AMERICAN STATISTICAL ASSOCIATION

1985
PROCEEDINGS

of the

BUSINESS AND ECONOMIC STATISTICS SECTION

Papers presented at the Annual Meeting of the American Statistical Association, Las Vegas, Nevada, August 5-8, 1985, under the sponsorship of the Business and Economic Statistics Section

AMERICAN STATISTICAL ASSOCIATION
806 15th Street, N.W. Washington, D.C. 20005

UNCERTAINTY AND POLICY AGGRESSIVENESS

Roger Craine and Douglas Steigerwald, University of California, Berkeley

Introduction

Decision theory analyzes decision making under risk. Agents do not know the precise outcome from their decision, but they know the probability distribution of the outcomes. A popular example in finance is the portfolio choice model. The mean and variance of the payoff depends on the agent's decision--a risk averse agent settles for a lower expected payoff in return for less risk. Brainard, in a seminal paper, applied portfolio analysis to macroeconomic policy decisions. He showed that if the policy multiplier were random, then a cautious (risk averse) policy was optimal.

Brainard also suggested that even if the multiplier were a parameter (i.e., not random), then a cautious policy would still be optimal because decision makers don't know the precise value of the true multiplier. Errors in inference cause the agent to make suboptimal decisions. If the multiplier is a parameter, however, then no mean-variance trade-off exists for the population values, the population variance does not depend on the agent's decision.

The unknown coefficient control problem lies in a no-man's land between estimation and control theory. Asymptotically the sampling errors disappear, but any finite sample inference errors affect decisions. Bayesian econometricians comfortably trod the no-man's land by assigning priors to the unknown parameter. Indeed, Zellner proposed an optimal Bayesian control rule that justifies Brainard's suggestion.

Statistical decision theory also can address the problem. Zaman examines the class of admissible decision rules and proposes several classical alternatives.

This paper reviews a simple unknown parameter control problem. Section 1 compares the random coefficient specification with the unknown parameter specification. Section 2 examines the unknown parameter control problem in more detail and proposes criteria for acceptable rules. Section 3 presents Monte Carlo evidence for three rules and shows that Zellner's suggested Bayesian rule does not dominate.

1. Random Versus Unknown

Random Coefficients. The outcomes in many decision problems are random, and the uncertainty usually affects the decision rule. We use a portfolio selection problem in this section to illustrate. Consider an agent who must decide how many shares (z) of a risky asset to purchase. The per share payoffs (d) vary randomly; in addition some idiosyncratic risk (v) also affects the total payoff (p),

$$(1.1)\qquad p = dz + v.$$

The agent knows the distribution of the random variables,

$$(1.2)\qquad \begin{pmatrix} d \\ v \end{pmatrix} \sim N\left(\begin{bmatrix} \delta \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_d^2 & 0 \\ & \sigma_v^2 \end{bmatrix}\right)$$

or alternatively we could state that the agent's information set (I) contains the distribution of the payoff conditional on the number of shares purchased,

$$(1.3)\qquad P(p|z) \sim N(\delta z,\ \sigma_d^2 z^2 + \sigma_v^2).$$

Notice both the mean and variance of the outcome depend on the shares purchased. The agent's decision rule depends on her attitude toward risk versus expected return and the parameters that characterize the conditional distribution. To illustrate, suppose her loss function is

$$(1.4)\qquad \begin{aligned} L^R &= E(p-p^*)^2 = (Ep-p^*)^2 + \sigma_p^2 \\ &= (\delta z-p^*)^2 + \sigma_d^2 z^2 + \sigma_v^2 \end{aligned}$$

where p^* represents a target payoff. The loss function is a typical control theory "tracking" criterion, rather than the more familiar expected utility function used in portfolio theory. However, the first term can be interpreted as the disutility resulting from deviations of the expected payoff from its target while the second term measures the disutility of risk. The "tracking" loss function implicitly assigns equal weights to both components of the loss.

The optimal decision is

$$(1.5)\qquad z^* = \delta p^*/(\delta^2 + \sigma_d^2).$$

Notice the decision rule is not random. The agent's attitude toward risk and the known parameters in the payoff distribution function affect the decision rule, but realizations of the random variable have no effect.

