Australian Journal of Agricultural Economics, Vol. 34, No. 1 (April 1990), pp. 21-38

A TEST OF BAYESIAN LEARNING FROM FARMER TRIALS OF NEW WHEAT VARIETIES*

R. LINDNER and M. GIBBS
Department of Agricultural Economics, University of Western Australia, Nedlands, WA 6009 and Department of Economics, University of Adelaide, Adelaide, SA 5000

In this study, elicited estimates of farmers' subjective beliefs about the mean and variance of wheat variety yields were used to test propositions about Bayesian learning developed in the recent literature on innovation adoption. A series of empirical tests of the Bayesian adoption model were conducted using beliefs elicited from farm surveys conducted in 1982, 1983 and 1984. The results of the analysis neither confirm nor reject the Bayesian approach as a model of how farmers revise subjective beliefs, but do raise serious doubts about its realism, and suggest some issues requiring further investigation. Shortcomings in the elicitation techniques are discussed and the assumptions of the Bayesian model are reviewed.

The original aim of this study was to test empirically the proposition that farmers revise beliefs about uncertain state parameters in a manner consistent with Bayesian learning as modelled in recent theoretical studies of innovation adoption. The origins for this idea can be traced back at least to Griliches (1960).[1] Some 10 years later, O'Mara (1971) was the first to use Bayesian decision theory to model innovation adoption as a learning process involving the collection and assimilation of information about innovation profitability into the decision makers' beliefs. This notion has become quite popular in the economic literature on innovation adoption (Lindner, Fischer and Pardey 1979; Stoneman 1980, 1981; Feder and O'Mara 1982; Jensen 1982, 1983; Tonks 1983).

The common thread running through most of the modern economic literature on process innovation adoption is the perception that the adoption decision *per se* is essentially one of technique choice under uncertainty. Some authors have focused on risk aversion and the conditions for optimal choice at a particular point in time. Such an approach is essentially static in nature, and typically takes beliefs about uncertain state parameters as a datum that requires no explanation. On the other hand, the studies cited above emphasise the dynamic nature of the adoption process, and investigate how beliefs change over time. The two distinguishing assumptions made in this branch of the literature are: at the time of innovation discovery, an information asymmetry exists (that is, the potential adopter's level of subjective uncertainty about innovation profitability exceeds that of the currently employed production process), and the central and universal feature of

* The authors gratefully acknowledge the financial support received from the Australian Research Grant Scheme which made this study possible.

[1] Griliches (1960, p. 356) suggested that the choice facing a potential adopter should be treated as a problem of decision making under uncertainty, in which the amount of information required to decide to adopt the innovation depends, *inter alia*, on actual innovation profitability, and on its variance.

the adoption process is one of learning about the innovation, thereby gradually eliminating the above-mentioned information asymmetry. All of the studies listed above used Bayesian decision theory to model this learning process, and in most cases make a number of simplifying assumptions.[2]

Although the learning process about innovation profitability takes place in both the evaluation stage (that is, preceding first use of the innovation) and trial use stage (that is, after first use, but before final acceptance or rejection of the innovation) of the adoption process, most of the current Bayesian adoption learning models in the literature are more appropriate to the latter than the former stage. In the evaluation stage some of the assumptions are not realistic because observations on innovation profitability come from other adopters or even less relevant sources.

This study is concerned exclusively with the process of belief revision during the trial use stage of the adoption process. Specifically, it concentrates on how farmers learn about mean yield of new wheat varieties on their own property as a result of growing them on a trial basis. By contrast, the theoretical adoption studies have modelled learning about innovation profitability. Early attempts to elicit subjective beliefs about innovation profitability revealed that most farmers have considerable difficulty in expressing such beliefs in the form of a probability distribution. This is not surprising given the conceptual difficulty of simultaneously allowing for uncertainty about the multiple determinants of profitability.

Analytical Approach

From a methodological point of view, theoretical learning models can be treated in various ways. For instance, the strongest hypothesis is that Bayesian adoption models literally describe the learning process employed by potential adopters (that is, potential adopters actually use Bayes' theorem to revise beliefs). With respect to trial use of wheat varieties by Australian farmers, this hypothesis is simply not credible given the widespread ignorance revealed in casual conversation of even the most basic premises of Bayesian decision theory. At the other extreme, the weakest hypothesis is that while the model may not be realistic, it predicts adoptive behaviour better than, or at least as well as, any alternative model. Evidence consistent with this hypothesis is reported in Lindner, Pardey and Jarrett (1982).

An intermediate hypothesis is to treat the model as a mathematical analogue for an intuitive, and possibly even unconscious, learning process in which farmers revise beliefs in a manner that is consistent with the learning process modelled in the recent theoretical adoption literature. To the extent that farmers learn in such an implicit Bayesian manner, the most plausible explanation is that they develop *ad hoc*

[2] The most common assumptions made are: innovation profitability is distributed normally across time and/or across the population of adopters, the only uncertain state parameter of relevance to the adoption decision is mean profitability of the innovation for the potential adopter (note, by implication, decision makers are risk-neutral), observation of innovation profitability is equivalent to sampling at random from a normally distributed random variable with mean equal to the unknown state parameter, and the variance about the true mean of the process generating individual observations on innovation profitability is known by the potential adopter.

processes, or 'rules of thumb', by experience, which generate outcomes approximating solutions of more formal optimising techniques.[3]

The general aim of this study, then, was to test the hypothesis that farmers are 'implicit' Bayesians in the sense described above. Unfortunately, there is a potential problem involved in testing such an hypothesis if in fact the Bayesian model of belief revision is only an analogue for the farmer's actual learning process. On the one hand, beliefs need to be elicited in the form of probability distributions to test hypotheses about Bayesian belief revision. On the other hand, unless farmers are explicit Bayesians, their beliefs are unlikely to be expressed as probability distributions; thus, attempts to elicit beliefs in this 'unnatural' form are likely to be subject to considerable noise.[4]

Another complication with the general aim defined above is the need to make specific assumptions about the detailed mechanisms of the belief revision process in order to construct specific hypotheses that are operationally testable. For these reasons, much of the empirical analysis reported below is in fact a joint test of the following set of hypotheses.

