UNIVERSITY OF WISCONSIN-MADISON

SSRI

Workshop Series

CERTAINTY-EQUIVALENCE AND THE THEORY OF THE FIRM UNDER UNCERTAINTY

Richard N. Clarke*

8119

SOCIAL SYSTEMS RESEARCH INSTITUTE

Preliminary Draft
Comments Invited

CERTAINTY-EQUIVALENCE AND THE THEORY OF THE FIRM UNDER UNCERTAINTY

Richard N. Clarke*

8119

*I am grateful to Gary Chamberlain, John Geweke and Don Hester for discussions on this topic.

June 1981

ABSTRACT

A firm's optimal decision rules may be certainty-equivalent when their objective function is quadratic with linear constraints. Whether or not a firm's actual decisions are certainty-equivalent depends on the technique it uses to estimate its true market environment. Maximum likelihood suggests the implausible result that decisions are insensitive to a firm's level of uncertainty.

In this paper I show that Bayes technique is a more reasonable model for firm decisionmaking. Bayes estimation, with quadratic profit functions, suggests that even if firms' decision rules are certainty-equivalent, their actions will not be. In particular, the accuracy of a firm's data is shown to influence its expected optimal action under uncertainty. Precise data lead to nearly the certainty action, less precise data cause the optimal action to diverge more sharply from the certainty action. A further feature is that the behavior generated by these Bayes estimation processes will resemble the behavior of a more typically risk-averse decisionmaker.

0. Introduction

Simple production theory supposes that firms behave as if their market environment were known with certainty. If the environment is stochastic, it is assumed that firms form point estimates of the market state, then make decisions as if these estimates were certain. Should firms have quadratic objective criteria with linear constraints, such "certainty-equivalent" behavior has been shown by Simon (1956) to be optimal. The validity of quadratic empirical models of firm behavior has been enhanced by this certainty-equivalence result.[1] The relevance, however, of certainty-equivalent models is questioned in the important contributions of Hymans (1966), McCall (1967), Baron (1970) and (1971), Sandmo (1971) and Leland (1972) among others. They point out that asymmetries such as managerial utility and risk aversion may make symmetric quadratic objectives unrealistic for the firm. They generally conclude that optimal firm policies are sensitive to stochasticity in the market environment. Such theoretic results accord with intuitive and empirical beliefs that uncertain firms choose actions that systematically differ from actions that would be chosen by certain firms.

Unfortunately, the desirability of the above authors' non-certainty-equivalent models is tempered by the fact that such models are intractable empirically due to our ignorance about firms' utility functions. As Papandreou (1952) points out:

---

[1] See Christensen, Jorgenson and Lau (1973) for example.

> "[Ideally], we should proceed to substitute general preference-function maximization for profit maximization. No doubt this procedure will reduce our chances of being wrong. This protection from error, however, is gained at a cost. It is much harder to derive operationally meaningful theorems concerning firm behavior from a construction which is directly based on preference-function maximization than to do so from the profit-maximization construction. The relative development of the theory of the firm (based on profit maximization) as contrasted to that of the consumer (based on utility-index maximization) testifies to the validity of this argument. If economists wish to replace profit maximization with preference-function maximization, they must take steps to make certain that their procedure will not be rendered meaningless from an empirical point of view. It is necessary to experiment with ideal types in which specific commitments are made about the shape of the preference function maximized by the entrepreneur. This calls, of course, for substantial work both in the general area of expectations (and the manner in which they are related to the flow of events)..."

In this paper, I show that it is not necessary to dispense with tractable quadratic models to ensure that firm behavior is realistically sensitive to uncertainty. This arises because even if linear-quadratic models suggest that decision rules are certainty-equivalent, it need not follow that the actual decisions are risk-invariant. The reason for this divergence comes from the firm's optimal Bayes process for using noisy data to estimate the market's true state. Since state estimates are used as arguments in the certainty-equivalent decision rule, the properties of the estimators generating the state estimate with respect to risk, will influence the firm's actual decision.[2] A further result is that properties of Bayes estimators suggest that firms behave in a "risk-averse" manner.

---

[2]This result holds *a fortiori* for non-quadratic objective functions.

The basic model and its structural decision rules are presented in Section 1, the statistical estimation process in Section 2. In Section 3 I analyze the behavioral characteristics of these models. My concluding remarks are in Section 4.

