Contemporary Mathematics
https://ojs.wiserpub.com/index.php/CM/

Research Article

Efficiency as a Quantile Process: Functional Profiles and Panel Heterogeneity in Firm Data

Vladislav Spitsin*, **Lubov Spitsina**, **Magerram Gasanov**, **Victoria Leonova**

Business School, Tomsk Polytechnic University, Tomsk, 634050, Russia
E-mail: spicynvv@tpu.ru

Received: 15 May 2026; **Revised:** 16 June 2026; **Accepted:** 26 June 2026

Abstract: This paper develops a functional extension of quantile-based frontier efficiency analysis for firm-level panel data. Instead of representing technical efficiency by a single benchmark-specific score, it defines a quantile-indexed efficiency profile that tracks how a firm-year observation changes as the benchmark moves across an ordered family of upper conditional quantiles. The framework is formulated in a log-log quantile frontier model with non-crossing benchmark curves, making it possible to study efficiency as a structured profile rather than an isolated scalar. Three interpretable profile summaries are introduced: the average profile level, measuring overall proximity to the frontier over the selected quantile interval; the tail fragility index, capturing the deterioration of relative performance under the most demanding upper-tail benchmarks; and the decline rate, summarizing profile steepness. The paper establishes key theoretical properties of the profile and its summaries, including boundedness, monotonicity with respect to benchmark strictness, dominance ordering, continuity, endpoint and approximation results, curvature-based shape interpretation, local classification stability, and plug-in consistency under uniform convergence of the estimated frontier. Empirically, the framework is applied to 2,464 Russian firms observed over 2019-2023. The estimated profiles are nondegenerate, heterogeneous in shape, and support an economically interpretable classification into frontier-near, intermediate, tail-fragile, and uniformly lagging types. Comparisons with endpoint benchmarking and Data Envelopment Analysis (DEA)-based scalar efficiency scores show that the profile contains information not recoverable from standard one-number or discrete quantile-specific measures. Fixed-effects regressions indicate that the average profile level behaves as a level-type efficiency measure, whereas tail fragility captures a distinct shape-related dimension. Robustness checks based on alternative quantile ranges, classification thresholds, model specifications, bootstrap confidence intervals, and lagged performance regressions support the empirical relevance of the proposed representation. Overall, replacing one-number efficiency measures with quantile efficiency profiles yields a richer and more informative representation of firm heterogeneity.

Keywords: quantile frontier, technical efficiency, efficiency profile, non-crossing quantile process, functional data analysis, firm heterogeneity

MSC: 62J05, 62G20, 62P20, 91B38

Copyright ©2026 Vladislav Spitsin, et al.
DOI: https://doi.org/10.37256/cm.75202610170
This is an open-access article distributed under a CC BY license
(Creative Commons Attribution 4.0 International License)
https://creativecommons.org/licenses/by/4.0/

1. Introduction

Efficiency measurement remains one of the central themes in production analysis, operations research, and applied statistics. The dominant frontier traditions—most notably Data Envelopment Analysis (DEA) and Stochastic Frontier Analysis (SFA)—have generated a large methodological literature and a broad range of empirical applications. Recent reviews show that frontier-based efficiency analysis continues to expand in both methodological sophistication and applied scope, including economic, financial, environmental, and operational settings [1–3]. At the same time, a common feature of much of this literature is that firm-level technical efficiency is ultimately reduced to a single score, even when the underlying benchmark relation is heterogeneous.

This scalar representation is useful for ranking firms and for second-stage regressions, but it can be too restrictive in heterogeneous firm populations. Consider two firms that look similar when evaluated against a moderately high benchmark, for example the 0.70 conditional quantile. If the benchmark is moved to the 0.95 conditional quantile, one firm may remain close to the frontier, while the other may fall sharply behind. A single efficiency score cannot distinguish these two cases. The first firm is robust to stricter benchmarking, whereas the second is fragile in the upper tail. This simple example illustrates the central motivation of the paper: efficiency may depend not only on how close a firm is to a benchmark, but also on how its relative position changes as the benchmark becomes progressively more demanding.

Quantile-based frontier methods are particularly attractive in this setting because they allow the benchmark to vary across the conditional distribution rather than forcing all firms to be evaluated against a single central tendency. This insight has been developed in both methodological and survey-oriented work on quantile stochastic frontier analysis [4, 5], while recent applied studies confirm that quantile choice matters for estimated efficiency and for the interpretation of technological heterogeneity [6–8]. However, even when several quantiles are estimated, the empirical endpoint often remains scalar: one selected quantile-specific efficiency score, or a small set of benchmark-specific scores reported side by side. Such multi-quantile benchmarking is useful for sensitivity analysis, but it is not the same as treating efficiency itself as a function of benchmark strictness.

The distinction is important. A set of separate quantile-specific scores answers the question of how efficiency changes at selected benchmark points. A quantile-indexed profile answers a different question: how the entire relative position of a firm-year observation evolves as the benchmark moves through an ordered upper-tail quantile process. This requires the benchmark family to be coherent. If estimated quantile frontiers cross, then higher quantiles no longer correspond to weakly more demanding benchmarks, and the economic interpretation of benchmark strictness breaks down. Recent developments in non-crossing quantile estimation therefore provide an important statistical foundation for ordered frontier processes [9], while functional data analysis provides the conceptual language for treating curves and profiles as primary statistical objects rather than as collections of unrelated scalar measurements [10, 11]. Related work on non-crossing convex quantile regression further reinforces the idea that quantiles should be treated jointly as an ordered object rather than as separately estimated benchmark points [12].

The present paper develops this idea in a log-log frontier setting for firm-level panel data. Rather than assigning firm i in period t a single efficiency value, we define a quantile-indexed efficiency profile

$$TE_{it}(\tau),\ \tau \in \mathscr{T}. \tag{1}$$

This profile records how the relative benchmark position of a firm-year observation changes as the benchmark moves from moderately high to increasingly demanding upper quantiles. Conceptually, the profile is the primary object; mathematically, it is represented as a function of quantile strictness; graphically, it can be visualized as a curve. This shift is modest at the level of notation, but substantial in its empirical interpretation. Instead of asking only how close a firm is to one benchmark, the profile asks whether the firm remains close to the frontier as the benchmark becomes stricter.

To summarize the profile while preserving its main economic content, we introduce three scalar functionals: the average profile level, the tail fragility index, and the decline rate. The average profile level measures overall proximity to the frontier over the selected quantile interval. The tail fragility index captures how strongly the firm's relative position

deteriorates when the benchmark moves into the most demanding upper-tail region. The decline rate summarizes the overall steepness of the profile. Together, these summaries separate firms that are uniformly close to the frontier from firms that appear competitive only under moderate benchmarks but become fragile when compared with the strongest performers.

The paper contributes on both the methodological and empirical sides. Methodologically, it does not merely estimate efficiency at several quantiles, as in standard multi-quantile benchmarking exercises. Instead, it treats the ordered quantile frontier family as the basis for a firm-year efficiency profile and studies the mathematical properties of the resulting object. Within this framework, we show that the profile is well defined and bounded, is monotone with respect to benchmark strictness under non-crossing quantile frontiers, induces a natural dominance ordering across firms, and yields scalar summaries that are continuous functionals of the underlying quantile process. We also establish plug-in consistency of the proposed summaries under uniform convergence of the estimated log-quantile frontier. In addition, we derive profile-specific results linking the summaries to endpoint efficiencies, establish explicit grid-based approximation bounds, and show that the relation between tail fragility and decline rate admits a geometric interpretation through the curvature of the profile.

Empirically, we apply the framework to a panel of Russian enterprises observed over 2019-2023. The empirical analysis asks three main questions. First, do firm-year efficiency profiles display substantial shape heterogeneity beyond what can be captured by a single-point efficiency score? Second, can firm-year observations be grouped into economically interpretable profile types? Third, are the resulting profile summaries systematically related to firm characteristics and subsequent performance in a way that supports meaningful second-stage analysis? These questions are motivated by recent evidence showing that distance-to-frontier measures are informative for firm-level productivity and innovation outcomes [13], and by methodological work emphasizing that second-stage efficiency analysis may be sensitive to how the first-stage efficiency object is constructed [14–16].

The empirical contribution is therefore not to revisit standard scalar-efficiency regressions with another benchmark-specific score. Rather, the paper examines whether replacing one-number efficiency measures with profile-derived summaries yields a richer representation of firm heterogeneity. The results show that estimated profiles are nondegenerate, economically interpretable, and heterogeneous in shape. The average profile level behaves similarly to conventional level-type efficiency measures, while tail fragility captures an additional shape-related dimension. Robustness checks based on alternative quantile ranges, DEA-based scalar benchmarks, classification stability, model-specification changes, bootstrap confidence intervals, and lagged performance regressions support the empirical relevance of the profile representation.

The paper is organized as follows. Section 2 reviews the relevant literature on frontier efficiency measurement, quantile-based frontier methods, and functional representations of heterogeneous statistical objects. Section 3 introduces the theoretical framework, including the log-log quantile frontier, the quantile efficiency profile, the associated scalar summaries, and the grid-based estimation setup. Section 4 develops the mathematical properties of quantile efficiency profiles, including boundedness, monotonicity, continuity, dominance ordering, stability, consistency, approximation, and shape-related results. Section 5 describes the data, estimation strategy, and econometric implementation. Section 6 presents the empirical results, including profile visualization, classification, robustness checks, and exploratory validation. Section 7 discusses the methodological and practical implications of the findings. Section 8 concludes.

2. Literature review

2.1 *Frontier efficiency analysis and the limits of single-score benchmarking*

Efficiency measurement remains one of the most active meeting points of applied mathematics, mathematical statistics, operations research, and empirical production analysis. The dominant traditions are DEA and SFA. Both traditions have evolved far beyond their original formulations: DEA now includes large families of cost, revenue, and profit efficiency models, while SFA has developed into a broad framework with panel specifications, latent heterogeneity, environmental adjustments, and more flexible distributional structures. The breadth of the modern DEA-based economic-efficiency literature, both methodologically and empirically, is well documented in recent surveys [1, 2].

This broader development is important because it shows that frontier analysis is no longer merely a device for assigning one benchmark score to each observation. It is increasingly concerned with how benchmark relations behave under heterogeneity, scale dispersion, and complex institutional environments [2, 3]. At the same time, a striking common feature of much of the frontier literature is that the final object assigned to each firm remains a single efficiency value. This scalarization is undeniably useful: it supports ranking, summary comparison, and second-stage regressions. Yet it also compresses information. In many empirical settings, especially with firm-level accounting data, an observation can be relatively close to one benchmark standard and far from another. The notion of "distance to the frontier" therefore depends not only on inputs and outputs, but also on the benchmark strictness embodied in the modelling choice.

A single score can conceal this dimension of heterogeneity. This is especially relevant in firm-level panels where one often observes many small and medium firms together with a thinner right tail of very large enterprises and a separate upper tail of high-performing firms. In such settings, benchmark difficulty is not an incidental modelling detail but a substantive feature of the data-generating environment. More generally, recent reviews of frontier-based economic-efficiency analysis emphasize that the field has become highly diversified in models and applications, while scalar efficiency summaries still remain the dominant empirical endpoint [1].

This limitation of one-number efficiency summaries has practical consequences. First, a single score cannot distinguish between firms that are uniformly near the upper boundary and firms that are only near a moderate benchmark but deteriorate sharply relative to the strongest performers. Second, scalarization may flatten the geometry of efficiency heterogeneity, making it harder to classify firms into meaningful profile types. Third, if the scalar score is then used in second-stage analysis, the interpretation of downstream regressions depends heavily on the chosen benchmark. These issues are not unique to one estimation tradition; they arise in both deterministic and stochastic benchmarking frameworks whenever heterogeneous performance is mapped into one scalar output.

A related issue concerns the role of second-stage efficiency analysis. In many applied studies, efficiency is first estimated and then regressed on firm characteristics or linked to outcomes such as profitability or growth. This strategy is useful but also fragile, because it depends on how the initial score is defined and on whether the scalar score captures all relevant variation in benchmark proximity. Recent methodological discussion continues to emphasize that downstream conclusions may be sensitive to the form of the efficiency measure and to the structure imposed in the second stage [14]. For the present paper, this motivates moving from benchmark-specific efficiency scores to a profile-based first-stage object that preserves more of the geometry of benchmark heterogeneity before any low-dimensional reduction is taken to panel regressions.

2.2 *Quantile frontiers, shape restrictions, and non-crossing estimation*

A natural way to enrich frontier analysis is to replace a single benchmark relation with an ordered family of quantile-specific frontiers. Quantile-based frontier methods are attractive because they define benchmarking relative to conditional upper-tail performance rather than to a conditional mean or to a deterministic envelope. In this sense, the benchmark becomes probabilistic and explicitly ordered by strictness. This is especially useful when the data exhibit substantial heterogeneity, outliers, or a meaningful distinction between moderate and highly demanding benchmark standards. Quantile stochastic frontier analysis provides an important formal basis for this perspective [17].

More recent work has expanded the area toward environmental applications, endogeneity, and more flexible quantile-based frontier formulations, reinforcing the point that the chosen quantile level is itself a substantive modelling dimension rather than a nuisance parameter [6–8]. This literature is directly relevant for the present study because it shows that efficiency may vary systematically with benchmark strictness. At the same time, much of this work still treats the quantile dimension mainly as a source of alternative benchmark-specific scores. In other words, several quantiles may be estimated, but the empirical endpoint typically remains one selected score or a small set of separate efficiency values.

A second line of development concerns economically coherent estimation under shape restrictions. In production analysis, benchmarking relations are rarely unconstrained: one often wants monotonicity, curvature restrictions, or economically admissible bounds on returns to scale. In the quantile setting, these requirements become intertwined with the ordered nature of the quantile family. Early work on nonparametric production-frontier estimation has considered economically motivated shape constraints, including monotonicity and concavity, in additive frontier functions [18].

Outside frontier analysis, recent work on quantile regression has developed methods for jointly estimating multiple non-crossing quantile curves under shape constraints [19].

For the present paper, this literature matters for two reasons. First, our profile construction assumes that the estimated frontier process is economically coherent: higher quantiles must correspond to weakly more demanding benchmark frontiers, and the fitted relation must obey restrictions compatible with production analysis. Second, once quantiles are interpreted as an ordered family rather than as unrelated estimation targets, efficiency can be indexed by benchmark strictness in a coherent way.

A third and closely related issue is quantile crossing. If quantiles are estimated pointwise and independently, fitted curves may cross, violating the basic ordering property of quantiles. In production and benchmarking contexts, this is not merely a technical inconvenience. It undermines the interpretation of upper-tail strictness and makes any quantile-indexed efficiency object incoherent. A firm cannot be meaningfully evaluated against an "ordered" benchmark family if the estimated quantiles fail to preserve the benchmark ordering itself.

Recent statistical work therefore treats the estimation of a non-crossing quantile process as a central objective. In particular, modern contributions develop fully nonparametric non-crossing quantile regression processes with theoretical guarantees, reinforcing the methodological idea that quantiles should be estimated as members of an ordered family rather than as unrelated pointwise objects [9]. Related constrained log-log quantile frontier work applies this same ordered benchmarking logic in a shape-restricted frontier setting and emphasizes coherent calibration and interpretation in empirical efficiency analysis [20]. The specific estimator used in the present paper is different, but the underlying lesson is the same: once the benchmark is viewed as an ordered quantile process, it becomes natural to ask whether the resulting firm-level efficiency object should itself also be treated as ordered and functional rather than merely benchmark-specific.

This distinction clarifies both the connection to and the difference from existing multi-quantile frontier studies. Applied quantile-frontier studies show that benchmark choice matters and may reveal heterogeneity that is not visible under mean-based or single-frontier specifications [6–8]. Methodological work on ordered and constrained quantile estimation further emphasizes non-crossing and shape coherence across the quantile family [9, 19, 20]. However, the usual empirical use of multiple quantiles remains comparative: researchers estimate several benchmark-specific scores, examine their sensitivity, and often return to one preferred scalar measure. The present paper takes a different step. It uses the ordered quantile frontier family not only as a set of alternative benchmarks, but as the domain on which a firm-year efficiency profile is defined. In this sense, the quantile dimension is not only a robustness dimension; it becomes the primary index of benchmark strictness.

Table 1 summarizes the main efficiency-analysis approaches most closely related to the present study and clarifies the specific role played by the proposed profile-based framework.

