Discussion Paper No. 1254

Calibration, Expected Utility and Local Optimality

by

Dean P. Foster

and

Rakesh V. Vohra

March 1999

Math Center web site:
http://www.kellogg.nwu.edu/research/math

Calibration, Expected Utility and Local Optimality

Dean P. Foster * Rakesh V. Vohra †

March 17, 1999

Abstract

We propose a framework for reconciling frequentist and subjectivist views of probability. In an environment with repeated trials we show that beliefs about the possible states of nature can be represented by probabilities. Second, these probabilities will correspond to long run frequencies. In particular they will be naively calibrated. Third, the actions chosen in each trial will be the ones that maximize expected utility on that trial. The expectation is with respect to the probabilities used to represent beliefs.

1 Introduction

There are many different interpretations of what a probability is, but only two command significant attention. The more venerable, called **frequentist**, holds that a probability is a long run frequency. The younger, called **subjectivist**, contends it is a matter of personal opinion or a manifestation

*Department of Statistics, The Wharton School, University of Pennsylvannia, Philadelphia, PA 19104

†Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208.

of ones preferences. Each claims a celebrated champion Fisher for the frequentist's, Savage for the subjectivists. Neither, it seems, reconcilable with the other.

The subjectivist perspective is perhaps the dominant one amongst scholars, but it is the frequentist view that plays in classroom and schoolbook. How can two such very different opinions of probability co-exist?

The frequentist view has the advantage that frequencies naturally satisfy the laws of probability. Nevertheless, it is not clear that anything resembling the repeated *independent* trials under *identical conditions* that are required to obtain such a frequency exist. Even if it did, the frequentist school, must, by definition, be silent about the uncertainty associated with a single trial.

Subjectivist's are not bound to the procrustean bed of infinitely repeated trials. However, it is not obvious that ones beliefs about the underlying uncertainty are best or even naturally modeled by probabilities. If probabilities are merely a device for representing personal preferences, then, there is no reason at all to expect them to coincide with frequencies of any kind.

In this paper we propose a framework for reconciling the two distinct schools of thought. We consider an environment with repeated but not necessarily identical trials. In each trial a subject must make a choice whose payoff depends on that choice as well as the realized state of nature. The subject has a preference for high (undiscounted) total payoffs. The subject's choices in each round are consistent with both their preferences as well as their beliefs about which states will be realized. Our first result is that the subject's beliefs about the possible states of nature can be represented by probabilities.

Second, these probabilities will have the property of being **calibrated**. Calibration is a measure of the closeness between a sequence of probability forecasts of an event and the events observed frequency. It accommodates the fact that the specified probability of the event may *change* from period to period. This would happen, for example, if the underlying distribution generating the event is not stationary.

An informal explanation of calibration might go something like this. Suppose each night we predict the chance of rain for the next day, say we announce a probability p of rain. On the subsequence of days that we made a forecast of p, compute the fraction of times it actually rained. Call this fraction $\rho(p)$. An obvious measure of the biasedness of our forecast is $\rho(p) - p$. Calibration, requires, roughly, that $\rho(p) - p$ be 0 for all p. That is the forecasts made are unbiased. The notion can be formalized in a number of different ways and we offer one later. Strengthenings of the calibration notion are described by Kalai, Lehrer and Smorodinsky [6].

Our third result shows that the choices made in each trial will be the ones that maximize expected utility on that trial. The expectation is with respect to the probabilities used to represent beliefs.

Our fourth result can be viewed as a justification of the notion of calibration itself. One reason for why probabilistic beliefs should be calibrated is offered by Schervish [9]. He shows that a forecast that is not calibrated can be 'rounded' into a calibrated forecast without losing any of the discriminatory power of the original. Now consider the payoffs that are obtained if one chooses at each period the action that maximizes payoffs with respect to the uncalibrated forecast. Compare this with the same for the calibrated counterpart. The time average of the payoffs generated by using the original forecast will be no larger than those those generated by its calibrated counterpart. Schervish notes however that the argument leaves something to be desired.[1] The calibrated counterpart is constructed after the fact, that is

[1]See Seidenfeld [10] as well.

after the original forecast and realization are known.

