# He Who Pays The Piper Must Know The Tune

Robin Hanson
Department of Economics
George Mason University*

April 2005 (First version May 2003)

**Abstract**

He who pays the piper calls the tune, but he can only succesfully call for a tune that he will recognize upon hearing. Previous models, of two candidates impressing a voter and of firm managers impressing stock speculators, found experts ignoring costly superior information in favor of client preconceptions. Similar result hold when we greatly generalize the agents, choices, information structures, and preferences. When experts must pay to acquire information, have no intrinsic interest in client topics, and can coordinate to acquire the same information, no expert ever pays to know more than any client will know when rewarding those experts.

## Introduction

If you can choose how much to pay the piper after he plays, you can call for him to play a particular tune that you know, a tune in a style that you know, a happy tune, or a tune that lasts an hour. But it will be hard for you to successfully call for an authentic fourteenth century Scottish ballad if the piper knows that you would not recognize such a ballad if you heard one. In this case he can get paid by just playing a tune that you cannot distinguish from such a ballad; he need not actually work to find an authentic ballad. If the piper knows that he who pays has a misconception about what such ballads sound like, the piper is usually better off playing a tune that fits the misconception than trying to correct it. While these lessons may seem obvious to some readers, their implications and generality may not yet be fully appreciated.

In two previous papers that were the immediate inspiration for this one, the authors noted the equivalent of pipers choosing their tunes based on the preconceptions of he who pays, neglecting their superior information. In a model by Heidhues and Lagerlöf (2003), two

**rhanson@gmu.edu* http://hanson.gmu.edu 703-993-2326 FAX: 703-993-2323 MSN 1D3, Carow Hall, Fairfax VA 22030-4444

candidates compete for the support of a representative voter and choose policy platforms before an election. Even though the candidates know more than the voter about policy consequences, the candidates may ignore that information in favor of the platform that looks better given the voter information. While candidates may sometimes correlate their platforms with their superior information, this only happens when candidates are indifferent to the platform they propose.

Similarly, in a model by Brandenburger and Polak (1996), a set of firm managers each try to raise their stock price in the eyes of stock speculators. Each manager makes a similar decision in sequence, observing the decisions of previous managers, and speculators then set the stock price of each firm based on the decision of that firm's manager. Even though the managers know more than the speculators, they may ignore that information in favor of the decision that looks better to the speculators. While managers may sometimes correlate their decisions with their superior information, this only happens when managers are indifferent to which decision they make.

We can consider these to be signaling models between a set of experts, such as candidates and managers, and a set of clients, such as voters and speculators.[1] Experts give advice, to which clients then respond. If we modify these models to give the experts a hidden choice about whether to pay a tiny amount to acquire their superior expert information, we find that the experts in these models will not pay anything to acquire (or retain) their expert information, and so their decisions will not depend on such information.

These models seem to call into question our common practice of relying on the superior information of professional experts. But as relatively specific examples, these models are limited in their ability to indicate how frequent such problems are, or what situational features mitigate them. In particular, these examples tell us little about whether clients can do better by playing experts off against one another. To make more progress on these issues, we might seek more general theoretical results.

In this paper I find that key results for these specific models also hold in a much more general model. Heidhues and Lagerlöf, and Brandenburger and Polak, consider binary expert choices, binary expert information signals, common symmetric priors, specific simple utility functions, and expert utility that is independent of the signals experts send.[2] In this paper I allow relatively general finite sets of experts, clients, choices, preferences, information structures, and information efforts. (This generality allow for general contracts as well.)

I find sufficient conditions under which, if there is a focal information level which each expert can achieve that contains all client information, no expert will pay any amount to acquire more than this focal information. These sufficient conditions have three main components. First, I assume that while client preferences can be state-dependent, making clients

[1]Many other papers have considered signaling models between experts and clients (Morris 2001, Krishna and Morgan 2001).

[2]Cummins and Nyman (2004) offer a modest generalization of Brandenburger and Polak. Ottaviani and Sorensen (2003) also have a similar result (in their proposition 12), for a single expert and client, a particular information structure, and expert and client utilities independent of expert actions. Also, Smart and Sturm (2003) have a related model, with a single expert and client, and hidden expert information on his binary preference type over expert actions.

directly interested in acquiring information about the state, experts do not intrinsically care about the state, and so have no direct reason to acquire information, other than to influence clients. Second, I assume that experts can sufficiently coordinate their actions, either by acting in sequence and observing all previous expert actions, or by having two experts with zero-sum payoffs choose simultaneously, since there are no gains from coordination in a two-person zero-sum game.

