The social value of fake experts[*]

Francisco Silva[†]

May 29, 2020

Abstract

Fake experts have no intrinsic ability or knowledge; they only know what they are told. Oftentimes, fake experts pretend to be real experts for personal gain, thus generating uncertainty over which experts are real and which experts are fake. I argue that i) fake experts are able to sustain a permanent reputation of being real even though all agents are assumed to have rational expectations and ii) uncertainty over whether each expert is fake or real might make the agents better off.

JEL classification: D82, C73

Keywords: imperfect monitoring, dynamic cheap talk, experts

[*]I would like to thank Martin Besfamille, Mustafa Dogan, Selman Erol, Nicolas Figueroa, Ju Hu, Jeanne LaFortune, George Mailath, Samir Mamadehussene, Joaquin Poblete, Francisco Ruiz-Aliseda, Bernardita Vial and Rakesh Vohra, as well as the seminar participants at PUC Chile, U. Diego Portales, U. de Chile, the Stony Brook's game theory conference and the North American Econometric Society meeting for their useful comments.

[†]Pontificia Universidad Catolica de Chile, Department of Economics, Vicuña Mackenna 4860, Piso 3. Macul, Santiago, Chile. (email: franciscosilva@uc.cl).

1 Introduction

The importance of experts in our world seems to be greater than ever. The technological progress mankind has gone through makes information incredibly valuable and normal people, who want access to that information, seek experts for advice on the more diverse matters (financial, political, religious, etc.). The power that some experts hold is naturally appealing to non-experts, who are sometimes tempted to claim to being experts. A lot of times, these "fake" experts are hard to spot by normal people, because the very fact that they are not experts prevents them from distinguishing fake from real experts.

Despite this, several of these charlatans have been uncovered throughout History even dating back to Ancient times. In Ancient Greece, it was common for people to consult oracles. The priests and priestesses of these oracles claimed to be able to talk with the Gods, who would tell them about the future. In the oracle of Delphi, the most successful of these oracles, God Apollo was supposed to induce the priestess into a transe that would lead her to make accurate predictions about the future. Modern historians seem more convinced that the transe was generated by hallucinogenic gases rather than divine intervention. More recently, there have been several documented high profile cases of fake financial experts; people who falsely claimed to have some intrinsic knowledge or ability to predict how financial markets would evolve in order to attract investment.

There is a general consensus that the existence of fake experts is bad. Fake experts lie to attract clients; they claim to know things they do not know, which induces those who consult them to make mistakes. They are especially dangerous, it is said, because they prey on those who are less sophisticated; those who are less equipped to identify a charlatan. If one scrutinizes this line of reasoning further, one realizes that it is not the fact that fake experts exist per se that is seen as negative; it is the fact that there is uncertainty over whether an expert is fake or not. If everyone could distinguish fake experts from real experts, all would be well, as presumably fake experts would not be consulted. In this paper, I argue against this consensus and suggest that uncertainty over whether an expert is fake might actually be a good thing for everyone (or almost everyone) involved - the expert herself and those who could potentially consult her.

My initial point of departure is the realization that, if indeed fake experts were generally harmful and gave bad advice to those who consult them, they would not last. It does not seem realistic to think that people would systematically get bad advice and never realize that they are being tricked. And yet we know that some fake experts have persisted throughout decades and even centuries, as is the case of the oracles of Ancient Greece. Therefore, if at least some fake experts do not disappear, it must be because they are providing good advice; advice that is good enough to support their claim of expertise.

At first glance, this seems impossible. Take the aforementioned oracle of Delphi for example. The priests argued that their predictions came directly from the Gods, something we now know (or at least strongly suspect) is false. They must have given good advice to the people who consulted them; if not, they would eventually stop being consulted. In fact, if the belief that the oracles did receive advice from the Gods was to be sustained over time, which, apparently, it was, their advice would have to be as good as the advice they would have given if they were actually in contact with the Gods. Some historians argue that this was possible because the priests were well connected (see, for example, Scott, 2014). People would travel large distances to visit the oracle of Delphi; it is said that very influential people, including the rulers of different lands, would visit the oracle seeking advice, and so it seems natural that the priests who ran the oracle would be people who would be better informed than most, and, as a result, would be in a position to actually give good advice. So, if one accepts this argument, while the priests of the oracles were not true experts, in the sense that they did not have any intrinsic knowledge that ordinary people did not have (i.e., they did not actually talk with the Gods), they did give good advice as a result of how well connected they were.

My argument builds on this idea. The way that a fake expert can look like a true expert is by being well connected. Her connections enable her to generate a virtuous circle: The expert gets people to consult her; those people make her well connected; those connections allow her to make good predictions, which cement her reputation of expertise; the enhanced reputation attracts more people, which makes the expert even more connected and so on.

Of course, by now the reader must be wondering the following: if, in order to survive in the long-run, fake experts must find a way to give as good advice as "real" experts, why lie? Why not just admit that the reason that they provide good advice is because they are well connected? Why did the oracles of Ancient Greece feel the need to say they received information from the Gods instead of admitting that the reason they gave good advice was because they were well connected? The reason is that there is an inherent fragility to those experts whose (main) source of expertise is how well connected they are, in that they are only as good as their connections. Imagine that, at some point in the virtuous circle described above, a group of people feel like the information the expert has from his connections is not relevant enough for them to go through the trouble of consulting her. The fact that these people choose not to consult the expert not only harms the expert directly but also makes her less connected (as compared to the event in which they would have consulted her), which makes her less able to give good advice going forward. By contrast, when agents believe that the expert's advice comes from some other intrinsic source that does not depend on the (number of) consumers she has, they are more likely to consult her, causing the virtuous circle to remain intact.

The argument for why uncertainty (over whether an expert is fake or not) might be good for (almost) everyone is that, without it, the fake expert faces more risk of losing her connections and, consequently, her ability to provide good advice. By convincing people that the reason she gives good advice is some intrinsic skill or ability, she mitigates that risk and is more likely to be able to maintain a virtuous circle of connections, which allow her to support her reputation of expertise.

Finally, while oracles have served as an inspiration for the argument, the argument itself - that fake experts might be welfare increasing because of how well connected they become - is reminiscent of ideas that have been discussed before in Economics and Finance. For example, Bodnaruk and Simonov (2015) find that some financial "experts" seem to do well only because of the private information they are able to obtain from their past and current clients, even though they often suggest otherwise. Another example is Bertrand, Bombardini and Trebbi (2014) who have found evidence that the value of political lobbyists is not their technical knowledge as they often claim, but their political connections.

The paper proceeds as follows. In section 2, I introduce the model; in section 3, I explain why an expert who simply gives good advice because of how well connected she is might face problems providing good advice in the long-run; in section 4, I show how those problems are mitigated when there is uncertainty with respect to whether the expert has some intrinsic knowledge or not. In section 5, I discuss the main results of the paper and the related literature.

