Computational Statistics and Data Analysis 83 (2015) 262–274

Contents lists available at ScienceDirect

Computational Statistics and Data Analysis

journal homepage: www.elsevier.com/locate/csda

High finite-sample efficiency and robustness based on distance-constrained maximum likelihood

Ricardo A. Maronna [a,*], Victor J. Yohai [b]

[a] Department of Mathematics, School of Exact Sciences, National University of La Plata, Argentina

[b] Department of Mathematics, School of Exact and Natural Sciences, University of Buenos Aires and CONICET, Argentina

ARTICLE INFO

Article history:
Received 29 November 2013
Received in revised form 7 October 2014
Accepted 8 October 2014
Available online 22 October 2014

Keywords:
Finite-sample efficiency
Robust regression
Robust multivariate location and scatter
Kullback–Leibler divergence

ABSTRACT

Good robust estimators can be tuned to combine a high breakdown point and a specified asymptotic efficiency at a central model. This happens in regression with MM- and τ-estimators among others. However, the finite-sample efficiency of these estimators can be much lower than the asymptotic one. To overcome this drawback, an approach is proposed for parametric models, which is based on a distance between parameters. Given a robust estimator, the proposed one is obtained by maximizing the likelihood under the constraint that the distance is less than a given threshold. For the linear model with normal errors, simulations show that the proposed estimator attains a finite-sample efficiency close to one while improving the robustness of the initial estimator. The same approach also shows good results in the estimation of multivariate location and scatter.

1. Introduction

Since the seminal work of Huber (1964) and Hampel (1971), one of the main concerns of the research in robust statistics has been to derive statistical procedures that are simultaneously highly robust and highly efficient under the assumed model. The efficiency of an estimator is usually measured by the asymptotic efficiency, that is, by the ratio between the asymptotic variances of the maximum likelihood estimator (henceforth MLE) and of the robust estimator. However if the sample size n is not very large, this asymptotic efficiency may be quite different from the finite sample size one, defined as the ratio between the mean squared errors (MSE) of the MLE and of the robust estimator, for samples of size n. However, it is obvious that for practical purposes only the finite sample size efficiency matters.

Consider for example the case of a linear model with normal errors. In this case the MLE of the regression coefficients is the least squares estimator (LSE). It is well known that this estimator is very sensitive to outliers, and in particular its breakdown point is zero. To overcome this problem, several estimators combining high asymptotic breakdown point and high efficiency have been proposed. Yohai (1987) proposed MM-estimators, which have 50% breakdown point and asymptotic efficiency as close to one as desired. Yohai and Zamar (1988) proposed τ-estimates, which combine the same two properties as MM-estimators. Gervini and Yohai (2002) proposed regression estimators which simultaneously have 50% breakdown point and asymptotic efficiency equal to one.

However, as will be seen in Section 2.1, when n is not very large the finite sample efficiency of these estimators may be much smaller than the asymptotic one. On the other hand, a 50% breakdown point does not guarantee that the estimator is

* Correspondence to: Departamento de Matemática, Facultad de Ciencias Exactas, C.C. 172, La Plata 1900, Argentina.
E-mail address: rmaronna@retina.ar (R.A. Maronna).

http://dx.doi.org/10.1016/j.csda.2014.10.015
0167-9473/

highly robust. In fact, this only guarantees that given $\varepsilon < 0.5$ there exists $K(\varepsilon)$ such that if the data are contaminated with a fraction of outliers smaller than ε, the norm of the difference between the estimator and the true value is smaller than $K(\varepsilon)$. However $K(\varepsilon)$ may be very large, which makes the estimator unstable under outlier contamination of size ε.

Bondell and Stefanski (2013) proposed a regression estimator with maximum breakdown point and high finite-sample efficiency. However, as it will be seen in Section 2.1, the price for this efficiency is a serious loss of robustness.

An alternative approach to robust estimation is proposed by Olive and Hawkins (2010, 2011); see also Zhang et al. (2012). Their estimators are consistent and have high breakdown point, but since they are not equivariant, comparisons with them are difficult.

The purpose of this paper is to present estimators which have a high finite sample size efficiency and robustness even for small n. Besides, these estimators are highly robust using a robustness criterion better than the breakdown point, namely, the maximum MSE for a given contamination rate ε.

The procedure to define the proposed estimators is very general and may be applied to any parametric or semiparametric model. However in this paper the details are given only to estimate the regression coefficients in a linear model and the multivariate location and scatter of a random vector.

To define the proposed estimators we need an initial robust estimator, not necessarily with high finite sample efficiency. Then the estimators are defined by maximizing the likelihood function subject to the estimate being sufficiently close to the initial one. Doing so we can expect that the resulting estimator will have the maximum possible finite sample efficiency under the assumed model compatible with proximity to the initial robust estimator. This proximity guarantees the robustness of the new estimator.

The formulation of our proposal is as follows. Let D be a distance or discrepancy measure between densities. As a general notation, given a family of distributions with observation vector $\mathbf{z}$, parameter vector $\boldsymbol{\theta}$ and density $f(\mathbf{z}, \boldsymbol{\theta})$, put $d(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) = D(f(\mathbf{z}, \boldsymbol{\theta}_1), f(\mathbf{z}, \boldsymbol{\theta}_2))$. Let $\mathbf{z}_i$, $i = 1, \ldots, n$ be i.i.d. observations with distribution $f(\mathbf{z}, \boldsymbol{\theta})$, and let $\widehat{\boldsymbol{\theta}}_0$ be an initial robust estimator. Call $L(\mathbf{z}_1, \ldots, \mathbf{z}_n; \boldsymbol{\theta})$ the likelihood function. Then our proposal is to define an estimator $\widehat{\boldsymbol{\theta}}$ as

$$\widehat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} L(\mathbf{z}_1, \ldots, \mathbf{z}_n; \boldsymbol{\theta}) \quad \text{with } d\left(\widehat{\boldsymbol{\theta}}_0, \boldsymbol{\theta}\right) \leq \delta \tag{1}$$

where δ is an adequately chosen constant that may depend on n. We shall call this proposal "distance-constrained maximum likelihood' (DCML for short).

Several dissimilarity measures, such as the Hellinger distance, may be employed for this purpose. We shall employ as D the Kullback–Leibler (KL) divergence, because, as it will be seen, it yields easily manageable results. Therefore the d in (1) will be

$$d_{\mathrm{KL}}(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2) = \int_{-\infty}^{\infty} \log\left(\frac{f(\mathbf{z}, \boldsymbol{\theta}_1)}{f(\mathbf{z}, \boldsymbol{\theta}_2)}\right) f(\mathbf{z}, \boldsymbol{\theta}_1)\, d\mathbf{z}.$$

In Sections 2 and 3 we apply this procedure to the linear model and to the estimation of multivariate location and scatter, respectively. In Section 4 the proposed estimators are applied to two data sets. Finally Section 5 summarizes the results.

2. Regression

Consider the family of distributions with $\mathbf{z} = (\mathbf{x}, y)$, with $\mathbf{x} \in R^p$ and $y \in R$, satisfying the model $y = \mathbf{x}'\boldsymbol{\beta} + \sigma u$, where $u \sim \mathrm{N}(0, 1)$ is independent of $\mathbf{x} \in R^p$. Here $\boldsymbol{\theta} = (\boldsymbol{\beta}, \sigma)$. Let $\widehat{\boldsymbol{\theta}}_0 = \left(\widehat{\boldsymbol{\beta}}_0, \widehat{\sigma}_0\right)$ be an initial robust estimator of regression and scale. We will actually consider σ as a nuisance parameter, and therefore we have

$$d_{\mathrm{KL}}\left(\boldsymbol{\beta}_0, \boldsymbol{\beta}\right) = \frac{1}{\sigma^2}\left(\boldsymbol{\beta} - \boldsymbol{\beta}_0\right)' \mathbf{C}\left(\boldsymbol{\beta} - \boldsymbol{\beta}_0\right) \tag{2}$$

with $\mathbf{C} = \mathrm{E}\mathbf{x}\mathbf{x}'$.

