

A Simultaneous Estimation and Variable Selection Rule

Amos Golan*

ABSTRACT

A new data-based method of estimation and variable selection in linear statistical models is proposed. This method is based on a generalized maximum entropy formalism, and makes use of both sample and non-sample information in determining a basis for coefficient shrinkage and extraneous variable identification. In contrast to tradition, shrinkage and variable selection are achieved on a coordinate by coordinate basis, and the procedure works well for both ill and well-posed statistical models. Analytical asymptotic results are presented and sampling experiments are used as a basis for determining finite sample behavior and comparing the sampling performance of the new estimation rule with traditional competitors. Solution algorithms for the non-linear inversion problem that results are simple to implement.

Keywords: Shrinkage Estimator; Maximum Entropy; Extraneous Variables; Squared Error Loss; Data Weighted Prior; Subset Selection

JEL Classification: C13, C14, C5

* American University.

Correspondence: Department of Economics, Roper 200, American University, 4400 Massachusetts Ave., NW, Washington, DC 20016, USA.

E-mail: agolan@american.edu; Tel: (202) 885-3783; Fax: (202) 885-3790.

This work is based on an earlier joint work with George Judge presented at the 1996 joint Statistical Meetings. In addition to thanking George for his infinite contribution and time put into an earlier version of this work, I also thank Leo Breiman, Douglas Miller, Dale Poirier, Arnold Zellner, an Associate Editor and the referee for their helpful suggestions and comments.

1. Introduction

Given a finite and non-experimental data set economists face two basic decisions. The first is the decision about the set of non-extraneous variables and related functions to include in the design matrix. The second is the choice of the estimation rule to be used in recovering estimates of the unknowns in the corresponding parameterized model. Traditionally, the model selection and estimation problems have been separated (Kempthorne, 1987) where the choice of the estimation rule is done after the choice of model and variable selection. In reality, however, the two decisions are interdependent as the sampling performance of the estimation procedure is conditioned by the model selection choice.

The objective of this paper is to construct an estimation rule that *simultaneously* considers the problems of statistical variable selection and parameter estimation. This semi-parametric estimation rule has its roots in information theory and builds on the generalized maximum entropy (GME) and generalized cross entropy (GCE) estimation rules proposed by Golan, Judge and Miller (1996). Two main properties of this proposed estimation and variable selection rule are worth noting. The first is that this method makes use of both sample and non-sample (prior) information on a coordinate basis. The second is that the prior for each coordinate (or variable) is determined endogenously during the optimization. That is, the optimization is done simultaneously with respect to both the posterior and an endogenously determined weight imposed on a convex combination of informative and non-informative priors.

The statistical model and variable selection rules are specified and reviewed in Section 2. The traditional maximum entropy formalism and the GCE estimators are reviewed in Section 3. In Section 4, the proposed flexible data weighted prior (DWP) estimation rule is formulated and its corresponding sampling characteristics are discussed. In Section 5, some sampling experiments are reported. These experiments demonstrate the empirical risk and variable identification performance of the DWP estimator. Section 6 summarizes and concludes this paper.

2. The Problem - Background and a Brief Review

To make exposition easier, consider the traditional linear statistical model. Let us assume that we are unable to directly measure the unknown K -dimensional parameter vector \mathbf{b} . Instead, we observe a T -dimensional vector of noisy sample observations $\mathbf{y} = (y_1, y_2, \dots, y_T)$ that are consistent with the underlying data generation process model

$$\mathbf{y} = X\mathbf{b} + \mathbf{e} \quad (2.1)$$

where X is a fixed $(T \times K)$ full column rank design matrix known to the experimenter.

Further, it is assumed that

Assumption A1. $\mathbf{b} \in B$ where B is a convex set.

Example 1: $B = \{\mathbf{b} \in \mathfrak{R}^K \mid \mathbf{b}_k \in (\underline{z}_k, \bar{z}_k), \quad k=1,2,\dots,K\}$

Given the data generation process described by (2.1), the objective is to derive an estimator that uses *minimum assumptions* on the likelihood structure and simultaneously

identifies the extraneous variables on a coordinate-wise basis. Before developing this estimator, we briefly review some of the current variable selection criteria and models.

Within the context of statistical model (2.1), the variable selection problem may be described in the following way. An investigator has a single and non-experimental sample of data that is known to have the linear functional form of (2.1). Suppose that some covariates are unrelated to the prediction of \mathbf{y} , so the true relationship may be characterized by a lower-dimensional parameter space \mathbf{b}_0 . Consequently, we visualize a K -dimensional parameter space that includes the set of K_0 relevant variables, plus an additional possible $K - K_0$ extraneous variables with coefficients of zero. Thus, the design matrix consistent with the data generation process is a proper subset of the included variables. In terms of variable selection, there are 2^K possible models that can be obtained from the general model (2.1). However, in most cases we employ our knowledge from economic theory to reduce our choice, of possible models, only to those remaining variables that include some uncertainties.

Traditionally, there are two elements in the criterion function for the various variable selection procedures. One element involves a measure of “goodness of fit” while the other involves a penalty for complexity, which is a function of the number of variables K_0 in one of the competing models. Following Zheng and Loh (1995), let $MG(K_0)$ be the measure of goodness of fit for the competing model K_0 , then the various sampling theory estimators \tilde{K}_0 are asymptotically equivalent to

$$\tilde{K}_0 = \arg \min_{0 \leq K_0 \leq K} \{MG(K_0) + h(K)\hat{\mathbf{S}}^2\} \quad (2.2)$$

where $\hat{\mathbf{S}}^2$ is a consistent estimate of \mathbf{S}^2 and $h(K)$ is some non-linear function(e.g., Hocking, 1976; Amemiya, 1980; Laud and Ibrahim, 1995; Mallows, 1973; Miller, 1990; Mitchell and Beauchamp, 1988)¹. For a recent review and development of both parametric and nonparametric approaches to variable selection, within a general discrimination framework, see Lavergne (1998).

In most cases, however, we have some prior knowledge and/or non-sample information that come from economic theory and from understanding our data. The variable selection methods discussed above do not “quantify” this knowledge. But some of this knowledge (priors) may be quantified and incorporated directly. This is addressed in the Bayesian approach to model selection, which also involves setting prior probabilities over the large class of models being considered together with setting the corresponding priors for the parameters of each model (see, for example, Zellner and Vandaele, 1975; George and McCulloch, 1993; Geweke, 1994; Kass and Raftery, 1995; Zellner, 1996b)². The model developed here also uses prior information, but introduces this knowledge in a different way. This is discussed in the next section.

¹ Among the more common methods are the C_p (Mallows, 1973) and the Akaike’s (1974) Information Criterion, AIC. For example, if $h(K) = 2K/T$ and $MG(K_0) = \|\mathbf{y} - \mathbf{X}\mathbf{b}_{K_0}\|^2$, Equation (2.2) reduces to the AIC. For more discussion of the statistical properties of these criteria see for example Shibata (1981) and Zhang (1992). Other criteria, such as Schwarz ‘s (1978) Criterion SC or the Zheng and Loh (1995) generalization of C_p as well as cross validation (Breiman and Spector, 1992; Stone, 1974) and penalized likelihood (Sin and White, 1996) are quite commonly used.

² For a good review and recent developments of the Bayes factors within Bayesian testing and model selection see Berger and Mortera (1999).

3. A Cross Entropy Estimator – Review, Background and Motivation

As an alternative to traditional frequentist shrinkage, pre-test and Bayes estimators for the location vector in (2.1), Golan, Judge and Miller (1996) proposed, for both the symmetric and non-symmetric cases and for both well and ill-posed problems, a new shrinkage estimation rule. This estimation rule is based on the entropy measure of Shannon (1948), a reformulation of the maximum entropy (ME) formalism of Jaynes (1957a; 1957b; 1984), Levine (1980), Shore and Johnson (1980), Skilling (1989) and Csiszár (1991), and the cross entropy principle of Gokhale and Kullback (1978), Good (1963), and Kullback (1959). Before developing the new entropy-based variable selection model, a brief review and background, for both the ME and GME, are presented.

3.1. The Classic Maximum Entropy Model

To provide a basis for understanding the philosophy of the ME approach, consider the following example. Let $\Theta = \{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M\}$ be a finite set and \mathbf{p} be a probability mass function on Θ . The Shannon's information criterion, called entropy, is

$$H(\mathbf{p}) = -\sum_{i=1}^M p_i \log p_i$$
 with $0 \log 0 \equiv 0$. This information criterion measures the uncertainty, or informational content, in Θ which is implied by \mathbf{p} . The entropy-uncertainty measure $H(\mathbf{p})$ reaches a maximum when $p_1 = p_2 = \dots = p_M = 1/M$ and a minimum with a point mass function. Given the entropy measure and structural constraints in the form of moments of the data (distribution), Jaynes (1957a, 1957b) proposed the maximum entropy (ME) method, which is to maximize $H(\mathbf{p})$ subject to these structural constraints. If no constraints (data) are imposed, $H(\mathbf{p})$ reaches its

maximum value and the distribution of the p 's is a uniform one. Thus, if we have partial information in the form of some moment conditions, Y_t ($t=1, 2, \dots, T$), where $T < M$, the maximum entropy principle prescribes choosing the $p(\mathbf{q}_i)$ that maximizes $H(\mathbf{p})$ subject to the given constraints (moments) of the problem. The solution to this underdetermined problem is

$$\hat{p}(\mathbf{q}_i) \propto \exp \left\{ - \sum_t \mathbf{1}_t Y_t(\mathbf{q}_i) \right\} \quad (3.1)$$

where $\mathbf{1}$ are the T Lagrange multipliers.

