# CEO Behavior and Firm Performance

Oriana Bandiera

*London School of Economics*

Andrea Prat

*Columbia University*

Stephen Hansen

*Imperial College Business School*

Raffaella Sadun

*Harvard University*

We develop a new method to measure CEO behavior in large samples via a survey that collects high-frequency, high-dimensional diary data and a machine learning algorithm that estimates behavioral types. Applying this method to 1,114 CEOs in six countries reveals two types:

This project was funded by Columbia Business School, Harvard Business School, and the Kauffman Foundation. We are grateful to our editor, Ali Hortaçsu, and to Morten Bennedsen, Nick Bloom, Robin Burgess, Wouter Dessein, Florian Englmaier, Bob Gibbons, Rebecca Henderson, Ben Hermalin, Paul Ingram, Amit Khandelwal, Renata Lemos, Nicola Limodio, Michael McMahon, Antoinette Schoar, Daniela Scur, Steve Tadelis, John Van Reenen, and seminar participants at Bocconi, Cattolica, Chicago, Columbia, Copenhagen Business School, Cornell, the Centre for Economic Policy Research (CEPR) Economics of Organization Workshop, the CEPR/Institute for Labour Economics (IZA) Labour Economics Symposium, Edinburgh, Harvard Business School, INSEAD, the London School of Economics, Massachusetts Institute of Technology, Munich, the National Bureau of Economic Research, Oxford, Politecnico di Milano, Princeton, Sciences Po, the Society for Institutional and Organizational Economics, Sydney, the Stanford Management Conference, Tel Aviv, Tokyo, Toronto, Uppsala, and Warwick for useful suggestions. Data are provided as supplementary material online.

Electronically published February 19, 2020
[*Journal of Political Economy*, 2020, vol. 128, no. 4]

"leaders," who do multifunction, high-level meetings, and "managers," who do individual meetings with core functions. Firms that hire leaders perform better, and it takes three years for a new CEO to make a difference. Structural estimates indicate that productivity differentials are due to mismatches rather than to leaders being better for all firms.

## I. Introduction

CEOs are at the core of many academic and policy debates. The conventional wisdom, backed by a growing body of empirical evidence (Bertrand and Schoar 2003; Bennedsen et al. 2007; Kaplan, Klebanov, and Sorensen 2012), is that the identity of the CEO matters for firm performance. This raises the question of what CEOs do and how differences in CEO behavior relate to differences in firm performance.

Scholars have approached these questions in two ways. At one end of the spectrum, Mintzberg (1973) and similar studies measure actual behavior by "shadowing" CEOs in real time through personal observation. These exercises produce a rich description of executives' jobs, but they are not amenable to systematic statistical analysis, as they are based on small samples.[1] At the other end of the spectrum, organizational economists have developed abstract categorizations of leadership styles that, however, are difficult to map into empirical proxies of behavior (Hermalin 1998, 2007; Dessein and Santos 2016).[2]

This paper develops a new methodology to scale up the shadowing methods to large samples, thereby combining richness of detail with statistical analysis. This presents two challenges: (1) how to shadow a large number of CEOs and (2) how to aggregate granular information on their activities into a summary measure that has a consistent meaning across subjects.

[1] Mintzberg (1973) shadows five CEOs for a week, and Porter and Nohria (2018) follow 27 CEOs for 3 months. Other authors have shadowed executives below the CEO level (for instance, Kotter [1990] studied 15 general managers). Some consulting companies, such as McKinsey, run surveys where they ask CEOs to report their overall time use, but this is done on the basis of their subjective aggregate long-term recall rather than on a detailed observational study.

[2] Hermalin (1998, 2007) proposes a rational theory of leadership, whereby the leader possesses private, nonverifiable information on the productivity of the venture that she leads. Van den Steen (2010) highlights the importance of shared beliefs in organizations, as these lead to more delegation, less monitoring, higher utility, higher execution effort, faster coordination, fewer influence activities, and more communication. Bolton, Brunnermeier, and Veldkamp (2013) highlight the role of resoluteness: a resolute leader has a strong, stable vision that makes her credible among her followers. This helps align the followers' incentives and generates higher effort and performance. Dessein and Santos (2016) explore the interaction between CEO characteristics, CEO attention allocation, and firm behavior: small differences in managerial expertise may be amplified by optimal attention allocation and result in dramatically different firm behavior.

We address the first challenge by shadowing the CEOs' diaries, rather than the individuals themselves, via daily phone calls with the CEOs or their personal assistants (PAs).[3] This approach allows us to collect comparable data on the behavior of 1,114 CEOs of manufacturing firms in six countries: Brazil, France, Germany, India, the United Kingdom, and the United States. Overall, we collect data on 42,233 activities covering an average of 50 working hours per CEO. In particular, we record the same five features for each activity: its type (e.g., meeting, plant/shop-floor visits, business lunches), its planning horizon, the number of participants involved, the number of different functions, and the participants' function (e.g., finance, marketing, clients, suppliers).

While this approach allows us to scale the data collection to a much larger sample of CEOs relative to earlier studies, this wealth of information is too high-dimensional to be easily compared across CEOs or correlated with other outcomes of interest, such as CEO and firm characteristics. To address this second challenge, we use a machine learning algorithm that projects the many dimensions of observed CEO behavior onto two "pure" behaviors—that is, groups of related activities that together reflect a coherent, underlying behavioral profile. The algorithm finds the combination of features that best differentiates among the sample CEOs. The first of the two pure behaviors is associated with (1) more time spent with employees involved with production activities and (2) one-on-one meetings with firm employees or suppliers. The second pure behavior is associated with more time spent with C-suite executives and in interactions involving several participants and multiple functions from both inside and outside the firm together. To fix ideas, we label the first type of pure behavior "manager" and the second "leader," following the behavioral distinctions described in Kotter (1990).[4] This approach allows us to generate a one-dimensional behavior index that represents each CEO as a convex combination of the two pure behaviors, which we use to study the correlation between CEO behavior and firm performance by merging the behavior index with firm balance-sheet data. We find that leader CEOs are more likely to lead more productive and profitable firms. The correlation is economically and statistically significant: one standard deviation in the CEO behavior index is associated with an increase of 7% in sales, controlling for labor, capital, and other standard firm-level variables.

[3] In earlier work (Bandiera et al. 2018), we used the same data to measure the CEOs' labor supply and assess whether and how it correlates with differences in corporate governance (and in particular whether the firm is led by a family CEO).

[4] In Kotter's work, management comprises primarily of monitoring and implementation tasks. In contrast, leadership aims primarily at the creation of organizational alignment and involves significant investments in interpersonal communication across a broad variety of constituencies.

These findings are consistent with two views. The first is that CEOs simply adapt their behavior to the firm's needs and that more productive firms need leaders. The second is that CEOs differ in their behavior and that this difference affects firm performance. We present three pieces of evidence that cast doubt on the view that the correlation is entirely due to CEOs adjusting their behavior to firm needs. First, while CEO behavior is correlated with firm traits—specifically, leader behavior is more common in larger firms, in multinationals, in listed firms, and in sectors with high R&D intensity and production processes denoted by higher incidence of abstract (rather than routine) tasks—these firm-level differences do not fully account for its correlation with firm performance. Second, firm performance before the appointment of the CEO is not correlated with differences in the CEO behavior index after appointment. Third, firms that hire a leader CEO experience a significant increase in productivity after the CEO appointment, but this emerges gradually over time. These findings cannot be reconciled with the idea that CEO behavior is merely a reflection of differential preappointment trends or firm-level, time-invariant differences in performance.

Taken together, these findings suggest that differences in CEO behavior reflect differences among CEOs, rather than merely firm-level unobserved heterogeneity. However, the association between the CEO behavioral index and firm performance does not necessarily imply that all firms would benefit from hiring a leader CEO. In fact, the performance correlations emerging for the data are consistent with both vertical differentiation among CEOs—that is, that all firms would be better off with a leader CEO—and horizontal differentiation with matching frictions—that is, some firms are better off with leaders and others with managers, but not all firms needing a leader CEO are able to appoint one.

We develop and estimate a simple model of CEO-firm assignment that encompasses both vertical and horizontal differentiation to test which is a better fit for the data. In the model, CEOs and firms have heterogeneous types, and a correct firm-CEO assignment results in better firm performance. The model estimation is consistent with horizontal differentiation of CEOs with matching frictions. In particular, while most firms with managers are as productive as those with leaders, overall the supply of managers outstrips demand, such that 17% of the firms end up with the "wrong" type of CEO. These inefficient assignments are more frequent in lower-income countries (36% vs. 5%). The productivity loss generated by the misallocation of CEOs to firms equals 13% of the labor productivity gap between high- and low-income countries.

Our measure of managerial behavior can be used to address questions at the core of organizational economics for which we have little or no evidence. For example, the coordinating role of entrepreneurs has been of interest to economics since Coase (1937), and Roberts (2004) emphasizes the critical role played by leadership behavior in complementing the organizational design tasks of general managers.[5]

Our results, however, should not be taken as evidence that all CEOs should behave like leaders, for two reasons. First, the evidence indicates that CEOs affect firm performance but that this effect is due to matching: that is, CEO behavior that maximizes performance is firm specific. Second, our data do not allow us to disentangle the effects of behavior—what CEOs do—from other CEO traits that are unobservable to us. For example, it may be that only CEOs with specific personality traits, say charisma or vision, can successfully implement the leadership behavior. If a CEO who does not possess those qualities tried to "play" the leader, firm performance might be even worse than it is when she behaves as a manager, as she may not possess the complementary qualities that make leader behavior effective. In that sense, the paper is consistent with an emerging literature studying CEO personality traits (Malmendier and Tate 2005, 2009; Kaplan, Klebanov, and Sorensen 2012; Kaplan and Sorensen 2016) or self-reported management styles (Mullins and Schoar 2016). We differ from this literature in the object of measure (behavior vs. traits) and in terms of methodology: behavior can be measured using actual diary data, while typically the assessment of personality measures must rely on third-party evaluations, potentially noisy self-reports, or indirect proxies for individual preferences.

