WORKING PAPERS

N° 1731

April 2026

“Robust Frontier and Efficiency Analysis with frontiles”

Abdelaati Daouia and Thibault Laurent

Toulouse School of Economics

1

Robust Frontier and Efficiency Analysis with `frontiles`

CONTENTS

This chapter discusses the current state of development of robust measures for evaluating firms' production performance, focusing on two prominent approaches: (i) partial order-m frontiers and related efficiency scores based on probability-weighted moments, and (ii) their competing order-α counterparts, which rely on quantiles. It provides a structured overview of the original concepts and their recently introduced robustified versions, analyzing their strengths and weaknesses in terms of axiomatic properties, estimation methods, and robustness. The `frontiles` package offers various functions for computing both order-α and order-m frontiers and efficiency scores, including their robustified analogs, in the general setting with multiple inputs and outputs. It supports different performance measurement directions, namely input, output and hyperbolic orientations. Additionally, `frontiles` includes procedures for inference and robustness assessment, notably through confidence intervals, gross-error sensitivity and breakdown points. It also provides diagnostic checks to assess the presence of outliers in the data and, accordingly, to guide the choice of suitable trimming levels. It further enables the visualization of robust surface estimators in three-dimensional settings involving two inputs and one output. The use of this package is illustrated with a number of empirical applications.

1.1 Introduction

In production theory [29], the set of attainable combinations of inputs $x \in \mathbb{R}^p_+$ and outputs $y \in \mathbb{R}^q_+$ is interpreted as

$$\Psi := \{(x, y) \in \mathbb{R}^{p+q} |\ x \text{ can produce } y\}.$$

The upper extremity of this set, which defines the production frontier and represents the set of the most efficient firms, can be described as

$$\Psi^\partial := \{(x, y) \in \mathbb{R}^{p+q} | (\gamma^{-1}x, \gamma y) \notin \Psi, \forall \gamma > 1\}.$$

A minimal assumption on Ψ is the free disposability of both inputs and outputs, meaning that if $(x, y) \in \Psi$, then $(x', y') \in \Psi$ for any (x', y') such that $x' \geqslant x$ and $y' \leqslant y$. This entails the monotonicity of the boundary Ψ^∂. The performance of a firm operating at $(x, y) \in \Psi$ can be assessed by the distance to its projection on this efficient boundary. When an output-oriented perspective is adopted, this distance yields the performance measure for the firm defined as

$$\lambda(x, y) := \sup \{\lambda > 0 \mid (x, \lambda y) \in \Psi\}.$$

In the input-orientation, efficiency is measured by

$$\theta(x, y) := \inf\{\theta > 0 | (\theta x, y) \in \Psi\},$$

while the hyperbolic measure of efficiency is given by

$$\gamma(x, y) := \sup\{\gamma > 0 | (\gamma^{-1}x, \gamma y) \in \Psi\}.$$

Assume that the input-usage of firms is generated by a random vector $X \in \mathbb{R}^p_+$ to be transformed into a random vector of output quantities $Y \in \mathbb{R}^q_+$, and let $(\Omega, \mathcal{A}, \mathbb{P})$ be the probability space on which the random pair (X, Y) is defined. The support of (X, Y) can be viewed as the feasible production set Ψ, and its joint distribution can be characterized, at a given production plan (x, y), by the probability

$$H_{XY}(x, y) := \mathbb{P}(X \leqslant x, Y \geqslant y)$$

that the firm operating at (x, y) is dominated by another firm using fewer inputs and producing more outputs. Under the free disposability assumption, the efficiency scores described above can then be defined equivalently as

$$\begin{aligned}
\lambda(x, y) &= \sup\{\lambda > 0 | H_{XY}(x, \lambda y) > 0\}, \\
\theta(x, y) &= \inf\{\theta > 0 | H_{XY}(\theta x, y) > 0\}, \\
\gamma(x, y) &= \sup\{\gamma > 0 | H_{XY}(\gamma^{-1}x, \gamma y) > 0\}.
\end{aligned}$$

In practice, these measures are unknown and must be estimated from a sample of observations $\mathcal{X}_n = \{(X_i, Y_i)\}_{i=1}^n$. The `frontiles` package [10] provides a collection of functions implementing nonparametric estimation techniques of particular interest in the statistical software environment R. Natural estimators follow by replacing $H_{XY}(x, y)$ with its empirical analog

$$\widehat{H}_{n,XY}(x, y) := n^{-1} \sum_{i=1}^{n} \mathbb{1}(X_i \leqslant x, Y_i \geqslant y),$$

with $\mathbb{1}(\cdot)$ being the indicator function, to get the classical Free Disposal Hull (FDH) estimators [22]; see *e.g.* [19] for a comprehensive presentation of these envelopment estimators and their full asymptotic properties in a general setup from the perspective of extreme value theory. Beyond providing computational implementations of these empirical estimators, `frontiles` makes a first practical contribution by offering dedicated tools for their graphical exploration. In particular, it enables the visualization of the associated FDH frontiers in three-dimensional settings, thereby facilitating both interpretation and empirical analysis.

When Ψ is assumed to be convex, it can be estimated by the convex hull of the FDH estimator [3], known as the Data Envelopment Analysis (DEA) estimator. In the context of frontier analysis for nonconvex Ψ, the linearly interpolated FDH (LFDH) frontier provides an enhanced variant of the FDH estimator [26]. Despite the fact that the FDH, LFDH and DEA estimators yield fitted boundaries of Ψ that satisfy monotonicity (and, for DEA, monotone concavity), they tend to be overly rough: they undersmooth the data and underestimate the true production frontier. To overcome these drawbacks, several alternative estimation approaches have been developed with the goal of producing smoother and more accurate frontier estimates while preserving the fundamental shape constraints. The survey article [11] provides an extensive bibliography of the best-known methods, and the `npbr` package [12] is, to our knowledge, the first freely available specialized software offering a broad suite of tools to implement many of these approaches. Among them, a particularly innovative proposal is due to [14], who combine spline smoothing with constrained estimation.

Most existing empirical and smooth estimators for frontier and efficiency analysis are, by construction, envelope-based and are therefore highly non-robust. To mitigate the influence of extreme values and outliers, the literature has developed two widely used robust approaches. Both proceed by first estimating a *partial frontier* (or an associated *partial efficiency measure*) lying well inside the data cloud, and then translating the resulting anchor estimate to recover an estimate of the full boundary Ψ^{∂} (or the corresponding *full efficiency measure*). In the output-orientation for instance, the first approach introduces a concept of *order-m* frontier, defined as the expected maximal output attainable among m (with $m \geqslant 1$) firms randomly drawn from those using no more than a given input level [2]. The second proposes the concept of *order-α* frontier (with $\alpha \in (0, 1]$), defined through a high-level conditional α-quantile of the output distribution [1, 18]. The former is determined by weighted expectations of Y given $X \leqslant x$, and its associated *order-m* efficiency scores exploit the data fully by incorporating distances to all observations, whereas the latter is determined by tail probabilities of Y given $X \leqslant x$, and its associated *order-α* efficiency scores rely on the frequency of tail observations. A similar approach can be adapted to the input-orientation by working with conditional expected minima or low-level quantiles of the cost X given $Y \geqslant y$. Because they are less extreme than the full frontier Ψ^{∂}, both order-m and order-α frontiers provide robust benchmarks for performance measurement. They also have a meaningful economic interpretation on their own and can be extrapolated to estimate the true frontier [5, 6, 9]. The review chapter 4 of [20] gives a nice summary of their theory and an extensive bibliography. Robustness and asymptotic comparisons between the two concepts are given in [8], [16] and [17].

The `frontiles` package provides a range of functions to compute the estimated order-m and order-α efficiency scores with multiple inputs and outputs, and to visualize the corresponding partial frontiers in three-dimensional settings. Similarly to the `FEAR` package [34], exact estimates are provided for the input- and output-oriented α-efficiency measures through the closed-form expressions established in [18]. The package also implements a fast exact computation of the hyperbolic α-efficiency scores based on the new explicit expression derived in [19], without resorting to the numerical bisection method proposed by

[32]. Likewise, `frontiles` implements the Monte Carlo approximation method of [2] for the computation of the estimated m-efficiency scores. Moreover, in contrast to `FEAR` which computes the exact order-m efficiency estimates suggested by [21], `frontiles` implements the exact L-statistic estimates for these scores derived in [8]. Chapter **??** in this volume provides a practical overview of the corresponding `FEAR` routines for robust partial frontiers and beyond, while also listing other related packages.

The use of the nonstandard conditional distribution of Y given $X \leqslant x$ in the output-orientation (or X given $Y \geqslant y$ in the input-orientation) is motivated by econometric considerations of tail monotonicity, but the resulting partial frontier estimators may exhibit disappointing behavior. In the output direction for example, both population order-m and order-α frontiers fulfill the monotonicity requirement, implied by the free disposability property for all trimming levels m and α, if and only if the distribution function of Y given $X \leqslant x$ is nonincreasing in x [2, 18]. The latter hypothesis, referred to as *tail monotonicity* of the conditional distribution, is natural in production theory, as it implies that the chance of producing less than some outputs decreases when firms use more inputs [2]. However, in finite samples, nothing guarantees tail monotonicity of the empirical conditional distribution, so the associated empirical order-m and order-α frontiers may not inherit the monotonicity enjoyed by their population counterparts. Moreover, both population partial frontiers and their empirical versions, including the monotonized estimators suggested in [15], tend to drift away from the true frontier Ψ^{∂} as the input level increases (resp. decreases) in the output (resp. input) orientation, and thus may fail to serve as reliable benchmark frontiers for large (resp. small) input values [7, 19]. A converse drawback is that partial frontier estimators tend to move ever closer to the non-robust FDH estimator as the input level decreases (resp. increases) in the output (resp. input) orientation, and may therefore become increasingly sensitive to outliers among firms with low (resp. high) inputs-usage [8]. The root of these limitations is the conditioning on the event $\{X \leqslant x\}$ or $\{Y \geqslant y\}$ used to define the two families of partial frontiers and their associated efficiency scores. Adapting this conditional approach of partial frontier and efficiency analysis to other orientations (*e.g.*, hyperbolic or directional distance functions) is straightforward; see the survey by [30]. Another consequence of adopting this conditional approach is that the resulting order-α (and likewise order-m) frontiers generally depend on the chosen orientation; they coincide across orientations only in the trivial, limiting full-frontier case $m \to \infty$ or $\alpha \to 1$.

A way to circumvent these difficulties was proposed for the quantile-based approach in [19] in the fully multivariate setting $(p, q \geqslant 1)$, through the introduction of *"unconditional"* partial order-α efficiency scores. These new scores converge smoothly to the full efficiency measure as $\alpha \to 1$, while sidestepping the drawbacks of the classical *"conditional"* concept of [2]. The key device is a *dimensionless* transformation of the $(p+q)$-dimensional production process (X, Y), which yields a coherent probabilistic formulation of the distance from any realization (x, y) of (X, Y) to the efficient support surface Ψ^{∂}, in any chosen direction. Related ideas were developed for expected optimal (maximum output/minimum input) production analysis in [7] for the single-output case $(q = 1)$ with multiple inputs $(p \geqslant 1)$, and more recently extended to the multi-input and multi-output framework in [13]. Both the resulting *robustified* order-α and order-m efficiency scores, as well as their empirical counterparts, inherit the same monotonicity properties as the corresponding full efficiency score. Here, `frontiles` is the first package to offer a suite of functions to compute estimated order-α and order-m efficiency scores in multi-input and multi-output settings, along with their associated asymptotic confidence intervals, and to visualize the associated partial frontiers in three-dimensional representations.

This chapter is structured as follows. Section 1.2 further motivates the consideration of quantile-type and expected maximum-output frontiers through three unrelated motivating data examples concerned with movie franchise finance (production budgets X vs. lifetime

gross earnings Y), top male tennis players (match wins X vs. ranking points Y), and microfinance institutions (operating expenses and personnel X vs. revenue Y). Then, Sections 1.3 and 1.4 review the respective approaches of (un)conditional quantile-based α-efficiency measures and (un)conditional weighted-moment m-efficiency measures. Section 1.5 compares the robustness of estimated order-α and order-m efficiency scores via local sensitivity (influence-type) and global breakdown analyses, then it provides routines to compute these measures, and offers practical guidance for selecting trimming levels when building diagnostic tools. Finally, Section 1.6 concludes.

1.2 Empirical motivation

Throughout this chapter, which is mostly of reviewing nature about robust estimation in frontier and efficiency analysis, we illustrate the use of the `frontiles` package through the following three datasets:

ATP 2017 Men: Match Wins vs. Ranking Points

The dataset "`atp`" from the `csr` package reports, for the top male tennis players in 2017, the total number of match wins and the corresponding ATP ranking points. ATP points are awarded for wins in official tour events (including Grand Slams), but the point value of a win depends strongly on the tournament and round. For instance, reaching a Grand Slam semi-final is worth 1200 points for the tournament, while winning the final adds 800 more. As a result, the association between wins and ranking points is strongly positive yet distinctly non-linear: a player may record fewer wins than another while accumulating more points by succeeding in higher-stakes matches. In this scoring structure, viewing the number of match wins as an input (X) and the corresponding ranking points as an output (Y), the joint support envelope can be interpreted as an efficient frontier, and directional (vertical, horizontal, or hyperbolic) shortfalls to this frontier as oriented efficiency scores. As shown in Figure 1.1(a), the presence of two dominant outliers (Nadal and Federer) makes a concavity assumption potentially misleading, as it can yield an implausible classification. For instance, it would then label the mid-ranked players (those with roughly 4000 to 6000 points) as highly inefficient. A monotone non-concave frontier provides a markedly tighter and more realistic envelopment there, but it also develops a pronounced change in slope around 50 wins, followed by a very steep segment. Using robust, partial (un)conditional frontier methods, one can obtain meaningful fits while leaving the two super-extreme top performers (Nadal and Federer) outside the estimated efficient frontier.

Blockbuster Franchise Movies: Budget vs. Lifetime Gross

The dataset "`movies`" provides basic financial characteristics of a collection of movie projects from major film franchises, originally compiled by [4]. In the `frontiles` package we retain a streamlined set of variables: the movie title, release year, runtime, total lifetime gross earnings, and production budget. The last two variables are particularly relevant for frontier analysis: the production budget can be viewed as an input (X), while lifetime gross earnings serve as an output (Y). The joint support boundary of earnings versus budget can be viewed as a production, profit, or cost frontier, with vertical, horizontal, or hyperbolic shortfalls to this boundary naturally yielding output-, input-, or hyperbolic-oriented efficiency scores. Imposing monotonicity is economically sensible, since higher production budgets should, in principle, not reduce box-office performance. As visualized in Figure 1.1(b), if the data call

for a non-concave global boundary, concavity still seems a reasonable approximation for most films with small to mid-range budgets.

Microfinance Institutions: Operating Expenses and Personnel vs. Revenue

The dataset "`mfi`" is an updated and extended version of the microfinance data studied by [24], originally drawn from the Microfinance Information Exchange (MIX) database and covering the years 2010–2018. We take revenue (in USD) as the output (Y), while the main inputs are personnel (number of employees X^1) and operating expenses (X^2 including personnel costs, depreciation, amortization, and administrative expenses). The cleaned sample retains 4833 observations out of 4854 institutions. Viewed through the lens of frontier analysis, the extreme boundary of revenue as a function of inputs can be interpreted as a production frontier, and directional shortfalls to this frontier as the corresponding efficiency scores. As illustrated in Figure 1.2 in the 3-dimensional framework, a key difficulty is that relatively few very large microfinance institutions populate the high-input region, so a sensible benchmark there requires robust methods that remain stable under sparsity.

Although the package is available on CRAN, we recommend installing the most recent development version from GitHub using the following code:

```
if (!require("devtools")) install.packages("devtools")
   devtools::install_github("tibo31/frontiles")
```

The code and examples presented in this chapter are available in the `README.md` file on the package's GitHub repository: `https://github.com/tibo31/frontiles`. To help users navigate the functionalities implemented in the `frontiles` package, Table 1.1 provides an overview of the estimation methods and their associated shape constraints, while Table 1.2 presents the main plotting functions and their key options. Table 1.3 then summarizes the robustness diagnostic routines available in the package, together with their purpose.

TABLE 1.1
Characteristics of the estimation methods in `frontiles`.

Function	Type of estimator	Allowed constraints
`alphascore`	Robust/envelope, conditional quantile	Unconstrained
`fdhscore`	Envelope, piecewise constant	Monotonicity
`ordermscore`	Robust, conditional weighted moment	Unconstrained
`ordermscore.boot`	Robust, conditional weighted moment	Unconstrained
`alphascore.star`	Robustified/envelope, unconditional quantile	Monotonicity
`ordermscore.star`	Robustified, unconditional weighted moment	Monotonicity

All estimation routines in `frontiles` share a common input–output structure. The first two arguments correspond to the observed inputs $X_1, \ldots, X_n$ (argument `xobs`) and the observed outputs $Y_1, \ldots, Y_n$ (argument `yobs`), which define the reference sample used to construct the production set. The third and fourth arguments are numeric vectors of evaluation points, inputs `xeval` and outputs `yeval`, at which the frontier estimator and the associated efficiency measure are to be computed. These evaluation points may coincide with the observed data points or correspond to new production plans whose performance is to be assessed relative to the estimated technology. The functions return a matrix with rows corresponding to the evaluation points and columns reporting the efficiency measures under the output, input, and hyperbolic orientations. Additional arguments may be required depending on the estimation method used.