2 Model

There is a long-run expert (she) who faces an infinite sequence of short-lived agents. In each period $t \geq 1$, agent t (he) must take an action $a_t \in \{L, R\}$. He receives a payoff of $u_t = 1$ if his action matches the state of the world $\omega_t \in \{L, R\}$, which is random and unobservable. Otherwise, he receives a payoff of $u_t = 0$.

Before deciding a_t, each agent t decides whether to consult the expert. If agent t consults the expert, then $b_t = 1$; otherwise, $b_t = 0$. Consulting the expert has a cost of $c > 0$, which represents the cost of travelling to where the expert is, the opportunity cost of the time that is spent, etc. The goal of the expert is simply to maximize the discounted sum of agents who consult her, i.e., the expert maximizes $E\left[\sum_{t=1}^{\infty} \delta^t b_t\right]$, for some $\delta \in (0, 1)$.

Each agent decides whether to consult the expert after having observed a private signal $s_t \in [0, 1]$ and after having observed which of the previous agents have consulted the expert and, if they did, whether they were successful. Formally, for each $t \geq 2$, let $h^t = \{b_\tau, b_\tau u_\tau\}_{\tau=1}^{t-1}$ denote the public history at the beginning of period t. Before choosing b_t, each agent t is assumed to observe h^t and s_t. Notice that allowing agent t to observe h^t makes it harder for a fake expert to build a reputation of expertise, because it gives each agent the ability to (indirectly) verify how good the expert's predictions were; if the agents who have consulted the expert keep making mistakes, that must mean that the expert's predicting power is not very high.

If agent t does consult the expert, their interaction is as follows. First, the agent sends the expert a message $m_t \in [0, 1]$, which the reader may interpret as revealing (truthfully or not) the agent's private signal s_t. Upon receiving message m_t, the expert gives the agent a recommendation $r_t \in \{L, R\}$, which the agent is free to follow or not. In addition to the information provided by the agents, the expert observes a private signal $\eta_t \in \{L, R\}$ in every period $t \geq 1$. Figure 1 displays the timing of the events within each period $t \geq 1$.

At period 0, the expert observes privately two additional random variables. On the one hand, she observes $s_0 \in [0, 1]$, a private signal possibly correlated with the following T future states of the world, i.e., s_0 contains only short-term information. On the other hand, she also observes $\theta \in \Theta$, which can be interpreted as her type. When discussing the model and the formal results, I avoid using the terms "real" and "fake" to avoid possible confusion with which type is realized. Instead, I assume that $\Theta = \{\theta_N\} \cup \Theta_S$, where type $\theta = \theta_N$ represents the *natural* type of the expert (instead of fake), while each type $\theta \in \Theta_S$ represents a *supernatural* type. The expert's type θ determines the distribution of all other random variables.

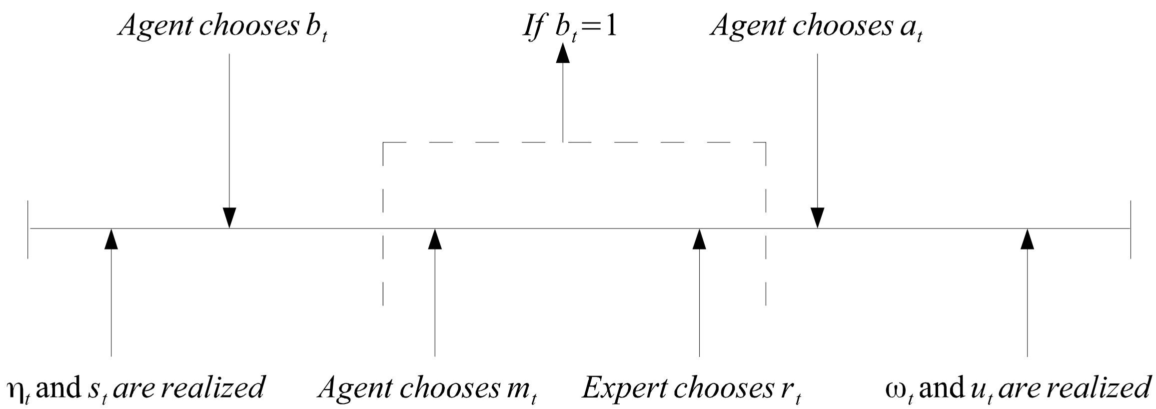

Figure 1: Timing within any period $t \geq 1$

The natural world: If $\theta = \theta_N$, I also say that the world is the natural world. In the natural world, the expert's private signals $\{\eta_t\}_{t=1}^{\infty}$ are completely irrelevant, i.e., they are independent of $\{\omega_t\}_{t=1}^{\infty}$ and of $\{s_t\}_{t=0}^{\infty}$. Therefore, in the natural world, the expert is just like an agent; in fact, it is as if she is agent $t = 0$, because she observes s_0, which, by assumption, is only correlated with the first few states of the world. She has no special intrinsic skills or knowledge, she does not talk with the Gods; she is just an ordinary person who receives people and gives advice. Without loss of generality, whenever $\theta = \theta_N$, I assume that $s_t \sim^{iid} U(0,1)$. Furthermore, I assume that, for each $t \geq 1$, ω_t is only correlated with past signals; the primitive of the model is $\Pr\{\omega_t | s_t, ..., s_0\}$.[1] The probability that $\theta = \theta_N$ is denoted by $\mu \in [0,1]$.

The supernatural worlds: If $\theta \in \Theta_S$, I say that the world is supernatural. Whenever $\theta \in \Theta_S$, signals $\{s_t\}_{t=0}^{\infty}$ are completely irrelevant; they are independent of $\{\omega_t\}_{t=1}^{\infty}$ and of $\{\eta_t\}_{t=1}^{\infty}$. Instead, the only signals that are correlated with the state of the world are the $\{\eta_t\}_{t=1}^{\infty}$, the private signals of the expert. So, if the world is a supernatural world, the expert really does have some intrinsic knowledge that ordinary people do not have. It might that these signals come from the Gods, or they might simply represent a superior cognitive ability or a special understanding of how the world works.

For simplicity, I assume that for each $\theta \in \Theta_S$,

$$\Pr\{\eta_t = L | \theta\} = \frac{1}{2},$$

[1] Notice that the assumption that $s_t \sim^{iid} U(0,1)$ is without loss of generality because there are no restrictions with respect to what $\Pr\{\omega_t | s_t, ..., s_0\}$ might be.

while

$$\Pr\{\omega_t = \eta_t | \theta, \eta_t\} = \frac{1}{2} + \lambda_t^\theta,$$

where, for each $t \geq 1$, $\lambda_t^\theta \in \left[0, \frac{1}{2}\right]$. Notice that this implies that

$$\Pr\{\omega_t = L | \theta\} = \frac{1}{2},$$

so that, if agent t knows $\theta \in \Theta_S$ but does not know η_t, there is only a $\frac{1}{2}$ chance that he makes the correct decision when choosing his action a_t. However, if the agent also knows η_t, that probability goes up to $\frac{1}{2} + \lambda_t^\theta$. So, λ_t^θ represents the marginal benefit of consulting the expert in the supernatural world $\theta \in \Theta_S$, assuming that the expert will be forthcoming in her recommendation.