The paper is also related to a growing literature documenting the role of management processes on firm performance (Bloom and Van Reenen 2007; Bloom, Sadun, and Van Reenen 2016). The correlation between CEO behavior and firm performance that we uncover is of the same order of magnitude as the correlation with management practices, but, as we show in using a subsample of firms for which we have both CEO time-use and management-practices data, management practices and CEO behavior are independently correlated with firm performance. More recently, the availability of rich longitudinal data on managerial transitions within firms has led to the quantification of heterogeneity in managerial quality and its effect on performance. Lazear, Shaw, and Stanton (2015) and Hoffman and Tadelis (2017), for example, report evidence of significant manager fixed effects within firms, with magnitudes similar to the ones reported in this paper. Unlike these studies, we focus on CEOs rather than middle managers. We share the objective of Lippi and Schivardi (2014) to quantify the output reduction caused by distortions in the allocation of managerial talent.

[5] More recently, Cai and Szeidl (2018) have shown that exogenous shifts in the interactions between an entrepreneur and his/her peers are associated with large increases in firm revenues, productivity, and managerial quality.

The paper is organized as follows. Section II describes the data and the machine learning algorithm. Section III presents the analysis of the relationship between CEO behavior and firm performance, looking, among other things, at whether firm past productivity leads to different types of CEOs being appointed. Section IV examines the extent to which CEO behavior merely proxies for observable or unobservable firm characteristics correlated with performance. Section V interprets the correlation between CEO behavior and firm performance by estimating a simple CEO-firm assignment model encompassing both vertical and horizontal differentiation in CEO behavior. Section VI concludes.

## II. Measuring CEO Behavior

### A. *The Sample*

The sampling frame is a random draw of manufacturing firms from Orbis,[6] in six of the world's 10 largest economies: Brazil, France, Germany, India, the United Kingdom, and the United States. For comparability, we chose to focus on established market economies and opted for a balance between high- and middle-to-low-income countries. We interview the highest-ranking individual who is in charge of the organization, has executive powers, and reports to the board of directors. While titles may differ across countries (e.g., managing director in the United Kingdom), we refer to these individuals as CEOs in what follows.

To maintain comparability of performance data, we restricted the sample to manufacturing firms. We then selected firms with available sales and employment data in the latest accounting year before the survey.[7] This yielded a sample of 6,527 firms in 32 two-digit SIC (Standard Industrial Classification) industries that we randomly assigned to different analysts. Each analyst would then call the companies on the list and seek the CEO's participation. The survey was presented to the CEOs as an opportunity to contribute to a research project on CEO behavior. To improve the quality of the data collected, we also offered CEOs the opportunity

[6] Orbis is an extensive commercial data set produced by Bureau Van Dijk that contains company accounts for more than 200 million companies around the world.

[7] We went from a random sample of 11,500 firms with available employment and sales data to 6,527 eligible ones after screening for firms for which we were able to find CEO contact details and that were still active. We could find CEO contact details for 7,744 firms, and of these, 1,217 later were found to be not eligible. Of these 1,217, 310 could not be contacted to verify eligibility before the project ended. Among this set, 1,009 were located in Brazil, 897 in Germany, 762 in France, 1,429 in India, 1,058 in the United Kingdom, and 1,372 in the United States. The lower number of firms screened in France and Germany is due to the fact that the screening had to be done by native-language research assistants based in Boston, of whom we could only hire one for each country. The sample construction is described in detail in app. A.

to learn about their own time use with a personalized time-use analysis, to be delivered after the data had been collected.[8]

Of the 6,527 firms included in the screened Orbis sample, 1,114 (17%) participated in the survey,[9] of which 282 are in Brazil, 115 in France, 125 in Germany, 356 in India, 87 in the United Kingdom, and 149 in the United States.

Table A.1 (tables A.1, A.2, B.1–B.3, and D.1–D.7 are available online) shows that sample firms have, on average, lower log sales (coefficient = 0.071, standard error = 0.011), but we do not find any significant selection effect on performance variables, such as labor productivity (sales over employees) and return on capital employed (see app. A for details; apps. A–D are available online). Table A.2 shows descriptive statistics on the sample CEOs and their firms. Sample CEOs are 51 years old, on average; nearly all (96%) are male and have a college degree (92%). About half of them have an MBA (Master of Business Administration degree). The average tenure is 10 years, with a standard deviation of 9.55 years.[10] Finally, sample firms are very heterogeneous in size and sales values. Firms have, on average, 1,275 employees and $222 million in sales (respectively 300 and $35 million at the median), but with very large standard deviations (6,498 for employment and $1,526 million for sales).

### B. *The Survey*

To measure CEO behavior, we develop a new survey tool that allows a large team of enumerators to record in a consistent and comparable way all the activities the CEO undertakes in a given day. Data are collected through daily phone calls with the CEO himself (43% of the cases) or with the CEO's PA. We record diaries over a week that we chose on the basis of an arbitrary ordering of firms. Enumerators collected daily information

[8] The report was delivered 2 years after the data collection and included simple summary statistics on time use but no reference to the behavioral classification across "leaders" and "managers" that we discuss below.

[9] This figure is at the higher end of response rates for CEO surveys, which range between 9% and 16% (Graham, Harvey, and Puri 2013). At first, 1,131 CEOs agreed to participate, but 16 dropped out before the end of the data collection week for personal or professional contingencies that limited our ability to reach them by phone. One CEO completed the survey for the whole week but provided incomplete information about the activities (i.e., the number and types of participants were missing from the agenda).

[10] The heterogeneity is mostly due to the distinction between family and professional CEOs, as the former have much longer tenures. In our sample, 57% of the firms are owned by families, 23% by dispersed shareholders, 9% by private individuals, and 7% by private equity. Ownership data are collected in interviews with the CEOs at the end of the survey week and independently checked using several internet sources, information provided on the company website, and supplemental phone interviews. We define a firm as owned by an entity if the entity controls at least 25.01% of the shares; if no single entity owns at least 25.01% of the shares, the firm is labeled as "dispersed shareholders."

on all the activities the CEO planned to undertake that day (in the morning) as well as those actually done (in the evening).[11] On the last day of the data collection, the enumerator interviewed the CEO to validate the activity data (if collected through his PA) and to collect information on the characteristics of the CEO and of the firm. Figure A.1 (figs. A.1, B.1, and D.1–D.3 are available online) shows a screenshot of the survey tool.[12] The survey collects information on all activities lasting longer than 15 minutes in the order they occurred during the day. To avoid under- (over) weighting long (short) activities, we structure the data so that the unit of analysis is a 15-minute time block.

Overall, we collect data on 42,233 activities of different durations, equivalent to 225,721 15-minute blocks, 90% of which cover work activities.[13] The average CEO has 202 15-minute time blocks, adding up to 50 hours per week.

### C. *The Data*

Figure 1*A* shows that the average CEO spends 70% of his time interacting with others (either face to face via meetings or plant visits or "virtually" via phone, videoconferences, or emails). The remaining 30% is allocated to activities that support these interactions, such as travel between meetings and time devoted to preparing for meetings. The fact that CEOs spend such a large fraction of their time interacting with others is consistent with the prior literature. Coase (1937), for example, sees as the main task of the entrepreneur precisely the coordination of internal activities that cannot otherwise be effectively regulated through the price mechanism. The highly interactive role of managers is also prominent in classic studies in management and organizational behavior, such as Drucker (1967) and Mintzberg (1973, 1979).[14]

The richness and comparability of the time-use data allow for a much more detailed description of these interactions, relative to prior studies. We use as primary features of the activities (1) their type (e.g., meeting, lunch), (2) their duration (30 minutes, 1 hour, etc.), (3) whether planned or unplanned, (4) the number of participants, and (5) the functions of participants, divided between employees of the firms, whom we define as "insiders" (finance, marketing, etc.), and nonemployees, or "outsiders" (clients, banks, etc.). Figure 1*B* shows that most of this interactive time is

[11] Of the surveyed CEOs, 70% worked 5 days, 21% worked 6 days, and 9% 7 days. Analysts called the CEO after the weekend to retrieve data on Saturdays and Sundays.

[12] The survey tool can also be found at www.executivetimeuse.org.

[13] The nonwork activities cover personal and family time during business hours.

[14] Mintzberg (1973), e.g., documents that in a sample of five managers, 70%–80% of managerial time is spent communicating.

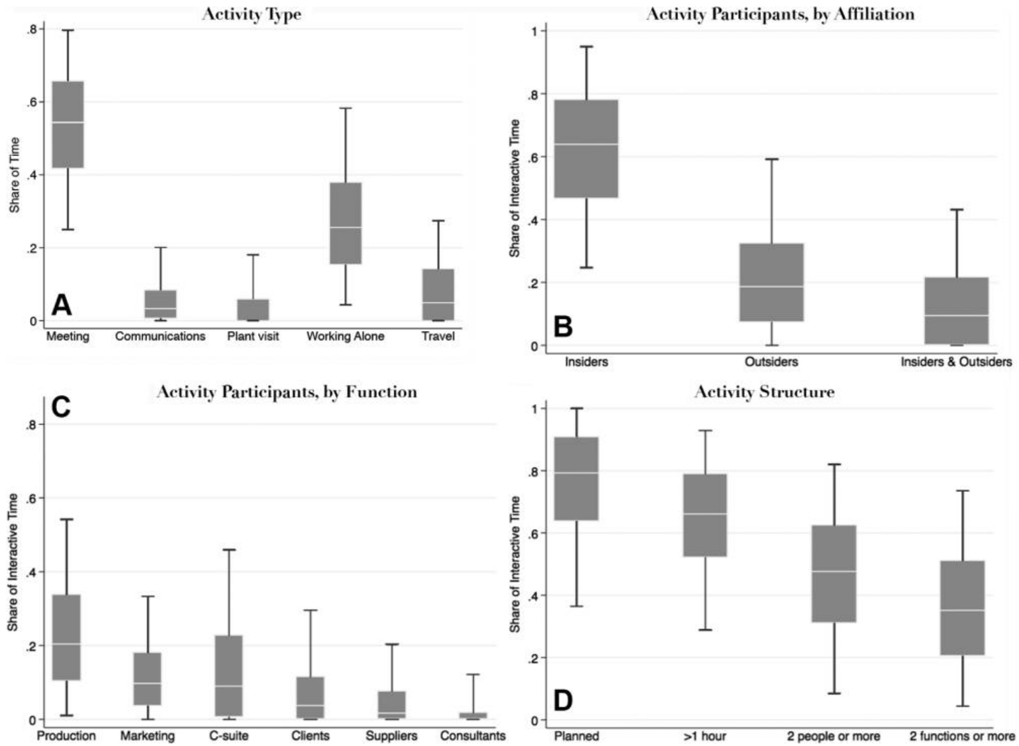


FIG. 1.—CEO behavior: raw data. For each activity feature, the figure plots the median (the line in the box), the interquartile range (the height of the box), and the interdecile range (the vertical line). The summary statistics refer to average shares of time computed at the CEO level. A color version of this figure is available online.

spent with insiders. This suggests that most CEOs chose to direct their attention primarily toward internal constituencies, rather than serving as "ambassadors" for their firms (i.e., connecting with constituencies outside the firm). Few CEOs spend time with insiders and outsiders together, suggesting that, if they do build a bridge between the inside and the outside of the firm, CEOs typically do so alone. Figure 1*C* shows the distribution of time spent with the three most frequent insiders—production, marketing, and C-suite executives—and the three most frequent outsiders—clients, suppliers, and consultants. Figure 1*D* shows that most CEOs engage in planned activities with a duration of longer than 1 hour with a single function. There is no marked average tendency toward meeting with one or more than one person. Another striking aspect of the data shown in figure 1 is the marked heterogeneity underlying these average tendencies. For example, CEOs in the bottom quartile devote just over 40% of their time to meetings, whereas those in the top quartile reach 65%; CEOs in the third quartile devote over three times as much time to production as their counterparts in the first quartile; and the interdecile ranges for time with two people or more and two functions or more are well over 50%. The evidence of such marked differences in behavior across managers is, to our knowledge, a novel and so far underexplored phenomenon.