TABLE 1.2
Graphical routines available in `frontiles`.

Function	Setting/Estimators/Arguments	Graphical output
`plot_frontier.2d()`	Bivariate (1 input, 1 output)/FDH frontier, m-frontiers, α-frontiers/ orientation, m, α, swap_axes, CI	Frontier curve in the plane
`plot_frontier.3d()`	Trivariate (2 inputs, 1 output)/FDH frontier, m-frontiers, α-frontiers/ panels, orientation, m, α	Surface plot or 2D choropleth/map; input-oriented isoquants
`diagnostic.trim()`	General/m-frontiers, α-frontiers/ `type = "alphascore"`, ...	Robustness diagnostic curves to help choose trimming levels

TABLE 1.3
Robustness diagnostic routines available in `frontiles`.

Function	Framework/arguments	Purpose
`gross.error.sensitivity()`	Robust, conditional/ type, α or m	Computes the estimated gross-error sensitivity
`gross.error.sensitivity .star()`	Robustified, unconditional /type, α or m	Computes the estimated gross-error sensitivity
`breakdown.point()`	Conditional quantile/α	Computes the finite-sample replacement breakdown point
`breakdown.point.star()`	Unconditional quantile/α	Computes the finite-sample replacement breakdown point

The use of our estimation routines is illustrated below via the `fdhscore` function. To demonstrate the computation of the classical FDH estimator in the output-, input-, and hyperbolic orientations, we employ the `atp` dataset. The reference sample is given by the observed inputs `xobs` (number of match wins) and observed outputs `yobs` (ranking points). We consider two evaluation points, denoted P1 and P2, shown in Figure 1.1(a) and representing two fictive new players whose efficiency is to be assessed. The function `fdhscore()` follows the standard interface of the package and returns, for each evaluation point and each direction, the corresponding oriented efficiency measure. The evaluation points are specified through the arguments `xeval` and `yeval`; if these are not provided, the function computes the efficiency measures for all observed units by default.[1]

```
library("frontiles")
data("atp")

xobs <- atp$wins
yobs <- atp$ranking_points

# Two new players (evaluation points only)
x_new <- c(46, 50)
```

[1]From an implementation perspective, the routine distinguishes between the general multivariate case with multiple inputs and outputs and the bivariate case considered here with one-input and one-output. In the latter setting, a dedicated and computationally efficient procedure is used, exploiting the stepwise structure of the FDH frontier. In the general case, FDH scores are obtained by calling `alphascore()` with $\alpha = 1$.

```
y_new <- c(3300, 6000)

fdh_atp_new <- fdhscore(xobs  = matrix(xobs), yobs  = matrix(yobs),
   xeval = matrix(x_new), yeval = matrix(y_new))
```

```
       output      input      hyper
P1   1.1530303 0.9347826 0.9347826
P2   0.8583333 1.0600000 1.0600000
```

The matrix returned by `fdhscore()` reports the output-, input-, and hyperbolic-oriented efficiency measures for each evaluation point. For P1, the output score equals 1.153, while the input and hyperbolic scores equal 0.935. For P2, the output score is 0.858, while the input and hyperbolic scores equal 1.060. This reflects a position above the frontier in the output direction. The FDH frontier can be represented graphically using the function `plot_frontier.2d()`. In the present case, we use:

```
plot_frontier.2d(input, output, type_frontier = "fdh",
   lty = 1, col = "darkgrey", add = TRUE, lwd = 1.8,
   xlab = "Number of match wins (input)",  ylab = "Ranking points (output)")
```

The resulting frontier is the same regardless of the chosen orientation (output, input or hyperbolic). What changes is the direction along which efficiency is evaluated. The function therefore computes the same boundary in the (x, y) plane and adapts only the projection mechanism when needed.[2] The same plotting routine is used for the other frontier estimators implemented in the package, ensuring a unified graphical representation across methods.

Panel (a) of Figure 1.1 presents a graphical illustration of the FDH frontier for the ATP data. It displays the scatterplot of observed players, where the input corresponds to the number of match wins and the output to ranking points. FDH-players who determine the upper envelope of the sample are highlighted in red. The corresponding FDH frontier is graphed using `plot_frontier.2d()`. The two additional points P1 and P2, shown in magenta, represent the evaluation points used for our illustrative purpose. For each of them, three projection directions are drawn: a vertical segment corresponding to the output-oriented projection, a horizontal segment corresponding to the input-oriented projection, and a diagonal segment corresponding to the hyperbolic projection. The intersection of each segment with the frontier identifies the target FDH projection. Panel (b) of Figure 1.1 replicates the same visualization strategy for the movie dataset, showing that the same plotting routine can be used for different datasets and frontier estimators in a unified manner.

We next apply the same methodology to the 3-dimensional microfinance dataset. The estimation procedure remains conceptually unchanged; the only difference is that we now consider two inputs and one output. The input variables are personnel and operating expenses, while the output variable is financial revenue. All variables are rescaled for numerical stability and readability (personnel in thousands, monetary variables in billions of USD). Observations with missing values or non-positive outputs are excluded. The data can be visualized in three dimensions using the `scatterplot3d` package, as shown in Figure 1.2(a).

```
data(mfi)
mfi <- na.omit(mfi[, c("personnel", "operatingexpense", "financialrevenue")])
mfi <- mfi[mfi$financialrevenue > 0, ]
```

[2]Internally, `plot_frontier.2d()` first constructs an augmented grid over the support of the data and then evaluates the frontier on this grid. In the FDH case, it does so by calling `fdhscore()` on a sequence of evaluation points to recover the stepwise upper envelope; the resulting coordinates are then joined to produce the frontier curve.

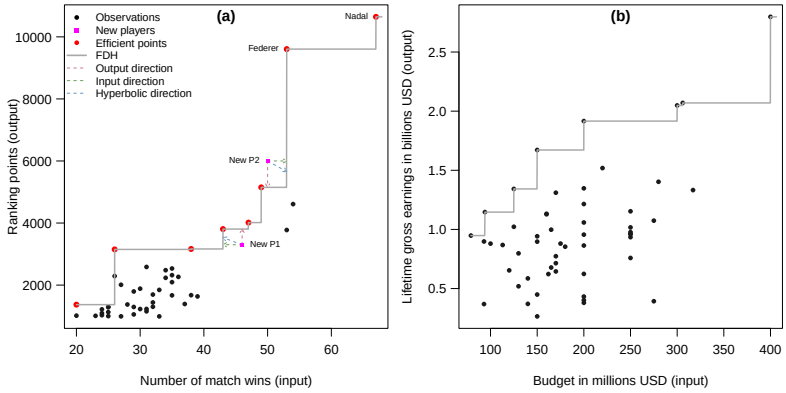

FIGURE 1.1
Scatterplots of the data with the FDH frontier. Panel (a) shows the ATP 2017 dataset, where efficient FDH observations are highlighted in red and the new evaluation points (P1 and P2) are displayed in magenta, together with their output-, input-, and hyperbolic-oriented projections onto the FDH frontier. Panel (b) presents the Blockbuster Franchise Movies dataset.

```
X <- cbind(mfi[,1] / 10^3, mfi[,2] / 10^9)
Y <- as.matrix(mfi[, 3] / 10^9)
library(scatterplot3d)
scatterplot3d(x1, x2, y)
```

We first compute the efficiency scores for the observed units in order to identify the FDH firms that characterize the FDH surface:

```
fdh_mfi_obs <- fdhscore(xobs = X, yobs = Y)
```

To visualize the frontier in three dimensions, we evaluate efficiency scores on a grid over the input space, spanning the range of efficient input combinations.[3] The `plot_frontier.3d()` function has been designed to automate this entire visualization workflow. Given a reference sample (X, Y) and a chosen estimator (here `type_frontier = "fdh"`), it constructs the evaluation grid, computes the corresponding efficiency measures, reconstructs the implied boundary output levels, and produces the requested graphical display. Depending on the selected panel, the estimated frontier can be shown either as a perspective surface or as a 2-dimensional projection onto the input space:

```
plot_frontier.3d(X, Y, type_frontier = "fdh", panels = "surface")
plot_frontier.3d(X, Y, type_frontier = "fdh", panels = "map")
```

Figure 1.2(b–c) displays the FDH estimates for the microfinance dataset. Panel (b) shows the FDH frontier as a perspective surface over the input grid, highlighting its characteristic stepwise structure and the boundary-efficient observations (shown in red). Panel (c) provides a 2-dimensional choropleth projection of the estimated frontier onto the input space. In this representation, color intensity reflects the magnitude of the estimated attainable output; selecting an appropriate color scale (*e.g.*, via `colorRampPalette`) is important to ensure clear visual separation of plateau regions and discontinuities.

[3] In practice, the grid extends beyond the observed points and includes intermediate values between efficient units. This refinement is important because the FDH estimator is piecewise constant: the frontier exhibits flat regions (plateaus) separated by vertical or horizontal jumps. Including intermediate grid points ensures that these stepwise transitions are properly captured in the graphical representation.

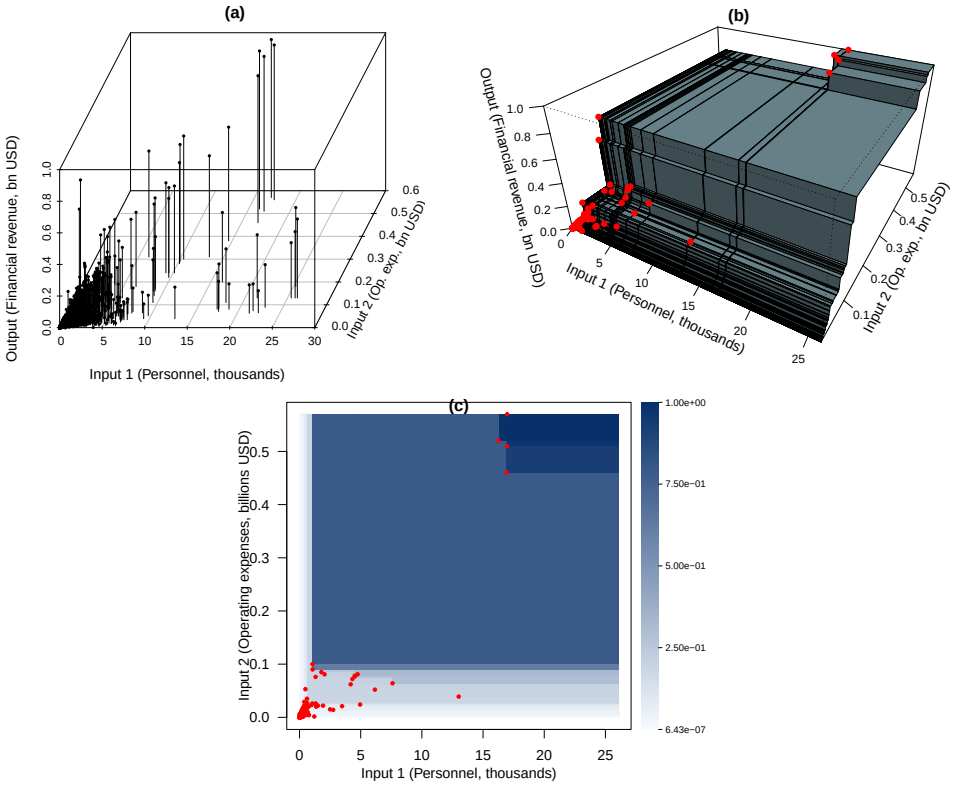

FIGURE 1.2
Microfinance institutions data. Panel (a) displays the three-dimensional scatterplot of the observations (X_i^1, X_i^2, Y_i). Panel (b) presents the estimated FDH production frontier as a perspective surface over the input space. Panel (c) provides a two-dimensional choropleth representation of the same FDH frontier, where color intensity increases from light to dark as the estimated attainable output rises. Efficient firms are highlighted in red in each panel.

1.3 Quantile-type partial efficiency scores

With multiple inputs $X \in \mathbb{R}^p_+$ and outputs $Y \in \mathbb{R}^q_+$, standard quantile regression of, for instance, Y conditional on $X = x$ (or X conditional on $Y = y$) is ill-suited to measuring economic efficiency in an output- (or input-) oriented sense, because $\mathbb{R}^q$ (or $\mathbb{R}^p$) lacks a natural total ordering as soon as $q > 1$ (or $p > 1$). To circumvent this difficulty, [18] proposed nonstandard *conditional* quantiles by conditioning on events of the form $X \leqslant x$ (output direction) or $Y \geqslant y$ (input direction). The resulting quantiles, however, depend intrinsically on the direction chosen *a priori* and might fail to satisfy basic monotonicity axioms. Section 1.3.1 reviews these quantile-type efficiency measures and offers practical guidance for their implementation. To address the limitations of the conditional quantile-based approach, [19] proposed an alternative construction based on *unconditional* quantiles, described in Section 1.3.2.

1.3.1 Robust α-efficiency measures

In what follows, we briefly outline the principle of benchmarking against conditional quantile-based frontiers in the *output orientation*; the corresponding input, hyperbolic, and directional versions are obtained by straightforward adaptations (see the survey in [30]).

Rather than targeting the full efficiency measure $\lambda(x, y)$ itself, [18] generalize the "univariate" construction of [1] and define the order-α partial efficiency score

$$\lambda_\alpha(x, y) = \sup\left\{\lambda > 0 \;\middle|\; \frac{H_{XY}(x, \lambda y)}{H_{XY}(x, 0)} > 1 - \alpha\right\}$$

for $\alpha \in (0, 1]$. The corresponding order-α output-oriented frontier can be characterized by

$$\Psi^\partial_\alpha = \{(x, y\,\lambda_\alpha(x, y))|(x, y) \in \Psi\} = \big\{(x, y) \in \Psi \,\big|\, \lambda_\alpha(x, y) = 1\big\}.$$

We have $\lambda_\alpha(x, y) \to \lambda_1(x, y) \equiv \lambda(x, y)$ as $\alpha \to 1$. Therefore, choosing α close to 1 yields a partial score $\lambda_\alpha(x, y)$ that closely approximates the full efficiency measure $\lambda(x, y)$. The empirical counterpart of $\lambda_\alpha(x, y)$ is

$$\widehat{\lambda}_\alpha(x, y) = \sup\left\{\lambda > 0 \;\middle|\; \frac{\widehat{H}_{n,XY}(x, \lambda y)}{\widehat{H}_{n,XY}(x, 0)} > 1 - \alpha\right\},$$

and it is immediate that $\widehat{\lambda}_\alpha(x, y) \to \widehat{\lambda}(x, y)$ as $\alpha \to 1$, for fixed n, where

$$\widehat{\lambda}(x, y) = \sup\Big\{\lambda > 0 \;\Big|\; \widehat{H}_{n,XY}(x, \lambda y) > 0\Big\}$$

stands for the FDH estimator of $\lambda(x, y)$. When the order $\alpha = \alpha(n)$ tends to 1 at the rate $n^{(p+q+1)/(p+q)}(1 - \alpha(n)) \to 0$ as $n \to \infty$, [18] show that $\widehat{\lambda}_\alpha(x, y)$ yields an alternative estimator of $\lambda(x, y)$ with asymptotic behavior akin to that of the FDH estimator (namely, a Weibull limiting distribution and convergence rate $n^{1/(p+q)}$). For finite n, one has $\alpha(n) < 1$, so the associated order-α frontier surface (which converges to the full frontier Ψ^∂) does not envelop all observations, and hence it is typically more robust to extreme points and outliers than the FDH estimator. From the standpoint of robustness theory, the statistical properties of $\widehat{\lambda}_\alpha(x, y)$ have been studied in [8] and [16]. From an axiomatic point of view, [18] show that $\lambda_\alpha(x, y)$ is monotone nondecreasing in x, for every α, if and only if the conditional distribution $F_{Y|X}(y|x) = \mathbb{P}(Y \leqslant y|X \leqslant x)$ satisfies the tail-monotonicity property

$$F_{Y|X}(y|x) \leqslant F_{Y|X}(y|x') \quad \forall (x, y), (x', y) \in \Psi,\ x \geqslant x'.$$

In finite samples, however, the estimator $\widehat{\lambda}_\alpha(x,y)$ does not enjoy this property as illustrated in [19]. This defect has been addressed in [17] for the case $q = 1$ by isotonizing the resulting estimate of the production function. It is also important to note that the corresponding order-α frontier is different from the ones obtained in other chosen orientations [32].

