1 \quad (\text{A3})$$

For $\Delta < 0$, $|E[x] - x_1| > |E[x] - x_2|$ and since q_p is normal, $q_p(x_1|s, \tau) < q_p(x_2|s, \tau)$ and eqn.(A3) is negative. Similarly, for $\Delta > 0$, $|E[x] - x_1| < |E[x] - x_2|$ and, therefore, $q_p(x_1|s, \tau) > q_p(x_2|s, \tau)$ and eqn.(A3) is positive. Hence, $a = E[x]$ minimizes the principal's loss.

Before the proof of Proposition 1, we introduce a technical Lemma.

Lemma A1 $\frac{\partial^2 V}{\partial \tau^2} < 0$

Recall that $\frac{\partial V}{\partial \tau}$ is defined as the marginal change in the principal's payoff given that the message reveals the signal.

Proof: The principal's payoff (excluding wages) is given by

$$V = - \int \int l(|x - a|) q(x|s, \tau) dx f(s, \tau) ds$$

with a chosen optimally. Then

$$\frac{dV}{d\tau} = - \frac{\partial V}{\partial a} \frac{\partial a}{\partial \tau} - \int \int l(|x - a|) \frac{dq(x|s, \tau)}{d\tau} dx f(s, \tau) ds + \int \int l(|x - a|) q(x|s, \tau) dx \frac{\partial f}{\partial \tau} ds \quad (\text{A4})$$

At the optimal action, $\frac{\partial V}{\partial a} = 0$. Note that V is a function of $|x - a|$ but not of x . Therefore, V is independent of s . Hence, $\int l(|x - a|) q(x|s, \tau) dx$ is independent of s and, therefore, the last term in eqn.(A4) is zero. Recall that $\frac{dq}{d\tau} = \frac{\partial q}{\partial \tau} + \frac{\partial q}{\partial E[x]} \frac{\partial E[x]}{\partial \tau}$ with $\frac{\partial q}{\partial E[x]} \frac{\partial E[x]}{\partial \tau} = (x - E[x]) \frac{\nu(s-\mu)}{\tau+\nu} q(x|s, \tau)$.

From Lemma 1, $a = E[x]$. Since l is symmetric around a with $a = E[x]$ and q is symmetric around $E[x]$, we have $\int l(|x - a|) \frac{\partial q}{\partial E[x]} \frac{\partial E[x]}{\partial \tau} dx = 0$. Then

$$\frac{dV}{d\tau} = - \int \int l(|x - a|) \left(\frac{1}{2(\tau + \nu)} - \frac{1}{2}(x - E[x])^2 \right) q(x|s, \tau) dx f(s, \tau) ds$$

and

$$\frac{d^2V}{d^2\tau} = - \int \int l(|x - a|) \left(-\frac{1}{2(\tau + \nu)^2} + \left(\frac{1}{2(\tau + \nu)} - \frac{1}{2}(x - E[x])^2 \right)^2 \right) q(x|s, \tau) dx f(s, \tau) ds$$

Write $\frac{d^2V}{d^2\tau}$ as $\frac{d^2V}{d^2\tau} = - \int \psi(s) f(s, \tau) ds$. By same reasoning as above, since $a = E[x]$, the interior integral $\psi(s)$ depends on x and $E[x]$ but not on s . Hence

$$\frac{d^2V}{d^2\tau} = - \int l(|x - a|) \left(-\frac{1}{2(\tau + \nu)^2} + \left(\frac{1}{2(\tau + \nu)} - \frac{1}{2}(x - E[x])^2 \right)^2 \right) q(x|s, \tau) dx$$

Note that $\int \left(\frac{1}{2(\tau + \nu)} - \frac{1}{2}(x - E[x])^2 \right)^2 q(x|s, \tau) dx$ can be written as

$$\frac{1}{4} \int \left(\frac{1}{(\tau + \nu)^2} - 2\frac{1}{(\tau + \nu)}(x - E[x])^2 + (x - E[x])^4 \right) q(x|s, \tau) dx.$$

Of course, $\int (x - E[x])^2 q(x|s, \tau) dx$ is the variance of x given s and is equal to $\frac{1}{(\tau + \nu)}$. Similar, $\int (x - E[x])^4 q(x|s, \tau) dx$ is the fourth centered moment of a normal distribution and is equal to $3\sigma^4$ with $\sigma^2 = \frac{1}{\tau + \nu}$. Therefore,