To sum up, in period t, if consulted by agent t, the expert will have observed all previous messages sent by the previous consulting agents (m_τ for all $\tau \leq t$ for which $b_\tau = 1$), her own private signals $\{\eta_\tau\}_{t=1}^{t}$, s_0 and θ.

3 No type-uncertainty

Let us first consider the case where the expert's type is commonly known. If the expert's type is a supernatural type $\theta \in \Theta_S$, the players' decisions are relatively straightforward, because the time periods are not linked. The expert has nothing to gain by misleading the agents, so she might as well report truthfully, i.e., report $r_t = \eta_t$ for any period t (and any possible history of play). In that way, by consulting the expert and then following her advice, the probability that each agent t is able to match the state is of $\frac{1}{2} + \lambda_t^\theta$. Given that this probability is only of $\frac{1}{2}$ if agent t chooses not to consult the expert, it follows that the expert is consulted if and only if $\lambda_t^\theta \geq c$. Experts provide a valuable service in that they provide advice to agents, which helps them make good decisions.

Things are more interesting when the expert's type is the natural type. While the natural expert does not have any intrinsic ability or knowledge, she is able to do something that agents cannot; she is able to interact with previous agents. And that is valuable. Imagine an agent at period $t = 10$. Should he consult the expert? He knows the expert's signal η_t is irrelevant, so he knows the expert does not "talk with the Gods". But, he also knows that the expert is able to talk with the agent of period

$t = 9$. And maybe the interaction between the two led the expert to learn something about the signal that agent $t = 9$ held - s_9 - which is informative about the state of the world in period $t = 10$. In that sense, the natural expert does not appear to be that different from the supernatural expert in that her advice is still valuable. Let me introduce some notation in order to further compare the two types of expert.

Consider the following (partial) strategy profile:[2]

Expert: Whenever consulted, the expert gives as good advice as she is capable of, i.e., the expert recommends $r_t = L$ if and only if the probability that $\omega_t = L$ given what she knows at the time is at least 50%.

Agents: Provided each previous agent has consulted the expert, for all $t \geq 1$, agent t consults the expert ($b_t = 1$), truthfully reveals his private signal s_t ($m_t = s_t$ for all $s_t \in [0, 1]$) and follows the advice of the expert ($a_t = r_t$).

Let H^t be the set of public histories h^t such that $b_\tau = 1$ for all $\tau < t$ for each period $t \geq 1$. Any $h^t \notin H^t$ is a public history that is off the path of play, according to the strategy profile described. For all $h^t \in H^t$, let $v_t^{h^t}(s_t) \in \left[\frac{1}{2}, 1\right]$ denote the probability that agent t is able to match the state of the world ω_t when his private signal is s_t, the public history is h^t and all players follow the strategy profile defined above. Finally, let

$$\beta_t^{h^t} \equiv \int_0^1 v_t^{h^t}(s_t)\, ds_t.$$

Notice that $\beta_t^{h^t}$ is the analogous to $\left(\frac{1}{2} + \lambda_t^\theta\right)$ in the supernatural world $\theta \in \Theta_S$; it represents the average probability that an agent at time t and history h^t matches the state after consulting the expert if all players play as specified. So, it would appear as though the nature of the expert does not matter per se; only her ability to make predictions, which is given by $\beta_t^{h^t}$ if she is of the natural type and by $\left(\frac{1}{2} + \lambda_t^\theta\right)$ if she is of the supernatural type. That turns out not to be the case, as the type of the expert plays a crucial role in the success of the expert as the following example illustrates.

3.1 An example

Suppose that

$$\Pr\{\omega_t = L | s_t, s_{t-1}\} = \begin{cases} s_t \text{ if } s_{t-1} \geq \frac{1}{2} \\ 1 - s_t \text{ if } s_{t-1} < \frac{1}{2} \end{cases}$$

[2] It is only partial, because I do not specify what happens "off-the-path of play".

for all $(s_t, s_{t-1}) \in [0,1]^2$ and for any $t \geq 1$. Therefore, in the natural world, the current state of the world is only correlated with the current signal and the one of the previous period. Fix the strategy profile described above: the expert always gives as good advice as possible when consulted and the agents always consult the expert, provided every preceding agent did so, report truthfully and follow the expert's advice.

Notice that

$$v_t^{h^t}(s_t) = \max\{s_t, 1 - s_t\}$$

for all $h^t \in H^t$, $s_t \in [0,1]$ and t, because agent t will gain access to s_{t-1} by consulting the expert.[3] As a result, it follows that

$$\beta_t^{h^t} \equiv \int_0^1 \max\{s_t, 1 - s_t\}\, ds_t = \frac{3}{4},$$

for all $h^t \in H^t$ and t, i.e., in this example, the natural expert is able to correctly predict each state of the world 75% of the times. For argument's sake, let us say that $\Theta_S = \{\theta_s\}$, where $\lambda_t^{\theta_s} = \frac{1}{4}$ for all t, so that the supernatural expert is also able to predict each state of the world with the same 75% probability. Despite this, the fate of both types of expert will be very different.

If the expert is supernatural (and that is publicly known), each agent will consult the expert provided $c < \frac{1}{4}$. Therefore, as long as c is sufficiently small, the supernatural expert will be consulted forever, which will make every agent better off (as compared to not having access to an expert).

However, in the case of the natural expert, there will be agents with signals s_t close to $\frac{1}{2}$ who will choose not to consult her even if every previous agent has, because

$$\lim_{s_t \to \frac{1}{2}} v_t^{h^t}(s_t) = \frac{1}{2}$$

for any $h^t \in H^t$ and t. What is even more disturbing is that once some agent t decides not to consult the expert, no other agent will ever consult the expert again, because the expert's "power" - her connections - will be gone and everyone will know about it. Therefore, under the same assumption that $c < \frac{1}{4}$, a natural expert would face more difficulties than the analogous supernatural expert and, in this example, would eventually disappear (with probability 1), even though every agent t would be better off from her existence; if no expert existed, each agent would only have a 50% chance

[3]This also applies to agent $t = 1$ because the expert is assumed to hold s_0. In general though, $v_t^{h^t}(s_t)$ depends on the history h^t.

