INSTITUTO DE ECONOMÍA

DOCUMENTO de TRABAJO

539
2020

An informational Ponzi-scheme

Francisco Silva.

www.economia.puc.cl • ISSN (edición impresa) 0716-7334 • ISSN (edición electrónica) 0717-7593

An informational Ponzi-scheme*

Francisco Silva†

August 14, 2019

Abstract

I show how "experts" who have no intrinsic ability or knowledge are able to sustain a permanent reputation that they do, even in a world where agents have rational expectations and access to an unlimited amount of data about the expert's predicting ability. The claim of having such knowledge attracts clients to the expert, allowing the expert to have access to the inside information they provide. That information can then be used by the expert to back up that claim.

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

Throughout History, there have always been charlatans who have proclaimed to be experts for personal gain. Typically, these "fake" experts are seen in a bad light, having been accused, among other things, of having caused the financial crisis of the 2000s. They are portrayed, both by the general public and by some of the economic literature, as people who give bad advice to their clients and who benefit from their lack of sophistication (for example, Spiegler (2006) models a market for "quacks" under the assumption that consumers are boundedly rational).

Some of the "schemes" that fake experts use are quite sophisticated. Consider the well known Ponzi scheme. In a Ponzi scheme, the fake expert uses the investments made by new clients to pay older clients (while keeping a part of the money for himself); by repaying previous clients, the fake experts builds up a reputation of expertise that attracts new clients; the investments of the new clients allow the fake expert to pay the older clients, which helps sustain the reputation of expertise and so on. While Ponzi schemes rely on the existence of clients with limited information, they do not necessarily require agents to be boundedly rational. A client with rational expectations, who is uncertain of the expert's business model would rationally believe the expert is not likely to be fake if he continuously generates large returns to his clients. However, Ponzi schemes do have a problem for the fake expert; they place a lot of pressure on the fake expert in not only getting more clients, but also getting them to make larger and larger investments so that the increasingly larger pool of clients can be paid. Therefore, while Ponzi schemes seem to work even with rational people, they are only self-sustainable in the short-run.

This observation raises the question of this paper: is it possible for a fake expert to generate a *long-run* reputation of expertise, i.e., is it possible that an ordinary person, with no particular talent or knowledge, is able to convince rational observers that she is an expert and never be caught, no matter how much time goes by?

There is a literature in economic theory, initiated by Foster and Vohra (1998), who has studied a similar question: is it possible to create a test, based on the actual predictions of the expert, which succeeds in separating fake experts from true experts? The general consensus is *no*; in particular, this literature finds that any test that can be passed by a true expert with probability ε can also be passed by a fake expert with probability ε (see, for example, Sandroni (2003), Olszewski and Sandroni (2008) and Olszewski and Sandroni (2009)). However, as I show in the following example, this pessimistic result is largely a consequence of taking a *prior-free* approach.

Example: Suppose a self-proclaimed expert claims to know the probability that it will rain tomorrow. Formally, say that there is a random variable $\omega \in \{0, 1\}$, where $\omega = 1$ represents the outcome "rain tomorrow", and that the probability that it rains tomorrow is given by $p \in [0, 1]$. If the expert is the true expert, he observes p; if he is the fake expert, he does not. An outside observer, in order to test whether the expert is real or fake, proposes the following test: if the expert reports that the probability of rain is $\widehat{p} \geq \frac{1}{2}$, he passes if and only it rains ($\omega = 1$); if the expert reports $\widehat{p} < \frac{1}{2}$, he passes if and only if it does not rain ($\omega = 0$).