Table 1. Comparison of the proposed profile-based framework with related efficiency-analysis approaches

Approach	Central empirical object	Role of quantile level	Role of shape/non-crossing constraints	Output used in second stage
Standard DEA/SFA application	One score per firm-year	Usually absent	May apply in the production set or parametric frontier	One scalar efficiency measure
Single-quantile frontier application	One score tied to one chosen quantile	Selected benchmark level	Sometimes imposed	One scalar efficiency measure
Multi-quantile frontier robustness analysis	Several benchmark-specific scores across selected quantiles	Alternative benchmark levels used for comparison or robustness	Sometimes imposed jointly	Usually one preferred scalar score or separate benchmark-specific comparisons
Proposed profile approach	One quantile-indexed efficiency function for each firm-year	Primary index of benchmark strictness defining the profile	Central to coherent profile construction	Curve summaries and optional profile classification

Source: Compiled by the authors based on research materials

Table 1 highlights that the proposed framework differs from related approaches mainly in its primary empirical object. Standard DEA/SFA applications, single-quantile frontier models, and multi-quantile robustness exercises all ultimately organize the empirical analysis around one score or a small set of benchmark-specific scores. By contrast, the profile approach treats efficiency as a function indexed by benchmark strictness and only then derives scalar summaries for interpretation and second-stage analysis. This sequence is important: the profile is not a decorative visualization of several scores, but the object from which the level, fragility, decline rate, and classification results are derived.

2.3 *From multi-quantile estimation to functional efficiency profiles*

Although the statistical literature on multiple quantiles is already extensive, the move from estimating multiple quantiles to representing efficiency as a function is conceptually nontrivial. In much of the existing quantile-frontier literature, the researcher estimates several quantiles either for robustness or to compare benchmark sensitivity. But the resulting empirical output is still usually a set of separate benchmark-specific efficiency values. This is not the same as treating efficiency itself as a function indexed by quantile strictness.

The difference is fundamental. A collection of quantile-specific scores remains a finite set of values attached to selected benchmark levels. A profile, by contrast, treats the ordered quantile dimension as part of the object itself. In that formulation, the question is no longer simply how efficient a firm is at one benchmark or another, but how its relative position evolves as the benchmark becomes progressively stricter. Once the quantile family is regarded as ordered and coherent, the natural mathematical object becomes a function of quantile level rather than a list of unrelated numbers.

The functional-data perspective can be understood in a simple non-technical way. If efficiency is estimated at several ordered quantile levels, these values are not merely separate numbers; they can be viewed as points on a curve. Functional Data Analysis (FDA) provides the language for interpreting such a curve through its level, slope, and shape, and recent overviews show the breadth of this framework for smoothing, classification, regression, and modern machine-learning extensions [10, 11]. In the present setting, the level of the curve reflects overall proximity to the benchmark, the slope reflects deterioration as the benchmark becomes stricter, and the upper-tail part of the curve reflects whether the firm remains close to the strongest performers. The important difference from standard FDA is that the curve is not directly observed. It is induced by an estimated ordered quantile frontier process. This is why non-crossing constraints and frontier coherence are essential for the economic interpretation of the profile.

At the same time, the object studied here differs from standard FDA applications in an important respect. In classical FDA, curves are usually observed directly or obtained by smoothing raw repeated measurements. Here, by contrast, the efficiency profile is a derived object induced by an estimated quantile frontier process. This gives rise to a specific methodological problem: one must justify not only the visual usefulness of the curve, but also its mathematical coherence as a transformation of the underlying frontier family. That is why the paper places substantial emphasis on Sections 3 and 4. The goal is not only to introduce a new graphical representation, but to show that the profile and its scalar summaries possess interpretable properties such as boundedness, monotonicity, continuity, approximation, and plug-in consistency.

The move from point scores to profiles also changes the logic of dimensional reduction. A scalar efficiency score compresses the benchmark relation into one number from the outset. A profile preserves more of the geometry of benchmark heterogeneity and only later allows reduction through interpretable functionals. The move from point scores to profiles also changes the logic of dimensional reduction. A scalar efficiency score compresses the benchmark relation into one number from the outset, whereas a profile preserves more of its geometry before reduction. This logic is consistent with the functional-data perspective, where infinite-dimensional curves can be represented through a finite number of leading functional components for dimension reduction and interpretation [21]. In the present paper, those summaries are the average profile level, the tail fragility index, and the decline rate. Their role is not to reintroduce scalarization mechanically, but to summarize different aspects of curve geometry.

This perspective is especially natural in panel settings. Firm-level panels already track repeated observations over time, so replacing a one-number efficiency score with a firm-year profile does not introduce an artificial object; rather, it enriches the existing unit of analysis. It becomes possible to ask whether firms differ only in their average benchmark proximity or also in how sharply they deteriorate as the benchmark becomes more demanding. This, in turn, creates

a direct bridge from quantile-process estimation to economically interpretable heterogeneity, classification, and second-stage analysis.

2.4 *Economic relevance: firm heterogeneity, distance to frontier, and outcome analysis*

Although the present paper is primarily methodological, its motivation is also economic. Recent firm-level research continues to show that distance-to-frontier concepts are highly informative for understanding productivity, innovation, and performance. For example, firm-level evidence from Europe indicates that frontier proximity can be strongly linked to economically meaningful productivity and innovation outcomes [13]. The broader implication is that frontier-derived objects are not merely abstract benchmarking devices; they can also proxy a firm's position relative to better-performing peers and therefore have explanatory value for real economic heterogeneity.

Yet most of this literature still works with a single frontier distance or one efficiency score. This is precisely where the proposed profile approach adds value. Two firms may have similar average efficiency levels and yet differ substantially in how their relative position changes as the benchmark becomes more demanding. One may remain close to the frontier across the whole upper quantile range, whereas another may appear competitive only at moderate benchmark levels and then deteriorate sharply in the upper tail. Such distinctions are economically relevant in settings where technological depth, resilience, and competitive proximity to top-performing firms matter.

This richer decomposition also matters for second-stage modelling. In standard two-stage analysis, one typically regresses the estimated efficiency score on firm characteristics or uses it to explain firm outcomes. But if the underlying benchmark relation is heterogeneous across quantiles, then a single downstream regressor may blur distinct mechanisms. Some factors may primarily affect the overall level of benchmark proximity, while others may be associated more strongly with upper-tail fragility or with the rate of profile deterioration as benchmark strictness increases. Recent methodological discussion around second-stage approaches continues to emphasize the importance of carefully choosing what exactly is brought from the first stage into the second stage [14]. The present paper responds to this concern not by abandoning second-stage analysis, but by enriching the first-stage object before it is reduced to a small number of interpretable summaries.

The paper does not aim to replace all existing frontier-based empirical pipelines with a fully functional approach. Rather, it asks whether replacing a one-number efficiency regressor with profile-derived summaries can preserve mathematical tractability while yielding more informative empirical distinctions. This shift allows the analysis to separate overall frontier proximity from upper-tail weakness and profile steepness, and therefore to ask more refined questions about firm heterogeneity than those available under a purely scalar efficiency representation.

Against this background, the contribution of the present paper can be stated more precisely. It combines an ordered log-log quantile frontier process with a functional representation of firm-year efficiency and then derives theoretically interpretable curve summaries that preserve benchmark ordering while supporting continuity, approximation, and plug-in consistency results. This perspective also gives the proposed summaries a practical role. The average profile level can be used as a broad diagnostic of overall frontier proximity, whereas tail fragility identifies firms whose performance deteriorates mainly under the most demanding benchmarks. For managers, this distinction can separate firms that are generally inefficient from firms that are competitive under moderate standards but lack resilience relative to the upper tail. For policymakers, the same distinction can support more targeted benchmarking, since support measures for uniformly lagging firms may differ from measures aimed at firms that are close to the frontier but fragile under stricter comparison. Thus, the profile framework is not only a methodological extension of quantile frontier analysis. It also offers a more differentiated language for interpreting firm heterogeneity. Sections 3 and 4 develop this framework formally. Together, they recast firm-year efficiency as a frontier-induced functional object with its own coherent mathematical structure before the paper turns to empirical implementation and panel evidence.

3. Theoretical framework

We consider a panel of firms $i = 1, \ldots, N$ observed over periods $t = 1, \ldots, T$. Let $Y_{it} > 0$ denote output and let

$$X_{it} = (X_{1,it}, \ldots, X_{m,it}) \in \mathbb{R}^m_{++} \tag{2}$$

denote the vector of production inputs. Throughout, we work with a log-log conditional quantile frontier and treat firm-year efficiency as a quantile-indexed profile over the benchmark quantile level.

3.1 *Log-log quantile frontier and efficiency profile*

For a given quantile level $\tau \in [\tau_L, \tau_U] \subset (0, 1)$, define the conditional log-quantile frontier by

$$q_{it}(\tau) := q_\tau(\log Y_{it} \mid \log X_{it}), \tag{3}$$

where, under a log-log specification,

$$q_{it}(\tau) = \beta_0(\tau) + \sum_{j=1}^{m} \beta_j(\tau) \log X_{j,it}. \tag{4}$$

The corresponding frontier benchmark in levels is

$$Y_{it}^{\tau} = \exp(q_{it}(\tau)). \tag{5}$$

We define the firm-year quantile efficiency profile as

$$TE_{it}(\tau) = \min\left\{1, \frac{Y_{it}}{\exp(q_{it}(\tau))}\right\} = \min\{1, \exp(\log Y_{it} - q_{it}(\tau))\}, \ \tau \in [\tau_L, \tau_U]. \tag{6}$$

Thus, the profile records how the relative benchmark position of a firm-year observation changes as the benchmark moves from moderately high to increasingly demanding upper quantiles. For each firm-year observation, efficiency is no longer represented by a single number but by an entire profile over the quantile range $[\tau_L, \tau_U]$. Mathematically, this profile is represented by the function $TE_{it}(\tau)$; graphically, it may be visualized as a curve. Economically, $TE_{it}(\tau)$ measures how close firm i in period t is to increasingly demanding upper-quantile benchmarks.

It is also convenient to introduce the associated log-gap profile

$$G_{it}(\tau) = -\log TE_{it}(\tau) = \max\{0, q_{it}(\tau) - \log Y_{it}\}. \tag{7}$$

The gap formulation is useful in theoretical arguments because it linearizes the distance to the frontier in log space.

3.2 *Scalar summaries of the profile*

To obtain economically interpretable summary measures, we define three scalar functionals of the profile.

The first is the average profile level

$$L_{it} = \frac{1}{\tau_U - \tau_L} \int_{\tau_L}^{\tau_U} TE_{it}(\tau)\, d\tau. \tag{8}$$

This quantity captures the average proximity of the firm to the frontier across the whole benchmark range.

The second is the tail fragility index

$$TF_{it} = \frac{1}{\tau_M - \tau_L} \int_{\tau_L}^{\tau_M} TE_{it}(\tau)\, d\tau - \frac{1}{\tau_U - \tau_M} \int_{\tau_M}^{\tau_U} TE_{it}(\tau)\, d\tau, \tag{9}$$

where

$$\tau_M = \frac{\tau_L + \tau_U}{2}. \tag{10}$$

This index measures how much the firm's relative position deteriorates when the benchmark moves from moderately high quantiles to the most demanding upper-tail quantiles.

The third is the decline rate

$$D_{it} = -\frac{TE_{it}(\tau_U) - TE_{it}(\tau_L)}{\tau_U - \tau_L}. \tag{11}$$

This measure summarizes the overall steepness of the efficiency profile.

In what follows, the profile $TE_{it}(\tau)$ is the primary object, whereas L_{it}, TF_{it}, and D_{it} are low-dimensional summaries used for interpretation. In the empirical analysis, the main emphasis is placed on L_{it} and TF_{it}, while D_{it} plays a complementary shape-oriented role. The theoretical analysis below therefore serves two purposes: to establish basic properties of the profile itself and to show that the proposed summaries inherit coherent and stable behavior from the underlying ordered quantile frontier.

3.3 *Assumptions*

We impose the following regularity conditions.

Assumption 1 (Positivity) For every firm-year pair (i, t) and every $\tau \in [\tau_L, \tau_U]$,

$$q_{it}(\tau) \in \mathbb{R}, \exp(q_{it}(\tau)) > 0, Y_{it} > 0. \tag{12}$$

Assumption 2 (Non-crossing quantile frontier) For every (i, t), the function $q_{it}(\tau)$ is nondecreasing in τ, i.e.,

$$\tau_1 < \tau_2 \Rightarrow q_{it}(\tau_1) \leq q_{it}(\tau_2). \tag{13}$$

Assumption 2 imposes coherent benchmark ordering: higher quantiles correspond to weakly more demanding frontiers than lower quantiles.

Assumption 3 (Integrability) For every (i, t), the profile $TE_{it}(\tau)$ is measurable and integrable on $[\tau_L, \tau_U]$.

These assumptions are mild in the present empirical setting and have direct production-economic interpretations. Assumption 1 reflects the fact that the log-log frontier framework requires positive output and positive benchmark values; observations with non-positive output or inputs cannot be interpreted within a multiplicative production frontier. Assumption 2 is essential for the economic meaning of benchmark strictness: a higher quantile should represent a weakly more demanding benchmark, not an easier one. Without this ordering, a decline or stability of the efficiency profile across τ would have no coherent interpretation. Assumption 3 is mainly technical but also natural empirically, because firm-year profiles are evaluated over a compact quantile interval and are constructed from a finite grid of estimated benchmark values in the empirical implementation.

3.4 *Constrained quantile-process estimation on a grid*

Although the profile is defined on the continuum $[\tau_L, \tau_U]$, the empirical implementation proceeds on a finite quantile grid

$$\mathscr{T} = \{\tau_1, \ldots, \tau_K\} \subset [\tau_L, \tau_U]. \tag{14}$$

This grid-based formulation is important both computationally and conceptually, because it links the continuous profile object to the finite set of fitted benchmark values used in estimation.

For each $\tau \in \mathscr{T}$, the log-log frontier is estimated subject to economically coherent shape restrictions and non-crossing constraints across quantiles. In particular, the empirical construction assumes that the fitted process satisfies the same ordering logic as the theoretical frontier family, namely that higher quantiles correspond to weakly higher benchmark frontiers. Denoting the fitted grid-based frontier by $\widehat{q}_{it}(\tau)$, the empirical profile is obtained by substitution into the profile mapping introduced above.

This section uses the grid only as an implementation device. The mathematical results in Section 4 are stated for the profile as a function on $[\tau_L, \tau_U]$, and the grid-based approximation results later show how the discrete implementation recovers the continuous objects of interest. In this way, the empirical implementation approximates the continuous profile object by a coherent finite-dimensional representation while preserving the benchmark-ordering logic required by the theory. Thus, Section 3.4 establishes the theoretical connection between the continuous profile and its grid-based implementation. The operational details of the constrained estimation algorithm, the industry-year construction of frontiers, and the empirical grid choices are described in Section 5.

4. Mathematical properties of quantile efficiency profiles

4.1 *Basic properties of the efficiency profile*

We begin with several basic but essential properties of the newly defined profile, including boundedness, monotonicity, continuity, and invariance under consistent changes of output units.

Proposition 1 (Well-definedness and boundedness)

Under Assumptions 1 and 3, the profile $TE_{it}(\tau)$ is well defined for all (i, t, τ), and

$$0 < TE_{it}(\tau) \leq 1. \tag{15}$$

Consequently,

$$0 < L_{it} \leq 1. \tag{16}$$

Proof. Since $Y_{it} > 0$ and $\exp(q_{it}(\tau)) > 0$, the ratio

$$\frac{Y_{it}}{\exp(q_{it}(\tau))} \tag{17}$$

is strictly positive. Applying $\min\{1, \cdot\}$ yields $0 < TE_{it}(\tau) \leq 1$. Integrating over $[\tau_L, \tau_U]$ implies $0 < L_{it} \leq 1$. □

Economically, Proposition 1 states that the profile always produces an interpretable efficiency value: a firm-year observation cannot have negative efficiency, and values above the estimated benchmark are capped at one.

Proposition 2 (Monotonicity of the profile in the quantile level)

Under Assumptions 1-2, the profile $TE_{it}(\tau)$ is nonincreasing in τ. That is,

$$\tau_1 < \tau_2 \Rightarrow TE_{it}(\tau_1) \geq TE_{it}(\tau_2). \tag{18}$$

Proof. By Assumption 2,

$$q_{it}(\tau_1) \leq q_{it}(\tau_2). \tag{19}$$

Exponentiation preserves order, hence

$$\exp(q_{it}(\tau_1)) \leq \exp(q_{it}(\tau_2)). \tag{20}$$

Therefore

$$\frac{Y_{it}}{\exp(q_{it}(\tau_1))} \geq \frac{Y_{it}}{\exp(q_{it}(\tau_2))}. \tag{21}$$

Since the mapping $x \mapsto \min\{1, x\}$ is nondecreasing, the claim follows. □

Economically, Proposition 2 means that a firm should not appear more efficient simply because it is compared with a stricter upper-tail benchmark; as τ increases, the benchmark becomes weakly more demanding and the efficiency profile can only remain unchanged or decline.