If prior information, q_i , concerning the unknown p_i exists, then one alternative to the ME approach is to minimize the Kullback-Leibler (K-L) entropy-distance between the post-data weights and the priors (Gokhale and Kullback, 1978). Under this criterion, known as cross entropy (CE), the problem of recovering \mathbf{p} , may be formulated as minimizing the CE subject to the relevant structural constraints (moments). The resulting solution is

$$\tilde{p}(\mathbf{q}_i) \propto q_i \exp \left\{ \sum_t \mathbf{1}_t Y_t(\mathbf{q}_i) \right\} . \quad (3.2)$$

When the prior information q_i has uniform mass, the optimal solutions of the ME and CE problems are identical.

To relate the ME formalism to the more familiar linear model (2.1), consider a special case of this model where B of A1 is $B = \{\mathbf{b} \in \mathcal{R}^1 | \mathbf{b}_k \in [0,1], \sum_k \mathbf{b}_k = 1\}$ and $\mathbf{e} = \mathbf{0}$

Thus,

$$\mathbf{y} = X\hat{\mathbf{a}} \equiv X\mathbf{p} \quad (3.3)$$

and \mathbf{p} is a K -dimensional proper probability distribution. The ME formulation is

$$ME = \begin{cases} \hat{\mathbf{p}} = \arg \max \left\{ - \sum_k p_k \log p_k \right\} \\ \text{s.t. } \mathbf{y} = X\mathbf{p} \text{ and } \sum_k p_k = 1 \end{cases} \quad (3.4)$$

Similarly, the CE formulation is just

$$CE = \begin{cases} \tilde{\mathbf{p}} = \arg \min \left\{ \sum_k p_k \log(p_k/q_k) \right\} \\ \text{s.t. } \mathbf{y} = X\mathbf{p} \text{ and } \sum_k p_k = 1 \end{cases} \quad (3.5)$$

where $I(\mathbf{p}, \mathbf{q}) = \sum_k p_k \log(p_k/q_k)$ is the Kullback-Leibler information, or CE, measure.

The exact CE solution is

$$\tilde{p}_k = \frac{q_k \exp\left(\sum_{i=1}^T \tilde{\mathbf{I}}_i x_{ik}\right)}{\sum_k q_k \exp\left(\sum_{i=1}^T \tilde{\mathbf{I}}_i x_{ik}\right)} \equiv \frac{q_k \exp\left(\sum_{i=1}^T \tilde{\mathbf{I}}_i x_{ik}\right)}{\Omega} \equiv \frac{q_k \exp(\tilde{\mathbf{h}}_k)}{\Omega(\tilde{\mathbf{h}})} \quad (3.6)$$

for some natural K -dimensional vector $\tilde{\mathbf{h}}$. The dual CE counterpart is

$$\inf_{\mathbf{p} \in P} I(\mathbf{p}, \mathbf{q}) = \sup_{\mathbf{I} \in D} \{\mathbf{I}'\mathbf{y} - \log \Omega(X\mathbf{I})\} = \sup_{\mathbf{I} \in D} \{\mathbf{I}'\mathbf{y} - F(\tilde{\mathbf{h}})\} \quad (3.7)$$

for $F(\mathbf{h}) = \log \Omega(\mathbf{h})$ and where $P = \{\mathbf{p}: X\mathbf{p} = \mathbf{y}\}$ is a set of proper (normalized)

distributions satisfying the linear constraints (3.3), and D is the set

$\{\tilde{\mathbf{e}} \in \mathfrak{R}^T : \Omega(X\tilde{\mathbf{e}}) \ll \infty\}$. Having solved for $\tilde{\mathbf{I}}$, one gets

$$\tilde{\mathbf{p}} = \frac{\mathcal{I}F(\tilde{\mathbf{h}})}{\mathcal{I}\tilde{\mathbf{h}}}. \quad (3.8)$$

3.2. The Linear GCE Model

Building on the above, we go back to model (2.1) and assumption A1. Let \mathbf{z}_k of

Example 1 be an M -dimensional vector $\mathbf{z}_k \equiv (\underline{z}_k, \bar{z}_k) = (z_{k1}, \dots, z_{kM})'$ for all k . Instead of

searching directly for the point estimates \mathbf{b} , we view \mathbf{b} as the expected value over some

reference set B . To do so, let \mathbf{p}_k be an M -dimensional proper probability distribution

defined on the set \mathbf{z}_k such that

$$\mathbf{b}_k = \sum_m p_{km} z_{km} \equiv E_{\mathbf{p}_k}[\mathbf{z}_k] \text{ or } \mathbf{b} = E_P[\mathbf{z}]. \quad (3.9)$$

In this way the observed data \mathbf{y} are viewed as the mean process \mathbf{z} with a probability

distribution P that is defined on B (or \mathbf{z}_k 's). Before proceeding, it helps to assume

Assumption A2. $\mathbf{e} \in V$ where V is a symmetric around zero convex set.

Example 2: $V = \{\mathbf{e} \in \mathfrak{R}^T \mid e_t \in (\underline{v}, \bar{v}), \quad t = 1, 2, \dots, T\}$.

Thus, Similar to the \mathbf{b} 's above, each error term is redefined as

$$e_t = \sum_j w_{ij} v_j \equiv E_{w_t}[\mathbf{v}] \text{ or } \mathbf{e} = E_w[\mathbf{v}]. \quad (3.10)$$

Having reparameterized \mathbf{b} and \mathbf{e} , the linear model can be specified as

$$y_t = \sum_{k=1}^K \sum_{m=1}^M z_{km} p_{km} x_{tk} + \sum_j v_j w_{tj}, \text{ or } \mathbf{y} = X E_P[\mathbf{z}] + E_W[\mathbf{v}], \text{ and the GCE rule is}$$

$$GCE = \begin{cases} \tilde{\mathbf{p}} = \arg \min_{\mathbf{p}, \mathbf{w}} \sum_k \sum_m p_{km} \log(p_{km}/q_{km}) + \sum_t \sum_j w_{tj} \log(w_{tj}/u_{tj}) \\ \text{s.t. } \mathbf{y} = X E_P[\mathbf{z}] + E_W[\mathbf{v}]; \sum_m p_{km} = 1; \sum_j w_{tj} = 1 \end{cases}. \quad (3.11)$$

The solution is

$$\tilde{p}_{km} = \frac{q_{km} \exp(\sum_t \tilde{\mathbf{I}}_t z_{km} x_{tk})}{\sum_m q_{km} \exp(\sum_t \tilde{\mathbf{I}}_t z_{km} x_{tk})} \equiv \frac{q_{km} \exp(\sum_t \tilde{\mathbf{I}} z_{km} x_{tk})}{\Omega_k(\tilde{\mathbf{I}})} \quad (3.12)$$

and

$$\tilde{w}_{tj} = \frac{u_{tj} \exp(\tilde{\mathbf{I}}_t v_j)}{\sum_j u_{tj} \exp(\tilde{\mathbf{I}}_t v_j)} \equiv \frac{u_{tj} \exp(\tilde{\mathbf{I}}_t v_j)}{\Psi_t(\tilde{\mathbf{I}})} \quad (3.13)$$

where the prior weights for \mathbf{p}_k are $\mathbf{q}_k = (q_{k1}, \dots, q_{kM})'$ and the corresponding prior weights for \mathbf{w} , consistent with the set of discrete points \mathbf{v} , are $\mathbf{u}_t = (u_{t1}, \dots, u_{tJ})'$. With these estimates we proceed to calculate the point estimates $\tilde{\mathbf{b}}_k \equiv \sum_m z_{km} \tilde{p}_{km}$ and $\tilde{\mathbf{e}}_t \equiv \sum_j v_j \tilde{w}_{tj}$.

Finally, the dual GCE is

$$\inf_{p \in P, w \in W} I(\mathbf{p}, \mathbf{w}; \mathbf{q}, \mathbf{u}) = \sup_{\mathbf{I} \in D} \left\{ \mathbf{I}' \mathbf{y} - \sum_k \log \Omega_k(X' \mathbf{I}) - \sum_t \log \Psi_t(\mathbf{I}) \right\}. \quad (3.14)$$

Solving the dual yields the optimal \mathbf{J} , which in turn yields the point estimates via (3.12) and (3.13). Finally, the Hessian matrix of the GCE problem is positive definite for $\mathbf{p}, \mathbf{w} \gg 0$, and thus satisfies the sufficient condition for a unique global minimum³.

3.3. Discussion

In terms of the specification of \mathbf{z}_k and \mathbf{v} , it is traditional to assume that the elements of \mathbf{b} and \mathbf{e} are finite, and that β_k and e_t are drawn from a finite sample and are usually bounded (see A1-A2). Furthermore, most sciences have a conceptual base for identifying and defining a relevant set of variables, along with the characteristics of the corresponding parameters consistent with a particular problem or data generation process. For example, equality and inequality constraints on β_k commonly arise in many econometric and other scientific applications, and theory or prior empirical results permit the parameters to be signed or specified within a bounded range $[a, b]$. In economics, behavioral and technical coefficients such as marginal productivity, marginal propensity to consume or price elasticities may be classified as nonnegative or positive and naturally bounded. In terms of \mathbf{v} , one possibility is to use the sample (empirical) variance of \mathbf{y} and the three-sigma rule. This is the approach taken here.

Next, one may ask how sensitive are the estimates to the specification of \mathbf{z}_k . The simple answer is that as long as the center of these supports remain unchanged, say zero, the estimates are practically insensitive to changes in the boundary points of the supports. For example, let $\mathbf{z}_k = (-C, 0, C)'$ for each k . Then, a symmetric change of these supports to $\mathbf{z}_k = (-100C, 0, 100C)'$ has practically no significant effect on the estimated \mathbf{b} 's. Since

³ Since for uniform prior distributions the GCE solution reduces to the GME solution, only the GCE

in this work the objective is to identify the extraneous and the non-extraneous variables, the choice of zero as the center point of the supports, for each k , seems to be natural. However, it is emphasized that for small data sets, the estimates may be sensitive to a change in the center of the supports z .

