Games and Economic Behavior 109 (2018) 271–293

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Smooth calibration, leaky forecasts, finite recall, and Nash dynamics ☆

Dean P. Foster [a,b], Sergiu Hart [c,*]

[a] Amazon Inc., New York City, United States
[b] University of Pennsylvania, Philadelphia, United States
[c] Institute of Mathematics, Department of Economics, and Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel

ARTICLE INFO

Article history:
Received 30 March 2017
Available online xxxx

JEL classification:
C7
D8
C1

Keywords:
Calibration
Nash dynamics
Fixed points
Deterministic calibration
Smooth calibration
Finite recall

ABSTRACT

We propose to smooth out the calibration score, which measures how good a forecaster is, by combining nearby forecasts. While regular calibration can be guaranteed only by randomized forecasting procedures, we show that *smooth calibration* can be guaranteed by *deterministic* procedures. As a consequence, it does not matter if the forecasts are *leaked*, i.e., made known in advance: smooth calibration can nevertheless be guaranteed (while regular calibration cannot). Moreover, our procedure has finite recall, is stationary, and all forecasts lie on a finite grid. To construct the procedure, we deal also with the related setups of online linear regression and weak calibration. Finally, we show that smooth calibration yields uncoupled finite-memory *dynamics* in n-person games—"smooth calibrated learning"—in which the players play approximate *Nash equilibria* in almost all periods (by contrast, calibrated learning, which uses regular calibration, yields only that the time averages of play are approximate correlated equilibria).

1. Introduction

How good is a forecaster? Assume for concreteness that every day the forecaster issues a forecast of the type "the chance of rain tomorrow is 30%." A simple test one may conduct is to calculate the proportion of rainy days out of those days for which the forecast was 30%, and compare it to 30%; and do the same for all other forecasts. A forecaster is said to be *calibrated* if, in the long run, the differences between the actual proportions of rainy days and the forecasts are small—no matter what the weather really was (see Dawid, 1982).

What if rain is replaced by an event that is under the control of another agent? If the forecasts are made public before the agent decides on his action—we refer to this setup as "*leaky forecasts*"—then calibration *cannot* be guaranteed; for example, the agent can make the event happen if and only if the forecast is less than 50%, and so the forecasting error (that is, the "calibration score") is always at least 50%. However, if in each period the forecast and the agent's decision are made "simultaneously"—which means that neither one knows the other's decision before making his own—then calibration

☆ Previous versions: July 2012, February 2015, March 2017. Research of the second author was partially supported by a European Research Council (ERC) Advanced Investigator grant. The authors thank Yakov Babichenko for useful comments, and the editor, associate editor, and referees for their very careful reading and helpful suggestions.

* Corresponding author.
E-mail addresses: dean@foster.net (D.P. Foster), hart@huji.ac.il (S. Hart).
URLs: http://deanfoster.net/ (D.P. Foster), http://www.ma.huji.ac.il/hart (S. Hart).

https://doi.org/10.1016/j.geb.2017.12.022
0899-8256/

can be guaranteed; see Foster and Vohra (1998). The procedure that yields calibration no matter what the agent's decisions are requires the use of *randomizations* (e.g., with probability 1/2 the forecaster announces 30%, and with probability 1/2 he announces 60%). Indeed, as the discussion at the beginning of this paragraph suggests, one cannot have a deterministic procedure that is calibrated (see Dawid, 1985 and Oakes, 1985).

Now the standard calibration score is overly fastidious: the days when the forecast was, say, 30.01% are considered separately from the days when the forecast was 29.99% (formally, the calibration score is a highly discontinuous function of the data, i.e., the forecasts and the actions). This suggests that one first combines all days when the forecast was *close to* 30%, and only then compares the 30% with the proportion of rainy days. If, say, there were 200 days with a forecast of 30.01%, out of which 10 were rainy, and another 100 days with a forecast of 29.99%, out of which 80 were rainy, then the forecaster is very far from being calibrated; however, he is smoothly calibrated, as his forecasts were all close to 30%, and there were $90/300 = 30\%$ rainy days. Undershooting at 29.99% and overshooting at 30.01% is now balanced out. Formally, what this amounts to is applying a so-called "smoothing" operation to the forecasting errors (which makes smooth calibration easier to obtain than calibration).[1]