Corollary 1

Under Assumptions 1-3,

$$TF_{it} \geq 0, D_{it} \geq 0. \tag{22}$$

Proof. By Proposition 2, $TE_{it}(\tau)$ is nonincreasing on $[\tau_L, \tau_U]$. Hence the average value of the profile on $[\tau_L, \tau_M]$ is not smaller than its average value on $[\tau_M, \tau_U]$, which implies $TF_{it} \geq 0$. Also,

$$TE_{it}(\tau_L) \geq TE_{it}(\tau_U), \tag{23}$$

which gives $D_{it} \geq 0$. □

Proposition 3 (Endpoint bounds for the profile summaries)

Under Assumptions 1-3, for every firm-year pair (i, t),

$$TE_{it}(\tau_U) \leq L_{it} \leq TE_{it}(\tau_L). \tag{24}$$

Moreover,

$$0 \leq TF_{it} \leq TE_{it}(\tau_L) - TE_{it}(\tau_U). \tag{25}$$

Proof. Since $TE_{it}(\tau)$ is nonincreasing on $[\tau_L, \tau_U]$ by Proposition 2, its average over the interval must lie between its endpoint values. Therefore,

$$TE_{it}(\tau_U) \leq \frac{1}{\tau_U - \tau_L} \int_{\tau_L}^{\tau_U} TE_{it}(\tau)\, d\tau \leq TE_{it}(\tau_L), \tag{26}$$

which proves the first claim.

For the second claim, let T_L and T_U denote the average values of $TE_{it}(\tau)$ on the left and right half-intervals, respectively. Then

$$TF_{it} = T_L - T_U. \tag{27}$$

By monotonicity,

$$TE_{it}(\tau_U) \leq T_U \leq T_L \leq TE_{it}(\tau_L), \tag{28}$$

hence

$$0 \leq TF_{it} = T_L - T_U \leq TE_{it}(\tau_L) - TE_{it}(\tau_U). \tag{29}$$

□

Proposition 4 (Continuity of the efficiency profile)

Suppose that for a given firm-year pair (i, t), the log-quantile frontier $q_{it}(\tau)$ is continuous on $[\tau_L, \tau_U]$. Then the efficiency profile

$$TE_{it}(\tau) = \min\{1, \exp(\log Y_{it} - q_{it}(\tau))\} \tag{30}$$

is continuous on $[\tau_L, \tau_U]$. Consequently, the log-gap profile

$$G_{it}(\tau) = \max\{0,\, q_{it}(\tau) - \log Y_{it}\} \tag{31}$$

is also continuous on $[\tau_L,\, \tau_U]$.

Proof. Since $q_{it}(\tau)$ is continuous on $[\tau_L,\, \tau_U]$, the mapping

$$\tau \mapsto \log Y_{it} - q_{it}(\tau) \tag{32}$$

is continuous, and therefore so is

$$\tau \mapsto \exp(\log Y_{it} - q_{it}(\tau)). \tag{33}$$

The function $x \mapsto \min\{1, x\}$ is continuous on $(0, \infty)$. Hence the composition

$$\tau \mapsto \min\{1, \exp(\log Y_{it} - q_{it}(\tau))\} \tag{34}$$

is continuous on $[\tau_L,\, \tau_U]$, proving continuity of $TE_{it}(\tau)$. Likewise, the function $x \mapsto \max\{0, x\}$ is continuous on $\mathbb{R}$, so

$$\tau \mapsto \max\{0,\, q_{it}(\tau) - \log Y_{it}\} \tag{35}$$

is continuous on $[\tau_L,\, \tau_U]$. □

Economically, continuity means that small changes in benchmark strictness do not generate artificial jumps in measured efficiency when the underlying quantile frontier varies smoothly.

Proposition 5 (Unit invariance of the profile and its summaries)

Fix a firm-year pair (i, t). Suppose output is rescaled by a positive constant $c > 0$, so that

$$\widetilde{Y}_{it} = cY_{it}, \tag{36}$$

and the log-quantile frontier is shifted consistently,

$$\widetilde{q}_{it}(\tau) = q_{it}(\tau) + \log c. \tag{37}$$

Then, for all $\tau \in [\tau_L,\, \tau_U]$,

$$\widetilde{TE}_{it}(\tau) = TE_{it}(\tau),\, \widetilde{G}_{it}(\tau) = G_{it}(\tau). \tag{38}$$

Consequently,

$$\widetilde{L}_{it} = L_{it},\ \widetilde{TF}_{it} = TF_{it},\ \widetilde{D}_{it} = D_{it}. \tag{39}$$

Proof. By definition,

$$\widetilde{TE}_{it}(\tau) = \min\{1, \exp(\log \widetilde{Y}_{it} - \widetilde{q}_{it}(\tau))\}. \tag{40}$$

Substituting $\widetilde{Y}_{it} = cY_{it}$ and $\widetilde{q}_{it}(\tau) = q_{it}(\tau) + \log c$, we obtain

$$\widetilde{TE}_{it}(\tau) = \min\{1, \exp(\log Y_{it} + \log c - q_{it}(\tau) - \log c)\} = \min\{1, \exp(\log Y_{it} - q_{it}(\tau))\} = TE_{it}(\tau). \tag{41}$$

Similarly,

$$\widetilde{G}_{it}(\tau) = \max\{0, \widetilde{q}_{it}(\tau) - \log \widetilde{Y}_{it}\} = \max\{0, q_{it}(\tau) - \log Y_{it}\} = G_{it}(\tau). \tag{42}$$

Since L_{it}, TF_{it}, and D_{it} are functionals of $TE_{it}(\tau)$, they are unchanged as well. □

This invariance property is important for empirical work because it ensures that the profile and its summaries do not depend on the units in which output is measured, provided the frontier is rescaled consistently.

4.2 *Ordering and stability of efficiency profiles*

The profile formulation naturally induces a partial ordering over firms observed under the same benchmark environment.

Definition 1 (Profile dominance)

For two firms a and b observed in the same period t, we say that firm a profile-dominates firm b, written

$$TE_{at} \succeq TE_{bt}, \tag{43}$$

if

$$TE_{at}(\tau) \geq TE_{bt}(\tau) \text{for all } \tau \in [\tau_L,\ \tau_U]. \tag{44}$$

Proposition 6 (Order preservation under profile dominance)

Suppose firm a profile-dominates firm b in period t. Then

$$L_{at} \geq L_{bt}. \tag{45}$$

If, in addition,

$$TE_{at}(\tau_L) = TE_{bt}(\tau_L), \tag{46}$$

then

$$D_{at} \leq D_{bt}. \tag{47}$$

Proof. The first statement follows directly by integrating the pointwise inequality

$$TE_{at}(\tau) \geq TE_{bt}(\tau) \tag{48}$$

over $[\tau_L,\ \tau_U]$. For the second statement,

$$D_{at} = -\frac{TE_{at}(\tau_U) - TE_{at}(\tau_L)}{\tau_U - \tau_L},\ D_{bt} = -\frac{TE_{bt}(\tau_U) - TE_{bt}(\tau_L)}{\tau_U - \tau_L}. \tag{49}$$

If the left endpoints coincide and $TE_{at}(\tau_U) \geq TE_{bt}(\tau_U)$, then $D_{at} \leq D_{bt}$. □

Proposition 6 formalizes an economically intuitive ranking principle: if one firm lies weakly above another firm at all benchmark levels, then its integrated efficiency must also be at least as high. If both profiles start from the same left-end benchmark, the flatter profile is less fragile.

We now establish that the efficiency profile and its scalar summaries are continuous functionals of the underlying log-quantile process. This property is important for both robustness and statistical approximation arguments.

Let $q_{it}(\tau)$ and $\widetilde{q}_{it}(\tau)$ be two candidate log-quantile frontier functions for the same firm-year pair (i, t). Define the corresponding profiles

$$TE_{it}(\tau) = \min\{1, \exp(\log Y_{it} - q_{it}(\tau))\}, \tag{50}$$

$$\widetilde{TE}_{it}(\tau) = \min\{1, \exp(\log Y_{it} - \widetilde{q}_{it}(\tau))\}. \tag{51}$$

Let

$$\| q_{it} - \widetilde{q}_{it} \|_{\infty} = \sup_{\tau \in [\tau_L,\ \tau_U]} | q_{it}(\tau) - \widetilde{q}_{it}(\tau) | . \tag{52}$$

Economically, profile dominance formalizes the idea that one firm is uniformly closer to the benchmark than another across all levels of benchmark strictness, rather than only at one selected quantile.

Proposition 7 (Uniform continuity of the efficiency profile)

For every firm-year pair (i, t),

$$\sup_{\tau\in[\tau_L,\,\tau_U]} |\, TE_{it}(\tau) - \widetilde{TE}_{it}(\tau)\,| \le \| q_{it} - \widetilde{q}_{it} \|_\infty . \tag{53}$$

Proof. Consider

$$h_y(z) = \min\{1,\, e^{y-z}\},\ y = \log Y_{it}. \tag{54}$$

For $z < y$, $h_y(z) = 1$; for $z \ge y$, $h_y'(z) = -e^{y-z}$, hence $|\, h_y'(z)\,| \le 1$. Thus h_y is 1-Lipschitz:

$$|\, h_y(z_1) - h_y(z_2)\,| \le |\, z_1 - z_2\,| . \tag{55}$$

Substituting $z_1 = q_{it}(\tau)$ and $z_2 = \widetilde{q}_{it}(\tau)$ gives

$$|\, TE_{it}(\tau) - \widetilde{TE}_{it}(\tau)\,| \le |\, q_{it}(\tau) - \widetilde{q}_{it}(\tau)\,| . \tag{56}$$

Taking the supremum over τ proves the result. □

Corollary 2 (Stability of the scalar summaries)
Under the notation of Proposition 7,

$$|\, L_{it} - \widetilde{L}_{it}\,| \le \| q_{it} - \widetilde{q}_{it} \|_\infty, \tag{57}$$

$$|\, TF_{it} - \widetilde{TF}_{it}\,| \le 2\,\| q_{it} - \widetilde{q}_{it} \|_\infty, \tag{58}$$

and

$$|\, D_{it} - \widetilde{D}_{it}\,| \le \frac{2}{\tau_U - \tau_L}\,\| q_{it} - \widetilde{q}_{it} \|_\infty . \tag{59}$$

Proof. The bound for L_{it} follows from Proposition 7 and

$$|\, L_{it} - \widetilde{L}_{it}\,| \le \frac{1}{\tau_U - \tau_L}\int_{\tau_L}^{\tau_U} |\, TE_{it}(\tau) - \widetilde{TE}_{it}(\tau)\,|\ d\tau. \tag{60}$$

The bound for TF_{it} follows by applying the same argument separately on the two half-intervals and using the triangle inequality. The bound for D_{it} follows from

$$|\, D_{it} - \widetilde{D}_{it}\,| \le \frac{|\, TE_{it}(\tau_U) - \widetilde{TE}_{it}(\tau_U)\,| + |\, TE_{it}(\tau_L) - \widetilde{TE}_{it}(\tau_L)\,|}{\tau_U - \tau_L}. \tag{61}$$

Applying Proposition 7 completes the proof. □

These stability bounds are important for the empirical analysis because they show that moderate perturbations of the estimated quantile frontier induce only controlled changes in the profile summaries.

4.3 *Consistency and approximation of the profile indices*

We next show that if the estimated log-quantile frontier converges uniformly to the true frontier on $[\tau_L, \tau_U]$, then the induced efficiency profile and all scalar summaries are consistent. We also relate the continuous profile summaries to their empirical grid-based approximations.

Let $q_{it}^0(\tau)$ denote the true log-quantile frontier and let $\widehat{q}_{it}(\tau)$ be its estimator. Define

$$TE_{it}^0(\tau) = \min\{1, \exp(\log Y_{it} - q_{it}^0(\tau))\}, \tag{62}$$

$$\widehat{TE}_{it}(\tau) = \min\{1, \exp(\log Y_{it} - \widehat{q}_{it}(\tau))\}. \tag{63}$$

Similarly, let $L_{it}^0, TF_{it}^0, D_{it}^0$ denote the population functionals, and $\widehat{L}_{it}, \widehat{TF}_{it}, \widehat{D}_{it}$ their plug-in estimators.

Proposition 8 (Plug-in consistency)

If

$$\sup_{\tau\in[\tau_L,\,\tau_U]} |\,\widehat{q}_{it}(\tau) - q_{it}^0(\tau)\,| \to 0, \tag{64}$$

then

$$\sup_{\tau\in[\tau_L,\,\tau_U]} |\,\widehat{TE}_{it}(\tau) - TE_{it}^0(\tau)\,| \to 0, \tag{65}$$

and

$$\widehat{L}_{it} \to L_{it}^0, \widehat{TF}_{it} \to TF_{it}^0, \widehat{D}_{it} \to D_{it}^0. \tag{66}$$

Proof. By Proposition 7,

$$\sup_{\tau\in[\tau_L,\,\tau_U]} |\,\widehat{TE}_{it}(\tau) - TE_{it}^0(\tau)\,| \leq \sup_{\tau\in[\tau_L,\,\tau_U]} |\,\widehat{q}_{it}(\tau) - q_{it}^0(\tau)\,|\,. \tag{67}$$

Hence the first convergence follows immediately. The consistency of $\widehat{L}_{it}, \widehat{TF}_{it}$, and $\widehat{D}_{it}$ then follows from Corollary 2 by substituting $\widetilde{q}_{it} = q_{it}^0$ and $q_{it} = \widehat{q}_{it}$. □

Economically, plug-in consistency means that if the estimated quantile frontier converges to the true ordered benchmark process, then the estimated profile summaries converge to their population counterparts. This justifies using the estimated summaries as empirical objects in second-stage analysis, because their sampling behavior is inherited from the consistency of the underlying ordered frontier estimator.

Proposition 9 (Grid approximation of the profile summaries)

Let

$$\tau_k = \tau_L + k\Delta,\ k = 0,\ 1,\ \ldots,\ K, \tag{68}$$

be an equally spaced grid on $[\tau_L,\ \tau_U]$, where $\Delta = (\tau_U - \tau_L)/K$, and assume that K is even so that $\tau_{K/2} = \tau_M$. Suppose that $TE_{it}(\tau)$ is continuous on $[\tau_L,\ \tau_U]$. Define

$$L_{it}^{(K)} = \frac{1}{\tau_U - \tau_L} \sum_{k=0}^{K-1} TE_{it}(\tau_k)\Delta, \tag{69}$$

$$TF_{it}^{(K)} = \frac{1}{\tau_M - \tau_L} \sum_{k=0}^{K/2-1} TE_{it}(\tau_k)\Delta - \frac{1}{\tau_U - \tau_M} \sum_{k=K/2}^{K-1} TE_{it}(\tau_k)\Delta, \tag{70}$$

and

$$D_{it}^{(K)} = -\frac{TE_{it}(\tau_K) - TE_{it}(\tau_0)}{\tau_U - \tau_L}. \tag{71}$$

Then

$$L_{it}^{(K)} \to L_{it},\ TF_{it}^{(K)} \to TF_{it},\ D_{it}^{(K)} = D_{it} \tag{72}$$

as $K \to \infty$.

Proof. The claims for $L_{it}^{(K)}$ and $TF_{it}^{(K)}$ follow directly from continuity of $TE_{it}(\tau)$ and standard convergence of Riemann sums to the corresponding integrals. The equality $D_{it}^{(K)} = D_{it}$ holds because D_{it} depends only on the endpoint values $TE_{it}(\tau_L)$ and $TE_{it}(\tau_U)$, which are included in the grid by construction. □

Proposition 9 provides the bridge between the continuous profile formulation and the empirical implementation on a finite quantile grid. In particular, it justifies the use of grid-based approximations for L_{it} and TF_{it}, while D_{it} is exactly recovered whenever the grid contains the endpoints.

Proposition 10 (Explicit discretization error bounds for grid-based profile summaries)

Let

$$\tau_k = \tau_L + k\Delta,\ k = 0,\ 1,\ \ldots,\ K, \tag{73}$$

be an equally spaced grid on $[\tau_L,\ \tau_U]$, where $\Delta = (\tau_U - \tau_L)/K$, and assume that K is even. Suppose that $TE_{it}(\tau)$ is Lipschitz continuous on $[\tau_L,\ \tau_U]$ with Lipschitz constant M_{it}, that is,

$$|\ TE_{it}(\tau) - TE_{it}(\tau')\ | \leq M_{it}\ |\ \tau - \tau'\ | \text{ for all } \tau,\ \tau' \in [\tau_L,\ \tau_U]. \tag{74}$$

Then the left-endpoint grid approximations satisfy

$$| L_{it}^{(K)} - L_{it} | \leq M_{it}\Delta, \tag{75}$$

and

$$| TF_{it}^{(K)} - TF_{it} | \leq 2M_{it}\Delta. \tag{76}$$

Proof. For each grid cell $[\tau_k, \tau_{k+1}]$, the Lipschitz condition implies

$$| TE_{it}(\tau) - TE_{it}(\tau_k) | \leq M_{it}\Delta \text{ for all } \tau \in [\tau_k, \tau_{k+1}]. \tag{77}$$

Integrating over each cell and summing gives

$$| \int_{\tau_L}^{\tau_U} TE_{it}(\tau)\, d\tau - \sum_{k=0}^{K-1} TE_{it}(\tau_k)\Delta | \leq (\tau_U - \tau_L) M_{it}\Delta. \tag{78}$$

After division by $\tau_U - \tau_L$, we obtain

$$| L_{it}^{(K)} - L_{it} | \leq M_{it}\Delta. \tag{79}$$

For TF_{it}, apply the same argument separately on the two half-intervals $[\tau_L, \tau_M]$ and $[\tau_M, \tau_U]$. Each half contributes at most $M_{it}\Delta$ after normalization by its own interval length, so by the triangle inequality

$$| TF_{it}^{(K)} - TF_{it} | \leq 2M_{it}\Delta. \tag{80}$$

□

This result is especially relevant for the empirical implementation because it makes explicit how the finite quantile grid approximates the continuous profile summaries under regularity.

