Aust. N. Z. J. Stat. **47**(4), 2005, 385–399

ANOMALIES, PROBINGS, INSIGHTS: KEN FOREMAN'S ROLE IN THE SAMPLING INFERENCE CONTROVERSY OF THE LATE 20TH CENTURY

KEN BREWER[1]

The Australian National University

Summary

This paper is based on the author's recollections of work done for and with E.K. Foreman, and later developments of that work. It describes most particularly the author's early intimations of the potential for using prediction models in survey sampling, the reasons why the use of such models, especially for inferential purposes, was so strongly resisted in the second half of the 20th century, the manner in which randomization and prediction inferences can be combined, and some advantages flowing from that combination.

Key words: design-based; inference; model-based; prediction; randomization; survey sampling.

1. Introduction

Some time last year I was reading the first review of my recently published book (Brewer, 2002). One of that book's principal characters had been modelled, at first unconsciously and later quite consciously, on Ken Foreman. The reviewer, herself a highly regarded survey statistician, described that character as 'an exacting but kindly supervisor' (Thompson, 2003). I was delighted, because that was a precise description of the man I knew. Admittedly at one time I might have described him as 'kindly but exacting', and occasionally even as 'infuriatingly exacting', but that would only have been because I was taking a long time to grow up and had no idea how fortunate a young man I then was.

Ken had many virtues, but one of the greatest was that he recognized the potential in every member of his staff, and always gave us opportunities to develop it. Many of those staff have become quite distinguished. Two of them are the present Australian Statistician and his immediate predecessor. Others have become Professors of Statistics. One of these, Warren Ewens, now at the University of Pennsylvania, was in 2002 awarded the University of Oxford's Weldon Memorial Prize for his work in Biometric Science. That may not mean much to many of my readers, but two of the previous recipients were J.B.S. Haldane and Sir Ronald Fisher. It is probably as close to a Nobel Prize as anyone can get for achievements in statistics. There are other prominent protégés of Ken's that I could name, but those three suffice to make my point.

Received November 2003; revised March 2004; accepted June 2004.
[1] School of Finance and Applied Statistics, Faculty of Economics and Commerce, LF Crisp Building (Building No. 26), The Australian National University, ACT 0200, Australia.
e-mail: Ken.Brewer@anu.edu.au
Acknowledgments. The author thanks Dr P.S. Kott for numerous invaluable suggestions which improved this paper. He also thanks the Australian Bureau of Statistics and the Canberra branch of the Statistical Society of Australia for their invitation to deliver the 2003 Foreman Lecture, from which this paper is derived. Some closely related material has also been included towards the end.

© 2005 Australian Statistical Publishing Association Inc. Published by Blackwell Publishing Asia Pty Ltd.

Now although I come some distance down the totem pole of achieving Foreman protégés, in this paper I mainly draw on my own career, partly because that is the one I am most familiar with, and mainly because it was so strongly influenced by Ken, both directly and indirectly, that my own history is, to some extent, his as well. First, however, let me recall a little anecdote about Ken's career in the Bureau before I joined his staff. I guess this incident would have occurred *circa* 1952. He had just come back from an extended period in the US Bureau of the Census, working under Morris Hansen and Bill Hurwitz, who were two of the half dozen or so brightest shining stars in survey research at the time. He straightaway set about designing appropriate samples for those of the Bureau's collections as stood most in need of them.

Before very long, however, he was summoned to the presence of Mr (later Sir) Stanley Roy Carver. At that time Mr Carver was both New South Wales Statistician and Acting Commonwealth Statistician, and therefore perpetually engaged in the curious task of writing letters to himself. He was definitely not in the habit, however, of interviewing subordinates below the rank that we might now describe as EL2. Ken was well below that rank. He would have been a Research Officer Grade 2, but almost certainly nothing more. Mr Carver asked him to recount what he had been doing since his return from the United States, and Ken, no doubt, gave him an enthusiastic description of all the resources that were being saved by converting several of the Bureau's collections from complete enumerations to sample surveys.

I understand that Mr Carver listened to his account in stony silence, cross-examined him at some length, and then sighed and said: 'I've spent my entire career getting rid of errors from statistics; and here you are, putting them back in again!' Ken would have had a ready answer, I imagine, along the lines that the sample errors he was putting in were substantially smaller than the non-sample errors he was taking out, but it still must have been a rather trying occasion for him. He later wrote, in the context of the discovery of serious discrepancies between figures in the then infant Survey of Retail Establishments and corresponding figures in the 1952–1953 Census of Retail Establishments: 'Mr Carver, who in any case distrusted sampling methods, concluded them to be unreliable and wished to discontinue all sampling operations' (Foreman, 1977).

2. Some early anomalies

I turn now to Ken's history of accomplishments as a supervisor, and, being highly selective, to events that led to a resolution of the controversy over design-based versus model-based sampling inference.

Quite early on, Ken rather trustingly handed me responsibility for the sample redesign of the Bureau's three business surveys: Capital Expenditure, Stocks, and Labour Turnover. These collections had actually been on a sample basis from 1947, even before Ken joined the Bureau. They had been allowed to be conducted that way because the survey estimates were not directly published by the Commonwealth Statistician, and were used only in the preparation of National Accounts, where they were aggregated with many other 'guesstimates' that were the best that could be put together in those early days.

This, then, was the first redesign of those surveys since Ken had returned from the US Bureau of the Census in 1951 with the 'gospel of survey sampling' according to Hansen and Hurwitz. I gave all three the standard treatment, stratifying the businesses by industry and size, and using wages (at that time the only available supplementary variable) for ratio-estimation purposes. I also estimated the stratum variances from past surveys, and used Neyman allocation to distribute the sample units among strata in an optimal fashion.

© 2005 Australian Statistical Publishing Association Inc.

In the course of that work, I came across three industries where subgroups of sample units had unusually high and unstable ratios of survey item values to wages. In the Capital Expenditure Survey, the industry called 'Other' contained such a subgroup. Up till then, not much attention had been paid to whether an observation was extreme, and many real extremes were being ignored, but when one came up that was extraordinarily extreme, the standard procedure was to delete it from both the sample and the sample frame, and add it on at the end, which gave it a sample weight of 1.

I did not feel comfortable with either of those options in this situation. The subgroup was recognizable and relevant, so could not be ignored, but its members were not all that exceptional, and it would have been incurring a large downward bias to give them each a sample weight of 1. So the National Accounts people checked the names in the framework list and found that they were all financial institutions (banks and insurance companies, mainly). That made good sense because, with the amounts of money that they were handling, they would have been able to undertake substantially larger capital expenditures than businesses that had similar payrolls but were engaged in other activities. So I created a separate industry for Financial Institutions and they were completely enumerated.

In the Stocks Survey there was a similar subgroup in the Food and Drink industry that had high ratios of stocks to wages. Their names were checked too, and they were all found to be wineries and distilleries, holding large amounts of maturing stocks as 'work in progress'. Same problem, same solution. There was also a subgroup of wholesalers in the Stocks Survey that had high ratios of stocks to wages. It looked like the same problem again. The wholesalers in question were wool brokers, so all the wool brokers in the sample frame were identified with the intention of constituting them into a separate industry.

Then it was discovered that most of the wool brokers in the sample itself did not have high ratios of stocks to wages after all. This anomaly took a little longer to sort out, but it was worth it. We found that while a small minority of wool brokers bought their clients' wool and resold it, most wool brokers received and stored wool and sold it *on behalf of* their clients. Consequently the large stocks physically held by most wool brokers did not need to appear in the survey returns; it was not their own wool. The solution was therefore to separate out and completely enumerate the group of wool-*buying* brokers. But that meant contacting every wool broker in Australia, to determine whether they actually bought their clients' wool, or only stored it for them. At that time, if I remember rightly, National Accounts just rang up every wool broker in the framework and asked them whether they bought their clients' wool. These days, with respondent burden a more serious issue, that way of going about things may not be possible. What could be done instead I am not certain, but I do know that knowledge of how potential respondents organize their businesses is a very useful thing to have, and that where there's a will there's usually a way. To ensure efficiency in the long run, we need to know and understand our collections, and that is also why we conduct pilot tests, to make our mistakes on a small scale now, rather than on a large scale later.

3. A more crucial anomaly

There was another important lesson from that redesign. Although it had no immediate application, it turned out in the long run to be more important than the other three instances put together. This time, the industry concerned was Extracting, Refining and Founding (or ERF). In the Capital Expenditure Survey, ERF had a large completely enumerated stratum, and in the most recent survey-round there had been a considerable incidence of non-response

© 2005 Australian Statistical Publishing Association Inc.

in it. At that time it was usual for a completely enumerated stratum with non-response to be treated as though it were a sampled stratum, and the respondents as though they had been randomly selected. In other industries I had ignored the additional variability introduced by that practice (let's call it the 'quasi-variance'), but for ERF it looked quite substantial, so I decided I had better estimate it as though it were a real variance, with a view to taking account of it explicitly in the redesign.

Alas, the ordinary variance estimator for that stratum gave me a figure that went clean through the roof. It seemed that in that completely enumerated stratum there would have to be a zero tolerance of non-response. The reason for this anomaly was obvious, and no doubt many readers already have guessed it. BHP, then known then as The Big Australian, dominated that stratum like a giant among pygmies. Since the quasi-variance had been estimated over all possible respondent sets containing the same number of businesses as the actual respondent set, and since many of those sets did not contain BHP, the quasi-variance was also enormous.

4. Prediction model variances

At the time I did not follow the anomaly right through. National Accounts assured me that they would never allow BHP to be a non-respondent under any circumstances, so the problem disappeared. It must have lingered in my subconscious, however, because years later, when I was writing a paper that explored a whole raft of related anomalies, I listed that one first, requiring (and here I quote) 'An expression for the conditional variance of a ratio estimator, subject to a particular sample of population units having been selected' (Brewer, 1963).

