Generalized S-Estimators

Christophe CROUX, Peter J. ROUSSEEUW, and Ola HÖSSJER*

In this article we introduce a new type of positive-breakdown regression method, called a generalized S-estimator (or GS-estimator), based on the minimization of a generalized M-estimator of residual scale. We compare the class of GS-estimators with the usual S-estimators, including least median of squares. It turns out that GS-estimators attain a much higher efficiency than S-estimators, at the cost of a slightly increased worst-case bias. We investigate the breakdown point, the maxbias curve, and the influence function of GS-estimators. We also give an algorithm for computing GS-estimators and apply it to real and simulated data.

KEY WORDS: Breakdown point; Influence function; Maxbias curve; Regression analysis; Robustness.

1. INTRODUCTION

In the linear model one often has to cope with outliers, which can make the classical least squares (LS) estimator highly unreliable. In fact, even a single outlier can destroy the LS estimate. Many alternative methods have been proposed. Commonly used are M- and GM-estimators (see, for example, Hampel, Ronchetti, Rousseeuw, and Stahel 1986), but their breakdown point goes down to zero when the dimension increases. The least median of squares (LMS) and least trimmed squares (LTS) estimators (Rousseeuw 1984) have a 50% breakdown point but a low asymptotic efficiency. A generalization is given by S-estimators (Rousseeuw and Yohai 1984), which can attain an efficiency up to 33% (Hössjer 1992). Both MM-estimators (Yohai 1987) and τ-estimators (Yohai and Zamar 1988) can attain arbitrarily high efficiency without losing their 50% breakdown point, but they pay for this with an increased bias. Note that MM-estimators need a high-breakdown start, for which we can use one of the estimators discussed herein.

In this article we introduce a new class of regression estimators, called *generalized S-estimators* (or *GS-estimators*), which can have a 50% breakdown point like S-estimators, but attain a much higher efficiency. As a special case of GS-estimators, we propose the *least quartile difference* (LQD) estimator, which we define as

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \ Q_n(r_1, \ldots, r_n), \qquad (1)$$

where r_i is the residual of the ith case and

$$Q_n = \{ |r_i - r_j|; i < j \}_{\binom{h_p}{2} \cdot \binom{n}{2}} \qquad (2)$$

is a scale estimator proposed by Rousseeuw and Croux (1993). Expression (2) means that Q_n is the $\binom{h_p}{2}$th order statistic among the $\binom{n}{2}$ elements of the set $\{ |r_i - r_j|; i < j \}$. Here $h_p = [(n + p + 1)/2]$, where p is the number of regression parameters. It turns out that the objective function (2) can be computed very quickly using an efficient algorithm. Therefore, the LQD can be easily implemented by adapting an existing LMS program. The Gaussian efficiency of the

LQD regression is shown to be 67.1%, which is more than twice the efficiency of any S-estimator with a 50% breakdown point. Moreover, the LQD does not require the choice of tuning constants.

An important property of GS-estimators is that their objective function does not depend on the intercept term. (The intercept can be estimated afterwards, with high statistical efficiency and consuming negligible computation time.) Another advantage of GS-estimators is that they are well suited for models with an asymmetric error distribution, unlike the usual S-estimators of which the objective function depends only on r_i through $|r_i|$; hence positive residuals are attached the same importance as negative residuals of the same size. Therefore, GS-estimators are more generally applicable. Note that the pairwise differences $r_i - r_j$ have a symmetric distribution even when the r_i themselves do not.

2. ROBUSTNESS AT FINITE SAMPLES

We will work with a linear model denoted by

$$y_i = \beta_1 x_{i1} + \cdots + \beta_{p-1} x_{i,p-1} + \alpha + \text{error},$$
$$\text{for} \quad i = 1, \ldots, n. \quad (3)$$

The parameter to be estimated is $\theta = (\beta, \alpha) \in \mathbb{R}^p$, where $\beta \in \mathbb{R}^{p-1}$ is the slope and α is the intercept. Our observations are of the form $\mathbf{z}_i = (\mathbf{x}_i, y_i) = (\mathbf{u}_i, 1, y_i) \in \mathbb{R}^{p+1}$. This means that the actual explanatory variables are combined in a vector, $\mathbf{u}_i \in \mathbb{R}^{p-1}$. (We require of course that $p \geq 2$.) As usual, we assume that $n/p > 5$ to avoid the curse of dimensionality.

We define a generalized S-estimator $\hat{\beta}$ as

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \ s_n(\beta), \qquad (4)$$

where $s_n(\beta)$ is based on the residuals $r_i = y_i - \beta' \mathbf{u}_i - \alpha$ through the equation

$$\binom{n}{2}^{-1} \sum_{i<j} \rho \left(\frac{r_i - r_j}{s_n(\beta)} \right) = k_{n,p}. \qquad (5)$$

To avoid having multiple solutions (or no solutions), it is better to define

$$s_n(\beta) = \sup \left\{ s > 0; \binom{n}{2}^{-1} \sum_{i<j} \rho \left(\frac{r_i - r_j}{s} \right) \geq k_{n,p} \right\}. \qquad (6)$$

* Christophe Croux is Research Assistant with the Belgian National Fund for Scientific Research, and Peter J. Rousseeuw is Professor, both at the Department of Mathematics and Computer Science, Universitaire Instelling Antwerpen (U.I.A.), B-2610 Antwerp, Belgium. Ola Hössjer is Lecturer, Department of Mathematical Statistics, Lund Institute of Technology, S-22100 Lund, Sweden. His work was supported by the Swedish Natural Science Research Council under Contract F-DP 6689-300. The authors wish to thank the referees for helpful remarks.

Note that $s_n(\beta)$ does not depend on α because $r_i - r_j$ does not. We will require the following condition:

(R) The function ρ is even, nondecreasing on the positive numbers, and continuous at 0 with $\rho(0) = 0$. There are only a finite number of points where ρ is not continuous or nondifferentiable. Furthermore, $0 < \rho(\infty) < \infty$ and $\rho(c) = \rho(\infty)$ for some $c > 0$.

We will denote $\lim_{n\to\infty} k_{n,p} = k$.

Because $s_n(\beta)$ is independent of the intercept, the latter must be estimated afterwards. We can estimate α by a location estimate based on the numbers $r_i(\hat{\beta}, 0) = y_i - \hat{\beta}'\mathbf{u}_i$; for instance, by using the median or a more efficient 50% breakdown estimator.

We will pay particular attention to the GS-estimator given by

$$\rho(x) = I(|x| \geq 1) \quad \text{and} \quad k_{n,p} = \left(\binom{n}{2} - \binom{h_p}{2} + 1\right)\bigg/\binom{n}{2}$$

with $h_p = [(n + p + 1)/2]$. Then we have that $s_n(\beta) = Q_n(\beta) = \{|r_i - r_j|; i < j\}_{\binom{h_p}{2}:\binom{n}{2}}$. (One should multiply $Q_n(\beta)$ with a certain constant factor to make it consistent as a scale estimator, but that is immaterial to our current goal of estimating β.) Note that $k = \frac{3}{4}$. The scale estimator Q_n was discussed by Rousseeuw and Croux (1993). We will denote the corresponding regression estimator (1) by LQD, because the objective function is approximately the first quartile of the pairwise differences of the residuals. It is instructive to compare this with the LMS objective function, which is given by

$$\{|r_i|; 1 \leq i \leq n\}_{h_p:n}. \tag{7}$$

We will prove that the LQD regression always exists and has the exact fit property and maximal breakdown point. Using these facts, we will show the existence, exact fit property, and maximal breakdown point for a whole class of generalized S-estimators.

Throughout this section we will assume the following property:

(H) No $\binom{h_p}{2}$ of the differences $(\mathbf{u}_i - \mathbf{u}_j, y_i - y_j)$ lie on the same vertical hyperplane in $\mathbb{R}^p$.

By "vertical hyperplane," we mean a hyperplane containing $(\mathbf{0}, 0)$ and $(\mathbf{0}, 1)$. Note that this condition is stronger than requiring that there are no h_p observations $(\mathbf{x}_i, y_i)$ lying on a vertical hyperplane in $\mathbb{R}^{p+1}$. In Figure 1a we see that no h_p observations $(\mathbf{u}_i, y_i)$ are lying on the same affine hyperplane, but nevertheless condition (H) is not satisfied. If the observations follow a continuous distribution, then (H) has probability 1.

Theorem 1. Under condition (H), there always exists a solution to $\operatorname{argmin}_\beta Q_n(\beta)$.

(All proofs are given in the Appendix.) To establish the breakdown point, we need another regularity condition.

Definition. We say that the differences of the $\mathbf{u}_i$ are in general position if no $\binom{p}{2}$ of the $\mathbf{u}_i - \mathbf{u}_j$ with $i < j$ belong to the same hyperplane in $\mathbb{R}^{p-1}$.
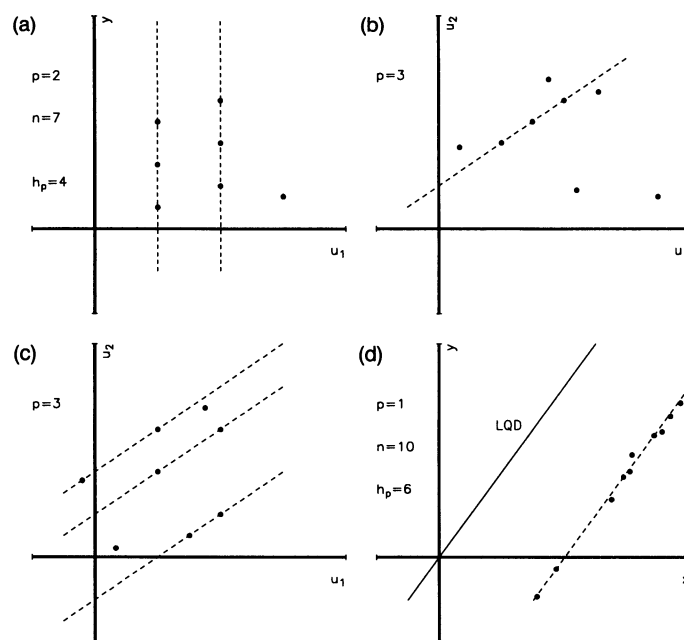

Figure 1. Examples Where (a) Condition (H) is not Satisfied; (b) Neither the u_i nor Their Differences are in General Position; (c) the u_i are in General Position but Their Differences are not; and (d) a Zero-Intercept Model is Inappropriate.