The portfolio problem is essentially a population or asymptotic theory of decision making under uncertainty. Agents act cautiously because they are risk averse. The population variance (eq. 1.3) depends on the agent's decision; therefore, agents can choose risk-expected return trade-offs. Uncertainty about outcomes affects their decisions, but they have full knowledge of the structure. A sample realization has no effect on the decision rule.

Unknown Coefficients. William Brainard applied portfolio theory to macroeconomic policy decision making. He compared the policy multipliers to the per share payoffs. Macroeconomic variables, say income (y), depend on policy instrument settings (x) via the "multiplier" (β) and some normally distributed independent random shocks (u).

$$(1.6) \qquad y = \beta x + u$$

Brainard carefully specified and analyzed the case with a random policy multiplier. He showed that the variance of the outcome depended on the squared instrument setting (as in eq. 1.4) indicating that the optimal policy is a "cautious" policy (as in eq. 1.5); policy makers trade expected return against variance.

Brainard also suggested, and many followed his suggestion, that even if the policy multiplier were a parameter (i.e., not a random variable) then a mean-variance trade-off still exists because the true multiplier is unknown. Policy makers must rely on a random estimate of the unknown parameter when making actual decisions,

$$(1.7) \qquad b = b(\underline{y}, \underline{x})$$

where $(\underline{y}, \underline{x})$ denote the $(n \times 1)$ vectors of sample observations.

This subtle change in the agent's information set creates a substantial change in the decision problem. In the random coefficient specification the agent's decision affects the expected outcome and the population variance (see equation 1.3)--a mean-variance trade-off exists for the decision maker. When the model specification is a fixed parameter specification (even if the parameters are unknown) the mean depends on the instrument setting, but the population variance is independent of the agent's decision,

$$(1.8) \qquad P(y|x) \sim N(\beta x, \sigma_u^2),$$

no mean-variance trade-off exists.

Ignorance of the true parameter value, however, can cause the agent to make suboptimal decisions and increase the value of the loss, but the variance of the outcome is independent of the agent's decision.[1]

Consider the decision problem

$$(1.9) \qquad L^u = \min_x E(y-y^*)^2$$

subject to the fixed parameter model outlined in equation (1.6), where the agent's information set is limited to observable data $(\underline{y}, \underline{x})$. Let x^* denote the value of the control that minimizes the population loss function, i.e., when β is an element of the agent's information set,

$$(1.10) \qquad x^* = y^*/\beta.$$

When the agent bases a decision on an estimate of the parameter, the control rule is a random variable that depends on the random sample. Let

$$(1.11) \qquad x = y^*/b(\underline{y}, \underline{x}) = x(b, y^*)$$

denote the random control rule. Then the loss function (1.9) can be written as

$$(1.12) \qquad L^u = \beta^2 E\left(x(b, y^*) - x^*\right)^2 + \sigma_u^2 .$$

Sampling error increases the loss. This is a small sample problem; the population variance does not depend on the control setting and no traditional risk-return tradeoff exists.

2. Control With An Unknown Parameter

This section considers the unknown parameter control problem in more detail. The decision maker's problem is to choose a control x (or estimator b) that minimizes the loss in equation (1.12). The necessary condition is

$$(2.1) \qquad E[x(b, y^*)] = x^*.$$

The problem is that the control is an inverse function of the random variables. Classical estimators of the unknown parameter which have desirable properties may yield a control rule with very undesirable properties. For example, substituting the maximum likelihood estimate

$$(2.2) \qquad b_{m\ell} = \Sigma x_i y_i / \Sigma x_i^2$$

for the unknown parameter in the control rule

$$(2.3) \qquad x_{m\ell} = y^* \Sigma x_i^2 / \Sigma x_i y_i$$

leads to an unbounded loss function since the reciprocal of the maximum likelihood estimate has no finite sample moments. The control rule is discontinuous at $b = 0$ and the loss function unbounded as b goes to zero. Figure 1 illustrates the problem.

