

ExtremeBounds: Extreme Bounds Analysis in R

Marek Hlavac
Harvard University

Abstract

This article introduces the R package **ExtremeBounds** to perform extreme bounds analysis (EBA), a sensitivity test that examines how robustly the dependent variable of a regression model is related to a variety of possible determinants. **ExtremeBounds** supports Leamer's EBA that focuses on the upper and lower extreme bounds of regression coefficients, as well as Sala-i-Martin's EBA which considers their entire distribution. In contrast to existing alternatives, it can estimate models of any size, use regression models other than Ordinary Least Squares, incorporate non-linearities in the model specification, and apply custom weights and standard errors. To alleviate concerns about the multi-collinearity and conceptual overlap of examined variables, **ExtremeBounds** allows users to specify sets of mutually exclusive variables, and can restrict the analysis to coefficients from regression models that yield a variance inflation factor within a pre-specified limit.

Keywords: extreme bounds analysis, robustness, sensitivity, regression, R.

1. Introduction

In this article, I introduce the R package **ExtremeBounds** to perform extreme bounds analysis (EBA), a sensitivity test that examines how robustly the dependent variable of a regression model is associated with a variety of possible determinants. EBA sifts through a large number of model specifications to answer the following questions:

- Which determinants are robustly associated with the dependent variable across a large number of possible regression models?
- Is a particular determinant robustly associated with the dependent variable?

Extreme bounds analysis has been used to examine a plethora of questions of interest to social scientists. Economists have used it to examine the determinants of long-term economic growth (Sala-i-Martin 1997; Levine and Renelt 1992; Sturm and de Haan 2005), regional growth rates (Reed 2009), foreign direct investment (Moosa and Cardak 2006), as well as investment in research and development (Wang 2010). Political scientists have analyzed democratization (Gassebner, Lamla, and Vreeland 2013), political repression (Hafner-Burton 2005) and lending decisions by the International Monetary Fund (Moser and Sturm 2011). Other examples of EBA in the social scientific literature include an examination of the relationship between wage inequality and crime rates (Fowles and Merva 1996), of the effects of concealed weapons laws in the United States (Bartley and Cohen 1998), and even of the determinants of Olympic medal counts (Moosa and Smith 2004).

The R package **ExtremeBounds** supports a demanding version of EBA, proposed by [Leamer \(1985\)](#), that focuses on the upper and lower extreme bounds of regression estimates, as well as a more flexible version by [Sala-i-Martin \(1997\)](#). Sala-i-Martin's EBA considers the entire distribution of regression coefficients. For Sala-i-Martin's version of extreme bounds analysis, **ExtremeBounds** estimates the normal model in which regression coefficients are assumed to be normally distributed across models, as well as the generic model which makes no such assumption.

Despite the relative ease with which EBA can be implemented in most statistical programming languages, only a handful of ready-made software solutions exist for its estimation. Available packages include the module **eba** ([Impavido 1998](#)) for **Stata** ([StataCorp 2013](#)), procedure **EBA** ([Doan 2004](#)) for **RATS** ([Estima 2013](#)) and program **MICRO-EBA** ([Fowles 1988](#)) for **Gauss** ([Aptech Systems 2013](#)) and **SHAZAM** ([SHAZAM Analytics 2013](#)).

These packages can only estimate Leamer's extreme bounds analysis ([Shiba 1992](#)) or, in the case of **EBA** for **RATS**, its slightly modified version suggested by [Granger and Uhlig \(1990\)](#). Recent scholarship has, however, tended to follow [Sala-i-Martin \(1997\)](#) and consider the entire distribution of coefficients that emerge from the extreme bounds analysis (e.g., [Gassebner et al. 2013](#); [Sturm and de Haan 2005](#)). In addition, researchers have used histograms to illustrate the distribution of estimated regression coefficients graphically (e.g., [Hegre and Sambanis 2006](#)). The packages listed above do not support Sala-i-Martin's EBA or generate graphical output. As a result, they do not provide a satisfactory implementation of extreme bounds analysis that would accurately reflect the method's contemporary use.

The **ExtremeBounds** package, by contrast, estimates both Leamer's and Sala-i-Martin's versions of extreme bounds analysis. It is the first EBA package for the R statistical programming language ([R Core Team 2014](#)). It can be installed free of charge from the [Comprehensive R Archive Network \(CRAN\)](#) ([2014](#)) in the usual way:

```
R> install.packages("ExtremeBounds")
```

The package provides users with a large number of important features that distinguish it from existing alternatives for EBA estimation. In particular, **ExtremeBounds** can estimate models of any size, use regression models other than Ordinary Least Squares, incorporate non-linearities in the model specification, and apply custom weights and standard errors. To alleviate concerns about the multicollinearity or conceptual overlap of examined variables, **ExtremeBounds** allows users to specify sets of mutually exclusive variables. It can also restrict the analysis to coefficients from regression models that yield a variance inflation factor within a given maximum limit.

In the next section, I briefly describe Leamer's and Sala-i-Martin's extreme bounds analysis. I then provide an overview of the **ExtremeBounds** package's capabilities in [Section 3](#). In [Section 4](#), I demonstrate them on an empirical example that involves the fuel efficiency of automobiles. [Section 5](#) concludes.

2. Extreme bounds analysis (EBA)

In this section, I present a brief description of extreme bounds analysis. My discussion is by no means intended to be exhaustive. Instead, I aim to familiarize the reader with the

fundamentals of this method’s estimation procedure. More detailed and rigorous treatments of extreme bounds analysis can be found in [Leamer \(1985\)](#), [Leamer and Leonard \(1983\)](#) and [Sala-i-Martin \(1997\)](#). For a critical perspective on EBA, please refer to [McAleer, Pagan, and Volker \(1985\)](#), [Breusch \(1990\)](#), [Hendry and Krolzig \(2004\)](#) or [Angrist and Pischke \(2010\)](#).

The basic idea of extreme bounds analysis is quite simple. We are interested in finding out which variables from the set \mathbf{X} are robustly associated with the dependent variable y . To do so, we run a large number of regression models. Each has y as the dependent variable and includes a set of standard explanatory variables \mathbf{F} that are included in each regression model. In addition, each model includes a different subset \mathbf{D} of the variables in \mathbf{X} . Following the convention in the literature, we will refer to \mathbf{F} as the free variables and to \mathbf{X} as the doubtful variables. Some subset of the doubtful variables \mathbf{X} might be so-called focus variables that are of particular interest to the researcher. The doubtful variables whose regression coefficients retain their statistical significance in a large enough proportion of estimated models are declared to be robust, whereas those that do not are labelled fragile.