If the differences of the $\mathbf{u}_i$ are in general position, then the $\mathbf{u}_i$ themselves also are in general position. The latter means that no p of the $\mathbf{u}_i$ lie on the same affine hyperplane in $\mathbb{R}^{p-1}$, which is equivalent to saying that no p of the $\mathbf{x}_i$ lie on the same hyperplane in $\mathbb{R}^p$. In Figure 1b we have a situation where neither the $\mathbf{u}_i$ nor their differences are in general position, whereas in Figure 1c we see that it can happen that the $\mathbf{u}_i$ themselves are in general position but their differences are not.

Note that this condition on the differences $\mathbf{u}_i - \mathbf{u}_j$ is more stringent than the condition on the individual $\mathbf{u}_i$, but not much more. If the $\mathbf{u}_i$ have a continuous distribution, then both conditions hold almost surely. From a semantic viewpoint, the $\mathbf{u}_i$ being in "general" position also precludes such linear relations between the $\mathbf{u}_i - \mathbf{u}_j$. (Actually, in computational geometry the phrase "general position" is often interpreted in this stronger sense.) When a data set contains a few exceptions to general position, this does not mean that GS-estimators can no longer be used, but merely that the expression in Theorem 2 will be slightly reduced.

Let us now look at the finite-sample breakdown point (Donoho and Huber 1983). The breakdown point of an estimator T at a sample Z is defined as

$$\varepsilon_n^*(T, Z) = \min\left\{m/n; \sup_{Z'}|T(Z) - T(Z')| = \infty\right\}, \tag{8}$$

where Z' is obtained by replacing any m observations by arbitrary points. (We will use the euclidean norm for $|\cdot|$.)

Theorem 2. If the differences of the $\mathbf{u}_i$ are in general position, then the breakdown point of the LQD estimator is given by $\varepsilon_n^*(\hat{\beta}, Z) = ([(n - p)/2] + 1)/n$, which is the maximal value for any regression equivariant estimator.

From the proof, it follows that we obtain the maximal breakdown point for any objective function $\{|r_i - r_j|; i < j\}_{q:\binom{n}{2}}$, where

$$\binom{[(n + p)/2] - 1}{2} + p - 1 \leq q$$

$$\leq \binom{[(n + p + 1)/2]}{2}. \quad (9)$$

We have chosen to define the LQD based on the largest rank $\binom{h_p}{2}$, where $h_p = [(n + p + 1)/2]$, which is also in accordance with the rank h_p used in the LMS objective (7).

From the general relation between the breakdown point and the exact fit property (Rousseeuw and Leroy 1987, p. 123), the next result immediately follows.

Corollary. If at least $h_p = [(n + p + 1)/2]$ of the observations satisfy $y_i = \beta_0' u_i + \alpha_0$ exactly and the differences of their u_i are in general position, then $\hat{\beta} = \beta_0$ no matter what the other observations are.

Remark 1. It is easy to see that if we take a high-breakdown estimator $\hat{\alpha}$ for the intercept, with $\varepsilon_n^*(\hat{\alpha}, Y) \geq ([n - p]/2] + 1)/n$ for any univariate sample Y, then we also have the maximal breakdown point and the exact fit property for $\hat{\theta} = (\hat{\beta}, \hat{\alpha})$.

Remark 2. It is not advisable to use GS-estimators for fitting a zero-intercept model to data that were actually generated *with* an intercept, because they estimate the slope of a point cloud regardless of the position of the origin. For instance, let us look at Figure 1d. Applying LQD yields an acceptable slope estimate, indicated by the dotted line, but because of the zero-intercept model the actual LQD regression (solid line) does not fit the data points. Note, however, that a study of the LQD-based residuals, all having the same sign and size, does reveal immediately that a zero-intercept model is inappropriate for these data (i.e., model misspecification).

Now we return to GS-estimators of the type (4) with general ρ-function.

Theorem 3. The existence, the maximal breakdown point, and the exact fit property hold for GS-estimators under the same conditions as for the LQD when

$$\frac{k_{n,p}}{\rho(c)} = \left(\binom{n}{2} - \binom{h_p}{2} + 1\right) / \binom{n}{2},$$

which implies that $k/\rho(c) = \frac{3}{4}$.

The most popular choice for a ρ-function is the biweight, $\rho(x) = \min(3x^2/c^2 - 3x^4/c^4 + x^6/c^6, 1)$. If we choose $c = .9958$ (and thus $k = .75$), we obtain a 50% breakdown regression estimator, which we will call the biweight GS-estimator. This estimator is consistent (Hössjer, Croux, and Rousseeuw 1994).

The usual S-estimators, which can be defined by the objective function

$$s_n = \sup\left\{s > 0; \frac{1}{n}\sum_{i=1}^n \rho\left(\frac{r_i}{s}\right) \geq k_{n,p}\right\}, \quad (10)$$

have maximal breakdown point if $k_{n,p}/\rho(c) = (n - h_p + 1)/n$. Therefore, we obtain the maximal breakdown point if $k/\rho(c) = \frac{1}{2}$, where $k = \lim_{n\to\infty} k_{n,p}$. The S-estimator with biweight ρ-function and 50% breakdown point (hence $c = 1.547$) will herein be called the biweight S-estimator.

3. MAXBIAS CURVES

In this section we follow the approach of Martin, Yohai, and Zamar (1989). Because the maxbias curve is an asymptotic notion, we first must determine the functional corresponding to a GS-estimator. We will add to condition (R) that $\rho(\infty) = 1$ (this is only a normalization) and drop the condition that "ρ is constant for large x." For any distribution F, we define the scale functional

$$s(F) = \sup\left\{s > 0; E_F\rho\left(\frac{r_1 - r_2}{s}\right) \geq k\right\}, \quad (11)$$

where r_1 and r_2 are iid according to F and $0 < k < 1$. We denote by $s(\beta, K)$, where K is the joint distribution of (u, y), the same scale functional evaluated at the distribution of the residuals $r(\beta) = y - \beta'u - \alpha$, where we note that this scale does not depend on α. The corresponding GS-estimator then has the functional version T given by $s(T(K), K) = \inf_\beta s(\beta, K)$; hence T is regression and affine equivariant.

Suppose that our model distribution K_0 of (u, y) is elliptical about the origin. We may assume without loss of generality that $T(K_0) = 0$ due to regression equivariance. Denote the distribution of the u_i by G_0 and that of the errors by F_0. Consider the contamination neighborhood $V_\varepsilon = \{K; K = (1 - \varepsilon)K_0 + \varepsilon K^*\}$, where K^* can be any distribution. Then the maxbias curve is given by

$$B_\varepsilon(T) = \sup\{|T(K)|; K \in V_\varepsilon\}. \quad (12)$$

The asymptotic breakdown point may then be defined as $\varepsilon^* = \inf\{\varepsilon; B_\varepsilon(T) = \infty\}$. Suppose that the following conditions hold

- (G) G_0 is spherical, $P_{G_0}(u'\beta) = 0$ for all $\beta \neq 0$ in $\mathbb{R}^p$, and for all β the distribution of $\beta'u$ is unimodal
- (F) F_0 has a unimodal, continuous, and symmetric density.

Define the functions

$$g(s, \beta) = E_{K_0 \times K_0}\rho\left(\frac{y_1 - y_2 - \beta'(u_1 - u_2)}{s}\right) \quad \text{and}$$

$$\tilde{g}(s, \beta) = E_{K_0}\rho\left(\frac{y - \beta'u}{s}\right). \quad (13)$$

Because G_0 is spherical, g and $\tilde{g}$ depend only on $|\beta|$ and s. It holds that g and $\tilde{g}$ are continuous, strictly increasing in $|\beta|$ and strictly decreasing in s (for $s > 0$). This can be proven as in lemma 3.1 of Martin et al. (1989) by using the fact that the distributions of $u_1 - u_2$ and $y_1 - y_2$ also satisfy (G) and (F). Therefore, we may define $g_1^{-1}(\cdot, |\beta|)$ as the inverse of g with respect to s, and $g_2^{-1}(\varepsilon, s)$ as the solution $|\beta|$ of $\tilde{h}(\varepsilon, s, |\beta|) = k$, where $k = \lim_{n\to\infty} k_{n,p}$ was defined following (6) and

$$\tilde{h}(\varepsilon, s, |\beta|) = (1 - \varepsilon)^2 g(s, |\beta|) + 2\varepsilon(1 - \varepsilon)\tilde{g}(s, |\beta|)$$

$$(0 < \varepsilon < 1). \quad (14)$$

By means of two lemmas, given in the Appendix, we now obtain the maxbias curve.

Theorem 4. Under conditions (G) and (F), we have

$$B_\varepsilon(T) = g_2^{-1}\left(\varepsilon, g_1^{-1}\left(\frac{k - 2\varepsilon + \varepsilon^2}{(1 - \varepsilon)^2}, 0\right)\right)$$

$$\text{for} \quad \varepsilon \leq \min(\sqrt{1 - k}, 1 - \sqrt{1 - k})$$

$$= \infty \quad \text{elsewhere.}$$

In the case where $(\mathbf{u}, y)$ is multivariate normal $N(\mathbf{0}, \mathbf{I}_p)$, we obtain

$$\tilde{g}(s, \gamma) = h\left(\frac{\sqrt{1 + \gamma^2}}{s}\right) \quad \text{and} \quad g(s, \gamma) = h\left(\frac{\sqrt{2(1 + \gamma^2)}}{s}\right),$$

where $h(\lambda) = E\rho(\lambda u)$ for $u \sim N(0, 1)$. Note that $h(\lambda)$ is increasing and continuous for $\lambda > 0$, so that h^{-1} is well defined. We can compute $B_\varepsilon(T)$ for any $\varepsilon < \varepsilon^*$ by first computing

$$s_1 = \sqrt{2}/h^{-1}\left(\frac{k - 2\varepsilon + \varepsilon^2}{(1 - \varepsilon)^2}\right).$$

Then $B_\varepsilon(T)$ is given by the solution of the following equation in γ:

$$(1 - \varepsilon)^2 h\left(\frac{\sqrt{2(1 + \gamma^2)}}{s_1}\right) + 2\varepsilon(1 - \varepsilon)h\left(\frac{\sqrt{1 + \gamma^2}}{s_1}\right) = k.$$

In the special case where ρ is a step function $\rho(x) = I(|x| > c)$, we find that $h(\lambda) = 2(1 - \Phi(c/\lambda))$ and thus $h^{-1}(t) = c/\Phi^{-1}(1 - t/2)$.

Remark 1. Theorem 4 implies that the breakdown point becomes

$$\varepsilon^* = \min(\sqrt{1 - k}, 1 - \sqrt{1 - k}). \quad (15)$$

So for the step function $\rho(x) = I(|x| > c)$ with $P_{F_0}(y_1 - y_2 \leq c) = 1 - k/2$, yielding the $(1 - k)$th quantile of the pairwise differences of residuals, we obtain (15). This gives the maximal breakdown point $\varepsilon^* = 50\%$ when $k = \frac{3}{4}$, which corresponds to the LQD.