Here we replace σ with its estimator $\widehat{\sigma}_0$. The natural estimator of $\mathbf{C}$ would be $\widehat{\mathbf{C}} = n^{-1}\mathbf{X}'\mathbf{X}$, where $\mathbf{X}$ is the $n \times p$ matrix with rows $\mathbf{x}_i'$. Since it is not robust, we will employ a robust version thereof. Put for $\boldsymbol{\beta} \in R^p$, $r_i(\boldsymbol{\beta}) = y_i - \mathbf{x}'\boldsymbol{\beta}$, the residuals from $\boldsymbol{\beta}$. All "smooth" robust regression estimators, like S-estimators (Rousseeuw and Yohai, 1984), MM- and τ-estimators satisfy the estimating equations of an M-estimator, which can be written as weighted normal equations, namely

$$\sum_{i=1}^{n} W\left(\frac{r_i(\boldsymbol{\beta})}{\widehat{\sigma}_0}\right)\mathbf{x}_i r_i(\boldsymbol{\beta}) = \mathbf{0}, \tag{3}$$

where W is a "weight function". Then we define, as in Yohai et al. (1991)

$$\mathbf{C}_w = \frac{1}{\sum_{i=1}^{n} w_i}\sum_{i=1}^{n} w_i \mathbf{x}_i \mathbf{x}_i', \tag{4}$$

with $w_i = W\left(r_i(\widehat{\boldsymbol{\beta}}_0)/\widehat{\sigma}_0\right)$. Put for any positive definite matrix V

$$\widehat{d}_{\mathrm{KL},V}\left(\boldsymbol{\beta}_0, \boldsymbol{\beta}\right) = \frac{1}{\widehat{\sigma}_0^2}\left(\boldsymbol{\beta} - \boldsymbol{\beta}_0\right)' \mathbf{V}\left(\boldsymbol{\beta} - \boldsymbol{\beta}_0\right). \tag{5}$$

It is immediate that (1) with $d = \widehat{d}_{\mathrm{KL},C_w}$ is equivalent to minimizing $\sum_{i=1}^n r_i\left(\boldsymbol{\beta}\right)^2$ subject to $\widehat{d}_{\mathrm{KL},C_w}\left(\widehat{\boldsymbol{\beta}}_0, \boldsymbol{\beta}\right) \le \delta$. Call $\widehat{\boldsymbol{\beta}}_{\mathrm{LS}}$ the LSE. Put for a general matrix $\mathbf{V}$:

$$\Delta_{\mathbf{V}} = \widehat{d}_{\mathrm{KL},V}\left(\widehat{\boldsymbol{\beta}}_0, \widehat{\boldsymbol{\beta}}_{\mathrm{LS}}\right).$$

Then using Lagrange multipliers, a straightforward calculation shows that in this case we have

$$\widehat{\boldsymbol{\beta}} = \begin{cases} \widehat{\boldsymbol{\beta}}_{\mathrm{LS}} & \text{if } \Delta_{\mathbf{C}_w} \le \delta \\ \left(\mathbf{X}'\mathbf{X} + \lambda \boldsymbol{C}_w\right)^{-1}\left(\mathbf{X}'\mathbf{X}\widehat{\boldsymbol{\beta}}_{\mathrm{LS}} + \lambda \boldsymbol{C}_w \widehat{\boldsymbol{\beta}}_0\right) & \text{else,} \end{cases} \tag{6}$$

where λ is determined from the equation $d_{\mathrm{KL},C_w}\left(\widehat{\boldsymbol{\beta}}_0, \boldsymbol{\beta}\right) = \delta$ and $\mathbf{C}_w$ is defined in (4). We thus see that $\widehat{\boldsymbol{\beta}}$ is a linear combination of $\widehat{\boldsymbol{\beta}}_0$ and $\widehat{\boldsymbol{\beta}}_{\mathrm{LS}}$.

Another approach is as follows. Define $\widehat{\boldsymbol{\beta}}$ as the minimizer of $\sum_{i=1}^n r_i\left(\boldsymbol{\beta}\right)^2$ subject to $\widehat{d}_{\mathrm{KL},\widehat{C}}\left(\boldsymbol{\beta}_0, \boldsymbol{\beta}\right) \le \delta$, where $\widehat{\mathbf{C}} = n^{-1}\mathbf{X}'\mathbf{X}$. In this case the solution is explicit:

$$\widehat{\boldsymbol{\beta}} = t\widehat{\boldsymbol{\beta}}_{\mathrm{LS}} + (1-t)\widehat{\boldsymbol{\beta}}_0, \tag{7}$$

where $t = \min\left(1, \sqrt{\delta/\Delta_{\widehat{\mathbf{C}}}}\right)$. Since $\Delta_{\widehat{\mathbf{C}}}$ is not robust, we now replace it with $\Delta_{\widehat{\mathbf{C}}_w}$, and therefore we choose

$$t = \min\left(1, \sqrt{\frac{\delta}{\Delta_{\widehat{\mathbf{C}}_w}}}\right). \tag{8}$$

The difference between both versions (6) and (7) showed to be negligible for all practical purposes.

It is easy to show that if $\widehat{\boldsymbol{\beta}}_0$ is regression- and affine-equivariant, so is $\widehat{\boldsymbol{\beta}}$.

2.1. Simulations

We now consider the model

$$y_i = \mathbf{x}_i'\boldsymbol{\beta} + \sigma u_i, \quad i = 1, \ldots, n, \tag{9}$$

with $\boldsymbol{\beta} \in R^p$ and $u_i \sim \mathrm{N}(0, 1)$ independent of $\mathbf{x}_i$. The performance of each estimator $\widehat{\boldsymbol{\beta}}$ will be measured by its prediction squared error, which is equivalent to $\left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\right)' \mathbf{C}_x \left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\right)$, where $\mathbf{C}_x = \mathrm{E}\mathbf{x}\mathbf{x}'$. Since all estimators considered are regression-equivariant, there is no loss of generality in taking $\boldsymbol{\beta} = \mathbf{0}$. In all cases, the distributions are normalized so that $\mathbf{C}_x = \mathbf{I}$, and therefore the criterion will be simply $\left\|\widehat{\boldsymbol{\beta}}\right\|^2$ where $\|\cdot\|$ stands for the Euclidean norm.

As initial estimator $\widehat{\boldsymbol{\beta}}_0$ for the DCML we chose the MM estimator with 85% asymptotic efficiency and bisquare ρ-function:

$$\rho_{\mathrm{bis}}(d) = 1 - \mathrm{I}(d \le 1)(1-d)^3, \tag{10}$$

where $\mathrm{I}(\cdot)$ denotes the indicator function. The reason for choosing 85% efficiency is that the maximum bias of the estimator is the same as that of the regression S-estimator, as explained in Section 5.9 of Maronna et al. (2006).

An S-estimator was also considered as an initial estimator. However, the asymptotic efficiency of these estimators is known to be less than 33%, and the finite-sample efficiency is still lower. Therefore to attain acceptable efficiencies for DCML the values δ should have to be substantially larger than the ones we employed (given in (12) below), which would entail a serious loss in robustness. These assertions were confirmed by the simulations and therefore MM was the estimator of choice.

The S estimator was computed by subsampling followed by the iteratively reweighted least squares (IRWLS) algorithm. We employed a modification suggested by Hawkins and Olive (2002), namely to add the LS and L_1 estimates to the candidates obtained by subsampling. The resulting estimator is a local (not necessarily global) minimum of the goal function defining the S-estimator. However, the resulting estimator is also strongly consistent under the model due to the following facts: (i) the LS and L_1 estimators are strongly consistent under general conditions, (ii) the M-scale of this S estimator is smaller than those of the LS and L_1 estimators, and (iii) the M-scale is a particular case of a τ-scale. Then the consistency of the S estimator follows from Theorem 4.1 of Yohai and Zamar (1988),

Using an approach similar to that in Section 5.2 of Maronna and Yohai (1993), it is possible to show that – roughly speaking – by taking a large enough number of subsamples, its breakdown point can be made as close as desired to the one of the

global minimum, with arbitrarily high probability. It follows from Yohai (1987) that these properties concerning the strong consistency and breakdown point are inherited by the MM estimator computed by IRWLS starting from the S estimator. It follows from Theorem 4.1 of Yohai (1987) that the MM estimator so computed (which is a local minimum of the respective objective function) has the same asymptotic distribution as the global minimum. Finally, the DCML estimator based on the subsampling-based MM estimator will also inherit these properties.