(a) Farmers act as implicit Bayesians when they learn about the yield potential of new wheat varieties.
(b) Farmers' beliefs about the yield of new wheat varieties can be accurately elicited in the form of probability distributions.
(c) The way in which the learning process has been modelled in the adoption literature using Bayesian decision theory is an accurate mathematical analogue for the way in which farmers utilise information on wheat yields to revise beliefs about the yield potential of new varieties.

A Bayesian Model of Variety Trial Learning

As applied to the revision of beliefs about wheat variety yields, the Bayesian learning models can be made to appear deceptively simple. For instance, if the only unknown characteristic of a new variety is the mean yield on the farmer's own property of that variety (denoted by μ), then one way for the farmer to learn about this unknown state parameter is to grow the variety on a trial basis. The yields obtained in any given year t from growing the variety (X_t) can be thought of as a signal (or message) providing some information about the true mean yield of the variety. As with most information-generating processes, this trial use signal is likely to be a stochastic transformation on the unknown state parameter, μ.

Hence, growers may employ inferential processes of varying degrees of sophistication to derive information in the form of an estimate of mean yield ($\hat{\mu}_t$) from the signal. The simplest possible inference is that observed trial use yield of the new variety is the best estimate of its

[3] Note that this hypothesis is less likely to be rejected the more actual experience the farmer has acquired in taking such decisions, and the more important the decision is to the farmer's welfare. For this reason, laboratory tests of Bayesian learning of the type reviewed by Hogarth (1975) and Slovic, Fischhoff and Lichtenstein (1977) are more likely to reject the hypothesis than tests involving 'real world' decision making. Therefore, the findings of psychological studies rejecting the hypothesis that decision makers do not revise beliefs in a manner consistent with Bayes' theorem were not treated as conclusive.

[4] For a discussion of sources of this noise, see, for example, Kahneman and Tversky (1979).

mean yield (that is, $\hat{\mu}_t = X_t$), and this has been the assumption employed in most theoretical studies.

For the time being, this naive approach will be assumed, and trial use will be treated as a process equivalent to independent random sampling from the population of variety annual yields, $\tilde{X}$. Furthermore, if the process of trial use generates no other information, and takes place on the farmer's property, then yield obtained would be an unbiased estimate of mean yield [that is, $E(X_t) = \mu$]. The informativeness, or amount of information conveyed by this signal, is measured by the inverse of the variance of the information measure, $\hat{\mu}_t$. In this simple case of trial use for a single year, the variance $\hat{\sigma}^2$ is the error variance σ_x^2 of the process generating the signal X_t. In Bayesian terms, this is the variance of the likelihood distribution, or the process variance for short.

In most Bayesian models of adoption [for example, O'Mara (1971), Lindner *et al.* (1979) and Stoneman (1980)], it is assumed that the process variance is known with certainty, and that both the likelihood distribution and prior beliefs about the unknown parameter μ take the form of normally distributed random variables,[5] in which case posterior beliefs formed according to Bayes' theorem will also be normally distributed.

Furthermore, the amount of information contained in posterior beliefs, as measured by the inverse of this distribution variance, δ_{t+1}^{-2}, which is simply the sum of the amount of information in prior beliefs (that is, δ_t^{-2}) and the informativeness of the message (that is, $\hat{\sigma}^{-2}$). That is,

$$(1) \qquad \delta_{t+1}^{-2} = \delta_t^{-2} + \hat{\sigma}^{-2}$$

Thus, an estimate of posterior variance can be calculated from prior variance (δ_t^2) and process variance $(\hat{\sigma}^2)$ using:

$$(2) \qquad \delta_{t+1}^2 = \delta_t^2 \hat{\sigma}^2 (\delta_t^2 + \hat{\sigma}^2)^{-1}$$

The above formulation of Bayesian learning also yields the following expression for the mean of the posterior distribution, γ_{t+1}, as a function of prior mean γ_t, and $\hat{\mu}_t (= X_t)$, denoting the signal, or observed yield from trial use in year t.

$$(3) \qquad \gamma_{t+1} = (\hat{\sigma}^2 \gamma_t + \delta_t^2 \hat{\mu}_t)(\hat{\sigma}^2 + \delta_t^2)^{-1}$$

which can be rewritten as:

$$(4) \qquad \gamma_{t+1} = (1 - \beta_t)\gamma_t + \beta_t \hat{\mu}_t$$

where

$$(5) \qquad \beta_t = \delta_t^2 (\delta_t^2 + \hat{\sigma}^2)^{-1} = \hat{\sigma}^{-2} (\delta_{t+1}^{-2})^{-1}$$

[5] The assumption that beliefs can be represented by normally distributed random variables has been, and continues to be, made for the sake of mathematical tractability. Such an assumption cannot be taken literally, especially where there are small samples of trial data, and it is largely an act of faith to treat this assumption as a reasonable approximation to reality that does no great violence to the results.

In this formulation of Bayesian learning, the posterior mean is a weighted average of the prior mean and the signal generated by trial use. Notice that the weight β attached to the signal in revising subjective beliefs can be shown by simple rearrangement of expression (5) to be the amount of information conveyed by the trial use signal (that is, $\hat{\sigma}^{-2}$) as a proportion of the amount of information in posterior beliefs (that is, δ_{t+1}^{-2}).