Dawid [2] establishes that calibration is a consequence of the Bayesian perspective.[2] Specifically, the posterior forecasts of a coherent Bayesian must be calibrated with respect to all sequences except those that occur with probability zero. The underlying measure is the one generated by the prior of the Bayesian forecaster. Thus the Bayesian forecaster will be calibrated but not on things that are deemed impossible by her philosophy!

In this paper we show that if a decision maker chooses in each period actions that maximize their expected utility and satisfy one other condition, called local optimality, then the probabilities used must be calibrated. The additional condition, local optimality, is our formulation of the idea that the decision maker has a preference for high aggregate utilities.

The next section introduces the set up and notation. A subsequent section provides the relevant mathematical arguments to support our claims. The last section is devoted to a discussion of the results.

2 The Set Up

There is a finite state space S, each of whose states is indexed by the integers $\{1, 2, 3, \ldots, n\}$. The decision makers beliefs will be modeled by a finite set $B = \{1, 2, 3, \ldots, b_{max}\}$. Call this the **belief set** of the decision maker.[3] A **bundle** of lotteries is a finite set of lotteries defined over S. It can be represented by a $m \times n$ matrix, L, where m is the number of lotteries in the bundle and the $\{ij\}^{th}$ entry, call it L_{ij}, is the utility to the decision maker

[2] Apparently much the same result was established by J. W. Pratt in an unpublished paper. This is mentioned in [10].

[3] The restriction to a finite state space and belief set, is for convenience only. The results can be extended to more general state and belief sets. The reader is referred to a companion paper, [5] for details.

from choosing lottery i in state j. So as to distinguish between an individual lottery and a bundle of lotteries we refer to an individual lottery as a **ticket**. Thus each row of L is a ticket and each column corresponds to a state of nature.

A **selection rule** is a function, f, that associates with each bundle L and each belief in B a unique ticket that should be selected from L. So, if L is an $m \times n$ matrix representing a bundle, and $b \in B$, then $f(L, b)$ is the index of one of the rows of L. A decision maker using selection rule f when faced with a choice from L first decides which belief in B is salient and then uses f to make a choice. Notice we do not specify how the decision maker decides what is a salient belief or how these will be modified with experience. We require that salience be a quality independent of L as well as the selection rule used. In other words, the beliefs are about states of S alone.

2.1 Local Optimality

In subjectivist developments it is common to specify the value that an acceptable selection rule must have for certain bundle and belief combinations. Ramsey [7], for example, has his ethically neutral propositions, de Finetti [3] his dutch book and Savage [8] his sure thing. We avoid doing this for two reasons. First, we wish to say as little as possible about the nature of the beliefs held by the decision maker. Second, in an environment with repeated trials what matters is the aggregate outcome of all choices made rather than any one choice. Instead we assume that our decision maker has a preference for selection rules that generate high aggregate utilities. This is made precise below.

Suppose at each time t, our decision maker must choose one ticket from

L and then obtains the relevant payoff. Let $L_{f(L,b^t),s^t}$ denote the utility achieved by this choice if belief $b_t \in B$ is salient and state $s_t \in S$ is realized in time t. We assume the decision maker has a preference for selection rules f that maximize the (undiscounted) time average of utilities, i.e:

$$\lim_{T\to\infty} \frac{\sum_{t=1}^{T} L_{f(L,b^t),s^t}}{T}.$$

In fact the limit above could be replaced with lim sup and nothing in the sequel is affected. Call such a rule **globally optimal**. Since neither we nor the decision maker knows what a globally optimal selection rule is, we settle for something weaker that we call **local optimality**.

Two distinct selection rules, f and g are said to be **neighbors** of each other if there is exactly one belief $k \in B$ and bundle M such that $f(L,r) = g(L,r)$ for all $(L,r) \neq (M,k)$. Selection rule f is **locally optimal** with respect to a realization of states if for all neighboring selection rules g and *all* bundles L defined over S we have

$$\lim_{T\to\infty} \inf \frac{\sum_{t=1}^{T} (L_{f(L,b^t),s^t} - L_{g(L,b^t),s^t})}{T} \geq 0.$$

The set of neighbors of a selection rule f is quite limited. It consists only of those rules that differ from f in one place. Thus local optimality is a mild condition. Notice that any globally optimal selection rule must be locally optimal.