Finally, I assume that information is costly, and that along each of a set of information dimensions, information levels can be ordered relative to the focal information level. That is, the total information of an expert is the combination of his information along each dimension, and every possible information level along each dimension is either strictly more informed than, less informed than, or exactly as informed as the focal information level along that dimension. Furthermore, every more informed level gives experts strictly less utility, holding all else constant.

We thus have a relatively general result to the effect that when information acquisition is costly, experts have no intrinsic interest in client topics, and experts can observe each other's actions enough to coordinate their actions and information levels, experts never want to know any more than clients will know when they reward the experts. Giving clients a relatively general ability to reward each expert according to the actions of other experts does not help if the experts have enough information to coordinate their actions.

Under a wider range of conditions than considered by Heidhues and Lagerlöf, candidates in two candidate elections do not want to know more than voters will know when they vote. Under a wider range of conditions than those considered by Brandenburger and Polak, firm managers who act sequentially and visibly do not want to know more than stock speculators will know when managers exercise their stock options.

The generality of this result raises questions about our relations to a wide range of other kinds of experts. We often seem to rely on doctors, newspaper reporters, university professors, and other experts to tell us things that we do not expect to be in a position to confirm when we reward them. We usually believe our doctor's diagnoses, even though we are reluctant to blame them for our poor health, nor are we eager to have their pay depend on our ultimate health. We rarely have any direct experience of events in foreign lands that we read about in newspaper articles, or of the events in ancient times we hear about in history lectures, yet we usually believe such newspaper articles and history lectures. The result of this paper suggests that we are either too trusting of such experts, that we do not much care whether their stories are true, or that we succeed in making them trustworthy by relying heavily either on expert non-optimization, on experts caring directly about being honest or about client topics, on visible or non-positive-cost expert information efforts, on the responses of the few clients who do get direct experience on expert claims, or on difficulties of expert coordination, such as when the easiest way for experts to tell clients the same thing is to tell them the truth.

In this paper I first review the models of Heidhues and Lagerlöf and of Brandenburger and Polak, placing them into a common notation and adding in an expert information effort. I then present a more general model, followed by two general results proven within that general model. Finally, I offer concluding comments.

# Prospective Elections

Heidhues and Lagerlöf (2003) have a model where two candidates commit to endorsing one of two policies, after which a single representative voter elects one candidate. That is, Heidhues and Lagerlöf consider a set $E = \{1, 2\}$ of two candidates $i$ each of which simultaneously makes a binary choice of an election platform $a_i \in B = \{0, 1\}$. A single representative voter ($i = 0$) sees both platforms, and then chooses the election winner $r \in E$.

The voter is best off when the winning candidate's platform $a_r$ matches a certain hidden binary value. To aid each candidate in making this match, he has available to him a noisy binary clue about this hidden value. For example, in the state $\omega = (1, 1, 0)$ the true hidden value is 1, candidate 1 gets a good clue of 1, and candidate 2 gets a misleading clue of 0. More generally, there are eight possible states $\omega$ which can be described by lists of three bits $b_i(\omega) \in B = \{0, 1\}$, so that $\omega = (b_0(\omega), b_1(\omega), b_2(\omega))$, where $b_0(\omega)$ is the hidden value that platforms should match, and $b_i(\omega)$ is the clue available to manager $i \in E$. The set of all states is $\Omega = B^3$.

At each state $\omega$ each agent $i$ initially has an information set $\pi_{i\omega} \subseteq \Omega$, which describes the set of states he considers to be possible. (Each collection $\pi$ of information sets $\pi_\omega$ form partitions, so that either $\pi_\omega = \pi_{\omega'}$ or $\pi_\omega \cap \pi_{\omega'} = \emptyset$.) Aside from hearing the candidate platforms, the voter is uninformed, so $\pi_{0\omega} = \Omega$. Candidates are initially uninformed as well, which we can write as $\tilde{\pi}_{i\omega} = \Omega$ for $i \in E$. Candidates then make a hidden binary choice of information effort $e_i \in B$, which determines a new information level $\pi_{i\omega}(e_i)$. (This information effort choice was not in Heidhues and Lagerlöf's version; I have added it for reasons that should soon become clear.) Candidates who choose zero effort remain uninformed, with $\pi_{i\omega}(0) = \Omega$, but candidates who choose unit effort learn their clue, with $\pi_{i\omega}(1) = \{\omega' \in \Omega : b_i(\omega') = b_i(\omega)\}$.