Remark 2. The derivative of the maxbias curve at 0 is infinite; hence the gross-error sensitivity of GS-estimators is infinite. This is a property of all regression estimators with dimension-free maxbias curve (He and Simpson 1993). But as Yohai and Zamar (1992) have proposed, one could define a modified gross-error sensitivity as $\gamma^{**} = \lim_{\varepsilon \downarrow 0} B_\varepsilon(T)/\sqrt{\varepsilon}$. In the case of a multivariate normal distribution, we obtain $\gamma^{**} = (2\sqrt{2}(1 - h(1))/h'(\sqrt{2}))^{1/2}$. For step functions $\rho(x) = I(|x| > c)$, we obtain $\gamma^{**} = (2\sqrt{2}(2\Phi(c) - 1)/(c\phi(c/\sqrt{2})))^{1/2}$, where $c = \sqrt{2}\Phi^{-1}(1 - k/2)$. In particular, for $k = .75$, we obtain $\gamma^{**}(\text{LQD}) = 2.399$. One can compare this with $\gamma^{**}(\text{LMS}) = 2.160$. For the biweight GS, we obtain $\gamma^{**} = 2.412$; for the biweight S-estimator $\gamma^{**} = 2.267$.
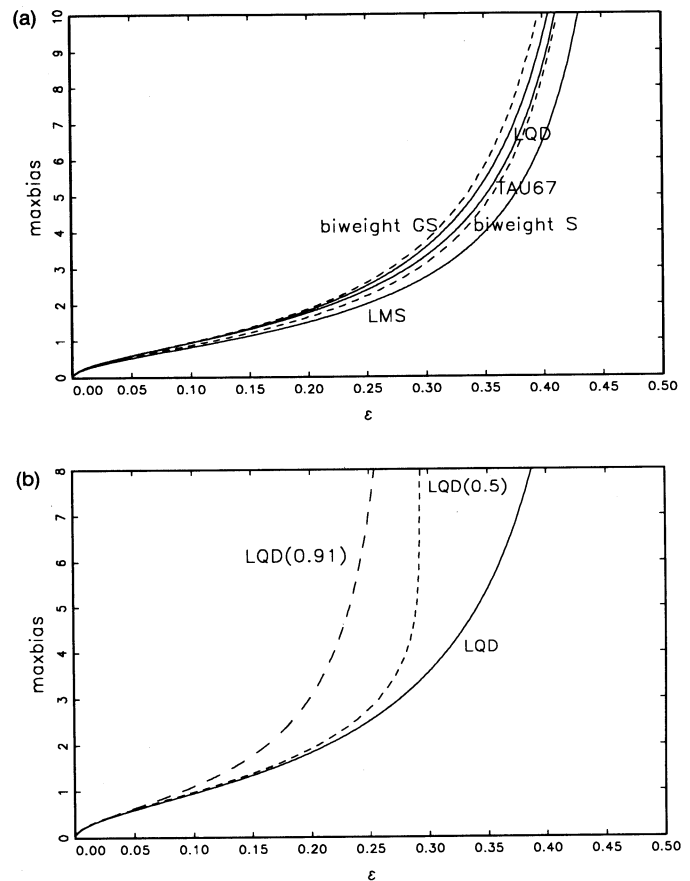

Figure 2. Maxbias Curves of (a) the LMS, Biweight S-, TAU67, LQD, and Biweight GS-Estimators (in increasing order); and (b) the LQD(0.5) and LQD(0.91) Estimators.

Figure 2a plots the maxbias curve of the LQD estimator, together with those of the LMS and the biweight S- and GS-estimators. We see that the bias of the LQD estimator is only slightly larger than that of the biweight S, whereas we will see that its efficiency is much better (67% instead of 28%). Switching from LQD to the biweight GS, which has a smooth ρ-function, again increases the bias only a little (but the efficiency gain will also be small). In Figure 2a we also see the maxbias curve of the TAU67-estimator (Yohai and Zamar 1988). The latter estimator is based on two biweight ρ-functions, the first with $c = 1.547$ to obtain a 50% breakdown point, and the second with $c = 3.26$ to obtain a Gaussian efficiency of 67.1%. (We use the TAU67 estimator here for comparison with the LQD, which has the same efficiency.) We see that the maxbias curve of TAU67 is very close to that of the LQD. Furthermore, $\gamma^{**}(\text{TAU67}) = 2.442$. Thus for small ε the LQD behaves slightly better. (Note that the maxbias curve of a τ-estimator with 95% efficiency is somewhat higher.)

In view of Remark 1, we note that there are often two different step functions ρ that yield GS-estimators with the same breakdown point. For example, $k_1 = .5$ (corresponding with the median interpoint distance) and $k_2 = 1 - (1 - \sqrt{.5})^2 \approx .91$ both yield a 29% breakdown estimator. It is interesting to see that their maxbias curves behave rather differently (see Figure 2b); the smaller value of k (which

corresponds with the higher quantile) is preferable. Note that the maxbias curve of the LQD(.5)-estimator is very close to the LQD for up to about 25% of contamination. Afterward, the LQD(.5) bias increases rapidly, whereas the LQD bias increases more slowly.

A referee asked to compute the breakdown rate (BR), as defined by Mazzi (1991) and Zamar (1992). For a 50% breakdown estimator T, the BR is given by

$$\text{BR}(T) = \lim_{\epsilon \uparrow .5} \frac{B_\epsilon(T)}{B_\epsilon(\text{LMS})}. \qquad (16)$$

Following the proofs and computations of Mazzi (1991), we obtain that BR(LQD) = $.5 + \sqrt{2}$ in the Gaussian case. The breakdown rate of the biweight S- and GS-estimators, and also of TAU67, equals infinity. Therefore, the maxbias of the LQD is lower than that of the biweight S and TAU67 in a neighborhood of the breakdown point.

4. INFLUENCE FUNCTION AND EFFICIENCY

The influence function (see Hampel et al. 1986) describes the (standardized) effect of a single outlier on the estimator. It is an asymptotic notion, based on the same vector-valued functional $T(K)$ as in the previous section. Let K_0 be a fixed distribution representing the central model, and let $K_\epsilon = (1 - \epsilon)K_0 + \epsilon\Delta_{\mathbf{u},y}$, where $\Delta_{\mathbf{u},y}$ is the distribution that puts all its mass at the point $(\mathbf{u}, y)$. Then the influence function is defined as

$$\text{IF}(\mathbf{u}, y) = \lim_{\epsilon \downarrow 0} \frac{T(K_\epsilon) - T(K_0)}{\epsilon}. \qquad (17)$$

Another definition is given by a von Mises expansion: If there exists a function IF: $\mathbb{R}^{p-1} \times \mathbb{R} \rightarrow \mathbb{R}^{p-1}$ (which depends on the estimator and underlying distribution K_0) such that

$$\sqrt{n}\left(\hat{\beta}_n - \beta_0 - \frac{1}{n}\sum_{i=1}^{n} \text{IF}(\mathbf{u}_i, y_i)\right) = o_P(1), \qquad (18)$$

then we call IF the influence function. Under regularity conditions, both definitions coincide. The latter definition has the advantage in that it readily implies that if $E[\text{IF}(\mathbf{u}, y)] = \mathbf{0}$ and $E|\text{IF}(\mathbf{u}, y)|^2 < \infty$, the estimator $\hat{\beta}_n$ is consistent and asymptotically normal with asymptotic covariance matrix

$$V = E[\text{IF}(\mathbf{u}, y)\text{IF}(\mathbf{u}, y)^t]. \qquad (19)$$

Hössjer, Croux, and Rousseeuw (1993) proved the asymptotic normality of GS-estimators, with the function IF specified in Theorem 5, under the following conditions:

(G′) The distribution of the $\mathbf{u}_i$ satisfies $E_{G_0}[\mathbf{u}] = 0$ and $E_{G_0}|\mathbf{u}|^3 < \infty$, and $E_{G_0}[\mathbf{u}\mathbf{u}^t]$ is positive definite.

(F′) The error distribution F_0 has a unimodal density f, which is twice differentiable with a bounded second derivative.

We may assume (due to equivariance) that $T(K_0) = \beta_0 = 0$ and that $s(\mathbf{0}, K_0) = 1$, in which case the influence function is given by the following theorem.

Theorem 5. If the model distribution K_0 satisfies conditions (G′) and (F′), and if ρ satisfies (R), then the influence function of the generalized S-functional is given by

$$\text{IF}(\mathbf{u}, y) = \frac{\bar{\psi}(y)}{E_{F_0}[\bar{\psi}'(y)]} (E_{G_0}[\mathbf{u}\mathbf{u}^t])^{-1}\mathbf{u},$$

where $\bar{\psi}(y) = E_{F_0}[\psi(y - Y)]$ and $\psi = \rho'$. (20)

From Equation (19), we can then compute asymptotic covariances. The efficiency of a GS-estimator at the Gaussian model is thus $e = (\int \bar{\psi}'(y) \, d\Phi(y))^2 / \int \bar{\psi}(y)^2 \, d\Phi(y)$. If f is symmetric, then the influence function of the LQD estimator becomes

$$\text{IF}(\mathbf{u}, y) = -\frac{f(y - c) - f(y + c)}{2\int f'(c + y)f(y) \, dy} (E_{G_0}[\mathbf{u}\mathbf{u}^t])^{-1}\mathbf{u}, \qquad (21)$$

where $P_{F_0}(y_1 - y_2 \le c) = \frac{5}{8}$. In Figure 3a we made a plot of this influence function when $(\mathbf{u}, y) \sim N(\mathbf{0}, \mathbf{I}_2)$. We see that for fixed $\mathbf{u}$, the influence function is redescending (in fact, it goes exponentially to zero for y tending to infinity). On the other hand, the function $\text{IF}(\cdot, y)$ is unbounded in $\mathbf{u}$; hence the overall influence function is not bounded. From (21), it follows that the total influence is small when $|y|$ is large, except when $|\mathbf{u}|$ is exponentially large compared to $|y|$. In the latter situation, the influence function is large at a point $(\mathbf{u}, y)$ lying in a direction with a very small inclination, whose effect on the actual estimated slope is thus negligible.