3.4. Comparison with Other Estimation Rules

First, we note that, unlike the usual Stein-like estimators (e.g., Stein, 1955, 1981; James and Stein, 1960; Judge and Bock, 1978; Brown, 1966; Bock, 1988), the GME-GCE estimators shrink estimates on a coordinate-by-coordinate basis. For example, if the support vector \mathbf{z}_k is centered at zero, the coefficients close to zero receive maximum shrinkage⁴.

Second, we note that the GCE formulation, that leads to post-data means, has many of the characteristics of the standard Bayesian conjugate analysis where the posterior mean is a matrix-weighted combination of the prior means and the OLS estimates. In the GCE approach however, the posterior means are weighted combinations of the data and the priors within the supports. Consequently, the posterior means are always within the bounds of the supports Z and \mathbf{v} ⁵.

Finally, it is important to note the relationship between the GCE-GME entropy based estimation rules and Zellner's seminal BMOM approach (e.g., Zellner 1996a, 1997; Tobias and Zellner, 1999). Like the GCE method, the objective behind the BMOM

method is presented here. But unlike traditional Bayes estimators, we always specify the support spaces that bound the possible parameter space.

⁴ For a detailed comparison and discussion of other entropy and non-entropy regularization methods, as well as the maximum entropy on the mean, see for example Donoho, Johnstone, Hoch and Stern (1992), Golan, Judge and Miller (1996, Chp. 8) and Bercher, Besnerais and Demoment (1996).

method is to estimate the unknown parameters with minimum assumptions on the likelihood function. As stated by Zellner (1997, p. 86), “The BMOM approach is particularly useful when there is difficulty in formulating an appropriate likelihood function. Without a likelihood function, it is not possible to pursue traditional likelihood and Bayesian approaches to estimation and testing. Using a few simple assumptions, the BMOM approach permits calculation of post-data means, variances and other moments of parameters and future observations”.

In the BMOM approach one starts by maximizing the continuous entropy function (of a density function) subject to some T side conditions (moments) and normalization. This yields the average log-height of the density function, which is the least informative density given these side-conditions. Thus, under the BMOM approach one works with the most uninformed, that is maxent, post data density for the parameters. A further advantage of the BMOM approach is that many types of side conditions can be (and have been) utilized in order to obtain post data densities for parameters and future observations. These side conditions include bounds on the parameters’ values, bounds on the error terms’ ranges, inequality restrictions, fractile conditions, and moment conditions. For more innovative applications of the BMOM see, for example, LaFrance (1999) and the discussion in Zellner (1999). For a full comparison of the traditional Bayes and the BMOM approaches see Zellner (1997, Table 1) and for an information theoretic derivation of Bayes’ Theorem that provides a link between maximum entropy procedures and Bayes’ Theorem see Zellner (1988) and Soofi (1996).

⁵ For other non-Bayesian methods see for example the recent work on the empirical likelihood (e.g., Owen, 1990, Qin and Lawless, 1994), weighted least squares and the GMM (e.g., Hellerstein and Imbens, 1999; Imbens, Johnson and Spady, 1998).

Even though the basic objectives of the BMOM and the GCE are similar, they differ in their inputs. Following Zellner, the basic inputs in the BMOM are the (i) data, (ii) prior information, (iii) mathematical form of the model, (iv) moments of the parameters and future values, and (v) the Maxent principle. The two basic differences in the GCE inputs are in the way prior information is incorporated and the assumption (input) on the moments of the parameters and future values. Specifically, under the GCE rule, the prior information is incorporated via three routes: the support spaces (\mathbf{z} and \mathbf{v}), the priors in the objective functional (\mathbf{q} and \mathbf{u}) and other restrictions that may enter as additional equality/inequality constraints. Further, the GCE rule is specified such that each data point enters directly in the optimization (rather than a quadratic form of the data or the moments' input). Therefore, the moment requirement can be thought of as a “weak” moments' requirement in the sense that the sample's moments may be different (up to a certain distance which is data dependent) from the underlying population's moments. Obviously, our choice of an estimator is problem and data specific and strongly depends on the amount of information we have with respect to a specific problem, model and data.

4. A Flexible, Prior, Data-Based Estimator

In this section the GCE estimator is extended so that it embraces both the model identification and estimation objectives discussed earlier. To accomplish this task the GCE estimator is reformulated such that the extraneous variables are identified and the **b**

parameter space may be truncated, or all variables are retained but the coefficients of extraneous variables are shrunk toward zero⁶.

Given the objective of identifying possible extraneous variables, we specify the discrete support space \mathbf{z}_k for each β_k to be symmetric around zero and employ the GCE estimator with a unit mass prior on zero. While effective for identifying extraneous variables, as an estimator of $\mathbf{\beta}$ under a squared error loss measure, this does not lead to a minimax rule. If, on the other hand, we wish to have estimators that, under a squared error loss measure, are minimax and thus superior to traditional estimators over all, or part, of the parameter space, the GCE estimator that uses uniform priors (or similarly, the GME) is a good rule. However, if our objective is to shrink but not necessarily eliminate the extraneous variables, and simultaneously produce an estimator that has a good sampling performance over the whole range of the possible parameter space, then we may combine the GME and GCE (or the GCE with uniform priors and the GCE with spike priors) estimators. This is the formulation proposed here.

We start by specifying each \mathbf{z}_k and \mathbf{v} to be *symmetric around zero*, with large lower and upper bounds for \mathbf{z} and the three-empirical-sigma-rule for \mathbf{v} , so that β_k and e_t are contained in a fixed interval with arbitrarily high probability. We also specify, as a possible alternative for each β_k , a “spike” prior, with a point mass at $z_{km} = 0$, for each $k = 1, 2, \dots, K$. Thus, a flexible, data-based prior is specified such that for each β_k coordinate either a spike prior at the $z_{km} = 0$, a uniform prior over the discrete support space \mathbf{z}_k , or any convex combination of the two can result. The goal is to produce a

⁶ For other approaches with similar objective see for example the Bayes-Stein estimators for the normal k-means reviewed and developed in Zellner and Vandaele (1975).

natural adaptive statistical method that is data based and free of subjective choices except for the bounds on the support spaces. Because we are interested in a formulation that puts pressure on the data by including both uniform and spike prior alternatives, or some convex combination of the two, we are not able to make use of the conventional cross entropy formalism. Consequently, given the problem at hand, and within the entropy approach, we consider an extended entropy formulation. But because the \mathbf{q}_k prior alternatives cannot be introduced as structural constraints, we must find a way to introduce them in the objective function. To accomplish this, consider the following variant of the GCE formulation of Section 3.

$$\begin{aligned} \underset{\mathbf{p}, \mathbf{p}^{\tilde{\mathbf{a}}}, \mathbf{w}}{\text{Min}} \quad I(\mathbf{p}, \mathbf{p}^{\tilde{\mathbf{a}}}, \mathbf{w}) \equiv & \sum_k (1 - \mathbf{g}_k) \left(\sum_m p_{km} \log p_{km} \right) / [(1 - \mathbf{g}_k)(-\log M) + \mathbf{g}_k \mathbf{x}] \\ & + \sum_k \mathbf{g}_k \sum_m p_{km} \log(p_{km} / q_{km}) + \sum_{k,h} p_{kh}^{\mathbf{g}} \log(p_{kh}^{\mathbf{g}} / q_{kh}^{\mathbf{g}}) + \sum_{t,j} w_{tj} \log(w_{tj} / u_{tj}) \end{aligned} \quad (4.1)$$

subject to

$$y_t = \sum_{k,m} z_{km} p_{km} x_{tk} + \sum_j v_j w_{tj} \quad (4.2)$$

and

$$\sum_m p_{km} = 1, \quad \sum_h p_{kh}^{\mathbf{g}} = 1, \quad \sum_j w_{tj} = 1. \quad (4.3)$$

and where the prior weight (prior mixture) \mathbf{g}_k is

$$\mathbf{g}_k \equiv \sum_h z_h^{\mathbf{g}} p_{kh}^{\mathbf{g}} \quad (4.4)$$

with $z_1^{\mathbf{g}} = 0$ and $z_H^{\mathbf{g}} = 1$ always and where $\mathbf{x} \equiv -\sum_m q_{km} \log q_{km}$. Further, *except for the point mass prior \mathbf{q}_k , all other priors (i.e., \mathbf{q}^γ and \mathbf{u}_i) are specified to be uniform.*

4.1 The Criterion

Having specified the new data-weighted prior (DWP) estimation rule, we now discuss the explicit structure of the objective function (4.1). There are four parts to this objective function. Just as in Section 3 (the upper part of Equation 3.11), the last element on the Right Hand Side (RHS) of (4.1) is the entropy of the noise component. Since \mathbf{v} is specified to be symmetric, and equally spaced, support around zero, this part of the objective function shrinks the noise components toward zero. The bounds for the errors' supports are just $\pm 3\bar{\mathbf{s}}_y$ where $\bar{\mathbf{s}}_y$ is the empirical standard deviation of the sample. The first two elements on the RHS of (4.1) relate to the uniform and informative (spike) priors respectively, with corresponding weight $\mathbf{g}_k \in [0,1]$, where $\gamma_k = 0$ implies use of a uniform (uninformative) prior, while $\gamma_k = 1$ implies use of a spike prior. In terms of the first RHS element, as noted in Section 3, when using a uniform prior over the support points there is no need to specify this prior explicitly. The third element in the objective function relates to the total entropy of the weights \mathbf{g} . Finally, the first element of the objective function is scaled by the convex combination of the negative of the entropy of the priors. If, for example, $\gamma_k = 0$, then it is scaled by the negative entropy of the uniform

distribution (for coordinate k), which is $-\ln(\mathbf{M})$. If, on the other hand, $\gamma_k = 1$, then it is scaled by $\mathbf{x} \equiv -\sum_m q_{km} \log q_{km}$, which equals zero iff \mathbf{q}_k is a spike prior. This scaling, or normalization, is needed because without it, the first two elements are of different magnitude and sign. Thus, without this normalization, the uniform prior *always* takes over, and the DWP reduces to the GME. This is because the first element of the objective function is always non-positive, while the second element is always non-negative. With this scaling, however, the two parts of the objective can “communicate” in the sense that both are of the same sign and magnitude, and there is a data-based choice of the prior for each coordinate that is most consistent with the sample information. In this way this estimator simultaneously chooses the mixture of the two alternative priors on a coordinate-by-coordinate basis and uses this information, along with the data, to determine the shrinkage and to provide estimates of the unknown parameters. We note in conclusion that if one wishes to avoid the scaling carried in (4.1), the first component on the RHS of (4.1) can be substituted for the GCE with uniform priors. However, the formulation used here, is computationally superior and more efficient.