$\int \left(\frac{1}{2(\tau + \nu)} - \frac{1}{2}(x - E[x])^2 \right)^2 q(x|s, \tau) dx = \frac{1}{2}\sigma^4 = \frac{1}{2(\tau + \nu)^2}$ which is, of course, the same

as $\int \frac{1}{2(\tau + \nu)^2} q(x|s, \tau) dx$. Hence, if l is constant, $\frac{d^2V}{d^2\tau} = 0$. To show that $\frac{d^2V}{d^2\tau} < 0$, define

$I_\tau = \left[E[x] - \frac{1}{\tau + \nu}, E[x] + \frac{1}{\tau + \nu} \right]$. For $x \in I_\tau$ we have $-\frac{1}{2(\tau + \nu)^2} + \left(\frac{1}{2(\tau + \nu)} - \frac{1}{2}(x - E[x])^2 \right)^2 \leq 0$

and for $x \notin I_\tau$ we have $-\frac{1}{2(\tau + \nu)^2} + \left(\frac{1}{2(\tau + \nu)} - \frac{1}{2}(x - E[x])^2 \right)^2 > 0$. Note that principal chooses

$a = E[x]$. Since $l' > 0$, for any $x \in I_\tau$, $l(|x - a|)$ is smaller than $l(|x - a|)$ for $x \notin I_\tau$. Hence,

$$\frac{d^2V}{d^2\tau} < 0.$$

Proof of Proposition 1

(i) The contract in Proposition 1 was derived under the restriction that all signals are revealed.

For a given precision, V is maximal if all signals are revealed. For a given expected wage, the contract in Proposition 1 maximizes the incentive to exert effort (i.e., the precision). Since all signals are revealed, the contract in Proposition 1 is optimal.

We show first that the agent announces the signal, i.e., that $m^*(s) = s$. Since $w(m, x) = 0$ if $m \neq x$, eqn.(1) can be rewritten as

$$\max_m q(m|s, \tau) \cdot w(m, m) \quad (\text{A5})$$

Substituting for q and w shows that the first-order condition that corresponds to eqn.(A5) holds with equality at $m = s$. The second-order condition is satisfied. Hence the agent announces the signal. Since the agent is risk-neutral, the wage is unique up to randomization, i.e., given signal s and optimal message $m^*(s) = s$, the expected wage $E[w]$ is unique.

It remains to show that the agent accepts the contract. Eqn.(5) can be rewritten as $E[w] = 2\tau^* c'(\tau^*)$. From $c(0) = 0$ and $c'' \geq 0$ follows that $c(\tau^*) \leq \tau^* c'(\tau^*)$. Hence, at the equilibrium, $E[w] - c(\tau^*) > 0$ and the agent accepts the contract.

(ii) Since the agent announces the signal, the principal's posterior is normal. From Lemma 1 follows $a = \frac{\tau s + \nu \mu}{\tau + \nu}$.

(iii) There exists a unique τ^* that maximizes the expected payoff of the agent. This follows from $\frac{\partial E[w]}{\partial \tau} = c'(\tau^*)$ and $E[w] = K\sqrt{\tau}$ (eqn.(6)).

There exists a unique τ^* that maximizes the payoff of the principal. The principal maximizes $V - E[w]$. We can write $E[w]$ implicitly as a function of the precision τ^* that the principal wants to implement with $E[w] = 2\tau^* c'(\tau^*)$. Hence, the principal wants to implement the precision τ^* for which

$$\frac{\partial V}{\partial \tau} = 2c'(\tau^*) + 2\tau^* c''(\tau^*) \quad (\text{A6})$$

and, therefore, chooses $E[w]$ (i.e. chooses k) such that $E[w] = 2\tau^* c'(\tau^*)$. Uniqueness of τ^* follows from $\frac{\partial^2 V}{\partial \tau^2} < 0$ (Lemma A1) and the fact that the RHS of eqn.(A6) is increasing in τ .

(iv) Eqn.(5) can be written as $E[w] = 2\tau^* c'(\tau^*)$. The equilibrium payoffs follow immediately.

Result 1 *The principal implements the same effort level regardless whether the signal is verifiable or not.*

Proof: If the signal is observed, the wage can be conditioned on x, s , and m . Since m does not contain any additional information, we can wlog. restrict attention to wages schedules that depend only on x and s . If the wage depends only on s , the agent is only rewarded if $s = \mu$ and at the equilibrium precision $\frac{\nu}{2\tau(\tau+\nu)}E[w] = c'(\tau)$. Since $\frac{\nu}{2\tau(\tau+\nu)} < \frac{1}{2\tau}$, this cannot be optimal.