In the literature of expert testing, the distribution of p is not specified, which is what makes this approach *prior-free*. This creates difficulties in evaluating the test. Presumably, a good test is passed more often by the expert if he is true. However, without knowing the distribution of p one cannot calculate the probability the expert passes if he is true; if the expert is true, he reports $\widehat{p} = p$ for all $p \in [0, 1]$, which means that, conditional on p, he passes with probability $\max\{p, 1-p\}$. Instead, what this literature does is that it uses the *worst case scenario*, i.e., it assumes that the probability that the true expert passes the test is given by

$$\min_{p \in [0,1]} \{\max\{p, 1-p\}\} = \frac{1}{2}.$$

As for the fake expert, who does not know the true p, he chooses a strategy that maximizes the probability of passing the test. For this particular test, he chooses to randomize using a uniform distribution over interval $[0, 1]$. In that way, the fake expert is able to pass the test with a probability of $\frac{1}{2}$ for any realized p.[1] Seeing as both types of the expert pass the test with the same probability, one would conclude that this test is bad.

I follow a different approach; I model the interaction between a long-lived expert and a sequence of short-lived rational agents, who, in each period, decide whether to consult the expert. That decision depends in large part of the expert's reputation of expertise, which, in turn, depends on his past success at predicting events. As a result, methodologically, my analysis resembles more the literature on reputations, in

[1] For any p, the probability that the fake expert passes the test when he reports $\widehat{p} \geq \frac{1}{2}$ is p, while the probability of passing the test when the report is $\widehat{p} < \frac{1}{2}$ is $(1-p)$. Seeing as the fake expert randomizes uniformly, it follows that the probability of passing the test is given by

$$\frac{1}{2}p + \frac{1}{2}(1-p) = \frac{1}{2}.$$

particular, the subset that studies imperfect monitoring (for example, Cripps, Mailath and Samuelson, 2004).[2]

The previous example is illustrative of the difficulties this approach brings. Imagine that one drops the prior-free assumption and assumes that $p \sim U(0,1)$. In that case, one can explicitly calculate the probability that the expert passes the test if he is the true type; it is

$$E(\max\{p, 1-p\}) = \frac{3}{4}.$$

Seeing as the probability that a fake expert passes the test is only $\frac{1}{2}$, the test is informative. In other words, the posterior belief that the expert is fake formed by a rational agent who has just observed that the test has been failed will be larger than the prior belief. For example, if the prior belief that the expert is fake is $\frac{1}{2}$, it will go up to $\frac{3}{4}$ after the test is failed. But then, one would think that over time, as more and more predictions are made, agents would be able to determine with certainty (or close to it) whether the expert is fake or not. Therefore, it seems impossible for fake experts to sustain a long-run reputation of expertise: on the one hand, true experts should know more things than fake experts by definition, but, on the other hand, if they know more, they should make better predictions, which should eventually reveal their true type.

Nevertheless, I find that it is indeed possible for a fake expert to sustain a permanent reputation of expertise. In particular, the main result of the paper is that, under certain conditions, the prior belief that the expert is fake is never updated regardless of how good the expert's predictions turn out to be empirically.

By definition, I cannot point to any empirical example to corroborate my claim that it is possible that fake experts sustain permanent reputations; that is precisely the point of the paper, it might be impossible to empirically identify fake experts. Despite this, there is a group of experts that is routinely labeled as fake which has served as an inspiration for the argument of the paper; the group made of oracles, psychics, fortune tellers, etc. In particular, the oracles of Ancient Greece, who would give predictions about the future, enjoyed a period of success that lasted hundreds of years. At the time, the oracles claimed that their predicting power came directly from contact with the Gods. Despite this assertion, historians have remained skeptic and have studied alternative explanations for their success. One argument that has been put forward is that oracles, particularly the larger ones like the oracle of Delphi, were successful because they 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

[2] In the conclusion, I discuss in more detail the connection with this literature.