For the biweight GS-estimator, we obtain a very similar plot (see Fig. 3b). This illustrates that the LQD estimator, with its nonsmooth ρ function, does have an influence function very similar to that obtained with a smooth ρ-function. There is also hardly any gain in efficiency: 68.4% for the biweight GS compared to 67.1% for LQD. This is very dif-
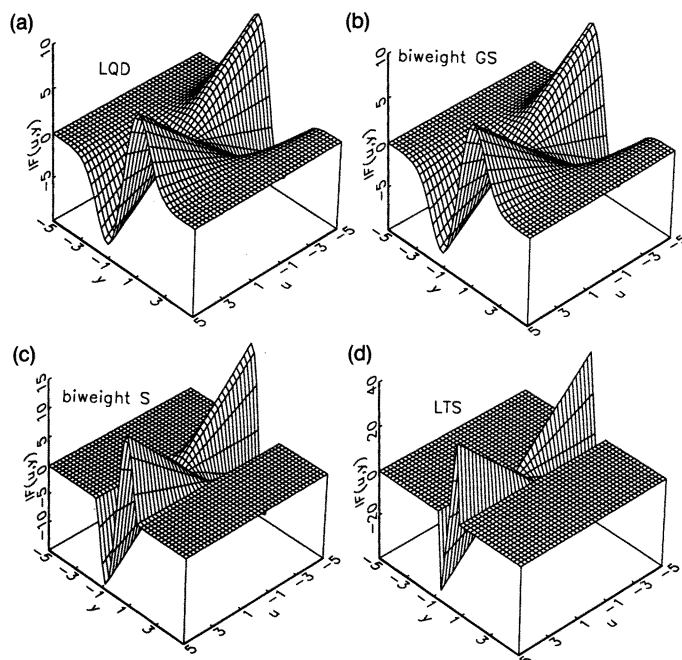

Figure 3. Influence Functions of the (a) LQD, (b) Biweight GS-, (c) Biweight S-, and (d) LTS Estimators.

ferent from the situation for usual S-estimators, where the quantile objective functions (like LMS) yield estimators converging at a lower rate.

In Figure 3c we see the influence function of the usual biweight S-estimator. It has the same shape as the GS-estimator but is steeper. That is the reason why its efficiency is lower (28.7%). In Figure 3d we plotted the influence function of the LTS, which is still steeper and corresponds to an even lower efficiency.

One might argue that the generalized S-estimator should be compared to an S-estimator with $\tilde{\rho}(y) = E\rho(y - Y) - E\rho(Y)$ for its ρ-function. This indeed yields an S-estimator with high efficiency, but with a lower breakdown point and a higher maxbias curve.

Formula (20) also holds for some unbounded ρ-functions. If we take, for example, $\rho(y) = y^2$, then we obtain the LS estimator (the objective function is the standard deviation). This estimator is extremely sensitive to outliers in $\mathbf{u}$ and in y. If we take $\rho(y) = |y|$, then we obtain Gini's average difference as objective function, corresponding to Wilcoxon scores. From Figure 4a we see that this estimator protects against vertical outliers but not at all against bad leverage points. The latter estimator can be seen as a smoothed version of the least absolute deviations estimator (L_1 estimator), which corresponds to a plain S-estimate with the same ρ-function. Its influence function is plotted in Figure 4b. Again, the efficiency increases (from 63.6% to 95.5%) when working on the pairwise differences instead of the individual residuals. Finally, in Figure 4c we see the influence function of the optimal robust 95% efficient Mallows estimator (see Hampel et al. 1986), which is bounded. Analogously, Figure 4d gives the IF of the 95% efficient Schweppe estimator.

Figure 4. Influence Functions of the Jaeckel Estimator Based on Wilcoxon scores (a), the L_1 Estimator (b), and the Optimal Robust 95% Efficient Mallows (c) and Schweppe (d) Estimators.

5. COMPUTATION AND SIMULATION

To compute a GS-estimator, we must minimize the objective function $s(\beta)$, where β is a p-dimensional vector. There has been a substantial amount of research on algorithms for S-estimators, especially the LMS. The same kind of techniques can be used for computing GS-estimators, including the LQD estimator.

The basic scheme for computing S-estimators is the p-subset algorithm (Rousseeuw and Leroy 1987), which minimizes the objective function over all β_J that correspond to fitting a subset J with p observations (out of the n available points). Note that the p-subset version of the LMS is itself a high-breakdown regression estimator (Rousseeuw and Bassett 1991), which is also true for the LQD. Therefore, we can use

$$\hat{\beta}^* = \underset{\beta_J}{\operatorname{argmin}} \, Q_n(y_i - \beta_J' \mathbf{u}_i), \qquad (22)$$

where β_J is determined by the p-subset J. If we use the efficient algorithm of Croux and Rousseeuw (1992) to compute Q_n, then this objective function merely needs $O(n \log n)$ operations, yielding an overall computation time of $O(n^{p+1} \log n)$ if all p-subsets are considered. In comparison, the exhaustive p-subset algorithm for LMS needs $O(n^{p+1})$ time, and also needs $O(n^{p+1} \log n)$ if the intercept is adjusted in every step. Therefore, the LQD needs no more compu-
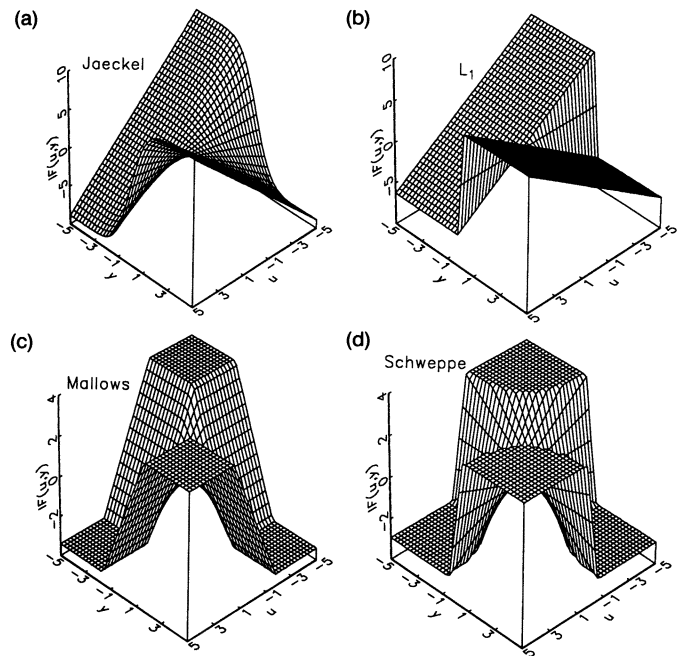
tation time than the LMS, while achieving a much better statistical efficiency.

Note that the p-subset algorithm can be modified to run much faster. (This holds for all estimators of this type, including LMS, LQD, S- and GS-estimators.) The idea is not to consider all $\binom{n}{p} = O(n^p)$ possible p-subsets, but rather to use only $O(n)$ such subsets according to a particular design (Rousseeuw 1993), which ensures that the regression estimator still has the deterministic 50% breakdown point. The resulting LQD algorithm needs only $O(n^2 \log n)$ operations.

Computing the objective function of the biweight GS takes $O(n^2)$ operations (using a fixed number of iterations to solve equation (5)), which is more time-consuming than the LQD. One can reduce the actual computation time, although it remains $O(n^2)$, in the following way. When considering m trial values β_J, we do not need to compute $s(\beta_J)$ each time. Indeed, suppose that $\tilde{s}$ is the currently best scale. Generalizing an observation of Yohai and Zamar (1991), we then have

$$s(\beta_J) \le \tilde{s} \Leftrightarrow \sum_{i<j} \rho\left(\frac{r_i - r_j}{\tilde{s}}\right) < k_{n,p}\binom{n}{2}. \qquad (23)$$

Therefore, we have to compute a new scale estimate only when (23) holds. This happens $O(\log m)$ times. At each new best estimate $\tilde{\beta}$, it is possible to carry out some local improvement, as was done by Ruppert (1992) for S-estimators. The smoothness of our objective function indicates that Newton steps can be useful. For this, we compute

$$\Delta(\beta) = \frac{s(\beta)}{d} \mathbf{\Lambda}^{-1} \sum_{i<j} \psi\left(\frac{r_i - r_j}{s(\beta)}\right)(\mathbf{u}_i - \mathbf{u}_j), \qquad (24)$$

where $\Lambda_{k,l} = \sum_{i<j} (u_{i,l} - u_{j,l})(u_{i,k} - u_{j,k})$ and $d = E\psi'(y_1 - y_2) = E_\Phi\psi'(y/\sqrt{2})$. We search for the smallest value of k

Table 1. Nitrogen Data Set With Robust Distances of (x_{i1}, x_{i2}, x_{i3}) Based on the MVE, as Well as Standardized Residuals $r_i/\hat{\sigma}$ From the LQD Regression

i	x_{i1}	x_{i2}	x_{i3}	y_i	RD_i	$r_i/\hat{\sigma}$
1	.00	61.45	.00	12.18	.72	.46
2	.00	58.11	.01	6.57	.63	-.50
3	.00	65.35	.01	6.99	.84	-.29
4	1.00	47.94	.22	10.69	.46	-.07
5	1.00	57.85	.13	13.75	.63	.60
6	1.00	35.23	.28	10.84	.61	-.29
7	4.00	44.12	.40	15.94	.41	.54
8	4.00	33.19	.39	9.41	.52	-.71
9	4.00	24.18	.40	17.81	.77	.46
10	19.00	25.03	2.40	17.46	.84	-.25
11	19.00	30.61	3.43	23.78	.46	.92
12	19.00	23.28	3.67	18.84	.84	.00
13	49.00	2.76	29.67	57.08	23.00	5.43
14	49.00	1.87	26.75	48.11	19.88	3.84
15	49.00	1.04	23.59	23.26	16.52	-.29
16	80.00	.87	26.08	35.37	13.38	.17
17	80.00	.44	31.00	28.82	18.64	-.64
18	80.00	.20	23.92	30.45	11.08	-.73
19	111.00	.42	18.99	44.63	.84	-.29
20	111.00	.63	23.73	51.75	5.14	1.08
21	111.00	.38	22.02	51.21	3.35	.90

(≤ 10) for which $s(\beta + 2^{-k}\Delta(\beta)) < \tilde{s}$, if any exists. An additional trick is to use

$$(r_i - r_j)(\beta + 2^{-(k+1)}\Delta(\beta))$$

$$= \frac{1}{2}(r_i - r_j)(\beta) + \frac{1}{2}(r_i - r_j)(\beta + 2^{-k}\Delta(\beta)) \quad (25)$$

to speed up the computation. Some experiments with this algorithm show that the objective function will be computed only a few times. The number m is obtained by a trade-off between robustness and computation speed. When computation time permits, carrying out the Newton steps at each β_J is even more accurate.