If the wage depends on s and x , the agent chooses effort to maximize

$$\int \int w(s, x)q(x|s, \tau)dx f(s, \tau)ds - c(\tau)$$

which results in the f.o.c.

$$\int \int \left(\frac{1}{2\tau} - \frac{1}{2}(x-s)^2 \right) w(s, x)q(x|s, \tau)dx f(s, \tau)ds = c'(\tau)$$

Given any expected wage payment, the incentive to exert effort is maximal if $w(s, x) = 0$ if $s \neq x$ and $w(s, x) \geq 0$ if $s = x$. Hence, the f.o.c. can be written as $\frac{1}{2\tau}E[w] = c'(\tau)$ which is the same as in the case where the signal is unobservable.

Result 2 *The convexity of the wage is increasing in the precision of the prior: $\frac{w''}{w'}$ increases (decreases) in ν for $m > \mu$ ($m < \mu$).*

Proof: The result follows immediately from the expression for $w(m, x)$.

Result 3 *The expected wage that is necessary to implement a given effort does not depend on the precision of the prior.*

Proof: Since $w(m, x) = 0$ for $m \neq x$, eqn.(5) can be rewritten as $\frac{1}{2\tau^*}E[w] = c'(\tau^*)$. This proves the claim.

Proof of Proposition 2 (i) If the messages of all agents reveal the signal, the precision of the posterior of the principal is $\tau^* + \nu$ and the posterior mean is $\frac{\sum \tau_i^* s_i + \nu \mu}{\tau^* + \nu}$. Signals contain noisy information about the state. Since the state is observed, the principal cannot gain from making the wages dependent on the signals of other agents. Therefore, contracts have the same form as in Proposition 1.

Let $E[w_i]$ be the expected wage of agent i . As shown above: $E[w_i] = 2\tau_i^* c'(\tau_i^*)$. Then n^* follows immediately. For each agent, the principal chooses the equilibrium precision to minimize the sum of expected wages. Then $\tau_i^* = \frac{\tau^*}{N}$ if $c'' > 0$ and $\tau_i^* = \tau^*$ if $c'' < 0$ follow immediately. If $c'' = 0$, the only restriction on τ_i^* is that $\sum \tau_i^* = \tau^*$ and that $\tau_i^* \geq 0$ for all i .

Note that $\frac{\partial V}{\partial \tau}$ is the derivative with respect to τ^* with $\frac{\partial V}{\partial \tau} = \sum \left[2c'(\tau_i^*) \frac{\partial \tau_i^*}{\partial \tau^*} + 2\tau_i^* c''(\tau_i^*) \frac{\partial \tau_i^*}{\partial \tau^*} \right]$. To see that at τ^* we have $\frac{\partial V}{\partial \tau} = 2c'(\tau_i^*) + 2\tau_i^* c''(\tau_i^*)$, note that either $n^* = 1$ (if $c'' < 0$) or that $\tau_i^* = \frac{\tau^*}{N}$ with $\frac{\partial \tau_i^*}{\partial \tau^*} = \frac{1}{N}$ (if $c'' > 0$) or that $\sum \frac{\partial \tau_i^*}{\partial \tau^*} = 1$ and $c'' = 0$.

(ii) Under simultaneous expertise, the principal offers the same contract as in Proposition 1. Under sequential expertise, the second agent knows the contract and the message of the first agent. Therefore, the second agent and the principal update their prior. The updated prior is that $x \sim N\left(\frac{\tau_1^* s_1 + \nu \mu}{\tau_1^* + \nu}, \frac{1}{\tau_1^* + \nu}\right)$. Hence, the optimal contract for the second agent is the same contract as in Proposition 1 except that μ is replaced with $\frac{\tau_1^* s_1 + \nu \mu}{\tau_1^* + \nu}$ and that ν is replaced with $\tau_1^* + \nu$. As shown in Proposition 1, the expected wage depends only on the implemented precision but not on μ and ν . Hence, the expected wage of the second agent is the same under simultaneous and under sequential expertise. From induction follows that for arbitrary vector $\tau_i^*, \tau_j^*, \dots$, payoffs and expected wages under simultaneous and sequential expertise are the same