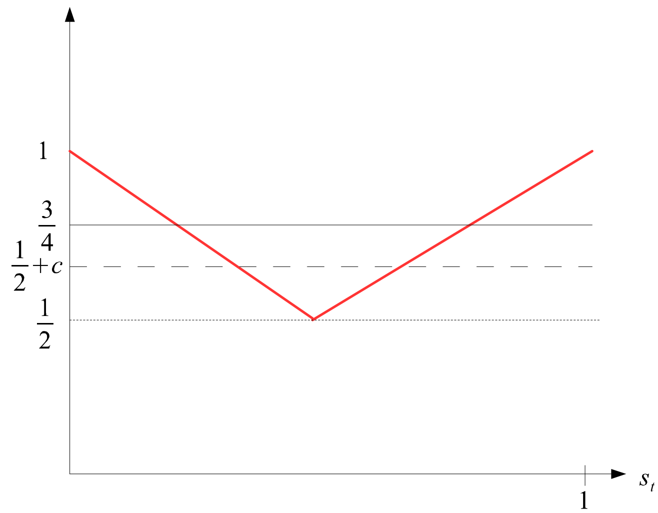

Figure 2: The red line represents function $v_t^{h^t}(s_t) = \max\{s_t, 1 - s_t\}$.

of matching the state, so the existence of the natural expert would make agents better off as long as $c < \frac{1}{4}$.

Figure 2 illustrates the challenges of being a natural expert. Recall that the benefit for any given agent t of consulting the expert when everyone else before him did is given by $v_t^{h^t}(s_t)$, which is represented by the red line. By contrast, the opportunity cost of the same agent t is given by $\frac{1}{2} + c$, represented by the dotted line.[4] The problem that the natural expert faces is with the agents who receive signals that are too close to $\frac{1}{2}$, because, whenever that happens, the red line is below the dotted line. By contrast, the benefit for agent t of consulting the supernatural expert is independent of his signal s_t; it is simply $\frac{3}{4}$, represented in the full line. The supernatural expert does not have the same problem of the natural expert because the benefit of consulting her is "flatter", i.e., even though the red line has an "average" of $\frac{3}{4}$, the fact that it tilts down causes the breakdown.

3.2 The fragility of being a natural expert

Let us formalize the phenomenon of the previous example. I make two assumptions.

[4]It follows that that the probability that agent t matches the state should he not consult the expert is equal to $\frac{1}{2}$ for any $h^t \in H^t$ and t, because the belief that agent t has of s_{t-1} is symmetric around $\frac{1}{2}$ for any $h^t \in H^t$.

Assumption A1: *There is $\varepsilon \in (0, c)$ and $\tau > 0$ such that, for all $t \geq 1$, there is an $I_t \in [0, 1]$ such that*

$$\int_{s_t \in I_t} s_t ds_t > \tau$$

and, for all $s_t \in I_t$,

$$\max \left\{ \Pr \left\{ \omega_t = L | s_t, s_{t-1}, ..., s_0 \right\}, 1 - \Pr \left\{ \omega_t = L | s_t, s_{t-1}, ..., s_0 \right\} \right\} < \frac{1}{2} + \varepsilon$$

for all $(s_{t-1}, ..., s_0) \in [0, 1]^t$.

Assumption A2: There is $T > 0$ such that each state ω_t is only correlated with the last T signals, i.e., for all $t \geq T$,

$$\Pr \left\{ \omega_t | s_t, s_{t-1}, ...s_{t-T}, s_{t-T-1}..., s_0 \right\} \text{ is independent of } \left(s_{t-T-1}..., s_0 \right).$$

Assumption A1 means that, in every period t, there could be a signal s_t such that agent t cares very little about the previous signals, i.e., for some signals s_t, knowing the previous signals only marginally improves the odds of agent t making a good decision. It implies that such an agent would not be willing to incur the cost c in order to consult the expert no matter what. In the example, this holds for signals that are sufficiently close to $\frac{1}{2}$.

Assumption A2 implies that each agent only cares about the last T signals. The logic is that signals that are too far into the past are not correlated with the current state and, thus, should be disregarded. In the example, $T = 1$.

Proposition 1 *Assume A1 and A2 hold and that θ is common knowledge.*

i) For all $\theta \in \Theta_S$ such that $\lambda_t^\theta \geq c$ for all t, there is a perfect Bayesian equilibrium (PBE) outcome where every agent consults the expert and follows her truthful advice.

ii) If $\theta = \theta_N$, then $b_t \rightarrow^{a.s.} 0$ in any PBE.

Proof. Part i: To construct the PBE, one can simply complete the description of the strategy profile described above with off-the-path behavior. For the expert, if $h^t \notin H^t$ and she is consulted, she always selects $r_t = L$ and $r_t = R$ with the same 50% probability. For the agents, if $h^t \notin H^t$, agent t does not consult the expert and, whenever consulting the expert, always reports $m_t = 0$ for all t. When deciding his action, agent t always chooses his preferred action given his current beliefs.

It follows trivially that the expert never wants to deviate: on the path of play, she is always consulted, while off the path, there is nothing she can do to be consulted again. As for each agent on the path of play, following the expert's advice is preferred because the expert's advice is valuable and consulting the expert is better than not to because $\lambda_t^\theta \geq c$. Finally, each agent is indifferent as to what to report (because each signal s_t is irrelevant in any supernatural world), so reporting $m_t = s_t$ is trivially optimal. Off the path of play, agents do not want to consult because the expert "babbles" and $c > 0$.

Part ii: Assumption A2 implies that if there is a sequence of T agents who choose not to consult the expert, then no one else will ever consult the expert, because they will know that the expert has no relevant information to disclose. Assumption A1 implies that, no matter the history, there is always a positive probability that such a sequence occurs. ■

Proposition 1 summarizes the fragility of the natural expert. Even when the probability of matching the state is the same as a supernatural expert, she is still at the mercy of a few agents deciding not to consult her because they think that obtaining past information is not worth the cost of visiting her. Whenever that happens, she becomes less connected and, thus, less appealing to future visitors, who then stop visiting, making her even less appealing and so on until she runs out of visitors. While there are ways out this vicious circle, none of them is particularly appealing to the natural expert. For example, whenever she runs out of visitors, she could start paying them in order to rebuild lost connections. But not only does that not seem very realistic, it might end up being very expensive for the natural expert. By contrast, a supernatural expert who has the same average probability of matching the state is not subjected to the same challenges and is guaranteed to have an endless supply of visitors.

Proposition 1 also implies that it might be better for the agents if they are tricked into thinking that the expert's type is supernatural when, in reality, it is not, because that would ensure that the natural expert would never disappear and would continue to provide valuable advice. In the next section, I show how that is possible even though agents have rational expectations.

4 The benefits of type uncertainty

The previous section explains why a natural expert might want to be seen as a supernatural expert. However, the challenge that she faces is that, not only do agents have rational expectations; they also constantly scrutinize the expert by observing how successful her past visitors were. Nevertheless, I show that under some conditions, it is possible for the natural expert to sustain the belief that she might be of a supernatural type indefinitely, which enables her to receive visitors in perpetuity even when assumptions A1 and A2 hold.