Remark. Stromberg (1993) has given an exact $O(n^{p+2})$ algorithm for the LMS, which can be generalized to the LQD estimator because

$$\min_{\beta} |r_i - r_j|_{\binom{h_p}{2}} = \min_{\beta} \min_{J \in \mathcal{C}_h} \max_{(i,j) \in J} |r_i - r_j|$$

$$= \min_{J \in \mathcal{C}_h} \min_{\beta} \max_{(i,j) \in J} |r_i - r_j|,$$

where $\mathcal{C}_h = \{J \subset \{1, \ldots, n\}^2; \text{ for all } (i, j) \in J : i < j \text{ and } \#J = \binom{h_p}{2}\}$. Thus we must compute Chebyshev fits β_c on $\{(\mathbf{u}_i - \mathbf{u}_j, y_i - y_j); (i, j) \in J\}$ for all possible J. Only at these values β_c must we compute the objective function, which needs $O(n \log n)$ operations. Now using a theorem of Cheney (1966, p. 36), it is sufficient to look at Chebyshev fits on collections of p observations $(\mathbf{u}_i - \mathbf{u}_j, y_i - y_j)$ with $i < j$. Because the computation time for such a fit depends only on p, we obtain a total time of $O(n^{2p+1}\log n)$, which of course is practical only for small values of n and p.

The data in Table 1 were obtained from T. Vos of the EPFL in Lausanne, Switzerland. The experiment went as follows. Labeled nitrogen (nitrogen-15) was administered to barley plants in the form of fertilizer (NH_4NO_3), to study the nitrogen cycle. The nitrogen is taken up by the plants and converted, after a certain time, to organic material in the soil. The purpose of the study was to explain the organic nitrogen by means of other variables. The variables included in the study are time (in days) after addition of nitrogen (x_1), nitrogen content in mineral form in the soil (x_2), nitrogen content in the plants (x_3), and nitrogen content in organic form in the soil (y).

Following Rousseeuw and van Zomeren (1990), we made a diagnostic plot (Fig. 5a) of the standardized robust residuals $r_i/Q(r_1, \ldots, r_n)$ obtained by the LQD method versus robust distances RD_i obtained with the Minimum Volume Ellipsoid (MVE) estimator. In this plot we can identify six good leverage points and two bad leverage points. The latter (cases 13 and 14) stand out considerably. (Note that also a designed experiment can yield leverage points!) A diagnostic plot of LS residuals versus Mahalanobis distances (see Fig. 5b) does not reveal outliers or leverage points. It would be possible to apply LS regression to this data without cases 13 and 14 (while keeping the good leverage points, because they augment the efficiency).

We also performed a small simulation study based on 1,000 samples $\{(u_i, y_i); i = 1, \ldots, n\}$ from a bivariate Gaussian distribution with unit covariance matrix, for various sample sizes n. For each sample, we computed the LMS,
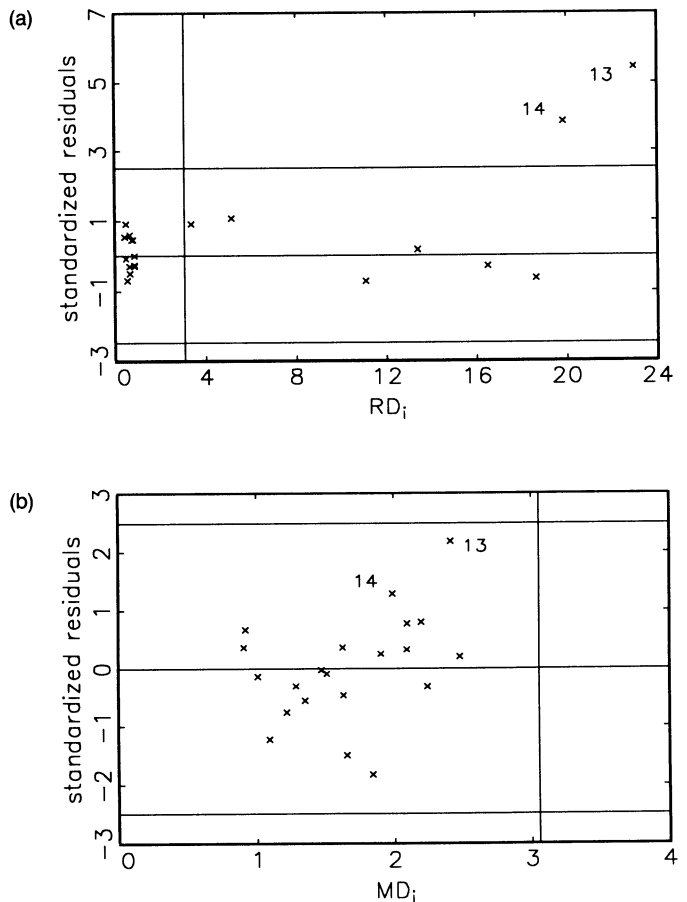
(a)

(b)

Figure 5. Diagnostic Plots of Nitrogen Data. (a) Standardized robust residuals obtained by the LQD method versus robust distances RD_i based on the MVE; (b) standardized LS residuals versus Mahalanobis distances.

LTS, LQD, S-, GS-, and TAU67 estimators by means of the exhaustive p-subset algorithm. Table 2 lists the resulting finite-sample efficiencies of these estimators, where those of the scale estimators were normalized as was done by Rousseeuw and Croux (1993).

For the slope, we note that LQD outperforms both LMS and LTS, the gain being larger for increasing n. The finite-sample efficiencies of LMS, LTS, and LQD all converge quite slowly to their asymptotic limits. (Also note that the LMS is more efficient than the LTS for a large range of sample sizes!) The finite-sample efficiencies of the biweight S are quite stable but are below those of LQD. The GS- and TAU67 estimators have the best performance overall.

For the corresponding estimators of the error scale, we see that the efficiencies of LMS and LTS are rather stable, whereas those of the others converge more slowly. Also here, the LQD, GS-, and TAU67 estimators outperform the plain S-estimator, both asymptotically and for finite samples.

6. OUTLOOK

Similarly to GS-estimators, we may construct other classes of high-breakdown estimators. For example, we can define a generalized R-estimator (or GR-estimator) as $\hat{\beta}_n = \arg\min_\beta D_n(r_1, \ldots, r_n)$, where

$$D_n(r_1, \ldots, r_n) = \sum_{i<j} a(R^+(r_i - r_j))|r_i - r_j|. \quad (26)$$

Here $R^+(r_i - r_j)$ stands for the rank of $r_i - r_j$ among the $\binom{n}{2}$ differences $\{r_i - r_j; i < j\}$. We assume that the scores are generated by a function $h^+ : [0, 1] \to \mathbb{R}_+$ using

$$a(i) = \int_{(i-1)/\binom{n}{2}}^{i/\binom{n}{2}} h^+(t) \, dt.$$

(When $h^+ = 1$, we obtain Wilcoxon scores.) Because the objective (26) is location invariant, we can estimate the intercept afterward.

If $h^+(u) = 0$ for all $u > \frac{1}{4}$ and $h^+(1/4) > 0$ then we obtain a 50% breakdown regression estimator. For instance, if $h^+ = \delta_{1/4}$, we obtain the LQD. When $h^+(u) = I(|u| < 1/4)$, we obtain the estimator

$$\hat{\beta}_n = \arg\min_\beta \sum_{k=1}^{\binom{h_p}{2}} \{|r_i - r_j|; i < j\}_{k:\binom{n}{2}}. \quad (27)$$

Note that this estimator cannot be written as a GS-estimator. An advantage of GR-estimators is that their objective func-

tion (26) is explicit, so one need not solve an equation. But in most cases (26) requires $O(n^2)$ computation time. We think that the maximal efficiency of a GR-estimator would not be much higher than that of the LQD (in fact, the efficiency of (27) is 66.04%), and that the LQD can be seen as a prototype of this class of estimators.

If we put $h(t) = h^+(2t - 1)$ for $t \in [\frac{1}{2}, 1]$ and $h(t) = -h^+(1 - 2t)$ for $t \in [0, \frac{1}{2}]$, then the influence function at the model distribution F will be given by

$$\text{IF}(\mathbf{u}, y) = \frac{E_F h(\tilde{F}(y - Y))}{B(h, F)} (E[\mathbf{u}\mathbf{u}^t])^{-1}\mathbf{u}, \quad (28)$$

where $\tilde{F}$ is the distribution of $y_1 - y_2$ when the y_i are iid according to F, and $B(h, F) = -\int h(\tilde{F}(y))\tilde{F}''(y)dy$. If we further denote $A(h, F) = \int (E_F h(\tilde{F}(y - Y)))^2 \, dF(y)$, then we obtain the asymptotic normality $n^{1/2}(\hat{\beta}_n - \beta) \to N(\mathbf{0}, (E[\mathbf{u}\mathbf{u}^t])^{-1}A(h, F)/B^2(h, F))$.

Instead of working with GS-estimators based on a kernel $\xi(r_i, r_j) = |r_i - r_j|$ of order two, one could also use higher-order kernels. If we use a generalized M-estimator (Serfling 1984) as objective function, then $s(\beta)$ is defined as the solution of the equation

$$\binom{n}{l}^{-1} \sum_{i_1 < \cdots < i_l} \rho\left(\frac{\xi(r_{i_1}, \ldots, r_{i_l})}{s(\beta)}\right) = k_{n,p}.$$

We want ξ to be scale equivariant and location invariant. If we take $\xi(r_{i_1}, \ldots, r_{i_l}) = \text{sdv}(r_{i_1}, \ldots, r_{i_l})$, where sdv stands for the standard deviation, then we obtain a 50% breakdown estimator if $k/\rho(\infty) = 1 - 2^{-l}$. A prototype of this class is

$$\hat{\beta}_n = \arg\min_\beta \{\text{sdv}(r_{i_1}, \ldots, r_{i_l}); i_1 < \cdots < i_l\}_{\binom{h_p}{l}:\binom{n}{l}}.$$

These higher-order estimators will have a higher efficiency but also a higher sensitivity to gross errors. We do not recommend these estimators in practice, because their computation time becomes too great.