4.2 Solution and Properties of the Estimator

The solution to the optimization problem (4.1–4.3) yields

$$\tilde{p}_{km} = \frac{q_{km}^{\tilde{g}_k/A_k} \exp\left(\frac{1}{A_k} \sum_t \tilde{\mathbf{I}}_t x_{tk} z_{km}\right)}{\sum_m q_{km}^{\tilde{g}_k/A_k} \exp\left(\frac{1}{A_k} \sum_t \tilde{\mathbf{I}}_t x_{tk} z_{km}\right)}, \quad (4.5)$$

$$\tilde{w}_{ij} = \frac{u_{ij} \exp(\tilde{\mathbf{I}}_t v_j)}{\sum_j u_{ij} \exp(\tilde{\mathbf{I}}_t v_j)} \equiv \frac{u_{ij} \exp(\tilde{\mathbf{I}}_t v_j)}{\Psi_t(\tilde{\mathbf{I}})} \quad (4.6)$$

and

$$\tilde{\mathbf{g}}_k \equiv \sum_h \tilde{p}_{kh} z_h^g \quad (4.7)$$

where $\mathbf{1}$ reflects the T -Lagrange multipliers for the data equations (4.2) and

$$A_k \equiv \frac{1 - \tilde{\mathbf{g}}_k}{(\tilde{\mathbf{g}}_k - 1) \log M + \tilde{\mathbf{g}}_k \mathbf{x}} + \tilde{\mathbf{g}}_k. \quad \text{As before, (3.9) and (3.10) provide the basis for}$$

recovering estimates of \mathbf{b} and \mathbf{e} . As $\gamma_k \rightarrow 0$ the designated prior becomes more uniform, with the estimates approaching those of the GME estimator. For large values of γ_k (above 0.49), the GCE estimator with an informative (in our case, point mass on zero) prior takes over.

The following conditions ensure consistency and asymptotic normality of the DWP estimation and variable selection rule.

- i) The errors' support \mathbf{v} is symmetric around zero (see A2).
- ii) The support space \mathbf{z}_k spans the true value of each one of the unknown parameters \mathbf{b} .

Further, the support has a finite lower and upper bounds z_{k1} and z_{kM} respectively (see A1).

- iii) The errors are iid.

- iv) $\lim_{T \rightarrow \infty} T^{-1} X' X$ exists and is nonsingular.

The proof of consistency and asymptotic normality follows directly from the empirical likelihood approach (Owen, 1990, 1991; Qin and Lawless, 1994; Golan and Judge, 1996). Similarly, these proofs can be established by following Mittelhammer and Cardell (1996).

In general, under the above four conditions, $\sqrt{T}(\hat{\mathbf{b}} - \mathbf{b}) \xrightarrow{d} N(\mathbf{0}, Q_{DWP})$ where $Q_{DWP} = \mathbf{S}^2 [\lim T^{-1} (X'X)^{-1}]$ is the asymptotic covariance matrix for the DWP. Since \mathbf{b} is a continuous function of \mathbf{I} (the Lagrange multipliers) this statement is an immediate extension of Qin and Lawless (1994, Lemma 1 and Theorem 1). Finally, $\mathbf{S}^2 = [1/(T - K)] \sum_t \hat{e}_t^2$ with $\hat{e}_t = \sum_j v_j \hat{w}_{tj}$ is a consistent estimator of the variance.

4.3. Diagnostics and Inference

Given this testing basis, we propose a test to compare the restricted DWP model with the unrestricted one. With identical, and symmetric around zero, supports \mathbf{z} and \mathbf{v} for both problems, we follow Owen (1990; 1991), Qin and Lawless (1994), and Golan and Judge (1996) and define an empirical entropy (or expected log-likelihood) ratio statistic

$$\ell_E \equiv L_{E(\mathbf{b}_0)} - L_{E(\tilde{\mathbf{b}})} \quad (4.8)$$

where $L_{E(\mathbf{b}_0)}$ applies to the optimal value of the DWP estimator's objective function when restricting $\mathbf{b} = \mathbf{b}_0 = \mathbf{0}$, while $L_{E(\tilde{\mathbf{b}})}$ is the optimal value of the DWP estimator's objective function when the \mathbf{b} 's are not restricted (i.e., model 4.1-4.3). Under the null, $2\ell_E$ is asymptotically distributed as a $\chi^2_{(K)}$. Similarly, one can test the hypothesis

whether each one of the K covariates is extraneous. In that case, the entropy-ratio statistic $2\ell_E$ has a limiting distribution of $\mathbf{c}_{(1)}^2$.

We now use the DWP rule to provide a normalized (information) entropy measure, which will be used as the criterion for identifying and eliminating the extraneous variables from the design matrix. Let the normalized entropy for each coordinate k be

$$S(\tilde{\mathbf{p}}_k) \equiv \frac{-\sum_{m=1}^M \tilde{p}_{km} \log \tilde{p}_{km}}{\log M} \quad (4.9)$$

where $S(\tilde{\mathbf{p}}_k) = 1$ implies maximum uncertainty and $S(\tilde{\mathbf{p}}_k) = 0$ implies full certainty.

Next, we can relate the normalized entropy, or information measure, S , with the \mathbf{c}^2 statistic. We start by pointing out that each component of (4.1) is a basic cross-entropy measure for some proper set of probabilities. Ignoring the γ 's for simplicity sake, let \tilde{p}_{km} be any appropriate estimator of p_{km} and let $\{\tilde{\mathbf{p}}_k\}$ be a set of M probabilities over the M -dimensional support space \mathbf{z}_k for each one of the K coordinates. Then, we define the statistics

$$I(\tilde{\mathbf{p}}; \mathbf{q}) \equiv \sum_{k=1}^K \sum_{m=1}^M \tilde{p}_{km} \log(\tilde{p}_{km} / q_{km}), \quad (4.10)$$

and

$$\mathbf{c}_{(M-1)}^2 \equiv M \sum_{m=1}^M \frac{1}{q_{km}} (\tilde{p}_{km} - q_{km})^2, \text{ for each } k=1, 2, \dots, K. \quad (4.11)$$

A second order approximation of $I(\tilde{\mathbf{p}}; \mathbf{q})$ is just the entropy-ratio statistic for evaluating $\tilde{\mathbf{p}}$ vs. \mathbf{q} discussed above. That is, for each k ,

$$I(\tilde{\mathbf{p}}_k; \mathbf{q}_k) \cong \frac{1}{2} \sum_{m=1}^M \frac{1}{q_{km}} (\tilde{p}_{km} - q_{km})^2, \text{ for each } k=1, 2, \dots, K \quad (4.12)$$

so $2M$ times the entropy-ratio statistic corresponds to $\mathbf{c}_{(M-1)}^2$. Given this relationship, we use the normalized entropy measure, for each covariate k , in order to test the hypothesis $H_0: \mathbf{b}_k = 0$ for each k . Thus, $\mathbf{c}_{(M-1)}^2 = 2MI(\tilde{\mathbf{p}}_k; \mathbf{q}_k)$ for the spike priors \mathbf{q}_k .

Based on the above we can define some rules to identify the extraneous variables. First, the following variable selection rule is defined. If $[1 - S(\tilde{\mathbf{p}}_k)] \geq 0.99$, then variable k is classified as extraneous. Conversely, if $[1 - S(\tilde{\mathbf{p}}_k)] < 0.99$, then a real classification of the variable is suggested. An $S(\tilde{\mathbf{p}}_k) = 0.99$ as opposed to 1.0 is used to allow for insignificant computer rounding effects.

Alternatively, a second identification criterion is proposed. This criterion is based on the weight $\hat{\mathbf{g}}_k$ in (4.1). If for a particular variable $\tilde{\gamma}_k < 0.5$, then a non-extraneous variable is identified. As $\tilde{\gamma}_k$ decreases, the strength of this identification is increased. Conversely, a $\tilde{\gamma}_k \geq 0.5$ suggests the variable belongs in the extraneous category.

Finally, we relate our choice of $1 - S(\tilde{\mathbf{p}}_k) \geq 0.99$ to the \mathbf{c}^2 measure which, in that case, yields,

$$\mathbf{c}_{(M-1)}^2 = 2MI(\tilde{\mathbf{p}}_k; \mathbf{q}_k^u) = 2M \ln(M)[1 - S(\tilde{\mathbf{p}}_k)] \quad (4.13)$$

where $\mathbf{q}_k'' = 1/M$ for all $m=1, 2, \dots, M$. Thus, for $M=5$ (the number of support points used in all the experiments presented in the following sections), the 0.99 rule implies

$$\mathbf{c}^2 = 15.9 > (\mathbf{c}_{(4)}^2)^c = 13.3 \text{ for } \mathbf{a}=0.01.$$

4.4 Comparison With the ML Approach

Having specified an estimator with a variable selection component and demonstrated its asymptotic properties, it is possible to compare its sampling performance with other conventional estimation rules and with other variable selection criteria. To achieve this comparison, instead of working with each data point directly, it is possible to follow common practice and transform the data representation (4.2) to a moment representation

$$\mathbf{X}'\mathbf{y} = \mathbf{X}'\mathbf{X}\mathbf{b} + \mathbf{X}'\mathbf{e} = \mathbf{X}'\mathbf{XZ}\mathbf{p} + \mathbf{X}'\mathbf{V}\mathbf{w} \quad (4.2a)$$

where all previous definitions of \mathbf{b} and \mathbf{e} follow and \mathbf{V} is a matrix composed of the T vectors \mathbf{v} .