2.2 Naive Calibration

The formalization of calibration we use is called **naive calibration** in [6]. To define it let X_t be a 0-1 vector in $R^{|S|}$ that indicates which state in S was realized at time t. So, if state j was realized, then $X_{t,j}$, the j^{th} of component of X_t, is 1 and the rest are all zero.

Let p_t be a probability forecast of X_t (made of course, before X_t is realized). It is a vector whose j^{th} component is a forecast of the probability that state $j \in S$ will be realized in period t. Denote by $N(p,t)$ the number of periods up to the t-th that a vector of forecasts equal to p was generated. Let $\rho(p,j,t)$ be the fraction of these periods for which state $j \in S$ was realized, i.e.,

$$\rho(p,j,t) = \begin{cases} 0 & \text{if } N(p,t) = 0 \\ \sum_{s=1}^{t} \frac{I_{p_s=p} X_{s,j}}{N(p,t)} & \text{otherwise} \end{cases}$$

The probability forecast is said to be **naively calibrated with respect to the sequence of realized states** if:

$$\lim_{T\to\infty} \sum_{j\in S} \sum_{\{p:p_t=p,t\leq T\}} (\rho(p,j,T) - p_j)^2 \frac{N(p,T)}{T} = 0.$$

The term

$$\sum_{j\in S} \sum_{\{p:p_t=p,t\leq T\}} (\rho(p,j,T) - p_j)^2 \frac{N(p,T)}{T}$$

is sometimes called the calibration component of the Brier score. See Blattenberger and Lad [1] for an exposition.

3 The Main Result

For each ticket r in L let Q^r be the set of probability vectors p over S such that

$$\sum_{j\in S} p_j L_{rj} \geq \sum_{j} p_j L_{ij}$$

for all tickets $i \neq r$. The diameter of Q^r will be the diameter of the smallest sphere to circumscribe Q^r. Denote the diameter of each Q^r by $diam(Q^r)$ and let $diam(L) = \max_r diam(Q^r)$.

Theorem 1 *Let $\{s_t\}_{t\geq 1}$ be a sequence of realized states and $\{b_t\}_{t\geq 1}$ the associated sequence of salient beliefs. Suppose the selection rule f is locally optimal on this sequence. Then there is a function c that associates with each $b \in B$ a probability vector over S such that*

$$f(L,b) = \arg\max_i \sum_{j=1}^{n} c_j(b) L_{ij}$$

for all b that are salient for a non-vanishing fraction of times and

$$\lim_{T\to\infty} \sum_{j\in S} \sum_{\{p:c(b_t)=p,t\leq T\}} (\rho(p,j,T) - p_j)^2 \frac{N(p,T)}{T} \leq (9/4) diam(L)^2.$$

Proof

Consider all the times up to T that belief b was salient, i.e., $\{t \leq T : b_t = b\}$. Let $n(b,T) = |\{t \leq T : b_t = b\}|$. Let r be the unique ticket such that $f(L,b) = r$. Without loss of generality we may assume that $n(b,T)/T$ does not vanish as $T \to \infty$. By local optimality we may assume that r is not dominated by any other lottery in L.

Let

$$p^b(T) = \frac{\sum_{\{t\leq T:b_t=b\}} X_t}{n(b,T)}$$

where X_t is a 0-1 vector in $R^{|S|}$ that indicates which state in S was realized at time t. So, if state j was realized, the j^{th} of component of X_t is 1 and the rest are all zero. Let A^b be the set of accumulation points of the sequence $\{p^b(t)\}$. We claim that $A^b \subset Q^r$. Suppose not. Then there is a point $q \in A$ such that $q \not\in Q^r$. Let $B_\epsilon(q)$ be an ϵ ball around q. For ϵ sufficiently small, $B_\epsilon(q) \cap Q^r = \emptyset$ and there is a ticket k such that

$$\sum_{j\in S} p_j L_{kj} > \sum_{j\in S} p_j L_{rj}$$

for all $p \in B_\epsilon(q)$.