From a computational point of view, it is easy to see that $\widehat{\lambda}_\alpha(x,y)$ is a simple empirical quantile. Let $N_x = n\widehat{H}_{n,XY}(x,0)$ be the number of inputs X_i such that $X_i \leqslant x$, and let $(Y_1(x),\ldots,Y_{N_x}(x))$ be the outputs Y_i such that $X_i \leqslant x$. Then, $\widehat{\lambda}_\alpha(x,y)$ is nothing but the αth quantile of the sample $(Z_1^{xy},\ldots,Z_{N_x}^{xy})$, where

$$Z_i^{xy} = Z^{xy}(X_i,Y_i) := \min_{j=1,\ldots,q}\left(\frac{Y_i^{[j]}(x)}{y^{[j]}}\right) \quad \text{for} \quad i = 1,\ldots,N_x,$$

with the superscript $[j]$ referring to the jth component of the vectors Y and y. More specifically, if $Z_{(1)}^{xy} \leqslant \cdots \leqslant Z_{(N_x)}^{xy}$ are the order statistics of the variables Z_i^{xy}, then

$$\widehat{\lambda}_\alpha(x,y) = Z^{xy}_{(\lfloor \alpha N_x \rfloor + 1)}$$

with $\lfloor \alpha N_x \rfloor$ being the integer part of αN_x.

The `alphascore()` routine in the `frontiles` package allows to compute the estimates $\widehat{\lambda}_\alpha(x,y)$ in the output-orientation for a set of evaluation points (x,y), but also their analogs

$$\widehat{\theta}_\alpha(x,y) = \inf\left\{\theta > 0 \;\middle|\; \frac{\widehat{H}_{n,XY}(\theta x,y)}{\widehat{H}_{n,XY}(\infty,y)} \geqslant 1-\alpha\right\}$$

that estimate the conditional α-quantile efficiency scores

$$\theta_\alpha(x,y) = \inf\left\{\theta > 0 \;\middle|\; \frac{H_{XY}(\theta x,y)}{H_{XY}(\infty,y)} \geqslant 1-\alpha\right\}$$

in the input-orientation [18, 31], as well as

$$\widehat{\gamma}_\alpha(x,y) = \sup\left\{\gamma > 0 \;\middle|\; \widehat{H}_{n,XY}(\gamma^{-1}x,\gamma y) > 1-\alpha\right\}$$

that estimate the unconditional hyperbolic-oriented efficiency scores [32]:

$$\gamma_\alpha(x,y) = \sup\left\{\gamma > 0 \;\middle|\; H_{XY}(\gamma^{-1}x,\gamma y) > 1-\alpha\right\}.$$

It is not hard to verify that the estimator $\widehat{\theta}_\alpha(x,y)$ also coincides with a conditional empirical quantile. Let $N_y = n\widehat{H}_{n,XY}(\infty,y)$ be the number of outputs Y_i such that $Y_i \geqslant y$, and let $(X_1(y),\ldots,X_{N_y}(y))$ be the inputs X_i such that $Y_i \geqslant y$. Then, $\widehat{\theta}_\alpha(x,y)$ is but the $(1-\alpha)$th quantile of the sample points

$$Z_i^{xy} = Z^{xy}(X_i,Y_i) := \max_{j=1,\ldots,p}\left(\frac{X_i^{[j]}(y)}{x^{[j]}}\right) \quad \text{for} \quad i = 1,\ldots,N_y, \tag{1.1}$$

that is, $\widehat{\theta}_\alpha(x,y) = Z^{xy}_{(\lfloor (1-\alpha)N_y \rfloor + 1)}$. As regards $\widehat{\gamma}_\alpha(x,y)$, [32] propose two sophisticated methods for its computation, including an algorithm based on the bisection method. Instead, below in Section 1.3.2, we present a more convenient formulation of this estimator as the usual αth quantile of an explicit sample, together with a simple associated confidence interval.

For our illustration purpose and to facilitate comparison with the results in the previous section, we reconsider the evaluation points P1 and P2 from the `atp` dataset and compute their order-α efficiency scores $\widehat{\lambda}_\alpha(x,y)$, $\widehat{\theta}_\alpha(x,y)$ and $\widehat{\gamma}_\alpha(x,y)$ with $\alpha = 0.975$:

```
alphascore(xobs = matrix(xobs), yobs = matrix(yobs),
    xeval = matrix(x_new), yeval = matrix(y_new), alpha = .975)

          output     input     hyper
P1    0.9590909 0.9347826 1.021739
P2    0.6691667 1.0600000 1.165049
```

Relative to the FDH benchmark where $\alpha = 1$, slight differences emerge when $\alpha = 0.975$, although the overall conclusions remain very close. For P1, the output-oriented score is below one ($\widehat{\lambda}_\alpha = 0.959$), which indicates efficiency in the output direction. The hyperbolic score exceeds one ($\widehat{\gamma}_\alpha = 1.022$), implying efficiency in the hyperbolic orientation. However, the input-oriented score remains below one ($\widehat{\theta}_\alpha = 0.935$), and so P1 does not achieve the desired efficiency in the input direction. In contrast, P2 is efficient in the three directions, since the output-oriented score is below one, while both the input- and hyperbolic-oriented scores exceed one. Overall, when α is chosen close to one, the α-frontier remains very close to the FDH benchmark, while providing a slight smoothing of the frontier and allowing a small proportion of extreme observations to lie beyond it, as illustrated below in Figure 1.3.

To compute and visualize the conditional quantile α-frontiers, we rely on the same graphical interface as for the FDH estimator, namely the functions `plot_frontier.2d()` and `plot_frontier.3d()`. Setting the argument `type_frontier = "alpha"` activates the order-α estimator, while the value of α is specified through the argument `alpha`. The evaluation direction is controlled by the argument `orientation`, which can take the values `"output"`, `"input"` or `"hyper"`. As in the FDH case, the function automatically constructs the evaluation grid, computes the corresponding efficiency scores, reconstructs the implied boundary, and produces the graphical representation. The frontier may either be drawn as a standalone curve or superimposed on an existing scatterplot by setting `add = TRUE`. Standard graphical arguments such as `col`, `lty` and `lwd` allow full customization of the display.

In the bivariate setting with a single input and a single output, the graphical interpretation depends on the chosen orientation. In the output direction, the production frontier corresponds to an upper α-quantile of the conditional output distribution and is represented in the usual (X, Y) plane. In the input direction, the cost frontier is defined by a lower α-quantile of the conditional input distribution. The axes are reversed in this case (via the argument `swap_axes = TRUE`), so that the α-frontier is represented as a function of the output, with efficient observations lying below it. In the hyperbolic direction, efficiency is evaluated along a multiplicative path that simultaneously contracts inputs and expands outputs, leading to a distinct partial boundary.

Figure 1.3 illustrates the three orientations of the α-frontier with $\alpha = 0.975$ for the `atp` and `movies` datasets. The output-oriented frontier (red) is constructed in the original (X, Y) plane, the input-oriented frontier (green) is displayed after permuting the axes, and the hyperbolic frontier (blue) reflects the multiplicative evaluation scheme described above. Each partial α-frontier is less influenced by extreme observations than the full FDH boundary. The implementation for the `atp` dataset proceeds as follows:

```
plot_frontier.2d(as.matrix(xobs), as.matrix(yobs), type_frontier = "alpha",
    orientation = "output", alpha = 0.975, col = "red")

plot_frontier.2d(as.matrix(xobs), as.matrix(yobs), type_frontier = "alpha",
  orientation = "input", alpha = 0.975, swap_axes = T, col = "green")

plot_frontier.2d(as.matrix(xobs), as.matrix(yobs), type_frontier = "alpha",
    orientation = "hyper", alpha = 0.975, col = "blue")
```

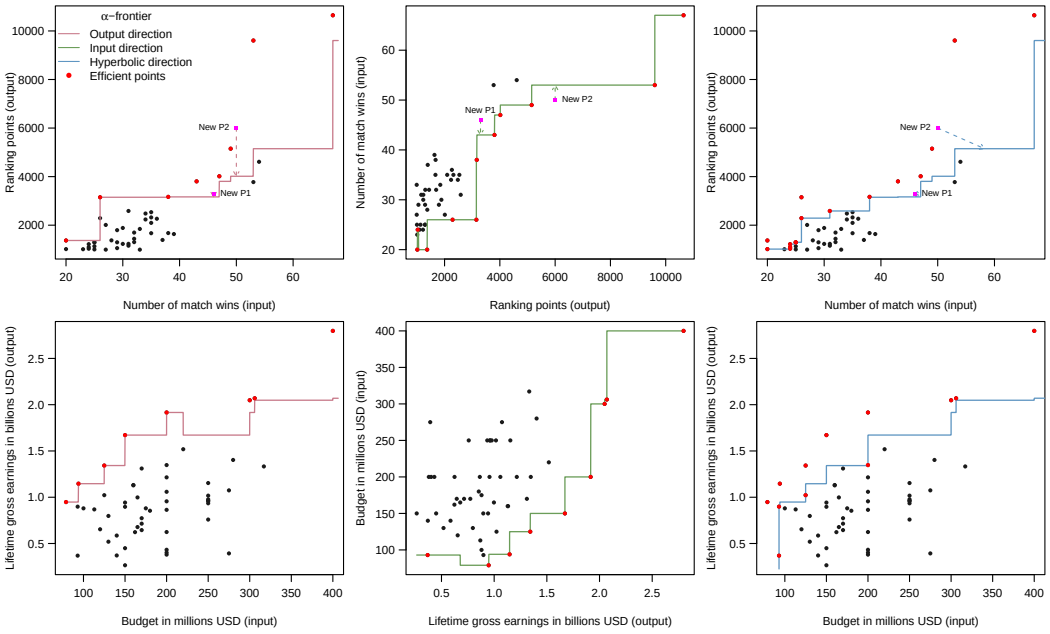

FIGURE 1.3
Quantile-based α-frontiers in the bivariate case for the `atp` dataset (top panels) and the `movies` dataset (bottom panels), shown from left to right in the output (red), input (green), and hyperbolic (blue) orientations, with $\alpha = 0.975$.

To better understand the role of the tail probability level α, we plot in Figure 1.4 its corresponding quantile-based frontier in the output direction for different values of the order $\alpha \in (0, 1]$ that controls the degree of extremeness of the frontier. For moderate α values, the partial frontier lies strictly below the FDH boundary and is less sensitive to extremes. As α increases, the frontier progressively moves outward and approaches the non-robust FDH frontier. When $\alpha = 1$, it coincides with the FDH envelope, and no trimming is applied.

When two inputs and one output are considered, the order-α frontier can be visualized using `plot_frontier.3d()`. The construction of the graphical representation depends on the chosen direction. In the output-oriented case, the frontier is reconstructed over a grid defined on the input space. More precisely, we first identify the efficient observations and construct sequences of evaluation points around them, supplemented with regularly spaced grid points covering the relevant input domain, in a spirit similar to FDH-type constructions. For each evaluation point (x_1, x_2) on this grid, the output-oriented order-α efficiency score $\widehat{\lambda}_\alpha(x, y)$ is computed. The estimated attainable output level is then recovered by rescaling the observed output according to this score, which allows us to reconstruct the output-oriented frontier as a surface in the 3-dimensional space (x_1, x_2, y). Figure 1.5(a) displays this surface for $\alpha = 0.999$. The perspective view highlights the piecewise structure of the estimated boundary. Alternatively, a 2-dimensional choropleth map could be used, projecting the attainable output levels onto the input space.

In the input-oriented case, the geometry of the frontier differs fundamentally. For a given output level y, the order-α frontier is defined as the set of input combinations that achieve this output with minimal proportional input contraction. The frontier therefore consists of isoquants in the input space rather than a single surface. Accordingly, the user specifies the output levels of interest, denoted by $\{y_1, \ldots, y_K\}$. For each selected y_k, we construct a

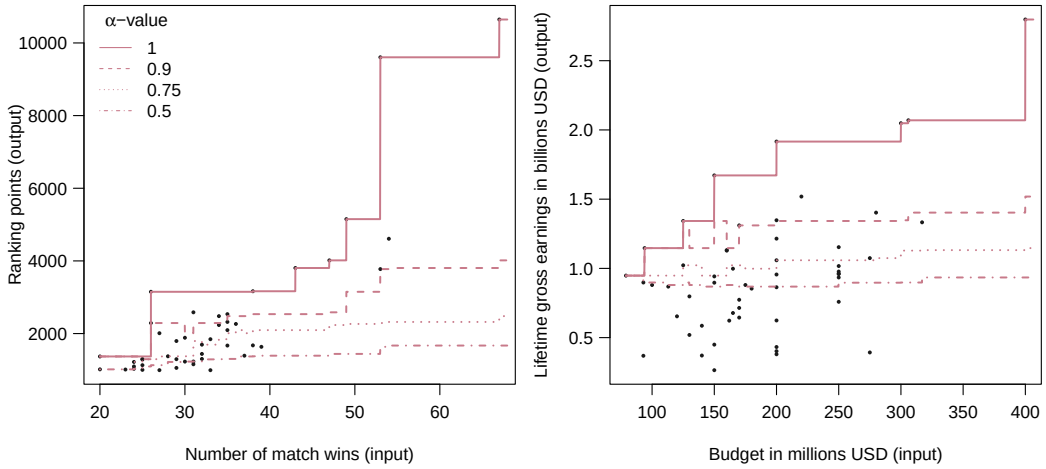

FIGURE 1.4
Output-oriented quantile frontiers in the bivariate case for the `atp` dataset (left panel) and the `movies` dataset (right panel). The α-frontiers are reported for $\alpha = 0.5, 0.75, 0.9$ and 1.

grid of evaluation points in the input space and compute the corresponding efficiency scores $\widehat{\theta}_\alpha(x, y_k)$. The boundary is then recovered by identifying the input combinations satisfying $\widehat{\theta}_\alpha(x, y_k) = 1$, which yields the estimated order-α isoquant associated with output level y_k. Figure 1.5(b) presents these isoquants for four output levels and $\alpha = 0.999$. Each curve represents the estimated input combinations required to attain a given output under the trimming level α. Efficient units, defined as observations satisfying $\widehat{\theta}_\alpha(x, y) \geqslant 1$ in the input orientation and $\widehat{\lambda}_\alpha(x, y) \leqslant 1$ in the output orientation, are highlighted to emphasize their position relative to the estimated frontier.

```
plot_frontier.3d(X, Y, type_frontier = "alpha", panels = "surface",
    orientation = "output", alpha = 0.999)

plot_frontier.3d(X, Y, type_frontier = "alpha", panels = "map",
    orientation = "input", alpha = 0.999, y_levels = c(0.1, 0.35, 0.72, 0.98))
```

1.3.2 Robustified α-efficiency measures

We now introduce the unconditional versions of both output- and input-oriented α-efficiency scores, along with an alternative formulation of technical α-efficiency for hyperbolic paths, as suggested in [19]. We will denote in the sequel by $F_Z^{\leftarrow}(\alpha) := \inf\{z | F_Z(z) \geqslant \alpha\}$ the quantile of order $\alpha \in (0, 1]$ of a random variable Z with distribution function F_Z, and by $Z_{(1)} \leqslant \cdots \leqslant Z_{(n)}$ the order statistics associated to a sample $Z_1, \ldots, Z_n$ form Z.