The data also show systematic patterns of correlation across these distributions, as we show in the heat map of table 1. This exercise reveals significant and intuitive patterns of cooccurrence. For example, CEOs who do more plant visits spend more time with employees working on production and suppliers. The data also show that they tend to meet these functions one at the time, rather than in multifunctional meetings. In contrast, CEOs who do more "virtual" communications engage in fewer plant visits, spend more time with C-suite executives, and interact with large and more diverse groups of individuals. They are also less likely to include purely operational functions (production and marketing—among inside functions—and clients and suppliers—among outsiders) in their interactions. These correlations are consistent with the idea that CEO time use reflects latent styles of managerial behavior, which we investigate in more detail in the next section.

The activities also appear to largely reflect conscious planning versus mere reactions to external contingencies. To assess this point, we asked whether each activity was undertaken in response to an emergency: only 4% of CEOs' time was devoted to activities that were defined as emergencies. Furthermore, we compared the planned schedule of the manager (elicited in the morning conversation) with the actual agenda (elicited in the evening conversation). This comparison shows that CEOs typically undertake all the activities scheduled for a given day—overall, just under 10% of planned activities were canceled.

## TABLE 1
CEO Behavior: Correlations in the Raw Data

| | Meeting | Plant visit | Communications | Planned | >1 Participant | >1 Function | Insiders | Outsiders | Insiders and Outsiders | C-Suite | Production | Marketing | Clients | Suppliers | Consultants |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Meeting | 1 | | | | | | | | | | | | | | |
| Plant visit | **−.0917** | 1 | | | | | | | | | | | | | |
| Communications | **−.337** | **−.1176** | 1 | | | | | | | | | | | | |
| Planned | **.1536** | **−.1009** | **−.1343** | 1 | | | | | | | | | | | |
| >1 participant | .0223 | −.0001 | **−.1317** | **.2862** | 1 | | | | | | | | | | |
| >1 function | **.0546** | **−.1688** | **−.0872** | **.2023** | **.5094** | 1 | | | | | | | | | |
| Insiders | **−.0628** | .0439 | **.0524** | **−.0989** | .0275 | −.0028 | 1 | | | | | | | | |
| Outsiders | **.0521** | −.0433 | −.0189 | .032 | **−.1855** | **−.4087** | **−.7116** | 1 | | | | | | | |
| Insiders and outsiders | **.0693** | **−.0798** | **−.0565** | **.1133** | **.211** | **.5439** | **−.4865** | **−.2236** | 1 | | | | | | |
| C-suite | **−.0533** | **−.1445** | **.1397** | **.1136** | **.1501** | **.1359** | **.3503** | **−.3265** | **−.052** | 1 | | | | | |
| Production | **−.0618** | **.3991** | **−.1048** | **−.1181** | .0219 | **−.1413** | **.3434** | **−.2937** | **−.1106** | **−.3045** | 1 | | | | |
| Marketing | .0049 | **−.1437** | .0233 | −.0244 | .0109 | **.1649** | **.1911** | **−.2699** | .0779 | **−.1892** | **−.1463** | 1 | | | |
| Clients | −.0409 | .0064 | **.1298** | .0124 | **−.1732** | **−.1403** | **−.4192** | **.4271** | **.0723** | **−.1796** | **−.1351** | −.0462 | 1 | | |
| Suppliers | **.0978** | **.1354** | **−.0539** | **−.0391** | **−.1718** | **−.1716** | **−.3295** | **.3488** | .0379 | **−.2199** | .0206 | **−.073** | .044 | 1 | |
| Consultants | **.1164** | −.0214 | **−.0544** | −.0188 | **−.0826** | **−.0258** | **−.2387** | **.2151** | **.0928** | −.0348 | **−.1436** | **−.0751** | **−.0609** | −.0088 | 1 |

Note.—Each cell reports the correlation coefficient between the variables listed in the row and those in the column. Each variable indicates the share of time spent by CEOs in activities denoted by the specific feature (these are the same data used to generate fig. 1). Boldface indicates that the null hypothesis of no correlation is rejected at $p \leq .10$; in other cells, the null cannot be rejected.

### D. The CEO Behavior Index

While the richness of the diary data allows us to describe CEO behavior in great detail, it makes standard econometric analysis unfeasible, because we have 4,253 unique activities (defined as a combination of the five distinct features measured in the data) and 1,114 CEOs in our sample.

To address this, we exploit the idea—based on the patterns of co-occurrence in time use shown in table 1—that the high-dimensional raw activity data are generated by a low-dimensional set of latent managerial behaviors. The next section discusses how we construct a scalar CEO behavior index employing a widely used machine learning algorithm.

#### 1. Methodology

To reduce the dimensionality of the data, we use latent Dirichlet allocation (LDA; Blei, Ng, and Jordan 2003), a hierarchical Bayesian factor model for discrete data.[15] Simpler techniques, such as principal components analysis (PCA, an eigenvalue decomposition of the variance-covariance matrix) or *k*-means clustering (which computes cluster centroids with the smallest squared distance from the observations), are also possible and indeed produce similar results, as we discuss below. The advantage of LDA relative to these other methods is that it is a generative model that provides a complete probabilistic description of time-use patterns.[16] LDA posits that the actual behavior of each CEO is a mixture of a small number of "pure" CEO behaviors and that the creation of each activity is attributable to one of these pure behaviors. Another advantage of LDA is that it naturally handles high-dimensional feature spaces, so we can admit correlations among all combinations of the five distinct features, which are potentially significantly more complex than the correlations between individual feature categories described in table 1. While LDA and its extensions are most widely applied to text data, where it forms the basis of much of probabilistic topic modeling, close variants have been applied to survey data in various contexts (Erosheva, Fienberg, and Joutard 2007; Gross and Manrique-Vallier 2014). Ours is the first application to survey data in the economics literature that we are aware of.

[15] LDA is an unsupervised learning algorithm and uncovers hidden structure in time use without necessarily linking it to performance. This allows us to first describe the most prominent distinctions among CEOs while staying agnostic on whether time use is related to performance in a systematic way. A supervised algorithm would instead "force" the time-use data to explain performance. Moreover, popular penalized regression models such as LASSO (least absolute shrinkage and selection operator) can be fragile in the presence of highly correlated covariates, which makes projecting them onto a latent space before regression analysis attractive.

[16] Tipping and Bishop (1999) have shown that one can provide probabilistic foundations for PCA via a Gaussian factor model with a spherical covariance matrix in the limit case where the variance approaches zero. Clearly, however, our survey data are not Gaussian, so PCA lacks an obvious statistical interpretation in our context.

To be more concrete, suppose that all CEOs have $A$ possible ways of organizing each unit of their time, which we define for short activities, and let $x_a$ be a particular activity. Let $X \equiv \{x_1, \ldots, x_A\}$ be the set of activities. A *pure behavior k* is a probability distribution $\boldsymbol{\beta}^k$ over $X$ that is common to all CEOs.[17]

We begin with the simplest possible case, in which there exist only two possible pure behaviors: $\boldsymbol{\beta}^0$ and $\boldsymbol{\beta}^1$. In this simple case, the *behavior* of CEO $i$ is given by a mixture of the two pure behaviors according to the weight $\theta_i \in [0, 1]$; thus, the probability that CEO $i$ generates activity $a$ can lie anywhere between $\beta_a^0$ and $\beta_a^1$.[18] We refer to the weight $\theta_i$ as the *behavior index* of CEO $i$.

Figure 2 illustrates the LDA procedure. For each activity of CEO $i$, one of the two pure behaviors is drawn independently, given $\theta_i$. Then, given the pure behavior, an activity is drawn according to its associated distribution (either $\boldsymbol{\beta}^0$ or $\boldsymbol{\beta}^1$). So the probability that CEO $i$ assigns to activity $x_a$ is $\chi_a^i \equiv (1 - \theta_i)\beta_a^0 + \theta_i\beta_a^1$.

If we let $n_{i,a}$ be the number of times activity $a$ appears in the time use of CEO $i$, then by independence the likelihood function for the model is simply $\prod_i \prod_a \chi_i^{n_{i,a}}$.[19] While in principle one can attempt to estimate $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ via direct maximum likelihood or the EM (expectation-maximization) algorithm, in practice the model is intractable because of the large number of parameters that must be estimated (which grows linearly in the number of observations). LDA overcomes this challenge by adopting a Bayesian approach and placing Dirichlet priors on the $\boldsymbol{\beta}$ and $\theta_i$ terms. For estimating posteriors we follow the Markov chain Monte Carlo approach of Griffiths and Steyvers (2004).[20] Here we discuss the estimated object of interest, which are the two estimated pure behaviors $\hat{\boldsymbol{\beta}}^0$ and $\hat{\boldsymbol{\beta}}^1$, as well as the estimated behavioral indices $\hat{\theta}_i$ for every CEO $i = 1, \ldots, N$.