The vertical axis shows the reciprocal of the maximum likelihood estimate $(1/b)$ and the horizontal axis the value of the estimated coefficient b. The shaded areas in quadrants II and IV show the admissible regions, while the rectangular hyperbolas represent the maximum likelihood control. For any value of β, such as β_0, the vertical distance between the hyperbola and the horizontal line through $1/\beta_0$ represents the error due to sampling. The failure of the maximum likelihood control to stay within the admissible region as b goes to zero indicates the potential for a superior estimator incorporating this information. The discontinuity at zero gives an unbounded loss.

Of course, small sample problems also frequently occur in econometrics. The small sample distribution for many estimators is intractable, or small sample moments don't exist, e.g., most

FIGURE 1

Estimator, Reciprocal and Admissible Region

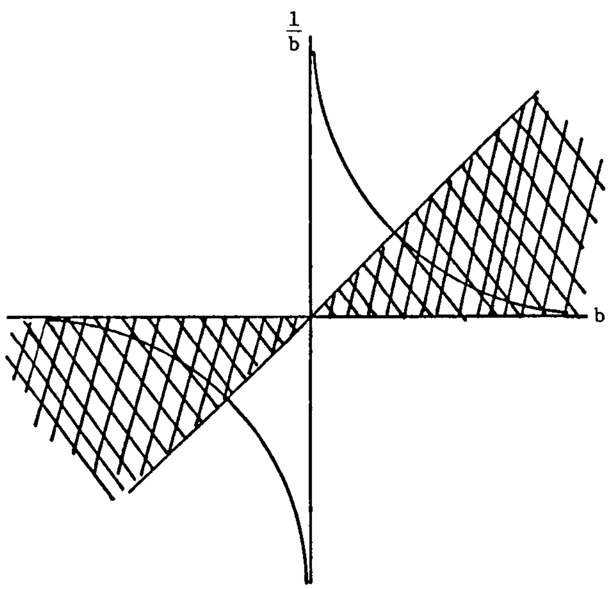

simultaneous equation estimators have no finite sample moments. As a result, econometricians often rely on asymptotic properties to rank estimators.

If we use asymptotic properties to rank control rules, then the control based on the maximum likelihood estimate in equation (2.3) is optimal. As long as the unknown parameter is not zero,[2] the reciprocal estimator $1/b$ is consistent and minimizes the asymptotic variance of the loss function. The control is proportional to the estimator, so no other consistent control reaches a lower probability limit of the loss function.

Thus we have a paradox: the maximum likelihood control leads to an unbounded finite sample loss and therefore is dominated by many rules (e.g., $x = 0$), but asymptotically no rule dominates the maximum likelihood rule. This suggests criteria that an acceptable rule must satisfy.

Let $x = x(\underset{\sim}{y}, y^*)$ denote the random control rule and $x^* = y^*/\beta$ the population value of the optimal control rule.

Criteria:

Large sample: (i) consistency, $\text{plim } x = x^*$, in the limit the control converges to the optimal population value; and (ii) asymptotic efficiency, $\text{plim} \sqrt{n}\,(x-x^*) = 0$, in the limit the control converges to the maximum likelihood control.

Small sample: bounded loss, the control has finite first and second moments so that $E(x-x^*)^2$ exists.

Section 3 proposes and presents Monte Carlo evidence from some rules that satisfy the criteria.

3. Control Rules and Evidence

Zellner & Geisel, and Zellner proposed a Bayesian solution to the unknown parameter control (reciprocal estimation) problem by describing the unknown parameter with a prior distribution. They chose a diffuse prior ($P(\beta) \propto - \infty < \beta < \infty$) leading to a posterior which is the sample distribution of the unknown parameter. Minimizing the resulting posterior expected loss function leads to the minimum expected loss (MELO) control:

$$(3.1) \qquad x_{MELO} = \frac{Eb_{m\ell}y^*}{Eb^2_{m\ell}} = \frac{\beta y^*}{\beta^2 + \sigma^2_{b_{m\ell}}}$$

This has the same form as the random coefficient control rule in equation (1.5).

An operational version of the control rule replaces the population moments with the sample moments leading to a "cautious" policy as Brainard suggested. A small sample mean-variance trade-off exists because the agent believes (as reflected in the prior distribution) the multiplier is random. As the sample gets large the maximum likelihood estimate dominates the prior and the MELO control converges to the maximum likelihood control.