But in the mind-set of the 1960s, such a demand made no sense. The only variance then considered was the one defined over all possible samples (that is, all samples containing a specified number of units), so to talk about a variance conditioned on a particular sample was a contradiction in terms. If I had sent that paper to a major journal, it would almost certainly have been rejected outright. I am greatly indebted to the then rather obscure *Australian Journal of Statistics* for having published it without demur, but I strongly suspect that the referee was not a sampling statistician.

It is a measure of Ken Foreman's readiness to accept new ideas, in this case the use of population modelling to define a variance conditional on a particular sample, that he listened patiently to what I had to say, grilled me to make sure I wasn't talking through my hat, and encouraged me to submit for publication. Later on we wrote a joint paper taking the same idea a few steps further (Foreman & Brewer, 1971). Later still, when Ken was writing his monumental book, *Survey Sampling Principles* (Foreman, 1991), he devoted half of Chapter 6 — one of the longest in the book — to this very topic.

Ken's attitude on this issue was, I might add, in stark contrast to that of Morris Hansen. Late in 1966 I took the opportunity to mention this matter to Morris at the US Bureau of the Census. But as soon as he realized I was talking about the use of population models, the conversation came to an abrupt end. 'We don't need *models*', he said, and it was quite obvious that he was not prepared to listen any further.

I don't blame him for having that attitude. For every good idea that comes in out of left field, I imagine there are between 20 and 50 that are pure rubbish. Morris would probably have had many such ideas put to him, and would have decided that none was ever worth listening to. I say more later about his controversy with Richard Royall in the 1970s. Instead of blaming Morris Hansen, I suggest we praise Ken for his openness of mind and readiness to listen.

© 2005 Australian Statistical Publishing Association Inc.

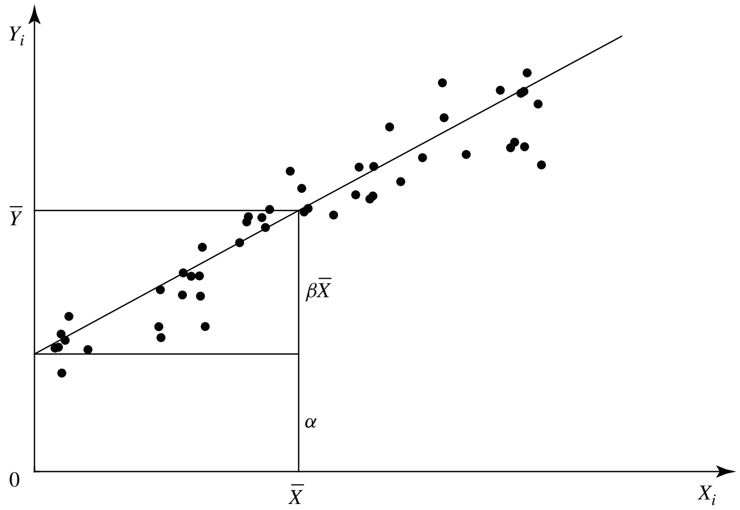

Figure 1. Expansion vs ratio estimation

5. Expansion vs ratio estimation

In Foreman & Brewer (1971), Ken and I together developed a solution to an important question raised in the early sampling textbooks, such as Hansen, Hurwitz & Madow (1953) and Cochran's first edition (1953). The question was essentially this: In the context of simple random sampling without replacement, when is the unbiased expansion estimator, $\hat{Y}_e = (N/n)\bar{y}$, of the population total $Y = \sum_{i=1}^{N} Y_i$ less efficient than its ratio estimator, $\hat{Y}_r = (\hat{Y}_e/\hat{X}_e)X$? Using design-based inference exclusively, they found that a necessary and sufficient condition for $\hat{Y}_r$ to be more efficient than $\hat{Y}_e$ was

$$\rho_{y,x} > \tfrac{1}{2}\frac{S_{Y_i}/\bar{Y}}{S_{X_i}/\bar{X}} = \tfrac{1}{2}\frac{\mathrm{CV}(Y_i)}{\mathrm{CV}(X_i)}, \tag{1}$$

where S_{Y_i} and $\mathrm{CV}(Y_i)$ were the population standard deviation of the Y_i and their coefficient of variation respectively, and a similar notation applied for the X_i.

Condition (1) is correct, but not easy to get a feel for. Foreman and Brewer, however, showed that if the population could be modelled by

$$\xi_0\colon\ Y_i = \alpha + \beta X_i + U_i\,,$$

where $\mathrm{E}_{\xi_0}(U_i) = 0$, $\mathrm{var}_{\xi_0}(U_i) = \sigma_i^2$ and $\mathrm{cov}_{\xi_0}(U_i, U_j) = 0$ for all $i \neq j$, then $\hat{Y}_r$ was more efficient than $\hat{Y}_e$, in terms of the mean squared error (MSE) over all possible samples and all possible realizations of the model, when

$$|\alpha| < |\beta\bar{X}|\,.$$

This inequality has a simple geometric interpretation (see Figure 1).

I should in fairness mention that Cochran, even as far back as 1953, was not averse to the use of prediction models such as ξ_0 when he saw a use for them. He just does not appear to have thought them useful in this instance.

© 2005 Australian Statistical Publishing Association Inc.

6. Model-assisted survey sampling

To return to that first outcome that I required back in 1963, 'An expression for the conditional variance of a ratio estimator, subject to a particular sample of population units having been selected'. You may already have recognized that by 'conditional variance' I meant the variance under a prediction model. The model I used was almost the same as ξ_0, but without the intercept term, α, and with a variance function thrown in, proportional to a power of the X_i. That model may be written:

$$\xi\colon\ Y_i = \beta X_i + U_i\,, \tag{2}$$

where $\mathrm{E}_\xi(U_i) = 0$, $\mathrm{var}_\xi(U_i) = \sigma_i^2 = \sigma_\gamma^2 X_i^{2\gamma}$ and $\mathrm{cov}_\xi(U_i, U_j) = 0$ for all $i \neq j$. The specification of the variance function as $\sigma_\gamma^2 X_i^{2\gamma}$ has a history going back to a paper written by Fairfield Smith (1938), in (of all places) Canberra, which then had a population of only 10 000 persons. Fairfield Smith was working for CSIRO at the time, and his specification related to soil samples. Large soil samples were more variable than small ones, so, for him, $\gamma > 0$. Jessen (1942) seems to have been the first to use this model in work we might describe as regular survey sampling.

These days, the model ξ and variants on it can be found almost everywhere in the sampling literature. Such models were used extensively by Särndal, Swensson & Wretman (1992) in their influential book *Model Assisted Survey Sampling*. As its title suggests, that book advocates the use of population models to improve, rather than to replace or even supplement, design-based inference. Earlier, however, Royall (1970, 1971) had used ξ (with $\gamma = \frac{1}{2}$) to make prediction model-based inferences on their own, ignoring the manner in which the sample was selected. For a thoroughgoing exposition of this approach see Valliant, Dorfman & Royall (2000).

7. Design-based and model-based inference

The dispute between those who use prediction models only to sharpen up design-based inference, and those who use such models as a direct source of inference in themselves, has been long and at times bitter. Each approach has its merits, and there are advantages in using both together. Consider how each of these inferences works.

First, design-based inference. Consider the general case where the inclusion probabilities π_i are known but may differ from unit to unit. In that case we can imagine the sampling statistician constructing a model of the population by looking at each of the sample units in turn and saying, 'Oh yes, you (the first unit) were included with one chance in 10, so my model of the population includes you and nine other non-sample units with the same Y_i value as you. But you (the second unit), you were included in a sample with one chance in two, so my model includes you and only one other unit like you. And you, the third, were included with certainty, so my model includes you, but no other units like you', and so on (see Figure 2 for a diagram based on an extremely small example of the above method — a population model of 13 units constructed from a sample of three).

The sum of the Y_i values in the model population so constructed is then identical with the Horvitz–Thompson estimator (HT) of the framework population total Y; and that estimator can be written

$$\hat{Y}_{\mathrm{HT}} = \sum_{i\in s} \frac{Y_i}{\pi_i} = \sum_{i=1}^{N} \delta_i \frac{Y_i}{\pi_i}\,, \tag{3}$$

© 2005 Australian Statistical Publishing Association Inc.

Inclusion probability	Sample	Imputed
1/10	○	○○○○○○○○○
1/2	△	△
1/1	□	

Figure 2. Randomization model

where $\delta_i = \mathrm{I}\,(i \in s)$ is an inclusion indicator, taking the value 1 if the ith unit is in the sample, and the value 0 otherwise.

So even design-based estimation can be thought of as being based on a model, but on a model quite different from the prediction models, such as ξ, that are favoured by the so-called 'model-based' school. More accurately that school should be described as 'prediction-based', and the 'design-based' school would also be better described as 'randomization-based'. Each school uses a model, but one uses a prediction model and the other a randomization model.

Prediction-oriented statisticians ridicule the use of randomization inference, because the π_i are chosen arbitrarily by the sample designer and are therefore unable (they say) to tell us anything about the population! Instead, they use the best linear unbiased estimator (BLUE) of the parameter β, namely

$$\hat{\beta}_{\text{BLUE}} = \frac{\sum_{i=1}^{N} \delta_i Y_i X_i/\sigma_i^2}{\sum_{i=1}^{N} \delta_i X_i^2/\sigma_i^2}\,,$$

to arrive at the best linear unbiased predictor (BLUP) of Y, which is

$$\hat{Y}_{\text{BLUP}} = \sum_{i \in s} Y_i + \sum_{i \notin s} \hat{Y}_{\text{BLUP},i} = \sum_{i \in s} Y_i + \sum_{i \notin s} X_i \hat{\beta}_{\text{BLUE}}\,. \tag{4}$$

Using δ_i, the inclusion indicator, we might continue along the following lines:

$$\begin{aligned} \hat{Y}_{\text{BLUP}} &= \sum_{i=1}^{N} \delta_i Y_i + \sum_{i=1}^{N} (1-\delta_i) \hat{Y}_{\text{BLUP},i} = \sum_{i=1}^{N} \delta_i Y_i + \sum_{i=1}^{N} (1-\delta_i) X_i \hat{\beta}_{\text{BLUE}} \\ &= \sum_{i=1}^{N} \delta_i Y_i + \left(X - \sum_{i=1}^{N} \delta_i X_i \right) \hat{\beta}_{\text{BLUE}}\,. \end{aligned} \tag{5}$$

The first terms in (4) and (5) are just the sample sums, and the second terms are the sums of the BLUPs for the non-sample units. The status of δ_i in prediction estimation is entirely different from its status under the randomization approach. There, it is a binomial random variable with mean π_i and variance $\pi_i(1-\pi_i)$, but under the prediction approach it is a fixed number (1 if the ith unit is in sample and 0 otherwise). For both inferences, $\delta_i Y_i$ is a random variable; but for randomization inference δ_i is a random variable and Y_i is fixed, while for prediction inference Y_i is the random variable and δ_i is fixed.