The initial scale $\widehat{\sigma}_0$ is a scale M estimator of the residuals, defined as the solution of

$$\frac{1}{n}\sum_{i=1}^{n}\rho\left(\frac{y_i-\mathbf{x}_i'\widehat{\boldsymbol{\beta}}_0}{\widehat{\sigma}_0}\right)=\gamma, \tag{11}$$

where ρ is the bisquare (10) function and $\gamma = 0.5\,(1-p/n)$ to attain the maximal breakdown point.

The constant δ in (1) is chosen as

$$\delta_{p,n}=0.3\frac{p}{n}. \tag{12}$$

To justify (12) note that under the model, the distribution of $nd_{\mathrm{KL}}\left(\widehat{\boldsymbol{\beta}}_0,\widehat{\boldsymbol{\beta}}_{\mathrm{LS}}\right)$ is approximately that of vz where $z\sim\chi^2_p$ and v is some constant, which implies that $\mathrm{E}d_{\mathrm{KL}}\left(\widehat{\boldsymbol{\beta}}_0,\widehat{\boldsymbol{\beta}}_{\mathrm{LS}}\right)\approx vp/n$. Therefore in order to control the efficiency of $\widehat{\boldsymbol{\beta}}$ it seems reasonable to take δ of the form Cp/n for some C. The value $C=0.3$ was arrived at after exploratory simulations aimed at striking a balance between efficiency and robustness. The behavior of the estimator is not very sensitive to the choice of the constant C; in fact, one may choose C between, say, 0.25 and 0.35 without serious effects.

2.1.1. Scenarios

Since the results may depend on the distribution of the predictors, we considered five cases, all of them including an intercept. Here each predictor vector has the form $\mathbf{x}=\left(1,x_1,\ldots,x_p\right)'$, where the x_js are i.i.d. random variables with distribution F. Note that here the number of parameters is $p+1$. In the first three cases F is standard normal, uniform in $[0,1]$ (short-tailed) and Student with four degrees of freedom (moderately heavy-tailed). In the other two, the x_js are the *squares* of standard normal and uniform variables. The Student distribution was excluded for in this case $\mathbf{C}_x=\mathrm{E}\mathbf{x}\mathbf{x}'$ does not exist since it involves the fourth moments of the t_4 distribution. We took $p=5$, 10 and 20, and $n=Kp$ with $K=5$, 10 and 20.

For each n and p we first computed the finite sample efficiency. Then to assess the estimators' robustness we contaminated the data as follows. For a contamination rate $\varepsilon\in(0,1)$ let $m=[n\varepsilon]$ where $[\cdot]$ stands for the integer part. Then for $i\le n-m$, $(\mathbf{x}_i,y_i)$ were generated according to model (9), and for $i>n-m$ we put $\mathbf{x}_i=(1,x_0,0,\ldots,0)'$ and $y_i=x_0K$, where the parameter K which regulates the slope of the contamination took on a range of values in order to determine the worst possible situations. The effect of the contamination would be to drag the first slope towards K. We took $x_0=5$ and K ranging between 0.5 and 2 with intervals of 0.1. We employed $\varepsilon=0.1$ and 0.2. The number of replications was $N_{\mathrm{rep}}=1000$.

For a given scenario and estimator $\widehat{\boldsymbol{\beta}}$ call $\widehat{\boldsymbol{\beta}}_k$, $k=1,\ldots,N_{\mathrm{rep}}$ the Monte Carlo values. As a measure of performance we employed the mean squared error: $\mathrm{MSE}=\mathrm{ave}_k\left\{\left\|\widehat{\boldsymbol{\beta}}_k\right\|^2\right\}$ where "ave" stands for the average.

2.1.2. Estimators

The estimators considered were: the Least Squares estimator, the regression S-estimator with bisquare scale (S–E), the MM estimator with bisquare loss function and 85% asymptotic efficiency, the Gervini and Yohai (2002) estimator (G–Y), the Bondell and Stefanski (2013) estimator (B–S), and the proposed estimator (DCML). Both versions (6) and (7) were considered, but since the latter yielded in general slightly better results, this is the one that is reported here. The code for B–S was kindly supplied by the authors.

2.1.3. Efficiency

We deal first with the efficiencies. In order to synthesize the results, for each combination (p,n) we took for each estimator the minimum efficiencies under normal errors over the five distributions, with respect to the MLE. The results are displayed in Table 1.

We note the following:

- The efficiency of S–E is low, as can be expected.
- When n/p is "small", the worst finite-sample efficiency of MM can be much lower than its nominal asymptotic one of 85%. The worst cases with $n/p=5$ corresponded to normal $\mathbf{x}_i$ with a quadratic term.
- The worst efficiency of G–Y is also low for small n/p.
- DCML outperforms both its initial estimator MM and G–Y.
- B–S shows the highest efficiencies in all cases.

Table 2 shows the efficiencies of the estimators with respect to the MLE for model (9) with the Student errors u_i with 3 and 5 degrees of freedom (" d.f.").

Here MM, G–Y, B–S and DCML exhibit high efficiencies, and none clearly dominates the others.

Table 1
Minimum efficiencies of estimators for normal errors over all x distributions.

p	n	S–E	MM	G–Y	B–S	DCML
5	25	0.315	0.666	0.668	0.959	0.857
	50	0.282	0.786	0.816	0.993	0.948
	100	0.279	0.822	0.875	0.997	0.983
10	50	0.290	0.699	0.718	0.991	0.922
	100	0.293	0.785	0.838	0.998	0.981
	200	0.266	0.821	0.904	0.999	0.992
20	100	0.305	0.716	0.784	0.997	0.951
	200	0.267	0.789	0.859	0.999	0.987
	400	0.260	0.837	0.928	0.999	0.998

Table 2
Efficiencies of estimators for the Student errors with 3 and 5 degrees of freedom, and normal predictors.

df	p	n	S–E	MM	G–Y	B–S	DCML
3	5	25	0.469	0.837	0.806	0.886	0.899
		50	0.458	0.929	0.876	0.890	0.924
		100	0.491	0.958	0.883	0.886	0.918
	10	50	0.421	0.860	0.842	0.893	0.903
		100	0.437	0.939	0.878	0.905	0.926
		200	0.463	0.954	0.880	0.895	0.911
	20	100	0.432	0.891	0.870	0.911	0.957
		200	0.434	0.942	0.893	0.899	0.912
		400	0.469	0.952	0.881	0.878	0.894
5	5	25	0.380	0.766	0.750	0.943	0.901
		50	0.417	0.933	0.899	0.923	0.955
		100	0.421	0.936	0.891	0.933	0.955
	10	50	0.376	0.822	0.797	0.952	0.960
		100	0.378	0.913	0.904	0.951	0.952
		200	0.396	0.942	0.924	0.943	0.949
	20	100	0.393	0.848	0.839	0.953	0.979
		200	0.349	0.900	0.883	0.927	0.947
		400	0.371	0.923	0.898	0.936	0.935

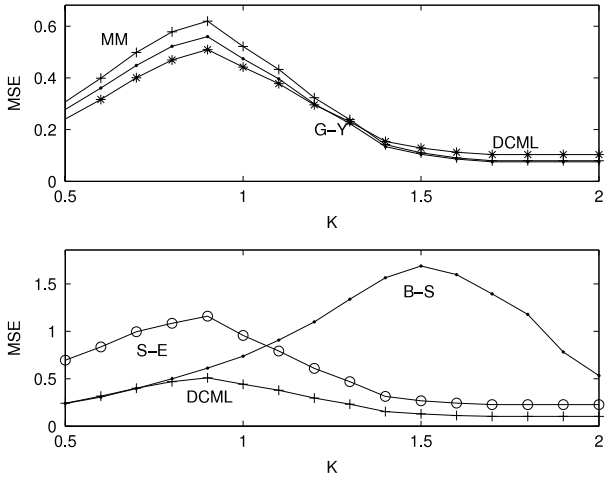

Fig. 1. MSEs of regression estimators as a function of outlier size K for normal $\mathbf{x}$, $p = 10$, $n = 200$ and $\varepsilon = 0.1$.

