

QUANTIFICATION OF ESTIMATION INSTABILITY AND ITS APPLICATION TO THRESHOLD SELECTION IN EXTREMES

Thomas L. Berning
Stellenbosch University
e-mail: *tberning@sun.ac.za*

Key words: estimation instability, extreme value index, threshold selection, bias-reduced estimator.

Summary: The main goal of this paper is to propose a measure which quantifies the instability of estimates over a range of chosen values of some other parameter. The measure is used to identify a region where the estimates are considered "stable". Methods for identifying stable regions are then developed. These methods are applied to threshold selection in an extreme value analysis context, where the perturbed Pareto distribution is fitted to observed relative excesses. As a result a more accurate estimator of the extreme value index is obtained. As a further application the instability measure is employed in second order parameter estimation, which leads to an adjustment of an existing estimator, having desirable properties as far as its role in threshold selection is concerned, as well as improved estimation of the second order parameter for large samples.

1. Introduction

In extreme value theory (EVT) the emphasis is on describing the tail of the underlying distribution. In this paper the focus will be on the right tail of the underlying distribution, which is assumed to be heavy-tailed.

A distribution is heavy-tailed if and only if its right tail function is regularly varying (de Haan and Ferreira, 2006). The definition is given below. See for instance Geluk, de Haan, Resnick and Stărică (1997).

Definition 1 Let X be a random variable with distribution function F concentrated on $[0, \infty)$ and tail function $\bar{F} = 1 - F$. \bar{F} is said to be regularly varying with index $-1/\gamma$, $\gamma > 0$, if

$$\lim_{t \rightarrow \infty} \frac{\bar{F}(tx)}{\bar{F}(t)} = x^{-\frac{1}{\gamma}} \quad (1)$$

for all $x > 0$.

The right hand side of 1 can be recognised as the tail function of the Pareto distribution. (It is defined for $x \geq 1$ in case of the Pareto distribution.)

The parameter γ is called the extreme value index (EVI). A positive EVI indicates a heavy right tail. The EVI is the crucial parameter, since accurate estimation of the EVI leads to accurate inferences concerning extreme quantiles. In this paper attention will only be given to EVI estimation. Once the methodology concerning the estimation of the EVI has been established, the extension to inferences concerning extreme quantiles is standard for maximum likelihood estimation (MLE). Refer for instance to Coles (2001).

Assume a sequence of independent, identically distributed random variables X_1, X_2, \dots, X_n with common heavy-tailed distribution F , concentrated on $[0, \infty)$. Denote the corresponding order statistics by $X_{1,n} \leq X_{2,n} \leq \dots \leq X_{n,n}$. Let k denote the number of observations exceeding a positive threshold t . The relation between t and k will be taken as $t = X_{n-k,n}$. If the random variable X has distribution F , and Z is defined as the relative (or multiplicative) excess above the threshold t , then $Z = X/t$ given $X \geq t$. For $j = 1, 2, \dots, k$, define $Z_j = X_i/t$ conditional on $X_i > t$, where i is the index of the j -th excess, with $Z_j \geq 1$.

Let \bar{F}_t be the tail function of Z . Then, applying Definition 1, it follows that

$$\bar{F}_t(z) = P(X/t > z | X \geq t) = P(X > tz | X \geq t) = P(X > tz) / P(X \geq t) = \bar{F}(tz) / \bar{F}(t) \rightarrow z^{-1/\gamma}$$

as $t \rightarrow \infty$ for all $z > 0$ (and in particular for $z \geq 1$), where $\gamma > 0$.

In words this result states that the distribution of the relative excesses from a heavy-tailed distribution is approximately Pareto, if the threshold is large.

The most natural way to estimate the EVI for relative excesses Z_1, Z_2, \dots, Z_k is to fit the Pareto distribution by means of maximum likelihood estimation. It is easy to show that this yields $H_{k,n} := \frac{1}{k} \sum_{i=1}^k \log Z_i$, which is the famous Hill estimator of the EVI (Hill, 1975). Note that the formula for the Hill estimator is derived as one would an MLE, but strictly speaking the variables Z_1, Z_2, \dots, Z_k are not independent.

The Hill estimator performs poorly as far as accuracy of estimation is concerned, since the estimate is extremely sensitive to the choice of k . For small k the estimates have low bias, but large variance. This is because the threshold is large enough for the first order approximation (Pareto distribution) to be valid, but the number of observations k on which the estimate is based, is small. Conversely large values of k lead to low variance, but large bias.

A variety of methods for choosing the threshold (or k) has been proposed in the literature, for instance the method by Drees and Kaufmann (1998), the method by Guillou and Hall (2001), and the method by Beirlant, Goegebeur, Segers and Teugels (2004), the last of which is based on the minimisation of the estimated asymptotic mean square error. Despite the improvements in estimation of the EVI brought about by these threshold selection methods, the Hill estimator still performs unsatisfactorily.

A proposed solution to the bias problem associated with the Hill estimator is to consider a bias-reduced estimator, obtained in this instance by fitting the perturbed Pareto distribution (PPD) to the excesses, instead of the (strict) Pareto as in the case of the Hill estimator. The motivation behind this approach is to obtain a larger range of values of k for which the bias is relatively low.

In the above setting, being able to *objectively* determine such a range of values of k for which the estimator is “stable” will not only be more scientific, but will also enable comparison of these methods by means of simulation studies. A method for quantifying the stability (or rather instability)

of estimates over a range of values of k is developed in this paper. The *stable region* is then defined as the region for which the instability measure yields the lowest value.

Even though the concept is only illustrated in the context of EVT, the same measure can be used in any application where the stability of estimates needs to be quantified over a range of some other parameter. Examples of such situations are mentioned in the concluding remarks at the end of the paper.

The purpose of this paper is to define a measure which quantifies the instability of estimates and applying this measure to develop a method of threshold selection when fitting the PPD to relative excesses. Note that the purpose is not an exhaustive comparative study of all estimators of the first and second order parameters or of all threshold selection methods.

The next section reviews some results from the literature. The necessity of bias-reduced estimators is demonstrated by considering the properties of the Hill estimator for a specific sample. Fitting the PPD to the excesses instead of the Pareto yields a second order generalisation of the Hill estimator. Some computational aspects are also highlighted.

In Section 3 the proposed instability measure is defined. It is then described how the instability measure can be used to locate *stable regions* (sets of consecutive estimates which yield a low value of the instability measure). Once a stable region has been identified, the EVI is estimated as the mean of the estimates over the stable region, leading to the concept of an implied threshold. Section 4 discusses six candidate external estimators of the second order parameter. The choice of this estimator crucially affects the accuracy of the resulting EVI estimation. In Section 5 a simulation study design is proposed which is applicable to any study of estimators of a positive EVI. This design is then applied to assess the performance of the second order parameter estimators, reducing the number of candidate external estimators to three. In Section 6 it is shown how the proposed EVI estimators are constructed. This section also includes the results of a simulation study comparing the performance of these estimators, including some results illustrating the effect of adaptive threshold selection for the Hill estimator. Some concluding remarks appear at the end of the paper.

2. Second order estimation

Consider a second order regular variation condition on the right tail function. Refer for instance to Geluk et al. (1997).

Definition 2 *A regularly varying tail function \bar{F} is said to be second order regularly varying with first order index $-1/\gamma$, $\gamma > 0$, and second order index ρ/γ , $\rho \leq 0$, if a function $A(t)$ exists which tends to 0 and is ultimately of constant sign as $t \rightarrow \infty$, such that*

$$\lim_{t \rightarrow \infty} \left(\frac{\bar{F}(tx)}{\bar{F}(t)} - x^{-1/\gamma} \right) (A(t))^{-1} = H(x)$$

for all $x > 0$, where $H(x) = dx^{-1/\gamma} \log x$ if $\rho = 0$ and $H(x) = dx^{-1/\gamma} \frac{x^{\rho/\gamma} - 1}{\rho/\gamma}$ if $\rho < 0$, with $-\infty < d < \infty$.

In this definition, γ is still the EVI (also called the *first order parameter*) and ρ is called the *second order parameter*.

Note that Geluk et al. (1997) specify that $d \neq 0$. de Haan and Ferreira (2006) also specify $d \neq 0$. This restriction is not applied in Beirlant et al. (2004).

Whether or not d may be zero is of theoretical rather than practical significance. Theoretically we should have $d \neq 0$, since if $d = 0$, ρ is undefined, and consequently the second order index is also undefined. The advantage of allowing d to be zero is that the Pareto distribution can be included as a special case. For the Pareto distribution $\bar{F}(tx)/\bar{F}(t) = x^{-1/\gamma}$, and therefore $d = 0$.

For all practical purposes one can assume that all heavy-tailed distributions of interest are first and second order regularly varying. Following limiting arguments similar to those shown above for the first order case, one can show that the second order condition implies that the relative excesses approximately follow the perturbed Pareto distribution (PPD), if the threshold is large. The definition of the PPD follows.

Definition 3 A random variable X is said to be distributed perturbed Pareto with parameters γ , ρ and c , denoted $X \sim \text{PPD}(\gamma, \rho, c)$, when

$$\bar{F}(x) = (1 - c)x^{-1/\gamma} + cx^{-(1-\rho)/\gamma},$$

where $x \geq 1$, $\gamma > 0$, $\rho < 0$ and $\rho^{-1} \leq c \leq 1$.

The restriction $\rho^{-1} \leq c \leq 1$ ensures that the density function $f(x) \geq 0$ for all $x \geq 1$. The Pareto distribution is a special case of the PPD when $c = 0$ (in which case ρ is not zero, but undefined).