My guess is that this 'role reversal' would have had a lot to do with the apparent inability of Hansen and Royall to understand each other's language at meetings of the Washington Statistical Society in the 1970s. Given that words such as 'expectation' and 'variance' conjured up utterly different concepts to the two people concerned, it is no surprise that they experienced extreme difficulty, and no little frustration, in attempting to communicate their ideas to each other.

© 2005 Australian Statistical Publishing Association Inc.

8. Comparing the inferences

Which is the better choice, the HT or the BLUP? The BLUP is the better if the prediction model holds exactly, and much the better if both the sample and the population are small, not only because (in an extreme case) the sample sum of the $1/\pi_i$ could easily turn out to be nothing like N, so that the randomization model of the population could have far too many or far too few units in it, but also because the BLUP would be using information that the HT was not using. However, as George Box has reminded us, 'all models are wrong' (Box, 1979 p. 202) and if the sample and population are made large enough, any imperfection in the prediction model could become detectable. So there is bound to be some sample size where the HT becomes preferable. Indeed, a particularly appealing feature of the randomization model is that it becomes more and more realistic as the population size N and the sample size n increase.

However, we should not really be comparing the BLUP with the purely randomization-based HT, but with the best estimator that randomization-based inference can come up with. The current randomization orthodoxy these days (model-assisted survey sampling) suggests that this 'best' might be the estimator here called the optimal regression estimator (or OREG), though other 'optimal criteria' are sometimes used for that purpose instead:

$$\hat{Y}_{\mathrm{OREG}} = \sum_{i=1}^{N} \delta_i \frac{Y_i}{\pi_i} + \left(X - \sum_{i=1}^{N} \delta_i \frac{X_i}{\pi_i} \right) \hat{\beta}_{\mathrm{BLUE}} , \tag{6}$$

where X is the population total of the supplementary or benchmark variable values X_i, and $\sum_{i=1}^{N} \delta_i X_i / \pi_i$ is its HT. The first term in (6) is the $\hat{Y}_{\mathrm{HT}}$ of (3). The second is a calibration term, ensuring that $\hat{Y}_{\mathrm{OREG}}$, like $\hat{Y}_{\mathrm{BLUP}}$ itself, is unbiased under the model ξ.

If the population is large, then, as the sample size increases, the ratio of $\sum_{i=1}^{N} \delta_i X_i / \pi_i$ to X approaches 1, so that the ratio of the calibration term to the HT approaches 0. The BLUP is a parametric estimator of Y, but the HT is non-parametric, while the OREG is a compromise, becoming less and less dependent on parametric inference as the sample size increases.

9. Combining the inferences

So far, we have considered only those estimators that are valid under one inference or the other, but it is possible to use an estimator that is valid under both. We have already observed that we can ensure unbiasedness under the model ξ and still use randomization inference. The relevant estimator, the $\hat{Y}_{\mathrm{OREG}}$ of (6), was a compromise between the fully non-parametric HT and the fully parametric BLUP, but the compromise did not go far enough to justify the claim that the OREG estimator was supported by prediction inference. We next see that the compromise can be pushed a little further with the resultant estimator being capable of being written both as a model-assisted regression estimator and as a model-based predictor. This combined estimator is, moreover, virtually the same as the OREG when the sample is large, and quite close to the BLUP of equations (4) and (5) regardless of whether the sample is large or small. We thus get nearly the best of both inferences.

We limit the derivation of such an estimator here to the case where there is only a single regressor variable (the same X_i as before). Extension to the case where there is more than one regressor variable is possible, but involves matrix algebra. Readers who wish to take this further step should refer to Brewer (1999) or to Brewer (2002 Chapter 8).

© 2005 Australian Statistical Publishing Association Inc.

This derivation can be started by writing down expressions based on the OREG and BLUP estimators, but generalized to the situation where the estimator of the model parameter β is not necessarily the BLUE. This yields

$$\hat{Y}_{\text{reg}} = \sum_{i=1}^{N} \delta_i \frac{Y_i}{\pi_i} + \Big(X - \sum_{i=1}^{N} \delta_i \frac{X_i}{\pi_i}\Big)\hat{\beta}_{\text{reg}} \tag{7}$$

and

$$\hat{Y}_{\text{pred}} = \sum_{i=1}^{N} \delta_i Y_i + \sum_{i=1}^{N} (1-\delta_i)\hat{Y}_{\text{pred},i} = \sum_{i=1}^{N} \delta_i Y_i + \Big(X - \sum_{i=1}^{N} \delta_i X_i\Big)\hat{\beta}_{\text{pred}}\,. \tag{8}$$

We can make these two estimators numerically identical and notationally compatible by requiring $\hat{Y}_{\text{reg}} = \hat{Y}_{\text{pred}} = \hat{Y}_{\text{c}}$ (say) and, correspondingly, $\hat{\beta}_{\text{reg}} = \hat{\beta}_{\text{pred}} = \hat{\beta}_{\text{c}}$. Solving for $\hat{\beta}_{\text{c}}$ we obtain

$$\hat{\beta}_{\text{c}} = \frac{\sum_{i=1}^{N} \delta_i Y_i \varpi_i}{\sum_{i=1}^{N} \delta_i X_i \varpi_i}\,.$$

The weights $\varpi_i = (1-\pi_i)/\pi_i$ are those that apply to the construction of the non-sample portion of the randomization model and they can conveniently be termed 'remainder weights'.

We can rewrite (7) and (8) as

$$\begin{aligned}\hat{Y}_{\text{c}} &= \sum_{i=1}^{N} \delta_i \frac{Y_i}{\pi_i} + \hat{\beta}_{\text{c}} \sum_{i=1}^{N} (1-\delta_i)\frac{X_i}{\pi_i} \quad \text{(regression form)}\\ &= \sum_{i=1}^{N} \delta_i Y_i + \hat{\beta}_{\text{c}} \sum_{i=1}^{N} (1-\delta_i) X_i \quad \text{(prediction form)}\,.\end{aligned}$$

When the selection is simple random sampling without replacement, the combined estimator, $\hat{Y}_{\text{c}}$, reduces to the classical ratio estimator,

$$\hat{Y}_{\text{r}} = \frac{\sum_{i=1}^{N} \delta_i Y_i}{\sum_{i=1}^{N} \delta_i X_i} X\,,$$

already widely used by both schools of inference.

The next few sections describe some unanticipated spin-offs from the use of combined estimation.

10. Near optimality of the combined estimator

Under fairly mild conditions, the combined estimator of β can be remarkably close to its BLUE, and the combined estimator of the population total is then close to the OREG when the sample is large, and to the BLUP regardless of whether the sample is large or small. Consider the behaviour of $\hat{Y}_{\text{c}}$ under the following three conditions:

(i) all the π_i are small compared with unity;
(ii) the π_i are optimized by choosing them to be proportional to the σ_i in the variance function (in a fashion analogous to Neyman allocation); and
(iii) the σ_i are proportional to the X_i, that is, $\gamma = 1$.

© 2005 Australian Statistical Publishing Association Inc.

Condition (i) usually holds reasonably well for all sampled strata. Condition (ii) is desirable and attainable, and holds approximately even for a stratified sample, provided it has been 'Neyman allocated'. Condition (iii) is also reasonably well approximated in practice. (The value of γ is almost always in the interval between $\frac{1}{2}$ and 1, and $\gamma = \frac{3}{4}$ is often a good enough approximation.)

Under these three conditions, assumed for the moment to hold exactly,

$$\hat{\beta}_{\mathrm{c}} \approx \frac{\sum_{i=1}^{N} \delta_i Y_i/\pi_i}{\sum_{i=1}^{N} \delta_i X_i/\pi_i} = \frac{\sum_{i=1}^{N} \delta_i Y_i/\sigma_i}{\sum_{i=1}^{N} \delta_i X_i/\sigma_i} = \frac{\sum_{i=1}^{N} \delta_i Y_i X_i/\sigma_i^2}{\sum_{i=1}^{N} \delta_i X_i^2/\sigma_i^2} = \hat{\beta}_{\mathrm{BLUE}} \,.$$

Thus $\hat{Y}_{\mathrm{c}}$ is not only supported by prediction-based inference, it can be made almost the same as the BLUP itself, and it approximates the OREG even more closely. So, even in practical situations where these three conditions are only approximately satisfied, prediction-oriented statisticians would have little to lose by using $\hat{Y}_{\mathrm{c}}$ instead of their existing best choices of estimator, and randomization-oriented statisticians would lose almost nothing.

In return for these small sacrifices, they would have an estimator that is

(i) supported by both inferences;
(ii) almost as efficient as the BLUP;
(iii) as robust (against model breakdown) as the OREG and therefore
(iv) particularly appropriate for large samples made up of small domains.

11. Some simple MSE or variance expressions, and their estimators

The combined estimator also has remarkably simple and mutually compatible formulae and estimators for the three relevant variances, i.e. the randomization-based variance, the prediction-based variance and Isaki & Fuller's (1982) 'anticipated variance', which is defined both over all possible samples and over all possible realizations of the prediction model.