Intuitively, LDA identifies pure behaviors by finding patterns of co-occurrence among activities across CEOs, so infrequently occurring activities are not informative. For this reason we drop activities in fewer

[17] Importantly, the model allows for arbitrary covariance patterns among features of different activities. For example, one behavior may be characterized by large meetings whenever the finance function is involved but small meetings whenever marketing is involved.

[18] In contrast, in a traditional clustering model, each CEO would be associated with one of the two pure behaviors, which corresponds to restricting $\theta_i \in \{0, 1\}$.

[19] While a behavior defines a distribution over activities with correlations among individual features (planning, duration, etc.), each separate activity in a CEO's diary is drawn independently, given pure behaviors and $\theta_i$. The independence assumption of time blocks within a CEO is appropriate for our purpose to understand overall patterns of CEO behavior rather than issues such as the evolution of behavior over time or other more complex dependencies. These are, of course, interesting but outside the scope of the paper.

[20] We set a uniform prior on $\theta_i$—i.e., a symmetric Dirichlet with hyperparameter 1—and a symmetric Dirichlet with hyperparameter 0.1 on $\boldsymbol{\beta}^k$. This choice of hyperparameter promotes sparsity in the pure behaviors. Source code for implementation is available from https://github.com/sekhansen.

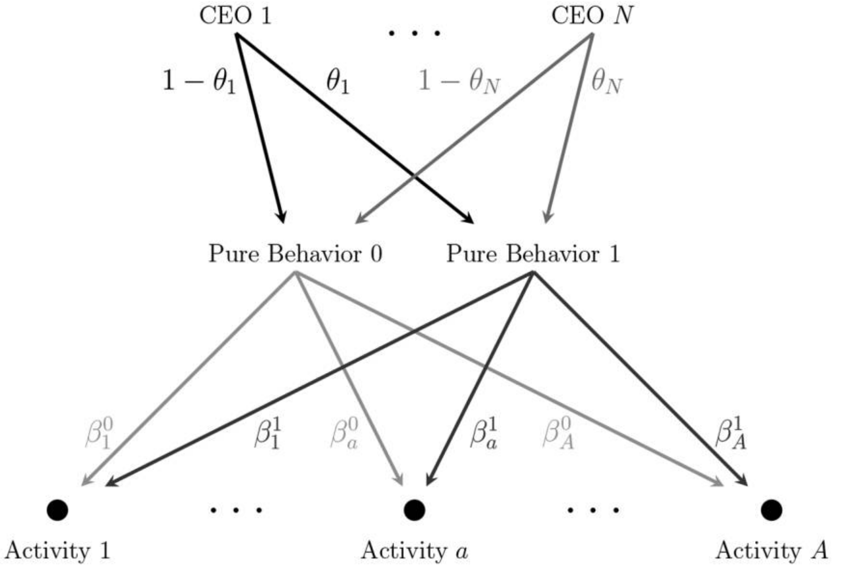


FIG. 2.—Data-generating process for activities with two pure behaviors. This figure provides a graphical representation of the data-generating process for the time-use data. First, CEO $i$ chooses—independently for each individual unit of his time—one of the two pure behaviors according to a Bernoulli distribution with parameter $\theta_i$. The observed activity for a unit of time is then drawn from the distribution over activities that the pure behavior defines. A color version of this figure is available online.

than 30 CEOs' diaries, which leaves 654 unique activities and 98,347 time blocks—or 78% of interactive time—in our baseline empirical exercise. In the appendix, we alternatively drop activities in fewer than 15 and 45 CEOs' diaries and find little effect on the main results (see table D.2).

### 2. Estimates

To illustrate differences in estimated pure behaviors, in figure 3 we order the elements of $X$ according to their estimated probability in $\hat{\boldsymbol{\beta}}^0$ and then plot the estimated probabilities of each element of $X$ in both behaviors. The figure shows that the combinations that are most likely in pure behavior 0 have low probability in pure behavior 1, and vice versa. Tables B.1 and B.2 list the five most common activities in each of the two behaviors.[21] To construct a formal test of whether the observed differences between pure behaviors are consistent with a model in which there

[21] Table B.3 displays the estimated average time that CEOs spend on the different categories in fig. 1, derived from the estimated pure behaviors and CEO behavioral indices. Reassuringly, there is a tight relationship between the shares in the raw data and the estimated shares.

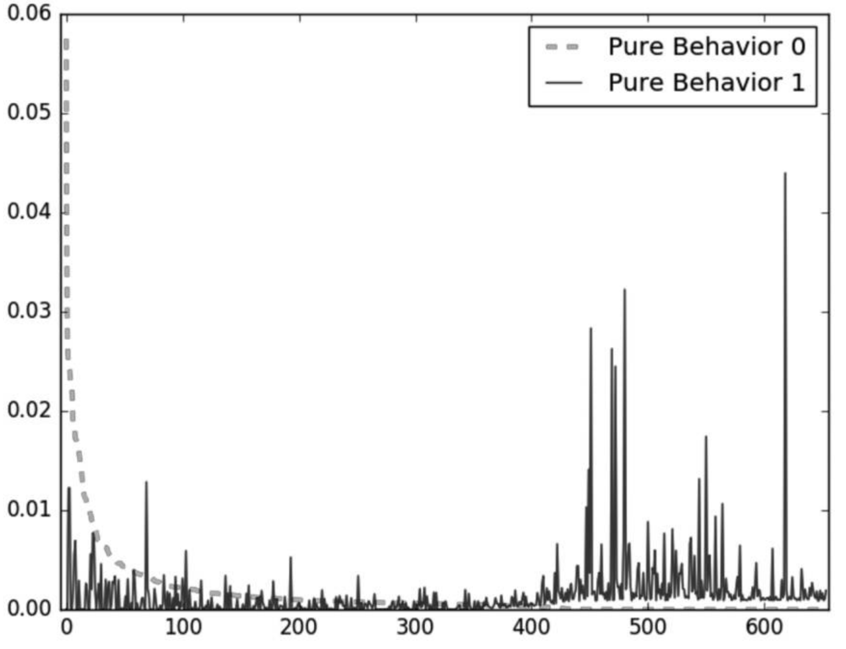


FIG. 3.—Probabilities of activities in estimated pure behaviors. The dashed line plots the estimated probabilities of different activities in pure behavior 0, and the solid line plots the estimated probabilities of different activities in pure behavior 1. The 654 different activities are ordered from left to right in descending order of their estimated probability in pure behavior 0. A color version of this figure is available online.

is only one pure behavior (i.e., a model with no systematic heterogeneity), we simulate data by drawing an activity for each time block in the data from a probability vector that matches the raw empirical frequency of activities. We then use these simulated data to estimate the LDA model with two pure behaviors as in our baseline analysis, and we find systematically less difference between pure behaviors than in our actual data (for further discussion, see app. B).

The two pure behaviors we estimate represent extremes. As discussed above, individual CEOs generate activities according to the behavioral index $\theta_i$ that gives the probability that any specific activity is drawn from pure behavior 1. Figure 4 plots both the frequency and cumulative distributions of the $\hat{\theta}_i$ (which we define as the “CEO behavior index”) estimates across CEOs. Many CEOs are estimated to be mainly associated with one pure behavior: 316 have a behavioral index less than 0.05, and 94 have an index greater than 0.95. As figure 4 shows, however, the bulk of CEOs lie away from these extremes, where the distribution of the index is essentially uniform. The mean of the index is 0.36 (standard deviation = 0.34).

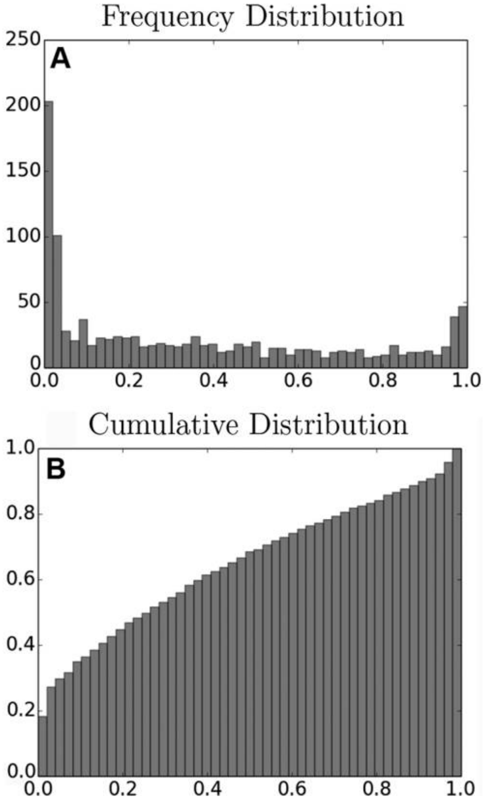


FIG. 4.—CEO behavior and index distribution. A, Number of CEOs with behavioral indices in each of 50 bins that divide the space [0,1] evenly. *B*, Cumulative percentage of CEOs with behavioral indices lying in these bins. A color version of this figure is available online.

Country and industry fixed effects together account for 17% of the variance in the CEO behavior index. This is due primarily to the fact that the CEO behavior index varies by country, and in particular it is significantly higher in rich countries (France, Germany, the United Kingdom, and the United States) than in low- and middle-income countries (Brazil and India). In contrast, industry fixed effects are largely insignificant.[22]

### 3. Results Using Alternative Dimensionality-Reduction Techniques

A question of interest is whether the CEO behavior index built using LDA could be reproduced using more familiar dimensionality-reduction

[22] See fig. D.1 and app. D.1 for more details.

techniques. To investigate this point, we examined the sensitivity of the classification to PCA and *k*-means analysis. For this analysis, we do not use the same 654-dimensional feature vector as for LDA, but rather six marginal distributions computed on the raw time-use data that capture the same distinctions that LDA reveals as important. For each CEO, we counted the number of engagements that (1) last longer than 1 hour, (2) are planned, (3) involve two or more people, (4) involve outsiders alone, (5) involve high-level inside functions, and (6) involve more than one function. The first principal component in the PCA explains 35% of the variance in this feature space and places a positive weight on five of these dimensions (all except 4). Meanwhile, *k*-means clustering produces one centroid with higher values on all dimensions except 4 (and, ipso facto, a second centroid with a higher value for dimension 4 and lower values for all others). Hence, the patterns identified using simpler methods validate the key differences from LDA with two pure behaviors. Note that LDA is still a necessary first step in this analysis because it allows us to identify the important marginals along which CEOs vary. We have also experimented with PCA and *k*-means on the 654-dimensional feature space over which we estimate the LDA model, but the results are much harder to interpret than the ones described above.