More formally, to find out whether a focus variable $v \in \mathbf{X}$ is robustly correlated with the dependent variable y , we estimate a set of regression models of the following form:

$$y = \alpha_j + \beta_j v + \gamma_j \mathbf{F} + \delta_j \mathbf{D}_j + \varepsilon \quad (1)$$

where j indexes regression models, \mathbf{F} is the full set of free variables that will be included in every regression model, \mathbf{D}_j is a vector of k variables taken from the set \mathbf{X} of doubtful variables, and ε is the error term. While \mathbf{D}_j has conventionally been limited to no more than three doubtful variables per model ([Levine and Renelt 1992](#); [Achen 2005](#)), the particular choice of k , the number of doubtful variables to be included in each combination, is up to the researcher.

The above regression is estimated for each of the M possible combinations of $\mathbf{D}_j \subset \mathbf{X}$. The estimated regression coefficients $\hat{\beta}_j$ on the focus variable v , along with the corresponding standard errors $\hat{\sigma}_j$, are collected and stored for use in later calculations. In the original formulation of extreme bounds analysis, the regressions were estimated by Ordinary Least Squares (OLS). In recent research, however, other types of regression models have also been used, such as ordered probit models ([Bjørnskov, Dreher, and Fischer 2008](#); [Hafner-Burton 2005](#)) or logistic models ([Hegre and Sambanis 2006](#); [Moser and Sturm 2011](#); [Gassebner *et al.* 2013](#)).

2.1. Leamer’s EBA

In order to determine whether a determinant is robust or fragile, Leamer’s extreme bounds analysis focuses only on the extreme bounds of the regression coefficients ([Leamer 1985](#)). For any focus variable v , the lower and upper extreme bounds are defined as the minimum and maximum values of $\hat{\beta}_j \pm \tau \hat{\sigma}_j$ across the M estimated regression models, where τ is the critical value for the requested confidence level. For the conventional 95-percent confidence level, τ will thus be equal to approximately 1.96. If the upper and lower extreme bounds have the same sign, the focus variable v is said to be robust. Conversely, if the bounds have opposite signs, the variable is declared fragile.

Leamer’s EBA relies on a very demanding criterion for robustness, since the results from a single regression model are enough to classify a determinant as fragile. In other words, a

focus variable will be declared fragile even if the extreme bounds have the same sign in all estimated models except one. Accordingly, [Sala-i-Martin \(1997\)](#) notes that “if the distribution of [regression coefficients] has some positive and some negative support, then one is bound to find one regression for which the estimated coefficient changes signs if enough regressions are run.” It should come as no surprise that studies that have employed Leamer’s EBA to test the robustness of determinants have generally concluded that most, if not all, examined variables are fragile ([Levine and Renelt 1992](#); [Levine and Zervos 1993](#); [Sala-i-Martin 1997](#)).

2.2. Sala-i-Martin’s EBA

In response to the perceived stringency of Leamer’s EBA, [Sala-i-Martin \(1997\)](#) proposes an alternative method for extreme bounds analysis that focuses on the entire distribution of regression coefficients, not just on its extreme bounds. Instead of applying a binary label of robust or fragile, he assigns some level of confidence to the robustness of each of the variables. In particular, [Sala-i-Martin \(1997\)](#) considers the value of $\text{CDF}(0)$, the fraction of the variable’s cumulative distribution that lies on each side of zero. To motivate his approach, he points out that “if 95 percent of the density function for the estimates of β_1 lies to the right of zero and only 52 percent of the density function for β_2 lies to the right of zero, one will probably think of variable 1 as being more likely to be correlated with [the dependent variable] than variable 2.”

Although the coefficients in each individual model have an asymptotic normal distribution, the coefficient estimates obtained from different regression models might be scattered in various ways. For this reason, [Sala-i-Martin \(1997\)](#) presents two variants of his extreme bounds analysis – a normal model, in which the estimated regression coefficients are assumed to follow a normal distribution across the estimated models, as well as a generic model, which does not assume any particular distribution of regression coefficients.

To estimate the normal model, Sala-i-Martin first calculates the weighted mean of the regression coefficients $\hat{\beta}_j$ and of the variances $\hat{\sigma}_j^2$:

$$\bar{\beta} = \sum_{j=1}^M w_j \hat{\beta}_j \quad (2)$$

$$\bar{\sigma}^2 = \sum_{j=1}^M w_j \hat{\sigma}_j^2 \quad (3)$$

where w_j represents weights that are applied to results from each estimated regression model. Once the weighted means of coefficients and standard errors are known, Sala-i-Martin calculates $\text{CDF}(0)$ based on the assumed normal distribution of regression coefficients such that

$$\beta \sim \mathcal{N}(\bar{\beta}, \bar{\sigma}^2) \quad (4)$$

In the generic model, Sala-i-Martin first uses the sampling distribution of the regression coefficient $\hat{\beta}_j$ to obtain an individual $\text{CDF}(0)$, denoted by $\phi_j(0 \mid \hat{\beta}_j, \hat{\sigma}_j^2)$, for each estimated regression model. He then calculates the aggregate $\text{CDF}(0)$ for β as the weighted average of all the individual $\text{CDF}(0)$ ’s:

$$\Phi(0) = \sum_{j=1}^M w_j \phi_j(0 \mid \hat{\beta}_j, \hat{\sigma}_j^2) \quad (5)$$

In both the normal and the generic model, Sala-i-Martin applies weights that are proportional to the integrated likelihood to give greater weight to models that provide a better fit:

$$w_j = \frac{L_j}{\sum_{i=1}^M L_i} \quad (6)$$

In principle, of course, the weights could be based on any other measure of the goodness of fit. Examples used in existing research literature include McFadden’s likelihood ratio index (McFadden 1974) used by Hegre and Sambanis (2006), or applying equal weights to each regression model (Sturm and de Haan 2005; Gassebner *et al.* 2013).

3. Overview of the **ExtremeBounds** package

The **ExtremeBounds** package consists of the main function `eba()`, which performs the extreme bounds analysis, as well as of two related methods – `print()` and `hist()` – which produce, respectively, histograms and text output to summarize the estimation results. In this section, I provide an overview of these functions’ capabilities, and highlight features that make **ExtremeBounds** the most versatile of the available EBA estimation tools. Users can obtain a more detailed description of the arguments and output of each function by typing, as appropriate, `?eba`, `?print.eba` or `?hist.eba` into the R console.

3.1. EBA estimation: Main function `eba()`

The main function `eba()` performs both Leamer’s and Sala-i-Martin’s versions of extreme bounds analysis using variables from a data frame specified in the function’s argument `data`. The user specifies the dependent variable (argument `y`), the free variables to be included in all regression models (`free`), the focus variables that are of interest (`focus`), as well as the full set of doubtful variables (`doubtful`). Note that the variables included in `focus` must be a subset of those in `doubtful`. If the user does not provide any focus variables, `eba()` will assume that all doubtful variables are of interest to the researcher.