2.1.4. Robustness

We begin with the results of a typical case, Fig. 1 displays the MSEs of the estimators for $p = 10$, $n = 200$, normal $\mathbf{x}$, and $\varepsilon = 0.1$, for different values of the outlier size K.

In the upper panel it is seen that G–Y and DCML have similar behaviors, and that their maximum MSEs are smaller than that of MM. The lower panel shows that the MSEs of S–E and B–S are generally larger than that of DCML, the one of B–S being remarkably high.

Since all cases show approximately this same pattern, we display only the maximum MSEs over K for normal $\mathbf{x}$. Table 3 shows the results.

Table 3
Maximum mean squared errors of estimators with normal predictors for contaminated data.

ε	p	n	S–E	MM	G–Y	B–S	DCML
0.1	5	25	1.80	1.11	1.08	2.09	0.97
		50	1.25	0.77	0.79	1.71	0.66
		100	0.93	0.56	0.48	1.61	0.45
	10	50	3.04	1.79	1.68	2.89	1.40
		100	1.59	0.78	0.70	1.83	0.69
		200	1.16	0.62	0.56	1.69	0.51
	20	100	2.28	1.35	1.32	3.37	1.12
		200	1.36	0.70	0.73	2.28	0.57
		400	0.94	0.50	0.44	1.92	0.44
0.2	5	25	13.26	9.53	9.43	28.49	8.20
		50	5.83	4.27	4.09	11.37	3.70
		100	3.51	2.57	2.42	7.94	2.23
	10	50	15.12	13.54	13.27	26.57	10.19
		100	6.75	4.57	4.51	13.03	4.14
		200	3.96	2.81	2.63	9.67	2.50
	20	100	7.14	5.86	5.91	31.23	4.86
		200	3.92	3.11	2.99	13.45	2.67
		400	3.03	2.29	2.08	10.85	2.17

Some comments are in order:

- The MSEs of G–Y and DCML are similar, the latter being lower in most cases. Both outperform MM, which in turn outperforms S–E.
- The price for the high efficiency of B–S is a high contamination bias.
- When $\varepsilon = 0.2$ and $n/p = 5$ all estimators have a remarkably high MSE.

As a closing comment, the joint consideration of Tables 1–3 suggests that DCML shows the best balance between efficiency and robustness.

2.2. Asymptotic results

Assume $y = \mathbf{x}'\boldsymbol{\beta} + u$, where u is independent of $\mathbf{x}$ and has distribution F. Call σ_0 to be the limit value of the M-scale applied to u and $\mathbf{C} = \mathrm{E}(\mathbf{x}\mathbf{x}')$. It is well known that under general conditions the following expansions hold for the MM-estimator $\widehat{\boldsymbol{\beta}}_0$ and the LS estimator $\widehat{\boldsymbol{\beta}}_{\mathrm{LS}}$.

$$n^{1/2}(\widehat{\boldsymbol{\beta}}_0 - \boldsymbol{\beta}) = \frac{\sigma_0}{n^{1/2}\mathrm{E}\psi'(u_i/\sigma_0)} \sum_{i=1}^{n} C^{-1}\psi\left(\frac{u_i}{\sigma_0}\right)\mathbf{x}_i + o\left(\frac{1}{n^{1/2}}\right),$$

and

$$n^{1/2}(\widehat{\boldsymbol{\beta}}_{\mathrm{LS}} - \boldsymbol{\beta}) = \frac{1}{n^{1/2}} \sum_{i=1}^{n} C^{-1}u_i\mathbf{x}_i + o\left(\frac{1}{n^{1/2}}\right)$$

It then follows from the Central Limit Theorem that the joint asymptotic distribution $J_{\mathbf{C},\mathbf{V}}$ of $n^{1/2}(\widehat{\boldsymbol{\beta}}_{\mathrm{LS}} - \boldsymbol{\beta}, \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)$ is $J_{\mathbf{C},\mathbf{V}} = \mathrm{N}_{2p}(\mathbf{0}, \mathbf{V} \otimes \mathbf{C}^{-1})$ where $\mathbf{V} = [V_{ij}]$ is a symmetric 2×2 matrix with elements

$$V_{11} = \mathrm{E}(u^2), \qquad V_{12} = V_{21} = \sigma_0 \frac{\mathrm{E}\left(u\psi\left(u/\sigma_0\right)\right)}{\mathrm{E}\left(\psi'\left(u/\sigma_0\right)\right)}, \qquad V_{22} = \sigma_0^2 \frac{\mathrm{E}\left(\psi^2\left(u/\sigma_0\right)\right)}{\mathrm{E}\left(\psi'\left(u/\sigma_0\right)\right)}. \tag{13}$$

Let $(\mathbf{z}_1, \mathbf{z}_2)' \in R^{2p}$ be a random vector with distribution $J_{\mathbf{C},\mathbf{V}}$ and define

$$\mathbf{z}_3 = t\mathbf{z}_1 + (1-t)\mathbf{z}_2 \quad \text{with } t = \min\left(1, \sqrt{\frac{0.3p}{(\mathbf{z}_2 - \mathbf{z}_1)'\mathbf{C}(\mathbf{z}_2 - \mathbf{z}_1)}}\right). \tag{14}$$

Then the distribution $H_{\mathbf{C},\mathbf{V}}$ of $\mathbf{z}_3$ is the same as the asymptotic distribution of $n^{1/2}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})$. Note that since $\mathbf{z}_3$ is a nonlinear function of $(\mathbf{z}_1, \mathbf{z}_2)$, H is not necessarily normal. The following theorem will be useful to determine the distribution of $n^{1/2}\mathbf{b}'(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ for any $\mathbf{b} \in \mathbf{R}^p$

Theorem 1. *If* $\mathbf{C} = \mathbf{I}$, *then the distribution of* $v = \mathbf{d}'\mathbf{z}_3$ *is the same for any* $\mathbf{d} \in R^p$ *with* $\|\mathbf{d}\| = 1$.

Table 4
Asymptotic efficiency of the proposed estimator for four error distributions.

	Efficiency of DCML with respect to								
	LS			MM			MLE-Student		
	$p = 5$	10	20	5	10	20	5	10	20
Normal	0.998	0.9997	0.9999	1.18	1.18	1.18			
t_3	1.84	1.84	1.84	0.97	0.97	0.97	0.92	0.92	0.92
t_5	1.19	1.19	1.19	1.01	1.01	1.01	0.95	0.95	0.95
Uniform	1.00	1.00	1.00	1.07	1.07	1.07			

Table 5
Probability of equality of DCML and LS estimators.