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 as true, 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, 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 the true expert is by being well connected. In some ways, the argument is very similar to the Ponzi scheme described above - it is a sort of informational Ponzi scheme: the expert gets people to consult him; those people make him well connected; those connections allow him to make good predictions, which cement his reputation of expertise; the enhanced reputation attracts more people, which makes the expert even more connected and so on. One crucial difference to standard Ponzi schemes is that the informational Ponzi scheme is not necessarily bad for the agents who consult the expert. The fake expert who employs an informational Ponzi scheme is actually able to give good advice; in fact, that is what allows him to maintain a long-run reputation. It is just that the good advice that is provided comes not from some intrinsic ability or knowledge; it comes from inside information.

The last part of the paper discusses the following question: if a fake expert actually gives good advice (because of his connections), why would he feel the need to pretend to be the true expert? Why not just announce that he is the fake expert? Suppose that the expert is known to be fake; he is only as good as his connections. In this case, there is some probability that some agent A chooses not to consult the expert because he believes that the potential information those connections would bring is not relevant enough for him to go through the trouble of consulting the expert. This makes future agents less willing to consult the expert, because the expert will not be as well connected as he would have, had agent A consulted him, which might generate a negative spiral of agents who choose not to consult the expert. I find that, under some plausible assumptions, this may actually lead to the expert not being consulted in the long-run, as opposed to ensuring that all of the agents consult him when he pretends to be the true expert.

Finally, while oracles have served as an inspiration for the argument, the argument itself - that fake experts can sustain a permanent reputation of expertise because of how well connected they become - is reminiscent of ideas that have been discussed before in the economics literature. 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.

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 agents were not allowed to observe the public history of the expert, the result of the next section would hold trivially.

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 η_t in every period $t \geq 1$. Figure 1 displays the timing of the events within each period $t \geq 1$.

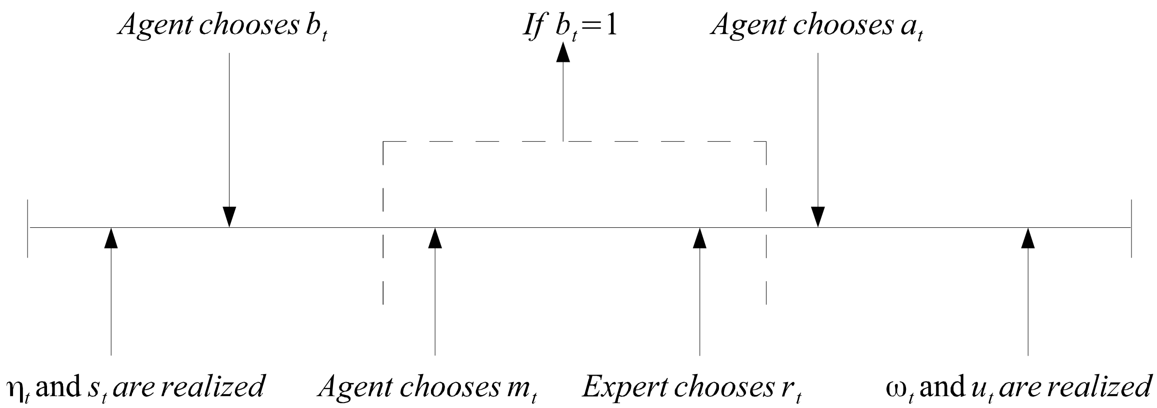

Figure 1: Timing within any period $t \geq 1$

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 future states of the world. I model s_0 mostly for technical reasons as will become clear in the next sections, but I limit its impact by assuming that there is some $T > 0$ such that s_0 is uncorrelated with ω_t for all $t \geq T$, i.e., s_0 might provide some information but only in the short term. More importantly, the expert is also assumed to observe $\theta \in \Theta$, which can be interpreted as his 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.

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 he is agent $t = 0$, because he observes s_0, which, by assumption, is only correlated with the first few states of the world. In other words, in the natural world, the expert is just an ordinary person, without any special intrinsic skill or knowledge, 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 past signals; the primitive of the model is $\Pr\{\omega_t | s_t, ..., s_0\}$.[3] The probability that $\theta = \theta_N$ is denoted by $\mu \in [0,1]$.