11.1. The unbiased prediction estimator or predictor

It is convenient to consider first the general form of the unbiased prediction estimator under the model ξ of (2). This has already been given in (8), but can also be written as

$$\hat{Y}_{\mathrm{pred},s} = \sum_{i=1}^{N} \delta_{is} w_{is} Y_i \quad \text{where} \quad \sum_{i=1}^{N} \delta_{is} w_{is} X_i = X \,. \tag{9}$$

Here, the subscript s denotes any subset of the N units in the population that happen to have been used as a sample, regardless of how that subset may have been chosen. Then δ_{is} is the inclusion indicator of the ith population unit for sample s, and w_{is} is the weight attached to that unit when estimating using sample s. The second equation in (9) is the condition that 'calibrates' the estimator $\hat{Y}_{\mathrm{pred},s}$ on the X_i, and in doing so ensures that it is unbiased under the model ξ.

Being a prediction estimator, $\hat{Y}_{\mathrm{pred},s}$ can only have a randomization variance if it is applied to a probability sample, but it always has a prediction variance over all possible realizations of ξ. Since $\mathrm{E}_\xi(\hat{Y}_{\mathrm{pred},s}) = \beta X$, that prediction variance is defined by

$$\mathrm{var}_\xi(\hat{Y}_{\mathrm{pred},s}) = \mathrm{E}_\xi\big(\{\hat{Y}_{\mathrm{pred},s} - \mathrm{E}_\xi(\hat{Y}_{\mathrm{pred},s})\}\big)^2 = \mathrm{E}_\xi\big((\hat{Y}_{\mathrm{pred},s} - \beta X)^2\big) \,.$$

© 2005 Australian Statistical Publishing Association Inc.

However, we are really interested in how well $\hat{Y}_{\mathrm{pred},s}$ estimates Y, not βX. In other words, we want to know the MSE of $\hat{Y}_{\mathrm{pred},s}$ when regarded as an estimator of Y, i.e.

$$\mathrm{MSE}_\xi(\hat{Y}_{\mathrm{pred},s}) = \mathrm{var}_\xi(\hat{Y}_{\mathrm{pred},s} - Y) = \mathrm{E}_\xi\big((\hat{Y}_{\mathrm{pred},s} - Y)^2\big) \leq \mathrm{E}_\xi\big((\hat{Y}_{\mathrm{pred},s} - \beta X)^2\big)\,. \quad (10)$$

This requires some explanation. MSEs are usually larger than variances, and always larger when a fixed value is being estimated. However, viewed from a prediction perspective, Y is a random variable positively correlated with $\hat{Y}_{\mathrm{pred},s}$, so the MSE is smaller than the variance.

Given (2), (9) and (10), it follows that

$$\begin{aligned}\mathrm{MSE}_\xi(\hat{Y}_{\mathrm{pred},s}) = \mathrm{var}_\xi(\hat{Y}_{\mathrm{pred},s} - Y) &= \mathrm{E}_\xi\bigg(\sum_{i=1}^N \delta_{is} w_{is} Y_i - \sum_{i=1}^N Y_i\bigg)^2 \\ &= \sum_{i=1}^N \delta_{is} w_{is}(w_{is} - 2)\sigma_i^2 + \sum_{i=1}^N \sigma_i^2 \\ &= \sum_{i=1}^N \delta_{is} w_{is}(w_{is} - 1)\sigma_i^2 + \bigg(\sum_{i=1}^N \sigma_i^2 - \sum_{i=1}^N \delta_{is} w_{is} \sigma_i^2\bigg)\,. \end{aligned} \quad (11)$$

For most practical purposes the last expression in brackets in (11) can be ignored. For a start, it is the difference between two terms of comparable size, each of which is an order of magnitude n/N smaller than the leading term. Next, if $\gamma = \frac{1}{2}$, $\sigma_i^2 \propto X_i$, so the two terms cancel exactly. But further, if the sample is a probability sample, the w_{is} asymptotes to $1/\pi_i$, and in the limit the expression in brackets is the difference between the population sum of the σ_i^2 and the HT of that population sum. As the sample size increases, that difference in its turn tends asymptotically to 0. For even quite small sample sizes, such as 10 or perhaps even fewer population units, the difference between the two is small compared with the leading term and can therefore be ignored. That gives us,

$$\mathrm{MSE}_\xi(\hat{Y}_{\mathrm{pred},s}) \approx \sum_{i=1}^N \delta_{is} w_{is}(w_{is} - 1)\sigma_i^2\,, \quad (12)$$

which can be estimated by replacing the σ_i^2 in (12) by appropriate sample estimators. It is possible to derive unbiased estimators of the individual σ_i^2 using the MINQUE of Rao (1970) and Chew (1970), but, since each of these is based on less than a single degree of freedom, it is usually preferable to use the MINQUE approach only to estimate a linear function of the sample σ_i^2, and then to estimate the individual σ_i^2 using their relative sizes, as prescribed by the variance function for the model given in (2).

11.2. The regression estimator

Consider the general form of the single-regressor regression estimator. This is

$$\hat{Y}_{\mathrm{reg},s} = \sum_{i=1}^N \delta_{is}\frac{Y_i}{\pi_i} + \bigg(X - \sum_{i=1}^N \delta_{is}\frac{X_i}{\pi_i}\bigg)\hat{\beta}_{\mathrm{reg},s}\,,$$

where $\hat{\beta}_{\mathrm{reg},s} = \sum_{i=1}^N \delta_{is} b'_{is} Y_i = \sum_{i=1}^N \delta_{is} b_{is} Y_i / \sum_{i=1}^N \delta_{is} b_{is} X_i$, so that $b'_{is} = b_{is}/\sum_{j=1}^N \delta_{js} b_{js} X_j$. Then $\hat{Y}_{\mathrm{reg},s} = \sum_{i=1}^N \delta_{is} w_{is} Y_i$, where $w_{is} = 1/\pi_i + (X - \sum_{j=1}^N \delta_{js} X_j/\pi_j) b'_{is}$, and $\hat{Y}_{\mathrm{reg},s}$ is a prediction estimator applied to a probability sample.

© 2005 Australian Statistical Publishing Association Inc.

The prediction MSE of $\hat{Y}_{\text{reg},s}$ is therefore of the same form as the right-hand side of (12) and can be estimated in the same manner.

When considering randomization and anticipated variances, however, it is convenient to assume that the selection procedure used is one of high entropy, such as systematic selection from a randomly ordered population, or Poisson sampling conditioned on the achievement of a specified sample size. (This assumption is analogous to assuming that a sample selected with equal probabilities had been selected using simple random sampling without replacement.)

For sufficiently large sample sizes, the anticipated variance of $\hat{Y}_{\text{reg},s}$ (for any given probability sample s) is then almost at the Godambe–Joshi lower bound (GJLB); see Godambe (1955) and Godambe & Joshi (1965). It can therefore be written

$$\text{antvar}(\hat{Y}_{\text{reg}}) = \text{E}_{\xi}\big(\text{var}(\hat{Y}_{\text{reg}})\big) \approx \sum_{i=1}^{N} \varpi_i \sigma_i^2 \,, \tag{13}$$

and estimated by

$$\widehat{\text{antvar}}_1(\hat{Y}_{\text{reg},s}) = \sum_{i=1}^{N} \delta_{is} \frac{\varpi_i}{\pi_i} \tilde{\sigma}_{is}^2 \,, \tag{14}$$

where $\tilde{\sigma}_{is}^2$ is a suitable estimator of σ_i^2 obtained from the sample s, as suggested above for the prediction estimator. The estimator (14) then has the same basic structure as the corresponding estimator for the prediction estimator, but with $1/\pi_i$ replacing w_{is}.

However, a better estimator of the anticipated variance is possible if the σ_i^2 are estimated for all the population units, again as suggested above for the prediction estimator; that is, first using the MINQUE to estimate a linear function of them, and then using their relative sizes as prescribed by (2). That estimator can then be written

$$\widehat{\text{antvar}}_2(\hat{Y}_{\text{reg},s}) = \sum_{i=1}^{N} \varpi_i \tilde{\sigma}_{is}^2 \,.$$

An exact expression for the randomization variance of $\hat{Y}_{\text{reg}}$ is not easily attainable, but it is often approximated by the HT of the residuals from the regression fits to the individual Y_i. We consider the HT's high-entropy randomization variance below.

The Horvitz–Thompson estimator (HT)

The HT is usually also a special case of the prediction estimator, because customarily the inclusion probabilities π_i are chosen to be proportional to the supplementary variable values X_i. The prediction MSE of the HT is therefore of the same form as (12), but with the w_{is} specified as taking the values $1/\pi_i$:

$$\text{MSE}_{\xi}(\hat{Y}_{\text{HT},s}) \approx \sum_{i=1}^{N} \delta_i \frac{\varpi_i}{\pi_i} \sigma_i^2 \,. \tag{15}$$

As before, this MSE can be estimated by replacing the σ_i^2 in (15) by a suitable estimator.

The HT is then exactly unbiased under both inferential approaches and its anticipated variance is exactly at the GJLB. In consequence we can write that variance as

$$\text{antvar}(\hat{Y}_{\text{HT}}) = \text{E}\big(\text{MSE}_{\xi}(\hat{Y}_{\text{HT},s})\big) = \text{E}\big(\text{var}_{\xi}(\hat{Y}_{\text{HT},s} - Y)\big) \approx \sum_{i=1}^{N} \varpi_i \sigma_i^2 \,, \tag{16}$$

© 2005 Australian Statistical Publishing Association Inc.

reflecting the right-hand side of (13) exactly. Again, we can estimate this expression either from the sample only, by

$$\widehat{\text{antvar}}_1(\hat{Y}_{\text{HT}}) = \sum_{i=1}^{N} \delta_{is} \frac{\varpi_i}{\pi_i} \tilde{\sigma}_{is}^2 , \tag{17}$$

or from the whole population, by

$$\widehat{\text{antvar}}_2(\hat{Y}_{\text{HT}}) = \sum_{i=1}^{N} \varpi_i \tilde{\sigma}_{is}^2 ,$$

in either case replacing the σ_i^2 in (16) by an appropriate estimator, written above as $\tilde{\sigma}_{is}^2$, because although σ_i^2 can be estimated even though the unit i may or may not be in the sample, its estimates vary from sample to sample.