The EVI γ is a parameter of the PPD which can be estimated using MLE, resulting in a second order generalisation of the Hill estimator.

Figure 1 below shows results pertaining to a sample of size $n = 200$ observations simulated from a Burr($\beta = 1, \gamma = 0.5, \rho = -0.7$) distribution. (Refer to Section 5 for more detail concerning the Burr distribution.) In the first graph the Hill estimates of the EVI were calculated for $k = 1, 2, \dots, 150$. In the second graph the PPD was fitted to k multiplicative excesses by means of MLE. The resulting estimates of the parameter γ (the EVI) is shown for $k = 1, 2, \dots, 150$.

The horizontal dashed line indicates the true value of the EVI, namely 0.5.

The first graph of Figure 1 clearly illustrates the sensitivity of the Hill estimate with respect to the choice of k . It is evident from the second graph that a larger value of k (and larger range of k values) can be used when fitting the PPD. This illustrates the typical (desired) effect of a bias-reduced estimator.

Some computational remarks are in order with respect to the MLEs of the PPD parameters.

The first concerns the estimation of ρ . One option is not to estimate ρ at all, but to fix it at some value, usually -1 . See for instance Beirlant et al. (2004) and Dierckx (2000). The second option is to estimate ρ simultaneously with the other two parameters, γ and c . The third option, called *external estimation of the second order parameter*, entails using some existing estimator of ρ , substituting the value of ρ in Definition 3 by its estimated value $\hat{\rho}$, and estimating the remaining two parameters using MLE.

The first option was popular before estimators of ρ existed which perform well. In the simulation studies which were carried out, the estimator of the EVI which performed best was obtained by external estimation of ρ . This is in line with theoretical results which state that external estimation of ρ leads to a smaller asymptotic variance in reduced bias estimators. See for instance Gomes, Figueiredo and Neves (2012). The choice of estimator of ρ will receive attention in later sections.




Figure 1: Comparison of first and second order estimates.

The second issue is the computational restrictions on the range of the parameter c . From extensive simulations it was found that applying the restriction $c \leq 0.5$ substantially improves the EVI estimation accuracy. This restriction also makes sense from a theoretical point of view. For $0 \leq c \leq 1$ the PPD survivor function is the weighted sum of a first order component $x^{-1/\gamma}$, and a second order component $x^{-(1-\rho)/\gamma}$. When $c > 0.5$ there is a greater weight on the second order component than on the first, which is not desirable.

3. An instability measure to determine regions of stable EVI estimates

Considering the second graph of Figure 1, one can see that values of k between 40 and 80 (roughly) would yield the most accurate estimates of the EVI. This is the region where the estimates are the most stable. What is meant by “stable” in this context is also clear: one needs to identify a region where the estimates do not vary excessively (low variance), and where the slope is more or less zero (no systematic change in bias). In summary, a region which optimises the variance-bias trade-off needs to be identified.

The proposed measure quantifies the instability of a quantity y with respect to chosen values of another quantity x .

Definition 4 For $m \geq 2$, let y_1, y_2, \dots, y_m denote the observed values of a quantity y corresponding to chosen values x_1, x_2, \dots, x_m of another quantity x , respectively.

The instability of y with respect to x is defined as

$$\vartheta^2 = \sigma^2 + b^2,$$

where σ^2 is the sample variance of the y values, and b is the slope of the simple least squares regression line of y on x .

In the context of threshold selection in EVT, $\{y_i\}$ is the set of estimated EVIs, and $\{x_i\}$ a set of values corresponding to the respective choices of k . The interpretation of ϑ^2 is similar to that of an MSE, namely that it is, in the sense described above, the sum of the variance and the square of the (change in) bias.

Applying the instability measure for the purpose of threshold selection, the assumption is made that the x values are equally spaced. Prior to calculating the instability measure, the x and y values are also normalised in such a way that the measure ensures location and scale invariance with respect to both the original x values and the original y values.

The chosen x values $x_1 < x_2 < \dots < x_m$ can be normalised by calculating $(x_i - x_1)/(x_m - x_1)$ for $i = 1, 2, \dots, m$. If the original x values are equally spaced and m is large, these values are close to $x_i^* = i/m$. The latter is not only used for the sake of simplicity, but performs better in practice. The reason is that it penalises small values of m . For instance, if $m = 2$ the formula $(x_i - x_1)/(x_m - x_1)$ yields adjusted values of 0 and 1, whereas the formula i/m yields $x_1^* = 0.5$ and $x_2^* = 1$, effectively doubling the value of the slope b , leading to larger value of ϑ^2 .

The observed y values are normalised by calculating $y_i^* = y_i/\bar{y}$, where \bar{y} denotes the mean of the y values.

Attention will now be given to procedures for identifying stable regions in the context of threshold selection in EVT. Consider again the second graph of Figure 1. To determine the most stable region in the set of estimates, a smoother plot than that of the right side of Figure 1 is required.

One method which works particularly well in practice is not to calculate the EVI estimates at all possible values of k (namely $1, 2, \dots, n - 1$), but rather only at $k = 5\%, 10\%, \dots, 95\%$ of n (rounded to the nearest integer). This leads to 19 estimates of the EVI, denoted in this context by y_1, \dots, y_{19} . The term *region* refers to a set of consecutive values of y . For instance, the set $\{y_3, y_4, \dots, y_9\}$ is a region of length 7.

The accuracy of EVI estimation is significantly improved by applying two procedures before applying methods of stable region selection. The first procedure rounds the values y_1, \dots, y_{19} to the nearest 5% of their mean to avoid insignificant fluctuations from having an effect when applying Algorithm 3.2 below. Specifically, the first procedure is the following:

Algorithm 3.1

factor := round($100 \times (0.05 \times \bar{y})/100$) Remark: 5% of mean rounded to two decimals

if factor = 0 then factor := 0.01

$\{y_1, \dots, y_{19}\} := \text{factor} \times \text{round}(\{y_1, \dots, y_{19}\}/\text{factor})$

The idea behind the second procedure is to remove completely from consideration the region of estimates where the bias becomes significant. This entails reducing the set of $\{y_1, \dots, y_{19}\}$ to $\{y_1, \dots, y_m\}$, where $m < 19$ if the original set contains non-decreasing values from some index onwards. Specifically, the algorithm is the following:

Algorithm 3.2

```

 $y_{20} := 2 \times y_{19}$  Remark: Definition of  $y_{20}$  required in final while loop
 $m := 19$ 
while ( $y_m \geq y_{m-1}$ ) and ( $y_{m-1} \geq y_{m-2}$ ) and ( $m > 2$ )
     $m := m - 1$ 
end
while ( $y_m = y_{m-1}$ ) and ( $m < 19$ ) Remark: One increase is retained
     $m := m + 1$ 
end
while ( $y_m = y_{m+1}$ ) and ( $m < 19$ )
     $m := m + 1$ 
end

```

An example of the application of Algorithm 3.2 will be given below. The methods for selecting an optimal (stable) region follow.

Method 0 only applies algorithms 3.1 and 3.2, regarding the remaining m values $\{y_1, \dots, y_m\}$ as the optimal region. No use is made of the instability measure.

For the sake of simplicity it will be assumed in the explanations below that m is still 19 after the application of Algorithm 3.2.

Method 1 entails fixing the region length. For a region length of 12, for example, ϑ^2 is calculated for regions $\{y_1, \dots, y_{12}\}$, $\{y_2, \dots, y_{13}\}$, up to $\{y_8, \dots, y_{19}\}$. The region for which ϑ^2 is a minimum, is regarded as the optimal region. If Algorithm 3.2 yields a value of m which is smaller than the region length, the region length is set to m . In other words, if the set of estimates which remains after applying Algorithm 3.2 has fewer elements than the specified region length, the stable region is simply taken as $\{y_1, \dots, y_m\}$.

Method 2 involves trimming the region. First ϑ^2 is calculated over the entire region $\{y_1, \dots, y_{19}\}$. Then either the first or the last value in the region is omitted, depending on which reduction of the region decreases ϑ^2 by the largest amount. This process is repeated until neither top nor bottom trimming of the region decreases the instability ϑ^2 .

Method 3 fixes the upper limit of the region and trims the region from the left. For an upper limit of 6, for example, ϑ^2 is calculated for $\{y_1, \dots, y_6\}$, $\{y_2, \dots, y_6\}$, up to $\{y_5, y_6\}$, and the region yielding the lowest ϑ^2 is regarded as the optimal region.

Simulation studies were carried out on all three methods (and for all possible parameters of each), across all sample sizes and distributions. The results indicated that algorithms 3.1 and 3.2 should always be applied (as mentioned earlier), and that two methods outperform the others on all counts, namely methods 0 and 3.

Simulation studies were carried out to obtain rules to decide between the methods, and to determine what the choice of upper limit should be in the case of Method 3. These and other issues are addressed in Section 6.

Once the stable region has been identified, the estimate of the EVI is taken as the mean of the *original* EVI estimates over the stable region, i.e. before applying the rounding of Algorithm 3.1.

As an example, consider again the Burr sample used earlier.




Figure 2: Stable region and EVI estimate for the Burr sample.

Figure 2 is the same as the right side of Figure 1, except that the estimates are calculated at larger intervals of k , namely $k = 10, 20, \dots, 150$. The plot therefore shows the values y_1, \dots, y_{15} . The vertical lines in Figure 2 indicate the optimal region which was found to be $k = 30, \dots, 60$. The dashed line indicates the final EVI estimate of 0.4727.