We start with the output-orientation by introducing an alternative representation of $\lambda(x, y)$ and its FDH estimator $\widehat{\lambda}(x, y)$. Specifically, define the transformed random variable

$$Z^{xy} := Z^{xy}(X, Y) = \min_{1 \leqslant j \leqslant q} \frac{Y^{[j]}}{y^{[j]}} \mathbb{1}(X \leqslant x). \tag{1.2}$$

This dimensionless transformation has distribution function

$$F_{Z^{xy}}(z) = \begin{cases} 1 - H_{XY}(x, zy), & \text{if } z \geqslant 0, \\ 0, & \text{if } z < 0, \end{cases}$$

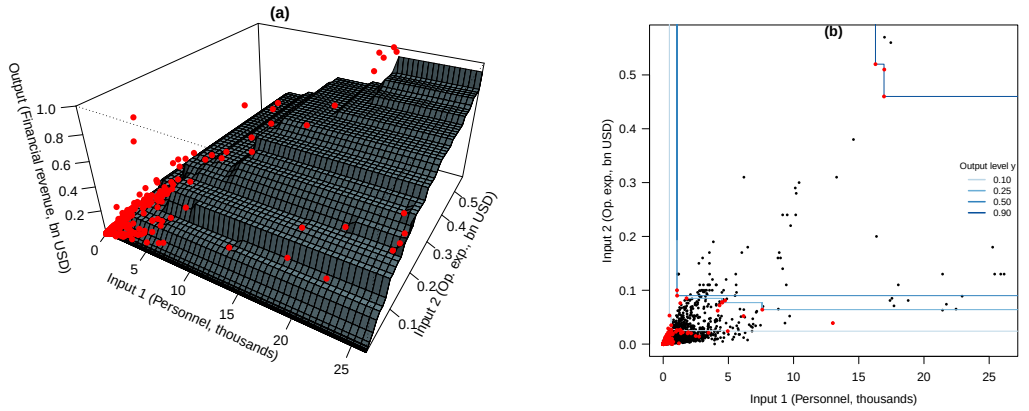

FIGURE 1.5
Quantile-type frontier of order $\alpha = 0.999$ for the microfinance institutions dataset. In panel (a), the output-oriented frontier is represented as a 3-dimensional surface showing the maximum attainable α-trimmed output for each input combination. In panel (b), the input-oriented frontier is depicted by isoquants associated with selected output levels, illustrating the minimal proportional inputs-usage needed to produce these outputs. Efficient observations are highlighted in red.

and yields the sample of n draws

$$Z_i^{xy} := Z^{xy}(X_i, Y_i) = \min_{1 \leqslant j \leqslant q} \frac{Y_i^{[j]}}{y^{[j]}} \, \mathbb{1}(X_i \leqslant x) \quad \text{for} \quad i = 1, \ldots, n. \tag{1.3}$$

While the efficiency measure $\lambda(x, y)$ coincides with the right endpoint $F^{\leftarrow}_{Z^{xy}}(1)$ of $Z^{xy}(X, Y)$, its FDH estimator $\widehat{\lambda}(x, y)$ coincides with the sample maximum $Z^{xy}_{(n)}$. This motivates an alternative quantile-type efficiency score and its estimator, replacing $\lambda_\alpha(x, y)$ and $\widehat{\lambda}_\alpha(x, y)$, defined by

$$\lambda^\star_\alpha(x, y) := F^{\leftarrow}_{Z^{xy}}(\alpha) \quad \text{and} \quad \widehat{\lambda}^\star_\alpha(x, y) := \widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha),$$

for the empirical distribution function $\widehat{F}_{Z^{xy}}(z) = \frac{1}{n}\sum_{i=1}^n \mathbb{1}(Z_i^{xy} \leqslant z)$ and its αth quantile $\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha) = Z^{xy}_{(\lfloor \alpha n \rfloor + 1)}$. Both $\lambda^\star_\alpha(x, y)$ and its estimator $\widehat{\lambda}^\star_\alpha(x, y)$ enjoy the desirable monotonicity properties in x and y. Note also that, for a firm operating at (x, y), the order parameter should satisfy $\alpha > 1 - \mathbb{P}(X \leqslant x)$; this condition provides a natural lower bound for selecting α in $\lambda^\star_\alpha(x, y)$. Using elementary extreme-value arguments, [19] derive the asymptotic distributional behavior of $\widehat{\lambda}^\star_\alpha(x, y)$ for estimating both full and partial efficiency scores. They further establish, via both infinitesimal and global robustness analyses, that $\widehat{\lambda}^\star_\alpha(x, y)$ enjoys enhanced robustness, in particular relative to its conditional counterpart $\widehat{\lambda}_\alpha(x, y)$. Moreover, applying the asymptotic approach of [28] developed in Section 2.6.3 (pp. 103–104), a $100(1-\tau)\%$ confidence interval for the αth quantile $\lambda^\star_\alpha(x, y)$ is given by

$$]\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha_{n1}), \widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha_{n2})[=]\widehat{\lambda}^\star_{\alpha_{n1}}(x, y), \, \widehat{\lambda}^\star_{\alpha_{n2}}(x, y)[,$$

where

$$\alpha_{n1} = \alpha - z_{1-\tau/2}\big(\alpha(1-\alpha)/n\big)^{1/2} \quad \text{and} \quad \alpha_{n2} = \alpha + z_{1-\tau/2}\big(\alpha(1-\alpha)/n\big)^{1/2}, \tag{1.4}$$

with $z_{1-\tau/2}$ being the quantile at level $1 - \tau/2$ of the standard Gaussian distribution; for example, to obtain a 95% confidence interval, $z_{1-\tau/2} = z_{0.975} \approx 1.959964$.

In the input-orientation, recast $Z^{xy}(X,Y)$ as

$$Z^{xy}(X,Y) = -\max_{1\leqslant j\leqslant p} \frac{X^{[j]}}{x^{[j]}} / \mathbb{1}(Y \geqslant y), \tag{1.5}$$

where dividing any nonnegative real number by 0 is taken to be $+\infty$. This transformed variable has support $[-\infty, 0]$ and distribution function

$$F_{Z^{xy}}(z) = \begin{cases} 1 - H_{XY}(-zx, y), & \text{if } z > -\infty, \\ 0, & \text{if } z = -\infty. \end{cases}$$

Here, $\theta(x,y) = -F^{\leftarrow}_{Z^{xy}}(1)$ and $\widehat{\theta}(x,y) = -\max_{1\leqslant i\leqslant n} Z^{xy}(X_i, Y_i)$. This motivates the introduction of unconditional countarparts of $\theta_\alpha(x,y)$ and its estimator $\widehat{\theta}_\alpha(x,y)$ as

$$\theta^\star_\alpha(x,y) := -F^{\leftarrow}_{Z^{xy}}(\alpha) \quad \text{and} \quad \widehat{\theta}^\star_\alpha(x,y) := -\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha),$$

with the associated asymptotic $100(1-\tau)\%$ confidence interval for $\theta^\star_\alpha(x,y)$ defined as

$$]-\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha_{n2}), -\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha_{n1})[=]\widehat{\theta}^\star_{\alpha_{n2}}(x,y),\, \widehat{\theta}^\star_{\alpha_{n1}}(x,y)[$$

for the levels α_{n1} and α_{n2} described in (1.4).

In the hyperbolic-orientation, redefine $Z^{xy}(X,Y)$ as

$$Z^{xy}(X,Y) = \min\left\{ \min_{1\leqslant j\leqslant p} \frac{x^{[j]}}{X^{[j]}},\ \min_{1\leqslant j\leqslant q} \frac{Y^{[j]}}{y^{[j]}} \right\} \tag{1.6}$$

whose distribution function is given by

$$F_{Z^{xy}}(z) = 1 - H_{XY}(z^{-1}x, zy) \quad \text{for } z > 0.$$

Then $\gamma(x,y) = F^{\leftarrow}_{Z^{xy}}(1)$ and $\widehat{\gamma}(x,y) = \max_{1\leqslant i\leqslant n} Z^{xy}(X_i,Y_i)$, suggesting to reformulate $\gamma_\alpha(x,y)$ and $\widehat{\gamma}_\alpha(x,y)$ by the respective quantiles

$$\gamma_\alpha(x,y) \equiv \gamma^\star_\alpha(x,y) := F^{\leftarrow}_{Z^{xy}}(\alpha) \quad \text{and} \quad \widehat{\gamma}_\alpha(x,y) \equiv \widehat{\gamma}^\star_\alpha(x,y) := \widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha),$$

and to consider the asymptotic $100(1-\tau)\%$ confidence interval $]\widehat{\gamma}^\star_{\alpha_{n1}}(x,y),\, \widehat{\gamma}^\star_{\alpha_{n2}}(x,y)[$ for $\gamma_\alpha(x,y)$, with the same levels α_{n1} and α_{n2} in (1.4).

A key appeal of the unconditional quantile-based approach is that the resulting order-α frontier is orientation-invariant. Specifically,

$$\Psi^\partial_\alpha = \big\{(x,y)\in\Psi \,\big|\, \lambda^\star_\alpha(x,y) = 1\big\} \equiv \big\{(x,y)\in\Psi \,\big|\, \theta^\star_\alpha(x,y) = 1\big\} \equiv \big\{(x,y)\in\Psi \,\big|\, \gamma^\star_\alpha(x,y) = 1\big\},$$

or equivalently

$$\begin{aligned} \Psi^\partial_\alpha &= \{(x,\, \lambda^\star_\alpha(x,y)y) \mid (x,y)\in\Psi\} \equiv \{(\theta^\star_\alpha(x,y)x,\, y) \mid (x,y)\in\Psi\} \\ &\equiv \{(\gamma^\star_\alpha(x,y)^{-1}x,\, \gamma^\star_\alpha(x,y)y) \mid (x,y)\in\Psi\}, \end{aligned}$$

so it is uniquely defined and retains the same economic meaning whether one adopts an output, input, or hyperbolic orientation. Naturally, the associated order-α efficiency scores (*i.e.*, the distances to this common frontier) depend on the chosen direction.

The `alphascore.star()` routine in the `frontiles` package computes the unconditional efficiency measures $\widehat{\lambda}^\star_\alpha(x,y)$, $\widehat{\theta}^\star_\alpha(x,y)$ and $\widehat{\gamma}^\star_\alpha(x,y)$ for a given set of evaluation points (x,y). Additionally, it provides the $100\times(1-\tau)\%$ confidence intervals associated with these estimates by setting the arguments `CI = TRUE` and specifying the desired `tau` value. To ensure comparability with the previous analysis, we again consider the evaluation points P1 and P2 from the `atp` dataset for which we set $\alpha = 0.975$ and compute the 95% confidence intervals for $\tau = 0.05$:

```
alphascore.star(as.matrix(xobs), as.matrix(yobs), xeval = matrix(x_new),
    yeval = matrix(y_new), alpha = 0.975,  CI = T, tau = 0.05)
```

```
        output    input    hyper output_low input_low hyper_low
P1 0.9585227 1.021739 1.024307  0.7809807 0.9347826 0.9347826
P2 0.6647917 1.340000 1.165049  0.5271118 1.0600000 1.0600000

     output_upp input_upp hyper_upp
P1   1.1530303  1.152174  1.047619
P2   0.8583333       Inf  1.340000
```

It should be first clear that, in the input-orientation and for the evaluation point $(x, y) =$ P2, there are only two observations (X_i, Y_i) such that $Y_i \geqslant y$. Consequently, the sample of transformed variables $Z^{xy}(X_i, Y_i)$ defined in (1.5) contains $n - 2$ values equal to $-\infty$ and only two finite negative values. This immediately implies that the 95% upper confidence bound for $\theta^{\star}_{\alpha}(x, y)$ is infinite, namely $\widehat{\theta}^{\star}_{\alpha_{n1}}(x, y) = -\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha_{n1}) \equiv \infty$, since $n = 46$ and $\alpha_{n1} = \alpha - 1.96\big(\alpha(1-\alpha)/n\big)^{1/2} \equiv 0.887$, and so $\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha_{n1}) = Z^{xy}_{(\lfloor \alpha_{n1} n \rfloor + 1)} \equiv Z^{xy}_{(41)} = -\infty$. It should also be noted that the proposed asymptotic confidence interval for the quantile $\widehat{F}^{\leftarrow}_{Z^{xy}}(\alpha)$ avoids the notoriously difficult task of estimating the value of the density function of $Z^{xy}(X, Y)$ at that quantile, but still requires a sufficiently large sample size to yield meaningful inference.

Compared with the conditional efficiency measures, noticeable differences arise under the unconditional specification. Both P1 and P2 are now identified as efficient in all orientations, whereas P1 was not deemed efficient in the input orientation under the conditional framework.

To visualize the robustified α-frontier, common to the output-, input- and hyperbolic-oriented paths, we use the function `plot_frontier.2d()` for 2-dimensional representations and `plot_frontier.3d()` for 3-dimensional representations, setting `type_frontier = "alpha.star"`. All remaining arguments are identical to those used with `type_frontier = "alpha"`, thereby ensuring full comparability between the unconditional and conditional graphical analyses. The resulting α-quantile frontiers are displayed in Figure 1.6 for the one-input–one-output setting and in Figure 1.7 for the two-input–one-output case. The monotonicity of the constrained unconditional frontier in Figure 1.6 is a rather powerful way of stability compared with its unconstrained conditional counterpart shown in Figure 1.3. Here, we provide a few illustrative code examples. The `plot_frontier.2d()` function can display confidence intervals using the arguments `CI = TRUE` and by specifying the value of the confidence level `tau`.

```
# 1-input and 1-output (output-oriented)
plot_frontier.2d(xobs = matrix(xobs), yobs = matrix(yobs),
   type_frontier = "alpha.star", orientation = "output", alpha = 0.975)

# 1-input and 1-output + Confidence Interval
plot_frontier.2d(xobs = matrix(xobs), yobs = matrix(yobs),
   type_frontier = "alpha.star", orientation = "output", alpha = 0.975,
   CI = T, tau = 0.05)

# 2-input and 1-output (output-oriented)
plot_frontier.3d(X, Y, type_frontier = "alpha.star", panels = "surface",
   orientation = "output", alpha = 0.999)
```

```
# 2-input and 1-output (input-oriented)
plot_frontier.3d(X, Y, type_frontier = "alpha.star", panels = "map",
   orientation = "input", alpha = 0.999 , y_levels = c(0.1, 0.35, 0.72, 0.9))
```

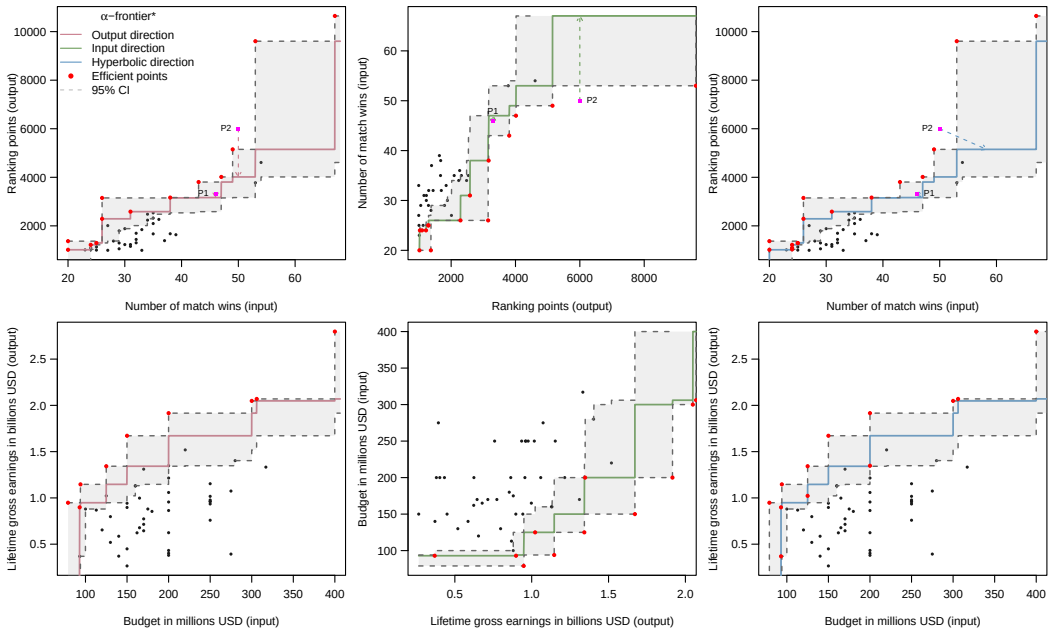

FIGURE 1.6
Robustified quantile-based α-frontiers, with `type_frontier = "alpha.star"`, in the bivariate case for the `atp` dataset (top panels) and the `movies` dataset (bottom panels), shown from left to right in the output (red), input (green), and hyperbolic (blue) orientations. The input-oriented frontier represents the same estimated attainable set, with the only difference being a permutation of the axes. The 95% confidence bands are displayed for each dataset and orientation.

1.4 Weighted-moment partial efficiency scores

Two competing notions to partial quantile-type efficiency, both based on suitable probability-weighted moments, were proposed in [7] under a conditional framework and in [13] through an unconditional construction. Below, Sections 1.4.1 and 1.4.2 present these two approaches from the dual perspective of defining the underlying concepts and describing their practical computation, respectively.

1.4.1 Robust m-efficiency measures

We focus the presentation on the output-orientation, but the results extend trivially to other directions. Fix a firm operating at (x, y) with $H_{XY}(x, 0) = \mathbb{P}(X \leqslant x) > 0$. Draw $m \in \{1, 2, \ldots\}$ independent outputs $Y_1(x), \ldots, Y_m(x)$ from the conditional distribution of

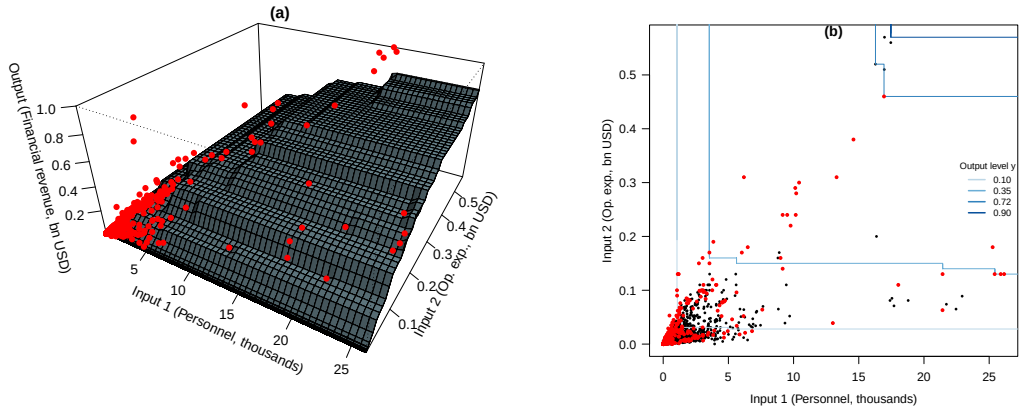

FIGURE 1.7
Robustified quantile-type frontier of order $\alpha = 0.999$ for the microfinance institutions dataset. In panel (a), the output-oriented frontier is represented as a 3-dimensional surface showing the maximum attainable trimmed output for each input combination. In panel (b), the input-oriented frontier is depicted by isoquants associated with selected output levels. Efficient observations are highlighted in red.