Many aspects of the above model of Bayesian learning from trial use of wheat varieties imply relatively naive information utilisation by farmers. First consider the assumption that trial use of a wheat variety is equivalent to sampling at random from a distribution of annual wheat yields. This would be a reasonable assumption if, in any given year, the farmer was totally ignorant of the magnitude and direction of any discrepancy between observed trial use yield (X_t) and (unknown) true mean yield (μ).

It is difficult to believe that farmers would employ such a simplistic approach because the most important determinants of the annual deviation of wheat yield from its long-term mean are various environmental factors about which farmers are far from ignorant. A more sophisticated approach would treat both the information about environmental conditions and the variety yield obtained as part of the signal generated by trial use. It is not obvious how this 'sophisticated' approach should be modelled because of the complex nature of seasonal environmental influences on wheat yields. One possible approach is to assume that the farmer also grows a 'benchmark' variety with known average yield, α. Thus, the signal obtained from trial use comprises the two-tuple, X_t^i, X_t^j, denoting realised yield in year t of the new variety, i, and of the benchmark variety, j, respectively. This can be used to estimate mean yield of the new variety relative to that of the benchmark variety, and can be combined with knowledge of benchmark mean yield, α, to obtain the following estimate

$$\hat{\mu}_t = X_t^i - (X_t^j - \alpha) \tag{6}$$

which is a sufficient statistic for the unknown state parameter μ.[6] If it is assumed that benchmark variety mean yield, α, is so well known that any residual uncertainty about its true value can be ignored,[7] the process variance $\hat{\sigma}^2$ associated with $\hat{\mu}_t$ as defined in equation (6) above is given by:

$$\hat{\sigma}^2 = \sigma_i^2 + \sigma_j^2 - 2\rho_{ij}\sigma_i\sigma_j \tag{7}$$

where σ_i^2, σ_j^2 = (seasonal) process variance of annual yields for the new variety, i, and for the benchmark variety, j, respectively, and ρ_{ij} = the correlation coefficient expressing the extent to which annual yields of the new and benchmark varieties covary.

Equations (1) to (5) above also apply to this more sophisticated model of learning provided that $\hat{\mu}_t$ and $\hat{\sigma}^2$ are defined as in equations

[6] Ratio estimates were not used because of the complexity of deriving the associated variances.

[7] If α is not known with certainty the variance of subjective beliefs about α need to be added to the expression for $\hat{\sigma}^2$ in equation (7).

(6) and (7) respectively (rather than by $\hat{\mu}_t = X_t^i$ and $\hat{\sigma}^2 = \sigma_i^2$ in the 'naive' model).

One contentious aspect of all of the above models is the assumption that the process variance is known with certainty. In theory, farmers should be aware of the likelihood of the possible range of seasonal conditions either from direct experience or from published records. Knowledge of potential variation in seasonal conditions does not guarantee that the variance of annual yields of a new wheat variety will be known with certainty, but it does account for the most important source of variability in wheat yields. A potentially more serious problem is the possibility that knowledge tends to atrophy over time (for example, people forget things and/or the state of nature changes).

Theoretical results have been derived for Bayesian learning based on an independent normal process when neither the process mean nor variance are known (Raiffa and Schlaifer 1961, pp. 298–310), but these results are not strictly applicable to the case under consideration here where variety trial use is for a single period only. Although it might be possible to adapt such results to this situation, it is beyond the scope of this paper to do so.

Survey Methods and Elicitation Procedures

The data for this study came from three consecutive surveys of a sample of 20 South Australian Eyre Peninsular wheat farmers chosen because they were planning to grow the new wheat variety, Aroona, for the first time during the 1982–83 growing season. This method of selection is clearly not random, and is almost certainly biased toward selection of more 'innovative' farmers. (Three of the 20 farmers failed to complete all three interviews.) Growers were interviewed for the first time in February/March 1982, prior to planting the 1982–83 wheat crop. Two follow-up surveys were conducted early in 1983 and 1984 respectively, and were timed so that growers were interviewed some time after the previous harvest, but before the next wheat crop was planted. Hence, the tests outlined below could be replicated over time, as well as across farmers, although tests utilising data from the 1983–84 surveys relate to trial use of Aroona for a second time.

All three questionnaires included questions relating to area sown, years of previous growing experience, actual yields of all varieties grown in the previous season, plus questions for eliciting subjective beliefs about yield of the following four categories of wheat variety (each of which was identified in the first survey):

(a) benchmark variety – a variety which the farmer had been growing for several years, and one which he nominated as a reference standard against which to evaluate new varieties;
(b) other sown variety – another variety which the farmer had been growing for several years in combination with the benchmark variety (note: some farmers did not have this other sown variety);
(c) new variety – the variety Aroona being grown by the farmer for the first time in 1982–3 (or for the second time 1983–4);
(d) disadopted variety – a variety previously grown for several years, but subsequently dropped from the portfolio of varieties grown.

Most of the tests reported below require, *inter alia*, data on farmers' subjective beliefs about wheat variety mean yield both before and after the variety has been grown on a trial basis. For the 'naive' model of the inferential process, the data required for these tests only concern the new variety and include elicited mean and variance of prior and posterior beliefs about variety mean yield, perceived variance of annual variety yields, and realised trial use yield. For the more sophisticated inferential model, perceived variance of annual yields of the benchmark variety, its realised yield in the year of trial use, plus perceived correlation between benchmark and new variety yields also need to be elicited.

Pilot testing of the visual impact method[8] (see Anderson, Dillon and Hardaker 1977) to elicit beliefs about possible yields in the following year raised doubts about how to interpret the variance measures so obtained. In theory, such variance measures should incorporate uncertainty about the growing conditions to be experienced in the following year (that is, 'environmental' uncertainty) as well as uncertainty about variety performance under all possible growing conditions (that is, 'genetic' uncertainty). However, it appeared that the variance measures elicited in this manner captured only, or mainly, environmental uncertainty, and by and large did not incorporate genetic uncertainty. This conclusion has been reinforced by subsequent experiments with elicitation procedures, which suggest that when farmers are questioned about possible future crop yields, they tend to 'anchor' on some best guess about mean crop yield, and then spread counters to reflect only uncertainty about next year's growing conditions.

