No. 2
ENERO de 2017

quantil | Documentos de Trabajo

ISSN 1234-5790 Edición electrónica

Screening Multiple Uninformed Experts

Francisco Barreras

quantil | matemáticas aplicadas

Serie Documentos de Trabajo Quantil, 2017-2
Edición electrónica.

ENERO de 2017

Carrera 7 # 77 - 07. Oficina 901, Bogotá, D. C., Colombia
Teléfonos: 3718132 – (310)6791459 – (320)8461236
E-mail: info@quantil.com.co
http://www.quantil.com.co

Impreso en Colombia – Printed in Colombia

Screening Multiple Uninformed Experts

Francisco Barreras*
Research Department
Quantil

January 30, 2017

Abstract

Testing the validity of claims made by self-proclaimed experts can be impossible when testing them in isolation, even with infinite observations at the disposal of the tester. However, in a multiple expert setting it's possible to design a contract that only informed experts accept and uninformed experts reject. The tester can pit competing theories against each other and take advantage of the uncertainty experts have about the other experts' type. This contract will work even when there is only a single data point to evaluate.

1 Introduction

This paper studies the relationship between two potentially uninformed experts who deliver forecasts to a principal, called Alice. Alice needs a mechanism to induce informed experts to reveal their knowledge honestly and to screen uninformed experts that would deliver useless and potentially harmful forecasts.

Olszewski and Sandroni (2007,2008) show the impossibility of screening uninformed experts when evaluating them in isolation. This motivates evaluating multiple experts at once, since the possibility of one forecast performing better than the other opens new possibilities for screening. Al-Najjar and Weinstein (2008) show that, when a true expert is present, there is a test that only informed experts can pass. However, it potentially requires infinite data points and does not generalize to the case where all experts are

*E-mail: francisco.barreras@quantil.com.co. This papes is based on my dissertation submitted in partial fulfillment of the requirements for the degree of Master of Science in the Faculty of Economics at the University of los Andes. I thank my advisor Alvaro Riascos for his guidance in this research and good ideas. I would also like to thank Alvaro Sandroni for helpful comments and suggestions.

potentially uninformed.

The results on this paper are a generalization of results by Sandroni (2014) were he models the interaction between a principal and a potential expert. By means of an amendment to a proper scoring rule, Sandroni finds a contract that compels informed experts to reveal the truth and uninformed experts to ďo no harm' in the sense that they predict prior odds held by Alice.

I extend Sandroni in the following way; Alice offers a contract to the experts that determines transfers based on the experts' forecasts and the future observed state of Nature. This contract specifies transfers according to how high each forecast scores on the Brier score (Brier, 1950) compared to the rival forecasts. This contract is accepted by informed experts and gives incentives to revealing the true odds, but it is rejected by uncertainty averse uninformed experts. This result holds even when there's only a single data point available. Moreover, this result can be extended to construct contracts that achieve screening of experts that are better informed than others. Results are presented for two experts, but can trivially be extended to an arbitrary number of experts.

2 Model

Let S be a finite set of states (e.g. a set of possible finite histories). Let $\Delta(S)$ be the set of probability distributions over S. Two experts, referred to as expert 1 and expert 2, deliver probabilistic forecasts f_1 and $f_2 \in \Delta(S)$ to Alice.

Alice creates a contract that specifies money transfers between her and each expert to elicit information. A contract is a payoff function $C : \Delta(S) \times \Delta(S) \times S \to \mathbb{R}$ whose value depends on the announced odds and the observed state. If any expert rejects the contract his payoff is 0. [1] Since the behaviour of both experts is symmetrical, focus on the behaviour of expert 1, who is offered a contract C_1. If both experts accept their respective contracts, they deliver odds f_1 and f_2 and when state s is observed, expert 1 receives (or gives) payoff $C_1(f_1, f_2, s)$.