The voter and candidates share a common prior $p_\omega$, which can be written

$$p_\omega = \begin{cases} \kappa(b_0)((1-\theta)^2 + \rho\theta(1-\theta)) & \text{if } b_0 = b_1 = b_2 \\ \kappa(b_0)(1-\rho)\theta(1-\theta) & \text{if } b_1 \neq b_2 \\ \kappa(b_0)(\theta^2 + \rho\theta(1-\theta)) & \text{if } b_0 \neq b_1 = b_2 \end{cases}$$

where $\kappa(1) = q$, $\kappa(0) = 1 - q$, $q \in (1/2, 1)$, $\theta \in (0, 1/2)$, and $\rho \in [0, 1)$. This gives the candidates equally informative private clues, with $\rho$ describing their degree of correlation.

Candidate utility for being elected and avoiding information effort can be written in terms of candidate platforms $a \equiv (a_1, a_2)$ and the election winner $r$ as

$$u_i(a, r, e_i) = \mathbf{1}[r = i] - \epsilon\, e_i,$$

where $\mathbf{1}[P]$ is 1 when the proposition $P$ is true, and 0 otherwise. Voter utility can be written

$$v_{0\omega}(a, r) = \mathbf{1}[b_0(\omega) = a_r].$$

The voter gets 1 when the winner's platform matches the hidden value, and 0 otherwise.

Heidhues and Lagerlöf (2003) in effect consider the case where $\epsilon < 0$, so that candidates are sure to become informed regardless of other factors. They find (in their proposition 1) that there are equilibria where one candidate sets his platform to his clue and always wins, while the other always picks an opposite platform from his clue and always loses. To exclude such implausible situations Heidhues and Lagerlöf require voter strategies to be independent of the candidate labels. They then find (in their proposition 2) that the only pure strategy equilibria are babbling and uninformative, so that a particular platform is always implemented, independent of candidates clues. Heidhues and Lagerlöf do find (in their propositions 3,4) mixed strategy equilibria where candidates are indifferent between which platform they choose, and choose platforms that are correlated with their clues.

Heidhues and Lagerlöf find it noteworthy that there is "a strong incentive for [candidates] to bias their messages toward the electorate's prior," and neglect their own clues. But we get even more striking results if we consider the more plausible case where information is costly, i.e., where $\epsilon > 0$. Since Heidhues and Lagerlöf find that when $\epsilon < 0$ the candidates only ever use their information when they are indifferent to using it or not, it is clear that when $\epsilon > 0$ the candidates will never pay to see their clues, and so candidate platforms will never be informative. We will see below that this result holds more generally.

# Firm Managers

Brandenburger and Polak (1996) have a similar model, where firm managers "cover their posteriors" by "making the decisions the market wants to see." We can describe their model in a similar framework.

Brandenburger and Polak consider a set $E = \{1, 2, \ldots, N\}$ of $N$ managers $i$, each of which manages a separate firm, and each of which makes a binary corporate decision $a_i \in B = \{0, 1\}$. These decisions are made in turn, so that each manager knows the decisions of previous managers. There are also stock markets (which we will describe by a single market agent $i = 0$) that see all manager decisions and in the end set the share price $\hat{r}_i$ of each firm to their expectation of that firm's value given its choice.

Ideally each manager decision $a_i$ would match a certain hidden value, and to aid this there is available to each manager a noisy clue about this hidden value. Thus states $\omega$ can be described by a sequence of bits $b_i(\omega) \in B = \{0, 1\}$, so that $\omega = (b_0(\omega), b_1(\omega), \ldots, b_N(\omega))$, where $b_0(\omega)$ is the hidden value that choices should match, and $b_i(\omega)$ is the clue available to manager $i \in E$. The set of all states is $\Omega = B^{N+1}$.

The markets are initially uninformed, so that $\pi_{0\omega} = \Omega$, and managers are initially uninformed as well, with $\tilde{\pi}_{i\omega} = \Omega$ for $i \in E$. Manager information effort is binary, with $e_i \in B$, and zero effort leaves managers uninformed, so that $\pi_{i\omega}(0) = \Omega$. Managers who exert effort learn their clue, so that $\pi_{i\omega}(1) = \{\omega' \in \Omega : b_i(\omega') = b_i(\omega)\}$. (This information effort choice was not in Brandenburger and Polak's version.) The markets and managers share a common prior $p_\omega$, which can be written

$$p_\omega = \begin{cases} \theta q^{b(\omega)}(1-q)^{N-b(\omega)} & \text{if } b_0(\omega)=1 \\ (1-\theta)q^{N-b(\omega)}(1-q)^{b(\omega)} & \text{if } b_0(\omega)=0, \end{cases}$$

where $b(\omega) \equiv \sum_{i=1}^{N} b_i(\omega)$, and $\theta, q \in (1/2, 1)$. This gives each manager an equally informative and independent clue.