Accordingly, it was decided in the present study to modify slightly the visual impact questions concerning next year's yield so as to reinforce this anchoring effect, with the aim of using these estimates of variance solely as measures of uncertainty about future annual variety yield relative to variety mean yield (that is, as a measure of environmental uncertainty, σ^2).

A great deal of time was spent with farmers explaining this approach and training them to apply it. By the end of this process, farmers were very comfortable with the use of counters to represent probabilities of different events and understood that their beliefs about the mean yield were not being questioned. Most elicited distributions were symmetrical, but a few were somewhat skewed and one or two were bimodal. In all cases, the farmers spread the counters quite widely to allow for extremes of season and did not seem to 'anchor' on the mean.

It was necessary to find a different procedure to elicit subjective uncertainty, δ^2, about mean yield (that is, genetic uncertainty). In the 1982 survey, an attempt was made to use the visual impact method to elicit separate subjective beliefs about variety mean yield, as well as about seasonal yield variation. Again, problems were encountered with respondents becoming confused or tending to 'anchor' on their best guess about variety mean yield.

[8] See Norris and Kramer (1989) for a discussion of the merits of this method *vis-à-vis* alternative methods.

Because no useable estimates of δ^2 were obtained from the 1982 survey, an indirect method of eliciting δ^2 was tried as an experiment in the 1983 survey, and then repeated in the 1984 survey. The actual procedure used was first to elicit the perceived probability that the true mean yield of the new variety exceeds that of the benchmark variety, and only subsequently to determine the grower's beliefs about the magnitude of the conditional expected superiority (that is, the partial expectation of difference in mean yield). This procedure seemed to overcome the problem of anchoring by first asking for a probability measure rather than an estimate of yield, and by then asking only beliefs about relative yield, thus avoiding the need for a respondent to express beliefs about mean yields in absolute terms. On the assumption that subjective beliefs are (at least approximately) normally distributed, the formula for partial expectation of the normal distribution can be used to estimate a measure of variance implied by the grower's responses to the above two questions.

Ideally, for this method to generate reasonable estimates, any residual subjective uncertainty about mean yield of the benchmark variety should be negligible so that it can be ignored. Evidence to be presented later in this paper suggests that farmers do not know the mean yield of the benchmark variety with certainty (that is, the subjective uncertainty associated with the benchmark variety cannot be ignored). In an attempt to obtain a proxy for this form of uncertainty, the three observations on the mean yield of the benchmark for each farmer were used to calculate a variance of farmer's beliefs about benchmark mean yield.[9] To allow for subjective uncertainty about the benchmark mean yield in all the tests of belief revision, the above estimates of this variance were added to the environmental uncertainty about the benchmark mean, and to the subjective uncertainty measures of all other varieties.[10]

Apart from offering a way around the problem of anchoring, the main advantage of the partial expectations method is that it economises on the number of questions which need to be asked. Against this must be set the conceptual complexity of the measures being elicited. In the end, though, the value of the procedure will depend on how accurately it measures subjective uncertainty about mean yield (or other innovation characteristics).

Certain tests also require a measure of the extent to which the grower perceives the yield of one variety to be correlated over time with the yield of another variety. The recommended procedure to obtain such a measure is first to elicit several distributions of subjective beliefs about possible yields from one variety, each being conditional on the other variety yielding some specified but different amount. A measure of

[9] Such a measure of variance is not a true estimate of farmers' subjective uncertainty about benchmark mean yield as it contains elicitation error variance and possibly other random factors as well. All except one farmer had a standard deviation of less than 1 bag/acre with the majority being less than 0·6 bags/acre. Given that most farmers had less than 9 years' experience sowing their benchmark variety (median 7 years), a certain, albeit minimal, level of uncertainty about the benchmark mean is not unexpected. Note though that these estimates of subjective uncertainty about benchmark mean yield are substantially lower than estimates of subjective uncertainty about mean yield of other varieties elicited by the partial expectations method (Table 4).

[10] The tests were also repeated excluding this variance measure, and were found to be robust in the sense that the results were insensitive to its exclusion.

yield covariance between the two varieties can then be calculated from these conditional distributions plus the marginal (that is, unconditional) distributions of annual yields already elicited. In the 1982 survey, the visual impact method was used to elicit conditional distributions for combinations of the benchmark variety with each of the other three variety types, and the associated covariance and correlation coefficient were calculated from these data. Apart from a few rare exceptions, all correlation coefficients were uniformly high (that is, $\geqslant 0{\cdot}9$); thus, for the purpose of this study, the median value of $\rho = 0{\cdot}97$ was used in equation (7) to calculate $\hat{\sigma}^2$ for each grower.[11]

Empirical Analysis

Belief revision for a new wheat variety

The relationships derived in the discussion of the Bayesian model of variety trial learning provide the basis for specific empirical tests of the proposition that farmers revise beliefs about the yield of a new wheat variety in a manner that is implicitly Bayesian. For instance, estimates of posterior mean yield derived from the theoretical relationships encapsulated in equation (3) can be tested for equality with the corresponding elicited values of posterior mean yield by employing a two-tailed test of whether the mean of the paired differences is significantly different from zero. Because the first attempt to measure the subjective uncertainty component of prior variance using the visual impact method failed, the mean and variance revision tests could be performed only for the 1983–84 data.

