

An Easier Way to Calibrate¹

May 16, 1995

revised December 7, 1996

Drew Fudenberg

David K. Levine²

© This document is copyrighted by the authors. You may freely reproduce and distribute it electronically or in print, provided it is distributed in its entirety, including this copyright notice.

¹ The authors are grateful for financial support from NSF grants SBR-9223320, SBR-9223175, SBR-9409180 and the UCLA Academic Senate. This paper benefited from conversations with Glen Ellison and Dean Foster.

²Departments of Economics Harvard and UCLA.

Introduction

Consider a forecaster who each period must predict a probability distribution over an exogenous set of outcomes such as the weather. The forecaster is said to be *calibrated* if in the asymptotic limit, the empirical frequencies in periods where each particular probability distribution is forecast exactly match the forecast probabilities. That is, calibration requires that (asymptotically) it should have rained in exactly 1/3 of the periods in which the forecast probability of rain was 1/3, and so on.³ Recently, Foster and Vohra (1994) proved the existence of random strategies that are approximately calibrated to any desired degree. That is, for any $\varepsilon > 0$, there is a randomized forecast rule that is calibrated with probability $1 - \varepsilon$. Their construction, which relies on a series of approximations, is somewhat complicated; this note provides a shorter and simpler one. Our proof shows that “ K -initialized myopic strategies” are approximately calibrated when K is large. These strategies first “initialize” by making each forecast exactly K times, and thereafter play, in each period t , the minmax strategy in a static game whose payoffs depend on the history until t in a way we define below.⁴

As a caveat to these results, note that calibration is a relatively weak notion: for example, it is a calibrated forecast to predict a 50% chance of heads when facing a deterministic alternation between heads and tails. However, Fudenberg and Levine (1996) show that calibration can be extended to conditional forecasts, provided that the number of events conditioned on does not grow too rapidly with time.⁵

³ Dawid (1982) defines a stronger notion of calibration, requiring that conditional frequency predictions be correct. Dawid (1985) and Oakes (1985) showed that no forecasting rule can be exactly calibrated for all possible sequences. This stronger notion of calibration is shown by Kalai, Lehrer and Smorodinsky (1996) to be equivalent to the more traditional version of merging.

⁴ There is a related decision-theoretic notion of calibration based on actions: a decision rule is calibrated if every action that is used with positive frequency earns nearly the most that was possible against the long-run frequency of opponents’ play in periods that action was used. Rules that are calibrated in this sense regardless of how opponents play are studied by Fudenberg and Levine (1996) and Hart and Mas-Colell (1996). The analysis of Hart and Mas-Colell shows the close connection of calibration to Blackwell’s (1956) notion of approachability.

⁵ See also the previous two footnotes.

The Model

Each period $t = 1, 2, \dots$, the player chooses from a finite set of announcements $a \in A$, then observes an outcome $x \in X$, a finite set. Mixed actions are denoted by $\alpha \in \Delta(A)$. To each announcement a corresponds a vector of probabilities over outcomes of $p(a) \in \Delta(X)$. We denote by λ the greatest distance (in the Euclidean norm) of any probability vector from the set of points $p(A)$. We will also write $e(x)$ for the unit vector that puts weight one on the outcome x . The player's choice in each period t (that is, his forecast rule) can depend only on the information his information at that time, which is the sequence of all previous pairs (a_τ, x_τ) of announcements and outcomes.

Given a history of announcements and outcomes up through and including date t , we define $n_t(a)$ to be the number of times that the action a has been chosen, and $P_t(a)$ to be the vector whose has components are the number of times that each outcome has occurred in the time periods when a was chosen. Thus $P_t(a) / n_t(a) \in \Delta(X)$ represents the empirical distribution of outcomes conditional on the action a .

For any action a we may define a quadratic cost associated with that action.

$$c_t(a) = n_t(a) \left\| \frac{P_t(a)}{n_t(a)} - p(a) \right\|^2$$

This measures the deviation of the empirical frequencies from the forecast probabilities, weighted by the number of observations. The overall cost is just $c_t = \sum_a c_t(a)$. The distribution of this random variable depends upon the strategy of both the player and nature. The idea is to show how to pick α_t so that $E_{t-1}c_t - c_{t-1}$ can be made smaller than any given positive ε , uniformly over histories, and regardless of the strategy chosen by Nature. A standard strong law of large numbers argument then implies that almost surely the average cost c_t / t is bounded asymptotically, and that this bound can be taken to be arbitrarily small. It is also straightforward to show (Foster and Vohra provide a proof)

that the existence of rules with arbitrarily low average cost implies the existence of rules that are approximately calibrated.