After rounding the values by using Algorithm 3.1, Algorithm 3.2 was applied. The first *while* loop in the algorithm yielded $m = 8$ (corresponding to $k = 80$). The second *while* loop yielded $m = 9$. Even though bias reduced estimation is applied, there is often a point from where the estimates increase steadily or, more specifically, are non-decreasing as a function of k . This point is at $k = 100$. The values corresponding to $k = 100, 110, \dots, 190$ are excluded before applying stable region methods 1, 2 or 3. The increase from $k = 80$ to $k = 90$ is regarded as variation, and not taken as the start of the region where the bias is regarded as significant.

Method 3 with an upper limit of 6 was applied to the remaining region $\{y_1, \dots, y_9\}$, yielding $\{y_3, \dots, y_6\}$, corresponding to $k = 30, \dots, 60$.

From this point onwards the original values y_3, \dots, y_6 (without the rounding of Algorithm 3.1) are used. The estimated EVI for this sample is calculated as the mean of y_3, \dots, y_6 , which is 0.4727.

The EVI in the optimal region which is closest to the final estimate of the EVI indicates the choice of threshold. In the example above, the single estimate in the set $\{y_3, \dots, y_6\}$ closest to 0.4727 is $y_5 = 0.4731$. This indicates a threshold of $k = 50$. The term *implied threshold* will be used to refer to the threshold selected in this manner. The resulting 50 excesses constitute the observations which will be used to construct confidence intervals, etc. In the concluding remarks more will be said of how to proceed from this point onwards, in the context of a specific financial application.

4. External estimation of the second order parameter

Attention will now be given to the exact method of obtaining the set of EVI estimates at a specified range of values of k . In Section 2 the need for external estimation of ρ was pointed out. In this section six estimators of ρ will be investigated as candidates for the external estimation of ρ when fitting the PPD to relative excesses using MLE. The first four estimators are those considered by Gomes and Martins (2002), all of which are easy to implement, both in terms of programming and computation time. The fifth estimator is one proposed by Gomes and Martins (2001). This estimator is chosen since it lends itself to the application of the instability measure, which yields an adjusted version of the estimator.

The first two estimators are

$$\hat{\rho}_1 := - \left| \log \left| \frac{1/M_n^{(1)}([n^{0.9}]) - 1/M_n^{(1)}([n^{0.5}])}{1/M_n^{(1)}([n^{0.95}]) - 1/M_n^{(1)}([n^{0.5}])} \right| / \log \frac{[n^{0.9}]}{[n^{0.95}]} \right|$$

by Hall and Welsh (1985), and

$$\hat{\rho}_2 := - \frac{1}{\log 2} \left| \log \left| \frac{M_n^{(2)}\left(\left[\frac{n}{2\log n}\right]\right) - 2\left(M_n^{(1)}\left(\left[\frac{n}{2\log n}\right]\right)\right)^2}{M_n^{(2)}\left(\left[\frac{n}{\log n}\right]\right) - 2\left(M_n^{(1)}\left(\left[\frac{n}{\log n}\right]\right)\right)^2} \right| \right|$$

by Peng (1998), where

$$M_n^{(j)}(k) := \frac{1}{k} \sum_{i=1}^k \left(\log \frac{X_{n-i+1,n}}{X_{n-k,n}} \right)^j, j \geq 1.$$

Note that $M_n^{(1)}(k)$ is the Hill estimator. $M_n^{(2)}(k)$ was introduced by Dekkers, Einmahl and de Haan (1989).

The following two estimators of ρ are special cases of a class of estimators proposed by Fraga Alves, Gomes and de Haan (2003). They are

$$\hat{\rho}_3 := - \left| \frac{3\left(T_n^{(0)}(k_1) - 1\right)}{T_n^{(0)}(k_1) - 3} \right|$$

and

$$\hat{\rho}_4 := - \left| \frac{3\left(T_n^{(1)}(k_1) - 1\right)}{T_n^{(1)}(k_1) - 3} \right|,$$

where $k_1 = \min(n-1, \lfloor 2n/\log \log n \rfloor)$, and

$$T_n^{(\tau)}(k) := \begin{cases} \frac{\left(M_n^{(1)}(k)\right)^\tau - \left(M_n^{(1)}(k)/2\right)^{\tau/2}}{\left(M_n^{(2)}(k)/2\right)^{\tau/2} - \left(M_n^{(3)}(k)/6\right)^{\tau/3}} & \text{if } \tau > 0 \\ \frac{\log\left(M_n^{(1)}(k)\right) - \frac{1}{2}\log\left(M_n^{(2)}(k)/2\right)}{\frac{1}{2}\log\left(M_n^{(2)}(k)/2\right) - \frac{1}{3}\log\left(M_n^{(3)}(k)/6\right)} & \text{if } \tau = 0 \end{cases}.$$

The estimator proposed by Gomes and Martins (2001) will be denoted by $\hat{\rho}_5$. The instability measure will be incorporated, leading to an adjusted estimator $\hat{\rho}_5^*$. Only the procedure to calculate $\hat{\rho}_5$ will be

stated here, together with some computational issues. For a full discussion of the estimator, refer to Gomes and Martins (2001).

Given a set of n positive, ordered observations $x_{1,n} \leq x_{2,n} \leq \dots \leq x_{n,n}$ assumed to be from a second order regularly varying distribution, define the following functions:

- For given choices of α and k , let $\hat{\gamma}_n^{(\alpha)}(k) = \frac{m_n^{(\alpha)}(k)}{\Gamma(\alpha+1)(H_{k,n})^{\alpha-1}}$, where $H_{k,n}$ denotes the Hill estimator, based on k excesses.
- For given choices of α , i_n and j_n , let $\chi(\alpha, i_n, j_n)$ be the sample median of $\hat{\gamma}_n^{(\alpha)}(k)$ for $i_n \leq k \leq j_n$.
- For given choices of α , i_n and j_n , let $\Sigma(\alpha, i_n, j_n) = \sum_{k=i_n}^{j_n} \left(\hat{\gamma}_n^{(\alpha)}(k) - \chi(\alpha, i_n, j_n) \right)^2$.

The procedure for estimating ρ is as follows:

1. Choose i_n and j_n which specify a range of k such that the estimates $\hat{\gamma}_n^{(\alpha)}(k)$ are “stable” for $i_n \leq k \leq j_n$.
2. Obtain an estimate of α as $\hat{\alpha} = \arg \min_{\alpha} \Sigma(\alpha, i_n, j_n)$.
3. Estimate ρ by $\hat{\rho}$ as the solution of $(1 - \hat{\rho})^{\hat{\alpha}-1} (1 + \hat{\rho}(\hat{\alpha} - 2)) = 1$. The restriction $\hat{\alpha} \geq 2$ applies. In particular, the value of $\hat{\alpha} = 2.09272$ corresponds to $\hat{\rho} = -10$, and $\hat{\alpha} = 15.02746$ to $\hat{\rho} = -0.01$. Computationally, let $\hat{\rho} = -10$, if $2 \leq \hat{\alpha} \leq 2.09272$, and $\hat{\rho} = -0.01$, if $\hat{\alpha} \geq 15.02746$. For $2.09272 < \hat{\alpha} < 15.02746$, $\hat{\rho}$ is solved numerically.

Simulations showed that $\Sigma(\alpha, i_n, j_n)$ as a function of α will typically have one of three shapes. The first shape (not shown) sees $\Sigma(\alpha, i_n, j_n)$ decreasing as α increases. In such a case take $\hat{\alpha} = 15$, yielding $\hat{\rho} = -0.01$.

The second and third typical shapes are shown in the following graphs:




Figure 3: Typical behaviours of $\Sigma(\alpha, i_n, j_n)$ as a function of α .

In the case of the first graph of Figure 3 the value of $\Sigma(\alpha, i_n, j_n)$ at $\alpha = 2$ is not regarded as a local minimum. Take $\hat{\alpha} = 15$, yielding $\hat{\rho} = -0.01$. The second graph of Figure 3 shows the situation where the local minimum of $\Sigma(\alpha, i_n, j_n)$ needs to be determined numerically. Note that $\hat{\gamma}_n^{(\alpha)}(k) \rightarrow 0$ (and consequently $\Sigma(\alpha, i_n, j_n) \rightarrow 0$) as $\alpha \rightarrow \infty$, and that $\hat{\rho} = 0$ is always a (trivial) solution of $(1 - \hat{\rho})^{\hat{\alpha}-1} (1 + \hat{\rho}(\hat{\alpha} - 2)) = 1$.

Gomes and Martins (2001) chose $i_n = 500$ and $j_n = 900$ for a Fréchet sample of size $n = 1000$ in their simulation study. Simulations performed in this study confirm that $i_n = 50\%$ of n and $j_n = 90\%$ of n is a good choice, robust over a wide variety of sample sizes and distributions as far as performance is concerned. The estimator $\hat{\rho}_5$ will imply this choice of i_n and j_n .

The adjusted estimator $\hat{\rho}_5^*$ chooses adaptively between three possible ranges which also covers 40% of the data, namely

1. $i_n = 10\%$ of n and $j_n = 50\%$ of n ,
2. $i_n = 30\%$ of n and $j_n = 70\%$ of n , and
3. $i_n = 50\%$ of n and $j_n = 90\%$ of n .