Mean revision tests. The test of the relationship between estimated and elicited values was conducted twice using two different estimates of the posterior mean. The first was calculated on the assumption that farmers employ a naive inferential process (in which variance merely equals perceived seasonal variance for Aroona) while the second presumed a more sophisticated approach [in which variance equals perceived variance of the difference in yield between Aroona and the benchmark variety; see equation (7)]. The mean of the paired differences between elicited posterior mean yield and estimated posterior mean yield was 0·34 bags/acre for the so-called naive approach, and 0·98 bags/acre for the sophisticated approach. The corresponding t values are 1·43 and 2·95 respectively, while the critical value of t at the 5 per cent level based on 15 degrees of freedom is 2·13. Thus, the null hypothesis that posterior mean yield estimated from equation (3) equals elicited posterior mean yield cannot be rejected if the naive inferential approach is employed, but is rejected for the sophisticated approach. Moreover, the percentage of variation in the estimated value 'explained' by regressing it on the elicited value is less than 60 per cent in both cases. This rather low level of 'explanation' could be due to the fact noted above that the estimated value of posterior variance is based,

[11] Actual correlations of yields for 11 selected pairs of varieties were computed from experimental data. For eight of the 11 pairs, the correlation coefficient was 0·96 and was less than 0·90 for only one of the 11 pairs. Three values of ρ (0·92, 0·97 and 0·99) were used initially in equation (7) to correspond to low, medium and high correlations. As the results were not sensitive to the value of ρ, they are reported for the median value only.

in part, on an elicited measure of prior variance which may be subject to substantial measurement error.

Variance revision tests. To conduct the corresponding variance revision test, two estimates of posterior variance for Aroona were calculated using equation (2). Again, the two estimates were calculated by assuming either a naive inferential process, or a sophisticated inferential process. For the naive approach, the mean difference between the elicited and estimated variances was only 0·14, and the associated t value of 0·3 was much less than the critical value of 2·13. While the null hypothesis of equivalence cannot be rejected on the basis of this test, visual inspection of the data revealed a wide scatter of observations. This apparent lack of any relationship between the variables was confirmed by the fact that the correlation coefficient was only −0·19, which is far from significant at the 5 per cent level. The correlation coefficient for the sophisticated approach was also not significant, and the t value of 3·65 for the test that the mean difference of 1·1 is significantly different from zero exceeds the critical value of 2·13.

Thus, these results clearly reject the hypothesis that the actual reduction in 'genetic uncertainty' is consistent with that predicted by the Bayesian learning model under test. Note that both the elicited value (of posterior variance) as well as its estimated value are likely to be adversely affected by the problems encountered in measuring subjective uncertainty referred to above. This may explain the even poorer performance of the variance revision test *vis-à-vis* the mean revision tests.

Less critical variance revision tests (that is, relatively tolerant of measurement error in prior and/or posterior variance) were also devised. For instance, from equation (1) it can be seen that trial use of a variety must reduce subjective uncertainty about its mean yield so that elicited posterior beliefs should contain more information than the corresponding elicited prior beliefs (that is, $\delta_{t+1}^{-2} > \delta_t^{-2}$). A test to determine whether trial use reduced farmers' 'genetic uncertainty' will be termed the weak variance revision test. Nine of the 17 farmers satisfied this weak test for the new variety Aroona in 1983–84 but, for the other eight cases, posterior variance was greater than prior variance.

Adjustment coefficient tests. The basis for a further set of tests which completely eliminates the need for elicited variances of subjective beliefs can be obtained by rearranging equation (4) as follows:

$$\gamma_{t+1} - \gamma_t = \beta_t(\hat{\mu}_t - \gamma_t) \tag{8}$$

From equation (5) it can be seen that, for Bayesian learning, the coefficient β measures the ratio of the amount of information in the signal generated by growing Aroona for one year to that in posterior beliefs about Aroona. As the denominator in this ratio is composed of the amount of information in the signal plus any information contained in prior beliefs (which must be ≥ 0), the estimated coefficient should satisfy the condition $0 \leq \hat{\beta} \leq 1$. Furthermore, irrespective of the learning process involved, β can be thought of as a version of the Nerlovian adjustment coefficient.[12] Whatever the underlying theory, the condition that $0 \leq \hat{\beta} \leq 1$ can be tested by eliciting prior and posterior mean beliefs, by inferring mean yield from trial use, and then calculating $\hat{\beta}_t = (\gamma_{t+1} - \gamma_t)(\hat{\mu}_t - \gamma_t)^{-1}$. Again, $\hat{\mu}_t$ can be calculated using either the naive or more sophisticated inferential model discussed above. Estimated values of β for belief revision about Aroona for both the 1982–83 and 1983–84 seasons of trial use are presented in the form of frequency distributions in Table 1. Because elicitation of subjective beliefs about mean yield may also be subject to a degree of measurement error, only values of $\hat{\beta} < -0{\cdot}25$ and $\hat{\beta} > 1{\cdot}25$ were treated as violating the condition for Bayesian learning.

TABLE 1

Distribution of Calculated $\hat{\beta}$ Values (Partial Adjustment Coefficients) for Aroona for 17 Aroona Growers who Completed All Interviews

	Naïve model[a]		Sophisticated model[b]	
$\hat{\beta}$	1982–83	1983–84	1982–83	1983–84
$< -0{\cdot}50$	1	1	6	7
$-0{\cdot}50$ to $-0{\cdot}25$	—	—	1	2
$-0{\cdot}25$ to $0{\cdot}00$	5	2	1	1
$0{\cdot}00$ to $0{\cdot}25$	6	7	3	1
$0{\cdot}25$ to $0{\cdot}50$	5	5	—	1
$0{\cdot}50$ to $0{\cdot}75$	—	1	2	—
$0{\cdot}75$ to $1{\cdot}00$	—	1	—	—
$1{\cdot}00$ to $1{\cdot}25$	—	—	—	2
$1{\cdot}25$ to $1{\cdot}50$	—	—	—	—
$> 1{\cdot}50$	—	—	4	3

[a] Naïve model $\hat{\beta} = (\gamma_{t+1} - \gamma_t)(Y_t - \gamma_t)$
where γ_t is the grower's prior mean estimate of the mean yield of Aroona, and Y_t is the actual yield of Aroona.
[b] Sophisticated model $\hat{\beta} = (\gamma_{t+1} - \gamma_t)(\hat{\mu}_t - \gamma_t)$
where $\hat{\mu}_t = Y_t(\text{Aroona}) - Y_t(\text{benchmark}) + \gamma_t(\text{benchmark})$.