Main Result

Notice that given $P_{t-1}(\cdot)$ and $n_{t-1}(\cdot)$, c_t is determined entirely by a_t, x_t . Thus at time t we may take

$$\begin{aligned} g_t(a, x) &= E_{t-1}[c_t - c_{t-1} | a_t = a, x_t = x] = E_{t-1}[c_t | a_t, x_t] - c_{t-1} \\ &= (n_{t-1}(a) + 1) \left\| \frac{P_{t-1}(a) + e(x)}{(n_{t-1}(a) + 1)} - p(a) \right\|^2 - c_{t-1}(a) \end{aligned}$$

to be the payoffs in a zero-sum game in which the player picks a , and the opponent picks x . If $\alpha \in \Delta(A), \gamma \in \Delta(X)$, we let $g_t(\alpha, \gamma)$ denote the expected value of $g_t(a, x)$. Our goal is to prove the existence of a strategy that ensures that this change in cost from period to period is always small.

The particular strategies we examine, called *K-initialized myopic strategies*, have the following special form: First, each pure strategy a is repeated K times. This is called the initialization phase. Following the initialization phase, at time t α_t is chosen to be a maxmin strategy in the time- t game with payoffs g_t . This strategy is characterized by the two parameters, K, λ where we recall that λ is the fineness of the grid $p(A)$.

Theorem: For any ε there exist K, λ so that the K -initial myopic strategy has $\limsup c_t / t \leq \varepsilon$ almost surely.

To prove this, we will show that after the initialization period the value of the period- t games defined by g_t remains uniformly small. Notice that these are not dynamic games. Rather, each g_t corresponds to a ‘‘stage game’’ which is a one-shot simultaneous-move game with fixed strategy spaces A, X and payoffs determined by the initial conditions n_{t-1}, P_{t-1} .

To find the bounds on the payoffs to the g_t , we can use the minmax theorem, which asserts that the player can guarantee himself exactly $\min_{\gamma \in \Delta(X)} \max_{\alpha \in \Delta(A)} g_t(a, \gamma)$ in each period. That is, we can suppose that the opponent's equilibrium strategy $\gamma_t \in \Delta(X)$ is chosen to minimize the player's payoff *on the assumption that the player correctly forecasts γ_t* .⁶

The proof shows that the (infeasible) strategy of first performing the initialization, and thereafter selecting $p(a)$ as close as possible to the actual γ_t , guarantees that g_t does not exceed the specified bound. It follows from the minmax theorem that there is a strategy in the static games that guarantees this result regardless of Nature's play. However, our analysis does not reveal what this strategy is.

Our theorem follows directly from the following lemmas.

Lemma 1: For each a and γ_t ,

$$g_t(a, \gamma_t) = - \left\| \frac{P_{t-1}(a)}{n_{t-1}(a) + 1} - \gamma_t(a) \right\|^2 - \frac{\|P_{t-1}(a)\|^2}{(n_{t-1}(a) + 1)^2 n_{t-1}(a)} + \frac{1}{n_{t-1}(a) + 1} + \|p(a) - \gamma_t(a)\|^2.$$

The proof of this, found in the Appendix, consists of algebraic manipulation.

Lemma 2: If a K -initial myopic strategy is used, then after the initialization phase $g_t \leq \frac{1}{K+1} + \lambda^2$

Proof: . Examining the expression for g_t in lemma 1, we see that the first two terms are always negative. The third term is bounded by $\frac{1}{K+1}$ after the initialization phase, since every action has been played at least this often. Intuitively, the potential period- t increment to the player's average prediction error is largest when there are some actions

⁶ Since we are discussing the stage game, which is finite, the minmax theorem applies immediately. Complications can arise in applying the minmax theorem to infinite games. Note that we are not asserting that the player can correctly forecast the realized value of x_t , but only its distribution.

that the player has only chosen rarely, for then the expected change in the empirical distribution $P_t(a)$ corresponding to a can be large; the initialization phase serves to bound this maximum increment. Finally, the fourth term in the expression is made less than or equal to λ^2 by picking the a which is closest on the grid to γ_t . This gives an upper bound on the value of the g_t game.

☑

Appendix

Lemma 1: For each a and γ_t ,

$$g_t(a, \gamma_t) = - \left\| \frac{P_{t-1}(a)}{n_{t-1}(a) + 1} - \gamma_t(a) \right\|^2 - \frac{\|P_{t-1}(a)\|^2}{(n_{t-1}(a) + 1)^2 n_{t-1}(a)} + \frac{1}{n_{t-1}(a) + 1} + \|p(a) - \gamma_t(a)\|^2.$$

Proof: In what follows we fix γ_t and a pure action a . From the definition of g_t , we write

$$\begin{aligned} g_t(a, \gamma_t) &= \sum_x \gamma_t(x) \left((n_{t-1} + 1) \left\| \frac{P_{t-1} + e(x)}{n_{t-1} + 1} - p \right\|^2 - c_{t-1} \right) \\ &= \left(\sum_x (n_{t-1} + 1) \left\| \frac{P_{t-1} + e(x)}{n_{t-1} + 1} - p \right\|^2 \gamma_t(x) \right) - n_{t-1} \left\| \frac{P_{t-1}}{n_{t-1}} - p \right\|^2 \end{aligned}$$

where for notational simplicity we have suppressed the dependence of n, P, p on a , which is being held fixed.

