The Best and the Rest:

Revisiting the Norm of Normality of Individual Performance

This Research Briefing is a service from BC HRMA's research group. Our aim is to make it easier and quicker for HR professionals to find and apply the latest and best people management insight to their challenges and projects. This paper contains a concise and practical summary of a recent academic finding that should shape your HR practices.

Summary

In the field of human resources management, organizational behaviour and industrial and organizational psychology, the assumption that individual performance follows a **Gaussian (normal) distribution** – the form of a bell curve with the majority of performers clustered around the mean – has long influenced organizational practice. The study, *"*The Best and the Rest: Revisiting the Norm of Normality of Individual Performance*"* by Ernest O'Boyle Jr. and Herman Aguinis, presents a new paradigm for understanding why applying a normal distribution, that assumes the majority of individuals will perform in an "average" manner, does not present an accurate picture of the way that individual performance unfolds in an organization. In fact, as some of us may witness in our own organizations, the findings of this research provide evidence that the majority of work is carried out by a small number of people that out-perform the rest. Thus, the "picture of performance" is better represented by a **Paretian distribution**, which resembles the shape of a ski slope, where under-performers are at the upper most starting point of the slope and high performers are at the tip of the end.

Key Definitions

Guasian Distribution - also known as a normal distribution, the Guasian distribtion follows the shape of a bell curve. That is to say that results cluster around a mean and then fan out into symetrical tails on either side. In the case of individual performance, a Guasian distribution would assume that there would be a small number of non-performers and a small number of high performers, while the majority of indivuals would be average performers, clustered around the mean. Performers not conforming to the normal distribution would be considered errors or outliers and excluded from analysis or forced to fit the distribution (e.g., "grading on a curve").

Paretian Distribution - also known as a power law distribution, Paretian distributions have unstable means and a greater portion of extreme events. They resemble a ski slope in shape. In the case of individual performance, the non-performers would be at the top of the slope and the high performers at the bottom. The Paretian distribution allows for more extreme events and has fatter tails than a normal distribution. In the case of individual performance, the majority of performers will fall below the mean.

The typical performer would fall below the mean or average result. That's right, in this model, the typical performer is "below average" (speaking from a statistical perspective).

Background

Normal distributions have long been used as a model for simplifying complex theories related to scientific disciplines (e.g., laws of thermo dynamics,) economics (e.g., stock pricing), human traits (e.g., height or intelligence), among others. Perhaps most evident in the study of performance and appraisal, we can see how both the design of studies and the statistical analysis used to measure performance yield to the assumptions of a normal distribution.

- **Forced Response Formats -** in this design situation, a study may, for example, force a performance rater to rate performance on a scale of 1 to 7, and to "normalize" judgments in their distribution.

- **Deletion or Correction of Outliers -** in statistical analysis, a researcher may either delete or correct outliers, based on the assumption that performance should fit a normal distribution and outliers are either the result of error or bias (e.g., the performance rater may be either too tough or too soft) or an anomaly (e.g., a highly successful sales person who simply has a natural talent in the sales business and is not a fair comparator for the average individual).

Perhaps this assumption of a normal distribution may be partially related to the lack of other models for understanding and / or predicting individual performance on the job, or lack of research negating the normal distribution. Regardless, whether individuals consciously believe that the workforce is primarily made up of "average" performers (performers that cluster around the mean in a normal distribution) or not, from recruitment to performance management to leadership development, the assumption that the majority of the workforce will be made up of average performers has had a tremendous impact on the decisions and practices that organizations make and apply on a daily basis (see sidebar).

Assumptions of the Normal Distribution of Performance

How does research based on the assumption of the normal distrubiton of individual performance translate into organizational practice? Here are a few examples:

Distribution of Work / Work Objectives - Work objectives (e.g., production targets) set according to the "average" production time, where the calculation of average excludes outliers (e.g., people perceived to vastly over or under perform) in order to maintain the normal distribution assumption.

Performance Appraisal - Appraisal systems which use rating systems and set limits on the distribution of ratings; thus enforcing a normal distribution.

Training & Development - Training initiatives that focuses on improving the productivity of the "average worker" or how to improve the performance of the majority (e.g., the team).