The 1982–83 season turned out to be a severe drought year while the 1983–84 season was one of the best on record. As a result, the signals generated by trial use in each year provided markedly different estimates of variety mean yield depending on the inferential model employed, and the results derived assuming naive inference are sharply differentiated from those based on the assumption that farmers employ sophisticated inferential processes. If the former assumption is correct, then in 1982–83, 11 farmers satisfied the criterion that $0 \leq \hat{\beta} \leq 1$, and for five of the six remaining farmers with negative $\hat{\beta}$ values, the level was small enough for the discrepancy to be ascribed to measurement error. Thus, only one grower was judged to have violated the theoretical constraints on β during the first year of trial use. In the second year of trial use, all but three of the 17 growers satisfied $0 \leq \hat{\beta} \leq 1$, and again only one grower had a value clearly inconsistent

[12] The adjustment coefficient measures how much trial use changes the mean of subjective beliefs about variety mean yield as a proportion of the discrepancy between the observed signal from trial use and the prior mean.

with this criterion. These results contrast strongly with those assuming sophisticated inference where 11 and 12 growers respectively had values during the 1982–83 and 1983–84 periods of trial use which fell outside a 'stretched' critical range from $\hat{\beta} = -0{\cdot}25$ to $\hat{\beta} = 1{\cdot}25$.

Belief revision for other wheat varieties

Due to the inconclusive results obtained from the above tests using data on revision of beliefs about a new variety, further tests were conducted on more established varieties in an attempt to determine whether or not farmers are implicit Bayesians.

Effect of experience on prior variance. Before any additional empirical evidence about such varieties can be assessed, the theoretical framework outlined above needs to be extended to derive propositions about the relationship between the farmer's years of experience growing a particular wheat variety and his beliefs about its yield potential. In principle, each year that a farmer grows a variety provides another opportunity for learning something about mean yield of that variety. Thus, each year's experience can be treated as a successive period of trial use yielding one more signal, and by recursive summation of equation (1) for t periods of trial use, we obtain:

$$\delta_t^{-2} = \delta_0^{-2} + t\hat{\sigma}^{-2} \tag{9}$$

which can be rewritten as:

$$\delta_t^2 = \hat{\sigma}^2 \delta_0^2 (\hat{\sigma}^2 + t\delta_0^2)^{-1} \tag{10}$$

where δ_0^2 is defined as 'initial' prior variance preceding the initial period of trial use.

Furthermore:

$$\partial \delta_t^2 / \partial t = -\hat{\sigma}^2 \delta_0^4 (\hat{\sigma}^2 + t\delta_0^2)^{-2} \tag{11}$$

and

$$\lim_{t \to \infty} \delta_t^2 = 0$$

In other words, subjective uncertainty about any variety's mean yield should be inversely related to the number of years the farmer has been growing that variety. As a corollary, subjective uncertainty should be negligible if the farmer has sufficient experience growing the variety.

In designing the survey it was tacitly assumed that farmers would only use a variety as a benchmark for evaluating other varieties after they had grown it for many years. For this reason, it was assumed that mean yield of the benchmark variety would be known with certainty. Hence, it was only possible to test the relationship postulated in equation (11) for 'other' sown varieties. Note that the validity of this test depends, *inter alia*, on the rather unlikely presumption that all farmers have the same initial prior variance, δ_0^2. The only extra data needed to test this hypothesis of experiential variance revision are the number of years a farmer has grown each variety at the time of elicitation of prior variances about mean yields.

The results of this test are presented in Table 2. While the coefficients are not significantly different from zero at the 5 per cent level, they do have the correct sign, suggesting that, as farmers gain more experience with particular wheat varieties, their elicited prior variances about varieties may tend to decrease. Such a finding is not inconsistent with the hypothesis that farmers are implicit Bayesians.

TABLE 2

Relationship between Elicited Prior Variance and Years Grown for Other Sown Varieties[a]

Year		Coefficient	Standard error	Student's t^b	R^2 (df)
1983	α_0	2·33	0·45	5·16**	
	α_1	−0·15	0·08	−1·88*	0·06 (52)
1984	α_0	0·42	0·12	3·47**	
	α_1	−0·04	0·02	−1·90*	0·16 (20)

[a] Regression models used were of the form: variance $= \alpha_0 + \alpha_1$ (years sown).
[b] *Statistically significant at the 10 per cent level; **statistically significant at the 5 per cent level.

Experience and adjustment coefficients. Partial adjustment coefficients can also be calculated for both the benchmark variety and the other sown varieties. Regardless of whether a farmer is sowing a new variety on a trial basis, or continuing to sow a variety he has been growing for several years, the values of these partial adjustment coefficients, β, should lie between zero and one if the farmer is a Bayesian learner. Furthermore, the more certain the farmer is about mean yield the closer they should be to zero. This can be shown by substituting equation (10) into equation (5) to obtain:

$$\beta_t = \delta_0^2[\hat{\sigma}^2 + (t+1)\delta_0^2]^{-1} \tag{12}$$

thus:

$$\partial\beta_t/\partial t = -\delta_0^4[\hat{\sigma}^2 + (t+1)\delta_0^2]^{-2} < 0 \tag{13}$$

and

$$\lim_{t\to\infty} \delta_t^2 = 0$$

Thus, the estimated adjustment coefficients of different varieties should be inversely related to length of the farmer's experience growing that variety if he is a Bayesian learner. As a corollary, the estimated coefficient should be equal to, or close to zero for the benchmark, or any other variety about which the farmer could be expected to be relatively well informed as a result of growing it for many years.