When expert 1 is informed, he maximizes his expected utility conditional on the other experts forecast. We say that expert 1 *accepts the contract* if for every $f_1, f_2 \in \Delta(S)$ we have $E^{f_1}\{C_1(f_1, f_2, \cdot)\} > 0$. That is, when revealing the truth gives him a positive payoff regardless of the other expert's forecast. Moreover, we say he *honestly reveal his beliefs* when for all $f_2 \in \Delta(S)$ and $f' \neq f_1 \in \Delta(S)$

$$E^{f_1}\{C_1(f_1, f_2, \cdot\,)\} > E^{f_1}\{C_1(f', f_2, \cdot\,)\}$$

[1]If only one expert remains and the tester is sure he's informed. He can receive a second contract that elicits true information when he's informed. See for example Sandroni (2014).

this is a property that proper scoring rules will guarantee and it ensures that informed experts won't misrepresent their beliefs. When expert 1 is uninformed, he evaluates his prospects using the minmax criteria [2] as in Gilboa and Schmeidler (1989). Considering both experts may announce their odds using random generators of theories ξ_1 and $\xi_2 \in \Delta(\Delta(S))$ this can be stated as saying that uninformed expert 1 only accepts contract C_1 when there exists a random generator of theories ξ_1 such that:

$$\min_{\substack{f \in \Theta_1 \\ \xi_2 \in \Delta\Delta(S)}} \iint_{\Delta(S)\Delta(S)} E^f C^1(f', f^*, \cdot)\xi_2(f^*)\xi_1(f') > 0, \tag{1}$$

where Θ_1 is a subset of $\Delta(S)$ that contains the theories expert 1 deems plausible. If expert 1 is uninformed, we say that he rejects the contract C_1 when there's no $\xi_1 \in \Delta(\Delta(S))$ that satisfies (1).

3 Main Result

PROPOSITION 1. In a multiple expert setting assume that Θ_1 contains at least two points. There's a contract C_1 such that expert 1 if informed accepts it and reveals his knowledge and if uninformed rejects it.

The intuition of the proof is simple. Design a contract that gives a payoff proportional to the difference of the Brier score plus a small enough ε. Informed experts can be assured to get paid at least ε since the Brier score is maximized with the true odds. Uninformed experts get negative payments when the other expert is informed and they forecast theories too distant to the true odds. In order to guarantee a positive payment they would need to randomize in a way that their forecasts are always close to the true odds and this is impossible. Because the contract depends on the set Θ_1, this solution doesn't give a single contract that screens all uninformed experts, but rather for each set of experts a contract that screens informed from uninformed experts exists.

This contract resembles the test in Al-Najjar and Weinstein (2008) in that it compares theories against each other in a way that the true theory will over perform the others. However, this result is valid even with a single data point and without the assumption that the tester knows there's an informed expert present. It's enough to assume that the experts are uncertain about each other's strategy. This assumption can't be relaxed, since uninformed experts with identical forecasts can always secure a positive payment in the proposed contract.

The result can be extended to a setting involving partially informed experts.An expert

[2]The uninformed expert is extremely averse to uncertainty and will accept a contract only if he gets a positive payoff in his worst case scenario.

is partially informed if he is uninformed and his set of plausible theories Θ takes the form $B_\delta(f^*) = \{f \in \Delta(S) : ||f - f^*|| \leq \delta\}$ for $f^* \in \Delta(S)$ [3]. We say that an expert is better informed than other when their sets of plausible theories are $B_{\varepsilon_1(f_1)}$ and $B_{\varepsilon_2}(f_2)$ and $\varepsilon_1 < \varepsilon_2$ for some pair of theories f_1 and f_2.

PROPOSITION 2. In a setting with two partially informed experts, where one is better informed than the other there's a contract that is accepted by the better informed expert and rejected by the other.

The assumption that there are perfectly informed experts is unrealistic. Proposition 2 shows us that even in the case when one expert is slightly better informed than the other, there is a contract that achieves perfect screening. Proposition 1 might be regarded as a degenerate case of proposition 2.