1. Model and Decision Rules

In this section I construct a simple quadratic model of a firm facing uncertainty and derive its optimal decision rules and profit functions. No particular realism is assigned this model; much richer and more general models could be constructed within this linear-quadratic framework.

The model is a two-period dynamic model. In the initial period, the firm must choose output levels q_0 and q_1 for both period 0 and period 1. A quadratic adjustment cost is incurred in period 0 for varying output between periods. Demand and unit costs are linear in each period with intercept terms that vary between periods.

Period 0

$$p_0 = a_0 - bq_0 \qquad \text{demand, } a_0\ b > 0$$

$$c_0 = (e_0 + fq_0)q_0 + \psi(q_1 - q_0)^2 \qquad \text{cost, } c_0 \geq 0$$

$$\pi_0 = p_0q_0 - c_0 \qquad \text{profit}$$

Evolution

$$a_1 = (1+\alpha)a_0 \qquad \text{demand intercept}$$

$$e_1 = (1+\beta)e_0 \qquad \text{unit cost intercept}$$

Period 1

$$p_1 = a_1 - bq_1 \qquad \text{demand}$$

$$c_1 = (e_1 + fq_1)q_1 \qquad \text{cost}$$

$$\pi_1 = p_1 q_1 - c_1 \qquad \text{profit}$$

$$\pi = \pi_0 + \rho\pi_1 \qquad \text{overall profit, discount factor } 0 < \rho < 1$$

The firm is uncertain about the state of market demand a_0, or unit cost level e_0. To condense exposition, rewrite the overall profit function as a general quadratic form in the choice (control) variables, $u = [q_0, q_1]'$, and random state variables, $x = [a_0, e_0]'$:

$$\pi(x,u) = u'Cx + \frac{1}{2}u'Du \ , \tag{1.1}$$

where some arithmetic shows that:

$$C = \begin{bmatrix} 1 & -1 \\ \rho(1+\alpha) & -\rho(1+\beta) \end{bmatrix}$$

$$D = \begin{bmatrix} -2(b+f+\psi) & 2\psi \\ 2\psi & -2(b+f+\psi) \end{bmatrix} .$$

Note that C is nonsingular (if $\alpha \neq \beta$), and D is negative definite.

Under uncertainty, the firm must choose a value for control vector u that maximizes its expected profit stream, conditional on θ, its available information.

(1.2) $$\max_u E[\pi(x,u)|\theta],$$

which may be written:

(1.3) $$\max_u u'CE(x|\theta) + \frac{1}{2} u'Du,$$

since x is the only stochastic variable.

First order conditions show the optimal control u* to be:

(1.4) $$u^* = -D^{-1}CE(x|\theta).$$

When u* is substituted into (1.2), we compute the expected maximal value for profits

(1.5) $$\max_u E[\pi(x,u)|\theta] = E[x'Sx|\theta] - E[(x-E(x|\theta))'S\ (x-E(x|\theta))|\theta],$$

where $S = -\frac{1}{2} C'D^{-1'}C$, positive definite. To simplify this expression, define:

(1.6) $$L = E(xx'|\theta),$$

(1.7) $$P = E[(x-E(x|\theta))(x-E(x|\theta))'|\theta] = Cov(x|\theta).$$

Thus (1.5) may be rewritten:

(1.8) $$\max_u E[\pi(x,u)|\theta] = tr\{SL\} - tr\{SP\}.$$

To review these results. The stochastic linear-quadratic optimization problem (1.2) yields (1.4) as an optimal control. This control rule is certainty-equivalent as the same control rule is generated under certainty.

Under certainty though, the rule's argument is the actual value of x rather than an estimate of x. In the control literature this result is sometimes called the Separation Theorem -- indicating that the optimal control rule may be computed independent of the process for estimating the value of x. Expected profits, though, remain directly related to the accuracy of the estimate of x through the estimator's variance-covariance matrix P. In the following section, I show how this characteristic determines the optimal choice of state estimator.

2. State Estimation Procedure

Expected profit maximization requires the firm to make two choices. The first, selection of an optimal control rule, is unaffected by risk in the linear-quadratic case. The second, is to select an optimal state estimator. Expected profits, conditioned on the use of an optimal control rule, are given by (1.8). Hence the profit-maximizing firm will choose a state estimator $d^*(\cdot)$ where the estimate $d^*(\theta) = E(x|\theta)$, is such that:

$$(2.1) \qquad d^*(\cdot) = \arg\min_{d} \operatorname{tr}\{SP\}.$$

This is equivalent to saying that the firm selects an estimator that is Bayes with respect to the estimation loss function induced by the firm's quadratic profits. While it is common for economic modellers to use non-Bayes estimators to represent the firm's expectations, I now show how the use of Bayes versus non-Bayes estimators alters the influence of uncertainty or the state estimate, and thus, on the firm's overall decision.