$p =$	5	10	20
Normal	0.85	0.91	0.96
t_3	0.02	0.001	0.00
t_5	0.14	0.05	0.01
uniform	1.00	1.00	1.00

Proof. Let $\mathbf{D}$ be an orthogonal matrix with first row equal to $\mathbf{d}'$ and let $\mathbf{v}_j = \mathbf{D}\mathbf{z}_j$, $1 \le j \le 3$, where the $\mathbf{z}_j$s are defined above. It is easy to check that $(\mathbf{v}_1, v_2)$ has the same distribution as $(\mathbf{z}_1, \mathbf{z}_2)$, and that $\mathbf{v}_3$ satisfies

$$\mathbf{v}_3 = t\mathbf{v}_1 + (1-t)\mathbf{v}_2.$$

Besides, we have

$$(\mathbf{z}_2 - \mathbf{z}_1)'\mathbf{C}(\mathbf{z}_2 - \mathbf{z}_1) = (\mathbf{v}_2 - \mathbf{v}_1)'\mathbf{C}(\mathbf{v}_2 - \mathbf{v}_1)$$

and therefore

$$t = \min\left(1, \frac{0.3p}{(\mathbf{v}_2 - \mathbf{v}_1)'\mathbf{C}(\mathbf{v}_2 - \mathbf{v}_1)}\right).$$

Then $\mathbf{v}_3$ has the same distribution as $\mathbf{z}_3$, and therefore $v_{3,1} = \mathbf{d}'\mathbf{z}_3$ has the same distribution as $z_{3,1}$ independently of $\mathbf{d}$. ■

Call $G_{\mathbf{V}}(z)$ the distribution function of $v_{3,1}$. Suppose now that we want the distribution of $w = \mathbf{b}'\mathbf{z}_3$ for an arbitrary $\mathbf{C}$. It is easy to see that $\mathbf{z}_3^* = \mathbf{C}^{-1/2}z_3$ has distribution $H_{\mathbf{I},\mathbf{V}}$ and therefore

$$w = \mathbf{b}'\mathbf{C}^{1/2}\mathbf{z}_3^* = \|\mathbf{C}^{1/2}\mathbf{b}\|\mathbf{d}'\mathbf{z}_3^*$$

where $\|\mathbf{d}\| = 1$. Then the distribution function of w is $G_{\mathbf{V}}(w/\|C^{1/2}\mathbf{b}\|)$.

To obtain the distribution $G_{\mathbf{V}}$ we can generate a very large sample of $(\mathbf{z}_1, \mathbf{z}_2)$ (say of size 10^6) from $H_{\mathbf{I},\mathbf{V}}$ and use the transformation (14) to generate a sample of $\mathbf{z}_3$ with distribution $G_{\mathbf{V}}$. In this way we can obtain estimates of the quantiles of $G_{\mathbf{V}}$ that can be used for asymptotic inference on any linear combination of the proposed estimator $\widehat{\boldsymbol{\beta}}$. To this end, the matrix $\mathbf{V}$ can be estimated through (13), replacing F by the residual empirical distribution.

This large-sample Monte Carlo can also be used to compute the asymptotic efficiencies of $\widehat{\boldsymbol{\beta}}$ for different error distributions F. We compute those of $\widehat{\boldsymbol{\beta}}$ with respect to the LS estimator (eff_{LS}) and with respect to the MM-estimator (eff_{MM}), defined by

$$\text{eff}_{\text{LS}} = \frac{\text{E}(\mathbf{z}_1'\mathbf{C}\mathbf{z}_1)}{\text{E}(\mathbf{z}_3'\mathbf{C}\mathbf{z}_3)}, \qquad \text{eff}_{\text{MM}} = \frac{\text{E}(\mathbf{z}_2'\mathbf{C}\mathbf{z}_2)}{\text{E}(\mathbf{z}_3'\mathbf{C}\mathbf{z}_3)}.$$

Since $\mathbf{z}_1$, $\mathbf{z}_2$ and $\mathbf{z}_3$ are spheric when $\mathbf{C} = \mathbf{I}$, these efficiencies do not depend on $\mathbf{C}$. We compute these efficiencies when F is normal, uniform, and Student's t with 3 and 5 degrees of freedom. For the latter we also compute the efficiency with respect to the MLE for the Student distribution with the respective degrees of freedom. For p we chose the values 5, 10 and 20 The results are shown in Table 4.

Also, using the same sample we also computed the probabilities that $\widehat{\boldsymbol{\beta}}$ coincides with $\widehat{\boldsymbol{\beta}}_{\text{LS}}$. The results are shown in Table 5.

Combining both tables we see that:

- For normal F, DCML has almost full efficiency with respect to LS. The fact that its efficiency with respect to MM is $1.18 \approx 1/0.85$ (which is the efficiency of MM) indicates that DCML behaves almost like LS.
- For uniform F DCML always coincides with LS.
- For the Student F the behavior of DCML is more similar to that of MM. It also has a high efficiency with respect to the MLE.

Table 6
Mean interval lengths and covering probabilities for asymptotic intervals with level 0.95, averaged over the six parameters.

n	Normal		t_3	
	Length	Coverage	Length	Coverage
20	0.930	86.0%	1.204	86.1%
50	0.562	92.6%	0.743	92.7%
100	0.392	93.3%	0.518	93.7%
200	0.280	94.6%	0.364	94.3%
500	0.176	94.8%	0.228	95.2%

Finally, to assess the behavior of the asymptotic approximation for finite n, we computed the average length and covering probabilities for the asymptotic intervals with nominal 0.95 confidence level, for a model with $p = 5$ variables plus intercept, and errors with the normal and Student distributions with 3 degrees of freedom. The values were averaged over the six parameters, and were computed through a simulation with 3000 replicates. To determine the confidence interval for a given data sample we took $m = 1000$ samples from the asymptotic distribution of the DCML estimator, where each element of the sample covariance matrix was replaced by an estimate. These estimators were obtained by replacing each theoretical expectation by the empirical one with the errors replaced by residuals, and the scales replaced by their estimators. Then for each coordinate j, $1 \le j \le p$ we computed the α and $1-\alpha$ quantiles, say $\widehat{Q}_\alpha$ and $\widehat{Q}_{1-\alpha}$, of the 1000 generated samples. Then the confidence interval for the regression coefficient β_j is $[\widehat{\beta}_i - \widehat{Q}_\alpha, \widehat{\beta}_i + \widehat{Q}_\alpha]$, where $\widehat{\beta}_j$ is the DCML estimator of β_j.

Table 6 shows the results. It is seen that the approximation is reasonably good for $n \ge 50$. The results are similar to those obtained by Bondell and Stefanski (2013, Table 7).

2.3. Breakdown point

It will be shown that for the estimators employed in this paper, the finite-sample replacement breakdown point of the DCML estimator $\widehat{\boldsymbol{\beta}}$ is that of the initial estimator $\widehat{\boldsymbol{\beta}}_0$.

Consider a data set $\mathbf{Z} = \{\mathbf{z}_i, i = 1, \ldots, n\}$ with $\mathbf{z}_i = (\mathbf{x}_i, y_i)$. Let m be such that $\varepsilon = m/n$ is less than the breakdown point ε^* of $\widehat{\boldsymbol{\beta}}_0$. Let S (the "outlier set") be any set of size m contained in $\{1, \ldots, n\}$. Let $\mathbf{Z}^* = \{\mathbf{z}_i^*,\ i = 1, \ldots, n\}$ where $\mathbf{z}_i^* = \mathbf{z}_i$ for $i \notin S$ and is arbitrary for $i \in S$. We have to prove that $\widehat{\boldsymbol{\beta}}$ is bounded as a function of $\mathbf{Z}^*$. The following assumptions will be needed.

(A) The initial scale $\widehat{\sigma}_0$ is a scale M estimator given by (11) where ρ is a "bounded ρ-function" in the sense of Maronna et al. (2006, p. 31), i.e., $\rho \in [0, 1]$, $\rho(0) = 0$, and $\rho(t)$ it is a nondecreasing function of $|t|$, which is strictly increasing for $t > 0$ such that $\rho(t) < 1$.
(B) The breakdown point of $\widehat{\sigma}_0$ is $\ge \varepsilon^*$.
(C) The weight function $W(t)$ in (3) is a nondecreasing function of $|t|$ which is "matched" to ρ in the sense that $W(t) = 0$ iff $\rho(t) = 1$. This is the case in the situations considered here, where ρ is the bisquare (10) function and $W(t) = \rho'(t)/t$.
(D) Finally we assume

$$n\left(1 - \varepsilon^* - \gamma\right) \ge p \tag{15}$$

with γ in (11).