Looking in a different direction, we note that GS-estimators can also be extended to high-breakdown estimation of scatter matrices. Consider a data set $x_1, \ldots, x_n$ of points in $\mathbb{R}^p$. Then a GS-estimator of its scatter can be defined as a symmetric positive definite matrix $\mathbf{C}$ that minimizes $\det(\mathbf{C})$ subject to

$$\binom{n}{2}^{-1} \sum_{i<j} \rho(\|\mathbf{x}_i - \mathbf{x}_j\|_\mathbf{C}) \geq \tilde{k}_{n,p}, \quad (29)$$

Table 2. Finite-Sample Efficiencies of the LMS, LTS, LQD, Biweight S, Biweight GS, and TAU67 Estimators

	Slope						Scale					
n	LMS	LTS	LQD	S	GS	τ	LMS	LTS	LQD	S	GS	τ
10	20.8	23.1	30.2	28.1	35.2	37.2	36.1	35.0	45.3	42.5	50.0	59.2
20	19.9	20.0	30.7	26.9	36.5	50.8	36.6	33.0	54.2	45.7	57.9	72.2
40	16.9	14.9	36.0	25.5	43.5	54.3	38.5	34.0	65.2	49.2	69.4	73.8
60	16.0	13.9	36.8	28.1	47.0	60.5	37.2	32.5	68.5	48.4	71.3	83.5
80	15.3	12.8	36.9	28.3	52.2	63.2	40.2	33.1	75.2	51.7	77.8	83.5
100	13.4	13.4	38.6	26.8	52.1	63.9	38.2	32.1	72.1	49.4	74.1	80.9
200	12.8	11.6	45.3	28.5	58.5	66.4	38.7	31.9	77.0	50.4	78.3	81.0
∞	0	7.1	67.1	28.7	68.4	67.1	36.7	30.7	82.3	53.9	82.9	82.7

where $\|\mathbf{x}_i - \mathbf{x}_j\|_{\mathbf{C}}$ stands for $((\mathbf{x}_i - \mathbf{x}_j)'\mathbf{C}^{-1}(\mathbf{x}_i - \mathbf{x}_j))^{1/2}$. Note that the constraint (29) is location invariant, unlike the usual S-estimators (Rousseeuw and Leroy 1987, p. 263), where it is necessary to estimate a location vector $\mathbf{T}$ simultaneously. A special case of (29) is the constraint

$$\{ \|\mathbf{x}_i - \mathbf{x}_j\|_{\mathbf{C}}; i < j \}_{\binom{n}{2}:\binom{n}{2}} \ge \tilde{k}, \tag{30}$$

yielding an analog to LQD regression. For the computation of the scatter matrices given by (29) and (30), one can adapt existing algorithms for the MVE and S-estimators.

APPENDIX: PROOFS

Proof of Theorem 1

Denote $M = \max_{i<j} |y_i - y_j|$. Due to condition (H), we have

$$\inf_{|\beta|=1} \{ |(\mathbf{u}_i - \mathbf{u}_j)'\beta|; i < j \}_{\binom{n}{2}:\binom{n}{2}} = \delta > 0.$$

For $|\beta| > 2M/\delta$, we obtain $|r_i - r_j| \ge \|(\mathbf{u}_i - \mathbf{u}_j)'\beta| - |y_i - y_j|\| \ge |(\mathbf{u}_i - \mathbf{u}_j)'\beta| - M \ge \delta|\beta| - M > M$ for at least $\binom{n}{2} - \binom{h_p}{2} + 1$ differences $|r_i - r_j|$. Thus for all $|\beta| > 2M/\delta$, it is true that $Q_n(\beta) > M \ge Q_n(\mathbf{0})$. Because $\beta \to Q_n(\beta)$ is continuous, $Q_n(\beta)$ will attain its minimum value inside the compact ball $\mathcal{B}(\mathbf{0}, 2M/\delta)$.

Proof of Theorem 2

Denote $\varepsilon_n^* = \varepsilon_n^*(\text{LQD}, Z)$. For any regression equivariant estimator, we have that $\varepsilon_n^* \le ([(n-p)/2] + 1)/n$ (Rousseeuw and Leroy 1987, p. 125), so it is sufficient to prove the reverse inequality. We can assume without loss of generality that $\mathbf{T}(Z) = 0$. Denote the contaminated sample obtained by replacing $k = [(n-p)/2]$ observations from Z by $\{(\mathbf{u}_i', y_i'); i = 1, \dots, n\}$ and denote the corresponding estimate by β_1. Denote $M = \max_i |y_i|$ and $r_i' := y_i' - \beta'\mathbf{u}_i'$. We will prove that $|\beta_1| < \mathbf{C}$, where $\mathbf{C}$ depends only on the original sample.

Note that $|r_i(0) - r_j(0)| = |y_i - y_j| \le 2M$ for all "good" points $(\mathbf{u}_i, y_i)$. Because $\binom{n-k}{2} = \binom{[(n+p+1)/2]}{2} \ge \binom{h_p}{2}$, we have that $Q_n(\mathbf{0}) \le 2M$. So it is sufficient to prove that for all $|\beta| \ge \mathbf{C}$, it holds that $Q_n(\beta) > 2M$, because then it is clear that $|\beta_1| \le \mathbf{C}$. Define

$$r = \frac{1}{2} \inf\{ \mu > 0; \text{ there is a } (p-2) \text{ dimensional subspace } V \text{ in}$$

$$(y = 0) \text{ such that } V^\mu \text{ contains } \binom{p}{2} \text{ differences } \mathbf{u}_i - \mathbf{u}_j \ (i < j) \},$$

where V^μ consists of the points with distance to V less than or equal to μ. Due to our condition of general position, we have $\tau > 0$. Denote $\rho = \tau/n$ and take $\mathbf{C} = 10M/\rho$. Take $|\beta| \ge \mathbf{C}$ and denote by H the hyperplane in $\mathbb{R}^p$ with equation $y = \beta'\mathbf{u}$. Then $L = H \cap (y = 0)$ has dimension $(p-2)$ in $\mathbb{R}^{p-1}$.

We will partition the good observations into classes, induced by the following equivalence relation:

$$(\mathbf{u}_i, y_i) \sim (\mathbf{u}_j, y_j) \Leftrightarrow \exists k \ (0 \le k \le n-2)$$

different observations $\mathbf{u}_{i_1}, \dots, \mathbf{u}_{i_k}$ such that

$$\mathbf{u}_i - \mathbf{u}_{i_1} \in L^\rho, \mathbf{u}_{i_2} - \mathbf{u}_{i_3} \in L^\rho, \dots, \mathbf{u}_{i_k} - \mathbf{u}_j \in L^\rho.$$

Denote the classes with more than one element by $B_1, \dots, B_m$ and denote the union of the other classes by B_0. One can see that

$$\mathbf{u}_i, \mathbf{u}_j \in B_l \Rightarrow \mathbf{u}_i - \mathbf{u}_j \in L^\tau \quad \text{for} \quad l \ge 1.$$

Due to the definition of τ, we have $\sum_{l=1}^m \binom{\#B_l}{2} < \binom{p}{2}$.

Now we will divide the "bad" points into subclasses. Denote by C_l ($1 \le l \le m$) the collection of bad points $(\mathbf{u}_j', y_j')$ for which there exists an element $(\mathbf{u}_i, y_i)$ in B_l such that $|r_i' - r_j'| \le (\rho|\beta| - 2M)/$

4. For each element $(\mathbf{u}_i, y_i)$ in B_0, we denote the collection of bad points $(\mathbf{u}_j', y_j')$ for which $|r_i' - r_j'| \le (\rho|\beta| - 2M)/4$ by C_{i+m} ($1 \le i \le m' = \#B_0$). Note that some of the C_{i+m} can be empty. The bad points that do not belong to any C_l ($1 \le l \le m + m'$) are put in C_0.

If $\mathbf{u} \notin L^\rho$, then $|\mathbf{u}'\beta| > \rho|\beta|$. Therefore, if two good observations $(\mathbf{u}_i, y_i)$ and $(\mathbf{u}_j, y_j)$ do not belong to the same class B_l ($l \ge 0$) or both belong to B_0 (and thus $\mathbf{u}_i - \mathbf{u}_j \notin L^\rho$), then

$$|r_i' - r_j'| = |(y_i - y_j) - \beta'(\mathbf{u}_i - \mathbf{u}_j)|$$

$$\ge \|y_i - y_j| - |\beta'(\mathbf{u}_i - \mathbf{u}_j)\| > \rho|\beta| - 2M,$$

where we have used that $|\beta| \ge 2M/\rho$ and $|y_i - y_j| \le 2M$. Now we can see that the collections C_l ($0 \le l \le m + m'$) are disjoint (using the triangular inequality).

From the foregoing, it follows that $|r_i' - r_j'| \le (\rho|\beta| - 2M)/4$ for at most

$$\sum_{l=1}^m \binom{\#B_l}{2} + \sum_{l=1}^m (\#B_l)(\#C_l) + \sum_{l=1}^{m'} (\#C_{l+m}) + \binom{\sum_{l=0}^{m+m'} \#C_l}{2}$$

couples $i < j$. Using the fact that each B_l ($l \ge 1$) contains at most $p - 1$ elements, we find that the preceding expression is less than or equal to

$$\binom{p}{2} - 1 + k(p-1) + \binom{k}{2} = \binom{k+p-1}{2} + p - 2.$$

Hence $|r_i' - r_j'| > (\rho|\beta| - 2M)/4 > 2M$ at least

$$\binom{n}{2} - \binom{[(n+p)/2]-1}{2} - p + 2$$

times. Because, if we assume the natural condition $n > p$,

$$\binom{[(n+p)/2]-1}{2} + (p-1) \le \binom{[(n+p)/2]-1}{2}$$

$$+ ([(n+p)/2]-1) = \binom{[(n+p)/2]}{2} \le \binom{h_p}{2},$$

we have that $Q_n(\beta) > 2M$.

Proof of Theorem 3

Looking at the proofs of the preceding propositions for the special case where $\hat{\beta}$ is the LQD, it is sufficient to prove that there exist constants $\gamma > 0$ and $\delta > 0$ such that

$$\gamma Q(r_1, \dots, r_n) \le s(r_1, \dots, r_n) \le \delta Q(r_1, \dots, r_n), \tag{A.1}$$

where $Q(r_1, \dots, r_n) = \{ |r_i - r_j|; i < j \}_{\binom{n}{2}:\binom{n}{2}}$.

We can take $\gamma = 1/c$. Indeed, suppose that $Q > cs$. Then there will be $\binom{n}{2} - \binom{h_p}{2} + 1$ differences $|r_i - r_j|$ greater than cs. For an $\varepsilon > 0$ small enough, we will have that

$$\binom{n}{2}^{-1} \sum_{i<j} \rho\left(\frac{r_i - r_j}{s + \varepsilon}\right) \ge \binom{n}{2}^{-1} \left(\binom{n}{2} - \binom{h_p}{2} + 1\right)\rho(c) = k_{n,p},$$

but then s no longer satisfies (6).