I make an additional assumption to the model:

Assumption A3: For all $\theta \in \Theta_S$ and for all $t \geq 1$, $\frac{1}{2} + \lambda_t^\theta \geq \beta_t^{h^t}$ for all $h^t \in H^t$.

Assumption A3 simply states that the average probability of matching the state for any supernatural expert is never smaller than that of the natural expert. Presumably, supernatural types for which the assumption does not hold would not be appealing for the natural expert to mimic. In the example, this simply requires $\lambda_t^\theta \geq \frac{1}{4}$ for all t and $\theta \in \Theta_S$.

Proposition 2 *If A3 holds and there is some $\gamma > c$ such that $\beta_t^{h^t} \geq \frac{1}{2} + \gamma$ for all $h^t \in H^t$ and t, there is some $\overline{\mu} \in (0,1)$ such that, whenever $\mu < \overline{\mu}$, there is a PBE where i) every agent consults the expert in every period, for any public history and private signal, ii) for any public history, the posterior distribution of θ is equal to the prior distribution of θ, i.e., $\Pr\{\theta = \theta_N | h^t\} = \mu$ for all h^t and for all $t \geq 1$.*

Proof. See appendix. ■

Let us look more closely at what proposition 2 states. On the one hand, it is assumed that $\lambda_t^\theta > c$ for all t. According to proposition 1, this implies that any supernatural expert would be able to attract every single visitor in every period for every state. Therefore, it guarantees that the natural expert has a justifiable desire to mimic any of the supernatural experts, especially when assumptions A1 and A2 hold. On the other hand, it is also assumed that $\beta_t^{h^t}$, the average probability that the natural expert matches the state if all agents visit her, is sufficiently large for all $h^t \in H^t$ and t, which is a direct consequence of rational expectations: the natural expert can only sustain a reputation of being a supernatural expert if her predictions are good enough.

Under the conditions stated, the prior belief that the agent's type is the natural type is never updated. Therefore, if it is sufficiently small to begin with, agents will be convinced the expert is likely to be one of the supernatural types and, as a result, will want to consult her no matter what else happens. Naturally, for this to happen, it must be that the natural type's *observed* behavior is indistinguishable to the agents from any of the supernatural types. The key word is precisely *observed*, because it is impossible for the natural type to always give the same exact recommendations as the various supernatural types; if she was capable of that, she would be a supernatural type. Nevertheless, agents do not observe past recommendations; only whether past visitors made good decisions or not (i.e., whether they were able to match the state). It turns out that it is that subtle difference that leads to the result. While the proof of the statement is in the appendix, I revisit the example to explain the intuition of proposition 2.

4.1 Example (revisited)

Consider the example from before but assume that agents do now know the expert's type. Notice that the conditions of proposition 2 are that $\lambda_t^{\theta} \geq \frac{1}{4} > c$ for all $\theta \in \Theta_S$ and for all $t \geq 1$. For now, let me continue to assume that $\Theta_S = \{\theta_s\}$ and that $\lambda_t^{\theta_s} = \frac{1}{4}$ for all $t \geq 1$, so that the expert's average probability of matching the state is the same regardless of her type, provided every preceding agent has visited her and has reported truthfully. Consider the same strategy profile from before: as long as every agent consults the expert, the expert gives as good a recommendation as she can, while each agent consults the expert, reports truthfully and follows the expert's advice.

Take some agent t on the path of play and let $\pi\left(h^t\right)$ denote the posterior probability that $\theta = \theta_N$ given public history $h^t \in H^t$. It follows that he consults the expert if and only if

$$\left(1 - \pi\left(h^t\right)\right)\frac{1}{4} + \pi\left(h^t\right)\left(v_t^{h^t}\left(s_t\right) - \frac{1}{2}\right) \geq c.$$

Given that $c < \frac{1}{4}$, it is straightforward that agent t consults the expert for any $s_t \in [0,1]$ provided $\pi\left(h^t\right)$ is sufficiently small.

What proposition 2 assumes is that the prior belief μ is small enough that agent 1 prefers to consult the expert for any private signal. As long as that happens, the belief will stay constant, i.e., $\pi\left(h^t\right) = \mu$ for all $h^t \in H^t$. To see why that is, suppose that agent t's belief about θ is given by $\pi\left(h^t\right) = \mu$. The probability that he matches the state, provided he consults the expert and follows her advice, depends on the expert's type θ. If $\theta = \theta_N$, that probability is given by $\max\{s_t, 1 - s_t\}$; if $\theta = \theta_s$, it is $\frac{3}{4}$. When agent $t+1$ looks back at what happened at period t, he does not observe s_t; he only observes whether agent t was able to match the state. Therefore, for agent $t+1$, the probability that agent t matches the state if $\theta = \theta_N$ is equal to

$$\int_0^1 \max\{s_t, 1-s_t\}\, ds_t = \beta_t^{h^t} = \frac{3}{4},$$

i.e., it is as likely for the agent to match the state in the natural world as it is in the supernatural world. As a result, while agent $t+1$ observes one more piece of evidence about the expert than agent t, he does not update his beliefs. Figure 3 illustrates.

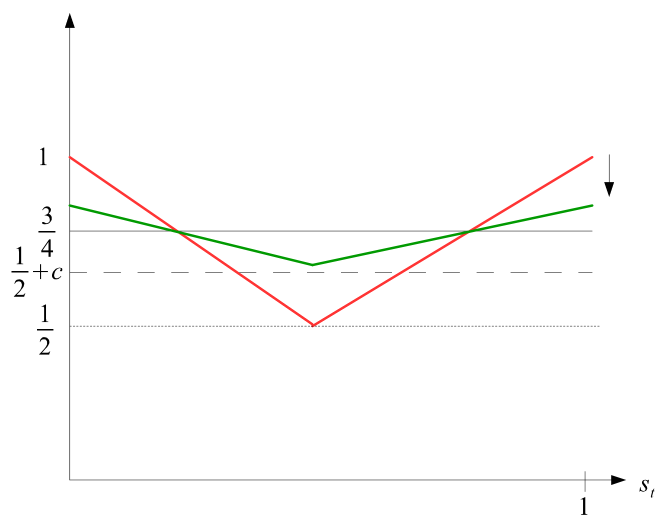

Figure 3: The green line represents a weighted average between the red line and $\frac{3}{4}$.

Recall from Figure 2 that the red line represents the benefit of consulting the natural expert on the path of play. The problem of the natural expert is that if s_t is close to $\frac{1}{2}$, the red line is below $\frac{1}{2}+c$. Uncertainty over the expert's type "flattens" the red line and transforms it into the green line. Essentially, with uncertainty, agents are unsure about whether they are facing the red line or the flat full line at $\frac{3}{4}$. The natural expert benefits from this because the green line is always above $\frac{1}{2}+c$, so she is able to attract every agent in every period for any signal.