Call h the maximum number of $\mathbf{x}_i$s in a subspace. The maximal breakdown point for $\widehat{\boldsymbol{\beta}}_0$ and $\widehat{\sigma}_0$ is: $\varepsilon^*_{\max} = 0.5(n - h - 1)/n$. Here we have $\gamma = 0.5(n-p)/n \le \varepsilon^*_{\max}$ since $h \ge p - 1$, which implies (15) since $\varepsilon \le \varepsilon^*_{\max}$.

We now proceed to the proof. Recall that $\widehat{\boldsymbol{\beta}}$ satisfies

$$\frac{1}{\widehat{\sigma}_0^2}\left(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}_0\right)' \mathbf{C}_w \left(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}_0\right) \le \delta,$$

where $\mathbf{C}_w$ is defined in (4). Recall that $\widehat{\boldsymbol{\beta}}_0$, $\widehat{\sigma}_0$ and $\mathbf{C}_w$ depend on $\mathbf{Z}^*$. Since $\varepsilon < \varepsilon^*$ there exist constants a, b, c such that for all S and $\mathbf{Z}^*$:

$$0 < a \le \widehat{\sigma}_0 \le b, \quad \left\|\widehat{\boldsymbol{\beta}}_0\right\| \le c.$$

Also, since $\varepsilon < \varepsilon^*$ there exists $\eta \in (0, 1)$ such that

$$n\left(1 - \varepsilon - \frac{\gamma}{1-\eta}\right) \ge p. \tag{16}$$

Let $t_0 > 0$ be such that $\rho(t_0) = 1 - \eta$, and put $w_0 = W(t_0)$. Then by (C) $|t| \le t_0$ implies $W(t) \ge w_0 > 0$. Let

$$N = N\left(\mathbf{Z}^*\right) = \#\left\{i \notin S : \rho\left(\frac{y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}_0}{\widehat{\sigma}_0}\right) \le 1 - \eta\right\}.$$

Then it follows from (11) that

$$n\delta \geq \sum_{i\notin S} \rho\left(\frac{y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}_0}{\widehat{\sigma}_0}\right) \geq (n - m - N)(1-\eta),$$

and therefore by (16), since $\varepsilon < \varepsilon^*$

$$N\left(\mathbf{Z}^*\right) \geq n - n\varepsilon - \frac{n\gamma}{1-\eta} \geq p \quad \forall\, \mathbf{Z}^*.$$

Call $\mathcal{A}$ the set of all subsets of $\{1, \ldots, n\}$ of size $h+1$. Put

$$\lambda_0 = \min_{A\in\mathcal{A}} \lambda_{\min}\left(\sum_{i\in A} \mathbf{x}_i\mathbf{x}_i'\right),$$

where $\lambda_{\min}$ denotes the smallest eigenvalue of a matrix. Then $\lambda_0 > 0$. For any vector $\mathbf{a}$ and all $\mathbf{Z}^*$ we have

$$\mathbf{a}'\mathbf{C}_w\mathbf{a} \geq \mathbf{a}'\left[\sum_{i\notin S} W\left(\frac{y_i - \mathbf{x}_i'\widehat{\boldsymbol{\beta}}_0}{\widehat{\sigma}_0}\right)\mathbf{x}_i\mathbf{x}_i'\right]\mathbf{a} \geq w_0\lambda_0\|\mathbf{a}\|^2,$$

and therefore we have for all $\mathbf{Z}^*$

$$\delta\widehat{\sigma}_0^2 \geq \left(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}_0\right)'\mathbf{C}_w\left(\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}_0\right) \geq w_0\lambda_0\left\|\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}_0\right\|^2,$$

which, in view of the boundedness of $\widehat{\boldsymbol{\beta}}_0$ and $\widehat{\sigma}_0$, implies that $\widehat{\boldsymbol{\beta}}$ is bounded.

3. Multivariate estimation

Consider observations $\mathbf{x}_i$, $i = 1, \ldots, n$ with a normal p-variate distribution $\mathrm{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Let $\left(\widehat{\boldsymbol{\mu}}_0, \widehat{\boldsymbol{\Sigma}}_0\right)$ be a robust estimator of multivariate location and scatter. We shall treat $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ separately.

For the estimation of $\boldsymbol{\Sigma}$ we have, considering $\boldsymbol{\mu}$ as a nuisance parameter:

$$d_{\mathrm{KL}}(\boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}) = \log|\boldsymbol{\Sigma}| - \log|\boldsymbol{\Sigma}_0| + \operatorname{trace}\left(\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_0\right) - p, \tag{17}$$

where $|\cdot|$ denotes the determinant. Our procedure amounts to

$$\widehat{\boldsymbol{\Sigma}} = \arg\min_{\boldsymbol{\Sigma}}\left[n\log|\boldsymbol{\Sigma}| + \sum_{i=1}^{n}(\mathbf{x}_i - \boldsymbol{\mu})\boldsymbol{\Sigma}^{-1}(\mathbf{x}_i - \boldsymbol{\mu})\right] \tag{18}$$

with $d_{\mathrm{KL}}\left(\widehat{\boldsymbol{\Sigma}}_0, \boldsymbol{\Sigma}\right) \leq \delta$.

Call $\widehat{\boldsymbol{\Sigma}}_{\mathrm{ML}}$ the MLE of $\boldsymbol{\Sigma}$, i.e. the sample covariance matrix. Put $d_0 = d_{\mathrm{KL}}\left(\widehat{\boldsymbol{\Sigma}}_0, \widehat{\boldsymbol{\Sigma}}_{\mathrm{ML}}\right)$. Then using Lagrange multipliers, a straightforward calculation shows that

$$\widehat{\boldsymbol{\Sigma}} = (1-t)\widehat{\boldsymbol{\Sigma}}_{\mathrm{ML}} + t\widehat{\boldsymbol{\Sigma}}_0, \tag{19}$$

where $t = 0$ if $d_0 \leq \delta$, and is otherwise determined from the equation $d_{\mathrm{KL}}\left(\widehat{\boldsymbol{\Sigma}}_0, \boldsymbol{\Sigma}\right) = \delta$, which is easily derived from (17)–(19).

We now turn to $\boldsymbol{\mu}$. We have

$$d_{\mathrm{KL}}(\boldsymbol{\mu}_0, \boldsymbol{\mu}) = \left(\boldsymbol{\mu} - \boldsymbol{\mu}_0\right)'\Sigma^{-1}\left(\boldsymbol{\mu} - \boldsymbol{\mu}_0\right).$$

The estimator is then defined by

$$\sum_{i=1}^{n}(\mathbf{x}_i - \boldsymbol{\mu})\boldsymbol{\Sigma}^{-1}(\mathbf{x}_i - \boldsymbol{\mu}) = \min \tag{20}$$

with $d_{\mathrm{KL}}\left(\boldsymbol{\mu}_0, \boldsymbol{\mu}\right) \leq \delta$. Let $\overline{\mathbf{x}}$ be the sample mean, and define

$$d_0 = \left(\overline{\mathbf{x}} - \widehat{\boldsymbol{\mu}}_0\right)'\widehat{\Sigma}_0^{-1}\left(\mathbf{x} - \widehat{\boldsymbol{\mu}}_0\right).$$

Then a straightforward calculation shows that

$$\widehat{\boldsymbol{\mu}} = t\overline{\mathbf{x}} + (1-t)\widehat{\boldsymbol{\mu}}_0 \tag{21}$$

with

$$t = \min\left(1, \sqrt{\frac{\delta}{d_0}}\right).$$

It is easy to show that if the initial estimators are affine-equivariant, so are the resulting ones.

Remark. Unlike the regression and location cases, $d_{\mathrm{KL}}(\boldsymbol{\Sigma}_0, \boldsymbol{\Sigma})$ is not symmetric in its arguments. Here we have chosen the form (17) because it yields the simple intuitive result (19), while the alternative order yields a more complicated result.