Let g be the selection rule that is identical to f except when f picks r, it picks k. Consider any $t' \in \{t : p^b(t) \in B_\epsilon(q)\}$. The fact that q is an accumulation point of the sequence $\{p^b(t)\}$ implies that there are an infinite number of such t''s. Then

$$\frac{\sum_{t \leq t'}[L_{f(L,b^t),s^t} - L_{g(L,b^t),s^t}]}{t'} = \frac{\sum_{\{t \leq t' : b^t = b\}}[L_{r,s^t} - L_{k,s^t}]}{t'}$$

$$= \sum_{j \in S} p_j^b(t')[L_{rj} - L_{kj}] < 0.$$

Since there are an infinite number of such t''s, this violates the local optimality of f. Thus we may assume that $A^b \subset Q^r$.

For each Q^r let q^r be any point in Q^r such that $d(q^r, q) \leq \frac{diam(Q^r)}{2}$ for all $q \in Q^r$ where d is Euclidean distance. Define $c(b_t) = q^r$ if $f(L, b_t) = r$. For beliefs b that were salient for a vanishing proportion of times, choose $c(b)$ arbitrarily. The first part of the theorem is obvious given the definition of $c(b)$. We now prove the second.

Now

$$\sum_{j \in S} \sum_{\{p : c(b_t) = p, t \leq T\}} (\rho(p, j, T) - p_j)^2 \frac{N(p, T)}{T} = \sum_{j \in S} \sum_{r \in S} (\rho(q^r, j, T) - q_j^r)^2 \frac{N(q^r, T)}{T}.$$

Notice that $n(b, T) = N(q^r, T)$ and $\rho(q^r, j, T) = p_j^b(T)$ for the $b \in B$ such that $f(L, b) = r$. Thus

$$\sum_{j \in S} \sum_{r \in S} (\rho(q^r, j, T) - q_j^r)^2 \frac{N(q^r, T)}{T} = \sum_{j \in S} \sum_{b \in B} (p_j^b(T) - q_j^{f(L,b)})^2 \frac{n(b, T)}{T}.$$

Also,

$$\sum_{j \in S} \sum_{b \in B} (p_j^b(T) - q_j^{f(L,b)})^2 \frac{n(b, T)}{T} = \sum_{b \in B} \frac{n(b, T)}{T} d(p^b(T), q^{f(L,b)})^2.$$

In the last summation we can ignore terms that correspond to b's salient for a vanishing proportion of times.

By the triangle inequality

$$d(p^b(T), q^{f(L,b)}) \leq d(p^b(T), q) + d(q, q^{f(L,b)})$$

for any $q \in A^b$. Since A^b is the set of accumulation points of the sequence $\{p^b(T)\}$ there is a T sufficiently large we can always choose $q \in A^b$ such that

$$d(p^b(T), q) \leq diam(Q^{f(L,b)}).$$

Hence

$$d(p^b(T), q^{f(L,b)}) \leq (3/2) diam(Q^{f(L,b)}).$$

Thus

$$\sum_{b \in B} \frac{n(b,T)}{T} d(p^b(T), q^{f(L,b)})^2 \leq \sum_{b \in B} \frac{n(b,T)}{T} (9/4) diam(Q^{f(L,b)})^2.$$

The last term is of course$\leq (9/4) diam(L)^2$.□

The first part of Theorem 1 says little more than that any sequence of choices (under mild consistency conditions) can be supported by appropriate probabilities. The content of Theorem 1 is that these probabilities can be chosen so as to approximate, in some sense, the realized frequencies. The next theorem shows what happens when we exploit the fact that local optimality must hold for *all* bundles defined over the state space S. The basic idea is to choose a bundle L with the property that $diam(L)$ is arbitrarily small.

Theorem 2 *Let $\{s_t\}_{t \geq 1}$ be a sequence of realized states and $\{b_t\}_{t \geq 1}$ the associated sequence of salient beliefs. Suppose the selection rule f is locally optimal on this sequence. Then there is a function c that associates with each $b \in B$ a probability vector over S such that*

$$f(L, b) = \arg\max_i \sum_{j=1}^{n} c_j(b) L_{ij}$$

for all bundles L defined over S and beliefs b that are salient for a non-vanishingly small proportion of times. Furthermore

$$\lim_{T\to\infty}\sum_{j\in S}\sum_{\{p:c(b_t)=p,t\leq T\}}(\rho(p,j,T)-p_j)^2\frac{N(p,T)}{T}=0.$$