Finally, notice that it is straightforward to extend the argument to the case where $\lambda_t^{\theta_s} > \frac{1}{4}$. In that case, the expert's strategy must be adjusted when $\theta = \theta_s$; on the path of play, the expert must randomize between giving as good advice as possible and giving bad advice. In that way, the supernatural expert is able to lower the probability that the agent matches the state to $\frac{3}{4}$, which is something that the expert of the natural world can also generate. Notice that neither type of expert wants to deviate, because following this strategy ensures that every agent consult the expert in every period, provided μ is sufficiently small. The same argument applies if there are several supernatural types, each with $\lambda_t^{\theta_s} \geq \frac{1}{4}$; in equilibrium, each of these supernatural types will randomize in such a way that the probability of matching the state if always equal to $\frac{3}{4}$.[5]

5 Discussion

5.1 Welfare

The welfare comparison is particularly striking when assumptions A1 and A2 hold. In that case, the natural expert will struggle to survive. Proposition 1 states that the fate of the natural expert is to disappear even if c is arbitrarily small. Not only would this be bad for the expert herself, it would be disastrous for the agents, who stop having access to valuable advice at a minimum cost. Even if the expert finds ways to mitigate the effects of proposition 1 by sometimes paying agents to consult her, she would still inevitably have less people visiting her as she would have had were she a supernatural expert with the same average predicting ability. This would reduce her predicting ability going forward, which would, in turn, harm her future visitors.

To avoid all of this, she can just pretend to be a supernatural expert of similar average predicting ability, as long as that average predicting ability is sufficiently high and the public belief in the existence of supernatural types is sufficiently large. This not only makes her better off, as she will have more visitors, it is also likely to make (almost) every agent better off. To see why that is, let us consider the problem of some future agent t. If the expert is known to be the natural expert, the expert's predicting ability will be close to $\frac{1}{2}$ by the time agent t comes along. Therefore, agent t will have to rely solely on his signal s_t in order to make his decision. Provided that $\Pr\{\omega_t = L | s_t\} = \frac{1}{2}$ for all t, an innocuous assumption that is completely compatible with all assumptions made before and holds in the example, his expected utility will be as low as it can be. By contrast, by pretending to be a supernatural expert, the natural expert's predicting ability will still be high at period t; in fact, it must be larger

[5]By slightly altering the model, it is possible to generate the same result that the natural expert is able to generate a permanent reputation of being a supernatural expert without requiring the supernatural experts to randomize between good and bad advice. I do this in section 5.2.

than $\frac{1}{2}+c$ for proposition 2 to hold. So, agent t's expected utility is guaranteed to be larger than $\frac{1}{2}$, i.e., uncertainty about the expert's type makes him better off. The only agents who are harmed by uncertainty are the first (finite) few.

The only downside of type-uncertainty is that it might make agents worse off in the event that the expert is indeed of a supernatural type for the histories $h^t \in H^t$ such that $\lambda_t^{\theta_s} > \beta_t^{h^t}$, because, should that happen, the supernatural expert would not be giving as good advice as she would have should her type be known. In the next section, I show how to slightly modify the model to ensure that this does not happen, so that agents are not harmed by type uncertainty even when the realized type is supernatural.

5.2 A different modelling assumption

The argument of proposition 2 may rely on supernatural types randomizing between giving good and bad advice; in the example of section 4.1., a supernatural type such that $\lambda_t^{\theta} > \frac{1}{4}$ must lower her ability to make predictions in order to make herself indistinguishable from the natural expert. That randomizing is not required, however, if one slightly alters the model.

Notice that randomization over advice is only necessary for supernatural types who have better average predicting abilities than the natural type; in the example of section 4.1., if $\Theta_S = \{\theta_s\}$, where $\lambda_t^{\theta_s} = \frac{1}{4}$, the PBE of proposition 2 is such that both expert's types report truthfully. Therefore, one could bypass this issue by allowing the expert to announce her type at period 0. In that way, the natural expert would announce her type to be the specific supernatural type who has the exact same average predicting ability than her.

To this effect, there are two changes one must make to the model. First, one must assume that there is some supernatural type $\theta' \in \Theta_S$ such that $\beta_t^{h^t} = \frac{1}{2} + \lambda_t^{\theta'}$ for all $h^t \in H^t$ and t. However, because $\beta_t^{h^t}$ is history dependent in general, this might force $\lambda_t^{\theta'}$ to also be history dependent, i.e., the predicting ability of the supernatural expert might have to depend on the history of visitors and their success level. In the example that is not required though as type θ' is such that $\lambda_t^{\theta'} = \frac{1}{4}$ for all t.

Second, one must assume that, before the game begins, at period 0, the expert is able to send a public cheap talk message $\widehat{\theta} \in \{\theta_N\} \times \Theta_S$, where she announces her type. This message is observed by every agent in perpetuity; it is the reason the expert gives as to why she is able to make predictions.

Under these conditions, and continuing to assume that $\lambda_t^\theta > c$ for all t and for all $\theta \in \Theta_S$, it is possible to find a PBE where every expert's type always gives as good advice as possible (on the path of play) and every agent consults the expert in every period even when assumptions A1 and A2 hold, provided μ is sufficiently small. Specifically, the experts reports $\widehat{\theta}$ as follows: if $\theta \in \{\theta_N, \theta'\}$, then the expert reports $\widehat{\theta} = \theta'$; otherwise, the expert reports $\widehat{\theta} = \theta$. Then, provided every preceding agent has chosen to visit the expert, each type of expert gives as good advice as possible, while each agent consults the expert, shares his signal truthfully and follows the expert's advice.

Basically, the natural expert initially announces which supernatural expert she is mimicking, which guides the agents' expectations and essentially removes the other expert types from consideration. Whenever the agents observe $\widehat{\theta} = \theta'$, we essentially have the model of the text but with $\Theta_S = \{\theta'\}$, in which case the construction of the proof of proposition 2 shows that no expert's type randomizes on her report. Indeed, this is the case I first cover in the example of section 4.1. If the agents observe $\widehat{\theta} \neq \theta'$, they just get to know the expert's type is supernatural type $\widehat{\theta}$. Seeing as that type is known to be a supernatural type, she may as well give as good advice as she is capable of.

5.3 Related Literature

One of the contributions of the paper is that it presents a way of thinking about fake experts that challenges some of the popular views on the subject; some of which have been expressed in the economic literature. In Spiegler (2010) and Szech (2011), the market fails to drive away fake experts (called quacks) because consumers are assumed to be boundedly rational. If consumers had rational expectations, they would realize that quacks provided no service (in their model) and, as a result, would never consult them. In that sense, quacks are bad; they make consumers worse off. By contrast, in my model, fake experts do give good advice, provided they are well connected. By convincing rational agents that they are true experts, fake experts not only help themselves but also (almost all) the agents. As a result, this paper questions the extent to which markets should be regulated, as excessive regulation that succeeds in driving fake experts away might be undesirable.