By contrast, if $\theta \in \Theta_S$, I say 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 has some intrinsic knowledge that ordinary people do not have. It might be a superior cognitive ability, a special understanding of financial markets or, in the example of the oracles, an ability to communicate with the Gods.

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

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

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 his 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.

The main result of the paper is about how the expert can systematically convince agents that they are in a supernatural world even when they are in the natural world.

[3] 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.

3 Main Result

For now, assume that $\mu = 1$, i.e., it is known that $\theta = \theta_N$, and consider the following strategy profile:

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 ($a_t = r_t$).

The previous strategy profile, while not necessarily a part of a perfect Bayesian equilibrium (PBE), does generate a path of play. More specifically, 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.

Assumption A1: For all $\theta \in \Theta_S$ and for all $t \geq 1$,

$$\frac{1}{2} + \lambda_t^\theta \geq \int_0^1 v_t^{h^t}(s_t)\, ds_t$$

for all $h^t \in H^t$.

The left hand side of the expression represents the probability that the expert of the supernatural world θ is able to correctly predict the state of world ω_t. The right hand side represents the probability that the expert of the natural world is able to correctly predict the state of the world ω_t, conditional on the public history h^t and given the strategy profile described above (notice that, even though the natural expert has no intrinsic knowledge apart from s_0, she is given the signals s_t of the various agents who consult her). Therefore, assumption A1 ensures that each supernatural world is good enough to make the expert of the natural world want to convince the agents that the real world is the supernatural world; in each supernatural world, the expert's predicting ability is weakly larger than in the natural world, even when the expert has access to all of the agents' private signals.

Proposition 1 *If assumption A1 holds and there is some $\gamma > c$ such that, for all $h^t \in H^t$,*

$$\int_0^1 v_t^{h^t}(s_t)\, ds_t \geq \frac{1}{2} + \gamma, \tag{1}$$

there is $\overline{\mu} \in (0,1)$ 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. ■

As is stated, the normal type expert is able to generate a long-run reputation of being a supernatural type if two things happen: the prior belief in supernatural types must be large enough and condition (1) must hold for any $h^t \in H^t$. Condition (1) can be interpreted as follows. Imagine a hypothetical scenario where it is known that the world is the natural world but where each agent decides whether to consult the expert *before* observing his private signal s_t. Condition (1) means that, in such a scenario, each agent wants to consult the expert for any history of play, provided the expert is forthcoming in her advice. Therefore, not only is the existence of the normal type expert positive for the agents; his ability to generate a long-run reputation of being a supernatural type depends directly on how good the advice he gives really is.

The complete proof is given in the appendix, but I illustrate the main intuition with the following example.

Example 1 *Suppose that*

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

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.

In order to determine $v_t^{h^t}(s_t)$, we assume that it is known that $\theta = \theta_N$ and calculate the probability that agent t matches the state after some history h^t on the path of play.

It follows that

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

because agent t will gain access to s_{t-1} by consulting the expert.[4] *Therefore, assumption A1 simply requires that*

$$\int_0^1 v_t^{h^t}(s_t)\, ds_t = \frac{3}{4} \leq \frac{1}{2} + \lambda_t^{\theta}$$

for all $\theta \in \Theta_S$ and for all $t \geq 1$, i.e., the probability that the expert of the supernatural world matches the state is at least $\frac{3}{4}$. For now, let us assume that $\Theta_S = \{\theta_s\}$ and that $\lambda_t^{\theta_s} = \frac{1}{4}$ for all $t \geq 1$. In this way, the probability that the supernatural expert matches the state is exactly $\frac{3}{4}$. Assume that $c < \frac{1}{4}$ so that the conditions of proposition 1 may be satisfied if μ is sufficiently small.