Proof

Let $\{L^i\}$ be a sequence of bundles over S. The i^{th} bundle has at least one more ticket than the $i-1$ first bundle. For each L^i let Q_i^r be the set of probability vectors over S for which ticket r in L^i has largest expected payoff. Call the sequence $\{L^i\}$ decreasing if for any $\epsilon > 0$, there is an i sufficiently large such that $diam(L^i) \leq \epsilon$.[4]

For each Q_i^r let $q^r(i)$ be the point in Q_i^r such that $d(q^r(i), q) \leq diam(Q_i^r)/2$ for all $q \in Q_i^r$. Define $c^i(b) = q^r(i)$ if $f(L^i, b) = r$. By Theorem 1 it follows that for T sufficiently large

$$\sum_{j\in S}\sum_{\{p:c^i(b_t)=p,t\leq T\}}(\rho(p,j,T)-p_j)^2\frac{N(p,T)}{T}\leq\epsilon.$$

Let $c(b)$ be an element of the set of limit points of the sequence $\{c^i(b)\}_{i\geq 1}$. For any $\epsilon > 0$ there is an i sufficiently large such that $d(c(b), c^i(b)) < \epsilon$ and $c(b) \in Q_i^r$ where $c^i(b) = q^r(i)$. By the triangle inequality it follows that for any probability vector ρ over S

$$d(c(b),\rho)\leq\epsilon+d(c^i(b),\rho).$$

Squaring both sides and using the fact that $d(x, y) \leq 2$ for any two probability vectors x and y we get

$$d(c(b),\rho)^2\leq 5\epsilon+d(c^i(b),\rho)^2.$$

[4]Existence of such is easy but tedious to establish. A proof is available upon request.

Now

$$\sum_{j\in S}\sum_{\{p:c(b_t)=p,t\leq T\}}(\rho(p,j,T)-p_j)^2\frac{N(p,T)}{T}=\sum_{\{p:c(b_t)=p,t\leq T\}}\frac{N(p,T)}{T}d(\rho(p),p)^2.$$

But the right hand side of the above is bounded above by

$$\leq 5\epsilon+\sum_{j\in S}\sum_{\{p:c^i(b_t)=p,t\leq T\}}(\rho(p,j,T)-p_j)^2\frac{N(p,T)}{T}\leq 6\epsilon.$$

Thus we may conclude that

$$\lim_{T\to\infty}\sum_{j\in S}\sum_{\{p:c(b_t)=p,t\leq T\}}(\rho(p,j,T)-p_j)^2\frac{N(p,T)}{T}=0.$$

Now suppose there is a bundle M, belief a and tickets k and h such that

$$k=f(M,a)\neq\arg\max_i\sum_{j=1}^{n}c_j(a)M_{ij}=h.$$

Let g be the selection procedure such that

$$g(M,b)=f(M,b)\ \forall b\in B\setminus a$$

and $g(M,a)=h$ otherwise. By local optimality of f we have for T sufficiently large that

$$\frac{\sum_{t\leq T}[M_{f(M,b^t),s^t}-M_{g(M,b^t),s^t}]}{T}=\sum_{j\in S}p_j^a(T)[M_{kj}-M_{hj}]>0.$$

Since $c(a)$ is a calibrated forecast it follows that $\lim_{T\to\infty}p^a(T)=c(a)$. Hence for T sufficiently large it follows by local optimality and calibration that

$$\sum_{j\in S}c_j(a)[M_{kj}-M_{hj}]>0$$

a contradiction.□

Theorem 2 shows that given local optimality one can represent B using probabilities that are naively calibrated. This does not preclude probability representations that are not naively calibrated. However, if we insist on local optimality and expected utility maximization, then the only probability representation consistent with these desiderata are those that are naively calibrated.

Theorem 3 *Let c be a function that associates with each $b \in B$ a probability vector of S. Let f be the selection rule defined as follows:*

$$f(L, b) = \arg\max_i \sum_{j=1}^{n} c_j(b) L_{ij}.$$

Suppose f is locally optimal on the sequence $\{s_t\}_{t \geq 1}$ of realized states. If $\{b_t\}_{t \geq 1}$ is the associated sequence of salient beliefs then $c(b_t)$ is naively calibrated.