No exact expression exists for the high entropy randomization variance of the HT, but a close approximation is given by

$$\text{var}(\hat{Y}_{\text{HT},s}) \approx \sum_{i=1}^{N} \pi_i (1 - c_i \pi_i) \Big(\frac{Y_i}{\pi_i} - \frac{Y}{n} \Big)^2 ;$$

see Brewer (2002 Chapter 9), Brewer & Donadio (2003). Three useful specifications of c_i are

$$c_i = \frac{n-1}{n - 2\pi_i + n^{-1} \sum_{j=1}^{N} \pi_j^2} , \tag{18}$$

$$c_i = \frac{n-1}{n - \pi_i} , \tag{19}$$

$$c_i = c = \frac{n-1}{n - n^{-1} \sum_{j=1}^{N} \pi_j^2} . \tag{20}$$

Of these, (18) is the most accurate, (19) the most convenient for variance estimation and (20) the most convenient in other ways, because its values do not vary from one unit to another. From Brewer & Donadio (2003), we can estimate the randomization variance of the HT by

$$\widehat{\text{var}}(\hat{Y}_{\text{HT}}) = \frac{n}{n-1} \sum_{i=1}^{N} \delta_{is} \Big(\frac{1}{c_i} - \pi_i \Big) \Big(\frac{Y_i}{\pi_i} - \frac{Y}{n} \Big)^2 . \tag{21}$$

Given the prediction model (2) and the specification (19) for c_i, it can be shown that there is a near equivalence between (21) and (17). The close parallels between many of the other formulae in this section are even more easily noticeable. These properties and the relative simplicity of the formulae themselves (as compared with the variance expressions and estimators derived without recourse to combined estimation) are all aspects of this second spin-off.

12. Avoidance of unacceptable sample weights

The results described so far have, for simplicity, been confined to the situation where there is only a single supplementary variable, with values X_i. When there is more than a

© 2005 Australian Statistical Publishing Association Inc.

single supplementary variable, the corresponding results are still obtainable, but they are best handled using matrix algebra; for details see Brewer (2002 Chapter 8).

Not infrequently, as for instance in surveys of farm operations, there are several, and occasionally up to a dozen, supplementary variables used for calibration. This proliferation of supplementary variables naturally leads to some instability in the sample 'case weights' (which are, by definition, the same for all survey variable values measured on a given unit, but vary from unit to unit).

As a result of this instability, many sample units can attract weights that are less than 1, and even less than 0. Weights greater than or equal to 1 are readily interpretable in terms similar to those used in Section 7 to describe the formation of randomization models of the population. A sample weight of 10 means that the unit is typical, in an easily understood sense. It represents itself and nine other similar units. A unit with weight unity is borderline. It represents itself, but no other unit(s). Sample units with weights less than 1 are problematic. The face meaning of such a weight is that they do not even represent themselves, but they are still genuine members of the population, and it is intuitively unacceptable that they should not make their full contribution. For units accorded negative weights, the paradox is even sharper and the unacceptability of the situation even more obvious.

If the sample weights are specified using prediction inference only, the situation is not easy to rectify. Those units having weights less than unity can be allotted that minimum weight as a new target, and the calibration can be recalculated; but the set of recalculated weights is then often nearly as problematic as the original. Recourse can also be had to tools such as ridge regression, but these sacrifice the exactness of the calibration process and therefore result in estimators that are no longer prediction unbiased.

Fortunately, when both forms of inference are used together, the situation is considerably more manageable. Because of the remainder weight, $\varpi_i = (1-\pi_i)/\pi_i = (1/\pi_i) - 1$, which then appears in the calculations, a unit that is assigned the weight unity retains that same weight throughout all subsequent calculations, regardless of whether it was included with certainty or not. The problem does not always disappear completely as a result, but it is always greatly ameliorated.

13. Concluding remarks

The drawing together of the two inferential paradigms, together with the spin-offs just described, would not have seen the light of day, or at least not in Australia, but for the consistently sympathetic support and ready collaboration of Ken Foreman. Ken's philosophical approach to his statistical tasks was always one that I would sum up as follows: 'Take your time to think about what you are trying to achieve. The more you know and understand about it, the better you will be able to tackle it in the long run.' Or more colloquially, and in the words of the master himself, 'I wouldn't even scratch myself without carrying out a pilot test first'.

References

BOX, G.E.P. (1979). Robustness in the strategy of scientific model building. In *Robustness in Statistics*, eds G.N. Wilkinson & R.L. Launer, pp. 201–236. New York: Academic Press.

BREWER, K.R.W. (1963). Ratio estimation and finite populations: some results deducible from the assumption of an underlying stochastic process. *Austral. J. Statist.* **5**, 93–105.

© 2005 Australian Statistical Publishing Association Inc.

BREWER, K.R.W. (1999). Cosmetic calibration with unequal probability sampling. *Survey Methodology* **25**, 205–212.

BREWER, K.R.W. (2002). *Combined Survey Sampling Inference: Weighing Basu's Elephants.* London: Arnold.

BREWER, K.R.W. & DONADIO, M.E. (2003). The high-entropy variance of the Horvitz–Thompson estimator. *Survey Methodology* **29**, 189–196.

CHEW, V. (1970). Covariance matrix estimation in linear models. *J. Amer. Statist. Assoc.* **65**, 173–181.

COCHRAN, W.G. (1953). *Sampling Techniques*, 1st edn. New York: Wiley.

FAIRFIELD SMITH, H. (1938). An empirical law describing heterogeneity in the yields of agricultural crops. *J. Agricultural Science* **28**, 1–23.

FOREMAN, E.K. (1977). Development and Co-ordination Division—Historical Note, Internal ABS document. Available from ABS Library, Canberra.

FOREMAN, E.K. (1991). *Survey Sampling Principles.* New York: Dekker.

FOREMAN, E.K. & BREWER, K.R.W. (1971). The efficient use of supplementary information in standard sampling procedures. *J. Roy. Statist. Soc. Ser. B* **33**, 391–400.

GODAMBE, V.P. (1955). A unified theory of sampling from finite populations. *J. Roy. Statist. Soc. Ser. B* **17**, 269–278.

GODAMBE, V.P. & JOSHI, V.J. (1965). Admissibility and Bayes estimation in sampling finite populations, 1. *Ann. Math. Statist.* **36**, 1707–1722.

HANSEN, M.H., HURWITZ, W.N. & MADOW, W.G. (1953). *Sample Survey Methods and Theory*, 2 Vols. New York: Wiley.

ISAKI, C.T. & FULLER, W.A. (1982). Survey design under the regression superpopulation model. *J. Amer. Statist. Assoc.* **77**, 89–96.

JESSEN, R.J. (1942). *Statistical Investigation of a Sample Survey for Obtaining Farm Facts.* Research Bulletin #304. Iowa State College of Agriculture.

RAO, C.R. (1970). Estimation of heteroscedastic variances in linear models. *J. Amer. Statist. Assoc.* **65**, 161–172.

ROYALL, R.M. (1970). On finite population sampling theory under certain linear regression models. *Biometrika* **57**, 377–387.

ROYALL, R.M. (1971). Linear regression models in finite population sampling theory. In *Foundations of Statistical Inference*, eds V.P. Godambe & D.A. Sprott, pp. 259–279. Toronto: Holt, Rinehart and Winston.

SÄRNDAL, C-E., SWENSSON, B. & WRETMAN, J. (1992). *Model Assisted Survey Sampling.* New York: Springer-Verlag.

THOMPSON, M.E. (2003). Review of Brewer (2002). In *Short Book Reviews*, **23**, 4.

VALLIANT, R., DORFMAN, A.H. & ROYALL, R.M. (2000). *Finite Population Sampling and Inference: A Prediction Approach.* New York: Wiley.

© 2005 Australian Statistical Publishing Association Inc.

OBITUARY

EUGENE KENNETH FOREMAN AM

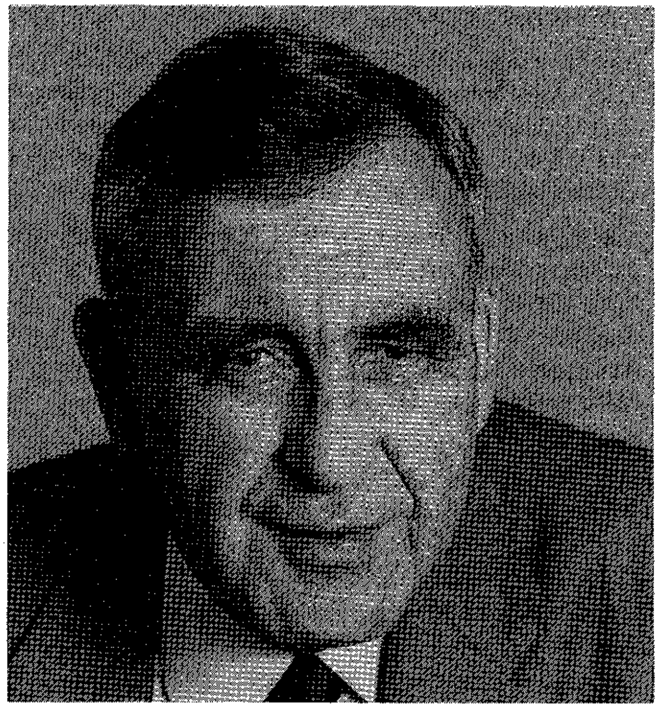

Ken Foreman was one of two survivors of a Lancaster shot down over Germany on the night of 3-4 March 1945, raiding the Dortmund-Ems canal. He was badly injured in bailing out of the aircraft, but lived for another 50 years.

He was born on 23 December 1923 at Coraki, NSW, the son of George and Kathleen Foreman, and was educated at Fort Street Boys' High School. He joined the RAAF in 1943, trained as a navigator/bomb aimer and was posted to 463 Squadron at Waddington, near Lincoln, crashing on his 18th mission. Ken frequently referred with gratitude to the careful attention he received from German and POW medical staff during the weeks between his crash and his liberation by the advancing Allies. The severity of his injuries occasioned a long convalescence, first in the UK and later in Australia. As soon as he was well enough, he enrolled at the University of Sydney, where he graduated with honours as a Bachelor of Economics in 1951, majoring in economics under Heinz Arndt and statistics under Stuart Rutherford.