For δ, we can take $\delta = 1/\rho^{-1}(\rho(c)/(\binom{h_p}{2} + 1))$. (We define $\rho^{-1}(u) = \sup\{ t > 0; \rho(t) \le u \}$, and then we have that $\rho(\rho^{-1}(u)) \le u$ for all u.) It holds that $\rho^{-1}(\rho(c)/(\binom{h_p}{2} + 1)) > 0$, because otherwise $\rho(t) > \rho(c)/(\binom{h_p}{2} + 1)$ for all $t > 0$. Due to the continuity of ρ at zero, this would yield $\rho(c) = 0$, which is a contradiction. Now suppose that $s\rho^{-1}(\rho(c)/(\binom{h_p}{2} + 1)) > Q$; then $\binom{h_p}{2}$ differences $|r_i - r_j|/s$ are less than $\rho^{-1}(\rho(c)/(\binom{h_p}{2} + 1))$ and thus, for an ε small enough,

$$\binom{n}{2}^{-1} \sum_{i<j} \rho\left(\frac{r_i - r_j}{s - \varepsilon}\right)$$

$$\le \binom{n}{2}^{-1} \rho(c)\left(\binom{h_p}{2}\right) \bigg/ \left(\binom{h_p}{2} + 1\right) + \binom{n}{2} - \binom{h_p}{2}\right) < k_{n,p},$$

which is contradictory to (6).

To prove Theorem 4, we need the following two lemmas.

Lemma 1. Let $(\mathbf{u}_1, y_1) \sim K_0$ and $(\mathbf{u}_2, y_2) \sim K^*$, where K^* can be any distribution. Then for all $s > 0$ and for all $\beta > 0$,

$$E_{K_0 \times K^*}\left[\rho\left(\frac{y_1 - y_2 - \beta'(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right] \geq E_{K_0}\left[\rho\left(\frac{y_1 - \beta'x_1}{s}\right)\right].$$

$$(A.2)$$

Proof of Lemma 1. Using symmetry and unimodality of y_1 and $\beta'\mathbf{u}_1$ we find that $z = (y_1 - \beta'\mathbf{u}_1)/s$ is unimodal and symmetric. Therefore, $E\rho(z - z^*) \geq E\rho(z)$, where z^* can be any stochastic variable. This proves (A.2).

Lemma 2. Let $(\mathbf{u}, y) \sim K_0$, $\mathbf{u}_n^*$ be uniformly distributed on the line segment $[\lambda_n\beta^*, 2\lambda_n\beta^*]$, and put $y_n = \mathbf{u}_n^t\beta^*$ and $K_n^* \sim (\mathbf{u}_n^*, y_n^*)$. Suppose that $\lambda_n \to \infty$, $\beta_n \to \tilde\beta$ and $|\tilde\beta| < |\beta^*|$. Then

$$\lim_{n \to \infty} E_{K_0 \times K_n^*}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right] = 1 \quad \forall s > 0 \quad (A.3)$$

and

$$\lim_{n \to \infty} E_{K_n^* \times K_n^*}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right] = 1 \quad \forall s > 0. \quad (A.4)$$

Proof of Lemma 2. Take $\varepsilon > 0$. We have that

$$\left|E_{K_0 \times K_n^*}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right] - 1\right|$$

$$= \left|\frac{1}{\lambda_n}\int_{\lambda_n}^{2\lambda_n} E_{K_0}\left[\rho\left(\frac{y - \beta_n^t\mathbf{u} - t(|\beta^*|^2 - \beta_n^t\beta^*)}{s}\right)\right] dt - 1\right|$$

$$= \left|E_{K_0}\left[\rho\left(\frac{y - \beta_n^t\mathbf{u} - t_n(|\beta^*|^2 - \beta_n^t\beta^*)}{s}\right)\right] - 1\right|,$$

where $\lambda_n \leq t_n \leq 2\lambda_n$. (In the last step we used the mean value theorem for integrals of positive continuous functions.) Denote H_n, the distribution of $(y - \beta_n^t\mathbf{u})/s$; because β_n is a bounded sequence, we can find a compact set C for which $H_n\{C\} > \sqrt{1 - \varepsilon}$. Denote further $x_n = t_n(|\beta^*|^2 - \beta_n^t\beta^*)/s$; then $x_n \to \infty$ (using that $\lim_{n \to \infty} |\beta^*|^2 - \beta_n^t\beta^* > 0$). So for n large enough, we have for all $x \in C$ that $\rho(x - x_n) > \sqrt{1 - \varepsilon}$. We can rewrite (A.5) as

$$1 - \int_C \rho(x - x_n) dH_n(x) - \int_{\mathbb{R}^p \setminus C} \rho(x - x_n) dH_n(x)$$

$$\leq 1 - \int_C \rho(x - x_n) dH_n(x) \leq 1 - \inf_{x \in C} \rho(x - x_n)H_n\{C\}$$

$$\leq 1 - (1 - \varepsilon) = \varepsilon.$$

This proves (A.3). We continue with the proof of (A.4):

$$\left|E_{K_n^* \times K_n^*}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right] - 1\right|$$

$$= \left|E_{K_n^* \times K_n^*}\left[\rho\left(\left(\frac{\beta^* - \beta_n}{s}\right)^t(\mathbf{u}_1 - \mathbf{u}_2)\right)\right] - 1\right|$$

$$= \left|\frac{2}{\lambda_n}\int_0^{\lambda_n}\left(1 - \frac{t}{\lambda_n}\right)\rho\left(\left(\frac{\beta^* - \beta_n}{s}\right)^t\beta^* t\right) dt - 1\right|, \quad (A.6)$$

where we worked out the distribution of $\mathbf{u}_1 - \mathbf{u}_2$. Denote $x_n = (|\beta^*|^2 - \beta_n^t\beta^*)/s$. Then $x_n > \delta > 0$ for n large. Choose L such that $\rho(x_n t) \geq \rho(t\delta) > 1 - \varepsilon/4$ for all $t > L$. For n large enough, we have that $\lambda_n > L$ and $2L/\lambda_n < \varepsilon/2$. Then (A.6) equals

$$\left|\frac{2}{\lambda_n}\int_0^L\left(1 - \frac{t}{\lambda_n}\right)(\rho(x_n t) - 1) dt + \frac{2}{\lambda_n}\int_L^{\lambda_n}\left(1 - \frac{t}{\lambda_n}\right)(\rho(x_n t) - 1) dt\right|$$

$$\leq \frac{2L}{\lambda_n} + \frac{2}{\lambda_n}(\lambda_n - L)\sup_{t > L}|\rho(x_n t) - 1| \leq \frac{2L}{\lambda_n}(1 - \varepsilon/4) + \varepsilon/2$$

$$\leq \frac{2L}{\lambda_n} + \varepsilon/2 \leq \varepsilon/2 + \varepsilon/2 = \varepsilon.$$

This proves (A.4).

Proof of Theorem 4

Let $c = g_2^{-1}(\varepsilon, s_1)$, where $s_1 = g_1^{-1}((k - 2\varepsilon + \varepsilon^2)/(1 - \varepsilon)^2, 0)$. Suppose that $\varepsilon < \min(1 - \sqrt{1 - k}, \sqrt{1 - k})$.

We will first prove that $B_\varepsilon(T) \leq c$. Take any distribution K of the form $K = (1 - \varepsilon)K_0 + \varepsilon K^*$ and consider a slope $|\beta| > c$. It is sufficient to prove that

$$s(\beta, K) > s(0, K). \quad (A.7)$$

Because $\tilde h(\varepsilon, s_1, c) = k$, we have that $\tilde h(\varepsilon, s_1, |\beta|) > k$. Using continuity, there must exist an $s_2 > s_1$ such that $\tilde h(\varepsilon, s_2, |\beta|) > k$. Using this last inequality, Lemma 1, and the positivity of ρ, we find

$$E_{K \times K}\left[\rho\left(\frac{y_1 - y_2 - \beta'(\mathbf{u}_1 - \mathbf{u}_2)}{s_2}\right)\right]$$

$$= (1 - \varepsilon)^2 g(s_2, \beta) + 2\varepsilon(1 - \varepsilon)$$

$$\times E_{K_0 \times K^*}\left[\rho\left(\frac{y_1 - y_2 - \beta'(\mathbf{u}_1 - \mathbf{u}_2)}{s_2}\right)\right]$$

$$+ \varepsilon^2 E_{K^* \times K^*}\left[\rho\left(\frac{y_1 - y_2 - \beta'(\mathbf{u}_1 - \mathbf{u}_2)}{s_2}\right)\right]$$

$$\geq (1 - \varepsilon)^2 g(s_2, \beta) + 2\varepsilon(1 - \varepsilon)\tilde g(s_2, \beta) = \tilde h(\varepsilon, s_2, \beta) > k.$$

Therefore,

$$s_2 \leq s(\beta, K). \quad (A.8)$$

Now for any $s > s_1$, we have that

$$E_{K \times K}\left[\rho\left(\frac{y_1 - y_2}{s}\right)\right] \leq (1 - \varepsilon)^2 g(s, 0) + 2\varepsilon(1 - \varepsilon) + \varepsilon^2$$

$$\leq (1 - \varepsilon)^2 g(s_1, 0) + 2\varepsilon(1 - \varepsilon) + \varepsilon^2 = k.$$

We can conclude that $s \geq s(0, K)$ and thus

$$s_1 \geq s(0, K). \quad (A.9)$$

Combining (A.8) and (A.9) yields (A.7).