It is often more convenient to specify the model variables through a single `formula` argument. The `eba()` function allows the user to pass on a multiple-part formula that specifies the dependent, free, focus and doubtful variables. These variables, along with the functional form of the model, are passed on to the `formula` argument in an object of class "Formula", as implemented by the **Formula** package (Zeileis and Croissant 2010). The model formula takes the following most general form: `y ~ free | focus | doubtful`. If all doubtful variables are of interest (i.e., all are focus variables), the user can pass on `y ~ free | focus`. Finally, one can also specify a model, in which there are no free variables and all doubtful variables are of interest: `y ~ focus`. When using the `formula` argument, the `doubtful` component of the right-hand side need not contain the `focus` variables. Section 4 provides a practical example of specifying the EBA model using both the single `formula` argument and the separate `y`, `free`, `focus` and `doubtful` arguments.

If no other arguments are specified, `eba()` will conduct a ‘standard’ extreme bounds analysis based on Ordinary Least Squares (OLS) regressions with unadjusted standard errors. All hypothesis tests will be performed at the conventional 95 percent confidence level against the

null hypothesis that the relevant parameter is equal to zero. Equal weights will be applied to results from each estimated model, and no maximum limit will be imposed on the variance inflation factor. Following [Levine and Renelt \(1992\)](#), `eba()` will include up to three doubtful variables in addition to the focus variable in each regression model.

The function `eba()`, however, also offers a great deal of flexibility for researchers who might be interested in conducting a different kind of extreme bounds analysis. Every aspect of EBA is fully customizable:

- `eba()` can estimate any type of regression model (argument `reg.fun`), not just Ordinary Least Squares (OLS). Researchers can thus easily perform EBA using, for instance, logistic or probit regressions. Attention can, furthermore, be restricted to regressions in which the variance inflation factor on the examined coefficient – a rule-of-thumb indicator of multicollinearity ([Mansfield and Helms 1982](#)) – does not exceed a set maximum (argument `vif`). Alternatively, users can request that `eba()` only use results from regression models that meet some other, user-specified condition (argument `include.fun`).
- Regression models can be specified very flexibly (argument `formula`). They can contain interaction terms, squared or cubic terms to account for non-linear relationships, as well as the lags, natural logarithms or other functions of the included variables. The function automatically transforms the variables and creates the appropriate model matrices. These regression models can be of any size, as the user can choose how many variables should be included in each specification (argument `k`).
- Users can specify sets of mutually exclusive variables that will never be included in the same regression model to alleviate concerns about regressor multicollinearity (argument `exclusive`). Specifying which doubtful variables cannot be included together can also be useful when several doubtful variables measure the same substantive concept.
- All hypothesis tests can be performed at any requested confidence level (argument `level`). In addition, the user can specify the null hypothesis value for each regression coefficient (argument `mu`). `eba()` can thus check whether the estimated coefficients are robustly different from any numerical value, not just from zero.
- If desired, weights can be applied to results from each estimated regression (argument `weights`). The weights can be based on the regression R^2 , adjusted R^2 , McFadden's likelihood ratio index ([McFadden 1974](#)) or calculated by a user-provided function.
- In conjunction with other R packages, `eba()` can apply various types of standard errors in its calculations (argument `se.fun`). It can, for instance, apply heteroskedasticity-robust ([Huber 1967](#); [Eicker 1967](#); [White 1980](#)), clustered ([Froot 1989](#); [Williams 2000](#)), and Newey-West standard errors ([Newey and West 1987](#)) provided by the `sandwich` package ([Zeileis 2004, 2006](#)), as well as panel-corrected standard errors ([Beck and Katz 1995](#)) from the `pcse` package ([Bailey and Katz 2011](#)). Additionally, users can also provide `eba()` with their own functions to calculate standard errors.
- To reduce the time required to complete the analysis, `eba()` can estimate a random sample of any given size drawn from the full set of the regression models (argument `draws`). This procedure yields unbiased estimates of the quantities of interest that can

give the researcher a good sense of the full EBA results within a reasonable time frame (Sala-i-Martin, Doppelhofer, and Miller 2004).

The function `eba()` returns an object of class "eba", which can then be passed on to the `print()` method to obtain a text summary of EBA results, or to the `hist()` method to produce histograms that will provide a graphical illustration of the estimation results. The object contains a `bounds` data frame with the results of both Leamer's and Sala-i-Martin's extreme bounds analysis, as well as a `coefficients` data frame with various quantities of interest: the minimum, maximum, mean and median values of coefficient estimates, along with the individual CDF(0)'s for Sala-i-Martin's generic EBA model.

In addition, the object stores the total number of doubtful variable combinations that include at least one focus variable (component `ncomb`), the number of regressions actually estimated in total (`nreg`) and by variable (`nreg.variable`), along with the number of coefficients used in the extreme bounds analysis (`ncoef.variable`). Importantly, the "eba"-class object also contains the estimation results from each regression model (component `regressions`). As a result, researchers can easily use the large number of regressions that EBA often produces in their own analyses, whether those be modifications of extreme bounds analysis or an entirely different statistical method.

3.2. Text summary: Method `print()`

Once the function `eba()` has completed its calculations, the user can obtain a text summary of the estimation results by passing the "eba"-class object to the `print()` method. The summary contains information about the number of regressions that `eba()` has estimated, as well as about the distribution of regression coefficient estimates. Most importantly, it provides a comprehensive summary of the analysis results for both Leamer's and Sala-i-Martin's EBA. The user can adjust the number of decimal places to which the all numerical figures in the output are rounded by changing the value of the `digits` argument.

At the top of the text output, the `print()` method reproduces the `eba()` function call, the confidence level, as well as the total number of variable combinations, the number of regressions that were actually estimated (in total, by variable, and as a proportion of the number of combinations) and the number of coefficients used in the EBA (by variable). The remainder of the output is divided into four parts, which I briefly summarize below:

- **Beta coefficients:** This part contains the weighted mean of the estimated regression coefficients and of their standard errors. Individual regression models receive a weight specified by the `eba()` function's argument `weights`. In addition, the `print()` method also reports the value of the lowest and highest regression coefficients across the estimated models, along with the corresponding standard errors.
- **Distribution of beta coefficients:** Here, the method reports the percentage of regression coefficients that are lower/greater than zero, along with the proportion that is statistically significantly different from zero at the specified confidence level. If a different value under the null hypothesis was specified, all regressions coefficients are compared with `mu` rather than with zero.
- **Leamer's Extreme Bounds Analysis:** Both the lower and upper extreme bounds, at

the specified confidence levels, from Leamer's extreme bounds analysis are reported. Based on these bounds, each variable is classified as either `robust` or `fragile`.