4.4 *Shape and classification of efficiency profiles*

The next result is useful for interpreting the relationship between the two shape-oriented indices TF_{it} and D_{it}.

Proposition 11 (Linear profile benchmark)

Suppose the efficiency profile is affine in τ:

$$TE_{it}(\tau) = a_{it} - b_{it}(\tau - \tau_L),\ b_{it} \geq 0. \tag{81}$$

Then

$$D_{it} = b_{it}, \tag{82}$$

and

$$TF_{it} = \frac{\tau_U - \tau_L}{2} D_{it}. \tag{83}$$

Proof. Since

$$TE_{it}(\tau_U) - TE_{it}(\tau_L) = -b_{it}(\tau_U - \tau_L), \tag{84}$$

we obtain

$$D_{it} = -\frac{TE_{it}(\tau_U) - TE_{it}(\tau_L)}{\tau_U - \tau_L} = b_{it}. \tag{85}$$

For a linear function, the average value on an interval equals the midpoint value. Thus the left-half average equals $TE_{it}((\tau_L + \tau_M)/2)$, and the right-half average equals $TE_{it}((\tau_M + \tau_U)/2)$. The difference between these two midpoint values is

$$b_{it} \cdot \frac{\tau_U - \tau_L}{2}. \tag{86}$$

Substituting $b_{it} = D_{it}$ yields the stated expression for TF_{it}. □

Proposition 11 shows that if the profile is approximately linear, then the fragility and decline-rate indices contain essentially the same information up to a scale factor. This helps explain why, in the empirical analysis, D_{it} is treated mainly as a complementary shape-oriented summary rather than as a coequal benchmark measure.

Proposition 12 (Curvature interpretation of tail fragility relative to the decline rate)

Suppose that, for a given firm-year pair (i, t), the profile $TE_{it}(\tau)$ is twice continuously differentiable on $[\tau_L, \tau_U]$. Then:

1. If TE_{it} is concave on $[\tau_L, \tau_U]$,

$$TF_{it} \leq \frac{\tau_U - \tau_L}{2} D_{it}; \tag{87}$$

2. If TE_{it} is convex on $[\tau_L, \tau_U]$,

$$TF_{it} \geq \frac{\tau_U - \tau_L}{2} D_{it}. \tag{88}$$

Equality holds in either case if $TE_{it}(\tau)$ is affine in τ.

Proof. Let

$$l(\tau) = TE_{it}(\tau_L) + \frac{TE_{it}(\tau_U) - TE_{it}(\tau_L)}{\tau_U - \tau_L}(\tau - \tau_L) \tag{89}$$

be the chord joining the two endpoint values of the profile.

If TE_{it} is concave, then

$$TE_{it}(\tau) \geq l(\tau) \text{ for all } \tau \in [\tau_L, \tau_U]. \tag{90}$$

Averaging over the left and right half-intervals preserves this inequality, so the difference between left-half and right-half averages of TE_{it} does not exceed the corresponding difference for the affine function $l(\tau)$. By Proposition 11 applied to $l(\tau)$,

$$TF(l) = \frac{\tau_U - \tau_L}{2} D(l). \tag{91}$$

Since $D(l) = D_{it}$, we obtain

$$TF_{it} \leq \frac{\tau_U - \tau_L}{2} D_{it}. \tag{92}$$

If TE_{it} is convex, then

$$TE_{it}(\tau) \leq l(\tau), \tag{93}$$

and the inequality reverses:

$$TF_{it} \geq \frac{\tau_U - \tau_L}{2} D_{it}. \tag{94}$$

If TE_{it} is affine, then it coincides with its chord, and equality follows from Proposition 11. □

Finally, since the empirical analysis classifies firm-year observations by profile type using (L_{it}, TF_{it}), we record a simple local stability property of such classifications.

Lemma 1 (Local stability of profile-based classification).

Suppose firm-year observations are classified by thresholding the pair (L_{it}, TF_{it}) at fixed reference values (c_L, c_{TF}). Let the estimated summaries be $\left(\widehat{L}_{it}, \widehat{TF}_{it}\right)$. If, for some $\delta > 0$,

$$| L_{it} - c_L | > \delta \text{ and } | TF_{it} - c_{TF} | > \delta, \tag{95}$$

and if

$$|\widehat{L}_{it} - L_{it}| < \delta, \ |\widehat{TF}_{it} - TF_{it}| < \delta, \tag{96}$$

then the class assigned using $\left(\widehat{L}_{it}, \widehat{TF}_{it}\right)$ coincides with the class assigned using (L_{it}, TF_{it}).

Proof. The assumptions imply that neither true summary lies within a δ-neighborhood of its classification threshold. Therefore, perturbations smaller than δ cannot change the sign of

$$L_{it} - c_L \text{ or } TF_{it} - c_{TF}. \tag{97}$$

Hence

$$\operatorname{sign}(\widehat{L}_{it} - c_L) = \operatorname{sign}(L_{it} - c_L), \ \operatorname{sign}(\widehat{TF}_{it} - c_{TF}) = \operatorname{sign}(TF_{it} - c_{TF}). \tag{98}$$

Since the profile class is determined only by these two signs, the estimated and true classifications coincide. □

The above results establish a set of properties needed for the subsequent empirical analysis. First, the efficiency profile is a coherent bounded object defined on the entire quantile range. Second, it is monotone in the benchmark strictness parameter τ, continuous under mild regularity conditions on the underlying log-quantile frontier, and invariant to consistent changes in output units. Third, the proposed scalar summaries preserve natural ranking information under profile dominance and admit explicit bounds in terms of endpoint efficiencies. Fourth, the profile and its summary indices are stable and plug-in consistent under uniform convergence of the estimated log-quantile frontier. Fifth, the continuous summaries L_{it} and TF_{it} admit natural grid-based approximations, together with explicit finite-grid error bounds under Lipschitz regularity, while D_{it} is exactly determined by endpoint values whenever the quantile grid contains τ_L and τ_U. Sixth, the relation between TF_{it} and D_{it} admits a useful geometric interpretation through the curvature of the profile, and simple threshold-based profile classifications are locally stable away from classification boundaries.

These properties justify the use of $TE_{it}(\tau)$ as the central firm-year object and of L_{it}, TF_{it}, and D_{it} as interpretable low-dimensional summaries. In the sections that follow, the empirical analysis focuses primarily on L_{it} and TF_{it}, while D_{it} is used mainly as a complementary summary of profile steepness.

5. Empirical design and econometric implementation

This section describes the empirical implementation of the proposed functional quantile-efficiency framework. We first outline the data source, sample structure, variable construction, and firm identifiers. We then explain how the log-log quantile frontier is estimated on a finite quantile grid, how firm-year efficiency profiles and their scalar summaries are computed, how profile types are defined, and how these objects enter second-stage panel regressions and validation checks.

5.1 *Data source, sample, and variables*

The empirical analysis is based on firm-level accounting data from the System for Professional Analysis of Markets and Companies (SPARK-Interfax), an information and analytical system that provides company registration information and financial reports for Russian enterprises [22]. The database offers broad coverage and consistent reporting formats across years, making it suitable for panel-based production and performance analysis. Industry affiliation is identified using the Russian Classification of Economic Activities (OKVED 2), which is largely harmonized with Statistical

Classification of Economic Activities in the European Communities (NACE Rev. 2) and therefore allows economically meaningful sector grouping.

The final estimation sample covers 2,464 Russian firms observed annually over 2019-2023, yielding 12,320 firm-year observations. The observation window is especially informative because it spans a period of elevated volatility, repeated macroeconomic disturbances, and heterogeneous firm-level adjustment. The final sample is restricted to firms with non-missing accounting information required to construct output, inputs, profile summaries, second-stage determinants, and robustness variables. Descriptive statistics for the main raw empirical variables in the final estimation sample are reported in Appendix A.

For calculating technical efficiency, output is proxied by sales revenue, while inputs are represented by fixed assets and the wage bill, which serve as parsimonious proxies for capital and labor, respectively. This choice follows the same production-analytic logic used in earlier log-log frontier applications and is especially convenient in a panel setting with heterogeneous firm sizes. It is also consistent with recent stochastic-frontier work emphasizing parsimonious production specifications under heterogeneous firm environments and panel structures [23, 24]. The strong dispersion of firm scale—many smaller firms and a smaller number of very large enterprises—also supports the use of logarithmic specifications.

For the baseline second-stage analysis, we focus on standard time-varying firm-level characteristics used as determinants of the estimated efficiency summaries. The main regressors are leverage and size. Leverage is measured as total debt divided by total assets, while size is measured as the logarithm of total assets. Additional firm-performance variables, including Return on Assets (ROA) and sales growth, are used in exploratory lagged validation regressions rather than in the baseline determinant regressions.

5.2 *Quantile grid, industry-year frontiers, and estimation algorithm*

Unlike the single-quantile approach, the present paper estimates an entire quantile-indexed frontier process on a finite grid,

$$\mathcal{T}_M = \{\tau_1, \ldots, \tau_M\}. \tag{99}$$

In the baseline specification, we use the grid

$$\mathcal{T}_M = \{0.70, 0.75, 0.80, 0.85, 0.90, 0.95\}. \tag{100}$$

This grid is wide enough to recover heterogeneity across benchmark strictness levels while avoiding the instability of very extreme upper quantiles in short firm panels. The baseline grid covers a broad upper-quantile range while stopping short of more extreme tail levels. This reduces exposure to the instability associated with sparse tail observations, a recognized difficulty in extreme-quantile estimation [25]. The use of upper quantiles as probabilistic frontier benchmarks is consistent with the broader quantile-frontier literature [4, 5].

As an additional robustness step, the empirical analysis also considers alternative subranges of the baseline quantile grid. In particular, the profile summaries are recomputed on the intervals [0.75, 0.95], [0.70, 0.90], and [0.75, 0.85] using the already estimated grid-specific efficiency values. This allows us to examine whether the main empirical patterns depend strongly on the exact choice of the upper-quantile range or remain qualitatively stable under plausible modifications of benchmark coverage.

The frontier estimation is industry-year specific. That is, for each industry s, year t, and quantile level $\tau \in \mathcal{T}_M$, the log-log frontier is estimated using observations belonging to the corresponding industry-year cell. This design avoids artificial technological pooling across heterogeneous sectors and allows the benchmark to reflect industry-specific production environments in each period. It should therefore be interpreted as a set of industry-year-specific frontiers rather than as a pooled frontier with industry-year dummy variables.

Let

$$y_{it} = \log Y_{it},\ k_{it} = \log K_{it},\ l_{it} = \log LB_{it}, \tag{101}$$

where Y_{it} denotes sales revenue, K_{it} mdenotes fixed assets, and LB_{it} denotes the wage bill. For each industry-year cell (s, t), the conditional log-quantile frontier is specified as

$$q_{\tau,s,t}(x_{it}) = \alpha_{\tau,s,t} + \beta_{K,\tau,s,t} k_{it} + \beta_{L,\tau,s,t} l_{it}. \tag{102}$$

The grid-specific quantile coefficients are estimated by minimizing the sum of quantile check losses within each industry-year cell:

$$\min_{\alpha_{\tau,s,t},\, \beta_{K,\tau,s,t},\, \beta_{L,\tau,s,t}} \sum_{\tau \in \mathscr{T}_M} \sum_{i \in \mathscr{I}_{s,t}} \rho_\tau \left(y_{it} - \alpha_{\tau,s,t} - \beta_{K,\tau,s,t} k_{it} - \beta_{L,\tau,s,t} l_{it} \right). \tag{103}$$

Here $\mathscr{I}_{s,t}$ is the set of firms observed in industry s and year t, and $\rho_\tau(u)$ is the standard quantile loss function,

$$\rho_\tau(u) = u\{\tau - \mathbf{1}(u < 0)\}. \tag{104}$$

The estimation follows the constrained quantile framework introduced in Sections 3 and 4. Specifically, the fitted quantile process is subject to two sets of restrictions. First, economic shape restrictions are imposed on the log-log technology, including monotonicity of output with respect to production inputs and economically admissible bounds on returns to scale. Second, non-crossing restrictions are imposed across the quantile grid:

$$\widehat{q}_{\tau_m,s,t}(x_{it}) \le \widehat{q}_{\tau_{m+1},s,t}(x_{it}),\ m = 1, \ldots, M-1,\ i \in \mathscr{I}_{s,t}. \tag{105}$$

These restrictions ensure that higher quantiles correspond to weakly higher benchmark frontiers. They are important because they combine production-theoretic coherence with ordered quantile benchmarking, in the spirit of recent work on shape-constrained and non-crossing quantile methods [9, 19]. The resulting fitted frontier process is therefore coherent both statistically and economically. It yields the grid-specific fitted frontier values

$$\widehat{q}_{\tau,s,t}(x_{it}),\ \tau \in \mathscr{T}_M, \tag{106}$$

from which all subsequent profile objects are computed.

For reproducibility, the empirical algorithm can be summarized as follows. First, the data are cleaned and restricted to observations with positive output, positive inputs, and non-missing second-stage variables. Second, a globally unique firm identifier is constructed by combining the industry code and the original firm identifier. Third, the sample is split into industry-year cells. Fourth, within each industry-year cell, log-log quantile frontiers are estimated jointly over the grid $\mathscr{T}_M$ under shape and non-crossing constraints. Fifth, fitted log-frontier values are exponentiated to obtain benchmark output levels. Sixth, firm-year efficiency values are computed at each quantile. Seventh, the quantile-specific efficiency

values are assembled into firm-year efficiency profiles. Eighth, the profile summaries and profile-type classifications are computed and then used in descriptive analysis, panel regressions, and robustness checks.

5.3 *Construction of firm-year efficiency profiles*

For each firm-year pair and each $\tau \in \mathscr{T}_M$, the estimated frontier benchmark in levels is

$$\widehat{Y}_{it}^{F}(\tau) = \exp\{\widehat{q}_{\tau,\, s(i),\, t}(x_{it})\}, \tag{107}$$

where $s(i)$ denotes the industry of firm i. The firm-year quantile efficiency profile is then defined as

$$\widehat{TE}_{it}(\tau) = \min\left\{1,\, \frac{Y_{it}}{\widehat{Y}_{it}^{F}(\tau)}\right\},\ \tau \in \mathscr{T}_M. \tag{108}$$

Thus, each firm-year observation is represented by an entire profile over the quantile range rather than by a single technical-efficiency score. Mathematically, this profile is represented by the function $\widehat{TE}_{it}(\tau)$; graphically, it can be visualized as a curve over benchmark strictness. This construction allows the empirical analysis to distinguish between firms that are uniformly close to upper-tail benchmarks and firms whose relative position deteriorates sharply as the benchmark becomes more demanding.

The three scalar profile summaries introduced in Section 3 are computed numerically on the grid $\mathscr{T}_M$. The first is the average profile level,

$$\widehat{L}_{it} = \frac{1}{M}\sum_{m=1}^{M} \widehat{TE}_{it}(\tau_m). \tag{109}$$

The second is the tail fragility index,

$$\widehat{TF}_{it} = \frac{1}{|\,\mathscr{T}^{L}\,|}\sum_{\tau_m \in \mathscr{T}^{L}} \widehat{TE}_{it}(\tau_m) - \frac{1}{|\,\mathscr{T}^{U}\,|}\sum_{\tau_m \in \mathscr{T}^{U}} \widehat{TE}_{it}(\tau_m), \tag{110}$$

where $\mathscr{T}^{L}$ and $\mathscr{T}^{U}$ denote the lower and upper parts of the selected quantile grid, respectively. The third is the decline rate,

$$\widehat{D}_{it} = \frac{\widehat{TE}_{it}(\tau_{\min}) - \widehat{TE}_{it}(\tau_{\max})}{\tau_{\max} - \tau_{\min}}. \tag{111}$$

In addition to these scalar summaries, the full profile vector

$$\left(\widehat{TE}_{it}(0.70),\, \widehat{TE}_{it}(0.75),\, \widehat{TE}_{it}(0.80),\, \widehat{TE}_{it}(0.85),\, \widehat{TE}_{it}(0.90),\, \widehat{TE}_{it}(0.95)\right)$$

is retained for graphical analysis, profile classification, endpoint comparisons, and bootstrap confidence-interval construction.