After graduating he joined the then Commonwealth Bureau of Census and Statistics, working briefly in Sydney and then in Canberra. In 1952 he was selected as a Commonwealth Public Service Scholar and spent a year at the US Bureau of the Census studying the latest developments in survey sampling under Morris Hansen and Bill Hurwitz. He was also much influenced by W. Edwards Deming during that period.

Almost immediately on his return to Canberra, he was set the task of designing a sample of Income Taxpayers to replace the complete enumeration hitherto used in the compilation of statistical tables. This enabled the Taxation Office to abolish 80 positions, and five new ones were created in the Bureau, bringing a Sampling Section into being with Ken as its head.

In 1958 he became Director of the then Methodology and Sampling Division, and was a First Assistant Statistician from 1973 until his retirement in 1984. In the same year he was made a member of the Order of Australia for his contributions to statistics both in Australia and in Papua New Guinea. He had planned PNG's first survey of indigenous agriculture and (on a sample basis) its first census involving the indigenous population.

On leaving the Commonwealth Public Service he established a consultancy practice in sample surveys. His Survey Sampling Principles was published in 1991. After a protracted illness, he died in Canberra on 19 March 1995.

Without his work, decisions important to all Australians would still have to be made "flying blind". The survey methods that he introduced remain the basis of most government economic and social surveys. The seasonal adjustment procedures that he long advocated and eventually introduced are now taken for granted as underpinning the analysis of many important time series. The people he recruited and trained remember him not only as a teacher and guide but as a thoughtful and convivial friend.

He is survived by his wife, Lucilla, two children and three grandchildren.

John Keany, Ken Brewer
and John Carroll

Austral. J. Statist. **38**(1), 1996, 1–14

THE PRODUCT OF THE AUSTRALIAN BUREAU OF STATISTICS

The 1995 Knibbs Lecture, Canberra, 28 November 1995

W. McLENNAN[1]

Australian Bureau of Statistics

1. Introduction

I should like to thank the Canberra Branch of the Statistical Society of Australia for inviting me to present the 1995 Knibbs Lecture.

For the last three years, until I returned home in April to take up the appointment of Australian Statistician, I was the Director of the UK Central Statistical Office and Head of the UK's Government Statistical Service. The experiences I had in the UK were challenging and rewarding, and I endeavour here to draw comparisons from my UK experiences where they are relevant.

In this talk about the Australian Bureau of Statistics (ABS) product I want to provide you first with some background by reflecting a little on Knibbs's contribution to Australian statistics, and also on the contribution of two other more contemporary people, and by describing official statistics. I then discuss at some length the ABS product, particularly: what it is, its relevance, the methodological challenges, its quality, its integrity, its privacy and respondent load, the delivery mechanisms, and dissemination and pricing policy.

2. Knibbs, Archer and Foreman

The Australian Constitution gives the Commonwealth Government power, to be shared with the States, over statistics. In March 1903, Federal Cabinet commenced an investigation into the establishment of a federal bureau of statistics, and the *Census and Statistics Act*, which established the Commonwealth Bureau of Census and Statistics, was given assent on 8 December 1905. (Incidentally, although it has been amended many times, key features of the original legislation are still with us today.) The position of Commonwealth Statistician was advertised in February 1906 and George Handley Knibbs, a surveyor by training and a man who had had little involvement with official statistics, was appointed in May 1906. His salary was £1000 per year.

Knibbs's legacy was the creation — in spite of many difficulties, not the least being the rivalry of other statisticians — of a successfully functioning national

Received and accepted January 1996.

[1]The Australian Statistician, Australian Bureau of Statistics, PO Box 10, Belconnen, ACT 2616.

statistical agency. He was responsible for the first Australian Yearbook and the conduct of the 1911 Census. He foresaw the need for an independent Bureau and was keenly involved in raising the standards of statistical work. He personally pushed for the amalgamation of the State and Commonwealth agencies and nearly succeeded in 1916. He was highly respected in international statistical circles. In 1921, he accepted the invitation of the Prime Minister to take up the directorship of the newly created Bureau of Science and Industries.

I would now like to say a few words about Keith Archer, who was Commonwealth Statistician from 1962 to 1972. I wish to do so, not only because Keith is a personal friend and one of the oldest surviving Statisticians, but also because Keith was the 'godfather' of an ambitious and far sighted scheme to recruit young graduates to work for the Commonwealth Bureau of Statistics. I am talking about the Statistics Cadetship scheme, which is still in existence today. Keith recognised the need both to recruit good graduates and to nurture their careers by providing appropriate and diverse challenges. He and his loyal deputy, Jack O'Neill, who succeeded Keith as Commonwealth Statistician, spent considerable effort and time to ensure the success of this innovative approach to securing and building leaders of the future.

In making these comments, I know that I am biased as I was included in one of the earlier waves of the scheme. I have no doubt that the training and encouragement I received in my early days in the work force contributed in no small way to my being appointed as Director of the UK Central Statistical Office and as Australian Statistician. Further, the cadetship scheme has produced a number of people who went on to be very successful, both inside and outside the Australian Public Service. Some names that spring to mind include Mike Codd, Mike Keating, Chris Higgins, Elizabeth Reid, Vince Fitzgerald and Andrew Podger.

In passing, I would like to mention Ken Foreman, a methodologist of world renown and a good friend, whose zeal and vision had a tremendous impact on many people currently holding senior positions in the ABS, including me. Sadly, Ken died recently, but I am proud that I have been asked to give the inaugural Ken Foreman lecture, jointly sponsored by the Statistical Society of Australia and the ABS, at the international conference to be held in Sydney in 1996.

Enough of history; let me now turn to the topic of this address.

3. Nature of Official Statistics

A formal answer to the question 'What are official statistics?' was given in the United Kingdom in July 1993 in the White Paper on Open Government where it said:

> Official statistics are collected by government to inform debate, decision making and research both within government and by the wider community.
>
> They provide an objective perspective of the changes taking place in national life and allow comparisons between periods of time and geographical areas.

> Open access to official statistics provides the citizen with more than a picture of society. It offers a window on the work and performance of government itself, showing the scale of government activity in every area of public policy and allowing the impact of public policies and actions to be assessed.
>
> Reliable social and economic statistics are fundamental to ... open government [and] it is the responsibility of government to provide them and to maintain public confidence in them.

I like this description, which is not surprising as I had a hand in writing it! In doing so I drew on the long-standing statement of the ABS's mission: 'We assist and encourage informed decision making, research and discussion within governments and the community, by providing a high quality, objective and responsive national statistical service' (Australian Bureau of Statistics, 1994 p.5).

The UK description contains some important elements, including: Government and wider use; to inform debate, decision making and research; an objective perspective; a window on society; a window on work and performance of government; fundamental to open government; and responsibility of government to provide statistics, and maintain public confidence in them.

I can almost hear you saying, 'So what?', because we have become accustomed to having an official statistical service in Australia with these characteristics. However, many other countries don't. In the UK for example, until recently, official statistics were considered to be primarily for the use of government, and public confidence in official statistics was not high either. In considering this fact I often remember some comments I heard Paul Keating make, when as Treasurer he was launching an Australian Year Book, along the lines that the effectiveness and the openness of a country's statistical system is a good indicator of the strength of the democratic processes in that country. I think his comments have a large amount of truth in them.

Within this context, as Australian Statistician I am, in effect, the managing director or chief executive officer of a large public company charged with showing a return (in terms of its statistical output) to its owners or shareholders, viz. the governments of Australia and the Australian public. All the providers and users of statistics expect that without fear or favour the best possible and most appropriate set of statistics is made available to governments and the community.

The independent status of the Australian Statistician's job is specified in law, and I am pleased to say the position and its functions have always received strong bi-partisan political support.

4. The ABS Product: What Is It?

The reason for the existence of the ABS is the product it puts out; we regularly produce a vast array of data, and have done so efficiently and effectively for many years. We exploit all the major avenues for the dissemination of statistics, including publications, computer readable media (magnetic tape, floppy disk, CD-ROM) and *online* electronic access. In all we produce about 1500 releases a year, or eight per work day. The arrival of the Internet and multimedia provides new opportunities for the broad distribution of information, and indeed we have already commenced using the Internet to distribute data to universities and more widely. The presentation of our output is certainly changing.

In addition, the ABS operates a central information service in each of its offices, to supply statistical information that is quickly and routinely available in response to personal, telephone and written inquiries. Also, the ABS has developed client service activities which target the information needs of individual clients on a fee for service basis. In doing this, the ABS has opened access not only to its substantial data resources, but to the professional statistical skills of its staff. I must stress that the ABS has 'public interest' obligations to ensure that basic statistics, at least, are both readily available and affordable, and we all take these obligations very seriously.

Even with the array of measures described above, the ABS cannot in practice meet the needs of all users of statistics direct. Various information intermediaries therefore play an important role in disseminating statistics. These include all branches of the media, libraries, commercial information networks, and business, academic and other research services. In this way ABS data are reaching further into our society, which is becoming more and more reliant on good information.

In particular, over the last three years our output has been changing, and it has been aimed at some important contemporary policy issues. The changes include:

- wide ranging statistics resulting from the very successful Aboriginal and Torres Strait Islander survey;
- substantially enhanced analytical work, including producing forward-looking economic indicators;
- improving and updating the Year Book, and its release on CD-ROM;
- enhanced social reporting, e.g. the Australian Social Trends;
- producing the first national balance sheet;
- valuing unpaid household work for the first time;
- producing uniform Australia-wide public finance statistics;
- the development of national crime and justice statistics.

We have recently engaged in several new projects which are directly aimed at emerging themes of economic and social concern. On the economic side we have committed ourselves to key initiatives in the areas of:

- prices and productivity including whole-economy price indexes, and better measures of inflation;
- a labour costs index where our existing measures of labour costs are often clouded by the impact of compositional change and changes in the mix of wage and non-wage costs, as well as changes in the actual costs of labour;
- environmental and natural resource accounts;

- enhanced services statistics, including steps to measure productivity in the service sectors of the economy.