- **Sala-i-Martin's Extreme Bounds Analysis:** Finally, the `print()` method reports results from Sala-i-Martin's EBA. In particular, for both the normal and generic models, the text output prints out the value of the aggregate CDF(0) (or CDF(`mu`), if appropriate), along with its complement $1 - \text{CDF}(0)$.

3.3. Histograms: Method `hist`

In addition to providing text output with EBA results, the **ExtremeBounds** package can produce histograms that summarize the distribution of estimated regression coefficients graphically for each examined variable. These histograms can, furthermore, be superimposed with curves that depict the corresponding kernel density or a normally distributed approximation to the coefficient distribution.

To produce EBA histograms, the user simply passes an "eba"-class object created by the main function `eba()` to the `hist()` method. The user can choose between histograms that represent frequencies or probability densities (argument `freq`). Unless the set of variables to be included is specified (argument `variables`), the method will produce histograms for all variables included in the extreme bounds analysis.

By default, the histograms include vertical lines that indicate the parameter value under the null hypothesis (argument `mu.show`), as well as a kernel density curve that relies on R's standard `density` method (argument `density.show`). The kernel density curve provides a non-parametric estimate of the EBA coefficients' probability density function. In addition, the `hist()` method for "eba" objects can overlay the histograms with a density curve for a normal distribution function that approximates the distribution of EBA coefficients (argument `normal.show`).

Many formatting options are available. Users can change the colors and widths of the vertical lines for null hypothesis values (arguments `mu.col` and `mu.lwd`), of the kernel density curves (`density.col` and `density.lwd`), and of the density curves for the normally distributed approximation (`normal.col` and `normal.lwd`). The user also has complete control over a variety of other visual properties of the histograms. These include the histograms' title labels (argument `main`), the range of values on the horizontal axis (`xlim`), as well as the color of the histogram bars (`col`).

4. Example: The fuel economy of automobiles

In this section, I demonstrate the capabilities of the **ExtremeBounds** package using an empirical example. In particular, I identify robust determinants of the fuel economy of selected automobiles using data from the data frame `mtcars`. This data frame is included in the **datasets** package, which is part the standard R distribution (R Core Team 2014) and is therefore readily available to the beginning **ExtremeBounds** user.

The information in `mtcars` was extracted from the 1974 *Motor Trend* magazine, and comprises the fuel consumption and ten aspects of vehicle design and performance for a selection of 1973–1974 automobile models. Existing research literature has already taken advantage of

these data for the purpose of demonstrating various statistical methods and procedures (Hocking 1976; Henderson and Velleman 1981). The `mtcars` data set is particularly well-suited for a demonstration of extreme bounds analysis, as its small size allows me to highlight the **ExtremeBounds** package's most important features without having to perform time-consuming estimations of a very large number of regressions.

The data frame `mtcars` contains 32 observations. Each observation is one particular model of automobile (e.g., "Mazda RX4", "Cadillac Fleetwood" or "Fiat X1-9"). For each model, the data frame contains information on its fuel economy, expressed as the vehicle's miles per gallon (variable `mpg`), and about its engine – the number of cylinders (`cyl`) and carburetors (`carb`), its displacement in cubic inches (`disp`), its gross horsepower (`hp`), as well as whether it is a V- or a straight engine (`vs`). In addition, we are given information about each model's rear axle ratio (`drat`), weight in thousands of pounds (`wt`), and quarter-mile time in seconds (`qsec`). The variable `gear` contains the number of forward gears, while `am` indicates whether the automobile has an automatic or manual transmission, with a value of 1 denoting a manual transmission.

4.1. Naïve EBA with all combinations of doubtful variables

First, I would like to get a basic sense of which determinants might be most robustly associated with the dependent variable `mpg` (miles per gallon). I therefore begin by conducting an EBA that estimates all possible combinations of doubtful variables across the ten automobile design and performance characteristics included in `mtcars`. Since I am interested in the robustness or fragility of all the doubtful variables, I regard all of them as focus variables.

From a statistical point of view, this type of EBA is somewhat naïve, as it does not take into account the possibility of high multicollinearity among the included variables. Neither does it account for the possibility that some variables measure similar concepts. The number of cylinders (`cyl`) and the gross horsepower (`hp`) might, for example, both be seen as measures of the engine's overall performance. Researchers interested in examining the fuel economy of automobiles would thus, in contrast to my naïve EBA, be unlikely to include both explanatory variables in the same regression model.

An extreme bounds analysis with all combinations of doubtful variables might nevertheless yield some valuable insights. In particular, it provides a particularly strong test for a determinant's robustness. As McAleer *et al.* (1985) suggest, such an EBA might indicate which variables should be treated as free, and therefore be included in all regression models in further EBA analyses.

I estimate the naïve EBA by calling the `eba()` function:

```
R> naive.eba <- eba(formula = mpg ~ cyl + carb + disp + hp + vs + drat + wt
+ qsec + gear + am, data = mtcars, k = 0:9)
```

Alternatively, I can use character vectors to specify the dependent and doubtful variables using the `y` and `doubtful` arguments, respectively:

```
R> naive.eba <- eba(data = mtcars, y = "mpg", doubtful = c("cyl", "carb",
"disp", "hp", "vs", "drat", "wt", "qsec", "gear", "am"), k = 0:9)
```

Since the `focus` argument is not specified, all the `doubtful` variables will be treated as focus variables. The argument `k = 0:9` ensures that, on top of the focus variable, up to nine doubtful variables will be included in each regression model. As a result, regressions with all possible combinations of the ten doubtful variables will be estimated. The `eba()` function will return an object of class `"eba"` that will be stored in `naive.eba`.

Next, I ask **ExtremeBounds** to produce a set of histograms that summarize the EBA estimation results. To do so, I simply pass the `naive.eba` object to the `hist()` method:

```
R> hist(naive.eba)
```

The resulting histograms are reproduced in Figure 1. The gray bins contain the Ordinary Least Squares (OLS) coefficients on each examined variable from all of the estimated regression models. Superimposed over each histogram is a thick blue curve that represents the corresponding kernel density, a non-parametric approximation of the regression coefficients' distribution. There is also a red vertical line at zero, the default coefficient value under the null hypothesis.

A visual inspection of the histograms allows me to get a quick overview of the EBA estimation results. If most of the histogram bins' area lies to the right of zero, a majority of the regression coefficient estimates on the corresponding variables are positive. A positive coefficient indicates that, holding all else equal, a higher value of the examined variable is associated with more miles per gallon, as given by the dependent variable `mpg`). The results of the naïve EBA suggests that straight engines (`vs`), a greater rear axle ratio (`drat`), a slower quarter-mile time (`qsec`), a greater number of forward gears (`gear`) and a manual transmission (`am`) are associated with greater fuel economy.