In the sequel I assume the random state vector x to have a normal distribution. This simplifies exposition without too much sacrifice of generality. The firm's information set θ contains two sources of information about the true state value, an a priori distribution and a noisy measurement:

$$x \sim N(\mu, M), \qquad \text{prior distribution,}$$

$$z \sim N(x, R), \qquad \text{noisy measurement.}$$

The optimal state estimate, $d^*(\theta)$ given in (2.1), may be computed through use of Bayes' Rule:

$$d^*(\theta) = E(x|\theta) = \mu + P^*R^{-1}(z-\mu) = P^*R^{-1}z + (I - P^*R^{-1})\mu \ , \tag{2.2}$$

where

$$P^* = (M^{-1} + R^{-1})^{-1} \tag{2.3}$$

is the error covariance matrix of the estimator d^*.

A firm's uncertainty about its market environment is measured by P. To rank this uncertainty, I make the following definition.

<u>Definition</u>: An estimator d_2 with covariance matrix P_2 is more accurate than d_1, with covariance matrix P_1 if $P_1 - P_2$ is positive definite.

This definition will prove useful in the following section where the degree of uncertainty is related to the level of expected profit.

From (2.2) and (2.3), it is easy to see that the level of uncertainty P will, in general, influence a state estimator d (and thus the firm's overall decision) unless:

1) $\mu \equiv z$

or

2) $M^{-1} \equiv 0.$

Since condition (1) is tantamount to saying that there is no uncertainty, i.e., prior $\mu = x$ with probability one, only condition (2) has operational importance. Verbally, condition (2) states that only when the firm's prior distribution for the state is completely uninformative (diffuse), will the state estimator be unaffected by the firm's level of uncertainty. Hence so long as the firm's prior is of some use in identifying the state ($M^{-1} \neq 0$), the firm's state estimator d, and thus its decision u, will depend on its uncertainty about the market environment. The influence of this property on firm behavior is described more fully in the following section.

3. Behavioral Implications

In this section I show how two types of estimators restrict the influence of uncertainty on a firm's decisions. Maximum likelihood estimation, which neglects prior information, suggests the unrealistic property that the location of a firm's distribution of decisions is independent of its uncertainty about the market state. Bayes estimation, on the other hand, which does not neglect prior information, has two appealing properties. It offers the firm a higher level of expected profit. And the degree of uncertainty facing the firm has an intuitively realistic influence on the location of the firm's distribution of decisions. This influence causes firm behavior to appear risk-averse.

To display these properties, first compute the posterior distribution of Bayes state estimates and the resulting control decisions. Integration under Bayes' rule shows that:

$$(3.1) \qquad d^* \sim N(\mu + P^*R^{-1}(z-\mu),\ P^*).$$

Since the firm's optimal decision is given by a linear feedback rule in the state estimate, $u^* = -D^{-1}Cd^*$, the decisions of the firm will be distributed:

$$(3.2) \qquad u^* \sim N(-D^{-1}C[\mu + P^*R^{-1}(z-\mu)],\ D^{-1}CPC'D^{-1'}).$$

For convenience, we can omit the notational baggage engendered by the $-D^{-1}C$ transformation by normalizing $-D^{-1}C = I$. Thus instead of considering decisions to vary proportionally with state estimates, we assume they vary equally.

The importance of both estimator accuracy and estimator type may be investigated through the influence of estimation error covariance P on the location and dispersion of the distribution of firm decisions.

Consider first the effects of estimator type and accuracy on the dispersion of firm decisions. The maximum likelihood estimator d_{ML} has covariance matrix $P_{ML} = R$. Thus if $M^{-1} \neq 0$, we have that:

$$(3.3) \qquad P_{ML} - P^* = R - (M^{-1} + R^{-1})^{-1} \text{ is positive definite.}$$

Hence the distribution of maximum likelihood estimates and induced decisions will be more dispersed than the distribution of Bayes estimates and induced decisions. The following property shows the profitability of maximum likelihood estimation to be less than Bayes estimation.