LEMMA 1: If we let $\mathbf{v}=\mathbf{0}$ and substitute the K pure moment conditions $\mathbf{X}'\mathbf{y} = \mathbf{X}'\mathbf{XZ}\mathbf{p}$ for the T data-consistency equations (4.2) then, under restriction (ii), the resulting estimates of the problem (4.1)–(4.3) are equivalent to the least squares (ML) estimates.

This proof is immediate since the constraints $\mathbf{X}'\mathbf{y} = \mathbf{X}'\mathbf{XZ}\mathbf{p}$, that must be satisfied, are just the first order conditions for the least squares (ML) estimates.

LEMMA 2: For any finite data set, the $approx \text{var}(\tilde{\mathbf{b}}_{k(DWP)}) \leq approx \text{var}(\tilde{\mathbf{b}}_{k(GCE)})$ for all k .

The logic for the proof is as follows. Golan, Judge and Miller (1996, Chap. 7)

show that for all $\mathbf{v} \neq \mathbf{0}$, the $approx \text{var}(\tilde{\mathbf{b}}_{k(GCE)}) < approx \text{var}(\hat{\mathbf{b}}_{k(LS/ML)})$ for all finite data sets and for all k . Following their logic, for expository purposes, let X be an orthonormal matrix and the error covariance $\Sigma_e = \sigma^2 \mathbf{I}_T$. Under these assumptions the approximate DWP covariance is

$$Cov \tilde{\mathbf{b}}_{(DWP)} = \tilde{\mathbf{S}}^2 \tilde{\Sigma}_z (\tilde{\Sigma}_z + \tilde{\Sigma}_v)^{-2} \tilde{\Sigma}_z \cong \tilde{\mathbf{S}}^2 (X'X)^{-1} \quad (4.14)$$

where $\tilde{\Sigma}_z$ and $\tilde{\Sigma}_v$ are the respective covariances for $\tilde{\mathbf{p}}$ and $\tilde{\mathbf{w}}$. The k th element of (4.14) is

$$Var(\tilde{\mathbf{b}}_{k(DWP)}) = \tilde{\mathbf{S}}^2 \left(\frac{\tilde{\mathbf{S}}_z^2}{\tilde{\mathbf{S}}_z^2 + \tilde{\mathbf{S}}_v^2} \right)^2 \quad (4.15)$$

where at the limit, $\tilde{\mathbf{S}}_v^2 = 0$, and (4.14) is just the LS (ML) variance $\hat{\mathbf{S}}^2$. To compare the DWP and GCE variances for the k th element, we need to evaluate $\tilde{\mathbf{S}}_z^2$, which is just

$$\tilde{\mathbf{S}}_z^2 = \sum_m \tilde{p}_{km} z_{km} - \left(\sum_m \tilde{p}_{km} z_{km} \right)^2. \quad (4.16)$$

The variance (4.16) reaches its maximum at a given \mathbf{z}_k for $\tilde{p}_{km} = 1/M$, for all $m=1, 2, \dots, M$. In Section 3, given the data constraints, the GCE objective (with uniform priors) makes the estimates \hat{p}_{km} as uniform as possible. Alternatively, the spike priors in the DWP estimator make the \tilde{p}_{km} the least possible uniform given the data.

Consequently, the less uniform the \tilde{p}_{km} , the smaller \tilde{S}_z^2 in (4.16). This establishes the relationship $\text{approx var}(\tilde{\mathbf{b}}_{k(DWP)}) \leq \text{approx var}(\hat{\mathbf{b}}_{k(GCE)})$.

To summarize, most variable selection models (criteria) require choosing among 2^K different models while imposing some unknown smoothing function [Zheng and Loh, (1995)] or smoothing parameter (e.g., C_p and AIC). Alternatively, the DWP is a data-driven penalty function estimator that is based on weak distributional assumptions. It does not require 2^K steps and/or a pre-specified smoothing parameter. It is a one step estimation rule that requires the pre-specified support spaces. Finally, note that both the popular AIC variable selection rule and DWP are based on the Kullback-Leibler information (cross entropy) measure. If we make use of the pure moment condition (4.2a) within the DWP problem (4.1)–(4.3), then it is easy to show a proportional relationship between the DWP and an AIC criterion for each of the 2^K models⁷.

5. Sampling Experiments

In this section we report some results of Monte Carlo sampling experiments to indicate the small sample performance of the DWP estimator and, under a squared error loss measure, compare it with other relevant traditional and shrinkage estimators. For the estimator comparisons, we continue to consider the multivariate estimation problem where the ML estimator is $\delta^0(\mathbf{b})$ and has risk $\rho(\mathbf{b}, \mathbf{d}) = \sigma^2 \text{tr}(\mathbf{X}'\mathbf{X})^{-1}$. We note here that in this work we restrict ourselves to the squared error loss measure and do not attempt to

analyze predictive performance⁸. For a comprehensive comparison of the predictive powers of GME and some Bayesian methods (MELO and BMOM) see the recent work of Perloff and Shen (1999).

5.1 The Symmetric Case

In order to provide an experiment that involves a mean square prediction error loss and a comparison with the well-known positive rule Stein shrinkage estimator, we use the orthonormal K -mean linear statistical model

$$\mathbf{y} = \mathbf{X}\mathbf{b} + \mathbf{e} = \mathbf{X}\mathbf{S}^{-1/2}\mathbf{S}^{1/2}\mathbf{b} + \mathbf{e} = \mathbf{A}\mathbf{q} + \mathbf{e} \quad (5.1)$$

where $\mathbf{S}^{1/2}$ is a positive definite symmetric matrix with $\mathbf{S}^{1/2}\mathbf{S}^{1/2} = \mathbf{S} = \mathbf{X}'\mathbf{X}$, $\mathbf{A}'\mathbf{A} = \mathbf{I}_K$ and $\mathbf{q} = \mathbf{S}^{1/2}\mathbf{b}$. In the experiment $\mathbf{e} \sim N(\mathbf{0}, \sigma^2\mathbf{I}_T)$, $\sigma = 1$ and $K = 4$. The risk

$$E\left[(\tilde{\mathbf{q}} - \mathbf{q})'(\tilde{\mathbf{q}} - \mathbf{q})\right] = E\left[(\tilde{\mathbf{b}} - \mathbf{b})' \mathbf{X}' \mathbf{X} (\tilde{\mathbf{b}} - \mathbf{b})\right] \quad (5.2)$$

yields a weighted loss function in terms of \mathbf{b} and results in the *mean square prediction error* criterion that is often used in econometrics to evaluate performance. The parameter space investigated is $\mathbf{q} = c\mathbf{d}_i$, where $\mathbf{d}_1 = (\mathbf{q}_1, 0, 0, 0)$, $\mathbf{d}_2 = (\mathbf{q}_1, \mathbf{q}_2, 0, 0)$ and $\mathbf{d}_3 = (\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3, 0)$. The scalar c is chosen so that the parameter vector length $(\mathbf{q}'\mathbf{q})^{1/2} = 0, 1, 2, \dots, 66$. For selected values of the $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space, 5,000 samples

⁷ For a nice discussion of the AIC criterion and its relationship to the CE criterion and other information criteria, within the model selection, see Lavergne (1998).

of size $T = 10, 30$ and 100 were generated and empirical estimator risks under a $\|\delta(\mathbf{y}) - \mathbf{q}\|^2$ measure were obtained. For the DWP estimator, $\mathbf{z}'_k = (-100, -50, 0, 50, 100)$ and $\mathbf{v}' = (-3\bar{\mathbf{S}}_y, 0, \bar{\mathbf{S}}_y)$ where $\bar{\mathbf{S}}_y$ is the empirical standard deviation of each sample. Further, $\mathbf{q}'_k = (0, 0, 1, 0, 0)$ for each k , which means we are putting point mass at zero for each θ in the case of the GCE estimator, and $\mathbf{u}' = (.33, .33, .33)$ for each t . For comparison purposes, the risk for the ML, positive rule Stein (PRS), GME and GCE estimators are reported. To make the ML approach fully comparable to the DWP, we need to use the constrained ML where the constraints specified (for each β) are the lower and upper bounds of \mathbf{z} . But because we use very wide bounds, the constrained solution is equivalent to the unconstrained solution and therefore we refer to it as ML.

The risk for the PRS over the $(\mathbf{q}, \mathbf{q})^{1/2}$ parameter space was numerically evaluated using an algorithm by Bohrer and Yancey (1984). The \mathbf{z}_k and \mathbf{v} support spaces noted above are also used for the GME and GCE estimators. It is worth noting that the support spaces for \mathbf{z}_k were chosen to reflect very wide bounds. Increasing these bounds did not change the estimates.

5.1.1 Experiment 1 - $\mathbf{d}'_1 = (\mathbf{q}_1, 0, 0, 0)$

5.1.1a Variable Identification

First, we focus on the variable selection objective with a design matrix involving one non-extraneous and three extraneous variables (Table 1). The last three columns of

⁸ In future work we will investigate the predictive performance of these estimation rules where the parameters of each method are restricted to fall within the same bounds.