Suppose $\mu \in (0, 1)$ and consider the problem of some agent t on the path of play. Let $\pi(h^t)$ denote the posterior probability that $\theta = \theta_N$ given the on-path public history h^t. He consults the expert if and only if

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

because, if $\theta = \theta_N$, the probability of matching the state after not consulting the expert can be shown to be equal to $\frac{1}{2}$.[5] *Given that $c < \frac{1}{4}$, it follows that agent t consults the expert for any $s_t \in [0, 1]$ provided $\pi(h^t)$ is sufficiently small.*

The argument is completed by showing that $\pi(h^t) = \mu$ for any on-path public history. Think of some agent t and suppose that his belief about θ is given by $\pi(h^t) = \mu$. The probability that he matches the state, provided he consults the expert and follows her advice, depends on 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

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

i.e., it is as likely for the agent to match the state in the natural world as it is in the

[4] 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 particular history h^t considered.

[5] This follows because, for any public history on the path of play, agent t believes the distribution of s_{t-1} is symmetric around $\frac{1}{2}$.

supernatural world. Therefore, while agent $t+1$ observes one more piece of evidence about the expert than agent t, he does not update his beliefs.

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 expert of the supernatural world 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. Neither type of expert wants to deviate because every agent would consult the expert in every period, provided μ is sufficiently large. 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 $\frac{3}{4}$.

4 Discussion

4.1 The natural world

Proposition 1 describes the conditions under which a fake or natural expert is able to generate a permanent reputation of being of a supernatural type. Under those conditions, by claiming to be a supernatural expert, she ensures that every agent consults her in every period for any history of play. However, in principle, it could be that the same is true if the expert is known to be the natural expert. If that was to happen, there would be no reason for the natural expert to want to pretend to be a supernatural expert. In this section, I argue that it is at least plausible that the expert is unable to attract every agent in every period when it is known that she is the natural type, even when the conditions of proposition 1 hold.

The argument is that, because the expert depends exclusively on her network, she is particularly vulnerable to agents believing that their private signal is enough for them to make a good guess about ω_t. In that case, not only does the expert lose directly by not receiving the agent in question, she also loses by not getting access to the agent's private signal, which limits her ability to attract future agents. The following two assumptions formalize this idea.

Assumption A2 *There is* $\varepsilon \in (0, c)$ *and* $\upsilon > 0$ *such that, for all* $t \geq 1$*, there is an* $I_t \in [0,1]$ *such that*

$$\int\limits_{s_t \in I_t} s_t ds_t > \upsilon$$

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

In words, assumption A2 means that, in every period t, there could be a signal s_t such that the agent cares very little about the previous signals, i.e., for some signals s_t, knowing the previous signals only marginally improves the odds of the agent making a good decision. As a result, if the world is known to be the natural world, such an agent would not be willing to incur the cost c in order to consult the expert no matter what. This implies that, for any period t, and for any PBE, the probability that agent t consults the expert is less than 1. On its own, this assumption should already be enough for the natural expert to want to pretend to be the supernatural expert. The following assumption makes it even more appealing.

Assumption A3 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 A3 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. Assumptions A2 and A3 together imply that the probability that agent t consults the expert converges to 0 in any PBE.

Proposition 2 *If assumptions A2 and A3 hold and* $\mu = 1$*, then* $b_t \rightarrow^{a.s.} 0$ *in any PBE.*

Proof. Assumption A3 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 A2 implies that, no matter the history, there is always a positive probability that such a sequence occurs. ■

The example from the previous section is also useful to illustrate the challenges of the expert if it is known that he is the normal type. Notice that $T = 1$ in the example and that assumption A2 holds because

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

which makes agents who receive a signal s_t close to $\frac{1}{2}$ unwilling to pay the cost c no matter what. Therefore, in the example, it follows that eventually no agent would visit the expert if she is known to be the natural expert. By contrast, by claiming to be the supernatural expert, she is able to attract every agent.