Now we will prove the other inequality,

$$B_\varepsilon(T) \geq c. \quad (A.10)$$

Take any $0 < c_1 < c$ and $|\beta^*| = c_1$. Let K_n^* be a contaminating distribution corresponding to $(\mathbf{u}_n^*, y_n^*)$, where $y_n^* = \beta^{*t}\mathbf{u}_n^*$ and $\mathbf{u}_n^*$ is uniformly distributed on the line segment $[\lambda_n\beta^*, 2\lambda_n\beta^*]$ and $\lambda_n \to \infty$. (In fact, we want both the location and spread of $\mathbf{u}_n^*$ to go beyond all bounds when n increases.) To prove (A.10), it is enough to show that

$$\sup_n |T(K_n)| \geq c_1, \quad (A.11)$$

where $K_n = (1 - \varepsilon) + \varepsilon K_n^*$. Suppose that (A.11) is not true; then we can construct a subsequence, which we shall still call K_n, for which $\lim_{n \to \infty} \beta_n = \tilde\beta$, where $T(K_n) = \beta_n$ and $|\tilde\beta| < |\beta^*| = c_1$. Now we have

$$E_{K_n \times K_n}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right]$$

$$= (1 - \varepsilon)^2 g(s, |\beta_n|)$$

$$+ 2\varepsilon(1 - \varepsilon)E_{K_0 \times K_n^*}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right]$$

$$+ \varepsilon^2 E_{K_n^* \times K_n^*}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right].$$

Using Lemma 2 and the definition of s_1 yields, for all $s < s_1$, that

$$\lim_{n \to \infty} E_{K_n \times K_n}\left[\rho\left(\frac{y_1 - y_2 - \beta_n^t(\mathbf{u}_1 - \mathbf{u}_2)}{s}\right)\right]$$

$$\geq (1 - \varepsilon)^2 g(s, 0) + 2\varepsilon(1 - \varepsilon) + \varepsilon^2$$

$$> (1 - \varepsilon)^2 g(s_1, 0) + 2\varepsilon(1 - \varepsilon) + \varepsilon^2 = k.$$

Therefore, $\lim_{n \to \infty} s(\beta_n, K_n) \geq s$ and thus

$$\lim_{n \to \infty} s(\beta_n, K_n) \geq s_1. \tag{A.12}$$

On the other hand, $\tilde{h}(\varepsilon, s_1, c_1) < \tilde{h}(\varepsilon, s_1, c) = k$. Due to the continuity of $\tilde{h}$, we can find an $s_2 < s_1$ such that $\tilde{h}(\varepsilon, s_2, c_1) < k$. Using the fact that $y_n^* = \beta^{*\,t}\mathbf{u}_n^*$ exactly, we have that

$$E_{K_n \times K_n}\left[\rho\left(\frac{y_1 - y_2 - \beta^{*\,t}(\mathbf{u}_1 - \mathbf{u}_2)}{s_2}\right)\right]$$

$$= (1 - \varepsilon)^2 g(s_2, c_1) + 2\varepsilon(1 - \varepsilon)\tilde{g}(s_2, c_1) = \tilde{h}(\varepsilon, s_2, c_1) < k.$$

Therefore,

$$s(\beta^*, K_n) \leq s_2. \tag{A.13}$$

Combining (A.12) and (A.13) shows that for n large enough, β_n does not minimize $s(\cdot, K_n)$, which gives a contradiction. Therefore, (A.11) must be true.

To complete the proof, we show that if $\varepsilon \uparrow \min(\sqrt{1-k}, 1 - \sqrt{1-k})$, then

$$g_2^{-1}\left(\varepsilon, g_1^{-1}\left(\frac{k - 2\varepsilon + \varepsilon}{(1 - \varepsilon)^2}, 0\right)\right) \to \infty.$$

This follows from $k - 2\varepsilon + \varepsilon \geq 0 \Leftrightarrow \varepsilon \leq 1 - \sqrt{1-k}$ and $(1 - \varepsilon)^2 + 2\varepsilon(1 - \varepsilon) \leq k \Leftrightarrow \varepsilon \leq \sqrt{1-k}$.

Proof of Theorem 5

Combining equations (11), (15), and (17) of theorem 1 of Hössjer et al. (1994) yields (20) according to definition (18). Here we give a proof using definition (17). The functional T is given by $T(K) = \operatorname{argmin}_\beta s(\beta, K)$, where $s(\beta, K)$ satisfies $E_{K \times K}[\rho((y_1 - y_2 - \beta^t(\mathbf{u}_1 - \mathbf{u}_2))/s(\beta, K)]$ and $(\mathbf{u}_1, y_1)$ and $(\mathbf{u}_2, y_2)$ are two independent variables drawn from K. Because $T(K)$ minimizes $s(\beta, K)$, we have

$$E_{K \times K}\left[\psi\left(\frac{y_1 - y_2 - T(K)^t(\mathbf{u}_1 - \mathbf{u}_2)}{s(T(K), K)}\right)(\mathbf{u}_1 - \mathbf{u}_2)\right] = 0.$$

Also, the contaminated distribution $K_\varepsilon = (1 - \varepsilon)K_0 + \varepsilon\Delta_{\mathbf{u},y}$ must satisfy

$$E_{K_\varepsilon \times K_\varepsilon}\left[\psi\left(\frac{y_1 - y_2 - T(K_\varepsilon)^t(\mathbf{u}_1 - \mathbf{u}_2)}{s(T(K_\varepsilon), K_\varepsilon)}\right)(\mathbf{u}_1 - \mathbf{u}_2)\right] = 0.$$

Working this out yields

$$(1 - \varepsilon)^2 E_{K_0 \times K_0}\left[\psi\left(\frac{y_1 - y_2 - T(K_\varepsilon)^t(\mathbf{u}_1 - \mathbf{u}_2)}{s(T(K_\varepsilon), K_\varepsilon)}\right)(\mathbf{u}_1 - \mathbf{u}_2)\right]$$

$$+ 2\varepsilon(1 - \varepsilon)E_{K_0}\left[\psi\left(\frac{y_1 - y - T(K_\varepsilon)^t(\mathbf{u}_1 - \mathbf{u})}{s(T(K_\varepsilon), K_\varepsilon)}\right)(\mathbf{u}_1 - \mathbf{u})\right] = 0.$$

Differentiating with respect to ε and evaluating in 0 gives

$$E_{K_0 \times K_0}\left[\psi'(y_1 - y_2)\left\{-\sum_{i=1}^{p} \mathrm{IF}_i(\mathbf{u}, y, K_0)(u_{1,l} - u_{2,l})\right.\right.$$

$$\left.\left. - (y_1 - y_2)\frac{\partial s(T(K_\varepsilon), K_\varepsilon)}{\partial \varepsilon}\bigg|_{\varepsilon=0}\right\}(u_{1,k} - u_{2,k})\right]$$

$$+ 2E_{K_0}[\psi(y_1 - y)(u_{1,k} - u_k)] = 0$$

for all $1 \leq k \leq p - 1$. Because y and $\mathbf{u}$ are independent at K_0 and $E_{G_0}(u_{1,k} - u_{2,k}) = 0$, we find that

$$-E_{F_0 \times F_0}[\psi'(y_1 - y_2)]\mathrm{IF}(\mathbf{u}, y)^t E_{G_0 \times G_0}[(\mathbf{u}_1 - \mathbf{u}_2)(u_{1,k} - u_{2,k})]$$

$$+ 2E_{F_0}[\psi(y_1 - y)]E[u_{1,k} - u_k] = 0$$

for all $1 \leq k \leq p - 1$. Because $E(\mathbf{u}_1) = 0$ and $E_{G_0 \times G_0}[(\mathbf{u}_1 - \mathbf{u}_2)(\mathbf{u}_1 - \mathbf{u}_2)^t] = 2E_{G_0}[\mathbf{u}_1\mathbf{u}_1^t]$, we obtain Equation (20).

REFERENCES

Cheney, E. W. (1966), *Introduction to Approximation Theory*, New York: McGraw-Hill.

Croux, C., and Rousseeuw, P. J. (1992), "Time-Efficient Algorithms for Two Highly Robust Estimators of Scale," in *Computational Statistics*, Vol. 1, eds. Y. Dodge and J. Whittaker, Heidelberg: Physika-Verlag, pp. 411–428.

Donoho, D. L., and Huber, P. J. (1983), "The Notion of Breakdown Point," in *A Festschrift for Erich Lehmann*, eds. P. Bickel, K. Doksum, and J. L. Hodges, Jr., Belmont, CA: Wadsworth, pp. 157–184.

Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., and Stahel, W. A. (1986), *Robust Statistics: The Approach Based on Influence Functions*, New York: John Wiley.

He, X., and Simpson, D. G. (1993), "Lower Bounds for Contamination Bias: Globally Minimax Versus Locally Linear Estimation," *The Annals of Statistics*, 21, 314–337.

Hössjer, O. (1992), "On the Optimality of S-Estimators," *Statistics and Probability Letters*, 14, 413–419.

Hössjer, O., Croux, C., and Rousseeuw, P. J. (1994), "Asymptotics of Generalized S-Estimators," submitted to *Journal of Multivariate Analysis*, Vol. 51.

Martin, R. D., Yohai, V. J., and Zamar, R. H. (1989), "Min-Max Bias Robust Regression," *The Annals of Statistics*, 17, 1608–1630.

Mazzi, S. V. (1991), "A New Measure of Quantitative Robustness," master's thesis, University of British Columbia, Dept. of Statistics.

Rousseeuw, P. J. (1984), "Least Median of Squares Regression," *Journal of the American Statistical Association*, 79, 871–880.

——— (1993), "A Resampling Design for Computing High-Breakdown Regression," *Statistics and Probability Letters*, 18, 125–128.

Rousseeuw, P. J., and Bassett, G. W. (1991), "Robustness of the p-Subset Algorithm for Regression With High Breakdown Point," in *Directions in Robust Statistics and Diagnostics*, Part II, eds. W. Stahel and S. Weisberg, New York: Springer-Verlag, pp. 185–194.

Rousseeuw, P. J., and Croux, C. (1993), "Alternatives to the Median Absolute Deviation," *Journal of the American Statistical Association*, 88, 1273–1283.

Rousseeuw, P. J., and Leroy, A. M. (1987), *Robust Regression and Outlier Detection*, New York: John Wiley.

Rousseeuw, P. J., and Yohai, V. J. (1984), "Robust Regression by Means of S-Estimators," in *Robust and Nonlinear Time Series Analysis. Lecture Notes in Statistics 26*, eds. J. Franke, W. Härdle, and R. D. Martin, New York: Springer-Verlag, pp. 256–272.

Rousseeuw, P. J., and van Zomeren, B. C. (1990), "Unmasking Multivariate Outliers and Leverage Points," *Journal of the American Statistical Association*, 85, 633–639.

Ruppert, D. (1992), "Computing S-Estimators for Regression and Multivariate Location/Dispersion," *Journal of Computational and Graphical Statistics*, 1, 253–270.

Serfling, R. J. (1984), "Generalized L, M, and R Statistics," *The Annals of Statistics*, 12, 76–86.

Stromberg, A. J. (1993), "Computing the Exact Least Median of Squares Estimate and Stability Diagnostics in Multiple Linear Regression," *SIAM Journal of Scientific Computing*, 14, 1289–1299.

Yohai, V. J. (1987), "High Breakdown Point and High Efficiency Robust Estimates for Regression," *The Annals of Statistics*, 15, 642–656.

Yohai, V. J., and Zamar, R. H. (1988), "High Breakdown Point Estimates of Regression by Means of the Minimization of an Efficient Scale," *Journal of the American Statistical Association*, 83, 406–413.

——— (1991), Discussion of "Least Median of Squares Estimation in Power Systems" by L. Mili, V. Phaniraj, and P. J. Rousseeuw, *IEEE Transactions on Power Systems*, 6, 520.

——— (1992), "Optimally Bounding a Generalized Gross-Error Sensitivity," unpublished manuscript.

Zamar, R. H. (1992), "A New Measure of Quantitative Robustness," paper presented at the Workshop on Data Analysis and Robustness, Ascona, Switzerland, June 29–July 3, 1992.