Zaman [1981a, b] extended Zellner's work on the reciprocal estimator, or unknown parameter control problem for a finite sample drawn from a normal population. Zaman [1981b] shows that admissible rules (i.e., rules that are not stochastically dominated for all values of the true parameter) must lie in the shaded regions of quadrants II and IV in figure 1. Zellner's MELO rule is admissible, but of course so are many others. We examine three rules that lie in the admissible region and satisfy the criteria in section 2.

The first is Zellner's MELO rule. This rule has the intuitively appealing property that the noisier the estimate (larger s^2 relative to b^2) the smaller the control. And of course, as the sample size grows the MELO control converges to the maximum likelihood control at rate $1/n$. Another important property, however, is the behavior of the control as a function of the parameter size. As the estimate gets large (b^2 relative to s^2) the MELO control goes to the maximum likelihood control $(1/b)$. More important is the behavior of the control as the estimate goes to zero. The control must converge to some value to bound the loss. The MELO control goes to zero since the estimated variance in the denominator is nonzero.

The second rule is a modified MELO rule (MMELO) that makes the control more cautious for small estimated coefficients.

$$(3.2) \qquad x(MMELO) = \frac{y^*}{b_{m\ell}}\left[1/\left(1 + s^2/b^k_{m\ell}\right)\right]; \quad k = 8$$

Econometrics emphasizes the behavior of the estimator as the sample size varies. In the decision problem one might emphasize the behavior of the control as the size of the unknown parameter varies. The maximum likelihood estimates of the unknown parameter are (normally) distributed around the true value. For large (absolute) values of the unknown parameter, reciprocals of the estimates reasonably approximate the reciprocal of the true value. For small (absolute) values of the unknown parameter, however, reciprocals of the estimates that lie between the true parameter and zero, or reciprocals of estimates with the wrong sign are very bad approximations to the reciprocal of the true parameter. Therefore, we modified the MELO rule to make it more cautious when the estimated coefficient is small.

The third rule we examine is a hybrid rule suggested by Zaman [1981b] that recognizes the discontinuity.

$$x(DIS) = \begin{cases} y^*/b_{m\ell} & |b| > s \\ y^* \cdot b_{m\ell} & |b| < s \end{cases} \tag{3.3}$$

We modified Zaman's rule to make the switching point a function of the estimated sample standard deviation. The rule uses the reciprocal to determine the control for large values of the estimated coefficient (when the risk of disastrous errors is small) and sets the control proportional to the estimated coefficient when the estimate is small (when the risk of disastrous errors is large). Since the estimate is consistent the rule converges to the maximum likelihood rule as the sample size gets large.

Since no rule dominates for all parameter values and sample sizes we ran Monte Carlo simulations to determine the sensitivity of the results to particular parameter values and sample sizes. We varied the parameter β from .001 to 100 and the sample size n from 10 to 100. We drew a sample $\underset{\sim}{y}(n)$ from a normal population with mean β and standard deviation of 2. For each sample we computed control values $x(\cdot)$ and evaluated the loss $L(\cdot) = \beta^2(x(\cdot) - x^*)^2$. We repeated the experiment 100 times and calculated the mean and standard deviation of the loss. Table 1 summarizes the results for selected parameter values and sample sizes.

Table 1 gives summary statistics for four rules. The first two columns of the table list the results for the maximum likelihood rule which violates the bounded loss criterion. The large standard deviations (these are Monte Carlo sample deviations and therefore must be bounded) indicate unboundedness is indeed a problem. The large sample standard deviations for the maximum likelihood control persisted even when we increased the sample size to 100.