Table 1 provide information that forms a basis for identifying the correct design matrix for different points in the $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space. The column labeled $\tilde{\mathbf{g}}$ identifies the weight between the uniform and the point mass prior for each coordinate at each point in the $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space. Note at the origin the weight is 0.5, which signifies that the informative prior with point mass at zero is the active prior. This choice remains active over the whole $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space for the extraneous variables. For the non-extraneous variable, as $(\mathbf{q}'\mathbf{q})^{1/2}$ increases, the weight on the point mass prior decreases and finally all weight is allocated to the uniform prior.

[Table 1 – about here]

Using the normalized entropy measure, column 5 reports the probability of correctly identifying each of the variables. Note at the origin, where all coordinates are extraneous, and the point mass prior is consistent with the location parameters, the probability of correct identification is about 0.98. Over the parameter space this identification probability is maintained for the extraneous variables. For the *non-extraneous* variable x_1 , over the range $(\mathbf{q}'\mathbf{q})^{1/2} \in (1, 2)$, the DWP rule underfits, and the probability of correct variable identification is on average less than 0.5. However, for $(\mathbf{q}'\mathbf{q})^{1/2} \geq 3$, the probability of correct identification, on average, approaches 1.0. In terms of the normalized entropy measure, these results mirror those reported under the $\tilde{\mathbf{g}}$ column. For the non extraneous variable, x_1 , the normalized entropy measure, which reflects the probability of being an extraneous variable, decreases as $(\mathbf{q}'\mathbf{q})^{1/2}$ increases, and thus reinforces the $\tilde{\mathbf{g}}_k$ measure in identifying this variable from its extraneous

counterparts. One reason for these nice results is that when all but a single variable are extraneous, a model that identifies these extraneous variables correctly, must also identify the single non-extraneous variable correctly. The next set of experiments investigate the more complex case of more than a single non-extraneous variable.

5.1.1b Empirical Risks

In this section, we assume that an investigator wishes to use the DWP formulation as an estimator that shrinks but does not eliminate. To show the statistical consequences of this rule, we present in Figure 1 empirical risks over the $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space. As a basis of comparison, we present the corresponding risks for the ML and PRS estimators. The comparison with the PRS points out the risk implications of Stein-like estimators where all coordinates are shrunk by the same amount versus estimators such as GCE, GME and DWP, where shrinkage may vary from coordinate to coordinate. Thus, in this experiment we contrast the performance of the estimators where extraneous variables exist and are shrunk but not identified and removed. Variability and bias results for this experiment (for the DWP) are presented in column 2 of Table 1.

[Figure 1 – about here]

Under this sample design the GCE estimator is risk superior to the PRS estimator over the whole range of the $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space. At the origin the empirical risks of the PRS estimator and the GCE estimator, with spike prior at zero, are 1.472 and 0.53, respectively. From this point in the $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space the risk of the positive part Stein increases sharply, reaches 2.8 at 10 and finally becomes equal to the ML risk of 4.0 around $(\mathbf{q}'\mathbf{q})^{1/2} = 60$. Alternatively, the GCE empirical risk increases more slowly and

finally reaches 3.37 as $(\mathbf{q}'\mathbf{q})^{1/2}$ approaches 60. Consequently, in these sampling experiments the GCE estimator is risk superior, over the range of the \mathbf{q} parameter space, to both the ML and PRS estimators. If, instead of using a point mass prior at zero, we had used a uniform prior over the elements in \mathbf{q}_k , the empirical risk of the GME estimator is about 3.90 over the whole $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space.

Alternatively, the empirical risk of the DWP estimator is 0.246 at the origin where all variable coordinates are correctly shrunk toward zero, and then reaches a maximum of 2.61 at $(\mathbf{q}'\mathbf{q})^{1/2} = 3$, where \mathbf{q}_1 continues to be shrunk toward zero and the bias is a major part of the risk component. From there on the risk decreases sharply and stabilizes at about 1.20 for the remainder of the parameter space. In this range of the parameter space, where \mathbf{q}_1 is shrunk relatively little, the coefficients of the extraneous variables $\mathbf{q}_2, \mathbf{q}_3$ and \mathbf{q}_4 continue to have a maximum shrinkage toward zero. The DWP estimator identifies and shrinks the extraneous variables correctly and is superior to all the other estimators over the whole parameter space⁹.

Finally, it is important to remember that under the GME, GCE and DWP rules, specific bounds are imposed on all of the parameters' values. Naturally, these bounds may have a significant impact on the sampling results. Such bounds can of course be introduced in traditional Bayes and other sampling theory approaches. To provide a fair comparison, the bounds used in this paper covered a very large range of the parameter spaces such that the restricted and unrestricted ML yielded the same results.

⁹ These experiments were repeated for different sample sizes but to save space are not reported here. For example, for T=30 the MSE for the DWP is about 3.2 time smaller than that of the ML over the whole parameter space. For T=100, the DWP maintains its superiority as well. In this case it is between 2.5 to 3.2 times smaller than the ML over the parameter space.

Nevertheless, one should keep in mind that the use of bounds can cause moments to exist and eliminate outlying estimates. For example, Perloff and Shen (1999) show that when estimating $\gamma=1/\beta$ with β restricted to the range $0 < a < \beta < b$, this restriction precludes estimates of β to take on values that are close to zero, implying that the effect on the estimated values could be quite large. Consequently, the results presented here should be interpreted with these qualifications in mind. To summarize, there is a need for much more research in order to identify and establish the exact impact of imposing bounds on the properties of any estimator. As the method developed here, as well as the GME/GCE, uses pre-specified bounds on the parameters, the effects of these bounds on the estimator's properties is a subject for future studies.

5.1.2 Experiment 2 - $\mathbf{d}'_3 = (\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3, 0)$

The previous experiment pretty much tells the estimation and variable selection story because the DWP results evolve on a coordinate-by-coordinate basis. Thus, risk and variable selection results for any mix of non-extraneous and extraneous in the full design matrix can be approximated from the results of Section 5.1.1.

To give a sense of the risk results for different ratios of non-extraneous to extraneous variables, we present the empirical risk functions for the DWP and ML rules in Figure 2. Note that, over a small portion of the parameter space (for small values of \mathbf{q}_k), the DWP rule with 25 percent extraneous variables underfits and is inferior to the ML and Stein rules. However, over most of the $(\mathbf{q}'\mathbf{q})^{1/2}$ parameter space, the DWP rule is risk superior. Finally, in terms of variable selection, the $\tilde{\mathbf{g}}_k$ weight parameter and the normalized entropy measure work as in Section 5.1.1a. Both the non-extraneous and

extraneous variables are identified, over the range of the $\boldsymbol{\eta}$ parameter space, with probabilities ≥ 0.95 .

[Figure 2 – about here]

5.1.3 Experiment 3 - The χ^2 Error Distribution

To demonstrate the robustness of the DWP estimator in a non-Gaussian setting, we repeated the experimental design of Section 5.1 with $\chi^2_{(4)}$ random errors, normalized to have a unit variance. The results for this case (not reported here)¹⁰ are similar in structure and magnitude to those presented in Figure 1 and Table 1. In terms of identifying the extraneous and non-extraneous variables, the DWP mirrors the results reported in Table 1. In particular, at least 93 percent of the extraneous variables are identified over all points in the $(\boldsymbol{\eta}\boldsymbol{\eta})^{1/2}$ parameter space.

5.2 The Non-Symmetric Case

5.2.1 Experiment 4 – High Condition Number

Consider for this experiment the general linear statistical model $\mathbf{y} = X\mathbf{b} + \mathbf{e}$, where the ML estimator $\hat{\mathbf{a}}^b$ is distributed as $N_K(\hat{\mathbf{a}}, \mathbf{S}^2(X'X)^{-1})$, and $X'X$ is a positive definite matrix. For constructing the design matrix we use the condition number's definition $\mathbf{k}(X'X) = \mathbf{p}_1/\mathbf{p}_2$ which is the ratio of the largest and smallest singular values of the design matrix X , with columns scaled to unit length. As $\mathbf{k}(X'X)$ increases, traditional estimators, such as the ML estimator, become unstable and have low precision. For a review of different regularization methods that exhibit relatively good risk characteristics

see, for example, Hoerl and Kennard (1970), O'Sullivan (1986), Breiman (1995), Tibshirani (1996), and Titterton (1985).

The sampling experiment for the non-symmetric case is similar to Experiment 1, reported in Section 5.1.1. However, to reflect a design matrix more in line with data that are by and large non-experimentally generated, instead of a condition number of $\mathbf{k}(X'X)=1$ for the design matrix, we now specify a very moderate condition number of $\mathbf{k}(X'X)=90$ ¹¹. Under the SEL, the ML risk is approximately 47.5 and the iterative ridge estimator risk is approximately 14.1. The risk performance of the DWP and GCE rules mirrors that for the well-posed case, where $\mathbf{k}(X'X)=1$. The DWP risk starts at about 0.21, increases to a maximum of 2.18 for $(\hat{\mathbf{a}}'\hat{\mathbf{a}})^{1/2} = 2$, and decreases rapidly to 1.0 from about $(\hat{\mathbf{a}}'\hat{\mathbf{a}})^{1/2} = 12$ in the parameter space. Unlike the DWP, the GCE risk increases monotonically with $(\mathbf{b}\mathbf{b})^{1/2}$ and is unbounded (or bounded by the support space Z). Selected points on the parameter space for the DWP estimator are reported in Table 2. Note in comparing the MSE for the well-posed design matrix in Table 1 to the moderately ill-posed design of Table 2, that the performance of the DWP rule actually improves as the condition number $\mathbf{k}(X'X)$ increases.

[Table 2 – about here]

The results of this case are basically similar in nature to the empirical risk results of Figure 2 and are in each case, greatly superior to the risk results for the ML and ridge

¹⁰ All results discussed in this Section are available, upon request, from the author.

estimators. The risk results from experiments with higher and lower condition numbers were virtually the same as those reported in the table and are not reported here.