Property 1: If S and P_1 are positive definite, and P_2 is positive semidefinite, then $tr\{SP_1\} > tr\{SP_2\}$ if and only if P_1-P_2 is positive definite.

Proof: in appendix

Analogously, if a firm using Bayes estimation has a more accurate (less dispersed) estimator than another firm using Bayes estimation, the first firm will gain a higher level of expected profit.

Varying the amount of uncertainty facing the firm, as measured by covariance P, may also cause a translation of the distributions for d and u. As shown in Section 2, the determining feature is whether $M^{-1} = 0$, i.e., P = R. If the maximum likelihood technique is selected, prior information is ignored and this condition for location-invariance is satisfied. From (2.2) we see the expected value of d_{ML} is x, hence:

(3.4) $$E(u_{ML}) \propto E(d_{ML}) = x.$$

Thus x is always the measure of central tendency determining the location of the distribution of a firm's maximum likelihood decisions. So while information accuracy affects the dispersion of the distributions of decisions of a firm using maximum likelihood estimation, it never affects the location of this distribution. Figure 1 illustrates a unidimensional view of this situation. Suppose a firm using maximum likelihood estimation has two prospective levels of uncertainty, i.e., $P_1 = R_1$, $P_2 = R_2$ with $M^{-1} = 0$. Let:

$P_1 - P_2$ be positive definite, then:

$$d^i_{ML} = \mu + P_i R_i^{-1}(z_i - \mu) = z_i , \qquad i=1,2$$

$$E(d^i_{ML}) = E(z_i) = x, \qquad i=1,2$$

Figure 1

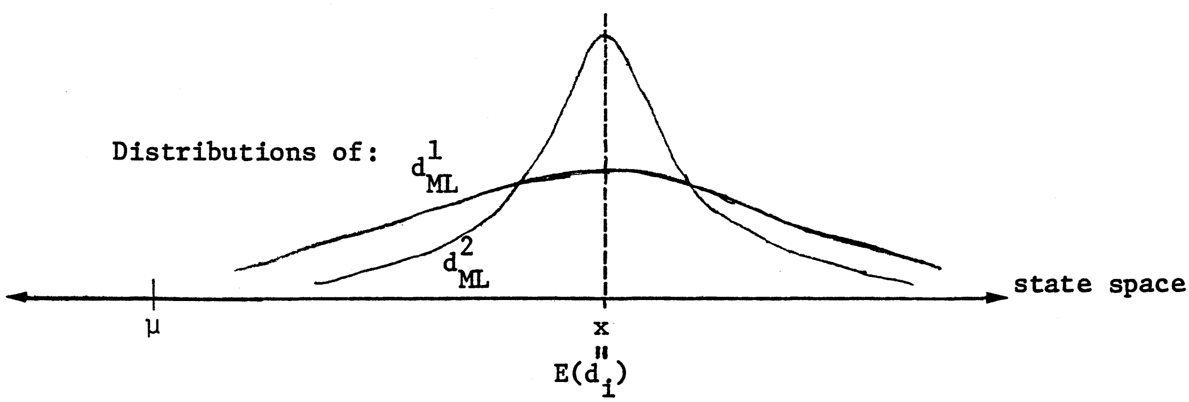

This unsatisfactory situation is remedied if a Bayes estimator is used, $M^{-1} \neq 0$. In this case, the relative weighting, P^*R^{-1} and $I-P^*R^{-1}$, on the data z and the prior μ are influenced by changes in uncertainty. Since $\mu \neq z$ in general, the expected value of d*, and the associated decision u*, shifts as data and prior accuracies change. Crudely described, an increase in data accuracy R^{-1}, keeping prior accuracy M^{-1} fixed, causes greater weight to be placed on the data z in determining d*, and less on the prior μ. As the data become perfectly precise, $R^{-1} \to \infty$ pointwise, the expected value of the state estimate d*, i.e., the location of the distribution of state estimates, converges to the true value x, and the dispersion of this distribution P* becomes degenerate. Formalized, this statement is called Property 2 and is proved in the appendix. Since the certainty decision is conditioned on the true value of the state, x, the Bayes decision u* converges to the certainty decision as the firm's data become completely accurate. In reverse fashion, as data become totally diffuse, $R^{-1} \to 0$ pointwise, the expected value of d* converges to the prior μ, and the expected Bayes decision E(u*) is equivalent to the decision that would be chosen if μ was considered the certain state value.