Compensation - Salary grades based on the assumption of a normal distribution of performance for individual performers in a job group or category which place the average performer at the mid point and distribute pay within the limits of a normal distribution.

This study provides an alternate model for considering the distribution of individual performance – that of a Paretian distribution, also known as a power law distribution. Paretian distributions have unstable means and a greater portion of extreme events. They resemble a ski slope in shape. In the case of individual performance, the non-performers would be at the top of the slope and the high performers at the bottom. In this type of distribution, most performers will typically fall below the mean.
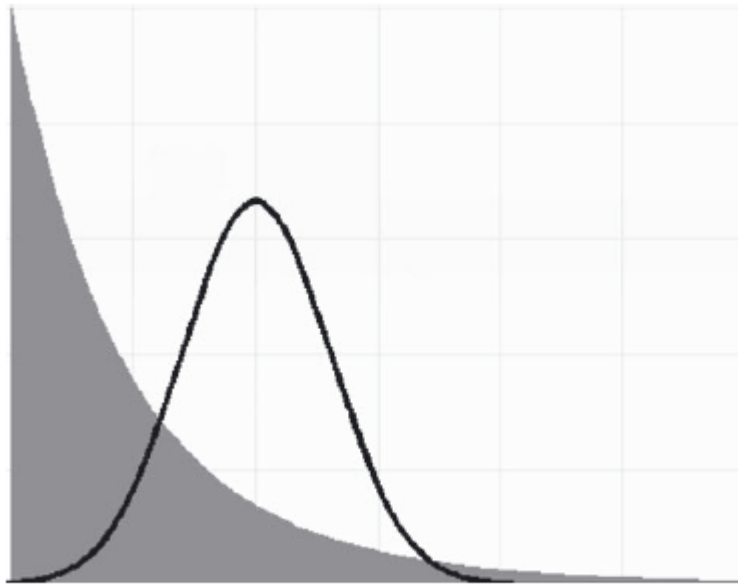

Figure 1: A Normal Distribution (Black) Overlaying a Paretian Distribution (Grey).

> This diagram shows the difference between a Gaussian (normal) distribution and a Paretian (power) distribution.
>
> The primary difference is that in a Paretian distribution the majority of the sample is below average.

Hypotheses and Results

The researchers posited the following hypothesis:

- The distribution of performance is better modeled with a Paretian curve than a Gaussian (normal) curve.

The hypothesis was applied to five different studies, in which a Chi-square analysis was used to force the data to "fit" a Gaussian (normal) curve and a Paretian curve. The studies included 198 samples and 633,263 researchers, entertainers, politicians, and amateur and professional athletes.

The results, based on the five separate studies, indicated that in over 93% of cases a Paretian distribution more closely represented the distribution of performance than a Gaussian distribution. That is to say, that most performance was generated by a small group of "superstars"; in other words, it's no coincidence that a large number of your favourite entertainment awards or athletic awards go to a small group of people.

In accordance with a Paretian distribution, the "picture of performance" resembles a ski slope. The typical performer would fall below the mean or average result. So there would be approximately 80% below average, 10% around the middle, and 10% exceeding; meaning that the assumption that the typical performer is average is a myth of the normal distribution assumption. The results lend themselves to the "80 / 20 rule", which assumes at 80% of the work is done by 20% of the people.

The researchers identified three main areas for examination based on their findings.

1. Performance Measurement & Management

The assumption of a normal distribution means that the majority of performance management systems are based on the mean or "average" worker (e.g., the performance of this individual should represent performance at the middle of the performance spectrum). Thus, if most people were to be rated below this level of performance, the result would be determined to be an anomaly and excluded, or either the rater or the instrument (or both) would be determined to be at fault. Not so, according to the results of the study. Following a Paretian distribution, the majority of performers would reside in the lower end of the performance spectrum.

To shift to a Paretian perspective, future research and innovation in performance measurement would benefit from the development of instruments that allow for the identification of top performers that make the greatest contributions to organizational results.

2. Utility Analysis

Utility analysis is built on the assumption of normality and specifically the standard deviation of performance. If we assume that the distribution of performance is normal, then we assume that performance increases at a steady and predictable rate. For example, if performance was measured in dollar values, and each standard deviation was equivalent to $20, the performance above one standard deviation from the mean would increase by $20, at two standard deviations from the mean it would increase to $40, etc. A Paretian distribution says that this is not the case, that the reality is that the performance of those individuals above the mean increases at a rate more rapid and extreme than this.