Proof of Proposition 3

(i) The proof is similar to the proof of Proposition 1. Using the revelation principle, let $m^*(s) = s$. Let $w_i(s_i, s_j)$ be the wage of agent i when the signals are s_i and s_j . Agent i chooses effort to maximize

$$\int \int w_i(s_i, s_j) q(s_j | s_i, \tau_i, \tau_j) ds_j f(s_i, \tau_i) ds_i - c(\tau_i)$$

where τ_i and τ_j are the precision of the signals and $q(s_j | s_i, \tau_i, \tau_j)$ is the conditional density of s_j . The signals are normal with mean x . Given s_i , we have $E[x | s_i] = \frac{\tau_i s_i + \nu \mu}{\tau_i + \nu}$. From the perspective of agent i , the signal of agent j is normal with mean $\frac{\tau_i s_i + \nu \mu}{\tau_i + \nu}$. Conditional on s_i , the variance of s_j

is the variance of $E[x|s_i]$ plus the unconditional variance of s_j , i.e. $\frac{1}{\tau_i + \nu} + \frac{1}{\tau_j} = \frac{\tau_j + \tau_i + \nu}{\tau_j(\tau_i + \nu)}$. Hence

$$q(s_j|s_i, \tau_i, \tau_j) = \frac{\sqrt{\tau_j(\tau_i + \nu)}}{\sqrt{2\pi}\sqrt{\tau_j + \tau_i + \nu}} \exp -\frac{1}{2} \left\{ \left(s_j - \frac{\tau_i s_i + \nu \mu}{\tau_i + \nu} \right)^2 \frac{\tau_j(\tau_i + \nu)}{\tau_j + \tau_i + \nu} \right\}$$

Then

$$\begin{aligned} \frac{\partial E[w_i]}{\partial \tau_i} &= \int \int \frac{\tau_j + \nu}{2\tau_i(\tau_j + \tau_i + \nu)} - \frac{1}{2} \left(s_j - \frac{\tau_i s_i + \nu \mu}{\tau_i + \nu} \right)^2 \frac{\tau_j^2}{(\tau_j + \tau_i + \nu)^2} \\ &\quad + \frac{\left(s_j - \frac{\tau_i s_i + \nu \mu}{\tau_i + \nu} \right) \tau_j \nu (s_i - \mu)}{(\tau_j + \tau_i + \nu)(\tau_i + \nu)} - \frac{1}{2} \frac{\nu^2 (s_i - \mu)^2}{(\tau_i + \nu)^2} w q d s_j f d s_i \end{aligned} \quad (\text{A7})$$

The principal decides for which signal s_j agent i is rewarded. Hence, the principal maximizes $\frac{\partial E[w_i]}{\partial \tau_i}$ with respect to s_j . Eqn.(A7) is maximal at $s_j = s_i + \frac{\nu(s_i - \mu)}{\tau_j}$. Let $\tilde{s}_j(s_i) = s_i + \frac{\nu(s_i - \mu)}{\tau_j}$.

Hence, $w_i(s_i, s_j) = 0$ if $s_j \neq \tilde{s}_j(s_i)$. Then

$$\begin{aligned} \frac{\partial E[w_i]}{\partial \tau_i} &= \int \int \left[\frac{\tau_j + \nu}{2\tau_i(\tau_j + \tau_i + \nu)} \right] w(s_i, \tilde{s}_j(s_i)) q(\tilde{s}_j(s_i)|s_i, \tau_i, \tau_j) d\tilde{s}_j(s_i) f d s_i \\ &= \frac{\tau_j + \nu}{2\tau_i(\tau_j + \tau_i + \nu)} E[w_i] \end{aligned} \quad (\text{A8})$$

Note that if $\tau_j = \infty$, the signal of agent j reveals the state and the contract is the same as if the state is observable with $\frac{\partial E[w_i]}{\partial \tau_i} = \frac{1}{2\tau_i} E[w_i]$.

Given that $w(m_i, s_j) = 0$ if $s_j \neq \tilde{s}_j(m_i)$, agent i chooses a message to maximize his expected payoff. Formally

$$\max_{m_i} q(\tilde{s}_j(m_i)|s_i, \tau_i, \tau_j) \cdot w_i(m_i, \tilde{s}_j(m_i))$$

Hence, for all m

$$\frac{\partial q}{\partial \tilde{s}_j(m_i)} \frac{\partial}{\partial m_i} \frac{1}{q} = - \frac{\partial w_i}{\partial m_i} \frac{1}{w(m_i, \tilde{s}_j(m_i))}$$

After substituting for q and solving the differential equation for w_i , we have

$$w_i(m_i, \tilde{s}_j(m_i)) = k \exp \left\{ \left(1 + \frac{\nu}{\tau_j} \right) \left(\frac{1}{2} \nu m_i^2 - \nu m_i \mu \right) \right\}$$

and, of course, $w_i(m_i, m_j) = 0$ whenever $m_j \neq \tilde{s}_j(m_i)$.

Whenever $\tau^* > 0$, the agent earns rent and, therefore, accepts the contract.

(ii) Follows from Lemma 1 and the fact that $m(s) = s$.