Table 7
Constants for the approximate computation of δ.

	a	b	c
$\boldsymbol{\Sigma}$	1.02	0.82	0.18
$\boldsymbol{\mu}$	0.55	0.88	−0.30

3.1. Simulations

As initial estimator we employ an S estimator (Davies, 1987) with bisquare scale, computed as described at the end of page 199 of Maronna et al. (2006). For the reasons given in Section 2.1, in order to ensure consistency the sample mean and covariance matrix were included as candidates for location and scale respectively.

This study includes $p = 2$, 5 and 10. The reason larger values of p are not included is the following. Rocke (1996) found out that the efficiency of S estimators with a monotone weight function increases with p, and therefore there is little to be gained with DCML when p is large.

We now define the S estimator. For $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denote the (squared) Mahalanobis distance of $\mathbf{x}$ as

$$d(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = (\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}).$$

Define a scale M estimator $\widehat{\sigma} = \widehat{\sigma}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ as the solution of

$$\frac{1}{n} \sum_{i=1}^{n} \rho \left(\frac{d(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma})^{1/2}}{\sigma} \right) = \gamma,$$

where ρ is the bisquare ρ-function (10), and $\gamma = 0.5\,(1 - p/n)$ which ensures maximal breakdown point. The S estimator is defined by

$$\left(\widehat{\boldsymbol{\mu}}_0, \widetilde{\boldsymbol{\Sigma}}\right) = \arg\min \left\{ \widehat{\sigma}(\mathbf{t}, \mathbf{V}) : \mathbf{t} \in R^p, |\mathbf{V}| = 1 \right\}$$

Since $|\widetilde{\boldsymbol{\Sigma}}| = 1$, we have to scale $\widetilde{\boldsymbol{\Sigma}}$ to make it a consistent estimator of the covariance matrix under normality. Put $d_i = d\left(\mathbf{x}_i, \widehat{\boldsymbol{\mu}}_0, \widetilde{\boldsymbol{\Sigma}}\right)$ and call χ^2_p the chi-squared distribution with p degrees of freedom. Then define

$$\widehat{\boldsymbol{\Sigma}}_0 = \frac{\text{median}_i\{d_i\}}{\text{median}(\chi^2_p)} \widetilde{\boldsymbol{\Sigma}}.$$

The constants δ in (18) and (20) were chosen as

$$\delta = a n^{-b} p^l, \tag{22}$$

with (a, b, c) given in Table 7.

The motivation for this choice is as follows. It was considered as reasonable to choose for each (p, n) , δ as some α-quantile of d_{KL} under the nominal model, i.e. the multivariate normal distribution. Exploratory simulations suggested α between 0.4 and 0.6. The quantiles were computed by simulation for p between 2 and 10 and n between $5p$ and 500. Then for each α the α-quantile was fitted by regression as a function of n and p of the form (22). Finally, after the simulation was completed, it was decided that $\alpha = 0.4$ yielded the best results. As was the case in regression, the results are not overly sensitive to the choice of α.

The values of c indicate that when p increases, the quantiles for $\boldsymbol{\Sigma}$ increase very slowly, and those for $\boldsymbol{\mu}$ decrease. This fact may seem counter-intuitive, but it is a consequence of the increasing efficiency of the S estimator: when p increases, the S estimator becomes "closer" to the classical one, which makes d_{KL} smaller.

For each n and p we generate $N_{\text{rep}} = 1000$ samples of size n from $\text{N}_p(\mathbf{0}, \mathbf{I})$. For a contamination rate ε, the first $m = [n\varepsilon]$ elements (rows) of $\mathbf{X}$ are replaced by $(K, 0, \ldots, 0)$ where K ranges between 1 and 10. For each sample three estimators were computed: the sample mean and covariance matrix, the S estimator, and the DCML estimator given by (19)–(21).

For each scenario, each estimator is evaluated by its "loss" defined as $\|\widehat{\boldsymbol{\mu}}\|^2$ for location and as $d_{\text{KL}}\left(\widehat{\boldsymbol{\Sigma}}, \mathbf{I}\right) = \text{trace}\left(\widehat{\boldsymbol{\Sigma}}\right) - \log|\widehat{\boldsymbol{\Sigma}}| - p$ for scatter, and the results were summarized by the respective mean losses. Table 8 shows the efficiencies, defined as the ratio of the mean losses of the classical and the robust estimator.

It is seen that DCML is able to substantially increase the efficiency of S–E, especially for $p = 2$. The efficiency for location is much higher than for scatter. The fact that the efficiency of S–E increases with p is also clear. Actually, for $p = 15$ the efficiency of S–E is ≥ 0.96.

Table 9 shows the maximum mean losses for contamination rate $\varepsilon = 0.1$. It is seen that in general the price for the increase in efficiency is at worst a small increase of the maximum loss and at best a decrease thereof. Fig. 2 compares the losses of S–E and DCML as a function of the outlier size K for $\varepsilon = 0.1$.

Table 8
Efficiencies of estimators.

p	n	$\boldsymbol{\Sigma}$		$\boldsymbol{\mu}$	
		S–E	DCML	S–E	DCML
2	10	0.417	0.645	0.714	0.893
	20	0.428	0.703	0.682	0.904
	40	0.426	0.779	0.608	0.881
5	25	0.783	0.933	0.903	0.980
	50	0.794	0.972	0.890	0.985
	100	0.791	0.995	0.871	0.983
10	50	0.947	0.996	0.961	0.997
	100	0.939	0.997	0.964	0.996
	200	0.932	0.998	0.960	0.999

Table 9
Simulation: maximum mean losses of estimators.

ε	p	n	$\boldsymbol{\Sigma}$		$\boldsymbol{\mu}$	
			S–E	DCML	S–E	DCML
0.1	2	10	1.02	1.15	0.37	0.39
		20	0.59	0.62	0.23	0.24
		40	0.31	0.34	0.12	0.14
	5	25	1.15	1.17	0.29	0.31
		50	0.73	0.84	0.22	0.25
		100	0.46	0.52	0.13	0.15
	10	50	3.26	3.57	0.49	0.58
		100	2.04	2.25	0.33	0.35
		200	1.56	1.92	0.24	0.31
0.2	2	10	1.59	1.66	0.57	0.58
		20	1.03	0.80	0.37	0.41
		40	0.77	0.57	0.32	0.34
	5	25	4.01	3.78	1.08	1.33
		50	2.76	3.70	0.76	0.98
		100	2.43	2.30	0.66	0.81
	10	50	12.51	12.32	2.74	3.19
		100	7.15	7.17	1.80	2.16
		200	6.29	6.26	1.60	1.93

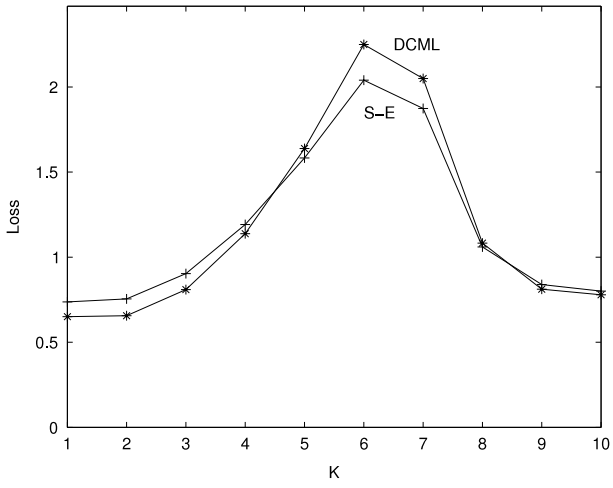

Fig. 2. Losses of scatter matrices for $p = 10$, $n = 100$ and 10% contamination, as a function of outlier size.