The multidimensional nature of uncertainty in this problem causes some difficulty, however. Without imposing restrictions on the R and M covariance matrices it is not possible to show uniform convergence (in the common Euclidean norm) of E(d*) to x or μ as R^{-1} varies between 0 and ∞.[3] This is due to the possible existence of non-zero covariance terms in R and M. There are, however, several restrictions that permit us to show uniform Euclidean convergence of E(d*) to μ or x. The first is to assume only one uncertain element: R and M scalars. In this case, convergence will be uniform along the line segment connecting μ with x. This same convergence path arises if R and M are scalar multiples of one another, thus E(d*) will still be located along $\overline{\mu x}$. There is also uniform, componentwise convergence if both R and M are diagonal. In this case, E(d*) always lies within the rectangle having $\overline{\mu x}$ as diagonal. These situations are illustrated in Figure 2. Without restrictions on R or M, E(d*) could have a convergence path like (1). Unidimensional uncertainty, or R and M scalar multiples, restricts the

[3]See Leamer (1973) and Chamberlain and Leamer (1976) for a discussion of this question.

Figure 2

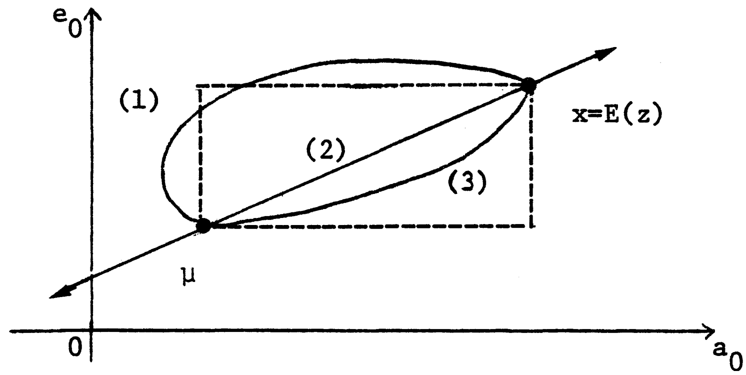

convergence path to (2). If both R and M are diagonal, the convergence path may resemble (3), but always remains within the dashed rectangle.[4]

While nonuniform convergence paths such as (1) may seem paradoxical, this is only because the Euclidean norm does not correspond to the profit norm motivating the firm. While nonzero covariances within the random vector may cause more accurate data to move the expected estimate further away, in the Euclidean sense, from the true value; this movement always improves the firm's expected profits. Thus points on the convergence path are uniformly ordered in terms of expected profit.[5]

To make the following analysis easier to follow, I confine attention to situations where convergence is uniform. This restriction is especially appropriate for making qualitative comparisons between this Bayes model and the models of Hymans (1966), McCall (1967), Baron (1970) and (1971), Sandmo (1971) and Leland (1972) who assume univariate uncertainty. If these authors

---

[4]The analysis surrounding Figure 2 is equivalent to that in portfolio theory for computing risk-return opportunity sets from bundles of securities with uncertain risk and return aspects. See a standard text like Francis and Archer (1979).

[5]Dhrymes (1964) notes similar problems in his investigation of a multi-product firm with uncertain demand.

had allowed unrestricted multivariate uncertainty, they would have encountered similar difficulties in characterizing firm actions with respect to uncertainty.[6]

Using a univariate example, Figure 3 displays how changes in uncertainty influence firm decisions. A fixed prior covariance M is assumed. The firm may have access to data with precision R_1^{-1}, or less accurate data with precision R_2^{-1}, i.e., $R_1^{-1} - R_2^{-1}$ positive definite, which implies $P_2^* - P_1^*$ is positive definite. Now if, without loss of generality we assume $\mu < x$, then