The utility of manager $i$ for a high stock price and low information effort can be described, in terms of the set of all decisions $a = (a_i)_{i \in E}$ and the set of all stock prices $r = (\hat{r}_i)_{i \in E}$, as

$$u_i(a, r, e_i) = \hat{r}_i - \epsilon e_i.$$

Each firm is worth 1 when $a_i = b_0(\omega)$, so that its manager's action matches the hidden value, and 0 otherwise. We can model the stock markets by a single market agent who sets his stock price responses $\hat{r}_i \in \hat{R}_i = \{p(S) : S \subseteq \Omega\}$ (where $p(S) = \sum_{\omega \in S} p_\omega$) to maximize a quadratic utility

$$v_{0\omega}(a, r) = -\sum_{i \in E} (\hat{r}_i - \mathbf{1}[a_i = b_0(\omega)])^2.$$

Such a market agent will set the share price $\hat{r}_i$ of each firm to the expectation of that firm's value given its choice, using this market agent's information.

Brandenburger and Polak (1996) in effect consider the case where $\epsilon < 0$, so that managers are sure to become informed regardless of other factors. For this case they find (in their Proposition 5) that "there is no pure strategy equilibrium in which any firm plays an informative strategy," i.e., which correlates their decision $a_i$ with their clue $b_i(\omega)$. While Brandenburger and Polak do find informative play when mixed strategies are allowed, in such equilibria managers are indifferent between their possible decisions.

Brandenburger and Polak find it noteworthy that manager decisions tend to follow "the market's 'prejudices' " instead of "those suggested by their own superior information." But we get even more striking results if we consider the more plausible case where information is costly, i.e., where $\epsilon > 0$. Since Brandenburger and Polak find that when $\epsilon < 0$ the managers only ever use their information when they are indifferent to using it or not, it is clear that when $\epsilon > 0$ the managers will never pay to see their clues, and so manager decisions will never be informative.[3] We will now see that this result holds more generally.

# General Model

Let us now consider a more general model, of which the above two models are examples. Consider a signaling game between a finite set $E$ of $N$ experts (or senders) who give advice

[3]The managers are better off in a mixed strategy equilibria where their decisions correlate with their clues. However, even if a manager expected all other managers to buy their information and then all other players to play according to such a mixed strategy equilibria, this manager is better off deviating and not buying his information.

to (and take actions on behalf of) a finite set $C$ of $M$ clients (or receivers), who then respond to (and as a result reward) those experts for their advice (and actions). Each expert $i$ gives advice (and takes actions) $a_{ij} \in A_{ij}$ for each client $j$, who then chooses a corresponding response (and reward) $r_{ij} \in R_{ij}$. The sets $A_{ij}$ and $R_{ij}$ are finite, and can be singletons, to model the possibility that expert $i$ and client $j$ do not interact.

The set of all advice given by expert $i$ is $a_i \equiv (a_{ij})_{j \in C} \in A_i \equiv \times_{j \in C} A_{ij}$, while the set of responses given by client $i$ is $r_i \equiv (r_{ji})_{j \in E} \in R_i \equiv \times_{j \in E} R_{ji}$. The set of all advice and actions of all experts is $a \equiv (a_i)_{i \in E} \in A \equiv \times_{i \in E} A_i$, while the set of all responses of all clients is $r \equiv (r_i)_{i \in C} \in R \equiv \times_{i \in C} R_i$.

Let each expert and client be given a unique ordered label. Let the expert set be written $E = \{1, 2, \ldots, N\}$, and the client set be written $C = \{N+1, \ldots, N+M\}$. We consider two variations on what experts know when they give advice. In a simultaneous variation, each expert does not know what advice the other experts give when he gives his advice. In the sequential variation, experts give their advice sequentially, so each expert knows what advice the previous experts have given when he gives his advice. That is, expert $i$ knows $a_i^- \equiv (a_j)_{j \in E: j < i} \in A_i^- \equiv \times_{j \in E: j < i} A_j$ when choosing $a_i$.

After experts choose their advice, clients make their responses sequentially, according to their numerical labels. When he makes his response, each client knows something about the advice of the experts and the responses of previous clients. Specifically, client $i$ knows the value of $F_{i\omega}(a, r_i^-)$ where $F_{i\omega}$ is a function from $A \times R_i^-$ onto a finite set, $r_i^- \equiv (r_j)_{j \in C: j < i} \in R_i^- \equiv \times_{j \in C: j < i} R_j$. By having $F_{i\omega}$ be one to one, then client $i$ can know all of $a$ and $r_i^-$, and by having $F_{i\omega}$ be constant client $i$ can know nothing. Otherwise he knows some intermediate amount. (The subscript $\omega$ in $F_{i\omega}$ allows clients to receive noisy state-dependent signals.)