The next six columns compare rules that satisfy the criteria. When the parameter is very small (.001), all the bounded loss rules cautiously do very little and the loss is one. As the parameter gets larger, the control rules get more active. The discontinuous control achieves the lowest mean loss at a parameter value of one but displays a higher standard deviation. The

TABLE 1

Mean and Standard Deviation of the Loss Function

Parameter Value	X(ML) μ	X(ML) σ	X(MELO) μ	X(MELO) σ	X(MMELO) μ	X(MMELO) σ	X(DIS) μ	X(DIS) σ
SAMPLE SIZE = 10								
0.001	1.0290	0.3020	1.0000	0.0003	1.0000	0.0003	1.0000	0.0012
1.000	13.0200	50.5200	0.6673	0.1860	0.6356	0.3330	0.3226	0.4442
10.000	0.0033	0.0049	0.0054	0.0057	0.0036	0.0048	0.0047	0.0082
SAMPLE SIZE = 25								
0.001	0.9985	0.0266	1.0000	0.0002	1.0000	0.0001	1.0000	0.0007
1.0000	4.8490	39.9600	0.6629	0.1051	0.6266	0.3202	0.1435	0.1626
10.000	0.0017	0.0027	0.0026	0.0029	0.0013	0.0017	0.0015	0.0021
SAMPLE SIZE = 50								
0.001	1.0200	0.1473	1.0000	0.0001	1.0000	5E-6	1.0000	0.0006
1.0000	0.1644	0.7334	0.6397	0.0735	0.6362	0.2728	0.0784	0.0938
10.000	0.0007	0.0010	0.0018	0.0019	0.0008	0.0011	0.0008	0.0010

modified MELO rule does slightly better at larger parameter values. As one would expect, increasing the sample size improves the accuracy of the estimates and the performance of all the rules. However, the distribution of 1/b does not seem to collapse rapidly. Increasing the parameter value also dramatically improves the performance of all the rules.

Conclusion

This paper examined the case for cautious policy when the decision maker is unsure about the policy multiplier. We contrasted a random policy coefficient with an unknown policy parameter. A random policy coefficient always leads to a cautious policy. An unknown policy parameter, however, leads to an asymptotically aggressive (certainty equivalence) policy.

In finite samples unknown parameters in general lead to intractable small sample distributions with no closed-form optimal rules. Bayesian representations manage to blend the small sample and asymptotic results by combining the prior and likelihood functions to form a posterior distribution. Bayesian rules based upon a diffuse prior yield a unique rule. Yet statistical decision theory shows that a broad class of rules are admissible for the policy decision problem. Our Monte Carlo results indicate that these Bayesian rules are not optimal. Most decision problems, as demonstrated by the examples in this paper where the discontinuity occurs at the origin, embody more information than is characterized by a diffuse prior.

Footnotes

[1]We choose to contrast an unknown (but fixed) parameter specification with a random coefficient specification because macroeconometric modelers typically specify relationships with fixed but unknown parameters. Control applications (if they acknowledge the parameter uncertainty) substitute a random coefficient specification for the unknown parameters. Of course, one could specify a random, but unknown, coefficient model. Then the population variance would depend on the agent's decision and the unknown variance of the random coefficient.

[2]If the unknown parameter is zero, the loss is independent of the control and any rule minimizes (maximizes) the loss.

Bibliography

Brainard, W., 1967, "Uncertainty and the 'Effectiveness of Policy,'" American Economic Review 57: 411-426.

Chow, G. C., 1973, "Effect of Uncertainty on Optimal Control Policies," International Economic Review 14: 532-645.

__________, 1981, Econometric Analysis by Control Methods, Chapter 3, New York: John Wiley & Sons.

Craine, R. and A. Havenner, 1977, "A Stochastic Optimal Control Technique for Models with Estimate Coefficients," Econometrica 45: 1013-1021.

Craine, R. and A. Havenner, 1983, "Classical versus Bayesian Models," International Journal of Systems Science 14: 871-877.

Zaman, A., 1981a, "Estimators without Moments," Journal of Econometrics 15: 289-298.

__________, 1981b, "A Complete Class Theorem for the Control Problem and Further Results on Admissibility and Inadmissibility," Annals of Statistics 9: 812-821.

Zellner, A., 1978, "Estimation of Functions of Population Means and Regression Coefficients including Structural Coefficients," Journal of Econometrics 8: 127-158.

Zellner, A. and M. S. Geisel, 1968, "Sensitivity of Control to Uncertainty and Form of the Criterion Function," in D. G. Watts, ed., The Future of Statistics, New York: Academic Press.