Table 3 contains the distribution of $\hat{\beta}$ values assuming a naive inferential process for both the benchmark and other sown varieties. These estimates were calculated for all surveyed farmers, including those not growing Aroona, who completed the relevant questionnaires. Again, we treat the range from $\hat{\beta} = -0{\cdot}25$ to $\hat{\beta} = 1{\cdot}25$ as the critical

TABLE 3

Distribution of Calculated $\hat{\beta}$ Values (Partial Adjustment Coefficients) for Farmers who Completed Some Interviews

$\hat{\beta}$	Benchmark variety — Naïve model 1982–83	Benchmark variety — Naïve model 1983–84
<-0.75	2	1
-0.75 to -0.25	2	1
-0.25 to 0.25	25	21
0.25 to 0.75	5	4
0.75 to 1.25	1	—
>1.25	—	—

$\hat{\beta}$	Other sown varieties — Naive model[a] 1982–83	Naive model[a] 1983–84	Sophisticated model[b] 1982–83	Sophisticated model[b] 1983–84
<-0.75	1	1	9	2
-0.75 to -0.25	2	1	1	5
-0.25 to 0.25	18	20	5	7
0.25 to 0.75	8	5	4	1
0.75 to 1.25	2	—	1	—
>1.25	—	—	4	4

[a] See footnote to Table 1.
[b] See footnote to Table 1.

range on the grounds that estimation of β is likely to be subject to considerable error. As is the case for Aroona, almost all naive $\hat{\beta}$ values fall inside this range, and most are very close to zero.

There is also very little difference between the benchmark variety and the other sown varieties in terms of the dispersion of these values, although it can be seen from Table 3 that the proportion of farmers with $\hat{\beta} > 0.25$ is slightly greater for the 'other' sown variety than for the benchmark. This evidence seems to support the contention that farmers have a high degree of certainty about the mean yield of established varieties, and are even more certain about the benchmark than about other established varieties.

However, such a conclusion cannot be supported if farmers also take account of extra information about seasonal conditions in the manner hypothesised above.[13] The results of estimating adjustment coefficients for the other sown variety on the assumption that farmers are sophisticated in their use of extra information are presented in the bottom section of Table 3, and show that less than half of these 'sophisticated' $\hat{\beta}$ values fall within the critical range.

Review of assumptions

The hypothesis that farmers know the mean yield of the benchmark with (near) certainty is not supported by the evidence presented in

[13] Given the use made of annual yields of the benchmark variety to estimate these seasonal influences, it was not possible to estimate sophisticated $\hat{\beta}$ values for the benchmark variety.

TABLE 4

Means of Growers' Beliefs About Wheat Varieties (Bags/Acre)

Belief	Year	Variety: Benchmark	Other Sown	New	Disadopted
Expected value of subjective beliefs about mean yield	1982	6·7	6·0	7·2	6·3
	1983	6·4	5·9	6·7	6·2
	1984	6·4	6·1	7·1	6·6
Process standard deviation	1982	2·5	2·5	2·5	2·8
	1983	2·9	2·9	2·9	2·9
	1984	2·7	2·6	2·7	2·9
Actual yield	1982	7·0	5·5	na	na
	1983	2·4	2·2	3·1	na
	1984	9·9	10·3	10·3	na
Prior standard deviation	1982	na	na	na	na
	1983	na	1·3	1·2	0·9
	1984	na	0·4	0·9	na

Table 4 which shows that, as a result of the drought, the population mean of beliefs for the benchmark variety changed by a comparable margin to that for the other variety categories. To test this hypothesis further, the values of β for the reference variety and for Aroona were compared and found to be not significantly different. Finally, values of naive β for the other sown varieties and the reference variety were pooled across years and regressed on the number of years the variety had been sown but no meaningful results were obtained.

In reviewing the reasons for the lack of support of the hypotheses developed above, it should be recalled that prior belief and likelihood distributions were assumed to be normally distributed. The postulate that farmers revise beliefs in accordance with Bayes' theorem also implies that they process information in a manner consistent with statistical inference. These assumptions may not be realistic, and should be the subject of further research. However, for reasons to be discussed next, they are not regarded as being of pivotal importance in determining the outcome of tests reported in this paper.

In order to understand better the belief revision process, it is also useful to consider changes in farmers' beliefs about all four different categories of wheat varieties as summarised in Table 4. Several general points can be made.

(a) The expectation of beliefs about mean yield differs from one category of varieties to another (from high to low, the order is: new variety, benchmark variety, disadopted variety, other sown variety).

(b) The expected beliefs about mean yield of all variety categories were revised downward following the 1982–83 drought, and then revised back upward following the bumper season in 1983–84 (except for the benchmark which remained the same).

(c) There is very little difference in the variance measures between varieties.

In particular, note that in 1983–84 growers revised expected beliefs about the disadopted variety to the same extent that they revised beliefs for the new variety, and by more than that for the benchmark variety. This behaviour occurred despite the fact that both the benchmark and new varieties were grown by survey respondents in 1983–84, while the disadopted variety was not. Clearly, this finding is not consistent with the Bayesian learning model as formulated above, and at best it suggests quite strongly that farmers' perception of climatic variability is more heavily influenced by recent events than by earlier ones.

As noted above, it is assumed in the model that the process variance is known with certainty. From Table 4, it can be seen that even the average for the sample of values of process standard deviation varies somewhat from year to year. Inspection of values on an individual farm basis revealed much greater variation. For instance, the perceived process variance for Aroona after trial use in 1982–83 exceeded that measured in the 1982 survey for 12 of the 17 growers and, for six of these 12 cases, the increase in perceived process variance from 1982 to 1983 was greater than 50 per cent. This evidence might be better explained by a model in which beliefs are characterised by fuzzy sets. Alternatively, the problem might lie with the state of the art for eliciting beliefs about uncertain events in the form of a probability distribution.