Each expert and client $i$ has a prior $p_{i\omega} > 0$ over possible states $\omega \in \Omega$ (finite), and at each state has an information set $\pi_{i\omega} \subseteq \Omega$ representing his information there. Information sets $\pi_\omega$ form partitions, and we say partition $\mu$, and the agent who holds it, is at least as informed as partition $\nu$ if for all $\omega \in \Omega$, $\mu_\omega \subseteq \nu_\omega$, and is more informed if in addition $\mu_\omega \subset \nu_\omega$ for some $\omega \in \Omega$. Partition $\nu$ is then, respectively, no more informed or less informed than $\mu$. (For most possible partitions $\mu$ and $\nu$, neither one is at least as informed as the other.)

When choosing his advice $a_i$, each expert $i$ will know the previous advice $a_i^-$ in the sequential variation, but not in the simultaneous variation. In addition, he will know an information set $\pi_{i\omega}(e_i)$, which depends on his information effort $e_i \in \mathcal{E}_i$ (finite). Experts choose their information efforts $e_i$ immediately before choosing advice $a_i$. Efforts $e_i$ are private, and so are not directly observed by clients or other agents. When choosing his response $r_i$, each client $i$ knows an information set $\pi_{i\omega}$, in addition to $F_{i\omega}(a, r_i^-)$. (For consistency, we assume $F_{i\omega}$ is constant across $\tilde{\pi}_{i\omega}$.)

Let us further assume that information effort $e_i$ is composed of $D$ dimensions $d$, so that $e_i = (e_{id})_{d \in D} \in \mathcal{E}_i = \times_{d \in D} \mathcal{E}_{id}$ (each finite), and $\pi_{i\omega}(e_i) = \bigcap_{d \in D} \pi_{id\omega}(e_{id})$. That is, the total information $\pi_{i\omega}$ of each agent simply combines his information $\pi_{id\omega}$ along each dimension $d$.

After the experts have chosen their efforts $e_i$ and advice $a_i$, and clients have chosen their responses $r_i$, each expert receives utility $u_i(a, r, e_i)$, while each client receives utility $v_{i\omega}(a, r)$. Note that while clients can directly want advice $a$ and responses $r$ to be related to the state

$\omega$, we assume experts only care about the state $\omega$ indirectly, via advice $a$ or responses $r$ depending on the state $\omega$. Not also that this formulation allows experts to be subject to relatively arbitrary contracts and incentives; the only requirement is that those who enforce such contracts and incentives must be either experts, with no direct preferences regarding states, or clients, with their limited information.

If effort $e_i$ makes expert $i$ at least as informed as effort $e_i'$, let us say the added information is costly to acquire when $\forall a \in A, r \in R, u(a,r,e_i) \leq u(a,r,e_i')$, and is strictly costly to aquire when this inequality is strict. Information $\pi_{id\omega}(e_{id})$ along a dimension is costly when it is costly to gain it, holding constant efforts along the other dimensions.

Thus the order of events and strategies are as follows. After nature chooses the true state $\omega \in \Omega$ according to a prior $p_{0\omega}$, each expert $i$ is in turn given his initial information set $\tilde{\pi}_{i\omega}$, and in the sequential variation is also told the advice $a_i^-$ of previous experts. This expert $i$ then chooses his information effort $e_i \in \mathcal{E}_i$, with a mixed strategy probability $\alpha_{i\omega}(e_i|a_i^-)$, after which he is given further information $\pi_{i\omega}(e_i) \subseteq \tilde{\pi}_{i\omega}$. Expert $i$ then chooses his advice $a_i \in A_i$, with mixed strategy probability $\beta_{i\omega}(a_i|a_i^-, e_i)$. (In the simultaneous version these strategies are written $\alpha_{i\omega}(e_i)$ and $\beta_{i\omega}(a_i|e_i)$, dropping the $a_i^-$ argument.)

Similarly, after the experts choose their advice, each client $i$ is in turn given his information set $\pi_{i\omega}$, is told the value of $F_{i\omega}(a, r_i^-)$, and then chooses his response $r_i \in R_i$ with a mixed strategy probability $\lambda_{i\omega}(r_i|F_{i\omega}(a, r_i^-))$. Note that while we have written these agent strategies $\alpha_{i\omega}(e_i|a_i^-)$, $\beta_{i\omega}(a_i|a_i^-, e_i)$ and $\lambda_{i\omega}(r_i|F_{i\omega}(a, r_i^-))$ as varying with the state $\omega$, they must in fact be constant over each agent's then relevant information set $\pi_{i\omega}$; no one can have their strategy depend on things that they do not know.