The procedure for the adjusted estimator is the following:

1. Determine $\hat{\alpha}_1, \hat{\alpha}_2$ and $\hat{\alpha}_3$, corresponding to ranges 1, 2 and 3 above, respectively.
2. Determine three sets of estimates: a set of $\hat{\gamma}_n^{(\hat{\alpha}_1)}(k)$ values, a set of $\hat{\gamma}_n^{(\hat{\alpha}_2)}(k)$ values, and a set of $\hat{\gamma}_n^{(\hat{\alpha}_3)}(k)$ values, where k assumes all integer values in ranges 1, 2 and 3, respectively.
3. Calculate the instability measure for each of the three sets of estimates.
4. Let $\hat{\alpha}$ denote the estimated value of α which corresponds to the range associated with the lowest instability measure.
5. Solve for $\hat{\rho}_5^*$ from $(1 - \hat{\rho}_5^*)^{\hat{\alpha}-1} (1 + \hat{\rho}_5^*(\hat{\alpha} - 2)) = 1$.

The performance of these six estimators will be evaluated in the next section.

5. Simulation study

In this section the aim is to provide an objective and comprehensive study design, which can be used for all simulation studies concerning positive EVI estimation.

The families of distributions should be those for which the expressions for their respective EVIs and second order parameters are known. The simulation study did not include observations simulated from the limiting distributions, for example the Pareto distribution, the generalised Pareto distribution, the PPD and the generalised extreme value distribution.

A reasonable range for the EVI is from 0 to 1. The choice of the lower bound of ρ is not straightforward. Seventeen papers focussing on the estimation of ρ were reviewed, of which only seven show estimates of ρ for real (case study) data. Fraga Alves (2002) calculated estimates of ρ of $-1.9973, -0.4989$ and -1.0609 for fire insurance claim data. Fraga Alves (2002) also shows graphs

of various estimators over a wide range of values of k . From those graphs it seems as though -2 is the lower limit for these estimates for that specific data set. Goegebeur and de Wet (2012) considered fire insurance claim data, and arrived at a median estimate of ρ of approximately -1.2 . Gomes, Caeiro and Figueiredo (2004), estimated the value of ρ as -0.69 for exchange rate data. Gomes, Martins and Neves (2007) obtained an estimate of -0.65 , also for exchange rate data. Gomes, Henriques-Rodrigues, Pereira and Pestana (2010) considered daily log-returns of indices and obtained $\hat{\rho} = -0.72$ for the Dow Jones Industrial, and $\hat{\rho} = -0.73$ for the NASDAQ. Gomes et al. (2010) studied males diagnosed with cancer of the tongue. Their data set yielded $\hat{\rho} = -0.654$. Gomes et al. (2012) examined a data set consisting of number of hectares burnt during wildfires, and obtained an estimate $\hat{\rho} = -0.39$.

For all the cases mentioned above, the estimates exceeded -2 . This does not mean that it is impossible to have $\rho < -2$ for real data, but -2 seems to be a reasonable lower bound. As far as the upper limit is concerned, $\rho = 0$ should be included, since for some heavy-tailed distributions $\rho = 0$.

The resulting set of distributions is as follows:

- The Burr distribution with distribution function $F(x) = 1 - \left(\frac{\beta}{\beta + x^\tau}\right)^\lambda$, defined for $x > 0$, and with β , τ and λ positive. This Burr type XII distribution is also called the Singh-Maddala distribution (Singh and Maddala, 1976). Since the EVI $\gamma = 1/\lambda\tau$ and the second order parameter $\rho = -1/\lambda$, the Burr can be reparameterised in terms of β , γ and ρ . The Burr distributions considered are the $Burr(\beta = 1, \gamma = 0.25, \rho = -2)$, the $Burr(\beta = 1, \gamma = 0.5, \rho = -2)$, the $Burr(\beta = 1, \gamma = 1, \rho = -2)$, the $Burr(\beta = 1, \gamma = 0.25, \rho = -0.5)$, the $Burr(\beta = 1, \gamma = 0.5, \rho = -0.5)$ and the $Burr(\beta = 1, \gamma = 1, \rho = -0.5)$.
- The Fréchet distribution with distribution function $F(x) = \exp(-x^{-\alpha})$, defined for $x > 0$, and with $\alpha > 0$. For the Fréchet $\gamma = 1/\alpha$ and $\rho = -1$. Three Fréchet distributions are considered in this study, namely those with $\alpha = 1$, $\alpha = 2$ and $\alpha = 4$, respectively.
- The student t distribution. Since a variable following the t distribution can assume negative values, the distribution is inflected on the positive half-line. The resulting distribution function is $F(x) = \frac{2\Gamma(\frac{n+1}{2})}{\sqrt{n\pi}\Gamma(\frac{n}{2})} \int_0^x \left(1 + \frac{w^2}{n}\right)^{-\frac{(n+1)}{2}} dw$, defined for $x \geq 0$, and with $n > 0$. For the $|t|$ distribution $\gamma = 1/n$ and $\rho = -2/n$. The t distributions considered in this study are the $|t_1|$, $|t_2|$ and $|t_4|$ distributions.
- The loggamma distribution with distribution function $F(x) = \frac{\lambda^\alpha}{\Gamma(\alpha)} \int_1^x w^{-\lambda-1} (\log w)^{\alpha-1} dw$, defined for $x > 1$, and with λ and α positive. For the loggamma distribution $\gamma = 1/\lambda$ and $\rho = 0$. Since $\alpha = 1$ yields the Pareto distribution, we take $\alpha = 2$. Three loggamma distributions will be considered, namely $\log\Gamma(\lambda = 1, \alpha = 2)$, $\log\Gamma(\lambda = 2, \alpha = 2)$ and $\log\Gamma(\lambda = 4, \alpha = 2)$.

The above yields fifteen distributions from four families: five distributions with $\gamma = 0.25$, five with $\gamma = 0.5$, five with $\gamma = 1$, four with $\rho = -2$, four with $\rho = -1$, four with $\rho = -0.5$, and three with $\rho = 0$.

Sample sizes of $n = 100, 200, 500, 1000, 2000$ and 5000 will be considered. As measure of error, the mean square error (MSE) will be used. In the tables which show results pertaining to estimators of the EVI, the values of the MSEs will be multiplied by a factor of 1000 to ease comparison.

All results shown in the following sections were based on 1000 samples from each combination of distribution and sample size. The standard errors of the MSEs were obtained by calculating 10000 bootstrap repetitions of the MSE.

The simulation study design described above was applied to the estimators of the second order parameter described in the previous section. Consistent with the results of Gomes and Martins (2002) it was found that (1) $\hat{\rho}_1$ and $\hat{\rho}_2$ do not perform well, (2) $\hat{\rho}_4$ generally performs better than $\hat{\rho}_3$ for small values of ρ , and (3) $\hat{\rho}_3$ has the best performance overall (of the first four estimators).

The estimators $\hat{\rho}_5$ and $\hat{\rho}_5^*$ yield in general more accurate estimates than $\hat{\rho}_3$ for large values of ρ , and outperform $\hat{\rho}_3$ overall for $n > 100$. The best overall performer for $200 \leq n \leq 2000$ is $\hat{\rho}_5$, and $\hat{\rho}_5^*$ for $n = 5000$. More extensive simulations showed $\hat{\rho}_5^*$ to be the best overall performer for sample sizes $n = 10000$, $n = 20000$ and $n = 50000$ as well.

Based on these results, only the estimators $\hat{\rho}_3$, $\hat{\rho}_5$ and $\hat{\rho}_5^*$ will be considered as candidates for the external estimation of ρ .

6. The PPD maximum likelihood estimators

The simulation results referred to in this section were obtained by estimating the EVI by fitting the PPD to relative excesses using MLE, after externally estimating the second order parameter by using $\hat{\rho}_3$, $\hat{\rho}_5$ or $\hat{\rho}_5^*$. The final result of this section is the construction of two estimators: one by applying $\hat{\rho}_5$ for EVI estimation and $\hat{\rho}_5^*$ for threshold selection, and one by applying $\hat{\rho}_3$ for both EVI estimation and threshold selection. The second estimator is included to provide an alternative estimator which is far easier to implement, and still yields good results.

The simulation study design of Section 5 was applied. For each of the estimators $\hat{\rho}_3$, $\hat{\rho}_5$ or $\hat{\rho}_5^*$ the EVI was estimated at 19 values of k (5%, 10%, ..., 95% of n) for 1000 samples from each combination of sample size and distribution.

All possible combinations of threshold selection techniques were applied to the resulting estimates: Method 1 with region lengths 2 to 19, Method 2, Method 3 with upper limits 2 to 19, each with and without applying Algorithm 3.2. Note that Method 1 with region length 19 yields Method 0.

It was found that $\hat{\rho}_5$ consistently delivers the lowest MSE for the estimated EVI, even for sample sizes $n = 100$ and $n = 5000$. This implies that even though an estimator of ρ yields the lowest MSE in terms of estimation of ρ , it does not necessarily perform the best in terms of estimation of the EVI when used as external estimator of ρ .

It was found that all methods across all combinations of sample size and distribution benefited from the application of Algorithm 3.2. Furthermore, Method 3 consistently yielded the best results, except for a couple of cases where Method 0 performed better.

The question of how to choose the optimal technique for a given sample size will now be addressed. The term *technique* denotes one of the 19 possibilities: Method 0 and Method 3 with each of the upper limits 2, ..., 19. Initially it was assumed that the choice of technique will depend on the family of the underlying distribution, and specifically the respective EVI. The results showed however that the optimal technique is for all practical purposes independent of the family of the underlying distribution and the underlying EVI, and only a function of the underlying value of ρ .