3.2. Breakdown point

It is easy to show that the replacement breakdown point of the DCML estimators is that of the initial ones. We give the details for $\widehat{\boldsymbol{\Sigma}}$, the case of $\widehat{\boldsymbol{\mu}}$ being similar. Consider a data set $\mathbf{X} = \{\mathbf{x}_i\ i = 1, \ldots, n\}$. Let m be such that $\varepsilon = m/n$ is less than the breakdown point ε^* of the initial estimator $\widehat{\boldsymbol{\Sigma}}_0$. Let $\mathbf{X}^*$ be a data set that coincides with $\mathbf{X}$ except for m elements which are arbitrary. We have to prove that, as a function of $\mathbf{X}^*$, the largest eigenvalue $\lambda_{\max}$ of $\widehat{\boldsymbol{\Sigma}}$ is bounded, and the smallest one

Table 10
Stack loss data: prediction RMSEs of estimators for "good" data.

Computed with	LS	S–E	MM	G–Y	B–S	DCML
Good data	1.095	1.416	1.126	1.095	1.095	1.095
Whole data	1.921	1.143	1.100	1.322	1.484	1.164

Table 11
Philips data: Kullback–Leibler distances between estimators and "true values".

	Computed with	MLE	S–E	DCML
Scatter	Good data		0.381	0.286
	Whole data	6.282	0.381	0.322
Location	Good data		0.051	0.039
	Whole data	1.067	0.051	0.044

$\lambda_{\min}$ is bounded away from zero. We know that this property holds for $\widehat{\mathbf{\Sigma}}_0$. Since by (17)

$$\log|\widehat{\mathbf{\Sigma}}| - \log|\widehat{\mathbf{\Sigma}}_0| + \operatorname{trace}\left(\widehat{\mathbf{\Sigma}}^{-1}\widehat{\mathbf{\Sigma}}_0\right) - p \le \delta,$$

it follows from the "trace" term that $\lambda_{\min}$ cannot tend to zero, and then it follows from the "log" term that $\lambda_{\max}$ cannot tend to infinity. ■

4. Real data

In this section we apply the proposed estimators to two published data sets. In both cases, the tuning parameters for DCML were chosen in the same way as in the simulations.

4.1. Regression

We consider the well-known stack loss data set with $n = 21$ and $p = 3$ plus intercept. Lacking a "true model" we have to employ alternative criteria for robustness and efficiency.

There seems to be a general agreement to consider observations 1, 3, 4 and 21 as atypical; see Rousseeuw and Leroy (1987). Call "good data" the data set without {1, 3, 4, 21}. The estimators were first computed using the good data, and the root mean squared prediction errors (RMSE: square root of the mean of the squared residuals) were computed for the same data. The comparison with LS was employed as a surrogate criterion for efficiency. For a surrogate criterion for robustness, the estimators were then computed for the whole data set, and the RMSE again computed *only* for the good data. Table 10 shows the results.

The first row shows that G–Y, B–S and DCML are here "fully efficient", S–E is rather inefficient, and MM has a high efficiency. The second row shows S–E, MM and DCML as most robust, followed by G–Y, and B–S as the less robust one.

The behavior of S–E is puzzling. It gives zero weights to some "good" observations. The estimator was recomputed several times to rule out the effect of the subsampling, but the results remained the same.

4.2. Multivariate estimation

Here we choose the Philips Mecoma data, employed in Problem 1 in Rousseeuw and Van Driessen (1999), with $n = 677$ and $p = 9$. Plotting the Mahalanobis distances from the S estimator shows a number of clear outliers, the sequence with indexes between 491 and 565 being the most outstanding ones. We defined as "bad data" the observations with Mahalanobis distances larger than 60, which yielded 80 observations. Lacking a criterion similar to prediction error like in the former example, we defined as the "true parameters" the MLE (mean and covariance matrix) applied to the "good" data, which will be called $\boldsymbol{\mu}_{\text{good}}$ and $\mathbf{\Sigma}_{\text{good}}$, respectively.

We then computed, as above, the estimators based on the "good" data and their Kullback–Leibler distances to the "truth"; and then did the same for the estimators based on the whole data. Namely, we computed

$$d = \operatorname{trace}\left(\mathbf{\Sigma}_{\text{good}}^{-1}\mathbf{V}\right) - p - \log|\mathbf{\Sigma}_{\text{good}}^{-1}\mathbf{V}|$$

for each scatter estimator $\mathbf{V}$, and

$$d = \left(\mathbf{t} - \boldsymbol{\mu}_{\text{good}}\right)' \mathbf{C}_{\text{good}}^{-1} \left(\mathbf{t} - \boldsymbol{\mu}_{\text{good}}\right)$$

for each location estimator $\mathbf{t}$. DCML was computed as in (12). Table 11 shows the results.

It is seen that here DCML outperforms S–E in all cases.

5. Conclusions

We propose an approach to improve on the finite-sample efficiency of robust estimators with a very small damage to their robustness. The approach is developed for linear regression and for the estimation of multivariate location and scatter. In both cases our proposal clearly outperforms its competitors. In the regression case, an asymptotic theory is developed, which can be employed for approximate inference.

Acknowledgments

This research was partially supported by Grants PIP 112-200801-00216 and 112-201101-00339 from CONICET, PICT 2011-397 from ANPCYT and W276 from the University of Buenos Aires, Argentina.

References

Bondell, H.D., Stefanski, L.A., 2013. Efficient robust regression via two-stage generalized empirical likelihood. J. Amer. Statist. Assoc. 108, 644–655.
Davies, P.L., 1987. Asymptotic behavior of S-estimators of multivariate location parameters and dispersion matrices. Ann. Statist. 15, 1269–1292.
Gervini, D., Yohai, V.J., 2002. A class of robust and fully efficient regression estimators. Ann. Statist. 30, 583–616.
Hampel, F.R., 1971. A general definition of qualitative robustness. Ann. Math. Stat. 42, 1887–1896.
Hawkins, D.M., Olive, D.J., 2002. Inconsistency of resampling algorithms for high-breakdown regression estimators and a new algorithm. J. Amer. Statist. Assoc. 97, 136–159.
Huber, P.J., 1964. Robust estimation of a location parameter. Ann. Math. Stat. 35, 73–101.
Maronna, R.A., Martin, R.D., Yohai, V.J., 2006. Robust Statistics: Theory and Methods. John Wiley and Sons, New York.
Maronna, R.A., Yohai, V.J., 1993. Bias-robust estimates of regression based on projections. Ann. Statist. 21, 577–1117.
Olive, D.J., Hawkins, D.M., 2010. Robust multivariate location and dispersion, available online at http://lagrange.math.siu.edu/Olive/pphbmld.pdf.
Olive, D.J., Hawkins, D.M., 2011. Practical high breakdown regression, available online at http://lagrange.math.siu.edu/Olive/pphbreg.pdf.
Rocke, D.M., 1996. Robustness properties of *S*-estimators of multivariate location and shape in high dimension. Ann. Statist. 24, 1327–1345.
Rousseeuw, P.J., Leroy, A.M., 1987. Robust Regression and Outlier Detection. John Wiley and Sons, New York.
Rousseeuw, P.J., Van Driessen, K., 1999. A fast algorithm for the minimum covariance determinant estimator. Technometrics 41, 212–223.
Rousseeuw, P.J, Yohai, V.J., 1984. Robust regression by means of S-estimators. In: Franke, J., Härdle y, W., Martin, R.D. (Eds.), Robust and Nonlinear Time Series. In: Lectures Notes in Statistics, vol. 26. Springer Verlag, New York, pp. 256–272.
Yohai, V.J., 1987. High breakdown–point and high efficiency estimates for regression. Ann. Statist. 15, 642–665.
Yohai, V.J., Stahel, W.A., Zamar, R.H., 1991. A procedure for robust estimation and inference in linear regression. In: Stahel, W., Weisberg, S. (Eds.), Directions in Robust Statistics and Diagnostics (Part II). In: The IMA Volumes in Mathematics and is Applications, Springer Verlag, New York, pp. 365–374.
Yohai, V.J., Zamar, R.H., 1988. High breakdown estimates of regression by means of the minimization of an efficient scale. J. Amer. Statist. Assoc. 83, 406–413.
Zhang, J., Olive, D.J., Ye, P., 2012. Robust covariance matrix estimation with canonical correlation analysis. Int. J. Stat. Probab. 1, 119–136.