The models we considered earlier, of competing candidates seeking a voter's approval and of firm managers seeking to impress speculators, are both examples of this more general model. Heidhues and Lagerlöf's model was a simultaneous version, with two expert candidates and a single voting client. Brandenburger and Polak's model was a sequential version, with $N$ expert managers and a single market speculator client. While those models were not originally formulated with expert information effort, such efforts are easy to add, and were included in the descriptions above.

# Analysis

Let us assume the existence of a *focal* information partition $\hat{\pi}_\omega = \bigcap_{d \in D} \hat{\pi}_{d\omega}$. By definition, a focal information set is at least as informed as each client set (so $\forall j \in C, \hat{\pi}_\omega \subseteq \pi_{j\omega}$), there is an effort $\hat{e}_i$ by which each expert can acquire this information ($\forall i \in E, \exists \hat{e}_i \in \mathcal{E}_i : \pi_{i\omega}(\hat{e}_i) = \hat{\pi}_\omega$). Furthermore, when we vary information effort along a single dimension $d$, all efforts $e_{id} \in \mathcal{E}_{id}$ along that dimension either make the expert at least as informed ($\pi_{id\omega}(e_{id}) \subseteq \hat{\pi}_{d\omega}$), no more informed ($\pi_{id\omega}(e_{id}) \supseteq \hat{\pi}_{d\omega}$), or exactly as informed ($\pi_{id\omega}(e_{id}) = \hat{\pi}_{d\omega}$). Also by definition, it is strictly costly to acquire more than the focal information along any dimension. It is easy to see that even if there are two or more distinct focal information partitions, there must be a unique least informed and so strictly least costly, or *cheapest*, focal information partition.

This focal information assumption is satisfied when there are efforts $\tilde{e}_i \in \mathcal{E}_i$ that allow experts to keep their initial information partitions ($\pi_{i\omega}(\tilde{e}_i) = \tilde{\pi}_{i\omega}$), which are the same ($\forall i, j \in E, \tilde{\pi}_{i\omega} = \tilde{\pi}_{j\omega}$), at least as informed as clients ($\forall i \in E, j \in C, \tilde{\pi}_{i\omega} \subseteq \pi_{j\omega}$), and strictly easier than other efforts ($u(a, r, e_i) < u(a, r, \tilde{e}_i) \, \forall e_i \neq \tilde{e}_i \in \mathcal{E}_i$). This was the case in the models of competing candidates and firm managers we considered above when $\epsilon > 0$.

The net result of the mixed strategies $\alpha_{i\omega}(e_i|a_i^-)$ and $\beta_{i\omega}(a_i|a_i^-, e_i)$ of each expert $i$ is a distribution $\gamma_{i\omega}(a_i|a_i^-) \equiv \sum_{e_i \in \mathcal{E}_i} \beta_{i\omega}(a_i|a_i^-, e_i) \, \alpha_{i\omega}(e_i|a_i^-)$. (In the simultaneous variation the $a_i^-$ arguments are dropped.) The net result of all client mixed strategies $\lambda_{i\omega}(r_i|F_{i\omega}(a, r_i^-))$ is a distribution $\lambda_\omega(r|a) \equiv \prod_{i \in C} \lambda_{i\omega}(r_i|F_{i\omega}(a, r_i^-))$. (For this finite extensive form game, we know that given common priors there exists a mixed strategy sequential equilibrium.)

Client strategies can be averaged over to produce an average expert payoff of

$$\bar{u}_{i\omega}(a, e_i) \equiv \sum_{r \in R} u_i(a, r, e_i) \, \lambda_\omega(r|a),$$

In the sequential variation, the strategies of later experts can be averaged over to produce an average expert utility of

$$U_{i\omega}(a_i|a_i^-, e_i) \equiv \sum_{a_i^+ \in A_i^+} \bar{u}_i((a_i^-, a_i, a_i^+), e_i) \prod_{j \in E: j > i} \gamma_{j\omega}(a_j|a_j^-),$$

where $a_i^+ \equiv (a_j)_{j \in E: j > i} \in A_i^+ \equiv \times_{j \in E: j > i} A_j$.