Y given $X \leqslant x$. Following [2], the classical order-m efficiency score is then defined as the expected maximum

$$\lambda_m(x,y) := \mathbb{E}\left[\max(Z_1^{xy},\ldots,Z_m^{xy})\right]$$

of the random variables

$$Z_i^{xy} = Z^{xy}(X_i,Y_i) := \min_{j=1,\ldots,q}\left(\frac{Y_i^{[j]}(x)}{y^{[j]}}\right) \quad \text{for} \quad i=1,\ldots,m, \tag{1.7}$$

or equivalently as

$$\lambda_m(x,y) = \lambda(x,y) - \int_0^{\lambda(x,y)}\left[1-\frac{H_{XY}(x,vy)}{H_{XY}(x,0)}\right]^m dv$$

which converges to the full efficiency score $\lambda(x,y)$ as $m\to\infty$. Its empirical counterpart

$$\widehat{\lambda}_m(x,y) := \widehat{\lambda}(x,y) - \int_0^{\widehat{\lambda}(x,y)}\left[1-\frac{\widehat{H}_{n,XY}(x,vy)}{\widehat{H}_{n,XY}(x,0)}\right]^m dv$$

converges to the FDH estimator $\widehat{\lambda}(x,y)$ as $m\to\infty$, for a fixed sample size n. A detailed account of the robustness and inferential properties of this estimator is provided in [2], [8] and [16]. As emphasized in [7], the main weaknesses of $\lambda_m(x,y)$ and its estimator $\widehat{\lambda}_m(x,y)$ stem from their conditional construction. In particular, the normalizing factors $H_{XY}(x,0)$ and $\widehat{H}_{n,XY}(x,0)$ may induce monotonicity violations in x, divergence from the true efficiency score $\lambda(x,y)$ for extreme values of x, and a consequent loss of robustness.

From a computational standpoint, since the pioneering work of [2], it has long been believed that no simple closed-form expression is available for the estimator $\widehat{\lambda}_m(x,y)$ in the presence of multiple inputs and outputs, and a Monte Carlo–based algorithm was therefore proposed for its computation. The `ordermscore.boot()` routine in the `frontiles` package

implements this bootstrap procedure to compute the output-oriented order-m efficiency estimates $\widehat{\lambda}_m(x,y)$ at a user-specified set of evaluation points (x,y), and it also returns the corresponding input-oriented counterparts

$$\widehat{\theta}_m(x,y) = \int_0^\infty \left[1 - \frac{\widehat{H}_{n,XY}(ux,y)}{\widehat{H}_{n,XY}(\infty,y)}\right]^m du$$

that estimate the conditional expected minimum [2, 31]:

$$\theta_m(x,y) = \theta(x,y) + \int_{\theta(x,y)}^\infty \left[1 - \frac{H_{XY}(ux,y)}{H_{XY}(\infty,y)}\right]^m du,$$

as well as the hyperbolic counterparts

$$\widehat{\gamma}_m(x,y) = \int_0^\infty \left[1 - \widehat{H}_{n,XY}(ux,u^{-1}y)\right]^m du$$

that estimate the unconditional expected minimum [33]:

$$\gamma_m(x,y) = \int_0^\infty \left[1 - H_{XY}(ux,u^{-1}y)\right]^m du.$$

As a matter of fact, the exact values of the estimators in the different directions can be obtained directly. In the output-orientation, it is straightforward to verify that $\widehat{\lambda}_m(x,y)$ admits the following simple explicit expression:

$$\widehat{\lambda}_m(x,y) = \sum_{i=1}^{N_x} Z_{(i)}^{xy}\left[\left(\frac{i}{N_x}\right)^m - \left(\frac{i-1}{N_x}\right)^m\right],$$

where $N_x = n\widehat{H}_{n,XY}(x,0)$ denotes the random number of inputs X_i such that $X_i \leqslant x$, and $Z_{(1)}^{xy} \leqslant \cdots \leqslant Z_{(N_x)}^{xy}$ are the order statistics of the variables Z_i^{xy} described in (1.7) for $i = 1,\ldots,N_x$, with $Y_i(x)$ therein being the N_x outputs Y_i such that $X_i \leqslant x$.

In the input-orientation, we have

$$\widehat{\theta}_m(x,y) = \sum_{i=1}^{N_y} Z_{(N_y-i+1)}^{xy}\left[\left(\frac{i}{N_y}\right)^m - \left(\frac{i-1}{N_y}\right)^m\right],$$

where $N_y = n\widehat{H}_{n,XY}(\infty,y)$ is the number of outputs Y_i such that $Y_i \geqslant y$, and the sample points $(Z_1^{xy},\ldots,Z_{N_y}^{xy})$ are defined in (1.1). Likewise, in the hyperbolic-orientation,

$$\widehat{\gamma}_m(x,y) = \sum_{i=1}^{n} Z_{(n-i+1)}^{xy}\left[\left(\frac{i}{n}\right)^m - \left(\frac{i-1}{n}\right)^m\right],$$

for the sample points

$$Z_i^{xy} := Z^{xy}(X_i,Y_i) = \max\left\{\max_{1\leqslant j\leqslant p}\frac{X_i^{[j]}}{x^{[j]}},\ \max_{1\leqslant j\leqslant q}\frac{y^{[j]}}{Y_i^{[j]}}\right\},\quad i = 1,\ldots,n.$$

The `ordermscore()` routine in the `frontiles` package computes $\widehat{\lambda}_m(x,y)$ and provides the corresponding input- and hyperbolic-oriented counterparts $\widehat{\theta}_m(x,y)$ and $\widehat{\gamma}_m(x,y)$. For ease of comparison with the previous sections, we return to the evaluation points P_1 and P_2 from the `atp` dataset and calculate the associated order-m efficiency scores $\widehat{\lambda}_m(x,y)$, $\widehat{\theta}_m(x,y)$ and $\widehat{\gamma}_m(x,y)$:

```
 ordermscore(xobs = matrix(xobs), yobs = matrix(yobs),
        xeval = matrix(x_new), yeval = matrix(y_new), m = 30)

          output     input     hyper
P1    1.0434340 0.9356373 0.9903968
P2    0.7444799 1.0600000 1.1711289
```

For the chosen trimming level $m = 30$, the results are quantitatively similar to those obtained with the quantile-type measures for $\alpha = 0.975$ in Section 1.3.1, indicating that these two parameter choices induce a comparable degree of trimming. Although P1 is no longer classified as efficient in the output and hyperbolic orientations under the order-m benchmark specification, the corresponding efficiency scores remain numerically close to the α-quantile scores. The distinction lies in the efficiency classification rather than in the magnitude of the performance measures. When it comes to compute and plot the order-m frontiers in the output, input and hyperbolic directions, we again employ the function `plot_frontier.2d()` for bivariate representations, using `type_frontier = "orderm"` and specifying the value of m (here `m = 30`). For example, when considering a case with single input and output in the output-orientation, we use

```
plot_frontier.2d(xobs = matrix(xobs), yobs = matrix(yobs),
   type_frontier = "orderm", orientation = "output", m = 30)
```

The order-m frontiers obtained in the three directions are graphed in Figure 1.8 for both `atp` and `movies` datasets. By construction, these partial frontiers approach the enveloping FDH boundary as m increases. While their overall shape remains close to that of the full envelopment benchmark, slight inward deviations are visible, especially in regions affected by extreme observations.

To better understand the role of m, we plot in Figure 1.9 the corresponding m-frontiers in the output direction for different values of m. For small m, the partial m-frontier is smoother and lies well below the FDH boundary. When m becomes large enough relative to the sample size, the m-frontier tends to coincide with the FDH frontier itself.

When two inputs and one output are considered, the order-m frontier can be visualized with `plot_frontier.3d()`. As in the quantile-based framework, the frontier is evaluated over the same grid of input combinations, but the corresponding efficiency scores are now computed from the expected minimum/maximum representation that defines the weighted-moment framework:

```
plot_frontier.3d(X, Y, type_frontier = "orderm", panels = "surface",
    orientation = "output", m = 500)

plot_frontier.3d(X, Y, type_frontier = "orderm", panels = "map",
    orientation = "input", m = 500, y_levels = c(0.1, 0.35, 0.72, 0.98))
```

Figure 1.10(a) displays the resulting frontier whose order $m = 500$ is chosen proportionally to the sample size $n = 4833$. In the input orientation, isoquants associated with selected output levels $\{y_1, \ldots, y_K\}$ are obtained in Figure 1.10(b) from $\widehat{\theta}_m(x, y_k) = 1$, similarly to the quantile framework in Figure 1.5(b). Efficient units satisfy $\widehat{\lambda}_m(x, y) \leqslant 1$ in panel (a) and $\widehat{\theta}_m(x, y) \geqslant 1$ in panel (b); they are highlighted in red in the figure.

1.4.2 Robustified m-efficiency measures

Similarly to the unconditional quantile-based approach [19] in the output-orientation, [13] transform the multivariate data generating process $(X, Y) \in \mathbb{R}_+^{p+q}$ into the dimensionless

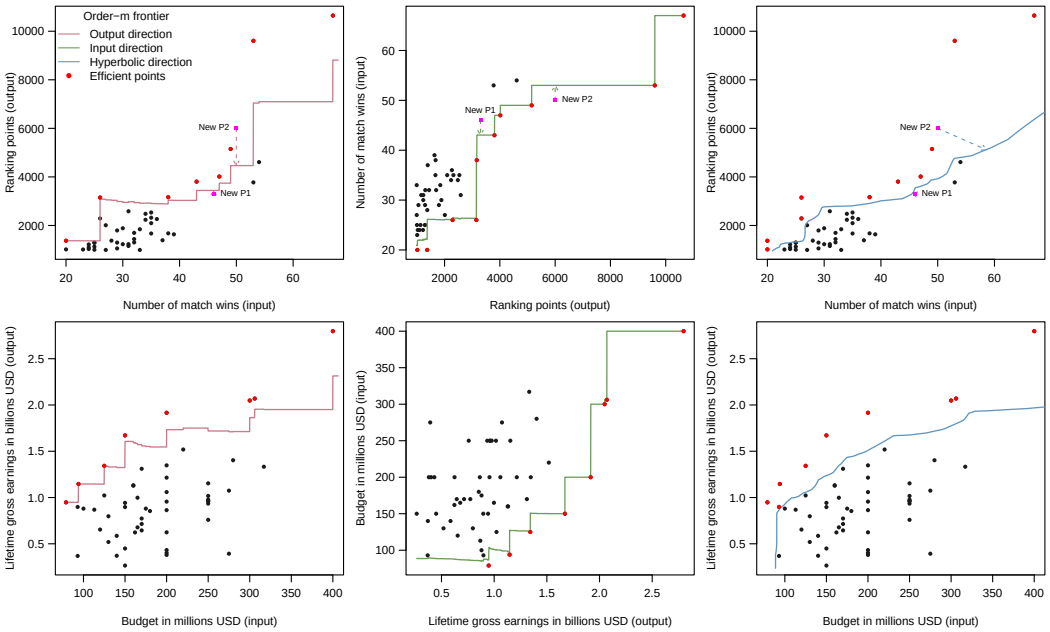

FIGURE 1.8
Weighted-moment m-frontiers in the bivariate case for the `atp` dataset (top panels) and the `movies` dataset (bottom panels), shown from left to right in the output (red), input (green), and hyperbolic (blue) orientations, with $m = 30$.

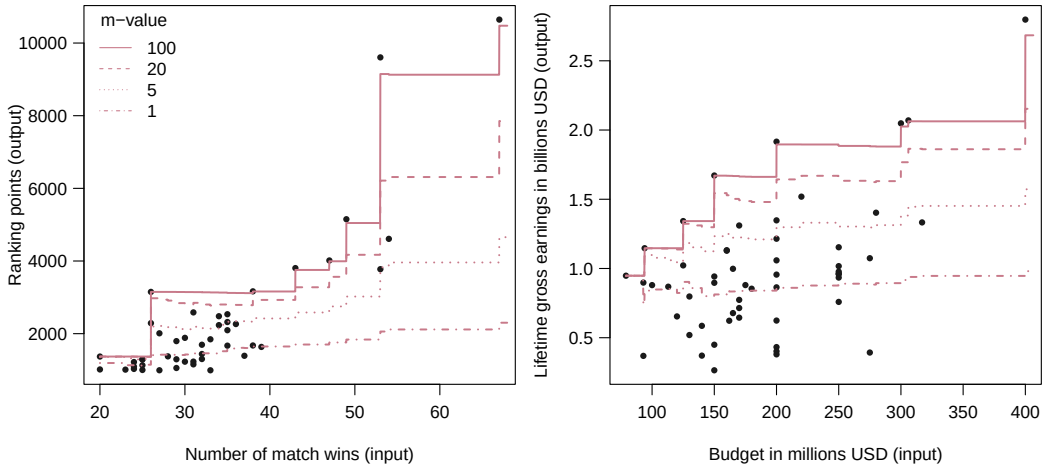

FIGURE 1.9
Output-oriented weighted-moment frontiers in the bivariate case for the `atp` dataset (left panel) and the `movies` dataset (right panel). The m-frontiers are reported for $m = 1, 5, 20,$ and 100.

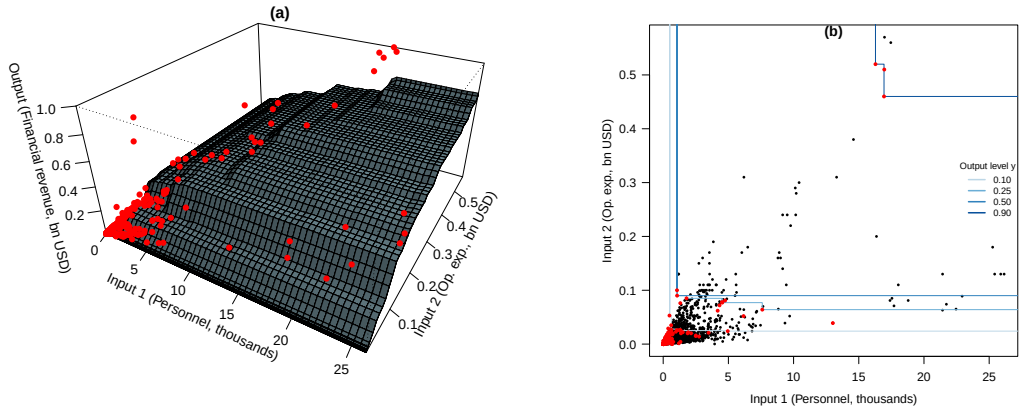

FIGURE 1.10
Weighted-moment frontier of order $m = 500$ for the microfinance institutions dataset. In panel (a), the output-oriented frontier is represented as a 3-dimensional surface showing the maximum attainable m-trimmed output for each input combination. In panel (b), the input-oriented frontier is depicted by isoquants associated with selected output levels. Efficient observations are highlighted in red.