$$E(d_2^*) < E(d_1^*) < x$$

where

$$E(d_i^*) = \mu + P_i^* R_i^{-1}(x-\mu).$$

Figure 3

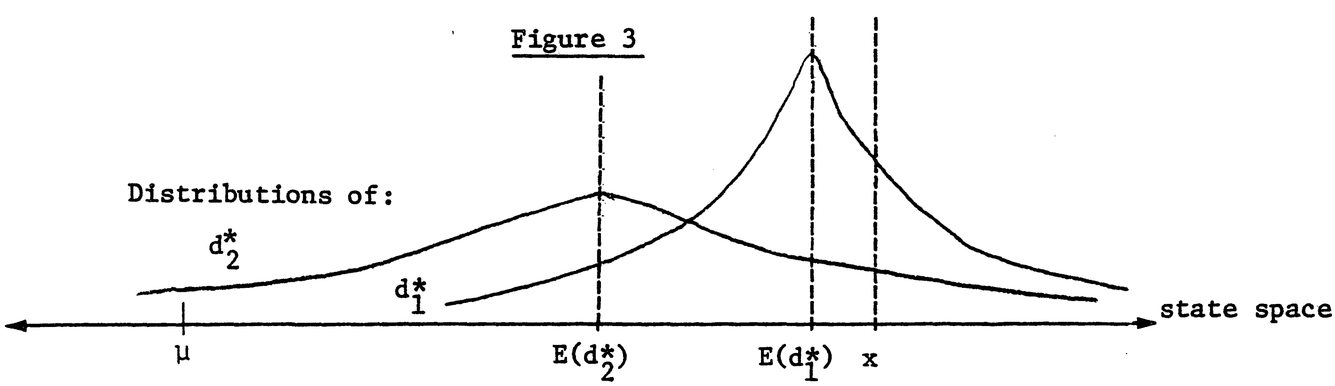

Thus more accurate data lead to a distribution of estimates (decisions) that is both less dispersed than if the data were less precise, but also located unidirectionally "closer" along the path of convergence ($\overline{\mu x}$) to the certainty value (decision).

---

[6]See Leamer (1973).

This behavioral structure is similar to that which arises from risk aversion caused by managerial utility. Variations in firm decisions under traditional risk aversion arise from differing degrees of concavity of the managers' utility function. The more concave the function, the less the firm alters its output decision on the basis of a noisy signal about market conditions. In the model presented here, Bayes decisionmaking causes the degree of a firm's uncertainty to order the amount by which it will react to innovations in the market environment. The key requirement is that firms share similar priors -- perhaps because priors are based past values of the market state that have since been revealed to all firms, while data accuracy differs among firms. In this event, a more knowledgeable firm will make decisions that are more heavily based on the volatile data, while a less knowledgeable firm will make decisions that are based more strongly on the prior. Figure 3 could represent a situation where demand suddenly increases. Firm 1, being more certain, is expected to raise production by more than firm 2 in response to this stimulus. The overall result is that, we should see the overall behavior of knowledgeable firms to be more variable than the behavior of less knowledgeable firms.[7]

4. Concluding Remarks

In this paper I have shown that expected profit maximization and Bayes estimation go hand-in-hand. The result is that even tractable quadratic profit functions with certainty-equivalent decision rules, suggest that a firm's behavior under uncertainty should differ from its certainty behavior.

---

[7] Turnovsky (1969) interprets Bayes learning and decisionmaking as a dynamic partial adjustment process.

The character of this behavior under uncertainty, may resemble risk aversion.

In the introduction I alluded to the complexities that managerial utility may create for econometric work. The reader may wonder whether I have just converted this into a dilemma by promoting Bayes estimation -- since it is widely believed that it is hard to make reasonable restrictions on the form of the prior distribution. My response is that full Bayes is a model that the firm instinctively uses in forming its expectations. However the modeler need not completely identify the firm's full Bayes process. There is a substantial and growing armory of Empirical Bayes techniques that may be used by the econometric modeler to represent the firm's Bayes estimation process.[8] These Empirical Bayes algorithms may be implemented with observable data, and proxy the full Bayes estimator being used by the firm, more faithfully than does maximum likelihood.

[8] See Maritz (1970) or Clarke (1980) for a description.

Appendix

Property 1: If S and P_1 are positive definite, and P_2 is positive semi-definite, then $tr\{SP_1\} > tr\{SP_2\} \Leftrightarrow P_1 - P_2$ is positive definite.

Proof:[9] From a theorem in Anderson (1958), we know there exists a nonsingular matrix R such that: $P_1 = R'R$ and $P_2 = R'\Lambda R$ where the elements λ_i of diagonal matrix Λ are roots to the equation $\det(P_2 - \lambda_i P_1) = 0$.

Thus: $tr\{SP_1\} = tr\{SR'R\} = tr\{RSR'\} = \sum_i d_{ii}$

$tr\{SP_2\} = tr\{SR'\Lambda R\} = tr\{\Lambda^{\frac{1}{2}} RSR' \Lambda^{\frac{1}{2}}\} = \sum_i \lambda_i d_{ii}$.