Proof

Suppose not. Then there is a belief h salient for a non-vanishing proportion of times such that

$$\sum_{j \in S} [\rho(c(h), j, T) - c_j(h)]^2 \leq \epsilon$$

for some positive ϵ and infinitely many T. Fix one such T. Notice that $c(h)$ and $\rho(c(h), j, T)$ are distinct points and so can be separated by a hyperplane.

Let $T(h) = \{t \leq T : b_t = h\}$. Let $r = f(L, h)$ and

$$k = \arg\max_i \sum_{j \in S} \rho(c(h), j, T) L_{ij}.$$

Define g to be the selection rule such that

$$g(L, b) = f(L, b) \ \forall b \in B \setminus h$$

and

$$g(L, h) = k$$

otherwise.

We now construct a lottery bundle M over S consisting of two tickets such that

$$\sum_{j \in S} [M_{1j} - M_{2j}] c_j(h) > 0$$

and

$$\sum_{j \in S} [M_{2j} - M_{1j}] \rho(c(h), j, T) > 0.$$

This can be done by choosing $M_{1.} - M_{2.}$ to be any one of the hyper-planes that separates $c(h)$ from $\rho(c(h), j, T)$.

Observe that $f(M, h) = 1$ and $g(M, h) = 2$. Then

$$\frac{\sum_{t \leq T} [M_{f(M,b^t),s^t} - M_{g(M,b^t),s^t}]}{T} = \frac{\sum_{t \in T(h)} [M_{1,s^t} - M_{2,s^t}]}{T}$$

$$= \sum_{j \in S} \rho(c(h), j, T)[M_{1j} - M_{2j}] < 0$$

which contradicts the local optimality of f.□

4 Discussion

The analysis above is predicated on the assumption that the decision maker cares about making choices that generate high total utility. An evolutionary argument can be offered to justify this, we don't. Our point is that if one accepts this then the requirement of local optimality is not unreasonable. From this it follows that any way of representing beliefs, intervals, propensities, words, images can be interpreted as probabilities. If the decision maker acts *as if* her beliefs about states are probabilistic it is not a giant leap to assert that her beliefs *are* probabilities. In this light the subjectivist view is very natural. So is the frequentist view, but not so directly—it relies on how calibration works. The idea behind calibration is to break an arbitrary sequence of events up into a collection of subsequences. Most of the subsequences (at least those that continue long enough and happen often enough) have limiting frequencies. These are exactly the sequences that the frequentist view requires to define probabilities. The only difference is that the subjectivist's would focus on cross sectional issues (consistency, dutch books, etc) while frequentist's would focus on long-run issues (limiting frequency, CLT, Law of large numbers, etc).

5 Acknowledgements

The research of the second author was supported in part by a Dean's summer fellowship from the Fisher College of Business of the Ohio State University. We thank Peter Klibanoff, Teo Chung Piaw, Alvaro Sandroni and H. Peyton Young for some useful comments.

References

[1] Blattenberger, G. and Lad, F. (1985). "Separating the Brier score into calibration and refinement components: A Graphical Exposition", *The American Statistician*, 39, 26 - 32.

[2] Dawid, A. P. (1982). "The well calibrated Bayesian", *J. of The Am. Stat. Assoc.*, 77, 605-613.

[3] de Finetti, B. 1972. *Probability, Induction and Statistcis*, New York: John Wiley and Sons.

[4] Foster, D. P., and Vohra, R. (1991). "Asymptotic Calibration",*Biometrika*, vol. 85, 2, 379-390.

[5] Foster, D. P. and Vohra, R. in preparation.

[6] Kalai, E., Lehrer, E. and Smorodinsky, R. (1996). "Calibrated Forecasting and Merging", *Games and Economic Behavior*, forthcoming.

[7] Ramsey, F. P. (1931) "Truth and Probability", in *The Foundations of Mathematics and other Logical Essays*, ed. R. B. Braithwaite, New York: Harcourt Brace.

[8] Savage, L. E. J. (1954). *The Foundation of Statistics*, New York: John Wiley and Sons.

[9] Schervish, M. (1989). "A general method for comparing Probability Assessors", The Annals of Statistics, vol. 17,4, 1856-1879.

[10] Seidenfeld, T. (1985) "Calibration, Coherence, and Scoring Rules", *Philosophy of Science*, 52, 274-294.