Supervisors and Performance Management Systems

In a recent study, Lazear, Shaw, and Stanton (2015) showed that supervisors of production teams are very influential in determining how their teams perform. They found that the productivity difference from replacing a bad boss with a good one, on average, is equivalent to adding an extra employee to a nine-member team.

Besides managing the day-to-day activities of their team members, supervisors must also evaluate the performance of their employees. Performance management systems govern the flow of information about employees from production teams to upper management. Crucially, performance ratings are meant to summarize the performance of employees[1]. Such performance ratings serve as a basis for determining the eligibility for promotions, bonuses, pay increases, and other important employee outcomes.

The use of performance management systems is sometimes controversial. Supervisor ratings inevitably involve a subjective assessment of employee performance. They therefore inject an element of arbitrariness that employees often intensely dislike. As such, firms may desire to place limits on performance ratings, e.g. by forcing supervisors to grade on a curve. At the same time, the supervisor's managerial ability can affect the performance of the team members, resulting in ratings that may reflect real productivity differences of workers assigned to better and worse bosses. A forced curve would miss these real performance differences. Either way, subjective ratings have the potential to vary widely across supervisors, depending on personality and managerial talents.

In a recent paper, we study the role of supervisors in the performance management system of a large Scandinavian service sector firm. Because the data contain information on who supervises whom as well as performance ratings, we can determine how much these ratings differ across supervisors. Over time, managers supervise many different employees, and employees are supervised by different supervisors. For example, employees who have been with the firm for 11 years have, on average, been evaluated by four different supervisors. This makes it possible to identify systematic differences across supervisors in how they rate employees, while holding constant the quality of the employees.

The results are striking. We find that some supervisors tend to give high ratings regardless of the quality of their incoming team, while others tend to give low ratings. In fact, working for a supervisor who is one standard deviation higher in terms of the tendency to give good ratings increases the probability of receiving a 'commendable' or 'outstanding' rating by 32 percent.

Supervisors are important for the rating employees receive. And so what? Many theoretical papers have addressed the potential biases that may produce the supervisor component in performance ratings (Prendergast, 1999). Maybe these differences occur because some supervisors are nice or lenient and others are not. Or maybe they exist because some supervisors are simply better managers, which is then reflected in performance ratings.

We can empirically distinguish between these stories. We do so by exploiting objective measures of performance (financial and other objective Key Performance Indicators) available at the level of the

[1] For an empirical investigation of performance management systems in practice, see Frederiksen, Lange, and Kriechel (2017).

teams that supervisors manage. If high-rating supervisors elicit high-performance from workers in their teams, their ratings most likely reflect performance rather than leniency. If high-rating supervisors, on the other hand, do not influence performance, then the variation in ratings is more likely to be explained by the fact that some managers are lenient with their ratings and others are not. Our results show that high-rating supervisors tend to supervise teams that perform well; this makes it plausible that high-rating supervisors are also good managers.

The hypothesis that high-rating supervisors in our firm are also the better managers is bolstered by the complementary observation that the subordinates of high-rating supervisors and the high-rating supervisors themselves are paid more. Our reasoning is that firms only compensate workers and supervisors for real productivity – not high ratings that come about because of a generous evaluation.

Empirically, we establish that employees working for high raters receive a higher base pay and higher bonuses; working for a supervisor who tends to give a one standard deviation higher rating leads to 2 percent higher earnings and 7 percent higher bonuses. That same supervisor also has 5 percent higher earnings on average.

All of our results support that supervisors who tend to give high ratings do so because their teams are more productive. The earnings data suggests that subordinates of these high raters share in the extra surplus generated by their teams, which makes it attractive to work for a high-rating supervisor. Additional evidence for this attractiveness can be obtained by analyzing worker mobility and data from employee job satisfaction surveys. We find some evidence that employees working for a high-rating supervisor are less likely to quit the company and tend to stay with their supervisor. Moreover, they rate their supervisors more highly on job satisfaction.

Having established that supervisors are important for the employees' careers, it may be relevant to ask: How important? To address this question fully, we take a career perspective that involves both pay and promotions. Consider a worker who is assigned for one year to a high-rating supervisor placed at the 90th percentile of the supervisor ratings distribution. This worker can expect an increase in the present discounted value of earnings at the firm equivalent to 7-14 percent of an average annual salary compared to a worker whose supervisor is at the 10th percentile. Thus, a one-period assignment to a good supervisor leads to substantial earnings gains over the career.

So what can we conclude? Subjective performance reviews are often criticized for being vulnerable to managerial biases. As a result, companies sometimes try to alleviate these biases by asking supervisors to grade employees on a curve. We urge caution when considering such changes to ratings systems. Because our results show that the grading habits by supervisors plausibly reflect productivity differences, a forced curve would constrain supervisors in their ability to make workers more productive. It might be better to train supervisors in the ratings process, enabling them to provide fair and accurate performance reviews, while still allowing them flexibility in how they use the ratings.

Literature

Frederiksen, Anders, Fabian Lange, and Ben Kriechel, 2017, Subjective Performance Evaluations and Employee Careers, *Journal of Economic Behavior and Organization,* Vol. 134, pp. 408-429.

Frederiksen, Anders, Lisa Kahn, and Fabian Lange, 2017, Supervisors and Performance Management Systems, http://www.cbs.dk/files/cbs.dk/frederiksen_kahn_and_lange_2017_2_1.pdf.

Lazear, Edward, Kathryn Shaw, and Christopher Stanton, 2015, The Value of Bosses, *Journal of Labor Economics,* Vol. 33, No. 1, pp. 823-861.

Prendergast, Canice, 1999, The Provision of Incentives in Firms, *Journal of Economic Literature*, Vol. 37, No. 1, pp. 7-63.

About the Brief

This CCP research brief is written by Anders Frederiksen (Aarhus University), Lisa B. Kahn (Yale School of Management) and Fabian Lange (McGill University) and is based on a presentation at the CCP Spring Meeting 2017. For further information, do not hesitate to contact the authors on afr@btech.au.dk.