On the social side there are a number of challenges, including:

- meeting the continuing, and growing, demand for regular benchmark data within each of a number of areas of social concern and particular segments of the population. Increasingly this involves statistics on sensitive issues. (It is in areas such as these that the balance between policy relevance and integrity needs to be carefully considered.)
- better integration of the data that are available from household surveys and administrative systems through the development of statistical classifications and standards and the better coordination of the statistical efforts of all data providers;
- better statistics on income distribution and the social wage, including statistics that can be compared internationally;
- a longitudinal survey of employment and unemployment patterns to provide information over a number of years on the labour force characteristics and experiences of jobseekers following the initiatives in *Working Nation*;
- assisting in the development and measurement of performance indicators of the delivery of government services as part of the Council of Australian Governments review and the development of more comprehensive sets of social indicators to monitor social progress.

5. The ABS Product: Relevance

As you can see from this brief and incomplete outline the ABS product is changing continually. This brings me to one of my major tasks in maintaining the relevance of the ABS product.

In a policy sense, statistics are not relevant on their own. To be relevant, they must be given a political and policy context, and shaped for a specific decision and decision maker. This of course is not the Statistician's job, though the Statistician can do a lot to help the process. Good statistical planning is crucial so that the right data are available at the right moment for decision making, or to inform community debate — official statistics need to be timely, with revisions kept to a minimum. Consistency across series is important in reducing uncertainty. Finally the ABS must analyse and interpret its data so that users are fully aware of what the data are saying.

Planning is the key element as it takes a considerable time to develop, test and implement a new statistical collection to a high standard. An essential element in good planning is the maintenance of effective and continuing contact with users. We use many ways of achieving this, including: involvement with the Australian Statistics Advisory Council; consultation with user groups; periodic evaluations of our programs; feedback from marketing activities; discussion and close association with key clients; maintaining an awareness of public policy issues; coordination with other providers of statistical information.

Although we pursue all these avenues, I often ask myself if we do it well in the sense of having an accurate and up-to-date understanding of user needs. The answer is probably, 'No', especially with respect to the last four or so avenues mentioned. This was one of the reasons behind a recent reorganisation I put through in the senior levels of the ABS in Canberra. Two senior positions were created to look after economic statistics, and population and social statistics. These Group Managers can devote more high level attention to strategic issues, to external contact and to the statistical coordination functions specified in the *ABS Act*.

6. The ABS Product: Methodological Challenges

Included in these current and proposed outputs is a myriad of methodological challenges. I have neither time nor the understanding to speak about each of them in any detail, so I shall just provide a general description of some of them. They include: how best to present data, both on paper and electronically (communications theory); how to develop electronic products (software/hardware); how to measure inflation (economics, price indexes); how to measure productivity (economics, statistics); how to measure labour costs (economics, price indexes); measurement of output and prices in service industries (economics); measurement of the distribution of income (social wage, poverty); producing the 'best' trend estimates (statistics); producing forward-looking economic indicators (econometrics); producing environmental and resource accounts (economics, accounting). These are all issues to which researchers outside the ABS could make important contributions.

7. The ABS Product: Quality

Quality encompasses many aspects, and can mean different things to different people. In my mind, statistical quality can be judged by asking how the data fit the use or uses to which they are being put. Therefore, for the ABS to understand data quality requirements, it has to understand the uses that are being made of data. This gets back to a need for us to maintain close relations with our clients, to obtain feedback from them on quality concerns.

Before turning to some of the specific plans that the ABS has for addressing the quality of its statistics, I should like to make a few general comments. First, quality is often closely related to the application of good statistical methods, so we must make sure that the latest statistical techniques are applied to our collections and compilations, and that they are consistently applied. I mentioned the current senior management restructure. An important focus of the restructure is on lifting the ABS's already considerable commitment to methodological development. One of the ways in which that is being achieved is by raising the prominence of the ABS's methodological unit by up-grading its Head to the First Assistant Statistician level and having that person answer directly to the Statistician.

Secondly, there are often trade-offs that need to be made, and the classic one is between timeliness and accuracy. Users understandably want statistics as soon as possible, particularly if the statistics are needed to monitor time-sensitive policy issues. However, for the statistics to be useful for decision making, they need to be accurate, and accurate statistics take time to compile. The traditional ABS response to this issue has been to release broad-detail statistics first, followed by a series of more detailed releases, with each successive release becoming (hopefully) more accurate. Doing this, though, raises a third dimension, the reliability of statistics, which is shown by the extent of the revisions made. We intend to focus in the future on having publicly-stated targets for these dimensions of quality for all our major releases.

I think that for us to come up with the right balance in the areas of timeliness and accuracy, an assessment of revisions history is essential. A recent case in point involving the quarterly state accounts highlighted the problem, as we were subjected to a significant amount of public criticism about the reliability of those statistics. After consultation with users, I decided to delay the release of those estimates by four weeks. The delay has enabled the ABS to incorporate more 'hard' data into the estimates and this should lead to their being more accurate. We were able, as well, to recast the publication in a more meaningful way.

Thirdly, increasingly we are seeing that quality issues are related across collections. By taking an holistic approach to quality we can often see issues that might have remained hidden if we had just looked at specific collections. An example is the 'statistical integration' project currently underway in the ABS where we are trying to ensure that the same or similar methodological approaches are adopted in all our business surveys; this applies to frame creation, selection, stratification, editing, imputation and estimation procedures, including the contribution of new businesses. To support the project, an in-depth analysis of business demographics has been undertaken and is continuing; such analysis is fundamental to our understanding of quality.

At this stage I should mention some of the specific quality initiatives that are currently being implemented, or that we are planning to implement. One of the most important of these is improvement to the quality of our central business register. The challenge for us is to improve the speed with which we list new businesses, and to improve our abilities to detect and subsequently correct erroneous information. Closer cooperation with the Australian Taxation Office is helping significantly.

Secondly, we are moving towards the use of common processing environments and systems in our economic and household surveys. Standardised processing systems facilitate the consistent application of statistical methodologies, and can take advantage of new technologies.

Thirdly, new technologies can contribute also to improved data quality in the area of data capture. We are increasingly using computer-assisted methods such as computer-assisted telephone interview and computer-assisted personal interview techniques. They can also reduce the costs of data collection and data input.

Fourthly, having appropriate classifications, and applying these classifications consistently across statistical collections, is another factor that contributes to statistics that are of high quality. We have expended and continue to expend considerable resources in these areas, in regard to both social and economic statistics. We are also continuing to endeavour to have these classifications applied outside the ABS, to improve both the quality of administrative data that are used by the ABS and the ease of relating other organisations' statistics to those of the ABS.

Fifthly, we can continue to improve the quality of our statistical products via our international relationships. The ABS is committed to using international standards in a number of areas. One of the benefits of this approach is that it enables us to take advantage of the wide range of expertise involved in the developments of those standards. Also, the application by the ABS of international standards leads to greater compatibility between Australia's statistics and those of other countries — an important outcome with the ever increasing globalisation of economic activity and the pursuit of world-wide social goals.

Sixthly, while I consider that our statistics rate highly when compared to overseas efforts, we can still learn from the experiences of other official statistical agencies. I am keen for us to develop and maintain good relationships both on a bilateral and multilateral basis. Because of this, our staff regularly visit overseas statistical agencies to assess best practices, and we welcome visitors from abroad to the ABS. Wherever appropriate, I am keen for these contacts to lead to cooperative efforts that can benefit both the ABS and overseas agencies. Multilaterally, the ABS has been increasingly involved in major international statistical decision-making forums, in particular with the United Nations Statistical Commission. Our thrust has been to emphasise the need for agreement on international standards, and their consistent implementation. A current example is ABS leadership in statistical developments following the World Social Summit in Copenhagen, and in work associated with the United Nations Statistical Commission.

You may have gathered that, as Statistician, I take quality very seriously. In order for the ABS to fulfil its mission our statistics must be perceived to be credible. If users lack confidence in any particular series then the credibility of the whole ABS is at risk.

8. The ABS Product: Integrity

It is important that the integrity of official statistics is maintained, and perceived to be maintained. That means that the ABS's data, and the analysis of them must be as accurate as possible, and that the ABS must always be objective. To achieve this, statisticians must follow appropriate professional and ethical standards. They must be open about the statistical processes, including adopting predetermined release dates and times and the publishing of their statistical and analytical methodologies. It follows they must be receptive to objective criticism, and be prepared to respond to it. The difference between the perceptions of integrity in Australia and in the UK are, and were, astounding.

9. The ABS Product: Privacy and Respondent Load

I should not talk about the ABS product without pointing out the importance of privacy and respondent load issues. We must always treat very seriously all privacy issues and have concern generally for the load placed on the providers of our data. In developing new collections, or in introducing changes to existing ones, these matters are always considered very carefully. Our questionnaires are always tested, and the reaction of respondents is a determining factor in the decisions made; all steps are taken to minimise respondent load. The Statistician must always face up to the question: is the load placed on respondents, and the possible intrusion of privacy which might occur, justified by the use made of the statistics collected?

In considering this question, it is worthwhile bearing in mind that the ABS has always rigorously maintained the confidentiality of the data it collects. The Australian public has shown it appreciates this by the way it willingly co-operates in providing data to the ABS. External review has supported this observation as well. For example, the Australian National Audit Office reported in 1993, 'The ABS is a good custodian of the personal and business information it collects and holds'. Also, in 1992, a representative of the Privacy Commissioner said, 'The ABS is probably the only Commonwealth agency whose assurances of confidentiality mean what they say ... The ABS appears to have an excellent record in relation to those assurances.' I can assure you all that these comments are well based, and these high standards are being maintained.

10. The ABS Product: Delivery Mechanisms

Having the right products is not of itself sufficient for us to meet our clients' needs. We also have to be able to deliver these products in an effective and appropriate manner so that the data are used. Without that happening all our efforts are for nothing.