## 4. Interpretation of the CEO Behavior Index: Leaders and Managers

We now turn to analyzing the underlying heterogeneity between pure behaviors that generate differences among CEOs, which is ultimately the main interest of the LDA model. To do so, we compute marginal distributions over each relevant activity feature from both pure behaviors. Table 2 displays the ratios of these marginal distributions (always expressed as the ratio of the probability for pure behavior 1 to that for pure behavior 0, for simplicity) for the activities that are more different across the two pure behaviors. A value of one indicates that each pure behavior generates the category with the same probability, a value below one indicates that pure behavior 1 is less likely to generate the category, and a value above one indicates that pure behavior 1 is more likely to generate the category.

Overall, the differences in the CEO behavior index indicate a wide heterogeneity in the way CEOs interact with others: pure behavior 0 assigns a greater probability to activities involving one individual at a time and activities (plant visits) and functions (production and suppliers) that are most related to operational activities. In contrast, pure behavior 1 places higher probabilities on activities that bring several individuals together, mostly at the top of the hierarchy (other C-suite executives), and from a

TABLE 2
MOST IMPORTANT BEHAVIORAL DISTINCTIONS IN CEO TIME-USE DATA

| Feature | Times Less/More Likely |
|---|---|
| Less likely in behavior 1: | |
| Plant visits | .11 |
| Just outsiders | .58 |
| Production | .46 |
| Suppliers | .32 |
| More likely in behavior 1: | |
| Communications | 1.90 |
| Outsiders and insiders | 1.90 |
| C-suite | 33.90 |
| Multifunction | 1.49 |

NOTE.—We generate the values in the table in two steps. First, we create marginal distributions over individual features in activities for each pure behavior. Then, we report the probability of the categories within features in behavior 1 over the probability in behavior 0 for the categories for which this ratio is largest.

variety of functions.[23] Higher values of the CEO behavior index $\hat{\theta}_i$ thus correspond to a greater intensity of these latter types of interactions.

While the labeling of the two pure behaviors is arbitrary, the distinctions between pure behavior 0 and pure behavior 1 map onto behavioral classifications that have been observed in the past by management scholars. In particular, the differences between the two pure behaviors are related to the behavioral distinction between "management" and "leadership" emphasized by Kotter (1990). This defines management primarily as monitoring and implementation tasks, entailing the creation of systems to enable the precise and efficient execution of plans. In contrast, leadership is needed to create organizational alignment and requires significant investment in communication across a broad variety of constituencies.[24]

Hereafter we refer to CEOs with higher values of the behavioral index as *leaders* and those with lower values as *managers*. In the next section, we investigate whether differences in the behavioral index—which are built exclusively on the basis of the CEO time-use data—correlate with firm

[23] We have constructed simulated standard errors for the differences in probabilities of each feature reported in the figure, based on draws from the Markov chains used to estimate the reported means. All differences are highly significant except time spent with insiders, as we discuss in the appendix.

[24] More specifically, leadership is "more of a communication problem. It involves getting a large number of people, inside and outside the company, first to believe in an alternative future—and then to take initiative based on that shared vision" (*Harvard Business Review* 2011). "Aligning invariably involves talking to many more individuals than organizing does. The target population can involve not only a manager's subordinates but also bosses, peers, staff in other parts of the organization" (Kotter 1990, 107).

performance, and provide a simple framework to assess the possible reasons behind the correlation.

## III. CEO Behavior and Firm Performance

To investigate whether the index of CEO behavior is correlated with performance, we match our CEO behavior data with accounting information extracted from Orbis. We were able to gather at least one year of sales and employment data in the period in which the CEOs were in office for 920 of the 1,114 firms in the CEO sample.[25]

### *A. Correlations with the Unidimensional Index*

#### 1. Productivity

We start by analyzing whether CEO behavior correlates with productivity, a key metric of firm performance (Syverson 2011). We begin with the simplest unidimensional measure of CEO behavior and follow a simple production-function approach that yields a regression of the form

$$y_{ifts} = \alpha\hat{\theta}_i + \delta^{E} e_{ft} + \delta^{K} k_{ft} + \delta^{M} m_{ft} + \zeta_t + \eta_s + \varepsilon_{ifts}, \tag{1}$$

where $y_{ifts}$ is the log sales (in constant 2010 US dollars) of firm $f$, led by CEO $i$, in period $t$ and sector $s$; $\hat{\theta}_i$ is the behavior index of CEO $i$; $e_{ft}$, $k_{ft}$, and $m_{ft}$ denote, respectively, the natural logarithms of the number of firm employees and, when available, capital and materials; and $\zeta_t$ and $\eta_s$ are period and three-digit SIC sector fixed effects, respectively.

The performance data includes up to the three most recent years of accounting data predating the survey, conditional on the CEO being in office.[26] To smooth out short-run fluctuations and reduce measurement error in performance, inputs and outputs are averaged across the cross sections of data included in the sample. The results are very similar when we use yearly data and cluster the standard errors by firm (table D.2, col. 2). We include country and year dummies throughout, as well as a set of interview noise controls.[27] The coefficient of interest is $\alpha$, which measures

[25] Of the 1,114 firms, 41 did not report sales and employment information; 64 were dropped when extreme values were removed from the productivity data; 89 had data only for years in which the CEO was not in office, or in office for less than one year, or not in any of the three years before the survey.

[26] We do not condition on the CEO being in office for at least 3 years to avoid introducing biases related to the duration of the CEO tenure; i.e., we include companies that have at least one year of data. We have three years of accounting for 58% of the sample, two years for 24%, and one year for the rest of firms.

[27] These are a full set of dummies to denote the week in the year in which the data were collected, a reliability score assigned by the interviewer at the end of the survey week, a dummy taking value one if the data were collected through the PA of the CEO rather than

the correlation between log sales and the CEO behavior index. Recall that higher values of the index imply a closer similarity to the pure behavior labeled as "leader."

Column 1 of table 3 shows the estimates of equation (1), controlling for firm size, country, year, and industry fixed effects, and noise controls. Since most countries in our sample report at least sales and number of employees, we can include in this labor productivity regression a subsample of 920 firms. The estimate of $\alpha$ is positive (coefficient = 0.343, standard error = 0.108), and we can reject the null of zero correlation between firm labor productivity and the CEO behavior index at the 1% level.

Column 2 adds capital, which is available for a smaller sample of firms (618). The coefficient of the CEO behavior index remains of similar magnitude (coefficient = 0.227, standard error = 0.111) and is significant at the 5% level in the subsample. A 1-standard-deviation change in the CEO behavior index is associated with a 7% change in sales—as a comparison, this is about 10% of the effect of a 1-standard-deviation increase in capital on sales.[28] In column 3 we add materials, which further restricts the sample to 448 firms. In this smaller sample, the coefficients on capital and materials have the expected magnitude and are precisely estimated. Nevertheless, the coefficient on the CEO behavior index retains a similar magnitude and significance. Column 4 restricts the sample to firms that, in addition to having data on capital and materials, are listed on the stock market and hence have higher-quality data (243 firms). The coefficient of the CEO behavior index is larger in magnitude (0.641) and significant at the 1% level (standard error = 0.278). In results reported in table D.2, we show that the coefficient on the CEO behavior index is of similar magnitude and significance when we use the Olley-Pakes estimator of productivity.

We have checked the robustness of the basic cross-sectional results in various ways. First, since the index summarizes information on a large set of activity features, a question of interest is whether this correlation is driven just by a subset of those features. To this purpose, in table D.1 we show the results of equation (1), controlling for the individual features used to compute the index separately. The table shows that each feature is correlated with performance on its own, so that the index captures their combined effect. Second, we have verified that the results are robust

from the CEO himself, and interviewer dummies. All columns are weighted by the week representativeness score assigned by the CEO at the end of the interview week. Errors are clustered at the three-digit SIC level. Since the data are averaged over 3 years, year dummies are set as the rounded average year for which the performance data is available.

[28] To make this comparison, we multiply the coefficient of the CEO behavior index in col. 2 (0.227) by the standard deviation of the index in the subsample, $(0.227 \times 0.33) = 0.07$, and express it relative to the same figures for capital $(0.387 \times 1.88 = 0.73)$.

to using more standard dimensionality-reduction techniques, such as *k*-means and principal components. In Table D.2, we show that these alternative ways of classifying CEOs do not fundamentally alter the relationship between CEO behavior and firm performance.

### 2. Management

What CEOs do with their time may reflect broader differences in management processes across firms rather than CEO behavior per se. To investigate this issue, we matched the CEO behavior index with management practices collected in the World Management Survey (Bloom, Sadun, and Van Reenen 2016).[29] We were able to gather management data for 191 firms in our CEO sample.

The CEO behavior index is positively correlated with the average management score: a 1-standard-deviation change in the management index is associated with a 0.054 increase in the CEO behavior index.[30] Management and CEO behavior, however, are independently correlated with firm productivity, as we show in column 5 of table 3, using the sample of 156 firms for which we could match the management and CEO behavior data with accounting information. The coefficients imply that a 1-standard-deviation change in the CEO behavior (management) index is associated with an increase of 0.16 (0.19) log points in sales.[31] Overall, these results imply that the CEO behavior index is distinct from other, firm-wide, management differences.

### 3. Profits

Column 6 of table 3 analyzes the correlation between CEO behavior and profits per employee. This allows us to assess whether CEOs capture all

[29] The survey methodology is based on semistructured double-blind interviews with plant-level managers, run independently from the CEO time-use survey.

[30] This is the first time that data on middle-level management practices and CEO behavior have been combined. The correlation between CEO behavior and management practices is driven primarily by practices related to operational practices, rather than by human resources– and people-related management practices. See table D.7 for details. Bender et al. (2018) analyze the correlation between management practices and employees' wage fixed effects and find evidence of sorting of employees, with higher fixed effects in better-managed firms. The analysis also includes a subsample of top managers, but because of data confidentiality it excludes the highest-paid individuals, who are likely to be CEOs.