In terms of identifying extraneous variables, the sampling performance for the case of a $k(X'X) = 90$ design is reported in Table 2. Again, the results for the ill-posed case, mirror those of the well-posed case reported in Table 1. Note the probability (or frequency) of identifying the extraneous variables exceeds .95 in all cases.

5.2.2 Experiment 5 - Ill-conditioned Design and t_3 Error Distribution

Within the linear model and the above framework, in this experiment we investigate the small sample performance of a moderately ill-conditioned design matrix and a non-normal error process. The design matrix consists of $K = 4$ covariates, generated such that $\mathbf{x}_4 = \mathbf{x}_1 + 0.15\mathbf{e}^*$, where \mathbf{e}^* is a normal $(0, 1)$ random vector. This experiment follows George and McCulloch (1993). The errors are generated from a student- t distribution with three degrees of freedom, normalized to a unit variance. The results plotted in Figure 3 exhibit in general the risk behavior of the previous experiments. The ML empirical risk is 46.167 and the GME risk is 38.43. Thus, the DWP estimator exhibits superior risk performance over the range of $(\mathbf{b}\mathbf{b})^{1/2}$ parameter space. If $\tilde{\mathbf{g}}$ or S are used as variable selection measures, the performance of the experiments in the previous sections is duplicated.

[Figure 3 – about here]

¹¹ For example, the commonly used “Longley” aggregated employment data (Longley, 1967) have a condition number of approximately 11,100; a highly collinear case. Our choice of 90 is completely arbitrary with the objective of maintaining a very moderate, but yet realistic, level of collinearity.

5.3. Experiment 6

To include experiments where the design matrix has more variables, we report two experiments that duplicate those carried out by Judge et al. (1987). The experiments involve the linear statistical model $\mathbf{y} = A\mathbf{q} + \mathbf{e}$ with nine location parameters, four extraneous variables, a (20×9) design matrix A , where $A'A = I_9$, an error process, \mathbf{e} , distributed as $N(\mathbf{0}, \mathbf{S}^2 I_{20})$ with $\mathbf{S}^2 = 1$ and 500 replications¹². The supports \mathbf{z}_k and \mathbf{v} are specified as before. The first experiment, Model 1, involves the location vector $\mathbf{q} = (5, 4, 3, 2, 1, 0, 0, 0, 0)'$ and the second, Model 2, involves $\mathbf{q} = (10, 8, 3, 2, 1, 0, 0, 0, 0)'$.

Table 3 compares the empirical risk of six estimators: ML; DWP; the positive part Stein (1981), \mathbf{q}^+ , that shrinks all coefficients toward zero; AIC (Akaike, 1974), SC (Schwarz, 1978); and C_p (Mallows, 1973). The ML estimator is close to the theoretical risk of 9. The traditional and the Stein \mathbf{q}^+ variable selection estimators have empirical risks that are superior to the ML but are *significantly inferior* to the DWP that shrinks but does not eliminate variables. In terms of variables selection, the DWP results are presented in Table 4.

[Tables 3-4 – about here]

The results presented here are consistent with the earlier experiments where the higher the parameter space $(\mathbf{q}'\mathbf{q})^{1/2}$, the better the DWP performs relative to the other estimators. Further, increasing the $(\mathbf{q}'\mathbf{q})^{1/2}$ to 37 where $\mathbf{b} = (25, 20, 15, 10, 5, 0, 0, 0, 0)'$ yields risk of 6.32 while the risk for \mathbf{q}^+ practically equals the ML risk, which is 8.8. In

terms of variable identification, the non-extraneous variables are *always* identified while the extraneous variables are identified at least 80 percent of the time.

6. Summary Remarks

In this work a new simultaneous variable selection and estimation rule is developed and investigated. This new rule provides a basis for identifying the non-extraneous and extraneous variables in the design matrix of a linear statistical model and simultaneously yields good estimates. The result is a simple, consistent, one-stage estimator, based on *one* sample of data and a K variable design specification that leads to a basis for semi-parametric entropy-based inference. This data-based procedure is based on weak distributional assumptions and it performs well for both ill and well-posed problems, non-Gaussian error distributions and small samples of data. In this approach coefficient shrinkage and variable elimination are data determined and done on an individual coordinate basis. Further, the choice of prior is data based and endogenously determined. Consequently, the method provides a simple way of introducing and evaluating prior information in the estimation and variable selection process. In contrast to other shrinkage-variable selection procedures, that require a tuning parameter for variable identification and/or determining the degree of shrinkage, upper and lower bounds on the estimated parameters are specified here.

This method is applicable to a wide range of econometric-statistical models (linear and nonlinear) and flexible enough to simultaneously cope with a variety of model specification uncertainties. Solution algorithms for these types of non-linear inversion

¹² Due to the larger size of each sample (relative to most of the previous experiments), the sampling experiment is performed for 500 samples rather than for 5000 samples for each point on the parameter space.

problems are available and easy to implement.

References

- Akaike, H. (1974), A New Look at the Statistical Model Identification, *IEEE Transactions on Automatic Control*, 19, 716–723.
- Amemiya, T. (1980), Selection of Regressors, *International Economic Review*, 21, 331–354.
- Bercher, J. F., G.Le Besnerais and G. Demoment, (1996), The Maximum Entropy on the Mean Method, Noise and Sensitivity, in Skilling J and Sibisi S (Ed.) *Maximum Entropy and Bayesian Studies* (Kluwer).
- Berger, J. O. and J. Mortera (1999), Default Bayes Factors for Nonnested Hypothesis Testing, *Journal of the American Statistical Association*, 94, 542–554.
- Bock, M.E. (1988), Shrinkage Estimators: Pseudo-Bayes Rules for Normal Mean Vectors, in S.S. Gupta and J.O. Berger (ed.), *Statistical Decision Theory and Related Topics IV*, Springer Verlag, New York, 281–298.
- Bohrer, R. and T.A. Yancey (1984), Algorithms for Numerical Evaluation of Stein-like and Limited Translation Estimators, *Journal of Econometrics*, 25, 235–239.
- Breiman, L. (1995), Better Subset Regression Using the Nonnegative Garrote, *Technometrics*, 37, 373–384.
- Breiman, L. and P. Spector (1992), Submodel Selection and Evaluation in Regression. The X-Random Case, *International Statistical Institute Review*, 60, 291-319.
- Brown, L.D. (1966), On the Inadmissibility of Invariant Estimators of One or More Location Parameters, *Annals of Mathematics and Statistics*, 37, 1087–1136.
- Csiszár, I. (1991), Why Least Squares and Maximum Entropy? An Axiomatic Approach to Inference for Linear Inverse Problems, *The Annals of Statistics*, 19, 2032–2066.

- Donoho, D. L., I. M. Johnstone, J. C. Hoch, and A. S. Stern (1992), Maximum Entropy and the Nearly Black Object, *Journal of the Royal Statistical Society, Ser. B*, 54, 41-81.
- George, E. and R. McCulloch (1993), Variable Selection in Gibbs Sampling, *Journal of the American Statistical Association*, 88(423), 881–889.
- Geweke, J. (1994), Variable Selection and Model Comparison in Regression, No. Working Paper 539, Federal Reserve Bank of Minneapolis.
- Gokhale, D.V. and S. Kullback (1978), *The Information in Contingency Tables*, Marcel Dekker, New York.
- Golan, A. and G. Judge (1996), A Maximum Entropy Approach to Empirical Likelihood Estimation and Inference, (unpublished paper), University of California, Berkeley (presented at the 1997 Summer Econometric Society meetings).
- Golan, A., G.G. Judge, and D. Miller (1996), *Maximum Entropy Econometrics: Robust Estimation with Limited Data*, John Wiley & Sons, New York.
- Good, I.J. (1963), Maximum Entropy for Hypothesis Formulation, Especially for Multidimensional Contingency Tables, *Annals of Mathematical Statistics*, 34, 911–934.
- Hellerstein, J.K., and G.W. Imbens (1999), Imposing Moment Restrictions from Auxiliary Data by Weighting, *The Review of Economics and Statistics*, 81, 1-14.
- Hocking, R.R. (1976), The Analysis and Selection of Variables in Linear Regression, *Biometrics*, 32, 1–51.
- Hoerl, A.E. and R.W. Kennard (1970), Ridge Regression: Biased Estimation for Non-orthogonal Problems, *Technometrics*, 1, 55–67.

- Imbens, G.W., Johnson, P. and R.H. Spady, Information-Theoretic Approaches to Inference in Moment Condition Models, *Econometrica* 66 (1998), 333-357.
- James, W. and C. Stein (1960), Estimation with Quadratic Loss, in Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, Berkeley, CA, 361–379.
- Jaynes, E.T. (1957a), Information Theory and Statistical Mechanics, *Physics Review*, 106, 620–630.
- Jaynes, E.T. (1957b), Information Theory and Statistical Mechanics II, *Physics Review*, 108, 171–190.
- Jaynes, E.T. (1984), Prior Information and Ambiguity in Inverse Problems, in D.W. McLaughlin (ed.), *Inverse Problems*, SIAM Proceedings, American Mathematical Society, Providence RI, 151–166.
- Judge, G., G. Yi, T. Yancey, and T. Teräsvirta (1987), The Extended Stein Procedure for Simultaneous Model Selection and Parameter Estimation, *Journal of Econometrics*, 35, 375–391.
- Judge, G.G. and M.E. Bock (1978), *The Statistical Implications of Pre-Test and Stein-Rule Estimators in Econometrics*, North-Holland, Amsterdam.
- Kass, R. and A.E. Raftery (1995), Bayes Factors, *Journal of the American Statistical Association*, 90, 773–795.
- Kempthorne, P.J. (1987), Estimating the Mean of a Normal Distribution with Loss Equal to Squared Error Plus Complexity Cost, *Annals of Statistics*, 15, 1389–1400.
- Kullback, J. (1959), *Information Theory and Statistics*, John Wiley & Sons, New York.