5.4 *Empirical classification of profile types*

Because the profile is the central empirical object, the analysis does not rely exclusively on scalar summaries. A useful intermediate step is to classify firms according to the shape of their efficiency curves.

The baseline procedure is as follows. First, we compute the profile vector $\widehat{TE}_{it}(\tau)$ for each firm-year. Second, profile types are classified by a simple rule-based procedure using the pair $(\widehat{L}_{it}, \widehat{TF}_{it})$. Specifically, observations are partitioned according to whether $\widehat{L}_{it}$ and $\widehat{TF}_{it}$ lie above or below their sample medians. This yields four classes: frontier-near, intermediate, tail-fragile, and uniformly lagging.

This classification is transparent, reproducible, and sufficient for the baseline article. It allows an economically meaningful separation between firms that remain close to the frontier across the quantile range, firms that deteriorate mainly under stricter benchmarks, and firms that remain far from the benchmark throughout the upper-tail range. More elaborate alternatives, such as functional clustering applied directly to the profile vectors, are potentially informative but methodologically heavier and are therefore better treated as extensions rather than as part of the baseline design. The stability of the classification is examined under alternative threshold rules in order to assess whether the four-type structure is robust to moderate changes in the cutoffs used for profile grouping.

5.5 *Second-stage econometric design*

The second-stage econometric design is organized around the idea that determinants and consequences of efficiency may differ depending on whether one studies the overall level of the profile, its upper-tail fragility, or its decline rate. The main empirical implementation reported in Section 6 focuses on the determinants of the profile summaries. In addition, exploratory lagged outcome regressions are estimated as a validation exercise linking profile summaries to subsequent firm performance. A more advanced functional-response specification is retained as a methodological extension.

Accordingly, the econometric framework contains three blocks: baseline regressions for the determinants of the profile summaries; exploratory lagged regressions linking profile summaries to firm outcomes; and a functional-response extension in which the entire efficiency profile is treated as the dependent object.

5.5.1 *Determinants of profile summaries*

This first block constitutes the baseline second-stage analysis implemented in the empirical results. The baseline specification explains scalar profile summaries by time-varying firm characteristics using fixed-effects panel regressions of the form

$$Z_{it} = a_i + \lambda_t + \beta_1 Leverage_{it} + \beta_2 Size_{it} + \varepsilon_{it},\ Z_{it} \in \{\widehat{TE}_{it}(0.70), \widehat{TE}_{it}(0.95), \widehat{L}_{it}, \widehat{TF}_{it}, \widehat{D}_{it}\}. \tag{112}$$

Here a_i denotes firm fixed effects and λ_t denotes year fixed effects. Standard errors are clustered at the globally unique firm level. Leverage is measured as total debt divided by total assets, while size is measured as the logarithm of total assets. This specification is designed to isolate within-firm variation over time and to distinguish whether changes in financial structure and firm scale are associated mainly with the average level of the profile or with its upper-tail fragility. Age is not included in the two-way fixed-effects specification because it evolves mechanically with calendar time and is therefore absorbed by the combination of firm and year effects.

To test whether the second-stage results are sensitive to model specification, two additional variants are estimated. The first removes $Size_{it}$ from the right-hand side in order to assess whether firm scale absorbs part of the association between leverage and efficiency summaries. The second replaces common year effects by industry-year effects:

$$Z_{it} = a_i + \lambda_{s(i),t} + \beta_1 Leverage_{it} + \beta_2 Size_{it} + \varepsilon_{it}, \tag{113}$$

where $\lambda_{s(i),t}$ captures industry-specific shocks in each year. This robustness specification is useful because the sample period includes large aggregate and sectoral disturbances, and because firms in different industries may face different benchmark shifts over the same calendar year.

5.5.2 *Profile summaries and firm outcomes*

The second class of models links profile summaries to firm performance. In contrast to the baseline determinant regressions, these models are used as exploratory validation checks rather than as causal tests. The general contemporaneous form is

$$Performance_{it} = a_i + \lambda_t + \delta_1 \widehat{L}_{it} + \delta_2 \widehat{TF}_{it} + \delta_3 Leverage_{it} + \delta_4 Size_{it} + u_{it}, \tag{114}$$

where $Performance_{it}$ denotes a firm-performance measure such as ROA or sales growth.

To reduce simultaneity concerns and focus on predictive associations, the reported validation regressions use lagged profile summaries and lagged controls:

$$Performance_{it} = a_i + \lambda_t + \delta_1 \widehat{L}_{i,t-1} + \delta_2 \widehat{TF}_{i,t-1} + \delta_3 Leverage_{i,t-1} + \delta_4 Size_{i,t-1} + u_{it}. \tag{115}$$

The dependent variables are winsorized ROA and winsorized log sales growth. Winsorization is applied within years in order to reduce the influence of extreme observations, which is particularly important in a short panel covering major external shocks. The reported specifications are estimated both with and without lagged size. The version without size is included because firm size may partly mediate the relationship between frontier proximity and subsequent performance; conditioning on size may therefore be conservative if scale is itself part of the adjustment channel.

This extension is conceptually aligned with recent methodological discussion emphasizing that the empirical meaning of downstream regressions depends crucially on what is brought from the first stage into the second stage [14–16]. At the same time, these models are interpreted descriptively. They do not identify causal effects of efficiency profiles on firm performance because profile summaries, financial structure, size, and outcomes may be jointly determined by unobserved shocks and managerial choices. Stronger causal interpretation would require additional identification strategies, such as external instruments, quasi-experimental shocks, difference-in-differences designs, or longer panels with richer dynamic structure.

5.5.3 *Functional-response extension*

As a more advanced extension, the profile itself may be treated as the dependent variable in a function-on-scalars regression:

$$\widehat{TE}_{it}(\tau) = a_i(\tau) + \lambda_t(\tau) + X_{it}'\beta(\tau) + e_{it}(\tau), \ \tau \in \mathscr{T}. \tag{116}$$

This specification is attractive because it reveals how explanatory variables shift different parts of the profile rather than only its scalar summaries. However, it is methodologically heavier and is not implemented in the baseline version of the paper. For that reason, we treat it as a possible extension and keep the main empirical analysis centered on $\widehat{L}_{it}$ and $\widehat{TF}_{it}$, with $\widehat{D}_{it}$ playing a complementary shape-oriented role.

5.6 *Robustness and validation design*

All panel regressions use heteroskedasticity-robust inference with standard errors clustered at the globally unique firm level. As a complementary robustness perspective, recent methodological work on multiple regression explicitly examines estimation under outlier contamination and outlier-management procedures [26]. The empirical analysis therefore considers several complementary robustness and validation checks.

First, the profile summaries are recomputed on alternative subranges of the baseline quantile grid in order to assess whether the main empirical patterns depend strongly on the exact choice of benchmark coverage. Second, the rule-based classification of firm-year observations is re-examined under alternative threshold rules to test whether the four-class structure is stable to moderate changes in the cutoffs used for grouping. Third, the profile summaries are compared with standard DEA-based efficiency measures in order to determine whether the proposed profile representation contains information beyond conventional scalar benchmarking. Fourth, model-specification robustness is examined by comparing the baseline two-way fixed-effects model with a no-size specification and with a specification that includes industry-year shocks. Fifth, 95% cluster-bootstrap confidence intervals are computed for the mean empirical efficiency profiles. The bootstrap is clustered at the firm level and is interpreted conditional on the first-stage estimated profiles. Sixth, exploratory lagged performance regressions are estimated for winsorized ROA and winsorized log sales growth in order to examine whether profile summaries contain predictive information about subsequent firm performance.

We do not use SFA as the main scalar benchmark in this robustness comparison. The empirical window covers two major external shocks for Russian firms: the pandemic shock of 2020 and the geopolitical and sanctions-related shock of 2022. In such a short panel, a parametric SFA comparison would require strong distributional assumptions to separate persistent inefficiency, transitory statistical noise, and shock-induced output variation. For this reason, the robustness comparison uses DEA-type scalar efficiency as a transparent nonparametric benchmark, while SFA-based validation is left for longer and more stable panels.

Taken together, these checks are intended not to replace the baseline analysis, but to verify that the main empirical distinction between level-type frontier proximity and profile-shape fragility remains meaningful under alternative empirical choices, benchmark ranges, classification rules, and validation exercises.

6. Empirical results

This section evaluates the empirical usefulness of the proposed quantile efficiency profile and its scalar summaries. The analysis proceeds in six steps. First, we document the general shape of the estimated profiles, summarize the distribution of the profile-based indices, and report bootstrap confidence intervals for the mean empirical profile. Second, we compare the profile representation with endpoint-based quantile benchmarking and DEA-based scalar efficiency scores. Third, we classify firm-year observations by profile type and examine the stability of the resulting four-class structure. Fourth, we examine whether the proposed summaries are systematically related to firm characteristics in panel regressions. Fifth, we report robustness checks with respect to alternative quantile ranges and second-stage model specifications. Sixth, we provide exploratory lagged performance regressions linking profile summaries to subsequent profitability and sales growth.

6.1 *General properties of the estimated efficiency profiles*

The first empirical finding is that the estimated profiles are well behaved and exhibit the expected ordering across quantile levels. For all 12,320 firm-year observations, the estimated efficiency curves are monotone in the benchmark strictness parameter: efficiency weakly declines as the benchmark moves from $\tau = 0.70$ to $\tau = 0.95$. This is consistent with the theoretical properties established in Sections 3 and 4 and confirms that the non-crossing quantile construction yields economically coherent profiles in the data.

At the lower end of the grid, the average efficiency equals 0.6956, while at the upper end it falls to 0.3616. The mean value of the average profile level is 0.5496, and the mean tail fragility index is 0.1935. Only about 4.9% of observations

display a flat frontier-at-one profile, indicating that the profile construction is empirically nondegenerate and that the upper part of the quantile range is substantially more demanding than the lower part.

Figure 1 illustrates several representative firm-year quantile efficiency profiles and provides a first visual impression of the heterogeneity in profile shape across observations.

Figure 1 shows that firms differ not only in average proximity to the frontier but also in the shape of their profiles. Some curves remain high and almost flat over the whole quantile range, whereas others decline sharply as the benchmark becomes stricter. To summarize this heterogeneity more systematically, Table 2 reports descriptive statistics for the main profile summaries.

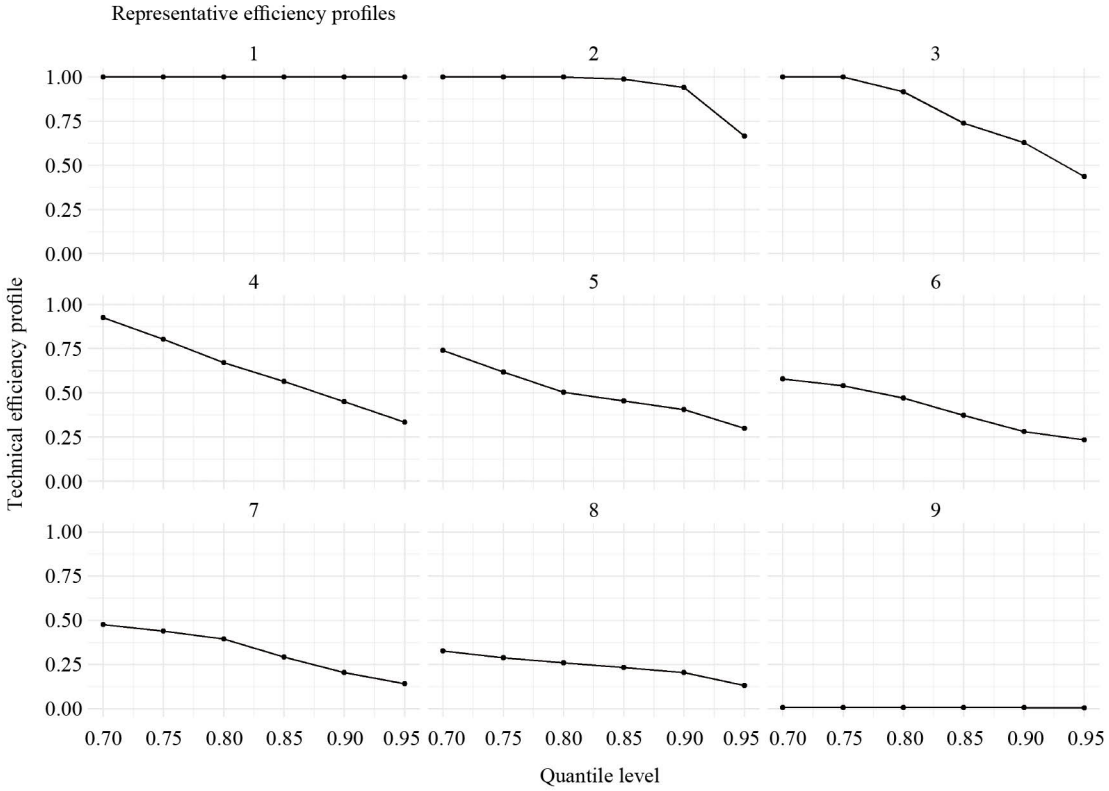


Figure 1. Representative quantile efficiency profiles
Source: Calculated by the authors based on research materials

Table 2 shows that both profile summaries have substantial dispersion and are far from degenerate. The average profile level ranges from almost zero to one, while the fragility index covers a wide empirical interval. This variation is essential for the empirical role of the paper because it implies that the proposed summaries distinguish firms not only by their average benchmark proximity but also by their upper-tail benchmark sensitivity.

Table 2. Descriptive statistics for the profile summaries

Variable	N	Mean	SD	Min	P25	Median	P75	Max
L	12,320	0.5496	0.2659	0.0071	0.3256	0.5043	0.7870	1.0000
TF	12,320	0.1935	0.1042	0.0000	0.1195	0.1895	0.2685	0.6562

Source: Calculated by the authors based on research materials

To reflect uncertainty in the estimated mean empirical profile, Figure 2 reports the average quantile efficiency profile together with 95% cluster-bootstrap confidence intervals. The bootstrap is clustered at the firm level and is interpreted conditional on the first-stage estimated profiles.

Figure 2 confirms that the average efficiency profile declines smoothly and substantially over the upper-quantile range. The confidence band is relatively narrow, indicating that the mean profile is estimated with high precision. At the same time, the decline from the moderate upper quantile to the most demanding upper-tail benchmark is economically large. This supports the main empirical premise of the paper: evaluating firms only at one benchmark point would miss an important part of benchmark-strictness heterogeneity.

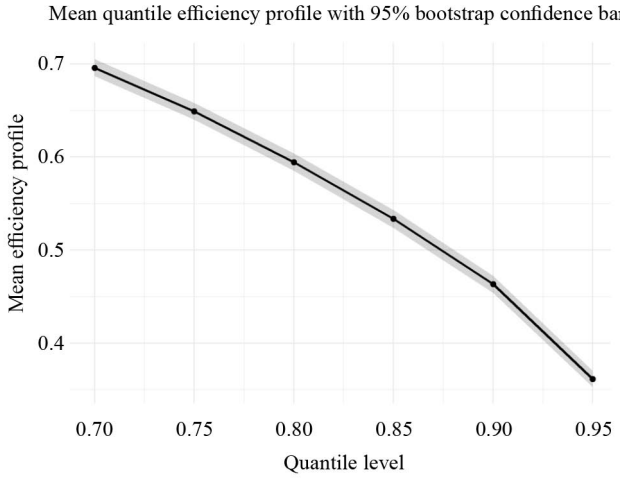


Figure 2. Mean quantile efficiency profile with 95% cluster-bootstrap confidence band
Source: Calculated by the authors based on research materials

6.2 *Comparison with endpoint benchmarking and DEA-based scalar measures*

A central empirical question is whether the profile representation provides information beyond a small number of quantile-specific efficiency scores. To address this issue, we first compare the proposed summaries with endpoint efficiency measures at the lower and upper ends of the selected quantile range. Specifically, we use $\widehat{TE}_{it}(0.70)$ and $\widehat{TE}_{it}(0.95)$ as benchmark-specific endpoint measures and compare them with the average profile level $\widehat{L}_{it}$, the tail fragility index $\widehat{TF}_{it}$, and the decline rate $\widehat{D}_{it}$.

Table 3 reports the correlation matrix for endpoint efficiency measures and profile summaries.