[31] The magnitude of the coefficient on the management index is similar to the one reported by Bloom, Sadun, and Van Reenen (2016) in the full management sample (0.15). When we do not control for the management (CEO) index, the coefficient on the CEO (management) index is 0.544 (0.199), significant at the 5% level in the subsample. When we also control for capital, the sample goes to 98 firms, but the coefficients on both the CEO index and management remain positive and statistically significant. Controlling for materials leaves us with only 56 observations, and on this subsample the CEO behavior and management are not statistically significant, even before we control for materials. See table D.7 for more details.

TABLE 3
CEO Behavior and Firm Performance

| | Dependent Variable: log(sales) | | | | | Dependent Variable: Profits/Employee |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| CEO behavior index | .343*** | .227** | .322*** | .641** | .506** | 10.029*** |
| | (.108) | (.111) | (.121) | (.278) | (.236) | (3.456) |
| Log | (.040) | (.066) | (.099) | (.152) | (.090) | (.734) |
| (employment) | .889*** | .555*** | .346*** | .339** | .784*** | −.284 |
| Log(capital) | | .387*** | .188*** | .194* | | |
| | | (.042) | (.056) | (.098) | | |
| Log(materials) | | | .447*** | .421*** | | |
| | | | (.073) | (.109) | | |
| Management | | | | | .179** | |
| | | | | | (.072) | |
| Observations (firms) | 920 | 618 | 448 | 243 | 156 | 386 |
| Observations used to compute means | 2,202 | 1,519 | 1,054 | 604 | 383 | 1,028 |
| Sample | All | With $k$ | With $k$ and $m$ | With $k$ and $m$, listed | With management score | With profits, listed |

Note.—We include at most 3 years of data for each firm and build a simple average across output and all inputs over this period. The number of observations used to compute these means are reported at the foot of the table. “Management” is the standardized value of the Bloom and Van Reenen (2007) management score. The sample in col. 1 includes all firms with at least one year with both sales and employment data. Columns 2–4 restrict the sample to firms with additional data on capital (col. 2) or capital and materials (cols. 3 and 4); the sample in col. 4 is restricted to listed firms. The sample in col. 5 is restricted to firms with a nonmissing management score. Columns 1–4 and 6 include a full set of country and year dummies, three-digit SIC industry dummies, and noise controls. Column 5 includes a full set of country dummies and two-digit SIC industry dummies. Noise controls in cols. 1–4 and 6 are a full set of dummies to denote the week in the year in which the data were collected, a reliability score assigned by the interviewer at the end of the survey week, a dummy taking value one if the data were collected through the PA of the CEO rather than from the CEO himself, and interviewer dummies. Noise controls in col. 5 are the reliability score assigned by the interviewer at the end of the survey week, a dummy taking value one if the data were collected through the PA of the CEO rather than from the CEO himself, the log of employment in the plant for which the management score is computed, an index measuring the reliability of the management score, dummies to denote the year in which the management interview was conducted, and the duration of the management interview. All columns are weighted by the week representativeness score assigned by the CEO at the end of the interview week. Errors clustered at the three-digit SIC level are in parentheses.

* Significant at the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

the extra rent they generate or whether firms profit from being run by leader CEOs. The results are consistent with the latter: the correlation between the CEO index and profits per employee is positive and precisely estimated. The magnitudes are also large: a 1-standard-deviation increase in the CEO behavior index is associated with an increase of approximately \$3,100 in profits per employee. Another way to look at this issue is to compare the magnitude of the relationship between the CEO behavior index and profits to the magnitude of the relationship between the CEO behavior index and CEO pay. We are able to make this comparison for a subsample of 196 firms with publicly available compensation data. Over this subsample, we find that a 1-standard-deviation change in the CEO behavior index is associated with an increase in profits per employee of \$4,939 (which, using the median number of employees in the subsample, would correspond to a \$2,978,000 increase in total profit) and an increase in annual CEO compensation of \$47,081. According to the point estimates above, the CEO keeps less than 2% of the marginal value he creates through his behavior. This broadly confirms the finding that the increase in firm performance associated with higher values of the CEO behavior index is not fully appropriated by the CEO in the form of rents.

### *B. Correlations with Multidimensional Indices*

Working with only two pure behaviors has the clear advantage of delivering a one-dimensional index, which is easy to represent and interpret. In contrast, when the approach is extended to $K$ rather than two pure behaviors, the behavioral index becomes a point on a $(K - 1)$-dimensional simplex. However, a natural question to ask is whether the simplicity of the two-behaviors approach may lead to significant loss of information, especially for the correlation between CEO behavior and firm performance. There are numerous model-selection approaches in the unsupervised-learning literature, and in appendix D.2.7 we detail two that we have implemented. The first is based on out-of-sample goodness of fit, and a range of models from $K = 5$ to $K = 25$ all appear to perform similarly. The second is a simulation-based analog of the Akaike information criterion. This criterion rewards in-sample goodness of fit, as measured by the average log likelihood across draws from Markov chains, and punishes model complexity, as measured by the variance of the log likelihood across the draws. It selects $K = 4$ as the optimal model.

Since the available methods do not univocally suggest a single optimal $K$, rather than wed ourselves to the idea of a single best model, we compare our baseline model with $K = 2$ to models with $K = 3$ through $K = 11$ (inclusive), as well as larger models with $K = 15$ and $K = 20$. First, we look at whether the use of a larger number of pure behaviors can better account for the observed variation in firm performance. To do so, table D.3 compares the $R^2$ of the regressions shown in table 3 when CEO behavior is summarized by these multidimensional indices. The first row displays the $R^2$ statistics from each of the six regressions in table 3 when we use the baseline scalar CEO behavior index. Each subsequent row then displays the $R^2$ from regressions in which we replace the scalar CEO behavior index with $K - 1$ separate indices that measure the time that each CEO allocates across $K$ pure behaviors. The main conclusion is that the explanatory power of CEO behavior for firm performance is remarkably constant across different values of $K$. While a model with a higher $K$ may better fit the variation in the time-use data, this better fit does not translate into a greater ability to explain firm performance.

Another question of interest is whether models with $K > 2$ identify the same behavioral distinction between leaders and managers that we emphasize above. To make the models comparable, for each CEO and value of $K$ we compute the similarity between the leader pure behavior estimated in the model with $K = 2$ (which here we denote $\hat{\boldsymbol{\beta}}^{\mathrm{L}}$) and the pure behaviors estimated in the richer model and use this as a weight to aggregate the different pure behaviors.[32] We then use this weighted average for each different value of $K$ in place of the CEO behavior index in the regressions in table 3. That is, we build a synthetic behavior index that aggregates across all the different pure behaviors while taking into account their (dis)similarity to (from) the pure leader behavior found in the $K = 2$ case. Table D.4 shows the results. In all cases the coefficient is positive, and in the large majority of cases it retains the same significance as the $K = 2$ case.[33] These results are reassuring in that they indicate that the distinction between leaders and managers remains an important source of variation even in models with higher $K$.

## IV. CEO Behavior and Firm Characteristics

The correlations presented in section III may simply reflect the fact that CEO behavior proxies for firm characteristics correlated with firm performance. To explore this idea, we proceed in two ways. First, we study the correlation between observable firm characteristics and CEO behavior and test whether these variables account for the correlation between CEO behavior and performance. Second, we use firm performance in the years predating the CEO appointment to test (1) whether differences

[32] The precise formula is $\Sigma_{k=1}^{K}\hat{\theta}_{i,k}[1 - H(\widehat{\boldsymbol{\beta}}^{k}, \widehat{\boldsymbol{\beta}}^{\mathrm{L}})]$, where $\widehat{\boldsymbol{\beta}}^{\mathrm{L}}$ is the pure behavior corresponding to the leader in the model with $K = 2$, $\widehat{\boldsymbol{\beta}}^{k}$ is the $k$th pure behavior in the model with $K > 2$, $\hat{\theta}_{i,k}$ is the share of time CEO $i$ is estimated to spend in pure behavior $k$, and $H$ is the Hellinger distance between the two.

[33] The main exception is in the reduced-sample regression in col. 5, which is based on the sample of 156 observations for which we have both the CEO behavior index and a firm-level management score drawn from the World Management Survey project.

in productivity trends before the CEO appointment predict the type of CEO that is eventually hired by the firm and (2) whether the CEO behavior index is associated with changes in productivity relative to the period preceding the appointment of the CEO. We can implement this latter test on the 204 firms that have accounting data within a 5-year interval both before and after CEO appointment.

### A. *Cross-Sectional Correlations*

Columns 1–6 of table 4 show that the CEO behavior index covaries positively with firm size, as proxied by number of employees, and dummies denoting firms listed on public stock exchanges, multinationals, and firms part of a larger corporate group. The index also varies across industries, with higher values in industries characterized by a greater intensity of managerial and creative tasks relative to routine tasks (which we identify using the industry-level measures built by Autor, Levy, and Murnane 2003) and greater R&D intensity (defined as industry business R&D divided by industry employment from National Science Foundation data). Conversely, the index is significantly lower in firms owned and managed by a family CEO, but this correlation turns insignificant when we control for the other variables (col. 6).

Overall, these correlations suggest that CEOs tend to spend a greater fraction of their time in coordinative rather than operational activities—which in our data would correspond to higher values of the CEO behavior index—when production activities are more complex and/or more skill intensive. These findings are consistent with the notion that coordination on the part of CEOs is particularly valuable in these circumstances. Drucker (1967, ch. 2.I), for example, mentions the importance of personal CEO meetings in the management of knowledge workers, arguing that the "relationships with other knowledge workers are especially time consuming."[34]

[34] According to Drucker, this is due to both status issues and information obstacles: "Whatever the reason—whether it is absence of the barrier of class and authority between superior and subordinate in knowledge work, or whether he simply takes more seriously—the knowledge worker makes much greater time demands than the manual worker on his superior as well as on his associates.. . .One has to sit down with a knowledge worker and think through with him what should be done and why, before one can even know whether he is doing a satisfactory job or not" (ch. 2.I). Similarly, Mintzberg (1979) emphasizes the importance of informal communication activities in the coordination of complex organizations. Mintzberg (1979, 3) refers to "mutual adjustments"—i.e., the "achievement of the coordination of work by simple process of informal communication"—in his proposed taxonomy of the various coordination mechanisms available to firms. Mintzberg states that mutual adjustment will be used in the very simplest of organizations, as well as in the most complicated. The reason is that this is "the only system that works under extremely difficult circumstances."