Conversely, if most of the bins' area lies to the left of zero, greater values of the corresponding variable are associated with lower miles per gallon, *ceteris paribus*, in most estimated regressions. Naïve EBA estimation results indicate that engines with more cylinders (`cyl`), carburetors (`carb`) and with greater gross horsepower (`hp`) achieve worse fuel economy. The vehicle's greater weight (`wt`) is consistently associated with greater fuel consumption. In fact, `wt` is the only variable for which all of the estimated regression coefficients have the same (in this case, negative) sign.

Engine displacement (`disp`) appears to be an interesting case, as the distribution of the regression coefficients appears to be bimodal. Some are negative, while others are positive. The bimodal nature of the distribution can be easily seen from the two different peaks of the histogram bins, as well as from the double hump of the kernel density curve.

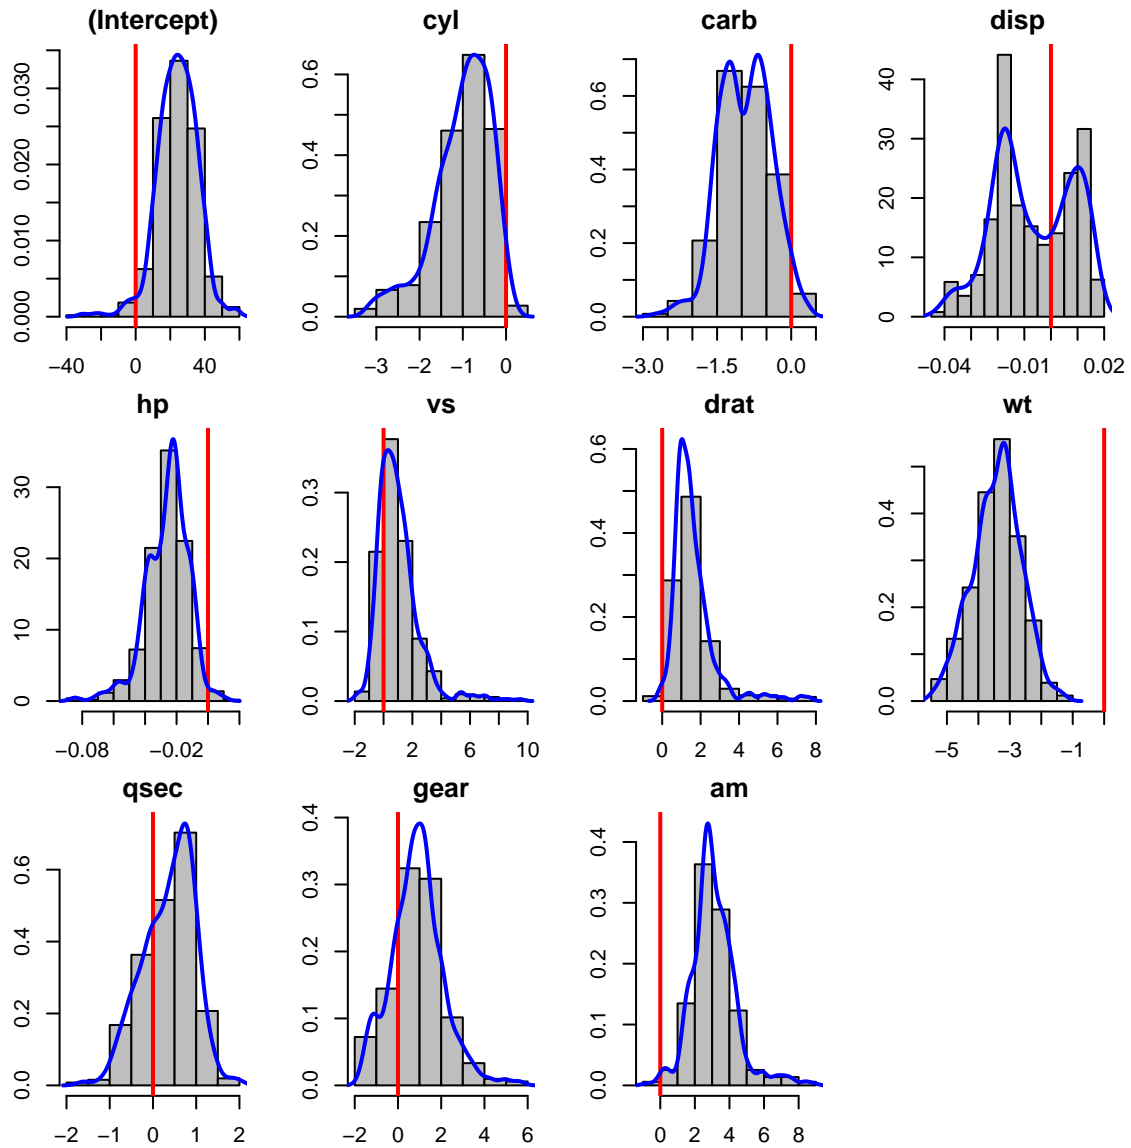


Figure 1: Histograms that summarize the estimation results of the naïve EBA. The magnitudes of regression coefficients are on the horizontal axis. The vertical axis indicates the corresponding probability density.

4.2. A more sophisticated EBA

Having inspected the results of my naïve EBA, I can now use the `ExtremeBounds` package's capabilities to make my analysis more sophisticated. In particular, I make the following changes to the `eba()` function call:

- Results from the naïve EBA have indicated that the regression coefficients on the vehicle's weight (variable `wt`) are negative, regardless of the model specification. For this reason, I treat `wt` as a free variable to be included in all regression models.
- Some of the doubtful variables measure similar concepts, and were therefore inappropriately included together in regression models by the naïve EBA. To prevent conceptual overlap, I use the `exclusive` parameter to specify two sets of mutually exclusive variables:
 - One set consists of variables that might be construed as measuring the performance of the engine: the number of cylinders (`cyl`), the number of carburetors (`carb`), engine displacement (`disp`) and the gross horsepower (`hp`).
 - The other set consists of the two doubtful variables that have to do with the car's transmission: the `am` indicator of an automatic vs. manual transmission, and the number of forward gears (`gear`).
- I am only interested in estimation results for the four variables that measure engine performance: `cyl`, `carb`, `disp` and `hp`. I therefore specify them as the focus variables.
- Rather than estimating all possible combinations of the doubtful variables as I did in the naïve EBA, I only add combinations of up to three doubtful variables to the focus variable in each specification. The value of the `k` argument will thus remain at its default value of 0:3.
- To eliminate the influence of coefficient estimates from model specifications that suffer from high multicollinearity, I specify a maximum acceptable variance inflation factor by setting `vif = 7`.
- I use heteroskedasticity-robust standard errors (White 1980), as calculated by the `sandwich` package (Zeileis 2004, 2006). To be able to do this, I define the `se.robust` function (below) that calculates the standard errors, and pass it to the `eba()` function:

```
library("sandwich")

se.robust <- function(model.object) {
  model.fit <- vcovHC(model.object, type = "HC")
  out <- sqrt(diag(model.fit))
  return(out)
}
```

- Finally, I give more weight to estimation results from regression models that provide a better fit to the data. More specifically, I set the argument `weights` to `"lri"`, and thus weight each model's results by its likelihood ratio index (McFadden 1974).