First we write out the squares:

$$\begin{aligned}
g_t(a, \gamma_t) &= \left(\sum_x (n_{t-1} + 1) \left\| \frac{P_{t-1} + e(x)}{n_{t-1} + 1} - p \right\|^2 \gamma_t(x) \right) - n_{t-1} \left\| \frac{P_{t-1}}{n_{t-1}} - p \right\|^2 \\
&= \left(\sum_x (n_{t-1} + 1) \left\| \frac{P_{t-1} + e(x)}{n_{t-1} + 1} \right\|^2 \gamma_t(x) \right) - 2 \left(\sum_x p \cdot (P_{t-1} + e(x)) \gamma_t(x) \right) + \sum_x \|p\|^2 \cdot \gamma_t(x) (n_{t-1} + 1) \\
&\quad - n_{t-1} \left\| \frac{P_{t-1}}{n_{t-1}} \right\|^2 + 2P_{t-1} \cdot p - n_{t-1} \|p\|^2
\end{aligned}$$

Next we apply $\sum_x (y \cdot e(x)) \gamma_t(x) = y \cdot \gamma_t$ for $y = p$ and $\sum_y \gamma_t(y) = 1$, $\|e(y)\|$ to find

$$\begin{aligned}
g_t(a, \gamma_t) &= \sum_x (n_{t-1} + 1) \left\| \frac{P_{t-1} + e(x)}{n_{t-1} + 1} \right\|^2 \gamma_t(x) - 2(p \cdot P_{t-1}) n_{t-1} - 2(p \cdot \gamma_t) + \|p\|^2 (n_{t-1} + 1) \\
&\quad - n_{t-1} \left\| \frac{P_{t-1}}{n_{t-1}} \right\|^2 + 2P_{t-1} \cdot p - n_{t-1} \|p\|^2
\end{aligned}$$

Examining the expression

$$\sum_x (n_{t-1} + 1) \left\| \frac{P_{t-1} + e(x)}{n_{t-1} + 1} \right\|^2 \gamma_t(x) - n_{t-1} \left\| \frac{P_{t-1}}{n_{t-1}} \right\|^2 - \|\gamma_t\|^2 + \|p - \gamma_t\|^2$$

we may check, by multiplying out the final square, that it is equal to the previous expression, and therefore to $g_t(a, \gamma_t)$. Multiplying out the first square, canceling n_{t-1} in the second term, and applying $\|e(y)\|^2 = 1$ leads to

$$g_t(a, \gamma_t) = \sum_x \frac{\|P_{t-1}\|^2 + 2P_{t-1} \cdot e(x) + 1}{n_{t-1} + 1} \gamma_t(x) - \frac{\|P_{t-1}\|^2}{n_{t-1}} - \|\gamma_t\|^2 + \|p - \gamma_t\|^2.$$

Again apply $\sum_x (y \cdot e(x)) \gamma_t(x) = y \cdot \gamma_t$, this time for $y = P_{t-1}$ and $\sum_y \gamma_t(y) = 1$ and combine terms to find

$$\begin{aligned}
g_t(a, \gamma_t) &= \frac{\|P_{t-1}\|^2}{n_{t-1} + 1} - \frac{\|P_{t-1}\|^2}{n_{t-1}} + \frac{2P_{t-1} \cdot \gamma_t}{n_{t-1} + 1} - \|\gamma_t\|^2 + \frac{1}{n_{t-1} + 1} + \|p - \gamma_t\|^2 = \\
&\quad - \left\| \frac{P_{t-1}}{n_{t-1} + 1} - \gamma_t \right\|^2 - \frac{\|P_{t-1}\|^2}{(n_{t-1} + 1)^2 n_{t-1}} + \frac{1}{n_{t-1} + 1} + \|p - \gamma_t\|^2. \quad \square
\end{aligned}$$

References

- Blackwell, D. (1956): "Controlled Random Walks," *Proceedings International Congress of Mathematicians*, III 336-338, Amsterdam, North Holland.
- Dawid, A.P.(1982): "The Well-Calibrated Bayesian," *Journal of the American Statistical Association*, 77, 605-613.
- Dawid, A.P. (1985): "The Impossibility of Inductive Inference," *Journal of the American Statistical Association*, 80, 340-341.
- Foster, D. and R. Vohra (1994) "Asymptotic Calibration," mimeo, Wharton.
- Fudenberg, D. and D.K. Levine (1995): "Universal Consistency and Cautious Fictitious Play," *Journal of Economic Dynamics and Control*, 19, 1065-1090. .
- Fudenberg, D. and D.K. Levine (1996): "Conditional Universal Consistency," mimeo, UCLA.
- Hart, S. and A. Mas-Colell (1996): "A Simple Adaptive Procedure Leading to Correlated Equilibrium," mimeo, Hebrew University of Jerusalem.
- Kalai, E., E. Lehrer and R. Smorodinsky (1996): "Calibrated Forecasting and Merging," Northwestern MEDS Discussion Paper 1144.
- Oakes, David (1985) "Self-Calibrating Priors Do not Exist," *Journal of the American Statistical Association*, 80.