Table 3. Correlations between endpoint efficiency measures and profile summaries

Variable	$TE_{0.70}$	$TE_{0.95}$	L	TF	D
$TE_{0.70}$	1.000	0.798	0.954	0.379	0.433
$TE_{0.95}$	0.798	1.000	0.926	−0.232	−0.198
L	0.954	0.926	1.000	0.107	0.167
TF	0.379	−0.232	0.107	1.000	0.964
D	0.433	−0.198	0.167	0.964	1.000

Source: Calculated by the authors based on research materials

The average profile level is, as expected, strongly correlated with the endpoint measures: its correlation is 0.954 with $\widehat{TE}_{it}(0.70)$ and 0.926 with $\widehat{TE}_{it}(0.95)$. This confirms that $\widehat{L}_{it}$ is a level-type summary of the profile. At the same time, the tail fragility index behaves differently. Its correlation with $\widehat{L}_{it}$ is only 0.107, which is a key empirical result: shape-related profile information is not absorbed by average benchmark proximity. Thus, the full profile contains information that goes beyond discrete endpoint benchmarking.

The decline rate $\widehat{D}_{it}$ is strongly correlated with $\widehat{TF}_{it}$, with a correlation of 0.964. This indicates that both measures capture closely related aspects of upper-tail deterioration. For this reason, the empirical interpretation below places the main emphasis on the pair $(\widehat{L}_{it}, \widehat{TF}_{it})$, while treating $\widehat{D}_{it}$ as a complementary robustness-oriented shape summary.

To place the profile-based representation in a broader benchmarking context, Table 4 additionally reports correlations between the profile summaries and two DEA-based efficiency measures, $\widehat{TE}_{it}^{out}$ and $\widehat{TE}_{it}^{in}$.

Table 4. Correlations between endpoint efficiency measures, profile summaries, and DEA-based efficiency scores

Variable	$TE_{0.70}$	$TE_{0.95}$	L	TF	D	TEout	TEin
$TE_{0.70}$	1.000	0.798	0.954	0.379	0.433	0.590	0.226
$TE_{0.95}$	0.798	1.000	0.926	−0.232	−0.198	0.754	0.371
L	0.954	0.926	1.000	0.107	0.167	0.690	0.308
TF	0.379	−0.232	0.107	1.000	0.964	−0.194	−0.205
$D4$	0.433	−0.198	0.167	0.964	1.000	−0.169	−0.187
TEout	0.590	0.754	0.690	−0.194	−0.169	1.000	0.618
TEin	0.226	0.371	0.308	−0.205	−0.187	0.618	1.000

Source: Calculated by the authors based on research materials

The results reinforce the earlier interpretation. The average profile level is positively correlated with both DEA measures, especially with the output-oriented score. By contrast, tail fragility remains only weakly and negatively related to the DEA-based scores. Thus, the profile summaries are not reducible to standard scalar efficiency measures: level-type information overlaps partly with conventional benchmarking, but the shape-related dimension captured by tail fragility remains largely distinct.

This comparison also clarifies the role of DEA and SFA in the empirical validation. DEA is used here as a transparent nonparametric scalar benchmark. SFA is not used as the main scalar comparison because, as discussed in Section 5.6, the short 2019-2023 panel includes major external shocks that complicate the parametric separation of inefficiency, statistical noise, and shock-induced output variation. The DEA comparison is therefore not intended to exhaust all possible scalar alternatives, but to provide a clear benchmark showing whether the proposed profile summaries overlap with or differ from a standard one-number efficiency measure.

To complement the correlation evidence, Figure 3 displays several firm-year observations with similar endpoint efficiency values but different interior profile shapes.

Figure 3 shows that observations with similar values of $\widehat{TE}_{it}(0.70)$ and $\widehat{TE}_{it}(0.95)$ may still have visibly different intermediate trajectories. In the reported examples, the selected observations have similar lower-end and upper-end efficiency values, yet their internal profile geometry differs substantially. This visual evidence is important because it shows directly why endpoint-based benchmarking alone does not recover the full geometry of benchmark sensitivity.

Overall, the endpoint and DEA comparisons support the main empirical claim of the profile approach. The average profile level behaves similarly to conventional level-type efficiency measures, whereas tail fragility captures a separate shape-related dimension. Consequently, the proposed profile representation is not merely a more elaborate notation for endpoint efficiency. It adds information about how relative performance changes inside the upper-quantile interval as the benchmark becomes progressively more demanding.

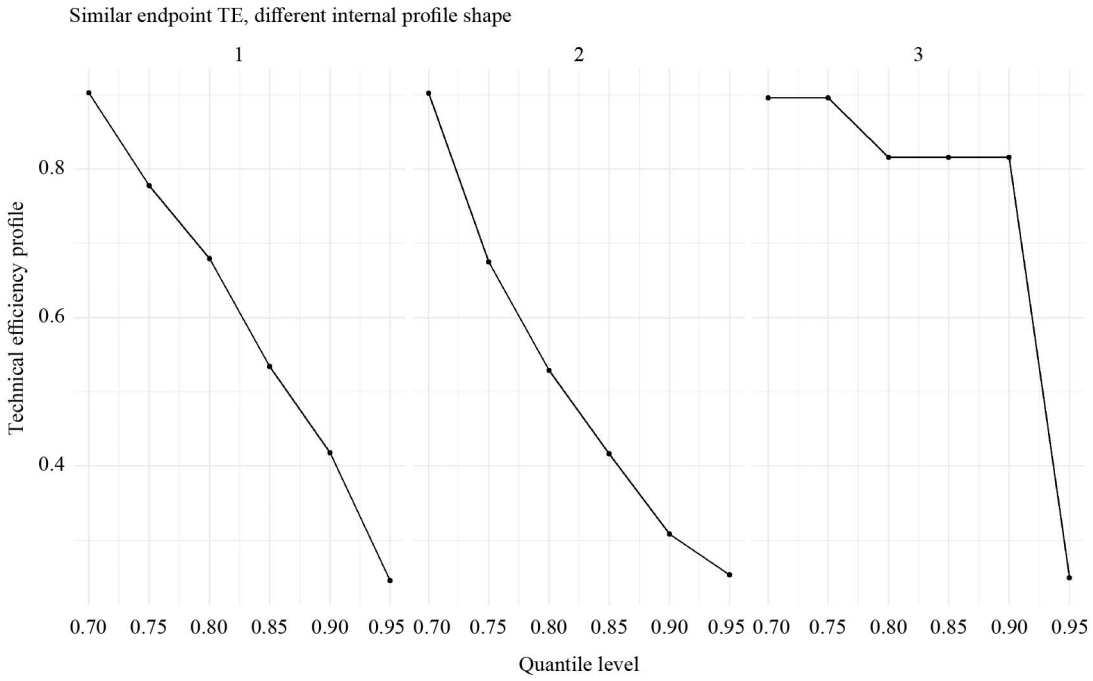


Figure 3. Firm-year observations with similar endpoint efficiency but different internal profile shape
Source: Calculated by the authors based on research materials

6.3 *Profile-based classification of firms*

The profile representation also supports a meaningful classification of firm-year observations. Using the summaries $\widehat{L}_{it}$ and $\widehat{TF}_{it}$, the sample can be partitioned into four profile-based classes. The classification is based on whether the average profile level and tail fragility index are above or below their sample medians. This rule yields four groups: frontier-near, intermediate, tail-fragile, and uniformly lagging firms.

Table 5 reports the resulting classes and their average profile characteristics.

Table 5. Profile-based classes of firm-year observations

Profile class	Grouping criteria	N	Mean L	Mean TF	Mean $TE_{0.70}$	Mean $TE_{0.95}$
frontier-near	$L \geq \text{med}_L$ and $TF < \text{med}_{TF}$	1,956	0.9523	0.0651	0.9872	0.8296
intermediate	$L \geq \text{med}_L$ and $TF \geq \text{med}_{TF}$	4,204	0.6949	0.2994	0.9097	0.4089
tail-fragile	$L < \text{med}_L$ and $TF \geq \text{med}_{TF}$	1,956	0.4229	0.2326	0.6158	0.2245
uniformly lagging	$L < \text{med}_L$ and $TF < \text{med}_{TF}$	4,204	0.2759	0.1291	0.3830	0.1604

Source: Calculated by the authors based on research materials

The resulting classes are economically interpretable and empirically distinct. The frontier-near class has the highest average profile level, the lowest tail fragility, and the highest upper-end efficiency. At the opposite end, the uniformly lagging class has a much lower average profile level and low efficiency even at the lower end of the quantile range, indicating persistent distance from the benchmark throughout the profile.

The intermediate class retains a relatively high average profile level but also displays the largest average tail fragility. This means that firms in this class may appear reasonably competitive under moderate benchmark strictness, but their

relative position deteriorates sharply when the benchmark moves toward the most demanding upper-tail quantiles. By contrast, the tail-fragile class has a lower average profile level than the intermediate class and also exhibits substantial deterioration across the quantile range. This distinction is important because it shows that profile shape cannot be reduced to a simple high-efficiency versus low-efficiency contrast.

The geometric distinctness of these classes becomes clearer when one considers their average efficiency curves, shown in Figure 4.

In the present results, the class means show four clearly distinct geometries. The frontier-near class remains close to one throughout most of the grid and declines meaningfully only at the highest quantiles. The intermediate class starts high but falls steadily, especially in the upper tail. The tail-fragile class begins at a moderate level and deteriorates rapidly across the grid. The uniformly lagging class remains low at every quantile level. This figure is one of the most important empirical outputs of the paper because it demonstrates that the profile-based classification reflects real geometric differences rather than arbitrary scalar grouping.

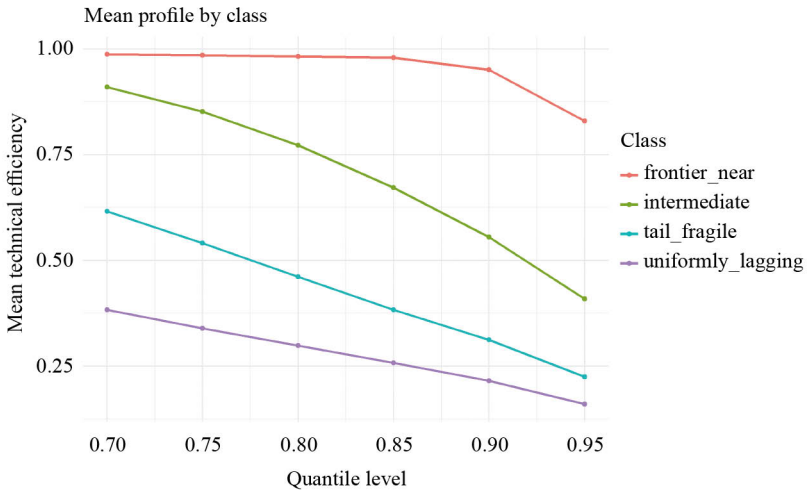


Figure 4. Mean quantile efficiency profile by class
Source: Calculated by the authors based on research materials

Overall, the classification exercise supports the main interpretation of the profile framework. It shows that firm-year observations differ not only in their average distance from the frontier but also in how their relative position changes as the benchmark becomes more demanding. In economic terms, the classification separates overall frontier proximity from upper-tail weakness; in statistical terms, it demonstrates that the profile contains shape information that is not captured by a one-dimensional efficiency measure.

6.4 *Panel regressions for endpoint efficiency measures and profile summaries*

The preceding sections show that the profile summaries are empirically distinct and support economically interpretable classification. We now examine whether these summaries are systematically related to firm characteristics in a fixed-effects panel framework. The purpose of this exercise is not to identify causal effects, but to assess whether the proposed profile-based measures behave in an economically coherent and empirically nonredundant way.

Table 6 reports fixed-effects panel regressions for the two endpoint efficiency measures, $\widehat{TE}_{it}(0.70)$ and $\widehat{TE}_{it}(0.95)$, and for the two main profile summaries, $\widehat{L}_{it}$ and $\widehat{TF}_{it}$. The regressions include firm fixed effects and year fixed effects, and standard errors are clustered at the firm level. The explanatory variables are leverage and size.

The baseline results show that the average profile level behaves as a level-type efficiency measure, while the tail fragility index retains a distinct shape-oriented interpretation. The coefficient on leverage is negative and statistically

significant in all four regressions. This indicates that higher leverage is associated with lower endpoint efficiency, lower average profile level, and lower tail fragility. The first part of this result is straightforward: more leveraged firms tend to be farther from the frontier on average. The result for tail fragility is more specific: conditional on firm and year effects, leverage is also associated with changes in the shape of the profile, not only with its average level.

The coefficient on size is positive and statistically significant in the regressions for $\widehat{TE}_{it}(0.70)$, $\widehat{TE}_{it}(0.95)$, and $\widehat{L}_{it}$. This suggests that larger firms tend to remain closer to the estimated frontier across the upper-quantile range. This finding is consistent with the interpretation that firm scale is associated with stronger average benchmark proximity, possibly because larger firms have more stable production organization, broader resource bases, or better ability to operate near high-performance benchmarks.

Table 6. Fixed-effects panel regressions for endpoint efficiency measures and profile summaries

Variables	$TE_{0.70}$	$TE_{0.95}$	L	TF
Leverage	−0.0010***	−0.0007***	−0.0008***	−0.0002***
	(0.0002)	(0.0001)	(0.0001)	(0.00008)
Size	0.0775***	0.0867***	0.0837***	−0.0070*
	(0.0072)	(0.0069)	(0.0067)	(0.0039)
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	12,320	12,320	12,320	12,320
R^2	0.8373	0.8494	0.8637	0.6447
Within-R^2	0.0345	0.0485	0.0460	0.0029

Source: Calculated by the authors based on research materials. Notes: clustered (firm-level) standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

The regression for $\widehat{TF}_{it}$ displays a different pattern. Size enters with a negative coefficient and is only weakly statistically significant. Thus, larger firms tend to have higher average profile levels, but they do not necessarily display stronger upper-tail fragility. This distinction is important because it shows that the profile summaries are not merely alternative transformations of the same scalar efficiency score. Rather, they separate average frontier proximity from the shape of deterioration as the benchmark becomes more demanding.

The comparison with endpoint efficiency measures is also informative. Both $\widehat{TE}_{it}(0.70)$ and $\widehat{TE}_{it}(0.95)$ behave similarly to $\widehat{L}_{it}$ in terms of coefficient signs: leverage enters negatively and size enters positively. This supports the interpretation of $\widehat{L}_{it}$ as a level-type profile summary. At the same time, $\widehat{TF}_{it}$ continues to behave differently, which is precisely the empirical distinction that motivates the profile-based framework.

More generally, the role of Table 6 is not to maximize explanatory power, but to show that the profile summaries remain empirically nonredundant under a strict panel specification. In particular, $\widehat{L}_{it}$ aligns closely with endpoint efficiency measures, whereas $\widehat{TF}_{it}$ captures a distinct dimension of variation that cannot be reduced to average or endpoint proximity alone. This supports the use of profile-based summaries as second-stage objects rather than relying exclusively on a single benchmark-specific efficiency score.

6.5 *Robustness checks and exploratory validation*

The final empirical step is to examine whether the main conclusions are robust to alternative empirical choices and whether the proposed profile summaries contain information relevant for subsequent firm performance. The robustness

analysis focuses on five dimensions: the choice of the quantile range, the stability of profile classification, comparison with scalar benchmarking, second-stage model specification, and lagged performance validation.

First, we test whether the determinant patterns for the profile summaries depend on the chosen quantile interval. The baseline profile is computed over the range [0.70, 0.95]. As a sensitivity check, the profile summaries are recomputed on three alternative subranges: [0.75, 0.95], [0.70, 0.90], and [0.75, 0.85]. Table 7 reports fixed-effects regressions for the average profile level and tail fragility index under these alternative definitions.

Table 7 shows that the determinant patterns for the average profile level remain qualitatively stable across all range choices. Leverage remains negative and statistically significant, while size remains positive and statistically significant in every specification. The results for tail fragility are also broadly stable. Leverage remains negative and statistically significant across the alternative ranges, whereas the coefficient on size remains negative but is weaker and less stable. These results indicate that the main empirical distinction between average profile level and upper-tail fragility is not driven by one specific choice of the upper-quantile interval.

Table 7. Robustness of fixed-effects regressions across alternative quantile ranges

Variables	Baseline L	Baseline TF	L 0.75–0.95	TF 0.75–0.95	L 0.70–0.90	TF 0.70–0.90	L 0.75–0.85	TF 0.75–0.85
Leverage	−0.0008***	−0.0002***	−0.0008***	−0.0002**	−0.0009***	−0.0002***	−0.0009***	−0.0002***
	(0.0001)	(0.00008)	(0.0001)	(0.00008)	(0.0001)	(0.00007)	(0.0001)	(0.00005)
Size	0.0837***	−0.0070*	0.0849***	−0.0041	0.0831***	−0.0072**	0.0833***	−0.0053**
	(0.0067)	(0.0039)	(0.0067)	(0.0039)	(0.0069)	(0.0033)	(0.0072)	(0.0024)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,320	12,320	12,320	12,320	12,320	12,320	12,320	12,320
R^2	0.8637	0.6447	0.8643	0.6325	0.8590	0.6091	0.8549	0.5297
Within-R^2	0.0460	0.0029	0.0463	0.0016	0.0420	0.0031	0.0393	0.0029

Source: Calculated by the authors based on research materials. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Second, we examine whether the four-class profile typology is sensitive to the particular median-based threshold rule used in the baseline classification. Table 8 compares the baseline median split with two nearby alternatives based on the 45th and 55th percentile rules.