4 Conclusion

Screening informed and uninformed experts can be difficult when evaluating a single expert, however, the presence of multiple experts brings strategical uncertainty to uninformed experts which can be exploited to design a contract that only informed experts would accept.

5 Appendix

Lemma 1. The Brier Score $B : \Delta(S) \times S \to \mathbb{R}$, defined as $B(f, s) = 2f(s) - \sum_{s' \in S}(f(s'))^2 - 1$ is such that:

$$E^f\{B(g, \cdot\ )\} = ||f||_2^2 - ||f - g||_2^2 - 1, \tag{2}$$

where $||\cdot||_2^2$ denotes the $\mathcal{L}_2$ norm squared.

Proof.

$$\begin{aligned} E^f\{B(g, \cdot\ )\} &= -\sum_{s \in S} f(s) \left(1 - 2g(s) + \sum_{s' \in S}(g(s'))^2\right) \\ &= \sum_{s' \in S}(g(s'))^2 + \sum_{s \in S} 2f(s)g(s) - 1 \\ &= ||f||_2^2 - ||f - g||_2^2 - 1 \end{aligned}$$

[3]A reasonable assumption is that the true odds be in $B_\varepsilon(f^*)$ but this is not necessary for the result

□

PROOF OF PROPOSITION 1: Let B be the Brier Score, as previously defined, and let f_x and f_y be two different elements of Θ_1. Define the contract C_1 for expert 1 as $C_1(f_1, f_2, s) = B(f_1, s) - B(f_2, s) + \varepsilon$, where $\varepsilon = \dfrac{||f-g||_2^2}{2}$.

The informed expert accepts the contract because, applying Lemma 1

$$\begin{aligned}&E^f\{C_1(f, f_2, \cdot\ )\} = \big(||f||_2^2 - ||f-f||_2^2 - 1\big) - \big(||f||_2^2 - ||f-f_2||_2^2 - 1\big) + \varepsilon\\ &= ||f-f_2||_2^2 + \varepsilon > 0.\end{aligned}$$

Moreover, he honestly reveals the truth since $\forall f_1 \neq f$:

$$E^f\{C_1(f, f_2, \cdot\ )\} = ||f-f_2||_2^2 + \varepsilon > ||f-f_2||_2^2 - ||f-f_1||_2^2 + \varepsilon = E^f\{C_1(f_1, f_2, \cdot\ )\}$$

If expert 1 is uninformed and forecasts using a random generator of theories ξ_1 begin by noting that his maxmin payoff is bounded above by the one obtained if the other expert forecasts the true probability (which he would if he's informed). Formally

$$\min_{\substack{f\in\Theta_1\\ \xi_2\in\Delta(\Delta(S))}} \iint_{\Delta(S)\Delta(S)} E^f C^1(f', f^*, \cdot) d\xi_2(f^*) d\xi_1(f') = \min_{\substack{f\in\Theta_1\\ \xi_2\in\Delta(\Delta(S))}} \iint_{\Delta(S)\Delta(S)} ||f-f_2||_2^2 - ||f-f_1||_2^2 + \varepsilon d\xi_2(f^*) d\xi_1(f') \tag{3}$$

$$\leq \min_{f\in\Theta_1} \int_{\Delta(S)} \varepsilon - ||f-f_1||_2^2 d\xi_1(f') = \min_{f\in\Theta_1} \int_{\Delta(S)} E^f C^1(f', f, \cdot) d\xi_1(f'). \tag{4}$$