This information is used to develop an estimator which adaptively chooses the appropriate technique. The idea is to determine $\hat{\rho}$, and if $\hat{\rho} < -1.5$ to use the optimal technique for distributions for which the underlying value of ρ is -2 . If $-1.5 \leq \hat{\rho} < -0.75$ use the optimal technique for $\rho = -1$. If $-0.75 \leq \hat{\rho} < -0.25$ use the optimal technique for $\rho = -0.5$. If $\hat{\rho} > -0.25$ use the optimal technique for $\rho = 0$.

Another surprising result was that, even though $\hat{\rho}_5$ yields the lowest MSEs for EVI estimation, $\hat{\rho}_5^*$ consistently performs better than $\hat{\rho}_5$ when choosing between the four categories mentioned in the previous paragraph. The “best” estimator is achieved by using the value of $\hat{\rho}_5^*$ to decide which technique to apply, and then applying the technique to the EVI estimates with $\hat{\rho}_5$ as external estimator. The results in the table which follow are also shown for the estimator resulting from using $\hat{\rho}_3$ to choose between the four categories, and applying the appropriate technique to the EVI estimates with $\hat{\rho}_3$ as external estimator.

Table 1: Optimal techniques for threshold selection.

ρ	\mathbf{n}	$\hat{\rho}_5, \hat{\rho}_5^*$	$\hat{\rho}_3$	ρ	\mathbf{n}	$\hat{\rho}_5, \hat{\rho}_5^*$	$\hat{\rho}_3$
-2	100	19	19	-0.5	100	5	5
-2	200	19	19	-0.5	200	5	5
-2	500	17	17	-0.5	500	5	5
-2	1000	17	17	-0.5	1000	5	5
-2	2000	15	15	-0.5	2000	4	4
-2	5000	14	14	-0.5	5000	<u>3</u>	<u>2</u>
-1	100	19	19	0	100	0	0
-1	200	18	18	0	200	0	0
-1	500	16	16	0	500	<u>0</u>	<u>9</u>
-1	1000	14	14	0	1000	9	9
-1	2000	12	12	0	2000	6	6
-1	5000	10	10	0	5000	3	3

The numbers in the $\hat{\rho}_5, \hat{\rho}_5^*$ and $\hat{\rho}_3$ columns refer to the upper limit for Method 3. A zero indicates Method 0. As an example, suppose $n = 5000$ and $\hat{\rho}_5^* < -1.5$. The appropriate technique is therefore to be found in the section of the table where $\rho = -2$. Method 3 with an upper limit of 14 (corresponding to $k = 70\%$ of n) is applied to the EVI estimates obtained with $\hat{\rho}_5$.

The upper limits were chosen in a way that it is as robust as possible with respect to the choice of external estimator, in this case only $\hat{\rho}_3$ and $\hat{\rho}_5$. No group MSE (average MSE for the group of distributions with the same value of ρ) obtained using the table differed by more than two standard errors from the actual observed minimum MSE. The two exceptions are underlined in Table 1. The first exception is for $\rho = -0.5$ and $n = 5000$. For such a large sample size one must bear in mind that there is a huge difference between $k = 10\%$ and $k = 15\%$ of n . The second exception sees $\hat{\rho}_3$ making a quicker transition than $\hat{\rho}_5$ from Method 0 to Method 3 (with an upper limit of 9) as the sample size increases.

In summary, the procedure for the proposed estimator, which will be denoted by $\hat{\gamma}_5$, is as follows:

1. Calculate $\hat{\rho}_5$ and $\hat{\rho}_5^*$ as described in Section 4.
2. Calculate 19 estimates of the EVI (at $k = 5\%, 10\%, \dots, 95\%$ of n) by fitting the PPD to the relative excesses using MLE, keeping ρ fixed at $\hat{\rho}_5$ and maximising the log likelihood over $\gamma > 0$ (EVI) and $1/\hat{\rho}_5 \leq c \leq 0.5$.
3. Choose the category for ρ . For $\hat{\rho}_5^*$ in $(-\infty, -1.5)$, choose $\rho = -2$, in $[-1.5, -0.75)$, choose $\rho = -1$, in $[-0.75, -0.25)$, choose $\rho = -0.5$, and in $[-0.25, 0)$, choose $\rho = 0$.
4. Use the $\hat{\rho}_5, \hat{\rho}_5^*$ column of Table 1 to determine which technique to use.
5. Apply algorithms 3.1 and 3.2 to the 19 estimates, followed by the technique resulting from Step 4 to determine the optimal region.
6. The estimate is calculated as the mean of the original estimates (not rounded) in the optimal region.

The estimator $\hat{\gamma}_3$ is defined in the same way, except that both $\hat{\rho}_5$ and $\hat{\rho}_5^*$ are replaced by $\hat{\rho}_3$. This estimator does not perform as well as $\hat{\gamma}_5$, but is considerably easier to program and computationally much less intensive.

Below are tables containing simulation results comparing the performance of EVI estimators in terms of MSE.

The effect of adaptive threshold selection was also investigated for the Hill estimator. The three threshold selection methods considered were those mentioned in the introduction, namely the method by Drees and Kaufmann (1998), the method by Guillou and Hall (2001), and the method by Beirlant et al. (2004). The method by Drees and Kaufmann (1998), yielding the estimator denoted by $\hat{\gamma}_{DK}^{Hill}$, performed best overall and is the only method for which the results are shown. For a discussion of these three adaptive threshold selection methods, please refer to the respective papers.

The alternative to adaptive threshold selection is to simply fix k at some value prior to estimation. In this study the value of k was fixed at $k_0 = 2\sqrt{n}$ for the Hill estimator. This is the choice made by Drees and Kaufmann (1998) to obtain their initial estimate $\hat{\gamma}_0^{Hill} = H_{2\sqrt{n}, n}$. The performance of $\hat{\gamma}_0^{Hill}$ is compared to that of $\hat{\gamma}_{DK}^{Hill}$.

The value of $k_0 = 2\sqrt{n}$ performs well as a default fixed choice of k for the Hill estimator, as can be seen from the results below. A similar formula might be useful when fixing the value of k when fitting the PPD. In the case of the PPD k/n should tend to zero slower than in the first order case. A reasonable choice would be $k_0 = 2n^{2/3}$. It turns out that this choice closely models that which was observed in the simulation studies. The estimator $\hat{\gamma}_0$ is defined as the estimate of the MLE of the EVI obtained by fitting the PPD to $k_0 = 2n^{2/3}$ relative excesses, with $\hat{\rho}_5$ as external estimator of the second order parameter.

The simulation study results are shown in the tables below. In each row the minimum value of the MSE is underlined.

Table 2: 1000× MSEs for estimates of the EVI for sample size $n = 100$.

Dist	γ	ρ	$\hat{\gamma}_0^{Hill}$	$\hat{\gamma}_{DK}^{Hill}$	$\hat{\gamma}_0$	$\hat{\gamma}_3$	$\hat{\gamma}_5$
Burr	0.25	-2	3.086 (0.144)	<u>2.626</u> (0.111)	3.253 (0.124)	2.893 (0.150)	3.669 (0.157)
Burr	0.25	-0.5	19.493 (0.647)	19.356 (0.744)	10.039 (0.469)	14.340 (0.614)	<u>8.349</u> (0.437)
Fréchet (4)	0.25	-1	3.299 (0.168)	3.376 (0.356)	3.196 (0.127)	<u>2.787</u> (0.158)	3.746 (0.165)
$ t_4 $	0.25	-0.5	36.255 (0.977)	30.310 (1.053)	20.352 (0.589)	15.796 (0.769)	<u>10.818</u> (0.549)
$\log\Gamma(4, 2)$	0.25	0	8.890 (0.432)	10.348 (1.342)	4.726 (0.210)	6.129 (0.282)	<u>4.319</u> (0.176)
Mean	0.25		14.205 (0.254)	13.203 (0.380)	8.313 (0.160)	8.389 (0.209)	6.180 (0.152)
Burr	0.5	-2	11.572 (0.560)	<u>9.455</u> (0.465)	12.135 (0.473)	10.663 (0.733)	13.973 (0.614)
Burr	0.5	-0.5	73.088 (2.434)	74.419 (2.647)	36.679 (1.604)	54.358 (2.316)	<u>31.727</u> (1.523)
Fréchet (2)	0.5	-1	14.676 (0.810)	<u>13.405</u> (0.622)	14.912 (0.534)	13.479 (0.692)	16.562 (0.656)
$ t_2 $	0.5	-1	24.513 (1.108)	27.493 (1.163)	<u>12.968</u> (0.768)	24.968 (1.106)	25.179 (1.147)
$\log\Gamma(2, 2)$	0.5	0	36.270 (1.589)	38.995 (2.234)	21.113 (0.906)	27.481 (1.255)	<u>19.343</u> (0.779)
Mean	0.5		32.024 (0.653)	32.754 (0.747)	19.562 (0.424)	26.190 (0.606)	21.357 (0.449)
Burr	1	-2	46.160 (2.220)	<u>40.643</u> (2.805)	50.740 (1.936)	43.836 (2.304)	58.115 (2.485)
Burr	1	-0.5	296.088 (10.427)	293.840 (10.573)	138.417 (5.762)	192.518 (7.391)	<u>118.216</u> (5.374)
Fréchet (1)	1	-1	58.165 (3.393)	65.676 (6.556)	51.091 (1.972)	<u>49.477</u> (2.589)	59.510 (2.497)
$ t_1 $	1	-2	53.228 (2.447)	<u>47.448</u> (2.142)	57.310 (2.240)	54.341 (2.690)	72.719 (2.858)
$\log\Gamma(1, 2)$	1	0	138.476 (6.442)	161.175 (9.671)	79.031 (3.318)	97.364 (4.003)	<u>67.987</u> (2.696)
Mean	1		118.423 (2.628)	121.756 (3.230)	75.318 (1.508)	87.507 (1.896)	75.309 (1.506)
Overall mean			54.884 (0.907)	55.904 (1.112)	34.397 (0.525)	40.695 (0.667)	34.282 (0.526)

Table 3: 1000× MSEs for estimates of the EVI for sample size $n = 200$.