I execute the following R code to estimate this more sophisticated extreme bounds analysis:

```
R> sophisticated.eba <- eba(formula = mpg ~ wt | cyl + carb + disp + hp |
+                           vs + drat + wt + qsec + gear + am,
+                           data = mtcars,
+                           exclusive = ~ cyl + carb + disp + hp | am + gear,
+                           vif = 7, se.fun = se.robust, weights = "lri")
```

I could, of course, achieve the same result by passing the sets of variables via the `y`, `free`, `doubtful` and `focus` arguments, as shown below. The `exclusive` argument can accept a list of character vectors in lieu of a multiple-part "Formula" object.

```
R> doubtful.variables <- c("cyl", "carb", "disp", "hp", "vs",
+                          "drat", "wt", "qsec", "gear", "am")
R> engine.variables <- c("cyl", "carb", "disp", "hp")
R> transmission.variables <- c("am", "gear")

R> sophisticated.eba <- eba(data = mtcars, y = "mpg", free = "wt",
+                          doubtful = doubtful.variables,
+                          focus = engine.variables,
+                          exclusive = list(engine.variables,
+                                           transmission.variables),
+                          vif = 7, se.fun = se.robust, weights = "lri")
```

Again, I produce a set of histograms that will allow me to get an initial sense of the results of my analysis. This time, I include the `hist()` method's `variables` argument to request histograms only for the four focus variables that I am interested in. Additionally, I use the `main` argument to make each histogram's main title more descriptive and the `normal.show` argument to request that `hist()` superimpose a density curve with a normally distributed approximation to the coefficient distribution.

```
R> hist(sophisticated.eba, variables = c("cyl", "carb", "disp", "hp"),
+       main = c(cyl = "number of cylinders", carb = "number of carburetors",
+               disp = "engine displacement", hp = "gross horsepower"),
+       normal.show = TRUE)
```

As the histograms in Figure 2 show, the more sophisticated EBA leads to more clear-cut predictions about the signs of regression coefficients than the naïve EBA estimated earlier. The coefficients on all four focus variables are consistently negative. This result suggests that, across a wide variety of reasonably well-specified regression models, a greater number of cylinders and carburetors, as well as greater engine displacement and gross horsepower, are associated with worse fuel economy (i.e., with fewer miles per gallon).

It is, moreover, evident that the blue kernel density curve and the green normally distributed approximation are quite different from each other. This lack of alignment suggests that regression coefficients are not normally distributed across model specifications.




Figure 2: Histograms that summarize the estimation results of the sophisticated EBA. The magnitudes of regression coefficients are on the horizontal axis. The vertical axis indicates the corresponding probability density.

In addition to histograms, **ExtremeBounds** allows users to examine EBA results through the text output produced by the `print()` method. This method produces a wealth of detailed information about the estimation results from both Leamer's and Sala-i-Martin's extreme bounds analysis. In the interest of clarity, I reproduce only the portions of `print()` output that are most relevant to my empirical example. **ExtremeBound** users can type `?print.eba` in the R console to obtain a complete description of the text output.

I obtain a text summary of my analysis results by passing the "eba"-class object to the `print()` method:

```
R> print(sophisticated.eba)
```

At the top of the `print()` method's text output, we find information about the number of specifications that were estimated. In total, there are 148 possible combinations of the examined doubtful variables that contain at least one focus variable. Since the parameter `draws` is not specified, no random sampling of estimated models occurs. As a result, `eba()` estimates all of the available combinations.

The free variable `wt` occurs, of course, in all 148 regression models. The focus variables are included in fewer specifications since each of them does not appear in some doubtful variable combinations. The four focus variables have, furthermore, been specified as mutually exclusive (in the `eba()` function's argument `exclusive`) and cannot therefore be included together in the same regression model. The `print()` method reports that the focus variables `cyl`, `carb`, `disp` and `hp` appear in 37 specifications each. The output also indicates that only 26 and 14 coefficient estimates (rather than all 37) were used in the extreme bounds analysis for the variables `cyl` and `disp`, respectively. The reduction in the number of coefficients used occurs because the variance inflation factors on some `cyl` and `disp` coefficients exceed the specified maximum of 7.

```
Number of combinations: 148
Regressions estimated: 148 (100% of combinations)
```

Number of regressions by variable:

(Intercept)	<code>wt</code>	<code>cyl</code>	<code>carb</code>	<code>disp</code>	<code>hp</code>
148	148	37	37	37	37

Number of coefficients used by variable:

(Intercept)	<code>wt</code>	<code>cyl</code>	<code>carb</code>	<code>disp</code>	<code>hp</code>
148	148	26	37	14	37

Next, the `print()` method reports the weighted means of coefficient and standard error estimates for the free and focus variables across the estimated regression models. The weights are specified in the `eba()` function's `weights` argument. In the sophisticated analysis I ran, I weight regression models by the likelihood ratio index.

The table reproduced below shows that, on average, a vehicle weight (`wt`) that is greater by a thousand pounds is associated with 3.623 fewer miles per gallon (`mpg`). An additional cylinder (`cyl`) or carburetor (`carb`) yields a fuel consumption that is lower by 1.370 and 0.822 miles per gallon, respectively. One cubic inch of additional engine displacement (`disp`) is associated with 0.016 fewer miles per gallon. Finally, a gallon of fuel yields 0.027 fewer miles with each unit of gross horsepower (`hp`). Relative to the size of the coefficients, the weighted standard errors are relatively small.

Beta coefficients:

	Type	Coef (Wgt Mean)	SE (Wgt Mean)
(Intercept)	free	26.199	6.286
wt	free	-3.623	0.902
cyl	focus	-1.370	0.403
carb	focus	-0.822	0.327
disp	focus	-0.016	0.008
hp	focus	-0.027	0.008

The text output from `print()` proceeds to summarize the distribution of coefficient estimates. It reports the proportions, expressed as percentages, of estimated regression coefficients that are lower or greater than zero. In our example, the coefficients on `wt`, `cyl`, `carb`, `disp` and `hp` are all negative, while the coefficients on the intercept term in the linear regression are always positive. These statistics are, of course, consistent with the information displayed in the histograms in Figure 2.