Figure 2 helps in demonstrating why the normal type expert benefits from the public uncertainty about his type. In red, I represent function $v_t^{h^t}(s_t)$, the marginal benefit of consulting the expert if all previous agents have and if the world is the natural world. The straight line at $\frac{3}{4}$ represents the marginal benefit of consulting the expert if the world is the supernatural world (under the assumption that $\Theta_s = \{\theta_s\}$ and that $\lambda_t^{\theta_S} = \frac{1}{4}$). Even though the two scenarios lead to the same expected marginal benefit, the normal type expert prefers agent t to believe that the real marginal benefit is constant and given by $\frac{3}{4}$ rather than given by the red line, i.e., the natural expert prefers the flatter line. In that way, provided $\frac{3}{4} > \frac{1}{2} + c$, the agent consults the expert for any signal s_t. By contrast, if the agent was to believe that the marginal benefit was given by the red line, he would not consult for some signals close to $\frac{1}{2}$, which would break the chain of agents who consult the expert. Uncertainty over the expert's type makes the agent average the two lines, which is represented by the green line. So, the prior belief on the supernatural world only needs to be large enough to make sure that the lowest point of the green line is above the threshold $\frac{1}{2} + c$.

The geometric nature of this argument makes it somewhat reminiscent of the Bayesian Persuasion literature (Kamenica and Gentzkow, 2011) in that, in a way, the expert uses the flexibility of the agents' rational expectations' assumption and the fact that previous signals are not public to increase her payoff. In order for the belief never to be updated, what is necessary is that the expected probability of matching the state is the same in both worlds. But because the conditional (on s_t) probability need not be the same, the natural type expert is able to convince the agents that the probability of matching the state is independent of their own private signal.

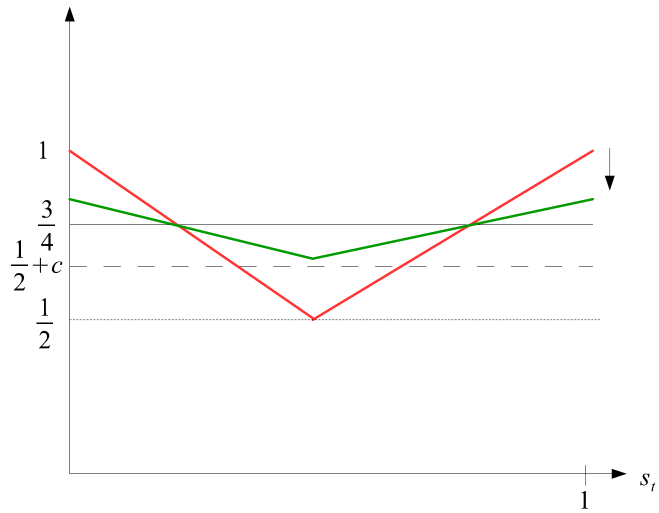

Figure 2: The red line represents function $\max\{s_t, 1 - s_t\}$, while the green line represents a weighted average between the red line and $\frac{3}{4}$.

4.2 Welfare

The existence of an expert, fake or not, does not harm the agents, provided they have rational expectations. What may or may or may not harm the agents is the uncertainty about the expert's type. In particular, what goes against the popular view about fake experts in this paper is that uncertainty over the expert's type might actually make the agents better off. In particular, conditional on the expert being the natural expert, almost all agents might be better off believing that she might be the supernatural expert. The argument is easy to understand from the example. If it is known that the expert is the natural expert, eventually the expert stops being consulted. So, an agent t in the future will essentially not have the option to access signal s_{t-1}, because agent $t-1$ will not have consulted the expert. But if the expert is believed to be the supernatural expert when he really is the natural expert, agent $t-1$ will have consulted the expert, so that agent t will be able to access s_{t-1} by also consulting the expert. While the uncertainty might not benefit all agents - for example, agent 1 would prefer to have no uncertainty - future agents almost surely benefit from it; in fact, recall that one of the conditions for a long-run reputation to be possible is precisely that the expected benefit of consulting the expert when he is the natural expert is larger than the cost of consulting her.