Discussion and Conclusions

To summarise, if farmers ignore available evidence about seasonal conditions when they make inferences about yield potential of a variety from the result of growing it on a trial use basis, then much but not all of the above evidence is not inconsistent with the hypothesis that farmers revise beliefs in a manner which is implicitly Bayesian. However, this interpretation is inconsistent with other evidence, not least of which is the degree of inferential sophistication displayed by farmers during discussions of how they judge the productivity of a newly released wheat variety.[14] The fact that farmers still decide to adopt new wheat varieties even when seasonal conditions during trial use are extremely poor suggests that they are not so naive in their use of variety trial data, and may use a learning process which is even more sophisticated than the so-called sophisticated model used in this study.

If in fact farmers are sophisticated in drawing inferences about variety mean yield from trial use results, then the empirical evidence presented above is not consistent with the implicit Bayesian hypothesis, or at least not consistent with published theoretical models of innovation learning. Furthermore, the systematic similarity in the way farmers revise their beliefs about mean yields of all varieties regardless of the amount of experience accumulated growing the variety and/or whether they grew the variety in the previous season poses a conundrum. Two possible explanations for this behaviour can be postulated.

[14] In the 1983 survey, an attempt was made to determine the process by which farmers revise their beliefs about the mean yield of a new variety using the information obtained from trial use. Typically farmers stated that they took account of prevailing seasonal conditions when they formed beliefs about yield potential of the variety from trial use results, but they were not able to describe the process whereby the nature of the season was incorporated into their beliefs.

One hypothesis that might be intuitively appealing is that farmers possess limited capacity to accumulate and retain information about all possible states of nature. As a result, at any point in time their subjective beliefs about possible yields of any given wheat variety will be distorted by their recent experience of seasonal growing conditions. Specifically, they will assign too high a probability to the possibility that next year will be similar to recently experienced seasonal conditions, and will discount the possibility of seasonal conditions which have not been experienced for some time. Thus, the occurrence of a drought year after a run of more or less normal seasons will lead to a downward revision of subjective beliefs about the yields of all varieties. This description fits the way survey farmers behaved in the 1982–83 growing season.

Another hypothesis is that, when farmers have very little information (as in the case with a new wheat variety), then they attribute their accumulated knowledge about cognate technologies (varieties) to the innovation. This hypothesis could explain the basic similarity of different classes of varieties with respect to both perceived mean yield and process standard deviation. Further empirical research is needed to test the validity of these hypotheses.

Finally, the possibility that farmers revise beliefs in a fuzzy manner which is not consistent with Bayesian learning has to be recognised. So too does the possibility that currently available elicitation techniques are so imprecise that the probability measures derived from them to test the Bayesian learning model bear little or no resemblance to farmers' actual perceptions of their decision-making environment. It is impossible to differentiate between these conflicting hypotheses on the basis of the data available from the current study. More detailed work on these issues may permit firmer conclusions to be drawn about the validity of the Bayesian learning model. For the time being, Bayesian adoption models are best treated as a mathematical analogue for the adoption process, rather than as a literal description of how farmers revise subjective beliefs about prospective innovations.

References

Anderson, J. R., Dillon, J. L. and Hardaker, J. (1977), *Agricultural Decision Analysis*, Iowa State University Press, Ames, Iowa.

Feder, G. and O'Mara, G. T. (1982), 'On information and innovation diffusion: a Bayesian approach', *American Journal of Agricultural Economics* 64(1), 145–7.

Griliches, Z. (1960), 'Congruence versus profitability: a false dichotomy', *Rural Sociology* 25, 354–6.

Hogarth, R. M. (1975), 'Rejoinder to comments by Winkler, Robert M. and Edwards, Ward', *Journal of the American Statistical Association* 70(350), 294.

Jensen, R. (1982), 'Adoption and diffusion of an innovation of uncertain profitability', *Journal of Economic Theory* 27, 182–93.

—— (1983), 'Innovation adoption and diffusion when there are competing innovations', *Journal of Economic Theory* 29(1), 161–71.

Kahneman, D. and Tversky, A. (1979), 'Prospect theory: an analysis of decision under risk', *Econometrica* 47(2), 263–91.

Lindner, R. K., Fischer, A. J. and Pardey, P. G. (1979), 'The time to adoption', *Economics Letters* 2, 187–90.

——, Pardey, P. G. and Jarrett, F. G. (1982), 'Distance to information source and the time lag to early adoption of trace element fertilisers', *Australian Journal of Agricultural Economics* 26(2), 98–113.

Norris, P. E. and Kramer, R. A. (1989), 'The elicitation of subjective probabilities with applications in agricultural economics', *Review of Marketing and Agricultural Economics* (in press).

O'Mara, G. T. (1971), 'A decision-theoretic view of the microeconomics of technique diffusion in a developing country', Unpublished Ph.D. thesis, Stanford University, Palo Alto.

Raiffa, H. and Schlaifer, R. (1961), *Applied Statistical Decision Theory*, Harvard Business School, Boston, MA.

Slovic, P., Fischhoff, B. and Lichtenstein, S. (1977), 'Behavioral decision theory', *Annual Review of Psychology* 28, 1–39.

Stoneman, P. (1980), 'The rate of imitation, learning, and profitability', *Economics Letters* 6, 179–83.

—— (1981), 'Intra-firm diffusion, Bayesian learning and profitability', *Economic Journal* 91, 375–88.

Tonks, I. (1983), 'Bayesian learning and the optimal investment decision of the firm', *Economic Journal* 93, 87–98.