(iii) The principal chooses τ_i, τ_j to minimize $E[w_i] + E[w_j]$ given that $\tau_i + \tau_j = \tau^*$. At the precision that maximizes the payoff of the agent, we have $\frac{\partial E[w]}{\partial \tau} = c'(\tau)$. From eqn.(A8) follows that $E[w_i] = \frac{2\tau_i(\tau_j + \tau_i + \nu)}{\tau_j + \nu} c'(\tau_i)$. Hence, the principal minimizes

$$\frac{2\tau_i(\tau_j + \tau_i + \nu)}{\tau_j + \nu} c'(\tau_i) + \frac{2\tau_j(\tau_j + \tau_i + \nu)}{\tau_i + \nu} c'(\tau_j) \quad \text{s.t. } \tau_i + \tau_j = \tau^* \quad (\text{A9})$$

Since $c''' \geq 0$, from eqn.(A9) follows immediately that $\tau_i^* = \tau_j^*$. Hence, $E[w_i] = \frac{2\tau_i^*(2\tau_i^* + \nu)}{\tau_i^* + \nu} c'(\tau_i^*)$. At the optimal precision that the principal implements we have $\frac{\partial V}{\partial \tau} = 2 \frac{dE[w_i]}{d\tau_i^*} \frac{\partial \tau_i^*}{\partial \tau^*}$. Hence, at τ^* we have $\frac{\partial V}{\partial \tau} = 4c'(\tau_i^*) - \frac{2\nu^2}{(\tau_i^* + \nu)^2} c'(\tau_i^*) + 2\tau_i^* c''(\tau_i^*) + \frac{2\tau_i^{*2}}{\tau_i^* + \nu} c''(\tau_i^*)$.

(iv) Follows immediately from substitution.

References

- Chevalier, J. and G. Ellison, "Career Concerns of Mutual Fund Managers," *Quarterly Journal of Economics*, 1999, 114(2), 389-432
- Cr mer, J. and F. Khalil, "Gathering Information before Signing a Contract," *American Economic Review*, 1992, 82(3), 566-578
- Cr mer, J. F. Khalil, and J-C. Rochet, "Strategic Information Gathering before a Contract is Offered," *Journal of Economic Theory*, 1998a, 81, 163-200
- Ehrbeck, T. and R. Waldmann, "Why are Professional Forecasters Biased? Agency versus Behavioral Explanations," *Quarterly Journal of Economics*, 1996, 111(1), 21-40
- Ewerhart, C., and P.W. Schmitz, "'Yes Men', Integrity, and the Optimal Design of Incentive Contracts," *Games and Economic Behavior*, 2000, 43, 115-125

- Gromb, D. and D. Martimort, "The Organization of Delegated Expertise," mimeo, London Business School, 2003
- Hong, H., J.D. Kubik, and A. Solomon, "Security Analysts's Career Concerns and Herding of Earning Forecasts," *RAND Journal of Economics*, 2000, 31(1), 121-144
- Krishna, V. and J. Morgan, "A Model of Expertise," *Quarterly Journal of Economics*, 2001, 116(2), 747-775
- Lamont, O.A., "Macroeconomic Forecasts and Microeconomic Forecasters," *Journal of Economic Behavior and Organization*, 2002, 48, 265-280
- Laster, D., P. Bennett, and I.S. Geoum, "Rational Bias in Macroeconomic Forecasts," *Quarterly Journal of Economics*, 1999, 114(1), 293-318
- Li, H., "A Theory of Conservatism," *Journal of Political Economy*, 2001, 109(3), 617-636
- Lewis, T. and D. Sappington, "Information Management in Incentive Problems," *Journal of Political Economy*, 1997, 105(4), 796-821
- Lewis, T. and D. Sappington, "Using Subjective Risk Adjusting to Prevent Patient Dumping in the Health Care Industry," *The Journal of Economics and Management Strategy*, 1999, 8(3)
- Osband, K., "Optimal Forecasting Incentives," *Journal of Political Economy*, 1989, 97(5), 1092-1112
- Ottaviani, M. and P.N. Sørensen, "Forecasting and Rank-Order Contests," mimeo, London Business School and University of Copenhagen, 2003
- Ottaviani, M. and P.N. Sørensen, "Professional Advice" mimeo, London Business School and University of Copenhagen, 2004
- Pesendorfer, W. and A. Wolinsky, "Second Opinions and Price Competition: Inefficiency in the Market for Expert Advice," *Review of Economic Studies*, 2003, 70(2), 417-439
- Prendergast, C., "A Theory of 'Yes Men'," *American Economic Review*, 1993, 83(4), 757-770
- Szalay, D., "The Economics of Clear Advice and Extreme Options," Working Paper, HEC University of Lausanne, 2003 (forthcoming *Review of Economic Studies*)