4.3 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, the paper questions the extent to which markets should be regulated, as excessive regulation that succeeds in driving fake experts away might be undesirable.[6]

The paper is also related to the literature that has studied financial experts who make public recommendations. Rudiger and Vigier (2019) describe a similar result to this paper's: that fake financial experts can sustain a permanent reputation of expertise. The argument is as follows. Experts give advice over whether to buy or sell some stock. Naturally, when people are sufficiently convinced of the expert's ability, they will always follow the expert's advice no matter what. This allows the expert to make self-fulfilling predictions even when her actual ability is small. For example, imagine that the expert predicts the price of a certain stock will rise; believing the expert is very able, every agent buys the stock, causing its price to rise, just as predicted, i.e., the price rose not because of any private information the expert might have had but because every agent believed her. This type of argument depends crucially on the reduced scrutiny that the experts are under, which is in contrast to my paper. In this literature, what is observed is only whether people follow the expert's advice; the underlying value of the asset, which presumably led to the expert's prediction is never made public. By contrast, in my paper, outside observers can judge experts' predictions after having observed the realization of the phenomenon in question; for example, if the expert recommends the agent use an umbrella the following day, an outside observer is able to judge how good that advice was by observing whether it has rained. This difference in the setup might also lead to experts not wanting to

[6]Berk and van Binsbergen (2017) also discuss whether fake experts increase welfare. While they find that fake experts generate inefficiencies, they also argue that they might make consumers better off (albeit for different reasons than the ones considered here); it is the true experts of the market that suffer.

give good advice if they are concerned with their reputation (Ottaviani and Sorensen (2006)), which is not true in this paper.

Finally, this paper is also related to 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 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.[7]

5 Conclusion

This paper builds on arguments about the oracles of ancient Greece to make the more general claim that it is possible that a fake expert is able to maintain a permanent reputation of expertise despite continuous scrutiny. Fake experts rely on a sort of "informational Ponzi scheme"; the agents who consult the expert allow the expert to be well connected, those connections allow the expert to give good predictions which builds a reputation of expertise, that reputation brings in more people, making the expert even more connected and so on. Even though fake experts give good advice, they are reluctant to announce that what makes them valuable are their connections;

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

doing so might lead to some agents refusing to consult the expert (if s_t is close to $\frac{1}{2}$ in the example), which not only has a direct impact on her (the agent will not consult) but also an indirect impact, because the expert becomes less connected. By pretending that the good advice they give has a different source, whether that is a superior cognitive ability or contact with the Gods, experts are able to mitigate that risk.

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 $\mu \leq \frac{1}{2}$.

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

$$\rho_t^{h^t}(\theta)\left(\frac{1}{2} + \lambda_t^\theta\right) + \left(1 - \rho_t^{h^t}(\theta)\right)\left(\frac{1}{2} - \lambda_t^\theta\right) \equiv \widehat{\rho}_t^{h^t}(\theta) = \int_0^1 v_t^{h^t}(s_t)\, ds_t,$$

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 just 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

$$(1 - \mu)\left(\int_0^1 v_t^{h^t}(s_t)\, ds_t - \frac{1}{2}\right) > c$$

because

$$\int_0^1 v_t^{h^t}(s_t)\,ds_t \geq \frac{1}{2}+\gamma.$$

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

$$\left(1-\pi\left(h^t\right)\right)\left(\int_0^1 v_t^{h^t}(s_t)\,ds_t - \frac{1}{2}\right) \geq c.$$

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 $\widehat{\rho}_t^{h^t}(\theta)$ if $\theta \in \theta_S$. Seeing as agent $t+1$ does not observe s_t, and seeing as, by definition,

$$\widehat{\rho}_t^{h^t}(\theta) = \int_0^1 v_t^{h^t}(s_t)\,ds_t,$$

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.