Distribution of beta coefficients:

	Type	Pct(beta < 0)	Pct(beta > 0)
(Intercept)	free	0	100
wt	free	100	0
cyl	focus	100	0
carb	focus	100	0
disp	focus	100	0
hp	focus	100	0

Even though all the free and focus variables' coefficients are negative, some of the point estimates may not be distinguishable from zero in a statistically significant way. The text output also includes the percentages of regression coefficients that are *both* statistically significant and lower/greater than zero. We find that only in the case of the free variable `wt` are all the coefficients statistically significant. A very large majority of coefficient estimates are significant for the `cyl` and `hp` variables (92.3 and 81.1 percent, respectively). By contrast, only about 60 percent of coefficients on `carb` and `disp` are statistically significant.

Distribution of beta coefficients:

	Type	Pct(signif & beta < 0)	Pct(signif & beta > 0)
(Intercept)	free	0.000	79.73
wt	free	100.000	0.00
cyl	focus	92.308	0.00
carb	focus	59.459	0.00
disp	focus	57.143	0.00
hp	focus	81.081	0.00

The `print()` method then goes on to summarize results from Leamer’s extreme bounds analysis. It produces a table that includes the lower and upper extreme bounds of regression coefficient estimates, and – based on these bounds – classifies determinants as robust or fragile. The only variable that is found to be robust using Leamer’s EBA is the free variable `wt`. Since the upper and lower extreme bounds of all the focus variables have opposite signs, they are declared to be fragile.

Leamer's Extreme Bounds Analysis (EBA):

	Type	Lower Extreme Bound	Upper Extreme Bound	Robust/Fragile?
(Intercept)	free	-19.521	55.021	fragile
<code>wt</code>	free	-7.495	-0.659	robust
<code>cyl</code>	focus	-2.295	0.101	fragile
<code>carb</code>	focus	-2.197	0.358	fragile
<code>disp</code>	focus	-0.034	0.009	fragile
<code>hp</code>	focus	-0.052	0.002	fragile

Finally, the text output includes results from Sala-i-Martin’s extreme bounds analysis. The histograms in Figure 2 suggest that the normally distributed approximation of the regression coefficients’ distribution does not provide a good fit to the data. For this reason, I focus on EBA results from the generic model, which does not assume any particular distribution of coefficient estimates across different specifications.

As the table presented below indicates, results from Sala-i-Martin’s EBA suggest very little fragility of the coefficient estimates. For variables `wt`, `cyl` and `hp`, more than 99 percent of the cumulative distribution of regression coefficients lies below zero. The same can be said of more than 95 percent of the cumulative distributions for variables `carb` and `disp`. According to results from Sala-i-Martin’s EBA, all of the free and focus variables appear to be robustly (and negatively) associated with the automobiles’ miles per gallon (`mpg`). In contrast to Leamer’s EBA in which a single insignificant coefficient implies fragility, the less stringent Sala-i-Martin’s EBA classifies more variables as robust determinants of fuel economy.

Sala-i-Martin's Extreme Bounds Analysis (EBA):

	Type	G: CDF(beta <= 0)	G: CDF(beta > 0)
(Intercept)	free	2.756	97.244
<code>wt</code>	free	99.957	0.043
<code>cyl</code>	focus	99.521	0.479
<code>carb</code>	focus	95.315	4.685
<code>disp</code>	focus	95.200	4.800
<code>hp</code>	focus	99.047	0.953

All in all, the extreme bounds analysis in this empirical example suggests that the examined automobiles’ weight and engine performance are robustly associated with the vehicles’ fuel economy. In conducting the analysis, I have demonstrated some of the most important features of the **ExtremeBounds** package. Specifically, I have shown how researchers can estimate Leamer’s and Sala-i-Martin’s EBA using the fully customizable `eba()` function. The estimation results can then be visually inspected with the help of the `hist()` method that produces

histograms of each examined variable's coefficient estimates. Finally, the `print()` method provides users with text output that contains a detailed summary of EBA estimation results. In the next section, I conclude.

5. Concluding remarks

In this paper, I have introduced the **ExtremeBounds** package for the R statistical programming language. The package allows researchers to perform extreme bounds analysis (EBA), a sensitivity test that calculates how robustly a regression model's dependent variable is associated with a variety of possible determinants.

ExtremeBounds represents a significant improvement over existing software implementations of extreme bounds analysis, as it supports not only Leamer's version of extreme bounds analysis, but also Sala-i-Martin's EBA. Furthermore, the package allows users to customize every aspect of the analysis: the type of regression model, its size and functional form, as well as the standard errors and weights.

I have showcased many of these customizable features through an empirical example that focused on the determinants of selected automobiles' fuel economy. Along the way, I have also demonstrated the package's ability to produce histograms of estimated regression coefficients via the `hist()` method, and to print out a detailed text summary of the EBA estimation results through the `print()` method.

Acknowledgments

I would like to thank Soo Jin Cho, Rebecca Goldstein, Daniel Yew Mao Lim and Christopher Lucas for helpful comments and suggestions.

References

- Achen CH (2005). "Let's Put Garbage-Can Regressions and Garbage-Can Probits Where They Belong." *Conflict Management and Peace Science*, **22**(4), 327–339.
- Angrist JD, Pischke J (2010). "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics." *Journal of Economic Perspectives*, **24**(2), 3–30.
- Aptech Systems (2013). *GAUSS Statistical Software: Release 14*. Aptech Systems, Inc., Chandler, AZ, USA. URL <http://www.aptech.com>.
- Bailey D, Katz JN (2011). "Implementing Panel-Corrected Standard Errors in R: The `pcse` Package." *Journal of Statistical Software*, **42**(1), 1–11. URL <http://www.jstatsoft.org/v42/c01/>.
- Bartley WA, Cohen MA (1998). "The Effect of Concealed Weapons Laws: An Extreme Bounds Analysis." *Economic Inquiry*, **36**(2), 258–265.