TABLE 4
CEO Behavior and Firm Characteristics

| | Dependent Variable: CEO Behavior Index | | | | | | Dependent Variable: log(sales) |
|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| CEO behavior index | | | | | | | .288** |
| | | | | | | | .116) |
| Log (employment) | .056*** | | | | | .048*** | .874*** |
| | (.009) | | | | | (.009) | (.038) |
| MNE (dummy) | | .105*** | | | | .075*** | .097 |
| | | (.025) | | | | (.024) | (.080) |
| Part of a group (dummy) | | | .132*** | | | .124*** | .047 |
| | | | (.023) | | | (.024) | (.086) |
| Listed (dummy) | | | | .104*** | | .043 | .141* |
| | | | | (.035) | | (.035) | (.084) |
| Family CEO (dummy) | | | | | −.066*** | −.007 | −.216** |
| | | | | | (.022) | (.022) | (.092) |
| Adjusted $R^2$ | .257 | .232 | .244 | .225 | .224 | .291 | .772 |
| Observations (firms) | 1,114 | 1,114 | 1,114 | 1,114 | 1,114 | 1,114 | 920 |
| Observations used to compute means | | | | | | | 2,202 |

Note.—"MNE (dummy)" is a variable taking value one if the firm is a domestic or foreign multinational enterprise. "Part of a group (dummy)" is a variable taking value one if the firm is affiliated to a larger corporate group. "Listed (dummy)" is a variable taking value one if the firm is listed on a public stock exchange. "Family CEO (dummy)" is a variable taking value one if the firm is owned by the founding family and the CEO is part of the owning family. All columns include a full set of country and year dummies, three-digit SIC industry dummies, and noise controls. Noise controls are a full set of dummies to denote the week in the year in which the data were collected, a reliability score assigned by the interviewer at the end of the survey week, a dummy taking value one if the data were collected through the PA of the CEO rather than from the CEO himself, and interviewer dummies. All columns are weighted by the week representativeness score assigned by the CEO at the end of the interview week. The sample in col. 7 includes all firms with at least one year of both sales and employment data. We include at most three years of data for each firm and build a simple average across output and all inputs over this period. The number of observations used to compute these means is reported at the foot of the table. Errors clustered at the three-digit SIC level are in parentheses.

* Significant at the 10% level.
** Significant at the 5% level.
*** Significant at the 1% level.

These findings raise the concern that CEOs may simply adapt their behavior to the characteristics of the firms they run—that is, that CEO behavior may simply be a proxy for firm characteristics correlated with firm performance. It is important to notice, however, that while some of the firm characteristics considered in table 4 are correlated with firm performance, they do not fully account for the correlation between CEO behavior and firm performance. To see this, consider column 7, in which we augment the specification of column 1 in table 3 with these additional variables. This shows that the coefficient on CEO behavior remains positive and significant, with a similar magnitude even when these additional controls are included.[35]

### B. *Exploiting Data before and after the CEO Appointment*

To consider the role of unobservable firm characteristics beyond the ones considered in table 4, we turn to the subsample of 200 firms for which we have firm performance data both before and after the CEO appointment.[36]

This analysis is presented in table 5. To start, column 1 shows that the set of firms with available data before and after CEO appointment are representative of the larger sample in terms of the correlation between the CEO behavior index and performance. The correlation is 0.360 (standard error = 0.132) for firms that do not belong to the subsample, and the interaction between the CEO behavior index and the dummy denoting the subsample equals −0.082 and is not precisely estimated.

We then test whether productivity trends before appointment can predict the type of CEO eventually hired by the firm. Column 2 shows that this is not the case—in the preappointment period, firms that eventually appoint a leader CEO have productivity trends similar to those of firms that hire managers.

Next, we investigate whether the correlation between CEO behavior and firm performance simply reflects time-invariant firm heterogeneity by estimating the following difference-in-differences model:

$$y_{ft} = \alpha A_t + \boldsymbol{\beta A_t \hat{\theta}_i} + \delta^{\mathrm{E}} \boldsymbol{e}_{ft} + \boldsymbol{\zeta}_t + \boldsymbol{\eta}_f + \varepsilon_{it}, \quad (2)$$

[35] Table D.6 repeats the same exercise for all the other columns of table 3. The data also show that CEO behavior varies systematically with specific CEO characteristics, namely, CEO skills (college or MBA degree) and experience abroad (see app. D.1 for more details). Note, however, that the correlation between CEO behavior and firm characteristics (and firm size in particular) remains large and significant even when we control for CEO traits. This points to the fact that observable CEO characteristics—i.e., what a board would observe by simply looking at the CV of the potential CEO—do not fully capture differences in CEO behavior. This can be one of the reasons why a mismatch between CEOs and firms may arise in equilibrium. We come back to this point in sec. V.

[36] We do not find this subsample of firms with before-and-after data to be selected in terms of the magnitude of the CEO behavior index or firm size. The subsample, however, tends to be skewed toward professional CEOs relative to family CEOs. This is because family CEOs tend to have longer tenures—therefore, the before-appointment period is typically not observed. The sample is also more skewed toward firms located in France, Germany, and the United Kingdom relative to US firms. This is due to the fact that accounting panel data for US private firms—of which our sample is primarily composed—are typically less complete than those for Europe.

TABLE 5

CEO Behavior and Firm Performance before and after CEO Appointment (Dependent Variable: log(sales))

| | Relative to Appointment of Current CEO | | | | |
|---|---|---|---|---|---|
| | After (1) | Before (2) | Before and After (3) | Before and After (with 2 "After" Subperiods) (4) | Before and After (with 2 "After Subperiods): CEO Tenure ≤ 3 Years at Time of the Survey (5) |
| CEO behavior index | .360*** | | | | |
| | (.132) | | | | |
| Firm is in balanced sample | .154 | | | | |
| | (.119) | | | | |
| Firm is in balanced sample × CEO behavior index | −.082 | | | | |
| Trend | | .006 | | | |
| | | (.018) | | | |
| Trend × CEO behavior index | | −.008 | | | |
| | | (.029) | | | |
| After CEO appointment | | | −.004 | | |
| | | | (.111) | | |
| After CEO appointment × CEO behavior index | | | .123** | | |
| | | | (.057) | | |
| After CEO appointment ($1 \leq t \leq 2$) | | | | .127 | .251 |
| | | | | (.111) | (.238) |

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| After CEO appointment ($3 \leq t \leq 5$) | | | | .190 | .203 |
| | | | | (.181) | (.394) |
| After CEO appointment ($1 \leq t \leq 2$) × CEO behavior index | | | | .052 | .139 |
| | | | | (.071) | (.122) |
| After CEO appointment ($3 \leq t \leq 5$) × CEO behavior index | | | | .215** | .428** |
| | | | | (.094) | (.182) |
| Log(Employment) | .888*** | .916*** | .785*** | .757*** | .736*** |
| | (.039) | (.091) | (.074) | (.061) | (.067) |
| Observations | 920 | 684 | 400 | 563 | 263 |
| Number of firms | 920 | 204 | 200 | 200 | 97 |
| Spell averages | Yes | No | Yes | Yes | Yes |
| Firm fixed effects | No | Yes | Yes | Yes | Yes |

NOTE.—All columns include the same controls used in col. 1 of table 3. The sample in cols. 1–4 includes the set of firms with at least one year of nonmissing productivity data in both of the 5-year intervals before and after CEO appointment; in col. 5 we also exclude CEOs who had been in their position for more than 3 years at the time of the survey. "Firms in balanced sample" is a dummy taking value one if the firm is part of this set. Productivity data in col. 1 are aggregated as in col. 1 of table 3. Column 2 uses all available yearly data within 5 years before CEO appointment. In col. 3, we build averages of output and inputs using data in the 5 years before CEO appointment and the 5 years after CEO appointment, combine the two cross sections, and include firm-level fixed effects. Columns 4 and 5 split the after-CEO-appointment period into two subperiods: 1–2 and 3–5 years after appointment. "After CEO appointment" is a dummy taking value one for the cross section computed in the years after CEO appointment. Errors clustered by industry in col. 1, by firm in col. 2, and by firm and before/after CEO appointment period in cols. 3–5 are in parentheses.

** Significant at the 5% level.

*** Significant at the 1% level.

where $t$ denotes whether the time period refers to the five years before or after the appointment of the CEO. Similarly to the results shown in table 3, inputs and outputs are aggregated across the two different subperiods, before and after CEO appointment. Here, $\eta_f$ are firm fixed effects, $A_t = 1$ after appointment, and $\hat{\theta}_i$ is the behavior index of the appointed CEO. The linear CEO behavior index term is omitted, since it is absorbed by the firm fixed effects. The coefficient of interest is $\boldsymbol{\beta}$, which measures whether firms that eventually appoint a CEO with higher levels of the CEO behavior index experience a greater increase in productivity after the CEO is in office relative to the years preceding the appointment.[37]

Column 3 shows that the coefficient $\boldsymbol{\beta}$ is positive and significant (coefficient = 0.123, standard error = 0.057). Given this coefficient, the within-firm change in productivity after the CEO appointment is 0, 0.03, and 0.11 log points for values of the CEO index at, respectively, the 10th, 50th, and 90th percentiles of the distribution of the CEO behavior index. In column 4, we provide more detail on the nature of the correlation between CEO behavior and performance by splitting the "after" period into two subperiods: 1–2 and 3–5 years after appointment. The results suggest that the correlation materializes only 3 years after appointment.

While the before-and-after results discussed so far control for time-invariant firm heterogeneity, CEOs may adjust their behavior in response to unobserved time-varying productivity shocks following their appointment. To investigate this issue, we restrict the sample to the 97 firms whose current CEOs had been in office for less than 3 years at the time of the survey—that is, we correlate the estimated CEO behavior with future changes in productivity. The results of this exercise are shown in column 5. The fact that the results hold, and are actually stronger, in this smaller sample of less experienced CEOs casts doubt on the hypothesis that the results are entirely driven by CEO learning effects, at least in the very first years after the appointment is made.

In sum, differences in time-invariant firm-level characteristics, time-varying shocks to performance predating the CEO appointment, or CEOs adapting their behavior to productivity shocks cannot fully account for the relationship between CEO behavior and firm performance. The evidence does not rule out that firms hire CEOs with specific behavioral traits in response to unobserved time-varying productivity shocks contemporaneous to the CEO appointment. Since the correlation materializes 3 years after the CEO is appointed, this would imply that corporate boards are able to predict performance 3 years in advance and to replace CEOs 3 years before the predicted performance effects actually occur.[38]

[37] Note that, since we do not know the behavior of the previous CEO, this is a lower bound on the effect of switching from manager to leader CEOs, since at least a fraction of these firms would have had already a leader CEO before the current appointment.