Dist	γ	ρ	$\hat{\gamma}_0^{Hill}$	$\hat{\gamma}_{DK}^{Hill}$	$\hat{\gamma}_0$	$\hat{\gamma}_3$	$\hat{\gamma}_5$
Burr	0.25	-2	2.339 (0.111)	<u>1.429</u> (0.064)	2.212 (0.092)	1.530 (0.083)	1.973 (0.101)
Burr	0.25	-0.5	11.028 (0.394)	12.476 (0.419)	<u>5.763</u> (0.274)	8.909 (0.358)	5.906 (0.293)
Fréchet (4)	0.25	-1	2.183 (0.101)	1.739 (0.080)	2.374 (0.099)	<u>1.560</u> (0.086)	2.367 (0.105)
$ t_4 $	0.25	-0.5	20.427 (0.570)	20.038 (0.630)	10.414 (0.315)	9.918 (0.414)	<u>7.008</u> (0.291)
$\log\Gamma(4, 2)$	0.25	0	6.297 (0.268)	7.190 (0.549)	4.375 (0.168)	5.199 (0.180)	<u>3.683</u> (0.128)
Mean	0.25		8.455 (0.152)	8.574 (0.188)	5.028 (0.094)	5.423 (0.118)	4.187 (0.091)
Burr	0.5	-2	8.745 (0.409)	<u>5.270</u> (0.247)	8.617 (0.372)	6.072 (0.327)	8.549 (0.409)
Burr	0.5	-0.5	46.050 (1.671)	50.958 (1.703)	23.165 (1.072)	35.136 (1.390)	<u>22.839</u> (1.071)
Fréchet (2)	0.5	-1	10.503 (0.509)	8.011 (0.383)	9.785 (0.383)	<u>7.502</u> (0.384)	10.075 (0.440)
$ t_2 $	0.5	-1	13.974 (0.708)	15.850 (0.714)	<u>9.710</u> (0.738)	15.646 (0.740)	15.575 (0.738)
$\log\Gamma(2, 2)$	0.5	0	26.677 (1.093)	27.398 (0.822)	18.727 (0.716)	21.408 (0.727)	<u>16.279</u> (0.604)
Mean	0.5		21.190 (0.443)	21.497 (0.414)	14.001 (0.316)	17.153 (0.361)	14.663 (0.311)
Burr	1	-2	36.532 (1.812)	<u>22.108</u> (1.023)	33.994 (1.394)	24.887 (1.355)	33.816 (1.619)
Burr	1	-0.5	181.968 (6.034)	202.763 (7.185)	93.315 (3.992)	143.099 (5.736)	<u>87.062</u> (3.972)
Fréchet (1)	1	-1	39.442 (2.016)	32.050 (1.368)	39.533 (1.588)	<u>28.668</u> (1.398)	39.296 (1.759)
$ t_1 $	1	-2	34.808 (2.058)	<u>26.771</u> (1.289)	42.175 (1.909)	30.064 (1.634)	44.497 (2.002)
$\log\Gamma(1, 2)$	1	0	98.580 (4.368)	107.525 (3.606)	69.386 (2.628)	79.382 (2.731)	<u>59.001</u> (2.130)
Mean	1		78.266 (1.638)	78.243 (1.664)	55.681 (1.113)	61.220 (1.368)	52.735 (1.096)
Overall mean			35.970 (0.568)	36.105 (0.575)	24.903 (0.387)	27.932 (0.473)	23.862 (0.381)

Table 4: 1000× MSEs for estimates of the EVI for sample size $n = 500$.

Dist	γ	ρ	$\hat{\gamma}_0^{Hill}$	$\hat{\gamma}_{DK}^{Hill}$	$\hat{\gamma}_0$	$\hat{\gamma}_3$	$\hat{\gamma}_5$
Burr	0.25	-2	1.355 (0.061)	<u>0.581</u> (0.026)	1.299 (0.057)	0.764 (0.045)	0.912 (0.053)
Burr	0.25	-0.5	5.926 (0.200)	7.746 (0.232)	<u>2.795</u> (0.137)	4.411 (0.176)	3.050 (0.136)
Fréchet (4)	0.25	-1	1.433 (0.071)	<u>0.854</u> (0.037)	1.522 (0.069)	0.872 (0.049)	1.232 (0.062)
$ t_4 $	0.25	-0.5	10.815 (0.309)	11.716 (0.357)	5.005 (0.203)	6.064 (0.250)	<u>3.918</u> (0.179)
$\log\Gamma(4,2)$	0.25	0	4.530 (0.191)	5.406 (0.134)	3.551 (0.141)	4.278 (0.116)	<u>3.077</u> (0.105)
Mean	0.25		4.812 (0.085)	5.260 (0.090)	2.835 (0.059)	3.278 (0.067)	2.438 (0.052)
Burr	0.5	-2	5.427 (0.260)	<u>2.455</u> (0.117)	5.490 (0.266)	3.160 (0.206)	3.949 (0.234)
Burr	0.5	-0.5	24.800 (0.878)	30.077 (0.969)	<u>12.066</u> (0.635)	18.161 (0.726)	12.854 (0.573)
Fréchet (2)	0.5	-1	6.002 (0.267)	<u>3.704</u> (0.149)	6.362 (0.267)	3.769 (0.199)	4.888 (0.234)
$ t_2 $	0.5	-1	7.675 (0.346)	8.255 (0.314)	<u>6.591</u> (0.428)	8.505 (0.432)	8.604 (0.405)
$\log\Gamma(2,2)$	0.5	0	18.454 (0.697)	22.668 (0.522)	15.466 (0.564)	17.604 (0.452)	<u>13.109</u> (0.409)
Mean	0.5		12.472 (0.246)	13.432 (0.232)	9.195 (0.205)	10.240 (0.200)	8.681 (0.176)
Burr	1	-2	23.276 (1.049)	<u>9.907</u> (0.533)	21.099 (0.901)	14.252 (0.835)	15.177 (0.896)
Burr	1	-0.5	103.445 (3.634)	116.008 (3.886)	49.345 (2.928)	74.587 (2.950)	<u>49.197</u> (2.309)
Fréchet (1)	1	-1	22.443 (0.992)	17.245 (0.760)	24.628 (1.159)	<u>15.934</u> (0.831)	20.271 (0.994)
$ t_1 $	1	-2	22.916 (1.111)	20.853 (9.036)	30.080 (1.775)	<u>15.524</u> (0.842)	20.996 (1.098)
$\log\Gamma(1,2)$	1	0	73.957 (2.972)	87.196 (2.149)	59.608 (2.194)	68.296 (1.818)	<u>51.872</u> (1.633)
Mean	1		49.207 (1.007)	50.242 (2.022)	36.952 (0.865)	37.719 (0.751)	31.503 (0.663)
Overall mean			22.164 (0.347)	22.978 (0.679)	16.327 (0.297)	17.079 (0.260)	14.207 (0.229)

Table 5: $1000 \times$ MSEs for estimates of the EVI for sample size $n = 1000$.

Dist	γ	ρ	$\hat{\gamma}_0^{Hill}$	$\hat{\gamma}_{DK}^{Hill}$	$\hat{\gamma}_0$	$\hat{\gamma}_3$	$\hat{\gamma}_5$
Burr	0.25	-2	1.009 (0.045)	<u>0.341</u> (0.015)	0.833 (0.034)	0.445 (0.024)	0.437 (0.023)
Burr	0.25	-0.5	4.176 (0.144)	5.690 (0.151)	<u>1.828</u> (0.092)	2.908 (0.111)	1.972 (0.084)
Fréchet (4)	0.25	-1	0.974 (0.039)	0.527 (0.022)	0.994 (0.045)	<u>0.510</u> (0.027)	0.687 (0.033)
$ t_4 $	0.25	-0.5	6.640 (0.194)	8.049 (0.217)	2.534 (0.095)	3.911 (0.139)	<u>2.248</u> (0.092)
$\log\Gamma(4, 2)$	0.25	0	3.200 (0.122)	4.717 (0.085)	2.886 (0.114)	3.476 (0.081)	<u>2.718</u> (0.084)
Mean	0.25		3.200 (0.055)	3.865 (0.056)	1.815 (0.037)	2.250 (0.040)	1.612 (0.031)
Burr	0.5	-2	3.909 (0.175)	<u>1.317</u> (0.062)	3.613 (0.191)	1.964 (0.117)	1.966 (0.114)
Burr	0.5	-0.5	16.367 (0.573)	21.214 (0.646)	<u>8.016</u> (0.497)	11.766 (0.483)	8.401 (0.417)
Fréchet (2)	0.5	-1	4.032 (0.187)	2.479 (0.100)	4.111 (0.179)	<u>2.137</u> (0.108)	2.848 (0.128)
$ t_2 $	0.5	-1	<u>4.948</u> (0.237)	5.330 (0.206)	5.454 (0.319)	5.028 (0.268)	5.093 (0.244)
$\log\Gamma(2, 2)$	0.5	0	14.281 (0.552)	19.911 (0.378)	11.880 (0.438)	14.414 (0.344)	<u>11.461</u> (0.347)
Mean	0.5		8.707 (0.174)	10.050 (0.157)	6.615 (0.156)	7.062 (0.134)	5.954 (0.124)
Burr	1	-2	14.867 (0.636)	<u>5.632</u> (0.251)	13.485 (0.560)	8.039 (0.437)	8.193 (0.479)
Burr	1	-0.5	60.773 (2.091)	74.966 (2.424)	<u>28.898</u> (1.532)	43.199 (1.700)	30.549 (1.342)
Fréchet (1)	1	-1	16.287 (0.752)	9.833 (0.378)	17.703 (0.836)	<u>9.077</u> (0.485)	12.514 (0.627)
$ t_1 $	1	-2	15.082 (0.705)	<u>5.984</u> (0.277)	20.733 (1.191)	8.443 (0.507)	9.886 (0.568)
$\log\Gamma(1, 2)$	1	0	54.459 (2.094)	78.858 (1.538)	48.269 (1.789)	57.814 (1.368)	<u>46.241</u> (1.420)
Mean	1		32.294 (0.639)	35.055 (0.584)	25.818 (0.565)	25.314 (0.467)	21.477 (0.436)
Overall mean			14.734 (0.222)	16.323 (0.202)	11.416 (0.196)	11.542 (0.162)	9.681 (0.152)