However, the expression in (4) is negative for every value of ξ_1 because

$$\begin{aligned}\min_{f\in\Theta_1} \int_{\Delta(S)} \varepsilon - ||f-f_1||_2^2 d\xi_1(f') &= \varepsilon - \max_{f\in\Theta_1} \int_{\Delta(S)} ||f-f_1||_2^2 d\xi_1(f')\\ &\leq \varepsilon - \max_{f\in\Theta_1} \left\| \int_{\Delta(S)} f - f' d\xi_1(f') \right\|_2^2\\ &= \varepsilon - \max_{f\in\Theta_1} \left\| f - \int_{\Delta(S)} f' d\xi_1(f') \right\|_2^2,\end{aligned}$$

the inequality comes from Jensen's inequality after noting that $||\cdot||_2^2$ is convex. If we define $\bar{f} := \int_{\Delta(S)} f' d\xi_1(f')$, we have that $\bar{f} \in \Delta(S)$ and then, using the triangular inequality, we get:

$$\varepsilon - \max_{f\in\Theta_1}\left\| f - \int_{\Delta(S)} f' d\xi_1(f')\right\|_2^2 = \varepsilon - \max_{f\in\Theta_1}\|f-\bar{f}\|_2^2$$
$$\leq \varepsilon - \frac{\|f_x-\bar{f}\|_2^2 + \|f_y-\bar{f}\|_2^2}{2}$$
$$\leq \varepsilon - \frac{\|f_x - f_y\|_2^2}{2} < 0,$$

so the uninformed expert never accepts the contract.

PROOF OF PROPOSITION 2: Let $\Theta_1 = B_{\varepsilon_1}(f_1)$ and $\Theta_2 = B_{\varepsilon_2}(f_2)$ for a pair of theories f_1 and f_2 in $\Delta(S)$. Without loss of generality assume $\varepsilon_2 > \varepsilon_1 > 0$. Let B be the Brier Score and let γ be such that $\varepsilon_2^2 > \gamma > \varepsilon_1^2$. Define contracts C_1 and C_2 for experts 1 and 2 respectively as $C_1(f_1, f_2, s) = B(f_1, s) - B(f_2, s) + \gamma^2$ and $C_2(f_2, f_1, s) = B(f_2, s) - B(f_1, s) + \gamma^2$.

As before, provided that experts may use random generators of theories $\xi_1, \xi_2 \in \Delta(\Delta(S))$, expert i will accept the contract if there exists $\xi_i \in \Delta(\Delta(S))$ such that

$$\min_{\substack{f\in\Theta_i\\ \xi_j\in\Delta(\Delta(S))}} \int\int_{\Delta(S)\Delta(S)} ||f'-f|| + (\gamma - ||f-f^*||_2^2) d\xi_i(f^*) d\xi_j(f') > 0, \tag{5}$$

for $i \neq j$.

It is simple to see that expert 1 accepts his contract, because he can get a positive payoff by making $\xi_1(\{f_1\}) = 1$. Formally

$$\min_{\substack{f\in\Theta_i\\ \xi_2\in\Delta(\Delta(S))}} \int_{\Delta(S)} ||f'-f|| + (\gamma - ||f-f_1||_2^2) d\xi_2(f') \geq \min_{f\in\Theta_i} \gamma - ||f-f_1||_2^2 \geq \gamma - \varepsilon_1^2 > 0.$$

The first inequality comes from the fact that the $||\cdot||$ is non-negative and the second from the fact that $\Theta_1 = B_{\varepsilon_1}(f_1)$. Since there are two theories f_x and f_y in Θ_2 such that $\gamma < \varepsilon_2^2 \leq \dfrac{\|f_x - f_y\|_2^2}{2}$ then expert 2 rejects the contract following the same argumentation as in the proof of proposition 1.

References

[1] Nabil I Al-Najjar and Jonathan Weinstein. Comparative testing of experts. *Econometrica*, pages 541–559, 2008.

[2] Glenn W Brier. Verification of forecasts expressed in terms of probability. *Monthly*

weather review, 78(1):1–3, 1950.

[3] Wojciech Olszewski and Alvaro Sandroni. Contracts and uncertainty. *Theoretical Economics*, 2(1):1–13, 2007.

[4] Wojciech Olszewski and Alvaro Sandroni. Strategic manipulation of empirical tests. 2008.

[5] Alvaro Sandroni. At least do no harm: The use of scarce data. *American Economic Journal: Microeconomics*, 6(1):1–4, 2014.