random variable $Z^{xy}(X, Y)$ and its associated sample of copies $Z_i^{xy} = Z^{xy}(X_i, Y_i)$, described in (1.2) and (1.3) respectively. While [19] define their unconditional α-efficiency score $\lambda_\alpha^\star(x, y) := F^{\leftarrow}_{Z^{xy}}(\alpha)$ as the αth quantile of the transformation $Z^{xy}(X, Y)$, [13] construct their unconditional m-efficiency score as the expected maximum

$$\begin{aligned}\lambda_m^\star(x, y) &:= \mathbb{E}\left[\max(Z_1^{xy}, \ldots, Z_m^{xy})\right] \\ &= \lambda(x, y) - \int_0^{\lambda(x,y)} \left[1 - H_{XY}(x, zy)\right]^m \, dz,\end{aligned}$$

which increases toward $\lambda(x, y)$ as the trimming level $m \to \infty$. Replacing H_{XY} with its empirical counterpart $\widehat{H}_{n,XY}$ then yields the natural plug-in estimator

$$\widehat{\lambda}_m^\star(x, y) := \widehat{\lambda}(x, y) - \int_0^{\widehat{\lambda}(x,y)} \left[1 - \widehat{H}_{n,XY}(x, zy)\right]^m \, dz$$

which converges monotonically to the FDH estimator $\widehat{\lambda}(x, y)$ as $m \to \infty$. Both $\lambda_m^\star(x, y)$ and $\widehat{\lambda}_m^\star(x, y)$ satisfy the key monotonicity axiom induced by the partial order: they are nondecreasing in x and nonincreasing in y. The estimator $\widehat{\lambda}_m^\star(x, y)$ is readily computable as it admits an explicit L-statistic representation:

$$\widehat{\lambda}_m^\star(x, y) \equiv \sum_{i=1}^{n} Z_{(i)}^{xy} \left[\left(\frac{i}{n}\right)^m - \left(\frac{i-1}{n}\right)^m\right],$$

with $Z_{(1)}^{xy} \leqslant \cdots \leqslant Z_{(n)}^{xy}$ being the order statistics of the sample $(Z_1^{xy}, \ldots, Z_n^{xy})$ introduced in (1.3). Its consistency, asymptotic normality and Berry-Esséen rate of uniform convergence are established in [13]. Other estimators with similar asymptotic properties to $\widehat{\lambda}_m^\star(x, y)$ are presented in [13], including an unbiased probability-weighted moment estimator. All these estimators of $\lambda_m^\star(x, y)$ are shown to be bias-robust and resistant to the influence of outliers

in the Y-direction. An asymptotic $100(1-\tau)\%$ confidence interval for $\lambda^{\star}_m(x,y)$ is given by $[\widehat{\lambda}^{\star}_m(x,y) \pm z_{1-\tau/2}\, \widehat{\sigma}_m(x,y)/\sqrt{n}]$, for the estimated asymptotic variance

$$\widehat{\sigma}^2_m(x,y) := m^2 \int_0^1 \int_0^1 (s \wedge t - s\,t)\,(st)^{m-1}\, d\widehat{F}^{\leftarrow}_{Z^{xy}}(s)\, d\widehat{F}^{\leftarrow}_{Z^{xy}}(t), \tag{1.8}$$

which has the exact finite-sample expression

$$\widehat{\sigma}^2_m(x,y) = \frac{m^2}{n^{2m}} \left\{ 2 \sum_{1 \leqslant i < j \leqslant n-1} i^m j^{m-1} (n-j) \left(Z^{xy}_{(i+1)} - Z^{xy}_{(i)}\right)\left(Z^{xy}_{(j+1)} - Z^{xy}_{(j)}\right) \right.$$
$$\left. + \sum_{i=1}^{n-1} i^{2m-1} (n-i) \left(Z^{xy}_{(i+1)} - Z^{xy}_{(i)}\right)^2 \right\}.$$

In the input-orientation, redefine $Z^{xy}(X,Y)$ as in (1.5). While the robustified, unconditional α-efficiency score $\theta^{\star}_{\alpha}(x,y) := -F^{\leftarrow}_{Z^{xy}}(\alpha) \equiv F^{\leftarrow}_{-Z^{xy}}(1-\alpha)$ is taken to be the negative αth quantile of $Z^{xy}(X,Y)$, its m-efficiency analog is defined as the negative expected maximum

$$\begin{aligned} \theta^{\star}_m(x,y) &:= -\mathbb{E}\left[\max(Z^{xy}_1, \ldots, Z^{xy}_m)\right] \\ &= \theta(x,y) + \int_{\theta(x,y)}^{0} \left[1 - H_{XY}(wx, y)\right]^m \, dw, \end{aligned}$$

which amounts to the expected minimum $\theta^{\star}_m(x,y) \equiv \mathbb{E}\left[\min(-Z^{xy}_1, \ldots, -Z^{xy}_m)\right]$. Its plug-in estimator is given by the L-statistic

$$\begin{aligned} \widehat{\theta}^{\star}_m(x,y) &:= \widehat{\theta}(x,y) + \int_{\widehat{\theta}(x,y)}^{0} \left[1 - \widehat{H}_{n,XY}(wx, y)\right]^m \, dw \\ &\equiv -\sum_{i=1}^{n} Z^{xy}_{(i)} \left[\left(\frac{i}{n}\right)^m - \left(\frac{i-1}{n}\right)^m \right]. \end{aligned}$$

The associated asymptotic $100(1-\tau)\%$ confidence interval for $\theta^{\star}_m(x,y)$ is $[\widehat{\theta}^{\star}_m(x,y) \pm z_{1-\tau/2}\, \widehat{\sigma}_m(x,y)/\sqrt{n}]$, for the asymptotic variance estimator $\widehat{\sigma}^2_m(x,y)$ described in (1.8).

For the hyperbolic path, using the transformed variable $Z^{xy}(X,Y)$ as in (1.6), for which the α-efficiency score is $\gamma_{\alpha}(x,y) \equiv \gamma^{\star}_{\alpha}(x,y) := F^{\leftarrow}_{Z^{xy}}(\alpha)$, the corresponding m-efficiency counterpart, initially introduced in [33] as an expected minimum $\gamma_m(x,y)$, can also be defined alternatively by the expected maximum

$$\begin{aligned} \gamma^{\star}_m(x,y) &:= \mathbb{E}\left[\max(Z^{xy}_1, \ldots, Z^{xy}_m)\right] \\ &= \gamma(x,y) - \int_{0}^{\gamma(x,y)} \left[1 - H_{XY}(z^{-1}x, zy)\right]^m \, dz \end{aligned}$$

whose empirical version is

$$\begin{aligned} \widehat{\gamma}^{\star}_m(x,y) &:= \widehat{\gamma}(x,y) - \int_{0}^{\widehat{\gamma}(x,y)} \left[1 - \widehat{H}_{n,XY}(z^{-1}x, zy)\right]^m \, dz \\ &\equiv \sum_{i=1}^{n} Z^{xy}_{(i)} \left[\left(\frac{i}{n}\right)^m - \left(\frac{i-1}{n}\right)^m \right]. \end{aligned}$$

Its associated asymptotic $100(1-\tau)\%$ confidence interval for $\gamma^{\star}_m(x,y)$ is $[\widehat{\gamma}^{\star}_m(x,y) \pm z_{1-\tau/2}\, \widehat{\sigma}_m(x,y)/\sqrt{n}]$, with $\widehat{\sigma}^2_m(x,y)$ as in (1.8).

The `ordermscore.star()` function in the `frontiles` package can be used to evaluate the three estimators $\widehat{\lambda}^{\star}_m(x,y)$, $\widehat{\theta}^{\star}_m(x,y)$ and $\widehat{\gamma}^{\star}_m(x,y)$ at any user–specified collection of points (x,y). Additionally, it can compute the $100(1-\tau)\%$ confidence intervals associated with these estimated measures by setting the arguments `CI = TRUE` and specifying the desired `tau` value. For illustration, we again consider the two reference points P_1 and P_2 used throughout the chapter to highlight the robustness of the order-m weighted moment approach. The corresponding efficiency measures remain very similar to those produced by the standard order-m approach in Section 1.4.1, while the associated frontier curves appear somewhat more pronounced graphically in the output orientation, as can be seen by comparing the left panels of Figures 1.8 and 1.11.

```
ordermscore.star(as.matrix(xobs), as.matrix(yobs),
  xeval = matrix(x_new), yeval = matrix(y_new), m = 30, CI = T, tau = 0.05)
```

```
       output     input     hyper   output_low  input_low  hyper_low
P1  1.0265013 1.007668 1.0132492  0.8559090 0.9034443 0.9386575
P2  0.7337338 1.204810 0.8656753  0.5670372 1.0176631 0.7544688

      output_upp input_upp hyper_upp
P1    1.1970936  1.111891 1.0878408
P2    0.9004303  1.391957 0.9768818
```

To generate the plots displayed in Figure 1.11, and more generally to compute the robustified order-m frontiers in the output, input and hyperbolic directions, we again use the functions `plot_frontier.2d()` and `plot_frontier.3d()` for bivariate and trivariate representations, respectively. This is done by setting `type_frontier = "orderm.star"` and specifying the desired order m, namely `m = 30` in the bivariate case and `m = 500` in the trivariate one. The corresponding estimated robustified m-frontiers for the `atp` and `movies` datasets are shown in Figure 1.11 for the 2-dimensional case and in Figure 1.12 for the 3-dimensional case. The `plot_frontier.2d()` function can display confidence intervals using the arguments `CI = TRUE` and by specifying the value of `tau`.

Compared with the traditional order-m frontiers in Figure 1.8, the robustified counterparts reported in Figure 1.11, obtained with `type_frontier = "orderm.star"`, produce monotone and slightly smoother estimates. In particular, the small local oscillations that can be observed in the classical output- and input-oriented m-frontiers tend to disappear in the robustified versions, which benefit from a natural regularization induced by monotonicity. It is also worth noting that the efficiency measure $\gamma_m(x,y)$ in the hyperbolic direction and its alternative formulation $\gamma^{\star}_m(x,y)$ produce the same robust frontier, which is considerably smoother than the output- and input-oriented frontier benchmarks.

Finally, a comparison between the quantile-based order-α frontiers in Figures 1.3 and 1.6 and their weighted-moment-based order-m competitors in Figures 1.8 and 1.11 reveals broadly similar patterns and leads to the same main conclusions for the datasets considered here. This suggests that both approaches achieve a comparable degree of robustness while preserving the overall geometry of the trimmed production frontier.

1.5 Robustness diagnostics

This section summarizes the robustness properties of the estimated order-α and order-m efficiency scores. Focusing on the output orientation, we compare the robust-conditional and

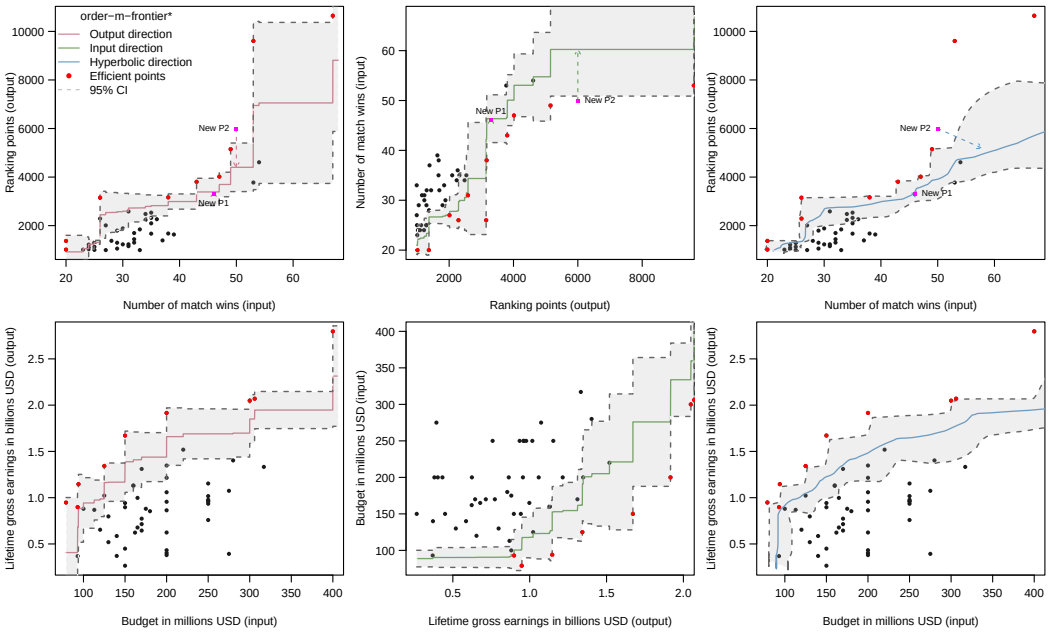

FIGURE 1.11
Robustified weighted-moment m-frontiers in the bivariate case for the `atp` dataset (top panels) and the `movies` dataset (bottom panels), shown from left to right in the output (red), input (green), and hyperbolic (blue) orientations, with $m = 30$. The 95% confidence bands are displayed for the `movies` dataset.

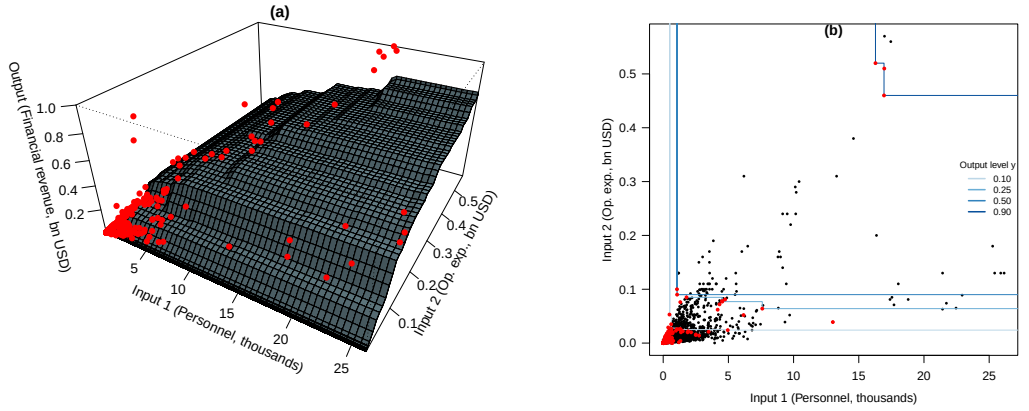

FIGURE 1.12
Robustified weighted-moment frontier of order $m = 500$ for the microfinance institutions dataset. Panel (a) shows the output-oriented m-surface of attainable expected maximum production, while panel (b) presents input-oriented isoquants for selected outputs. Efficient observations are highlighted in red.

robustified-unconditional estimators through two complementary robustness lenses: local sensitivity (influence-type) analysis and global breakdown considerations. We also provide the corresponding routines for computing these robustness measures. Altogether, these results offer concrete guidance for constructing diagnostic tools that help practitioners choose suitable trimming levels α and m in applications.

1.5.1 Quantile-type partial efficiency scores

We consider two important robustness measures of the conditional quantile estimator $\widehat{\lambda}_\alpha(x,y)$ and its unconditional variant $\widehat{\lambda}^\star_\alpha(x,y)$, namely their influence functions and sample breakdown points. The influence function (see [25], Definition 1, p. 84) quantifies the impact of an infinitesimal contamination placed at a production unit (x_0,y_0) on the estimator—scaled by the contamination mass—and, in parallel, provides a way to gauge the relative influence of each observation (X_i,Y_i) on the estimator. If it is unbounded, a single outlier may have a disproportionately large effect. Its supremum absolute value over $(x_0,y_0)\in\mathbb{R}^{p+q}_+$, known as the gross-error sensitivity (GES), quantifies the worst-case effect of contamination. If the GES is finite, the impact of any single outlying observation (X_i,Y_i) on the estimator is necessarily bounded. This robustness requirement—equivalently, finiteness of the GES—is termed B-robustness [27]. It has been found in [19] that, as $\alpha\to 1$,

$$\mathrm{GES}\Big(\widehat{\lambda}_\alpha(x,y)\Big)=\frac{\alpha}{f_{xy}(\lambda_\alpha(x,y))}>\mathrm{GES}\Big(\widehat{\lambda}^\star_\alpha(x,y)\Big)=\frac{\alpha}{f_{xy}(\lambda^\star_\alpha(x,y))},$$

where $f_{xy}(\cdot)$ stands for the probability density function of the transformed random variable $Z^{xy}(X,Y)$ defined in (1.2). Consequently, when α is taken large enough, the estimator $\widehat{\lambda}^\star_\alpha(x,y)$ can be interpreted as having a stronger infinitesimal robustness than the classical counterpart $\widehat{\lambda}_\alpha(x,y)$. The `gross.error.sensitivity()` and `gross.error.sensitivity.star()` routines in the `frontiles` package allow one to compute the respective estimates

$$\widehat{\mathrm{GES}}\Big(\widehat{\lambda}_\alpha(x,y)\Big)=\frac{\alpha}{\widehat{f}_{xy}\Big(\widehat{\lambda}_\alpha(x,y)\Big)}\quad\text{and}\quad\widehat{\mathrm{GES}}\Big(\widehat{\lambda}^\star_\alpha(x,y)\Big)=\frac{\alpha}{\widehat{f}_{xy}\Big(\widehat{\lambda}^\star_\alpha(x,y)\Big)},$$

by using `npudensbw()` from the `np` package to obtain the Parzen–Rosenblatt density estimates $\widehat{f}_{xy}$ with a chosen kernel and a bandwidth selected via cross-validation. To compute the estimated GES values at the two evaluation points P_1 and P_2 considered throughout the analysis, we use the following commands:

```
gross.error.sensitivity(xobs = matrix(xobs), yobs = matrix(yobs),
  xeval = matrix(x_new), yeval = matrix(y_new), type = "alphascore",
  alpha = 0.975)
```

```
P1   3.461625
P2   3.340155
```

```
gross.error.sensitivity.star(xobs = matrix(xobs), yobs = matrix(yobs),
  xeval = matrix(x_new), yeval = matrix(y_new), type = "alphascore",
  alpha = 0.975)
```

```
P1   2.982849
P2   2.415782
```

Despite their B-robustness, the two estimators $\widehat{\lambda}_\alpha(x,y)$ and $\widehat{\lambda}^\star_\alpha(x,y)$ may remain quite sensitive to small but finite perturbations of the data. A more global assessment of robustness is offered by the finite-sample breakdown point [23], which quantifies the smallest proportion of contamination in the original sample that can drive an estimator arbitrarily far from its value computed on the uncontaminated sample. While the sample replacement breakdown point of the unconditional quantile $\widehat{\lambda}^\star_\alpha(x,y)$ attains

$$\mathrm{RB}\Big(\widehat{\lambda}^\star_\alpha(x,y)\Big) = \begin{cases} \big(n(1-\alpha)+1\big)/n & \text{if } n\alpha = \lfloor n\alpha \rfloor, \\ \big(n - \lfloor n\alpha \rfloor\big)/n & \text{otherwise,} \end{cases}$$

as reported in [19], the corresponding breakdown value for the conditional quantile $\widehat{\lambda}_\alpha(x,y)$ is given in [8] by

$$\mathrm{RB}\Big(\widehat{\lambda}_\alpha(x,y)\Big) = \begin{cases} \big(n\widehat{F}_X(x)(1-\alpha)+1\big)/n & \text{if } n\widehat{F}_X(x)\alpha = \lfloor n\widehat{F}_X(x)\alpha \rfloor, \\ \big(n\widehat{F}_X(x) - \lfloor n\widehat{F}_X(x)\alpha \rfloor\big)/n & \text{otherwise,} \end{cases}$$

where $\widehat{F}_X(x) := \widehat{H}_{n,XY}(x,0)$ denotes the empirical version of the marginal distribution function $F_X(x) := H_{XY}(x,0)$ of X. Of particular interest is the limiting regime $\alpha = \alpha(n) \to 1$ with $n(1-\alpha) \to \infty$ (meaning that the robust α-frontier is extreme but well inside the sample), where it holds with probability one that

$$\mathrm{RB}\Big(\widehat{\lambda}^\star_\alpha(x,y)\Big) \sim 1-\alpha \geqslant (1-\alpha)H_{XY}(x,0) \sim \mathrm{RB}\Big(\widehat{\lambda}_\alpha(x,y)\Big).$$

Hence, the proportion of adverse outliers that the conditional efficiency score $\widehat{\lambda}_\alpha(x,y)$ can tolerate is strongly driven by the input level x, whereas the alternative estimator $\widehat{\lambda}^\star_\alpha(x,y)$ enjoys a larger, more uniform robustness behavior, governed solely by the sample size n and the trimming level α.