The equilibrium concept we will use is that each expert's strategy optimizes his choice of $e_i$ and $a_i$ in contexts (such as $a_i^-$) that are reached with positive probability. Clients need *not* optimize, however; a client may use any strategy that is constant over his information set $\pi_{i\omega}$. When an expert optimizes, $\beta_{i\omega}(a_i|a_i^-, e_i) > 0$ implies that $a_i$ maximizes

$$\bar{U}_{i\omega}(a_i|a_i^-, e_i) \equiv \sum_{\omega' \in \pi_{i\omega}(e_i)} U_{i\omega'}(a_i|a_i^-, e_i) p_{i\omega'} / p_i(\pi_{i\omega}(e_i)),$$

and $\alpha_{i\omega}(e_i|a_i^-) > 0$ implies that $e_i$ maximizes

$$\bar{\bar{U}}_{i\omega}(a_i|a_i^-, e_i) \equiv \sum_{\omega' \in \tilde{\pi}_{i\omega}} \bar{U}_{i\omega'}(a_i|a_i^-, e_i) p_{i\omega'} / p_i(\tilde{\pi}_{i\omega}).$$

In the simultaneous variation, the definitions of $\bar{U}_{i\omega}(a_i|e_i)$ and $\bar{\bar{U}}_{i\omega}(a_i|e_i)$ are the same, except that the $a_i^-$ arguments are dropped. The only other definition change in the simultaneous variation is to average expert utility, which becomes

$$U_{i\omega}(a_i|e_i) \equiv \sum_{a_i^+ \in A_i^+, a_i^- \in A_i^-} \bar{u}_i((a_i^-, a_i, a_i^+), e_i) \prod_{j \in E: j \neq i} \gamma_{j\omega}(a_j).$$

Let us say that two experts $i, j$ have *zero-sum payoffs* when there exists some real-valued function $f_{ij}(u; r, e_i, e_j)$ strictly decreasing in its first real-valued argument such that

$$u_i(a, r, e_i) = f_{ij}(u_j(a, r, e_j); r, e_i, e_j).$$

And let us say that events which happen with zero probability *never* happen. Given these definitions, we can express our two main results (proofs in an appendix).

**Proposition 1** *With simultaneous choice by two experts with zero-sum payoffs, optimizing experts never acquire more than the cheapest focal information.*

**Proposition 2** *With sequential expert choice, optimizing experts never acquire more than the cheapest focal information.*

We thus have some relatively general results to the effect that when information acquisition is costly, experts have no intrinsic interest in client topics, and experts can observe each other's actions enough to coordinate their actions and information levels, experts never want to know any more than clients will know when they reward the experts. Proposition 1 is similar to the main result of Heidhues and Lagerlöf (2003), showing relatively generally that in two candidate elections, candidates will not want to know more than voters will know when they vote. Similarly, proposition 2 is similar to the main result of Brandenburger and Polak (1996), showing relatively generally that firm managers who move sequentially and visibly will not want to know more than stock speculators will know when managers exercise their stock options and related rewards. (In both applications, the no information level is the focal level.)

There are several obvious ways to generalize the above results, which I did not show here to avoid introducing a more complex and hence opaque notation. For example, the same results should obtain if each expert and client could take actions at more than one time, if clients could choose information efforts, or if all expert and client utilities depended on all information efforts.

# Conclusion

In this paper I reviewed models of Heidhues and Lagerlöf (2003), and of Brandenburger and Polak (1996), placing them into a common framework and adding in an expert information effort. These models considered binary expert choices, binary expert information signals, common symmetric priors, specific utility functions, and expert utility that is independent of the signals they send. After reviewing these models, I considered more general finite sets of experts, clients, choices, preferences, and information structures, and found sufficient conditions under which, if there is a focal information level which each expert can achieve that contains all client information, no expert will pay any positive amount to acquire more information than this level. This gives us general results to the effect that when information acquisition is costly and hidden, experts have no intrinsic interest in client topics, and experts can observe each other's actions enough to coordinate their actions and information levels, experts never want to know any more than clients will know when they reward the experts.

The apparent generality of this result raises questions about our relations to a wide range of other kinds of experts. We often seem to rely on doctors, newspaper reporters, university professors, and other experts to tell us things that we do not expect to be in a position to confirm when we reward them. The result of this paper suggests that we are either too trusting of such experts, that we do not much care whether their stories are true, or that we succeed in making them trustworthy by relying heavily either on expert non-optimization, on experts caring directly about being honest or about client topics, on visible or non-positive cost expert information efforts, on the responses of the few clients who do get direct experience on expert claims, or on difficulties of expert coordination, such as when the easiest way for experts to tell clients the same thing is to tell them the truth.

# Proof Appendix

## Proof of Proposition 1

Since each client strategy $\lambda_{i\omega}(r_i|F_{i\omega}(a, r_i^-))$ is constant across his information set $\pi_{i\omega}$, and each focal information set is contained with any client information set ($\hat{\pi}_\omega \subset \pi_{i\omega}$), then each client strategy $\lambda_{i\omega}(r_i|F_{i\omega}(a, r_i^-))$ is constant across the focal set $\hat{\pi}_\omega$. Thus expert payoffs $\bar{u}_{i\omega}(a, e_i)$, which average over client responses $r$, are constant across the focal set.