Table 8. Stability of profile-based classification under alternative threshold rules

Threshold rule	Frontier-near	Intermediate	Tail-fragile	Uniformly lagging	Share reassigned vs. baseline
Baseline median	1,956	4,204	1,956	4,204	–
45th percentile rule	1,877	4,899	1,877	3,667	0.096
55th percentile rule	2,024	3,520	2,024	4,752	0.096

Source: Calculated by the authors based on research materials.

The reassignment rates remain below 10% in both alternative cases. This indicates that the proposed classification is not merely an artifact of one specific threshold choice. Instead, the four-type structure reflects a stable empirical pattern

of profile heterogeneity. This result is especially important because the classification is based on simple and transparent summary rules rather than on a complex clustering algorithm.

Third, the profile summaries are compared with DEA-based scalar efficiency measures. As shown in Section 6.2, the average profile level overlaps partly with conventional scalar benchmarking, whereas tail fragility remains only weakly related to DEA scores. This reinforces the interpretation that the proposed framework adds information beyond standard one-number efficiency measures. The use of DEA as the scalar comparison is also consistent with the design choice discussed in Section 5.6: in a short panel covering the 2020 pandemic shock and the 2022 geopolitical and sanctions-related shock, SFA would require strong distributional assumptions to separate inefficiency, transitory noise, and shock-induced variation.

Fourth, we examine whether the baseline fixed-effects regression results are sensitive to second-stage model specification. Appendix C compares three specifications for $\widehat{TE}_{it}(0.70)$, $\widehat{TE}_{it}(0.95)$, $\widehat{L}_{it}$, and $\widehat{TF}_{it}$: the baseline model with firm and year fixed effects, a no-size specification, and a specification with firm fixed effects and industry-year shocks. The results are stable. Leverage remains negative and statistically significant across all dependent variables and specifications. Size remains positive and statistically significant for endpoint efficiency and the average profile level. The specification with industry-year shocks yields the same qualitative conclusions as the baseline model. This supports the interpretation that the main results are not driven by common macro-year effects or by broad industry-year disturbances.

Fifth, we examine uncertainty in the mean empirical profiles. Figure 2 reports the overall mean quantile efficiency profile with a 95% cluster-bootstrap confidence band. In addition, Appendix B reports class-specific mean profiles with 95% cluster-bootstrap confidence intervals. The class-specific intervals are narrow and therefore visually close to the profile curves. This is consistent with the large number of observations within the profile classes and supports the conclusion that the differences between profile types are systematic rather than driven by sampling noise alone.

Finally, we estimate exploratory lagged performance regressions. These models use winsorized ROA and winsorized log sales growth as dependent variables and lagged profile summaries as the main explanatory variables. The goal is not to identify causal effects, but to evaluate whether the profile summaries contain predictive information about subsequent firm performance. Appendix D shows that the lagged average profile level is positively associated with subsequent winsorized ROA and negatively associated with subsequent winsorized log sales growth. The first result is consistent with the interpretation that firms closer to the frontier tend to have stronger subsequent profitability. The second result is consistent with a catch-up interpretation: firms already closer to the frontier may have less room for rapid sales expansion. Tail fragility is not statistically significant in these exploratory regressions, suggesting that it is more informative as a shape-related heterogeneity measure than as a direct short-run predictor of profitability or sales growth.

Overall, the robustness and validation checks strengthen the empirical interpretation of the paper in five respects. First, the main determinant patterns remain stable when profile summaries are recomputed on alternative quantile ranges. Second, the four-type classification is robust to moderate changes in threshold rules. Third, the DEA comparison shows that the average profile level overlaps partly with conventional scalar efficiency measures, whereas tail fragility captures a distinct dimension of profile shape. Fourth, the second-stage results are stable across alternative model specifications, including industry-year shocks. Fifth, the exploratory lagged performance regressions show that the average profile level is meaningfully related to subsequent firm outcomes. Taken together, these checks support the view that the proposed profile representation captures systematic empirical heterogeneity rather than a purely formal reparameterization of scalar efficiency scores.

7. Discussion

The results of the paper suggest that the main gain from the proposed framework does not lie in producing yet another efficiency score, but in changing the representation of efficiency itself. In the standard frontier literature, the empirical endpoint is still typically a scalar efficiency measure attached to the firm or firm-year observation. Recent reviews of DEA-based economic-efficiency analysis document the continued use of such scalar efficiency measures across a wide range of applications [1]. Work in stochastic frontier analysis has substantially expanded the modelling of panel structure

and unobserved heterogeneity [23, 24, 27]. Other recent work addresses skewness, endogeneity, and related distributional complications in stochastic frontier models [28, 29]. Against this background, the present paper proposes a different unit of analysis: not a point score, but a quantile-indexed efficiency profile. This shift is substantial at the level of interpretation. Instead of asking only how close a firm is to one benchmark, the profile asks how its relative position changes as the benchmark becomes progressively stricter.

This difference is especially important when comparing the present framework with the usual single-point technical efficiency approach. In the scalar setting, the researcher selects one frontier benchmark and all heterogeneity is compressed into a single value. That representation is often sufficient for ranking, but it cannot distinguish between firms that are uniformly close to the frontier and firms that are close to the frontier only under moderate benchmark strictness. The empirical results show that this distinction is not merely conceptual. The estimated profiles are heterogeneous, the average profile level is highly correlated with endpoint efficiency measures, and the tail fragility index captures a separate shape-related dimension. This means that the profile does not merely restate scalar information in more elaborate notation; it reveals an additional dimension of heterogeneity related to upper-tail benchmark sensitivity.

A more demanding comparison is with discrete multi-quantile frontier analysis. In that literature, the benchmark is no longer a single frontier but a finite collection of quantile-specific frontiers. Recent applied work on quantile stochastic frontiers shows that benchmark choice matters for technological and efficiency interpretation [6–8]. A broader methodological formulation makes the quantile dimension explicit at the stochastic-frontier model level [17]. Broader methodological treatments of quantile stochastic frontier analysis reinforce the same point: quantile choice is substantive rather than incidental [4, 5]. Related semiparametric work shows how multiple quantiles can be modeled jointly under a non-crossing condition [30]. These developments go beyond single-point benchmarking, but the empirical object in frontier applications is still usually a small set of quantile-specific scores rather than a function with its own geometry. In parallel, recent SFA research also shows that scalar efficiency estimates may be sensitive to distributional misspecification [31]. These studies represent an important step beyond single-point benchmarking because they recognize that efficiency changes across quantile levels. However, the empirical object is still usually a small set of quantile-specific scores. In such a framework, multiple quantiles are estimated, but efficiency is not yet treated as a function with its own geometry.

This is where the present paper departs most clearly from the existing quantile-frontier literature. The contribution is not simply to estimate more quantiles. Rather, it is to treat the ordered quantile family as the basis for a functional efficiency object and then to derive interpretable summaries from that object. This distinction is worth emphasizing relative to recent work on constrained log-log quantile frontiers as smooth probabilistic benchmarking tools [20]: there the frontier itself is the main object, whereas here the ordered frontier family serves as the basis for a profile representation of firm-year efficiency. In other words, the difference between the present paper and the usual multi-quantile approach is analogous to the difference between observing a handful of points on a curve and treating the curve itself as the primary object.

From a mathematical-statistical perspective, this also clarifies why the paper is not merely an empirical extension of quantile regression. Recent statistical work increasingly treats quantiles as members of a jointly estimated ordered process, rather than as unrelated pointwise regressions [9, 12]. The present paper imports that logic into frontier analysis. The profile is defined not from isolated quantile estimates, but from an ordered log-quantile frontier process. This is what makes it meaningful to study boundedness, monotonicity, dominance ordering, continuity, approximation, and plug-in consistency. The ordered-process construction is therefore not only a technical convenience; it is the condition under which benchmark strictness receives a coherent economic interpretation.

The empirical results also speak to the role of second-stage efficiency analysis. A recurring concern in that literature is that downstream conclusions depend heavily on what exactly is brought from the first stage into the second stage. Work in this area argues that the first-stage object should be chosen carefully because second-stage inferences may be sensitive to the form of the efficiency measure itself [14, 15]. Related contributions on individualized second-stage corrections and Bayesian second-stage designs point in the same direction by showing that second-stage inference can be materially affected by how the first-stage efficiency object is constructed and adjusted [16, 32]. The present paper contributes to that discussion by showing that the issue is not only how to estimate the first-stage score, but also whether the first-stage

object should be a score at all. In heterogeneous firm settings, it may be more informative to treat the first-stage object as a curve and only then reduce it to summaries that preserve distinct aspects of its geometry.

Beyond second-stage design, recent frontier research has continued to broaden the statistical and distributional treatment of efficiency models. A recent survey documents the development of statistical approaches in non-parametric frontier analysis [33]. Distributional stochastic frontier models relax conventional restrictions by allowing richer modelling of the frontier and composed-error structure [34]. Recent decision-tree-augmented SFA further addresses technological heterogeneity through segmented local frontiers for heterogeneous firm subpopulations [35].

The robustness and validation checks strengthen this interpretation. The main determinant patterns remain qualitatively stable when the profile summaries are recomputed on alternative quantile subranges. The four-type profile classification remains stable under nearby threshold rules. The comparison with DEA-based scalar efficiency scores shows that the average profile level overlaps partly with conventional scalar benchmarking, whereas tail fragility remains only weakly related to DEA scores. The model-specification checks further show that the main second-stage patterns survive the exclusion of size and the inclusion of industry-year shocks. The cluster-bootstrap confidence intervals for the mean profiles indicate that the main profile patterns are estimated with high precision. Together, these results support the view that the proposed profile representation captures systematic empirical heterogeneity rather than a purely mechanical transformation of scalar scores.

The exploratory performance regressions provide an additional practical validation. The lagged average profile level is positively associated with subsequent winsorized ROA, suggesting that firms closer to the frontier tend to display stronger subsequent profitability. At the same time, the lagged average profile level is negatively associated with subsequent winsorized log sales growth. This pattern is consistent with a catch-up interpretation: firms already closer to the frontier may have less room for rapid sales expansion, whereas firms farther from the frontier may experience stronger growth when they adjust toward better-performing peers. Tail fragility is not statistically significant in these exploratory regressions, which suggests that it should be interpreted primarily as a shape-related heterogeneity measure rather than as a direct short-run predictor of profitability or growth. This does not weaken the role of tail fragility; rather, it clarifies its empirical meaning.

The framework also has practical implications for managers and policymakers. For managers, the average profile level provides a broad diagnostic of overall proximity to upper-tail benchmarks, while tail fragility indicates whether the firm's relative position deteriorates when compared with the most demanding performers. This distinction is useful because two firms with similar average efficiency may require different managerial responses: one may need general productivity improvement, while another may need targeted upgrading to remain competitive under stricter frontier comparisons. For policymakers, profile-based classification can support more differentiated benchmarking and support policies. Uniformly lagging firms may require broad capability-building measures, whereas tail-fragile firms may benefit more from interventions aimed at technological depth, resilience, and movement toward best-practice performance standards.

At the same time, the results should be read with caution. The second-stage regressions are descriptive and should not be interpreted as causal evidence. Firm efficiency profiles, leverage, size, profitability, and growth may be jointly shaped by unobserved managerial quality, market shocks, demand conditions, or financing constraints. The lagged performance regressions reduce simultaneity concerns but do not by themselves solve endogeneity. Stronger causal interpretation would require additional identification strategies, such as external instruments, quasi-experimental shocks, difference-in-differences designs, or longer dynamic panels in which changes in profile shape can be separated more clearly from unobserved firm-level shocks.

Several methodological limitations also define the boundary of the present implementation. First, the empirical model uses a single-output log-log frontier specification. This is appropriate for a transparent first implementation but does not exhaust the possibilities of the profile approach. Multi-output settings would require extending the profile idea to output-distance, directional-distance, or revenue/cost frontier formulations. Second, the current frontier is estimated on a fixed upper-quantile grid. Although robustness checks indicate that the main conclusions are stable across alternative subranges, future work could study data-driven grid selection or adaptive quantile ranges. Third, the baseline summaries

are intentionally parsimonious. They are chosen for interpretability, but richer functional summaries, functional principal components, or clustering methods could recover additional profile-shape information.

The framework is also generalizable beyond the specific Russian firm-panel application. The empirical setting used here is useful because it contains substantial heterogeneity and covers a period with major external disturbances, but the profile logic is not specific to Russia. Any setting in which firms, plants, banks, hospitals, universities, or regions are evaluated against ordered upper-tail benchmarks may benefit from a profile representation. Cross-country applications would be especially useful for studying whether firms in different institutional environments differ mainly in average frontier proximity or in their upper-tail fragility. Similarly, dynamic profile models could examine whether firms improve by shifting the whole profile upward, by reducing tail fragility, or by changing the curvature of the profile over time.

Overall, the discussion reinforces the central point of the paper. Relative to standard point analysis, the proposed approach preserves richer information about benchmark strictness. Relative to discrete multi-quantile analysis, it replaces a collection of selected efficiency values with a functional object that has its own theoretical and empirical structure. Relative to conventional second-stage efficiency analysis, it provides a richer first-stage object from which level-type, shape-type, and classification-based measures can be derived. The resulting framework remains parsimonious enough for panel implementation, but it offers a more informative representation of firm heterogeneity than a single benchmark-specific efficiency score.

8. Conclusion

This paper develops a functional extension of quantile-based frontier efficiency analysis for firm-level panel data. Instead of reducing technical efficiency to a single benchmark-specific score, the proposed framework represents each firm-year observation by a quantile-indexed efficiency profile. The profile records how the relative position of a firm changes as the benchmark moves from moderately high to increasingly demanding upper conditional quantiles. This makes it possible to study efficiency not only as distance from one frontier point, but also as a pattern of sensitivity to benchmark strictness.

The methodological contribution of the paper is to treat the ordered quantile frontier family as the basis for a functional efficiency object. Within a log-log quantile frontier setting, the paper defines the efficiency profile and derives three interpretable summaries: the average profile level, the tail fragility index, and the decline rate. The theoretical analysis shows that the profile is well defined, bounded, monotone under non-crossing benchmark frontiers, stable under perturbations, and consistent under uniform convergence of the estimated quantile frontier. These properties provide a coherent foundation for using profile-based summaries in empirical efficiency analysis.

The empirical application to Russian firms over 2019-2023 shows that the proposed profile representation is informative in practice. The estimated profiles are nondegenerate and heterogeneous in shape. The average profile level behaves similarly to conventional level-type efficiency measures, whereas tail fragility captures an additional shape-related dimension of heterogeneity. The classification of firm-year observations into frontier-near, intermediate, tail-fragile, and uniformly lagging types provides an economically interpretable representation of this heterogeneity. Robustness checks based on alternative quantile ranges, classification thresholds, model specifications, DEA-based scalar benchmarks, bootstrap confidence intervals, and lagged performance regressions support the empirical relevance of the profile approach.

The results also have practical implications. For managers, the average profile level can serve as a diagnostic of overall proximity to upper-tail benchmarks, while tail fragility identifies firms whose relative position deteriorates under the most demanding comparisons. For policymakers, profile-based classification can support more differentiated benchmarking and support measures, since uniformly lagging firms and tail-fragile firms may require different interventions.

Several limitations remain. The empirical implementation uses a single-output log-log frontier and a fixed upper-quantile grid, while the second-stage regressions are descriptive rather than causal. Future research may extend the framework to multi-output, semi-parametric and non-parametric frontiers, cross-country applications, adaptive quantile

grids, and dynamic profile models. Overall, the paper shows that replacing one-number efficiency measures with quantile-indexed efficiency profiles yields a richer and more informative representation of firm heterogeneity, especially when the meaning of efficiency depends on how demanding the benchmark is.

Use of AI tools declaration

During the preparation and revision of this work, the authors used Writefull for text embellishment, including improvement of language clarity and readability. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Author contributions

Conceptualization, V.S. and M.G.; Methodology, V.S.; Software,V.S. and V.L.; Validation, V.L.; Formal analysis, V.S., M.G. and L.S.; Investigation, V.S., M.G. and L.S.; Data curation, L.S. and V.L.; Writing–original draft, V.S., L.S., M.G. and V.L.; Writing–review & editing, V.S., M.G. and L.S.; Visualization, V.L.; Supervision, M.G. All authors have read and agreed to the published version of the manuscript.