Traditionally, the main way in which the ABS has disseminated its information is via paper publications. Today these publications retain an important and still prominent role in delivering the ABS message. Over the years, ABS publications have become sophisticated products — equal to anything produced by overseas statistical agencies — and we have witnessed significant growth in the content and detail of the publications. Also, in recent times the ABS has expanded its range of analytical publications. In particular we have developed a number of thematic or 'all about' publications, such as the Australian Women's Yearbook and a suite of publications on different aspects of families and family life. I think that we will continue to seek to bundle our information for delivery in this manner. Another way in which we are improving our publications is by changing the way the information is actually presented, based on the results of extensive communications research, in order to make our publications easier to follow and more readable. The quarterly national accounts publication was the first to be produced in the new format, and within the next year or so most of our publications will move across to this new 'house style'.

The most exciting prospects for enhanced presentation are, of course, in new areas of electronic delivery. The ABS already has extensive experience in this field. For example, data from the 1986 and 1991 Censuses were made available on a CD-ROM product called CDATA. In fact, the ABS was one of the pioneers of the use of CD-ROM technology in Australia. Continuing this thrust, more recently a comprehensive regional database has been established on CD-ROM, the Integrated Register Data Base, and a CD-ROM version of the 1995 Year Book was developed, which uses the search facilities *etcetera* inherent in the Windows environment. As more and more microcomputers incorporate CD-ROM drives and supporting software, we can expect the range of ABS data available in that format to increase significantly. This method of delivery enables not only rapid access to vast amounts of statistical data, but also allows for the incorporation of search features that make particular information easy to find. CD-ROM also is suited to the development of multimedia products, which are definitely 'flavour of the month'. It's not clear though how the ABS could use multimedia to improve the effectiveness of the delivery of our products, but that's something to which we shall be giving thought.

The ABS is also expanding its on-line delivery of statistical data. PC-AUSSTATS, which was first released in 1992, was a significant improvement on the old mainframe-based AUSSTATS system. Last month, the Windows version of PC-AUSSTATS was launched, and has been received very well. Earlier this year, the ABS joined the information 'superhighway' with the introduction of an ABS Internet service. A homepage was established on the World Wide Web offering general information about the ABS and its services, publication release advices, and basic statistics. Currently, about 1600 accesses are made to this homepage each day. Also, now available on the Internet is a service that provides ABS time series data, effectively the AUSSTATS data base, to the staff and students serviced by subscribing Australian Universities Libraries. We are continuing to investigate the opportunities offered by the Internet, and I expect that the range of products available through this mechanism will grow substantially in the next few years.

Over time, we expect people increasingly to substitute electronic access of information for traditional methods of delivery. This will undoubtedly have implications for the ABS publishing program, but it is not yet clear what these will actually be. I know many people are forecasting the demise of printed output, but I do not expect this to come to pass, at least not in the foreseeable future.

You will be interested to know that the ABS is developing a large database to hold all its statistical data as well as data from other sources, and, most importantly, the associated data descriptions — that is, meta-data. This will permit data from a variety of sources to be brought together easily for use and/or dissemination. In terms of data 'relatability', and of being able to meet specific requests from clients quickly, this will be a big step forward.

The ABS has a public-use obligation to ensure that basic statistics, at least, are both readily available and affordable. This obligation is one we take very seriously. To meet it, publications are made available on a complimentary basis to parliamentarians, major news media organisations, and libraries. Through the Library Extension Program (LEP) the ABS makes its publications available at no cost to over 560 libraries, and assists those libraries to display and use this vast array of information. This service, which is unequalled anywhere else in the world, helps the ABS to fulfil its mission, not just to produce good numbers, but to get them used! The ABS recently won a public administration award for the service.

The ABS has a role not only to provide statistical information, but to organise and deliver that information in a way that facilitates analysis and understanding. The latter task is particularly challenging and we have worked hard in recent years to improve our performance in this area. Nevertheless, I feel that we can, and need, to make further improvements if we are to fulfil our mission of assisting and encouraging informed decision making.

One way in which the ABS now regularly facilitates analysis is by releasing unidentifiable unit records, or microdata. The *Census and Statistics Act* requires that the ABS does not release statistics 'in a manner that is likely to enable the identification of a particular person or organisation'. A Ministerial Determination made under the *CAS Act* gives me discretionary power to release unidentifiable individual statistical records providing a recipient gives a legally binding undertaking that: no attempt will be made to identify particular persons or organisations; the information will be used for statistical purposes only; and the information will not be disclosed to any other person or organisations. If I consider it necessary, I can impose further conditions.

This Determination is an enabling Determination; it does not compel me to release unit record data. Before I authorise the release of microdata I need to be satisfied that ABS guarantees in respect of confidentiality are not only met, but are seen to be met. As you can appreciate, particularly where personal data are concerned, this is a very sensitive area. To assist me in reaching decisions, which are made on a case by case basis, I am advised by a Microdata Review Panel which rigorously assesses all proposals for the release of unit record information to ensure that the data are not likely to enable identification. The Panel takes into account the level of detail in each record, in particular the level of geographic detail, and the extent of disclosure measures taken (e.g. releasing values not as collected but as classes, and randomly perturbing values by some small number).

The release of microdata is just one way that the ABS can support analysis and research, and this approach has been well accepted by the research and general user community; often it has been suggested that such releases should be seen as being in the 'public interest'. While this approach provides maximum flexibility, users sometimes express concern that disclosure avoidance techniques significantly reduce the value of the data that are released. Also, in some cases it may not be possible for the ABS to construct a useful microdata file that would protect the confidentiality of responses. In cases such as these a number of alternatives are available , including:

- ABS undertaking the analysis and research on behalf of clients;
- for analysis and research required to produce an ABS output, use staff from key clients or academia with the requisite skills. Such work would preferably be done on ABS premises under ABS supervision. Analysis and research not for ABS purposes cannot be done this way;
- providing users with sufficient meta-data and a small sample of unidentifiable unit records to enable them to develop analytical models that the ABS would then run against the full data set.

11. The ABS Product: Dissemination and Pricing Policy

Open access to official statistics on the economic and social condition of a country and its population is essential. As the chief executive of a publicly funded national statistical service, I attach considerable importance to easy and widespread access by all levels of government, and the community generally, to the basic official statistics.

The ABS strives to have its basic statistics widely distributed to the community through the media and libraries and from the marketing of its products and services. Market signals through the price mechanism are used by the ABS to assist with the efficient supply and distribution of its statistical information to anyone wanting to have their own copies of publications, and access to other unpublished statistics or to other ABS products.

The ABS policy of charging is intended to serve three main purposes:

- to enable the demand for ABS products and services to be used as a more reliable indicator of how ABS resources should be used;
- to encourage users to address their real needs for ABS products, both statistics and services; and
- to relieve the general taxpayer of those elements of the cost of the statistical service which have a specific and identifiable value to particular users.

The balancing of 'public good' obligations and 'user pays' underpins the ABSs pricing policy. To meet its public good obligations, the main findings of statistical collections and statistical reports on matters of public interest are made available free of charge to the community via the media and by access to copies of publications and other selected special releases provided to libraries located throughout Australia. ABS publications are also made available on a complimentary basis to parliamentarians. Simple enquiries for data from the public to the ABS are handled free of charge. The cost of producing the public good copies of publications is funded from the budget appropriation made to the ABS.

Copies of ABS publications obtained for private use are sold via subscription and over the counter. Such copies are priced to recover collectively their costs of production, distribution and marketing beyond the cost of production of the copies distributed as public goods. Prices are determined by the application of a formula based on a fixed base price and the size of the publication.

Other standard products containing the more detailed statistics of widespread interest are priced to recover their cost of production and distribution beyond the creation of the basic statistics from which they are produced.

Users wanting information more detailed than that published in standard products are required to pay for all costs (including overheads) incurred beyond the costs of collection and production of clean unit record files from which the information is produced.

The extensive range of information collected from providers and the resulting findings can be used to produce value-added products and services sold commercially to satisfy identified needs within the market for statistical information. To satisfy these needs the ABS needs to invest resources in the development, production and delivery of such products and services, and the price of such products and services is set at market prices where quantifiable, but at least to recover the full costs involved, including amortisation of capital and a reasonable allowance for contingencies and risk involved.

Where an organisation seeks to on-sell ABS statistics and statistical products for profit, the ABS seeks licensing arrangements whereby in effect a contribution is obtained from each sale towards the cost of collecting and producing the data provided.

The charging policy represents a government decision to shift some of the costs from the general taxpayer to the user of statistics. The charges are not designed to cover the substantial costs associated with data collection and processing to the clean data stage, nor for the cost of preparing manuscripts for publications which are produced to meet public good obligations. These costs are funded by appropriation.

In addition, the ABS conducts surveys funded in full or part by users to produce official statistics for public benefit. It also provides other services which are cost-recovered such as statistical consultancy, outposts of skilled ABS staff, training courses, seminars and the funding of statistical units responsible for producing statistics on specific fields of statistics which cut across the ABS collection base. Such services are generally for government agencies and are charged on a cost recovery basis in accordance with Department of Finance guidelines.

12. Conclusion

The ABS is, I believe, one of the best statistical offices in the world with products that are second to none. However, making our product relevant, accessible, and of high quality is a continuing challenge. This challenge was taken up by George Knibbs and it is one that I am pleased to accept today. The future for official statistics is exciting, and I believe the ABS is well placed to meet the requirements of our wide range of clients. I am sure that you will agree with me that we have an impressive list of issues on our plate and my role is to make sure that we continue to move forward with the challenges of the times. It is a demanding and responsible role, but in my view a very satisfying one.

There is one not so satisfactory angle which I should mention. We have been disappointed for some time by the lack of interest the research community has shown in methodological issues surrounding official statistics. Even when we offer financial support, via our Research Fellowship Scheme, the take up is not very encouraging. I must say this contrasts very starkly with the interest academics show in official statistics in the UK, the USA and Canada. I can say, though, that this is one direction which the new First Assistant Statistician of our Methodology Division will be actively pursuing, with my full support and, I hope, yours also.

Reference

AUSTRALIAN BUREAU OF STATISTICS (1994). *Corporate Plan.*