In addition, note that since RSR' is positive definite, all diagonal terms d_{ii} must be positive.

($\Rightarrow$) $tr\{SP_1\} > tr\{SP_2\}$

$\sum_i d_{ii} > \sum_i \lambda_i d_{ii} \quad \Rightarrow \quad \lambda_i < 1$ for all i.

But $\det(P_2 - \lambda_i P_1) = 0 \quad \Rightarrow \quad c'(P_2 - \lambda_i P_1)c = 0$ for $c \neq 0$.

Now $\lambda_i < 1 \Rightarrow c'P_1c > c'P_2c \Rightarrow c'(P_1 - P_2)c > 0$

only if $P_1 - P_2$ is positive definite

($\Leftarrow$) $P_1 - P_2$ positive definite $\Rightarrow c'(P_1 - P_2)c > 0$ for $c \neq 0$

or $c'P_1c > c'P_2c$. But since we also know:

$c'(P_2 - \lambda_i P_1)c = 0$ for $c \neq 0$, we have $c'P_2c = \lambda_i c'P_1c$

Thus $\lambda_i < 1$ which implies $\sum_i d_{ii} > \sum_i \lambda_i d_{ii}$

or $tr\{SP_1\} > tr\{SP_2\}$

QED

[9] I am indebted to John Geweke for this proof.

<u>Property 2</u>: If M^{-1} is fixed and $R^{-1} \to \infty$ pointwise, then $E(d^*) \to x$ and $P^* \to 0$.

<u>Proof</u>: $R^{-1} \to \infty$ pointwise with M^{-1} fixed implies: $P^{*-1} = R^{-1} + M^{-1} \to R^{-1}$, hence $P^* \to 0$ pointwise and $P^*R^{-1} \to I$ pointwise. From (2.2) we have:

$$E(d^*) = P^*R^{-1}E(z) + (I-P^*R^{-1})\mu$$
$$= P^*R^{-1}x + (I-P^*R^{-1})\mu$$

But $P^*R^{-1} \to I \Rightarrow E(d^*) \to x$

QED

References

Anderson, T. W., An Introduction to Multivariate Statistical Analysis, New York: Wiley, 1958.

Baron, D. P., "Price Uncertainty, Utility, and Industry Equilibrium in Pure Competition," International Economic Review 11, (October 1970): 463-80.

__________, "Demand Uncertainty in Imperfect Competition," International Economic Review 12, (June 1971): 196-208.

Chamberlain, G. and E. E. Leamer, "Matrix Weighted Averages and Posterior Bounds," Journal of the Royal Statistical Society (B) 38, (1976): 73-84.

Christensen, L. R., D. W. Jorgenson and L. J. Lau, "Transcendental Logarithmic Production Frontiers," Review of Economics and Statistics 55, (February 1973): 28-45.

Clarke, R. N., "Differential Game Models of Oligopoly," unpublished Ph.D. dissertation, Harvard University, 1980.

Dhrymes, P. J., "On the Theory of the Monopolistic Multiproduct Firm Under Uncertainty," International Economic Review 5, (September 1964): 239-257.

Francis, J. C. and S. H. Archer, Portfolio Analysis (2nd. ed.), Englewood Cliffs, NJ: Prentice-Hall, 1979.

Hymans, S. H., "The Price-Taker: Uncertainty, Utility, and the Supply Function," International Economic Review 7, (September 1966): 346-56.

Leamer, E. E., "Multicollinearity: A Bayesian Interpretation," Review of Economics and Statistics 55, (1973): 371-380.

Leland, H. E., "Theory of the Firm Facing Uncertain Demand," American Economic Review 62, (June 1972): 278-91.

Maritz, J. S., Empirical Bayes Methods, London: Methuen, 1970.

McCall, J. J., "Competitive Production for Constant Risk Utility Functions," Review of Economic Studies 34, (October 1967): 417-420.

Papandreou, A. G., "Some Basic Problems in the Theory of the Firm," in A Survey of Contemporary Economics, Vol. II, B. F. Haley (ed.). Homewood, Illinois: Irwin, 1952.

Sandmo, A., "Theory of the Competitive Firm Under Price Uncertainty," American Economic Review 61, (March 1971): 65-73.

Simon, H. A., "Dynamic Programming Under Uncertainty With a Quadratic Criterion Function," Econometrica 24, (January 1956): 74-81.