Within the `frontiles` package, the `breakdown.point()` and `breakdown.point.star()` functions compute the corresponding breakdown measures $\mathrm{RB}\Big(\widehat{\lambda}_\alpha(x,y)\Big)$ and $\mathrm{RB}\Big(\widehat{\lambda}^\star_\alpha(x,y)\Big)$, respectively. To display the sample breakdown values at the two evaluation points P_1 and P_2 considered previously, we use the following instructions:

```
breakdown.point(xobs = matrix(xobs), yobs = matrix(yobs),
   xeval = matrix(x_new), yeval = matrix(y_new), alpha = 0.975)

P1   0.02173913
P2   0.04347826

breakdown.point.star(xobs = matrix(xobs), yobs = matrix(yobs),
   xeval = matrix(x_new), yeval = matrix(y_new), alpha = 0.975)

P1   0.04347826
P2   0.04347826
```

In practice, the ability to compute the gross-error sensitivity and the breakdown point is useful for comparing the two competing estimators, assessing their reliability in the presence of outliers or measurement errors, justifying the use of one of them from a robustness viewpoint, and tuning the trimming level α (or m for the weighted-moment procedure) by examining the trade-off between robustness and efficiency.

1.5.2 Weighted-moment partial efficiency scores

Both the conditional order-m estimator $\widehat{\lambda}_m(x,y)$ and its unconditional version $\widehat{\lambda}^{\star}_m(x,y)$ were found to be B-robust (see [7] and [13]) with the following gross-error sensitivity values:

$$\begin{aligned}
\mathrm{GES}\Big(\widehat{\lambda}_m(x,y)\Big) &= \frac{m}{F_X(x)} \cdot \max\big\{\lambda(x,y) - \lambda_m(x,y), \lambda_m(x,y) - \lambda_{m-1}(x,y)\big\},\\
\mathrm{GES}\Big(\widehat{\lambda}^{\star}_m(x,y)\Big) &= m \cdot \max\big\{\lambda(x,y) - \lambda^{\star}_m(x,y), \lambda^{\star}_m(x,y) - \lambda^{\star}_{m-1}(x,y)\big\}.
\end{aligned}$$

The troublesome border/divergence effect exhibited by the partial efficiency estimators $\widehat{\lambda}_m(x,y)$ and their associated frontiers, which stems from conditioning on the event $\{X \leqslant x\}$, is manifested through the appearance of potentially small values of $F_X(x)$ in the GES denominator. Again, the routines `gross.error.sensitivity()` and `gross.error.sensitivity.star()` can be used to obtain the corresponding estimates:

$$\begin{aligned}
\widehat{\mathrm{GES}}\Big(\widehat{\lambda}_m(x,y)\Big) &= \frac{m}{\widehat{F}_X(x)} \cdot \max\big\{\widehat{\lambda}(x,y) - \widehat{\lambda}_m(x,y), \widehat{\lambda}_m(x,y) - \widehat{\lambda}_{m-1}(x,y)\big\},\\
\widehat{\mathrm{GES}}\Big(\widehat{\lambda}^{\star}_m(x,y)\Big) &= m \cdot \max\big\{\widehat{\lambda}(x,y) - \widehat{\lambda}^{\star}_m(x,y), \widehat{\lambda}^{\star}_m(x,y) - \widehat{\lambda}^{\star}_{m-1}(x,y)\big\}.
\end{aligned}$$

To evaluate these estimates at the two points P_1 and P_2 considered previously, we use the following instructions:

```
gross.error.sensitivity(xobs = matrix(xobs), yobs = matrix(yobs),
  xeval = matrix(x_new), yeval = matrix(y_new), type = "ordermscore", m = 20)
```

```
P1  3.678808
P1  3.574237
```

```
gross.error.sensitivity.star(xobs = matrix(xobs), yobs = matrix(yobs),
  xeval = matrix(x_new), yeval = matrix(y_new), type = "ordermscore", m = 20)
```

```
P1  3.590447
P2  3.508565
```

As regards the finite sample breakdown point, similarly to Theorem 2.1 of [8], it is not hard to show that even a contamination by a single outlying observation is sufficient for breakdown of the L-statistics $\widehat{\lambda}_m(x,y)$ and $\widehat{\lambda}^{\star}_m(x,y)$, as is the case for the FDH estimator $\widehat{\lambda}(x,y)$, *i.e.*,

$$\mathrm{RB}\Big(\widehat{\lambda}_m(x,y)\Big) = \mathrm{RB}\Big(\widehat{\lambda}^{\star}_m(x,y)\Big) = \mathrm{RB}\Big(\widehat{\lambda}(x,y)\Big) = \frac{1}{n} \quad \text{for any order } m \geqslant 1.$$

In contrast, with a suitable choice of the tail probability level α, the associated quantile estimators $\widehat{\lambda}_{\alpha}(x,y)$ and $\widehat{\lambda}^{\star}_{\alpha}(x,y)$ can tolerate arbitrary contamination affecting a substantial fraction of the sample, without suffering catastrophic deterioration in performance. However, once these empirical quantiles break down at very high levels of α, their resistance to outliers can fall below that of competing order-m efficiency estimators (see *e.g.* [8]).

1.5.3 Duality between order-α and order-m efficiencies

Given that order-m efficiency scores are driven by tail expectations, whereas their order-α counterparts are governed by tail quantiles, a natural question in this literature on robust frontier and efficiency analysis is the following: Once the practitioner has fixed a trimming order m based on suitable considerations, what principled rule should be used to determine the associated level $\alpha = \alpha(m)$ so that the corresponding (un)conditional order-m and order-α efficiency scores estimate comparable benchmark functionals and can therefore be analyzed and contrasted on an equal footing—much like the classical mean–median duality in robust statistics?

The families $\{\lambda_m(x,y) : m \geqslant 1\}$ and $\{\lambda_\alpha(x,y) : \alpha \in (0,1]\}$—and likewise their unconditional counterparts $\{\lambda^\star_m(x,y) : m \geqslant 1\}$ and $\{\lambda^\star_\alpha(x,y) : \alpha \in (0,1]\}$—have been introduced in the econometric literature as two distinct notions of partial production-performance measures. Their estimators are therefore not directly comparable, since they target different population quantities, except in the limiting case $m \uparrow \infty$ and $\alpha \uparrow 1$, where both scores converge to the true full technical efficiency. Nevertheless, Proposition 2.2 of [8] shows that these concepts are tightly connected in the interesting sense that, for every $m \geqslant 1$ there exists a specific level $\alpha = \alpha(m) = (1/2)^{1/m}$ such that the population order-m and order-α scores correspond, respectively, to the mean and the median of the same distribution—namely, the distribution of the maximum/minimum of m i.i.d. draws Z_i^{xy} from the transformed variable $Z^{xy} = Z^{xy}(X,Y)$. This mean–median duality also helps explain why order-α estimators typically dominate their order-m counterparts in terms of finite-sample breakdown point and gross-error sensitivity, although this added robustness may come at the expense of statistical efficiency.

1.5.4 Trimming parameter selection

We now propose a simple diagnostic to help decide whether the benchmark should be the full envelopment support frontier or, instead, a robust partial frontier. For clarity, we focus on the "robustified" unconditional quantile-type efficiency measure $\widehat{\lambda}^\star_\alpha(x,y)$, for which a single tuning level α can be applied uniformly across all evaluation units (x,y). To this end, we consider the heuristic criterion

$$D^\star_n(\alpha) := \max_{1 \leqslant i \leqslant n} \Big\{ \widehat{\lambda}^\star_1(X_i,Y_i) - \widehat{\lambda}^\star_\alpha(X_i,Y_i) \Big\},$$

which quantifies, uniformly over the data points, the largest gap between the nonrobust FDH techical efficiency $\widehat{\lambda}^\star_1 \equiv \widehat{\lambda}$ and the resistant partial one $\widehat{\lambda}^\star_\alpha$. We track the evolution of $D^\star_n(\alpha)$ along the discretized grid $\alpha = \alpha(k) = 1 - k/n$, for $k = 0, \ldots, n-1$. In large samples of the order of several thousands of observations, it is typically enough to evaluate this diagnostic for k between 0 and *e.g.* 200 in order to keep the computational cost moderate. When the data contain no marked anomalies, the sequence $D^\star_n(\alpha(k))$ should evolve gently as a "staircase" in the integer index k. In that case, using the full efficiency scores associated with the extreme order $\alpha = \alpha(0) = 1$ (equivalently, $k = 0$) is the most efficient choice. Conversely, if the curve displays a pronounced upward jump between two indices, say k_ℓ and k_u, and then levels off for some values $k \geqslant k_u$, this pattern suggests the presence of potential outliers: the estimates $\widehat{\lambda}^\star_\alpha(X_i,Y_i)$ remain globally robust for orders $\alpha \leqslant \alpha(k_u)$, but would break down at $\alpha(k_u - 1)$. In such a situation, a prudent robust choice is to take the threshold $\alpha = \alpha(k_u)$, for which $\widehat{\lambda}^\star_\alpha(\cdot,\cdot)$ still reflects the magnitude of valuable extreme firms while remaining resistant to the effect of isolated outliers.

It is the `diagnostic.trim()` routine within the `frontiles` package which implements the diagnostic $k \mapsto D^\star_n(\alpha(k))$ for the unconditional order-α estimator $\widehat{\lambda}^\star_\alpha$, but also the

diagnostic $k \mapsto D_n(\alpha(k)) := \max_{1 \leqslant i \leqslant n} \left\{ \widehat{\lambda}(X_i, Y_i) - \widehat{\lambda}_{\alpha(k)}(X_i, Y_i) \right\}$ for the conditional version $\widehat{\lambda}_\alpha$. For the motivating `atp` dataset, we use the instructions

```
diagnostic.trim(xobs = matrix(xobs), yobs = matrix(yobs), type = "alphascore")
diagnostic.trim(xobs = matrix(xobs), yobs = matrix(yobs), type = "alphascore.star")
```

to get the diagnostic curves shown in Figure 1.13: Moving from left to right, the function $k \mapsto D_n(\alpha(k))$ in the top-left panel displays a sharp increase when k moves from $k_\ell = 1$ to $k_u = 2$ (indicated by the vertical purple line), followed by a relatively mild increase as k grows further to 8 (indicated by the vertical green line). The choices $k = 2$ and $k = 8$ correspond to $\alpha(2) = 1 - k/n = 0.9565$ and $\alpha(8) = 0.8260$, respectively. As can be seen in the bottom-left panel, the corresponding conditional quantile frontiers, represented by the purple and green curves, trace the most efficient observations, with the purple curve remaining more conservative and the green curve deviating more substantially from the largest observations. These features support selecting the tail probability level as $\alpha = \alpha(2)$. Although the resulting purple frontier is fairly robust to the influence of the top players, including Nadal and Federer, it still envelops the bulk of the sample—namely, 39 out of 46 observations—while failing to satisfy the desired monotonicity axiom. Turning to the diagnostic function $k \mapsto D_n^\star(\alpha(k))$ for the unconditional order-α frontier estimator in the top-right panel, we observe a large initial jump as k increases from $k_\ell = 1$ to $k_u = 2$ (marked by the vertical purple line), followed by a smaller upward jump when k is increased to 4 (vertical green line). A further substantial jump appears between $k = 14$ and $k = 15$, although it is less pronounced than the initial one and occurs much farther to the right. Setting $k = 2$ and $k = 4$ yields $\alpha(k) = 0.9565$ and 0.9130, respectively. The corresponding unconditional quantile frontiers are graphed in purple and green curves in the bottom-right panel. Our recommendation is to choose the trimming level as $\alpha = \alpha(k_u) \equiv \alpha(2)$, corresponding to the initial sharp jump. This choice mitigates both the boundary and monotonicity defects of the conditional approach, while still yielding a reasonably robust frontier for efficiency benchmarking.

For the `movies` dataset, the diagnostic checks in Figure 1.14 suggest selecting the tail probability levels $\alpha(3) = 0.9444$ for the conditional quantile-type frontier and $\alpha(1) = 0.9814$ for the unconditional counterpart. Finally, for the `mfi` dataset, the diagnostic curves in Figure 1.15 point towards the choice of $\alpha(20) = 0.9958$ for the conditional frontier and $\alpha(17) = 0.9964$ for the unconditional counterpart.

1.6 Concluding remarks

This chapter has presented a unified overview of robust nonparametric methods for frontier and efficiency analysis, with particular emphasis on quantile-based order-α approaches, weighted-moment order-m approaches, and their robustified unconditional counterparts. These notions of partial benchmark frontiers and efficiency measures offer flexible and robust alternatives to full envelopment techniques, especially in the presence of extreme observations, outliers, or sparsely populated boundary regions.

While the conditional order-α and order-m frontiers already provide useful partial measures of performance, their unconditional robustified analogs often yield more satisfactory benchmark frontiers by improving monotonicity, reducing the undue influence of extremes, and offering a more coherent global representation of the attainable set. Rather than com-

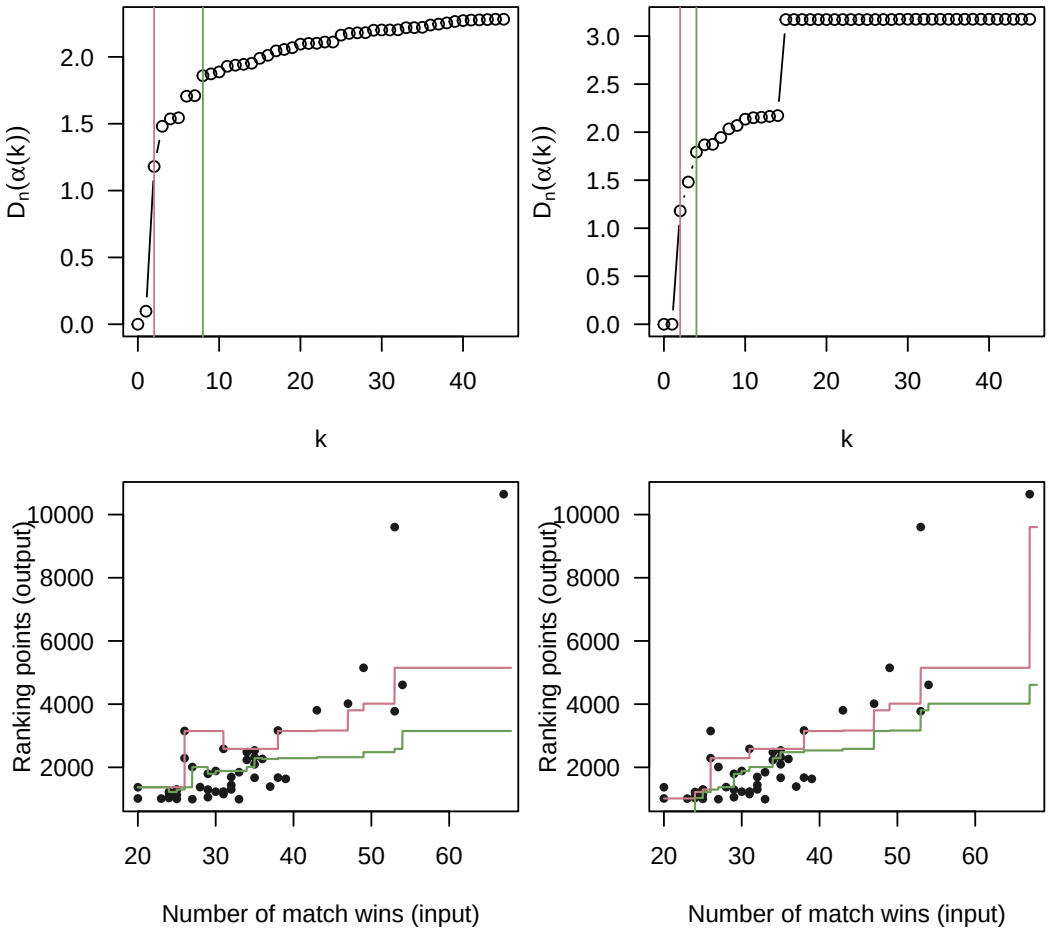

FIGURE 1.13
Results for the `atp` dataset: (Top panels) Diagnostic curves $k \mapsto D_n(\alpha(k))$ on the left and $k \mapsto D_n^\star(\alpha(k))$ on the right; (Bottom left) The conditional order-$\alpha(k)$ frontier estimators for $k = 2$ (purple) and $k = 8$ (green); (Bottom right) The unconditional order-$\alpha(k)$ frontier estimators for $k = 2$ (purple) and $k = 4$ (green).