Table 6: $1000 \times$ MSEs for estimates of the EVI for sample size $n = 2000$.

Dist	γ	ρ	$\hat{\gamma}_0^{Hill}$	$\hat{\gamma}_{DK}^{Hill}$	$\hat{\gamma}_0$	$\hat{\gamma}_3$	$\hat{\gamma}_5$
Burr	0.25	-2	0.672 (0.028)	<u>0.187</u> (0.008)	0.543 (0.022)	0.250 (0.015)	0.251 (0.015)
Burr	0.25	-0.5	2.748 (0.094)	4.195 (0.102)	1.294 (0.087)	2.089 (0.075)	<u>1.231</u> (0.060)
Fréchet (4)	0.25	-1	0.679 (0.031)	<u>0.333</u> (0.013)	0.711 (0.035)	0.361 (0.019)	0.424 (0.022)
$ t_4 $	0.25	-0.5	4.596 (0.137)	5.877 (0.146)	1.476 (0.068)	2.999 (0.097)	<u>1.471</u> (0.064)
$\log\Gamma(4, 2)$	0.25	0	2.720 (0.092)	4.399 (0.062)	2.393 (0.088)	2.857 (0.061)	<u>2.361</u> (0.064)
Mean	0.25		2.283 (0.039)	2.998 (0.038)	1.283 (0.029)	1.711 (0.028)	1.148 (0.022)
Burr	0.5	-2	2.753 (0.125)	<u>0.771</u> (0.034)	2.364 (0.109)	1.025 (0.057)	1.134 (0.073)
Burr	0.5	-0.5	10.218 (0.353)	14.722 (0.426)	<u>4.498</u> (0.239)	7.940 (0.292)	4.770 (0.206)
Fréchet (2)	0.5	-1	3.110 (0.137)	<u>1.505</u> (0.056)	2.895 (0.148)	1.583 (0.081)	1.814 (0.095)
$ t_2 $	0.5	-1	3.432 (0.163)	3.309 (0.125)	4.507 (0.241)	<u>2.832</u> (0.131)	3.475 (0.161)
$\log\Gamma(2, 2)$	0.5	0	10.146 (0.371)	18.076 (0.282)	9.624 (0.368)	11.292 (0.262)	<u>9.239</u> (0.270)
Mean	0.5		5.932 (0.114)	7.677 (0.106)	4.778 (0.107)	4.934 (0.085)	4.087 (0.079)
Burr	1	-2	11.327 (0.511)	<u>3.277</u> (0.155)	9.841 (0.465)	4.821 (0.306)	5.127 (0.361)
Burr	1	-0.5	41.164 (1.459)	50.202 (1.606)	19.393 (1.229)	30.540 (1.158)	<u>19.035</u> (0.911)
Fréchet (1)	1	-1	11.635 (0.543)	6.621 (0.237)	11.756 (0.602)	<u>6.362</u> (0.361)	7.553 (0.418)
$ t_1 $	1	-2	12.535 (0.538)	<u>3.900</u> (0.201)	16.609 (0.917)	6.132 (0.374)	7.106 (0.432)
$\log\Gamma(1, 2)$	1	0	43.365 (1.596)	73.104 (1.144)	<u>38.197</u> (1.410)	47.759 (1.008)	38.903 (1.059)
Mean	1		24.005 (0.470)	27.421 (0.400)	19.159 (0.444)	19.123 (0.330)	15.545 (0.313)
Overall mean			10.740 (0.162)	12.699 (0.139)	8.407 (0.152)	8.589 (0.114)	6.926 (0.108)

Table 7: $1000 \times$ MSEs for estimates of the EVI for sample size $n = 5000$.

Dist	γ	ρ	$\hat{\gamma}_0^{Hill}$	$\hat{\gamma}_{DK}^{Hill}$	$\hat{\gamma}_0$	$\hat{\gamma}_3$	$\hat{\gamma}_5$
Burr	0.25	-2	0.460 (0.020)	<u>0.090</u> (0.004)	0.410 (0.025)	0.118 (0.009)	0.129 (0.010)
Burr	0.25	-0.5	1.670 (0.056)	2.808 (0.061)	0.692 (0.043)	1.168 (0.046)	<u>0.627</u> (0.030)
Fréchet (4)	0.25	-1	0.441 (0.020)	0.198 (0.008)	0.446 (0.022)	<u>0.182</u> (0.009)	0.205 (0.011)
$ t_4 $	0.25	-0.5	2.444 (0.070)	3.621 (0.083)	<u>0.611</u> (0.038)	1.532 (0.062)	0.671 (0.029)
$\log\Gamma(4, 2)$	0.25	0	2.260 (0.074)	3.862 (0.044)	2.047 (0.069)	2.337 (0.047)	<u>2.044</u> (0.048)
Mean	0.25		1.455 (0.024)	2.116 (0.023)	0.841 (0.019)	1.067 (0.018)	0.735 (0.013)
Burr	0.5	-2	1.709 (0.073)	<u>0.327</u> (0.015)	1.405 (0.067)	0.429 (0.027)	0.453 (0.029)
Burr	0.5	-0.5	5.938 (0.217)	8.981 (0.252)	2.895 (0.237)	4.236 (0.193)	<u>2.561</u> (0.161)
Fréchet (2)	0.5	-1	1.808 (0.082)	0.922 (0.033)	1.805 (0.097)	<u>0.790</u> (0.041)	0.866 (0.047)
$ t_2 $	0.5	-1	1.725 (0.079)	1.629 (0.064)	3.190 (0.147)	<u>1.512</u> (0.058)	2.002 (0.087)
$\log\Gamma(2, 2)$	0.5	0	7.999 (0.258)	15.302 (0.186)	<u>7.392</u> (0.248)	9.006 (0.188)	7.793 (0.189)
Mean	0.5		3.836 (0.073)	5.432 (0.064)	3.337 (0.078)	3.195 (0.056)	2.735 (0.054)
Burr	1	-2	7.560 (0.338)	<u>1.311</u> (0.058)	6.261 (0.261)	2.140 (0.154)	2.180 (0.159)
Burr	1	-0.5	24.769 (0.851)	31.584 (0.914)	10.124 (0.565)	16.946 (0.682)	<u>9.408</u> (0.446)
Fréchet (1)	1	-1	7.302 (0.352)	3.358 (0.123)	7.388 (0.413)	<u>3.184</u> (0.160)	3.695 (0.203)
$ t_1 $	1	-2	6.657 (0.292)	<u>1.632</u> (0.073)	9.254 (0.624)	2.341 (0.130)	2.582 (0.155)
$\log\Gamma(1, 2)$	1	0	33.432 (1.102)	59.177 (0.782)	<u>30.625</u> (1.044)	35.967 (0.719)	31.001 (0.744)
Mean	1		15.944 (0.301)	19.412 (0.243)	12.730 (0.285)	12.116 (0.205)	9.773 (0.184)
Overall mean			7.078 (0.103)	8.987 (0.084)	5.636 (0.099)	5.459 (0.071)	4.415 (0.064)

The following observations can be made:

1. All methods based on the fitting of the PPD perform significantly better than all methods based on the Hill estimator. This indicates that the choice of first or second order model outweighs all other concerns, at least for this limited number of estimators considered.
2. The threshold selection method by Drees and Kaufmann (1998) performs significantly worse than fixing the threshold at $2\sqrt{n}$, despite the fact that it is the best performing method of the three methods considered for the Hill estimator. The only exception is the distributions with an EVI of 0.25, for samples of size $n = 100$. This might be an indication that, if an estimator performs extremely poorly, the very underlying mechanics of an adaptive threshold selection method are undermined by the instability of the estimator. In such a case a reasonable choice of a fixed threshold may perform better on average.
3. The estimator $\hat{\gamma}_0$ outperforms $\hat{\gamma}_3$ consistently, except when the sample size and underlying EVI is large. In particular $\hat{\gamma}_3$ performed better than $\hat{\gamma}_0$ for $n = 1000$ and $n = 2000$ for the group with an EVI of 1, and for $n = 5000$ for groups with EVI 0.5 and 1. The fact that the estimator $\hat{\gamma}_0$ outperforms $\hat{\gamma}_3$ in general shows that threshold selection is of secondary importance relative to the choice of external estimator.
4. The estimator $\hat{\gamma}_5$ has the best average performance for each combination of EVI group and sample size. The only two exceptions are the groups with EVI 0.5, for sample sizes of both $n = 100$ and $n = 200$, where the estimator of choice was $\hat{\gamma}_0$.