Berk and van Binsbergen (2017) also argue that regulation that forces fake experts out of the market might be harmful for consumers because of the ensuing changes in the competitive environment. In particular, if there are less fake experts in the market, the prices charged by the true experts are likely to increase not only because the perceived quality of their advice is larger but also because the number of competitors is smaller.

Another related paper is Rudiger and Vigier (2019) who study financial experts who make public recommendations. Like me, they find that fake financial experts can sustain a permanent reputation of expertise. The crucial assumption is that agents can only observe the expert's past recommendations; they are not able to see whether those recommendations were actually good or made the clients better off. In that sense, experts are under more scrutiny in my work, because, while agents cannot observe the expert's past recommendations, they can observe whether the previous visitors of the expert made good decisions or not. In Rudiger and Vigier (2019), a not-so-able expert is able to sustain a permanent reputation of being more able by mimicking the empirical frequency of recommendations that a more able expert would give. Another key difference is that, in contrast to this paper, uncertainty over the expert's type is detrimental for the agents.

On the technical side, the model of the paper fits into the literature on reputation games with imperfect monitoring. Cripps, Mailath and Samuelson (CMS) (2004) consider a repeated game of imperfect monitoring between a long-lived player with two possible types, a normal type and a commitment type, and several short lived players. They show that, under certain conditions, it is impossible for the normal type to build a permanent reputation that he is the commitment type. That result seems to be in contrast with the main result of this paper if one thinks of the commitment type as the supernatural type. While there have been several papers that have followed CMS (2004) who have shown that permanent reputations are possible, they assume that either the type of the long-run player is not permanent (Mailath and Samuelson (2001) and Ekmekci, Gossner and Wilson (2012)), or that the access to past data is either costly (Liu (2011)) or limited (Ekmekci (2011) and Hu (2016)). Instead, in this paper, the fake expert is able to generate a long-run reputation for expertise even though types are assumed to be permanent and access to public data is unrestricted. What causes the difference to CMS (2004) is, on the one hand, the fact that the moral hazard element of CMS (2004) is not present in this model (in CMS (2004), the normal type has a short-run incentive to deviate from the behavior of the commitment type) and the fact that, by design, it is possible for one type to take different actions than another type and yet generate a distribution of signals that is empirically indistinguishable.[6]

[6] In particular, the conditions that are violated from CMS (2004) are that a) the distribution of

Finally, the paper is also related to the literature on Bayesian Persuasion and information design (Kamenica and Gentzkow, 2011 among many others) in that the natural expert uses the flexibility of the agents' rational expectations' assumption and the fact that each agent does not observe previous signals to increase her payoff. The geometric nature of the argument that is depicted in figure 2 is particularly reminiscent of this literature.

6 Appendix

In this appendix, I consider one extension and then provide the proof of proposition 1.

6.1 Extension: small prior

In this section, I address one concern the reader might have; that the prior probability of the natural world, denoted by μ, may have to be too small for a permanent reputation that the world is supernatural to persist. In particular, using the example, I show that, even when μ is large, there might be a PBE where the expert creates a permanent reputation that the world is a supernatural world with positive probability.

Consider the example from before but assume that $\Theta_S = \{\theta_s\}$, where $\lambda_t^{\theta_s} = \frac{1}{2}$ for all $t \geq 1$. Assume also that $c = \frac{1}{8}$ and consider the following strategy profile: each agent t consults the expert, provided the preceding agents also did, sends message $m_t = s_t$ and follows the expert's recommendation; if $\theta = \theta_N$, the expert always gives the best possible advice, while if $\theta = \theta_s$, she recommends $r_t = \eta_t$ with probability $\frac{3}{4}$. As discussed in the text, on the path of play, the posterior belief that $\theta = \theta_N$ will always be equal to μ. The issue is that, for each agent to want to consult the expert, μ has to be small enough. In particular, agent t consults the expert if and only if

$$(1-\mu)\frac{1}{4} + \mu\left(\max\{s_t, 1-s_t\} - \frac{1}{2}\right) \geq \frac{1}{8},$$

so that for agent t to want to consult the expert even when $s_t = \frac{1}{2}$, it must be that $\mu \leq \frac{1}{2}$.

the public signals u_t only depends on the actions chosen by the long-lived player at period t, and b) different actions chosen by the long-lived player at period t lead to different distributions over u_t.

Suppose instead that $\mu = 0.95$, i.e., there is a 95% chance that the world is the natural world. The idea of the argument is to "build up" the belief that $\theta = \theta_s$ by manipulating the strategy of the expert. Consider the following alternative strategy profile: whenever the agent consults the expert, he sends message m_t and follows the expert's advice; if $\theta = \theta_s$ and $t \leq 12$, the expert recommends $r_t = \eta_t$; if $\theta = \theta_s$ and $t > 12$, she recommends $r_t = \eta_t$ with probability $\frac{3}{4}$, if $\theta = \theta_N$, the expert gives the best possible advice.

Under this strategy profile, agent $t = 1$ consults the expert if and only if

$$0.05 * \frac{1}{2} + 0.95 * \left(\max\{s_1, 1 - s_1\} - \frac{1}{2} \right) \geq \frac{1}{8},$$

i.e., whenever

$$s_1 \in [0, 0.394\,7] \cup [0.6053, 1].$$

In period 2, agent 2 updates his beliefs about θ. If agent 1 has not been able to match the state, the belief that $\theta = \theta_s$ will drop to 0, because that should never happen when $\theta = \theta_s$. If, however, agent 1 has been able to match the state, then the belief that $\theta = \theta_s$ will increase because the probability that agent 1 matches the state when $\theta = \theta_N$ is only equal to $0.802\,65 < 1$. In particular, at period 2, if agent 1 has consulted the expert and has been able to match the state, the public belief that $\theta = \theta_s$ goes up to 0.06153. By this logic, it is only a matter of time until the belief that $\theta = \theta_s$, conditional on all preceding agents having consulted the expert and having been able to match the respective states reaches $\frac{1}{2}$. In particular, that happens at period 13. After period 13, the expert of the supernatural world starts randomizing so that the probability that the agents match the state is the same in both worlds going forward. It is easy to see that such a profile is a PBE and that there is a positive, albeit small, probability that the first 12 agents all consult the expert and manage to match each state. Whenever that happens, the belief over θ will stay constant forever after period 13 and all future agents will consult the expert.

6.2 Proof of Proposition 1

Consider the following strategy profile for the agents.