- Beck N, Katz JN (1995). “What to Do (and Not to Do) with Time-Series Cross-Section Data.” *American Political Science Review*, **89**, 634–647.
- Bjørnskov C, Dreher A, Fischer JAV (2008). “Cross-country Determinants of Life Satisfaction: Exploring Different Determinants Across Groups in Society.” *Social Choice and Welfare*, **30**, 119–173.
- Breusch TS (1990). *Modelling Economic Series*, chapter Simplified Extreme Bounds. Oxford University Press, New York, NY, USA.
- Comprehensive R Archive Network (CRAN) (2014). “Comprehensive R Archive Network (CRAN).” URL <http://cran.us.r-project.org/>.
- Doan T (2004). “**EBA**: RATS Procedure to Perform Extreme Bounds Analysis.” *Statistical Software Components*.
- Eicker F (1967). “Limit Theorems for Regression with Unequal and Dependent Errors.” *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, pp. 59–82.
- Estima (2013). *RATS Statistical Software: Release 8.3*. Estima, Evanston, IL, USA. URL <http://www.estima.com/>.
- Fowles R (1988). “Micro Extreme Bounds Analysis.” *The American Statistician*, **42**(4), 274.
- Fowles R, Merva M (1996). “Wage Inequality and Criminal Activity: An Extreme Bounds Analysis for the United States, 1975-1990.” *Criminology*, **34**(2), 163–182.
- Froot KA (1989). “Consistent Covariance Matrix Estimation with Cross-Sectional Dependence and Heteroskedasticity in Financial Data.” *Journal of Financial and Quantitative Analysis*, **24**, 333–355.
- Gassebner M, Lamla MJ, Vreeland JR (2013). “Extreme Bounds of Democracy.” *Journal of Conflict Resolution*, **57**(2), 171–197.
- Granger CW, Uhlig HF (1990). “Reasonable Extreme Bounds Analysis.” *Journal of Econometrics*, **44**, 159–170.
- Hafner-Burton EM (2005). “Right or Robust? The Sensitive Nature of Repression to Globalization.” *Journal of Peace Research*, **42**(6), 679–698.
- Hegre H, Sambanis N (2006). “Sensitivity Analysis of Empirical Results on Civil War Onset.” *Journal of Conflict Resolution*, **50**(4), 508–535.
- Henderson HV, Velleman PF (1981). “Building Multiple Regression Models Interactively.” *Biometrics*, **37**(2), 391–411.
- Hendry DF, Krolzig HM (2004). “We Ran One Regression.” *Oxford Bulletin of Economics and Statistics*, **66**(5), 799–810.
- Hocking R (1976). “The Analysis and Selection of Variables in Linear Regression.” *Biometrics*, **32**(1), 1–49.

- Huber PJ (1967). “The Behavior of Maximum Likelihood Estimates Under Nonstandard Conditions.” *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, pp. 221–233.
- Impavido G (1998). “EBA: Stata Module to Perform Extreme Bound Analysis.” *Statistical Software Components*.
- Leamer EE (1985). “Sensitivity Analysis Would Help.” *American Economic Review*, **57**(3), 308–313.
- Leamer EE, Leonard H (1983). “Reporting the Fragility of Regression Estimates.” *Review of Economics and Statistics*, **65**(2), 306–317.
- Levine R, Renelt D (1992). “A Sensitivity Analysis of Cross-Country Growth Regressions.” *American Economic Review*, **82**(4), 942–963.
- Levine R, Zervos SJ (1993). “What We Have Learned About Policy and Growth from Cross-Country Regressions?” *The American Economic Review*, **83**(2), 426–430.
- Mansfield ER, Helms BP (1982). “Detecting Multicollinearity.” *The American Statistician*, **36**(3), 158–160.
- McAleer M, Pagan AR, Volker PA (1985). “What Will Take the Con Out of Econometrics?” *American Economic Review*, **75**(3), 293–307.
- McFadden DL (1974). *Frontiers in Econometrics*, chapter Conditional Logit Analysis of Qualitative Choice Behavior. Academic Press, New York, NY, USA.
- Moosa IA, Cardak BA (2006). “The Determinants of Foreign Direct Investment: An Extreme Bounds Analysis.” *Journal of Multinational Financial Management*, **16**, 199–211.
- Moosa IA, Smith L (2004). “Economic Development Indicators as Determinants of Medal Winning at the Sydney Olympics: An Extreme Bounds Analysis.” *Australian Economic Papers*, **43**(3), 288–301.
- Moser C, Sturm JE (2011). “Explaining IMF Lending Decisions After the Cold War.” *Review of International Organizations*, **6**(2–3), 307–340.
- Newey WK, West KD (1987). “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix.” *Econometrica*, **55**(3), 703–708.
- R Core Team (2014). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org>.
- Reed WR (2009). “The Determinants of U.S. State Economic Growth: A Less Extreme Bounds Analysis.” *Economic Inquiry*, **47**(4), 685–700.
- Sala-i-Martin X (1997). “I Just Ran Two Million Regressions.” *American Economic Review*, **87**(2), 178–183.
- Sala-i-Martin X, Doppelhofer G, Miller RI (2004). “Determinants of Long-Term Growth: A Bayesian Averaging of Classical Estimates (BACE) Approach.” *American Economic Review*, **94**(4), 813–835.

- SHAZAM Analytics (2013). *SHAZAM Statistical Software: Release 11*. SHAZAM Analytics Ltd, London, UK. URL <http://www.econometrics.com>.
- Shiba T (1992). “MICRO-EBA: Leamer’s Extreme Bounds Analysis on Gauss.” *Journal of Applied Econometrics*, **7**(1), 101–103.
- StataCorp (2013). *Stata Statistical Software: Release 13*. StataCorp LP, College Station, TX, USA. URL <http://www.stata.com>.
- Sturm JE, de Haan J (2005). “Determinants of Long-Term Growth: New Results Applying Robust Estimation and Extreme Bounds Analysis.” *Empirical Economics*, **30**(3), 597–617.
- Wang EC (2010). “Determinants of R&D Investment: The Extreme-Bounds-Analysis Approach Applied to 26 OECD Countries.” *Research Policy*, **39**(1), 103–116.
- White H (1980). “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, **48**, 817–838.
- Williams RL (2000). “A Note on Robust Variance Estimation for Cluster-Correlated Data.” *Biometrics*, **56**, 645–646.
- Zeileis A (2004). “Econometric Computing with HC and HAC Covariance Matrix Estimators.” *Journal of Statistical Software*, **11**(10), 1–17. URL <http://www.jstatsoft.org/v11/i10/>.
- Zeileis A (2006). “Object-Oriented Computation of Sandwich Estimators.” *Journal of Statistical Software*, **16**(9), 1–16. URL <http://www.jstatsoft.org/v16/i09/>.
- Zeileis A, Croissant Y (2010). “Extended Model Formulas in R: Multiple Parts and Multiple Responses.” *Journal of Statistical Software*, **34**(1), 1–13. URL <http://www.jstatsoft.org/v34/i01/>.

Affiliation:

Marek Hlavac
Harvard University
John F. Kennedy School of Government
79 John F. Kennedy Street
Cambridge, MA 02138
United States of America
E-mail: hlavac@fas.harvard.edu