Overall the results suggest that the most important contributing factors to the accuracy of estimation of the EVI are, in decreasing order of importance, the following:

1. Whether a first or second order model is fitted to the relative excesses.
2. The method of parameter estimation, for a given threshold. This refers to the choice of external estimator of the second order parameter, as well as to limits on the ranges of the parameters.
3. The method of threshold selection.

Concluding remarks

Frequently a situation is encountered where an estimator of the parameter of interest depends on the choice of some nuisance parameter. More often than not suboptimal ranges of the nuisance parameter exist, within which the estimates of the parameter of interest are extremely sensitive to the slightest change in the value of the nuisance parameter. Failure to avoid these problem ranges may have severe adverse effects on estimation accuracy. This paper proposes a new approach to circumnavigate this problem to a large extent by defining a measure which quantifies the instability of an observed quantity over a range of chosen values of another parameter.

Two applications of the proposed approach were considered, both in an extreme value analysis context. One application was to refine the second order parameter estimator proposed by Gomes and Martins (2001), by using the instability measure to distinguish between possible ranges of estimates. An improvement in the accuracy of second order parameter estimation was obtained for large samples, specifically when the sample size exceeds 5000. More importantly, the adjusted estimator

performed better in the ability to choose between different threshold selection techniques, based on the range of the underlying value of ρ .

The main application of the instability measure involved threshold selection when estimating the EVI by fitting the PPD to relative excesses. This was accomplished by determining a region where the EVI estimates can be regarded as most stable. It led to the construction of the estimator $\hat{\gamma}_5$, which shows significant improvement on the corresponding estimator which uses a fixed threshold.

The rest of the discussion will focus on some possible applications of the results presented in this paper. Financial risk managers need to estimate, inter alia, the value at risk (VaR) of portfolios, which is the level below which the future portfolio will drop with a specified small probability. EVT is used to calculate the maximum (extreme) losses that can occur in a given time period. Consider again the results at the end of Section 3, where it was found that for a data set of size $n = 200$ the EVI was estimated as 0.4727, yielding an implied threshold of $k = 50$. The estimated value of ρ (determined externally), as well as the MLEs of γ and c when fitting the PPD to the 50 excesses, yields a fitted model for observations exceeding the threshold. Estimated values of extreme quantiles can be obtained from this model for a given small exceedance probability. If the data represented insurance claims, for example, the corresponding extreme claim amount can be estimated. The accuracy of this estimated claim amount depends heavily on the accuracy of the estimated EVI. The techniques presented in this paper improve on EVI estimation by proposing an improved method of threshold selection.

There are also several possible areas of application of the instability measure outside the realm of EVT. Generally the variance is used as a measure of stability. The main advantage of the proposed instability measure is that it also takes into account the slope of the estimates over the specified region. It therefore does not only detect high variability but also systematic change.

The following are some examples of possible areas of application. Many applications are simply assessing stability of estimates over time, where there is no nuisance parameter which has to be chosen.

- In statistical process control (SPC), both the variance and mean of a production process need to be in-control for the process to run effectively. The instability measure can be applied here as an additional measure which needs to be in-control. Alternatively, some threshold can be decided upon which, if exceeded by the instability measure, can be used to flag the possibility of future out-of-control signals.
- The amount of health data measured on each individual is expected to increase exponentially in the near future with the availability of applications on electronic hand-held devices which record these data automatically. A similar concept to that of the application of the measure in SPC can be applied here to give early warning signals if measurements pertaining to health are becoming unstable.
- Stacy, Guarino, Reckase and Wooldridge (2013) investigate in their working paper the effect of student characteristics on the value-added estimates of teachers. The purpose of these estimates is to evaluate teacher performance, but inter-year instability of estimates is a problem. The problem is partially solved by taking into account in the regression model various variables which explain the variance in the teachers' performance, and also by calculating moving

averages. However, the authors warn that care should be taken when applying these measures. In this setting the proposed instability measure can easily be applied to a large data base in order to identify estimates which are unstable over time. The measure also takes systematic change into account. The cases with the most extreme instability can be identified timeously and objectively. Investigating these cases can lead to new insights as to the causes of unstable performance estimates. These insights can be used to assist with policy decisions, or to help identify additional variables which should be included in the model.

Acknowledgements

The author would like to express his sincere gratitude towards the following persons:

- Professor Tertius de Wet, for his invaluable contributions, specifically his advice and suggestions during the research leading up to this paper, as well as for his long-standing guidance and support.
- Professor Yuri Goegebeur, for his constructive criticism and advice which led to a major revision of the content and presentation of the material.
- The reviewer and associate editor, for their recommendations and suggestions which enhanced the quality of this paper.

References

- BEIRLANT, J., GOEGBEUR, Y., SEGERS, J., AND TEUGELS, J. L. (2004). *Statistics of extremes*. Chichester: John Wiley & Sons.
- COLES, S. G. (2001). *An introduction to statistical modeling of extreme values*. London: Springer.
- DE HAAN, L. AND FERREIRA, A. (2006). *Extreme value theory: an introduction*. Boston: Springer.
- DEKKERS, A. L. M., EINMAHL, J. H. J., AND DE HAAN, L. (1989). A moment estimator for the index of an extreme-value distribution. *Annals of Statistics*, **17**, 1833–1855.
- DIERCKX, G. (2000). *Estimation of the extreme value index*. Phd thesis, Leuven: Katholieke Universiteit Leuven.
- DREES, H. AND KAUFMANN, E. (1998). Selecting the optimal sampling fraction in univariate extreme value estimation. *Stochastic Processes and their Applications*, **75**, 149–172.
- FRAGA ALVES, M. I. (2002). Estimation of first and second order parameters in heavy tails. Technical report, Lisbon: University of Lisbon.
- FRAGA ALVES, M. I., GOMES, M. I., AND DE HAAN, L. (2003). A new class of semi-parametric estimators of the second order parameter. *Portugaliae Mathematica*, **60** (2), 193–213.
- GELUK, J., DE HAAN, L., RESNICK, S., AND STĂRICĂ, C. (1997). Second-order regular variation, convolution and the central limit theorem. *Stochastic Processes and their Applications*, **69**, 139–159.
- GOEGBEUR, Y. AND DE WET, T. (2012). Local estimation of the second order parameter in extreme value statistics and local unbiased estimation of the tail index. *Communications in Statistics - Theory and Methods*, **41**, 1–33.

- GOMES, M. I., CAEIRO, F., AND FIGUEIREDO, F. (2004). Bias reduction of a tail index estimator through an external estimation of the second-order parameter. *Statistics: A Journal of Theoretical and Applied Statistics*, **38** (6), 497–510.
- GOMES, M. I., FIGUEIREDO, F., AND NEVES, M. M. (2012). Adaptive estimation of heavy right tails: resampling-based methods in action. *Extremes*, **15** (4), 463–489.
- GOMES, M. I. AND HENRIQUES-RODRIGUES, L. (2010). A heuristic choice of tuning parameters in a location invariant reduced-bias estimation of the extreme value index: application to financial log-returns and simulated data. In *Proceedings of the ITI 2010 32nd Int. Conf. on Information Technology Interfaces*. Cavtat, Croatia, 21–24 June, p. 527.
- GOMES, M. I., HENRIQUES-RODRIGUES, L., PEREIRA, H., AND PESTANA, D. (2010). Tail index and second-order parameters' semi-parametric estimation based on the log-excesses. *Journal of Statistical Computation and Simulation*, **80** (6), 653–666.
- GOMES, M. I. AND MARTINS, M. J. (2001). Generalizations of the Hill estimator - asymptotic versus finite sample behaviour. *Journal of Statistical Planning and Inference*, **93**, 161–180.
- GOMES, M. I. AND MARTINS, M. J. (2002). “asymptotically unbiased” estimators of the tail index based on external estimation of the second order parameter. *Extremes*, **5** (1), 5–31.
- GOMES, M. I., MARTINS, M. J., AND NEVES, M. M. (2007). Improving second order reduced bias extreme value index estimation. *REVSTAT Statistical Journal*, **5** (2), 117–207.
- GUILLOU, A. AND HALL, P. (2001). A diagnostic for selecting the threshold in extreme value analysis. *Journal of the Royal Statistical Society*, **63**, 293–305.
- HALL, P. AND WELSH, A. H. (1985). Adaptive estimates of parameters of regular variation. *Annals of Statistics*, **13**, 331–341.
- HILL, B. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, **3**, 1163–1174.
- PENG, L. (1998). Asymptotically unbiased estimator for the extreme-value index. *Statistics and Probability Letters*, **38** (2), 107–115.
- SINGH, S. K. AND MADDALA, G. S. (1976). A function for the size distribution of incomes. *Econometrica*, **44**, 963–970.
- STACY, B., GUARINO, C. M., RECKASE, M. D., AND WOOLDRIDGE, J. M. (2013). Does the precision and stability of value-added estimates of teacher performance depend on the types of students they serve? *IZA Discussion Paper*, **7676**.