Perhaps surprisingly, once we consider smooth calibration, there is no longer a need for randomization when making the forecasts: we will show that there exist *deterministic* procedures that guarantee smooth calibration, no matter what the agent does. In particular, it follows that it does not matter if the forecasts are made known to the agent before his decision, and so smooth calibration can be guaranteed even when forecasts may be leaked.[2] This may come as a surprise, because, as pointed out above, an agent who knows the forecast *before* deciding on the weather will choose rain when the forecast is less than 50% and no rain otherwise, giving a calibration error of 50% or more, no matter what the forecaster does. However, against such an agent one can easily be *smoothly* calibrated, by forecasting 50.01% on odd days and 49.99% on even days (the resulting weather will then alternate between rain and no rain, and so half the days will be rainy days—and all the forecasts are indeed close to 50%). What this proves is only that one can be smoothly calibrated against this specific strategy of the agent (this is the strategy that shows that it is impossible to have calibration with deterministic leaky procedures); our result shows that one can in fact *guarantee* smooth calibration with a deterministic strategy, against *any* strategy of the agent.

The forecasting procedure that we construct and that guarantees smooth calibration has moreover finite recall (i.e., only the forecasts and actions of the last R periods are taken into account, for some fixed finite R), and is stationary (i.e., independent of "calendar time": the forecast is the same any time that the "window" of the past R periods is the same).[3] Finally, we can have all the forecasts lie on some finite fixed grid.

The construction starts with the "online linear regression" problem, introduced by Foster (1991), where one wants to generate every period a good linear estimator based only on the data up to that point. We provide a finite-recall stationary algorithm for this problem; see Section 3. We then use this algorithm, together with a fixed-point argument, to obtain "weak calibration", a concept introduced by Kakade and Foster (2004) and Foster and Kakade (2006); see Section 4. Section 5 shows that weak and smooth calibration are essentially equivalent, which yields the existence of smoothly calibrated procedures. Finally, these procedures are used to obtain dynamics ("smoothly calibrated learning") that are uncoupled, have finite memory, and are close to Nash equilibria most of the time (while the similar dynamics that are based on regular calibration yield only the time average becoming close to correlated equilibria; see Foster and Vohra, 1997).

1.1. Literature

The *calibration problem* has been extensively studied, starting with Dawid (1982), Oakes (1985), and Foster and Vohra (1998); see Olszewski (2015) for a comprehensive survey of the literature. Kakade and Foster (2004) and Foster and Kakade (2006) introduced the notion of *weak calibration*, which shares many properties with smooth calibration. In particular, both can be guaranteed by deterministic procedures, and both are of the "general fixed point" variety: they can find fixed points of arbitrary continuous functions (see for instance the last paragraph in Section 2.3).[4] However, while weak calibration may be at times technically more convenient to work with, smooth calibration is the more natural concept, easier to interpret and understand; it is, after all, just a standard smoothing of regular calibration.

The *online regression problem*—see Section 3 for details—was introduced by Foster (1991); for further improvements, see J. Foster (1999), Vovk (2001), Azoury and Warmuth (2001), and the book of Cesa-Bianchi and Lugosi (2006).

2. Calibration: model and result

In this section we present the calibration game in its standard and "leaky" versions, introduce the notion of smooth calibration, and state our main results.

[1] Corollary 12 in Section 4 will formally show that regular calibration implies smooth calibration.

[2] When the forecasting procedure is deterministic it can be simulated by the agent, and so it is irrelevant whether the agent *observes* the forecasts, or just *computes* them by himself, before taking his action.

[3] Another, seemingly less elegant, way to obtain this is by restarting the procedure once in a while; see, e.g., Lehrer and Solan (2009).

[4] They are thus more "powerful" than the standard calibration procedures (such as those based on Blackwell's approachability), which find *linear* fixed points (such as eigenvectors and invariant probabilities).