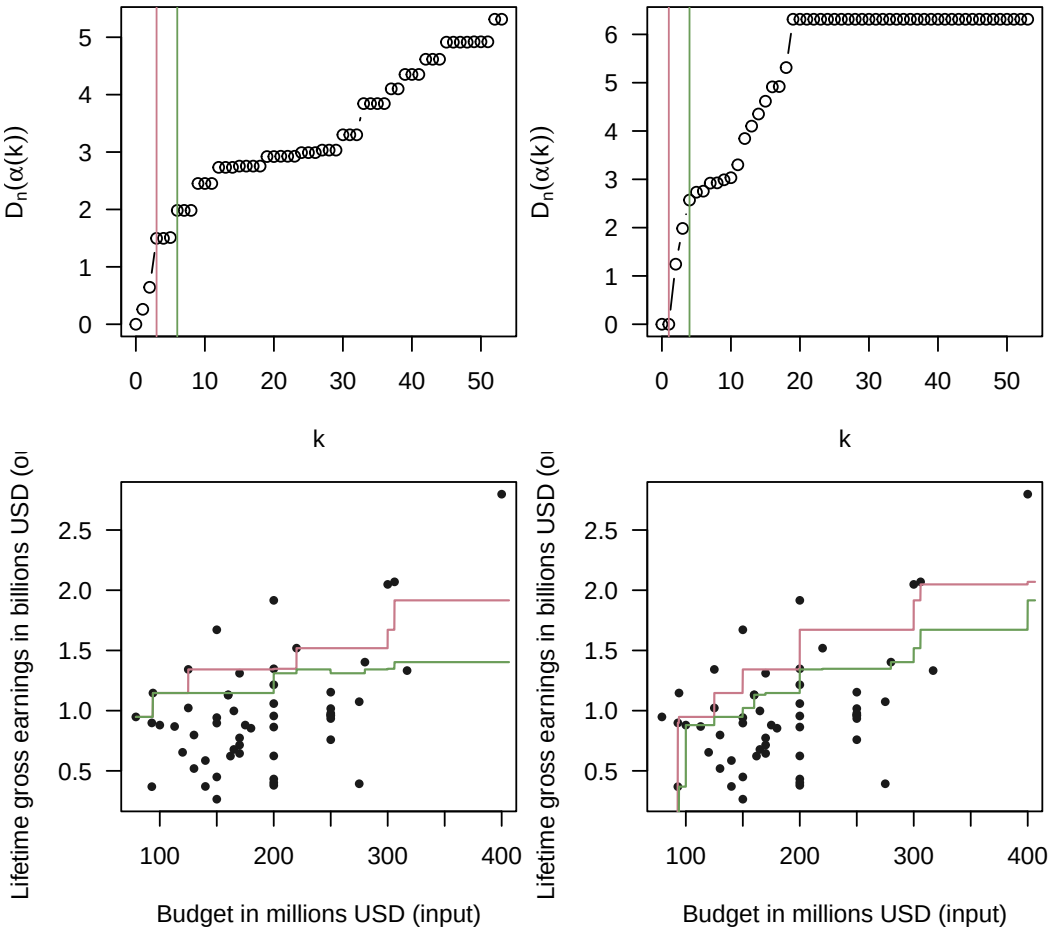

FIGURE 1.14
Results for the `movies` dataset: (Top panels) Diagnostic curves $k \mapsto D_n(\alpha(k))$ on the left and $k \mapsto D_n^\star(\alpha(k))$ on the right; (Bottom left) The conditional order-$\alpha(k)$ frontier estimators for $k = 3$ (purple) and $k = 6$ (green); (Bottom right) The unconditional counterparts for $k = 1$ (purple) and $k = 4$ (green).

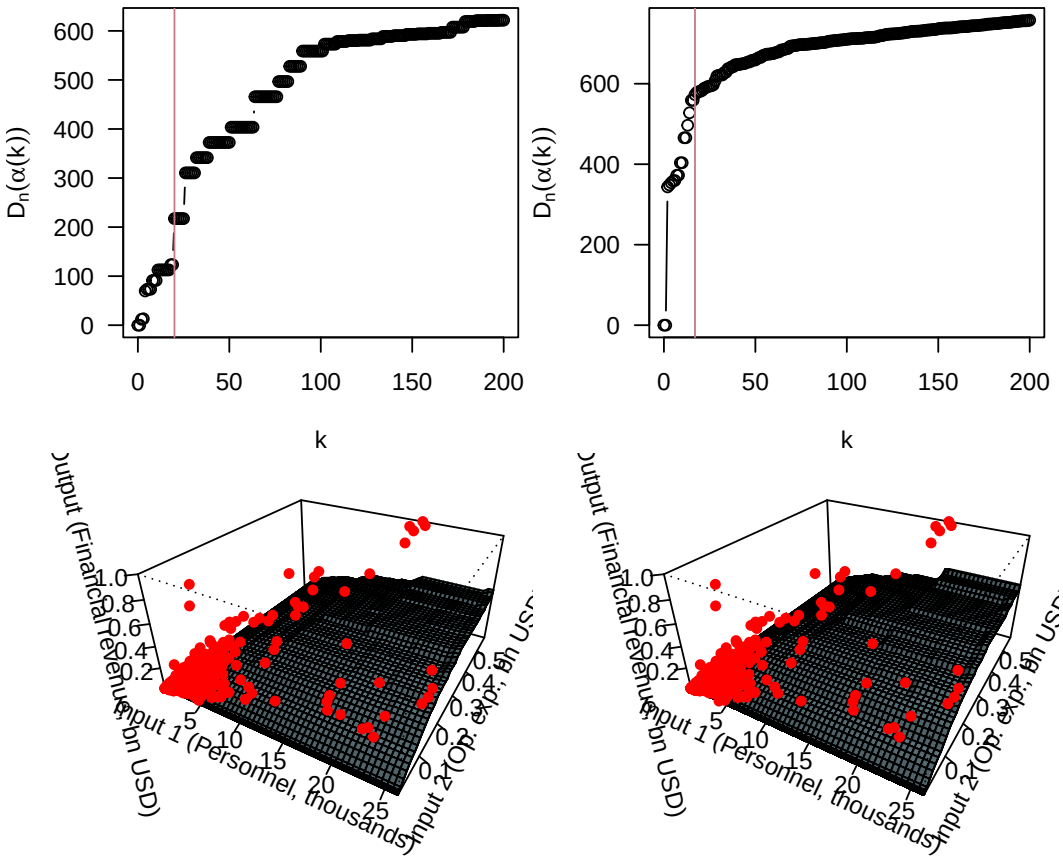

FIGURE 1.15
Results for the `mfi` dataset: (Top panels) Diagnostic curves $k \mapsto D_n(\alpha(k))$ on the left and $k \mapsto D_n^\star(\alpha(k))$ on the right; (Bottom left) The conditional order-$\alpha(k)$ surface for $k = 20$; (Bottom right) The unconditional counterpart for $k = 17$.

peting methods, the quantile- and weighted-moment-based approaches are best viewed as complementary tools, reflecting different points along the robustness–efficiency trade-off.

The `frontiles` package brings these methodological advances together in a coherent and user-friendly framework. It makes it possible to compute, visualize and compare robust frontier and efficiency estimators in several orientations, to assess their sensitivity through confidence intervals and robustness diagnostics, and to guide the choice of trimming parameters in empirical applications. The examples considered in this chapter illustrate that these tools can provide informative and stable efficiency benchmarks in a variety of settings.

Overall, robust partial frontier methods now form a mature and practically relevant toolbox for modern efficiency analysis. We hope that the combination of theoretical insight, diagnostic guidance and software implementation presented here will facilitate their broader use in applied work and stimulate further developments in high-dimensional and more complex production environments. In our ongoing research agenda, we also plan to extend the `frontiles` package in two important directions. The first is to incorporate regression frameworks with environmental covariates that affect production performance but are not under the control of firms. The second concerns situations in which prior knowledge, or the diagnostic checks developed in Section 1.5.4, indicate the absence of potential outliers in the data. In such settings, the constrained spline smoothing approach of [14] offers particularly innovative envelopment fits under free disposability and/or monotone concavity. For computational convenience, however, the most promising cubic spline smoother is currently implemented in the `npbr` package only through a sufficient condition for monotonicity; see Section 3.1, p. 17 of [14]. In the present development of `frontiles`, we are working toward a satisfactory resolution of this limitation by imposing monotonicity through standard second-order cone constraints, using the exact formulation proposed in [14], which provides an equivalent characterization of monotonicity.

Acknowledgments

This research was supported by the French National Research Agency under the grant ANR-17-EURE-0010 (EUR CHESS) and the TSE-HEC ACPR Chair “Regulation and systemic risks”.

Bibliography

[1] Yves Aragon, Abdelaati Daouia, and Christine Thomas-Agnan. Nonparametric frontier estimation: A conditional quantile-based approach. *Econometric Theory*, 21:358–389, 2005.

[2] Catherine Cazals, Jean-Pierre Florens, and Léopold Simar. Nonparametric frontier estimation: a robust approach. *Journal of Econometrics*, 106:1–25, 2002.

[3] Abraham Charnes, William W. Cooper, and Edwardo Rhodes. Evaluating program and managerial efficiency: An application of data envelopment analysis to program follow through. *Management Science*, 27(6):668–697, 1981.

[4] E. Culwell. moviefranchises. `https://data.world/eculwell/moviefranchises`, 2022. data.world dataset. Accessed 2025-06-05.

[5] Abdelaati Daouia, Jean-Pierre Florens, and Léopold Simar. Frontier estimation and extreme value theory. *Bernoulli*, 16:1039–1063, 2010.

[6] Abdelaati Daouia, Jean-Pierre Florens, and Léopold Simar. Regularization of nonparametric frontier estimators. *Journal of Econometrics*, 168:285–299, 2012.

[7] Abdelaati Daouia, Jean-Pierre Florens, and Léopold Simar. Robustified expected maximum production frontiers. *Econometric Theory*, 37:346–387, 2021.

[8] Abdelaati Daouia and Irène Gijbels. Robustness and inference in nonparametric partial frontier modeling. *Journal of Econometrics*, 161:147–165, 2011.

[9] Abdelaati Daouia, Stéphane Girard, and Armelle Guillou. A Γ-Moment Approach to Monotonic Boundary Estimation. *Journal of Econometrics*, 178:727–740, 2014.

[10] Abdelaati Daouia and Thibault Laurent. *`frontiles`: A Package for Robust Frontier and Efficiency Analysis in R*, 2026. R package version 1.3.1. `https://CRAN.R-project.org/package=frontiles`.

[11] Abdelaati Daouia, Thibault Laurent, and Hohsuk Noh. `npbr`: A package for nonparametric boundary regression in R. *Journal of Statistical Software*, 79(9):1–43, 2017.

[12] Abdelaati Daouia, Thibault Laurent, and Hohsuk Noh. *`npbr`: Nonparametric Boundary Regression in R*, 2023. R package version 1.8. `https://CRAN.R-project.org/package=npbr`.

[13] Abdelaati Daouia, Carlos Martins-Filho, and Léopold Simar. A multivariate weighted-moment measure of production performance. Manuscript, 2026.

[14] Abdelaati Daouia, Hohsuk Noh, and Byeong U. Park. Data envelope fitting with constrained polynomial splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(1):3–30, 2016.

[15] Abdelaati Daouia and Byeong U. Park. On projection-type estimators of multivariate isotonic functions. *Scandinavian Journal of Statistics*, 40:363–386, 2013.

[16] Abdelaati Daouia and Anne Ruiz-Gazen. Robust nonparametric frontier estimators: Qualitative robustness and influence function. *Statistica Sinica*, 16:1233–1253, 2006.

[17] Abdelaati Daouia and Léopold Simar. Robust nonparametric estimators of monotone boundaries. *Journal of Multivariate Analysis*, 96:311–331, 2005.

[18] Abdelaati Daouia and Léopold Simar. Nonparametric efficiency analysis: a multivariate conditional quantile approach. *Journal of Econometrics*, 140:375–400, 2007.

[19] Abdelaati Daouia, Léopold Simar, and Paul W. Wilson. Measuring firm performance by using nonparametric quantile-type distances. *Econometric Reviews*, 36:156–181, 2017.

[20] Cinzia Daraio and Léopold Simar. *Advanced Robust and Nonparametric Methods in Efficiency Analysis: Methodology and Applications*. Springer, New York, 2007.

[21] Cinzia Daraio, Léopold Simar, and Paul W. Wilson. Fast and efficient computation of directional distance estimators. *Annals of Operations Research*, 288(2):805–835, 2020.

[22] Dominique Deprins, Léopold Simar, and Henry Tulkens. Measuring labor-efficiency in post offices. In Maurice Marchand, Pierre Pestieau, and Henry Tulkens, editors, *The Performance of Public Enterprises: Concepts and Measurement*, pages 243–267. North-Holland, Amsterdam, 1984.

[23] David L. Donoho and Peter J. Huber. The notion of breakdown point. In Peter J. Bickel, Kjell A. Doksum, and Joseph L. Hodges, editors, *A Festschrift for Erich L. Lehmann*. Wadsworth, Belmont, CA, 1983.

[24] Fatou S. Fall, Henri Tchakoute Tchuigoua, Anne Vanhems, and Léopold Simar. Gender effect on microfinance efficiency: A robust nonparametric approach. *European Journal of Operational Research*, pages 744–757, 2021.

[25] Frank R. Hampel, Elvezio M. Ronchetti, Peter J. Rousseeuw, and Werner A. Stahel. *Robust Statistics: The Approach Based on Influence Functions*. John Wiley & Sons, New York, 1986.

[26] Seok-Oh Jeong and Léopold Simar. Linearly interpolated FDH efficiency score for nonconvex frontiers. *Journal of Multivariate Analysis*, 97(10):2141–2161, 2006.

[27] Peter J. Rousseeuw. A new infinitesimal approach to robust estimation. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 56:127–132, 1981.

[28] Robert J. Serfling. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, New York, 1980.

[29] Ronald W. Shephard. *Theory of Cost and Production Functions*. Princeton University Press, Princeton, NJ, 1970.

[30] Léopold Simar and Paul W. Wilson. Estimation and inference in nonparametric frontier models: Recent developments and perspectives. *Foundations and Trends in Econometrics*, 5(3–4):183–337, 2013.

[31] Léopold Simar and Paul W. Wilson. Statistical approaches for non-parametric frontier models: A guided tour. *International Statistical Review*, 83(1):77–110, 2015.

[32] David C. Wheelock and Paul W. Wilson. Non-parametric, unconditional quantile estimation for efficiency analysis with an application to federal reserve check processing operations. *Journal of Econometrics*, 145:209–225, 2008.

[33] P. W. Wilson. Asymptotic properties of some non-parametric hyperbolic efficiency estimators. In Ingrid Van Keilegom and P. W. Wilson, editors, *Exploring Research Frontiers in Contemporary Statistics and Econometrics*, pages 115–150. Springer-Verlag, Berlin, 2011.

[34] Paul W. Wilson. Fear: A software package for frontier efficiency analysis with r. *Socio-Economic Planning Sciences*, 42(4):247–254, 2008.