Agent $t = 1$: The agent consults the expert for any signal s_1, always reports truthfully ($m_1 = s_1$ for all $s_1 \in [0, 1]$) and always follows the advice of the expert ($a_1 = r_1$ for all $r_1 \in \{L, R\}$).

Agent $t > 1$: If every preceding agent has consulted the expert, agent t consults the expert for any signal s_t, always reports truthfully ($m_t = s_t$ for all $s_t \in [0,1]$) and always follows the advice of the expert ($a_t = r_t$ for all $r_t \in \{L, R\}$). Otherwise, agent t does not consult the expert, always reports $m_t = 0$ when consulting the expert and ignores the recommendation of the expert when deciding a_t (i.e., the agent reports what he would find best had he not consulted the expert).

Notice that, given the strategy profile of the agent, set H^t represents the set of public histories on the path of play at period t. Consider the following strategy for the expert.

Expert when $\theta \in \Theta_S$: For all $h^t \notin H^t$, if consulted, the expert always reports $r_t = L$ and $r_t = R$ with equal probability. For all $h^t \in H^t$, if consulted, the expert reports $r_t = \eta_t$ with probability $\rho_t^{h^t}(\theta) \in \left[\frac{1}{2}, 1\right]$, where $\rho_t^{h^t}(\theta)$ is such that

$$\rho_t^{h^t}(\theta)\left(\frac{1}{2} + \lambda_t^{\theta}\right) + \left(1 - \rho_t^{h^t}(\theta)\right)\left(\frac{1}{2} - \lambda_t^{\theta}\right) = \beta_t^{h^t},$$

for all $t \geq 1$ and $\theta \in \Theta_S$. The existence of $\rho_t^{h^t}(\theta) \in \left[\frac{1}{2}, 1\right]$ is guaranteed by Assumption A1.

Expert when $\theta = \theta_N$: For all $h^t \notin H^t$, if consulted, the expert always reports $r_t = L$ and $r_t = R$ with equal probability. For all $h^t \in H^t$, if consulted, the expert reports L if and only if that $\Pr\{\omega_t = L | s_t, ..., s_0\} \geq \frac{1}{2}$.

Notice that the off-the-path behavior of the players represents "babbling", just like in Crawford and Sobel (1986), so that no player wants to deviate. On the path of play, it follows trivially that the expert never wants to deviate on his recommendation, because by not deviating, the expert is able to attract every agent. It is also clear that each agent wants to report truthfully his private information and follow the recommendation of the expert; in that way they get the best available advice.

Let

$$\overline{\mu} \equiv \frac{\gamma - c}{\gamma}.$$

Notice that, for all $\mu < \overline{\mu}$,

$$(1 - \mu)\gamma > c,$$

which implies that

$$(1 - \mu)\left(\beta_t^{h^t} - \frac{1}{2}\right) > c$$

for all $h^t \in H^t$ and t, because

$$\beta_t^{h^t} \geq \frac{1}{2} + \gamma.$$

Let $\pi(h^t)$ denote the posterior probability that $\theta = \theta_N$ given the public history h^t. It follows that after history h^t, agent t consults the expert if

$$\left(1 - \pi\left(h^t\right)\right)\left(\beta_t^{h^t} - \frac{1}{2}\right) \geq c.$$

Therefore, if $\pi(h^t) \leq \mu < \overline{\mu}$, agent t consults the expert. I complete the proof by showing that if $\mu < \overline{\mu}$, then $\pi(h^t) = \mu$ for all $h^t \in H^t$ and for all $t \geq 1$.

Let $\pi(h^1) \equiv \mu$. Take any h^t and assume that $\pi(h^t) = \mu$. Notice that, for each s_t, the probability that agent t matches the state of world ω_t after consulting the expert is given by $v_t^{h^t}(s_t)$ if $\theta = \theta_N$ and by $\beta_t^{h^t}$ if $\theta \in \theta_S$. Seeing as agent $t+1$ does not observe s_t, and seeing as, by definition,

$$\beta_t^{h^t} = \int_0^1 v_t^{h^t}(s_t)\, ds_t,$$

it follows that, from the point of view of agent $t+1$, agent t is just as likely to match the state when $\theta = \theta_N$ then when $\theta \in \Theta_S$. Therefore, it follows that $\pi(h^{t+1}) = \mu$ for all h^{t+1} consistent with h^t.

References

[1] Berk, J.B., & van Binsbergen, J.H. (2017). Regulation of Charlatans in High-Skill Professions. Working paper.

[2] Bertrand, M., Bombardini, M., & Trebbi F. (2014). Is It Whom You Know or What You Know? An Empirical Assessment of the Lobbying Process. The American Economic Review, 104(12), 3885-3920.

[3] Bodnaruk, A., & Simonov, A. (2015). Do financial experts make better investment decisions?. Journal of Financial Intermediation, 24(4), 514-536.

[4] Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. Econometrica: Journal of the Econometric Society, 1431-1451.

[5] Cripps, M. W., Mailath, G. J., & Samuelson, L. (2004). Imperfect monitoring and impermanent reputations. Econometrica, 72(2), 407-432.

[6] Ekmekci, M. (2011). Sustainable reputations with rating systems. Journal of Economic Theory, 146(2), 479-503.

[7] Ekmekci, M., Gossner, O., & Wilson, A. (2012). Impermanent types and permanent reputations. Journal of Economic Theory, 147(1), 162-178.

[8] Foster, D. P., & Vohra, R. V. (1998). Asymptotic calibration. Biometrika, 85(2), 379-390.

[9] Hu, J. (2016). Biased Learning and permanent reputation. Working paper.

[10] Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. The American Economic Review, 101(6), 2590-2615.

[11] Liu, Q. (2011). Information acquisition and reputation dynamics. The Review of Economic Studies, 78(4), 1400-1425.

[12] Mailath, G. J., & Samuelson, L. (2001). Who wants a good reputation?. The Review of Economic Studies, 68(2), 415-441.

[13] Olszewski, W., & Sandroni, A. (2008). Manipulability of Future-Independent Tests. Econometrica, 76(6), 1437-1466.

[14] Olszewski, W., & Sandroni, A. (2009). A nonmanipulable test. The Annals of Statistics, 1013-1039.

[15] Ottaviani, M. & Sorensen, P.N. (2006). Reputational cheap talk. RAND Journal of Economics, 37(1), 155-175.

[16] Rudiger, J. & Vigier, A. (2019). Learning about analysts. Journal of Economic Theory, 180, 304-335.

[17] Sandroni, A. (2003). The reproducible properties of correct forecasts. International Journal of Game Theory, 32(1), 151-159.

[18] Scott, M. (2014). Delphi: a history of the center of the ancient world. Princeton University Press.

[19] Spiegler, R. (2006). The market for quacks. The Review of Economic Studies, 73(4), 1113-1131.

[20] Szech, N. (2011). Becoming a bad doctor. Journal of Economic Behavior & Organization, 80(1), 244-257.