- LaFrance, J. T. (1999) Inferring the Nutrient Content of Food with Prior Information, *American Journal of Agricultural Economics* , 81: 728-734.
- Laud, P.W. and J.G. Ibrahim (1995), Predictive Model Selection, *Journal of the Royal Statistical Society, Ser. B*, 57, 247–262.
- Lavergne P. (1998), Selection of Regressors in Econometrics: Parametric and Nonparametric Methods, *Econometric Reviews*, 17, 227–273.
- Levine, R.D. (1980), An Information Theoretical Approach to Inversion Problems, *Journal of Physics, A*, 13, 91–108.
- Longley, J. (1967), An Appraisal of Least Squares Programs from the Point of the User, *Journal of the American Statistical association*, 62, 819–841.
- Mallows, C.L. (1973), Some Comments on C_p , *Technometrics*, 15, 661–665.
- Miller, A.J. (1990), *Subset Selection in Regression*, Chapman and Hall, London.
- Mitchell, T.J. and J.J. Beauchamp (1988), Bayesian Variable Selection in Linear Regression (with discussion), *Journal of the American Statistical Association*, 83, 1023–1036.
- Mittelhammer, R. and S. Cardell (1996), On the Consistency and Asymptotic Normality of the Data Constrained GME Estimator of the GML, (mimeo), Washington State University, Pullman, WA.
- O'Sullivan, F. (1986), A Statistical Perspective on Ill-Posed Inverse Problems, *Statistical Science*, 1, 502–527.
- Owen, A. (1990), Empirical Likelihood Ratio Confidence Regions, *The Annals of Statistics*, 18, 90–120.

- Owen, A. (1991), Empirical Likelihood for Linear Models, *The Annals of Statistics*, 19, 1725–1747.
- Perloff, J. and Z. Shen (1999), Maximum Entropy and Bayesian Approaches to the Ratio Problem, (unpublished paper), University of California, Berkeley.
- Qin, J. and J. Lawless (1994), Empirical Likelihood and General Estimating Equations, *The Annals of Statistics*, 22, 300–325.
- Schwarz, G. (1978), Estimating the Dimensions of a Model, *The Annals of Statistics*, 6, 461–464.
- Shannon, C.E. (1948), A Mathematical Theory of Communication, *Bell System Technical Journal*, 27, 379–423.
- Shibata, R. (1981), An Optimal Selection of Regression Variables, *Biometrika*, 68, 45–54.
- Shore, J.E. and R.W. Johnson (1980), Axiomatic Derivation of the Principle of Maximum Entropy and the Principle of Minimum Cross-Entropy, *IEEE Transactions on Information Theory*, IT-26(1), 26–37.
- Sin, C-Y and H. White (1996), Information Criteria for Selecting Possibly Misspecified Parametric Models, *Journal of Econometrics*, 71, 207-225.
- Skilling, J. (1989), The Axioms of Maximum Entropy, in J. Skilling (ed.), *Maximum Entropy and Bayesian Methods in Science and Engineering*, Kluwer Academic, Dordrecht, 173–187.
- Soofi, E. (1996) Information Theory and Bayesian Statistics, , in D. A. Berry, K. M. Chaloner and J. K. Geweke (eds.), *Bayesian Analysis in Statistics and Econometrics: Essays in Honor of Arnold Zellner*, John Wiley, New-York, 179-189.

- Stein, C. (1955), Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution, in Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, Berkeley, CA, 197–206.
- Stein, C. (1981), Estimation of the Mean of a Multivariate Normal Distribution, The Annals of Statistics, 9, 1135–1151.
- Stone, M. (1974), Cross-validatory Choice and Assessment of Statistical Predictions, Journal of the Royal Statistical Society, Ser. B, 36, 111–147.
- Tibshirani, R. (1996), Regression Shrinkage and Selection via the Lasso, Journal of the Royal Statistical Society, Ser. B, 58, 267–288.
- Titterton, D.M. (1985), Common Structures of Smoothing Techniques in Statistics, International Statistical Review, 53, 141–170.
- Tobias, J. and A. Zellner (Forthcoming), Further Results on the Bayesian Method of Moments Analysis of Multiple Regression Model, International Economic Review.
- Zellner, A. (1988) Optimal Information Processing and Bayes's Theorem, American Statistician, 42: 278-284.
- Zellner, A. (1996a) Bayesian Method of Moments/Instrumental Variables (BMOM/IV) Analysis of Mean and Regression Models, in J. C. Lee, A. Zellner and W. O. Johnson, eds., Prediction and Modelling Honoring Seymour Geisser, Springer Verlag..
- Zellner, A. (1996b) Models, Prior Information, and Bayesian Analysis, Journal of Econometrics, 75: 51-68.

- Zellner, A. (1997) The Bayesian Method of Moments (BMOM): Theory and Applications, in T. Fomby & R. Hill, eds., *Advances in Econometrics*, 12: 85-105.
- Zellner, A. (1999) New Information-Based Econometric Methods in Agricultural Economics: Discussion, *American Journal of Agricultural Economics* , 81: 742-746.
- Zellner, A. and W. Vandaele (1975) Bayes-Stein Estimators for k-means, Regression and Simultaneous Equation Model, in S. E. Fienberg & A. Zellner, eds., *Studies in Bayesian Econometrics and Statistics, In Honor of Leonard J. Savage*, North-Holland.
- Zhang, P. (1992), On the Distributional Properties of Model Selection Criteria, *Journal of the American Statistical Association*, 87, 732–737.
- Zheng, X. and W.-Y. Loh (1995), Consistent Variable Selection in Linear Models, *Journal of the American Statistical Association*, 90, 151-156.

Table 1

Performance of the DWP estimator for selected points where $\mathbf{q} = (q, 0, 0, 0)$ and x_2, x_3, x_4 are extraneous and the ML $MSE(\mathbf{d}^b) = 4$

$(q/q)^{1/2}$	MSE	$tr\{Cov(\tilde{\mathbf{q}})\}$	\mathbf{g}				Frequency of Correctly Identifying				$1 - S(\tilde{p}_k)$			
			x_1	x_2	x_3	x_4	x_1	x_2	x_3	x_4	x_1	x_2	x_3	x_4
0	0.246	0.246	.50	.5	.5	.5	.98	.97	.98	.98	.999	.999	.999	.999
1	1.159	.50	.49	.5	.5	.5	.12	.97	.98	.97	.99	.999	.999	.999
2	2.576	1.30	.47	.5	.5	.5	.47	.97	.97	.97	.97	.999	.999	.999
3	2.617	1.62	.43	.5	.5	.5	.84	.97	.97	.97	.94	.999	.999	.999
4	2.182	1.46	.40	.5	.5	.5	.98	.96	.96	.97	.91	.999	.999	.999
5	1.891	1.29	.38	.5	.5	.5	1.00	.97	.96	.97	.89	.999	.998	.999
6	1.802	1.30	.35	.5	.5	.5	1.00	.96	.96	.97	.87	.999	.999	.999
10	1.548	1.28	.26	.5	.5	.5	1.00	.96	.96	.97	.81	.998	.998	.999
20	1.276	1.22	.14	.5	.5	.5	1.00	.96	.95	.96	.69	.998	.998	.998
60	1.202	1.20	.01	.5	.5	.5	1.00	.96	.96	.97	.58	.998	.998	.999

Table 2

Performance of the DWP estimator for $\mathbf{b} = (\mathbf{b}_1, 0, 0, 0)$, $T = 10$ and $\mathbf{k}(X'X) = 90$

$(\mathbf{b}\mathbf{b})^{1/2}$	MSE	$tr\{Cov(\tilde{\mathbf{q}})\}$	$\boldsymbol{\xi}$				Frequency of Correctly Identifying				$1 - S(\tilde{\rho}_k)$			
			x_1	x_2	x_3	x_4	x_1	x_2	x_3	x_4	x_1	x_2	x_3	x_4
0	0.209	0.209	.50	.5	.5	.5	.97	.999	.98	.96	.998	1.0	.999	.999
1	1.096	0.52	.49	.5	.5	.5	.16	.998	.97	.96	.99	1.0	.999	.999
2	2.180	1.26	.46	.5	.5	.5	.58	.995	.96	.95	.97	1.0	.999	.998
3	2.051	1.38	.43	.5	.5	.5	.91	.995	.96	.95	.94	1.0	.998	.998
4	1.689	1.21	.40	.5	.5	.5	.99	.999	.96	.95	.91	1.0	.998	.998
6	1.521	1.19	.34	.5	.5	.5	.999	.999	.96	.95	.87	1.0	.998	.998
10	1.238	1.06	.26	.5	.5	.5	1.0	1.0	.95	.95	.80	1.0	.998	.998
20	1.074	1.02	.13	.5	.5	.5	1.0	1.0	.95	.95	.69	1.0	.998	.998
60	0.996	0.996	.02	.5	.5	.5	1.0	1.0	.96	.95	.58	1.0	.998	.998

Table 3

Empirical risks of a range of competing estimators for the symmetric case

	Empirical Risk of					
	ML	\hat{q}^+	AIC	SC	C_p	DWP
Model 1	8.77	8.18	8.50	8.73	8.64	7.25
Model 2	8.89	8.63	8.03	8.40	8.88	7.00

Table 4

Variable selection results of the DWP estimator for the symmetric case

Model	Estimator	Frequency of Identifying Correctly									Risk
		x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	
Model 1	DWP	1.00	.99	.94	.89	.39	.83	.82	.80	.77	7.25
Model 2	DWP	1.00	1.00	.93	.73	.39	.81	.82	.82	.82	7.00