Many leadership and training and development practices that are grounded in the Gaussian distribution have focused on shifting the performance of the majority of workers (those clustered around the mean), rather than the workers that do the majority of the work as displayed in a Paretian distribution. Practices that fail to influence those top performers that generate the majority of performance are likely to result in minimal impacts in the organization. Despite small numbers, the impact of increasing the performance of individuals at the extreme positive end of performance will result in a much greater impact to the organization.

Further, additional research needs to be conducted to understand how the performance of top performers interact with and are impacted by the performance of the rest of the team in a variety of team composition scenarios (e.g., primarily high performers on a team, a single high performer, etc.), and how the rest of the team interacts with and impacts the performance of top performers.

3.   Prediction of Performance

Basing prediction-of-performance research on the Gaussian (normal) distribution has resulted in limited ability to gain insights into the antecedents of strong performance. Perhaps this is due to the failure of a normal curve to represent trends in individual performance and, therefore, performance improvements due to interventions. The researchers suggest that future investigations should first determine if a normal distribution of a Paretian distribution best fits the data set (e.g., using a Chi-square analysis as in this research). After this determination, analytical techniques in line with the best fit model should be selected.

Implications for HR

Given the findings of this research, do organizations now shift their focus away from the majority of the current and potential workforce and to the superstars that have the greatest organizational impact? The findings of the study further bring into conflict traditional perceptions of fairness and equality in the workplace and the ongoing drive to increase individual, and therefore organizational, performance. The pay-for-performance debate demonstrates the challenges in shifting away from Gaussian assumptions and toward Paretian assumptions.

From the findings of this research we can examine a number of implications for HR practice as a whole. From a high level perspective, consider the following:

Recruitment & Retention

The findings of this research do not mean that you should only employ top performers. From an availability standpoint and a total rewards perspective this is simply not realistic. What the research does imply is that you should have a strong understanding of the organizational structure and the critical / high impact roles in your organization so that you can selectively source superstars for key positions where they will make the most impact to the bottom line. An understanding of your organizational structure will further help you to understand where a solid second string player will be sufficient. Through this understanding you will facilitate your organization to manage the retention demands (e.g., compensation, perquisites, flexibility) that will support you to keep that critical group of top performers. If you need an analogy for selection and retention, consider building your sports team under the salary cap.

Performance Management

The findings of this research further speak to the need to know who your top performers are so that you can appropriately develop and take action to retain them. Evaluate your performance management systems and make adjustments to ensure that the systems don't impose the traditional bell curve (e.g., ask raters to distribute performers in this manner). Allow the performance system to showcase top performers where they exist so that efforts to develop them and retain them can be targeted. Ensure that managers and employees are trained to use the system thus limiting bias, and support the success of the system with clear definitions of ratings (in systems that are ratings based).

For HR professionals, growing evidence that the majority of the workforce is, in fact, performing below the statistical mean may be alarming. In a Paretian curve, typical no longer means average, and this has implications for the way that organizations make decisions and the HR practices that they employ. It does not mean that the overall workforce is not vital. All members of your organization are valuable, but the results of the research provide further evidence to suggest that the way that we design and focus our HR efforts should be examined.

Performance Appraisals: Gaussian vs. Paretian

Consider a performance management system that uses a five point rating scale. Assuming a Guassian (normal) distribution of individual performance, the mid-point of the scale, a rating of three, would be considered an average rating. While the ideal for the organization would be for performance to be above the average, the expectation would be that most performers in an organization would be clustered around the "three" rating, with fewer and fewer fanning out to the more extreme scores.

If one assumes a Paretian distribution of individual performance then the ratings and their descriptors need to shift. The expectation would be that the typical performer exhibits performance that is below the mean. Acceptable performance would be rated as a two (below the mid-point of the scale). Points three, four and five would demonstrate the performance that individuals can strive for.

This article is based on the following research paper:

"The Best & the Rest: Revisiting the Norm of Normality of Individual Performance"
Ernest O'Boyle Jr., Herman Aguinis,
Publisher: Personnel Psychology, 2012, 65, 79-119