Given the assumption of zero-sum expert payoffs and of simultaneous expert choice, then holding constant their information efforts $e_i$ and client responses $r$, these two experts are playing a zero-sum normal form game in their choice of advice $a_i$. And since a mixture of zero-sum games is also a zero-sum game, averaging over client responses to give $\bar{u}_{1\omega}((a_1, a_2), e_1)$ and $\bar{u}_{2\omega}((a_1, a_2), e_2)$ still leaves a zero-sum game in the two advice choices $a_1, a_2$. Similarly, mixed strategies in efforts $e_i$ would still leave a zero-sum game in advice $a_i$.

As is well known, while zero-sum normal form games can have multiple equilibria, each player's payoff must be the same in every pure strategy equilibria. Thus any combination of mixtures over such equilibrium pure strategies is also an equilibrium of the game, giving the same payoffs to each player. Since the average payoffs $\bar{u}_{i\omega}(a, e_i)$ are constant within the focal set, an expert therefore gains no advantage from having his strategy $\beta_{i\omega}(a_i|e_i)$ vary by state $\omega$ within the focal set, even when his opponent so varies his strategy $\beta_{j\omega}(a_j|e_j)$.

Since it is strictly costly to acquire more than the focal information by varying his effort $e_{id}$ along any particular dimension $d$, the only reason an expert might do so would be to let his strategy $\beta_{i\omega}(a_i|a_i^-, e_i)$ vary within the focal set $\hat{\pi}_\omega$. Since this gives no advantage, experts choose no more than the focal information. QED.

## Proof of Proposition 2

Each client strategy $\lambda_{i\omega}(r_i|F_{i\omega}(a, r_i^-))$ is constant across the focal set $\hat{\pi}_\omega$, making payoffs $\bar{u}_{i\omega}(a, e_i)$ constant as well. If each expert $j$ that moves after expert $i$ (so $j > i$) does not choose more than the focal information along any dimension, then by the definition of focal information, he must choose the same or less information along that dimension. In this case the focal information set is contained within expert $j$'s information set, with $\hat{\pi}_\omega \subseteq \pi_{j\omega}(e_j)$, and so this other expert's strategy $\gamma_{j\omega}(a_j|a_j^-)$ is constant across the focal information set. Thus in this case the last term in the definition of $U_{i\omega}(a_i|a_i^-, e_i)$ above, $\prod_{j\in E:j>i} \gamma_{j\omega}(a_j|a_j^-)$, is constant across the focal set.

Thus if each later expert $j > i$ chooses no more than the focal information along any dimension, all terms in the definition of $U_{i\omega}(a_i|a_i^-, e_i)$ are constant across the states $\omega \in \hat{\pi}_\omega$ in the focal information set, and so $U_{i\omega}(a_i|a_i^-, e_i)$ is itself constant over this set. Thus the set of advice $a_i$ that maximizes $\bar{U}_{i\omega}(a_i|a_i^-, e_i)$ is constant over the focal information set. But since it is strictly costly to acquire more than the focal information along any dimension, the only reason for an expert to do so would be to allow his strategy $\beta_{i\omega}(a_i|a_i^-, e_i)$ to vary within this set. Thus if each later expert $j > i$ chooses no more than the focal information, the possible strategy for expert $i$ of acquiring more than the focal information is dominated by the strategy of acquiring no more than the focal information along any dimension. And so by recursion over experts in reverse order, in equilibrium each expert chooses no more than the focal information. QED.

# References


Brandenburger, A., and B. Polak (1996): "When managers cover their posteriors: making the decisions the market wants to see," *RAND Journal of Economics*, 27(3), 523–541.

Cummins, J. G., and I. Nyman (2004): "The dark side of competitive pressure," *RAND Journal of Economics*, forthcoming.

Heidhues, P., and J. Lagerlof (2003): "Hiding information in electoral competition," *Games and Economic Behavior*, 42, 48–74.

Krishna, V., and J. Morgan (2001): "A Model of Expertise," *Quarterly Journal of Economics*, 116(2), 747–777.

Morris, S. (2001): "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.

Ottaviani, M., and P. N. Sorensen (2003): "Professional Advice: The Theory of Reputational Cheap Talk," london.edu/faculty/mottaviani/PA.pdf.

Smart, M., and D. Sturm (2003): "Term Limits and Electoral Accountability," .