Funding

This study was supported by the Russian Science Foundation (RSF), project No. 25-28-00731 "Technological leadership and digital technologies as key factors for the development of Russian firms in the context of economic instability: analysis and modeling". The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Data availability

The firm-level data used in this study were obtained from the SPARK-Interfax information system. The underlying raw data cannot be publicly redistributed by the authors. The analysis code is not publicly available due to data licensing restrictions but may be provided by the authors upon reasonable request.

Conflict of interest

The authors declare no competing financial interest.

References

[1] Camanho AS, Silva MC, Piran FS, Lacerda DP. A literature review of economic efficiency assessments using data envelopment analysis. *European Journal of Operational Research*. 2024; 315(1): 1-18. Available from: https://doi.org/10.1016/j.ejor.2023.07.027.

[2] Sousa R, Camanho AS, Silva MC, da Silveira GJC, Arabi B. Best practices, performance advantage and trade-offs: New insights from frontier analysis. *Journal of Productivity Analysis*. 2024; 62(1): 91-110. Available from: https://doi.org/10.1007/s11123-023-00714-y.

[3] Guillen MD, Aparicio J, Zofío JL, España VJ. Improving the predictive accuracy of production frontier models for efficiency measurement using machine learning: The LSB-MAFS method. *Computers & Operations Research*. 2024; 171: 106793. Available from: https://doi.org/10.1016/j.cor.2024.106793.

[4] Papadopoulos A, Parmeter CF. Quantile methods for stochastic frontier analysis. *Foundations and Trends in Econometrics*. 2022; 12(1): 1-120. Available from: https://doi.org/10.1561/0800000042.

[5] Jradi S, Parmeter CF, Ruggiero J. Quantile estimation of stochastic frontiers with the normal-exponential specification. *European Journal of Operational Research*. 2021; 295(2): 475-483. Available from: https://doi.org/10.1016/j.ejor.2021.03.002.

[6] Tsionas MG, Tzeremes NG. Eco-efficiency estimation with quantile stochastic frontiers: Evidence from the United States. *Journal of Environmental Management*. 2022; 320: 115876. Available from: https://doi.org/10.1016/j.jenvman.2022.115876.

[7] Cabas Monje J, Guesmi B, Ait Sidhoum A, Gil JM. Measuring technical efficiency of Spanish pig farming: Quantile stochastic frontier approach. *Australian Journal of Agricultural and Resource Economics*. 2023; 67(4): 688-703. Available from: https://doi.org/10.1111/1467-8489.12518.

[8] Fusco E, Benedetti R, Vidoli F. Stochastic frontier estimation through parametric modelling of quantile regression coefficients. *Empirical Economics*. 2023; 64(2): 869-896. Available from: https://doi.org/10.1007/s00181-022-02273-x.

[9] Shen G, Jiao Y, Lin Y, Horowitz JL, Huang J. Nonparametric estimation of non-crossing quantile regression process with deep ReQU neural networks. *Journal of Machine Learning Research*. 2024; 25(88): 1-75. Available from: https://www.jmlr.org/papers/volume25/22-0488/22-0488.pdf.

[10] Gertheiss J, Rügamer D, Liew BXW, Greven S. Functional data analysis: An introduction and recent developments. *Biometrical Journal*. 2024; 66(7): e202300363. Available from: https://doi.org/10.1002/bimj.202300363.

[11] Wang S, Zhang W, Cao G, Huang Y. Functional data analysis using deep neural networks. *WIREs Computational Statistics*. 2024; 16(4): e70001. Available from: https://doi.org/10.1002/wics.70001.

[12] Dai S, Kuosmanen T, Zhou X. Non-crossing convex quantile regression. *Economics Letters*. 2023; 233: 111396. Available from: https://doi.org/10.1016/j.econlet.2023.111396.

[13] da Silva Marioni L, Rincon-Aznar A, Venturini F. Productivity performance, distance to frontier and AI innovation: Firm-level evidence from Europe. *Journal of Economic Behavior & Organization*. 2024; 228: 106762. Available from: https://doi.org/10.1016/j.jebo.2024.106762.

[14] Holý V. Ranking-based second stage in data envelopment analysis: An application to research efficiency in higher education. *Operations Research Perspectives*. 2024; 12: 100306. Available from: https://doi.org/10.1016/j.orp.2024.100306.

[15] Frýd L, Sokol O. Hidden heterogeneity in measuring production factors: Implications for two-stage efficiency analysis. *Socio-Economic Planning Sciences*. 2026; 104: 102418. Available from: https://doi.org/10.1016/j.seps.2026.102418.

[16] Henriques IC, Sobreiro VA, Kimura H, Mariano EB. Two-stage DEA in banks: terminological controversies and future directions. *Expert Systems with Applications*. 2020; 161: 113632. Available from: https://doi.org/10.1016/j.eswa.2020.113632.

[17] Tsionas MG. Quantile stochastic frontiers. *European Journal of Operational Research*. 2020; 282(3): 1177-1184. Available from: https://doi.org/10.1016/j.ejor.2019.10.012.

[18] Wang L, Xue L, Yang L. Estimation of additive frontier functions with shape constraints. *Journal of Nonparametric Statistics*. 2020; 32(2): 262-293. Available from: https://doi.org/10.1080/10485252.2020.1721494.

[19] Muggeo VMR, Sottile G, Cilluffo G. Joint modelling of non-crossing additive quantile regression via constrained B-spline varying coefficients. *Statistical Modelling*. 2023; 23(5-6): 540-554. Available from: https://doi.org/10.1177/1471082X231181734.

[20] Spitsin V, Martyushev N, Spitsina L, Gasanov M, Leonova V. DEA-inspired constrained log-log quantile frontier: Smooth benchmarking, calibration, and dynamic interpretation. *Contemporary Mathematics*. 2026; 7(3): 2770-2801. Available from: https://doi.org/10.37256/cm.7320269045.

[21] Li Y, Qiu Y, Xu Y. From multivariate to functional data analysis: fundamentals, recent developments, and emerging areas. *Journal of Multivariate Analysis*. 2022; 188: 104806. Available from: https://doi.org/10.1016/j.jmva.2021.104806.

[22] Interfax. *SPARK-Interfax information and analytical system*. 2026. Available from: https://spark-interfax.ru/ [Accessed 6th June 2026].

[23] Parmeter CF, Kumbhakar SC. The generalized panel data stochastic frontier model: A review and nonparametric estimation. *Journal of Productivity Analysis*. 2025; 64: 321-339. Available from: https://doi.org/10.1007/s11123-025-00769-z.

[24] Bellio R, Grassetti L. Efficient estimation of true fixed-effects stochastic frontier models. *Journal of Productivity Analysis*. 2024; 62: 111-118. Available from: https://doi.org/10.1007/s11123-024-00725-3.

[25] Velthoen J, Dombry C, Cai JJ, Engelke S. Gradient boosting for extreme quantile regression. *Extremes*. 2023; 26(4): 639-667. Available from: https://doi.org/10.1007/s10687-023-00473-x.

[26] Sinsomboonthong J, Phaeobang A, Sinsomboonthong S. Adjusted optimal trimmed Theil-Sen method for multiple regression model, with outliers detection and management. *Contemporary Mathematics*. 2025; 6(3): 2939-2967. Available from: https://doi.org/10.37256/cm.6320255546.

[27] Cai J, Horrace WC, Lee Y. Identification and estimation of panel semiparametric conditional heteroskedastic frontiers with dynamic inefficiency. *Econometric Reviews*. 2024; 43(5): 238-268. Available from: https://doi.org/10.1080/07474938.2024.2328905.

[28] Haschka RE. Endogeneity in stochastic frontier models with 'wrong' skewness: copula approach without external instruments. *Statistical Methods & Applications*. 2024; 33(3): 807-826. Available from: https://doi.org/10.1007/s10260-024-00750-4.

[29] Haschka RE, Wied D. Skewness issues in quantifying efficiency: Insights from stochastic frontier panel models based on closed skew normal approximations. *Computational Economics*. 2025; 66: 4381-4416. Available from: https://doi.org/10.1007/s10614-025-10857-9.

[30] Catania L, Luati A. Semiparametric modeling of multiple quantiles. *Journal of Econometrics*. 2023; 237(2): 105365. Available from: https://doi.org/10.1016/j.jeconom.2022.11.002.

[31] Kumbhakar SC, Peresetsky A, Shchetynin Y, Zaytsev A. Technical efficiency and inefficiency: reliability of standard SFA models and a misspecification problem. *Econometrics and Statistics*. 2025; 36: 55-72. Available from: https://doi.org/10.1016/j.ecosta.2021.12.006.

[32] Costa MA, da Silva AV, Mineti LB. Geographical Bayesian second stage analysis for operating efficiency of Brazilian electricity distribution system operators. *Energy Economics*. 2025; 144: 108371. Available from: https://doi.org/10.1016/j.eneco.2025.108371.

[33] Moradi-Motlagh A, Emrouznejad A. The origins and development of statistical approaches in non-parametric frontier models: a survey of the first two decades of scholarly literature (1998-2020). *Annals of Operations Research*. 2022; 318: 713-741. Available from: https://doi.org/10.1007/s10479-022-04659-7.

[34] Schmidt R, Kneib T. Multivariate distributional stochastic frontier models. *Computational Statistics & Data Analysis*. 2023; 187: 107796. Available from: https://doi.org/10.1016/j.csda.2023.107796.

[35] Deng MY, Kutlu L, Mao X. Decision tree-augmented stochastic frontier analysis. *Economics Letters*. 2026; 264: 112943. Available from: https://doi.org/10.1016/j.econlet.2026.112943.

Appendix A. Descriptive statistics of the main raw empirical variables

Table A1 reports descriptive statistics for the main raw empirical variables used in the empirical analysis. Monetary variables are reported in million RUB. Leverage, ROA, sales growth, and log sales growth are reported in percent. Size is measured as the logarithm of total assets.

Table A1. Descriptive statistics of the main raw empirical variables

Variable	Unit	N	Mean	SD	Min	P25	Median	P75	Max
Sales revenue	million RUB	12,320	2,474.12	9,986.13	50.17	359.57	777.92	1,756.69	330,882.00
Fixed assets	million RUB	12,320	540.34	2,586.16	10.01	42.06	103.06	294.70	64,741.80
Wage bill	million RUB	12,320	338.19	2,114.58	5.00	49.47	103.53	244.16	210,450.00
Total assets	million RUB	12,320	2,738.46	11,333.90	23.27	328.10	727.37	1,775.67	448,324.00
Equity/capital	million RUB	12,320	1,297.69	6,005.07	−54,219.60	118.03	309.08	846.06	183,292.00
Total debt	million RUB	12,320	1,440.78	6,841.61	0.88	113.60	303.93	838.82	323,486.00
Net profit	million RUB	12,320	229.86	1,477.79	−58,454.20	9.47	49.84	172.01	50,313.40
Leverage	%	12,320	50.10	29.54	0.31	24.92	49.22	72.45	194.78
Size = log(total assets)	log points	12,320	20.38	1.31	16.96	19.61	20.40	21.30	26.83
ROA	%	12,320	10.41	13.45	−77.10	2.04	8.00	16.97	78.79
Sales growth	%	9,856	9.00	35.62	−96.35	−10.70	4.52	24.12	198.76
Log sales growth	%	9,856	2.88	36.45	−330.96	−11.31	4.42	21.61	109.45

Source: Calculated by the authors based on research materials

Appendix B. Class-specific mean profiles with 95% cluster-bootstrap confidence intervals

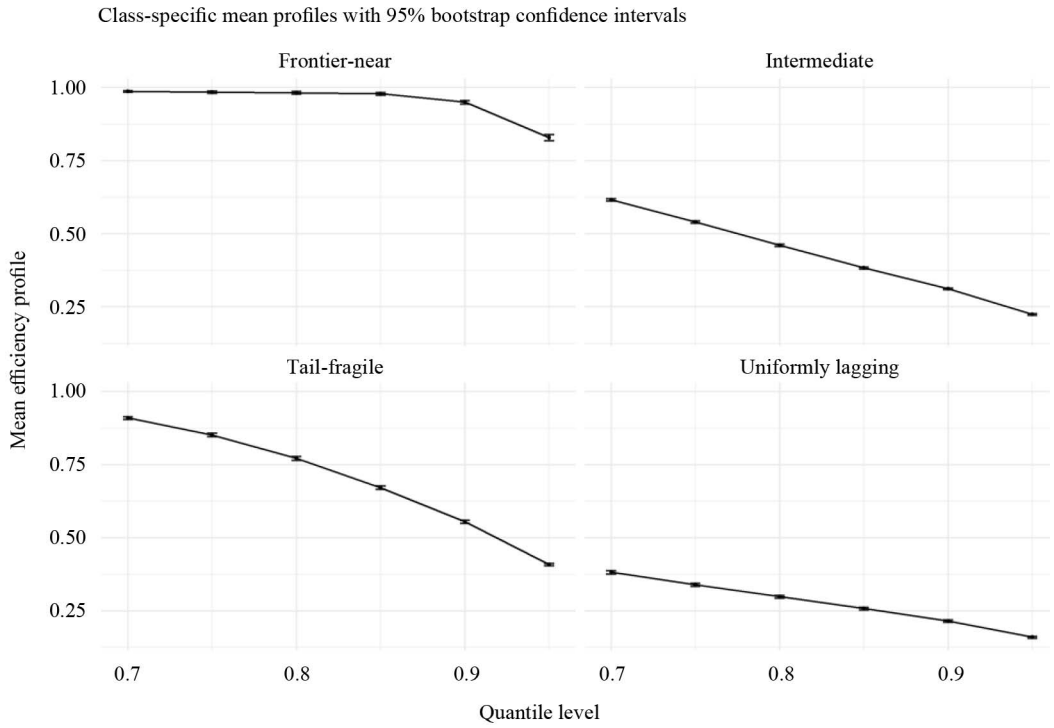


Figure B1. Mean quantile efficiency profile by class with 95% cluster-bootstrap confidence intervals
Source: Calculated by the authors based on research materials

Appendix C. Robustness to second-stage model

Table C1. Robustness to second-stage model specification

Dependent variable	Specification	Leverage	Size	Firm FE	Year FE	Industry-year FE	Observations	R_2	Within R_2
$\widehat{TE}_{it}(0.70)$	Baseline	−0.0010***	0.0775***	Yes	Yes	No	12,320	0.8373	0.0345
	No Size	−0.0007***	—	Yes	Yes	No	12,320	0.8323	0.0047
	Industry-year shocks	−0.0010***	0.0769***	Yes	No	Yes	12,320	0.8385	0.0339
$\widehat{TE}_{it}(0.95)$	Baseline	−0.0007***	0.0867***	Yes	Yes	No	12,320	0.8494	0.0485
	No Size	−0.0004**	—	Yes	Yes	No	12,320	0.8420	0.0014
	Industry-year shocks	−0.0007***	0.0914***	Yes	No	Yes	12,320	0.8558	0.0549
$\widehat{L}_{it}$	Baseline	−0.0008***	0.0837***	Yes	Yes	No	12,320	0.8637	0.0460
	No Size	−0.0005***	—	Yes	Yes	No	12,320	0.8576	0.0031
	Industry-year shocks	−0.0009***	0.0846***	Yes	No	Yes	12,320	0.8648	0.0468
$\widehat{TF}_{it}$	Baseline	−0.0002***	−0.0070*	Yes	Yes	No	12,320	0.6447	0.0029
	No Size	−0.0003***	—	Yes	Yes	No	12,320	0.6445	0.0021
	Industry-year shocks	−0.0002***	−0.0086**	Yes	No	Yes	12,320	0.6642	0.0031

Source: Calculated by the authors based on research materials

Appendix D. Lagged performance regressions

Table D1. Lagged performance regressions

Variable	ROA_w, full	ROA_w, no Size	LogGrowth_w, full	LogGrowth_w, no Size
$\widehat{L}_{i,t-1}$	8.398***	6.702***	−100.2***	−104.5***
	(1.192)	(1.202)	(5.048)	(5.132)
$\widehat{TF}_{i,t-1}$	0.9179	1.597	−5.558	−3.834
	(1.755)	(1.747)	(6.861)	(6.815)
$Leverage_{i,t-1}$	0.1086***	0.0932***	0.3500***	0.3109***
	(0.0146)	(0.0146)	(0.0465)	(0.0449)
$Size_{i,t-1}$	−3.505***	—	−8.902***	—
	(0.5785)	—	(2.103)	—
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Observations	9,856	9,856	9,856	9,856
R^2	0.6346	0.6314	0.3734	0.3704
Within R^2	0.0254	0.0169	0.1199	0.1158

Source: Calculated by the authors based on research materials. Notes: Dependent variables are winsorized ROA and winsorized log sales growth. Clustered firm-level standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$