[38] Table D.5 replicates the table, using the weighted average of the pure behaviors from models with higher $K$ discussed in sec. III.B.

### C. Summary

Taken together, the results discussed in this section suggest that, while correlated with firm traits associated with firm performance, CEO behavior does not appear to be fully endogenous to firm performance. These findings open the door to the possibility that the behavior of CEOs itself could be a possible driver of firm performance, rather than just its mere reflection. In the next section, we present a simple model that illustrates the different channels through which this effect may arise in the data.

## V. Vertical or Horizontal Differentiation in CEO Behavior?

The findings in section IV show that the CEO behavior is not a mere reflection of firm traits. However, the fact that the appointment of a leader CEO is associated with an increase in performance for the average firm does not necessarily imply that all firms would benefit from hiring a leader CEO—that is, that CEOs are vertically differentiated in terms of their behavior. In fact, a positive correlation between CEO behavior and performance may also arise in the case in which CEOs are horizontally differentiated—some firms are better off with leaders and others with managers—if matching frictions are sufficiently large.

We illustrate this point through a simple assignment model consisting of CEOs with different behaviors who are matched to firms with different characteristics. In the case of vertical differentiation, leaders are preferred by all firms, and those who are able to hire one perform better. In the horizontal case, some firms prefer managers, but if managers are relatively more abundant than the demand for their services, some of the firms that should be matched with leaders instead end up with managers and consequently suffer a performance penalty.

### A. Simple Assignment Model

#### 1. Setup

CEO $i$ can adopt one of two possible behaviors: $x_i = \mathrm{m}$ (“manager”) and $x_i = \mathrm{l}$ (“leader”). Once a CEO is hired, he decides how he is going to manage the firm that hired him. CEO $i$ has a type $\tau_i \in \{\mathrm{m}, \mathrm{l}\}$. Type m prefers behavior m to behavior l; that is, he incurs a cost of 0 if he selects behavior m and a cost of $c > 0$ if he selects behavior l. Type l is the converse: he incurs a cost of 0 if he selects behavior l and cost of $c$ if he selects behavior m. The cost of choosing a certain behavior can be interpreted as coming from the preferences of the CEO (i.e., he may find one behavior more enjoyable than the other) or his skill set (i.e., he may find one behavior less costly to implement than the other).

Firms also have types. The type of firm $f$ is $\tau_f \in \{\mathrm{m}, \mathrm{l}\}$. The output of firm $f$ assigned to CEO $i$ is

$$y_{fi} = \lambda_f + \left(I_{\tau_f = x_i}\right)\Delta, \tag{3}$$

where $I$ is the indicator function and $\Delta > 0$. Hence, firm $f$'s productivity depends on two components. The first is a firm-specific component that we denote $\lambda_f$. In principle, this can depend on observable firm characteristics, unobservable firm characteristics, and more generally the firm's "innate" type. We include this term to build the unobserved firm heterogeneity issues discussed in section IV explicitly into the model and its subsequent estimation. The second component is specific to the behavior of the CEO; that is, if the CEO's behavior matches the firm's type, then productivity increases by a positive amount $\Delta$. This captures the fact that different firms require different behaviors: there is not necessarily a "best" behavior in all circumstances, but there is scope for horizontal differentiation. We assume that $c < \Delta$, so that it is efficient for the CEO to always adopt a behavior that corresponds to the firm's type.

To introduce the possibility of matching frictions, we must discuss governance. Firms offer a linear compensation scheme that rewards CEOs for generating good performance. The wage that CEO $i$ receives from employment in firm $f$ is

$$w\left(y_{fi}\right) = \overline{w} + B\left(y_{fi} - \lambda_f\right) = \overline{w} + BI_{\tau_f = x_i}\Delta,$$

where $\bar{w}$ is the fixed part and $B \geq 0$ is a parameter that can be interpreted directly as the performance-related part of CEO compensation or indirectly as how likely it is that a CEO is retained as a function of his performance (in this interpretation, the CEO receives a fixed per-period wage but is more likely to be terminated early if firm performance is low).

The total utility of the CEO is equal to compensation less behavior cost, that is, $w\left(y_{fi}\right) - I_{\tau_i \neq x_i} c$. After a CEO is hired, he chooses his behavior. If the CEO is hired by a firm with the same type, he will obviously choose the behavior that is preferred by both parties. The interesting case is when the CEO type and the firm type differ. If $B > c/\Delta$, the CEO will adapt to the firm's desired behavior, produce an output of $\lambda_f + \Delta$, and receive a total payoff of $\overline{w} + B\Delta - c$. If instead $B < c/\Delta$, the CEO will choose $x_i = \tau_i$, produce output $\lambda_f$, and receive a payoff $\overline{w}$. We think of $B$ as a measure of governance. A higher $B$ aligns CEO behavior with the firm's interests.

### 2. Pairing Firms and CEOs

Now that we know what happens once a CEO begins working for a firm, let us turn our attention to the assignment process. There is a mass 1 of firms. A proportion $\phi$ of them are of type l, the remainder are of type m. The pool of potential CEOs is larger than the pool of firms seeking a CEO. There is a mass $P \gg 1$ of potential CEOs. Without loss of generality, assume that a proportion $\gamma \leq \phi$ of CEOs are of type l. The remainder are of type m. From now on, we refer to type l as the "scarce" CEO type and type m as the "abundant" CEO type. We emphasize that scarcity is relative to the share of firm types. So it may be the case that the share of type l CEOs is actually larger than the share of type m firms. Note that the model nests the case of pure vertical differentiation, where no firm actually wants a type m CEO, that is, when $\phi = 1$.

The market for CEOs works as follows. In the beginning, every prospective CEO sends his application to a centralized CEO job market. The applicant indicates whether he wishes to work for a type m or a type l firm. All the applications are in a large pool. Each firm begins by downloading an application meant for its type. Each download costs $k$ to the firm. After receiving an application, firms receive a signal about the underlying type of the CEO who submitted it. If the type of the applicant corresponds to the type of the firm, the signal has value one. If the type is different, the signal is equal to zero with probability $\rho \in [0, 1]$ and to one with probability $1 - \rho$. Thus, $\rho = 1$ denotes perfect screening and $\rho = 0$ represents no screening.[39] This last assumption distinguishes our approach from existing theories of manager-firm assignment, where the matching process is assumed to be frictionless, and the resulting allocation of managerial talent achieves productive efficiency.[40]

Potential CEOs maximize their expected payoff, which is equal to the probability that they are hired times the payoff if they are hired. Firms maximize their profit less the screening cost (given by the number of downloaded application multiplied by $k$). Clearly, if $k$ is low enough, firms download applications until they receive one whose associated signal indicates that the CEO type matches the firm type, which we assume holds in equilibrium.

Define residual productivity as total productivity minus type-specific baseline productivity: $y_{fi} - \lambda_f$.

PROPOSITION 1. Firms led by the type l CEOs and those led by the type m CEOs have equal residual productivity if at least one of the following

[39] The implicit assumption is that CEOs have private information about their types, while firms' types are common knowledge. However, we could also allow firms to have privately observed types; in equilibrium, they will report them truthfully. Moreover, if CEOs have limited or no knowledge of their own type, it is easy to see that our mismatch result would hold a fortiori.

[40] See, e.g., Gabaix and Landier (2008), Tervio (2008), and Bandiera et al. (2015). An exception in the literature is Chade and Eeckhout (2016), who present a model in which agents' characteristics are realized only after a match is formed, which leads to a positive probability of mismatch in equilibrium.

conditions is met: (i) neither CEO type is sufficiently scarce, (ii) screening is sufficiently effective, or (iii) governance is sufficiently good.

Each of the three conditions guarantees efficient assignment. If there is no scarce CEO type ($\gamma = \phi$), a CEO has no reason to apply to a firm of a different type. If screening is perfect ($\rho = 1$), a CEO who applies to a firm of the other type is always caught (and hence he will not do it). If governance is good ($B < c/\Delta$), a CEO who is hired by a firm of the other type will always behave in the firm's ideal way (and hence there will either be no detectable effect on firm performance or CEOs will apply only to firms of their type).

In contrast, if any of conditions i–iii are not met, CEO behavior and firm performance will be correlated because of inefficient assignments. The following proposition characterizes how the latter can occur in equilibrium and the implications of the mismatches for observed performance differentials.

Proposition 2. If the screening process is sufficiently unreliable, governance is sufficiently poor, and one CEO type is sufficiently abundant,[41] then in equilibrium:

- all scarce-type CEOs are correctly assigned;
- some abundant-type CEOs are misassigned;
- the average residual productivity of firms run by abundant-type CEOs is lower than those of firms run by scarce-type CEOs.

*Proof.* See appendix C.

The intuition for this result is as follows. If all abundant-type CEOs applied to their firm type, they would have a low probability of being hired, and they would prefer to apply to the other firm type and try to pass as scarce-type CEOs. In order for this to be true, it must be that the share of abundant types is sufficiently larger than the share of scarce types and that the risk that they are screened out is not too large. If this is the case, then in equilibrium some abundant-type CEOs will apply to the wrong firm type, up to the point where the chance of getting a job is equalized under the two strategies. In the extreme case of vertical differentiation where $\phi = 1$, that is, when no firm demands type m CEOs, abundant-type CEOs reduce productivity in all firms.

What does proposition 2 imply for productive efficiency? Recall that in this economy the pool of scarce-type potential CEOs is sufficiently large to cover all firms (because $P \gg 1$). Thus, productive efficiency could be achieved, but it is not if the conditions for proposition 2 are satisfied.[42]

[41] Formally, this is given by the conditions $B < c/\Delta$ and $\rho < (\phi - \gamma)/(\phi - \gamma\phi)$.

[42] If side transfers were feasible, this would also be a Pareto improvement, as a type l CEO assigned to type m firm generates a higher bilateral surplus than a type m CEO matched with a type l firm, and the new firm-CEO pair could therefore compensate the now-unemployed type m CEO for her job loss.

### 3. From Theory to Data

As described in equation (3), the output of firm $f$ assigned to CEO $i$ depends on firm type and CEO behavior. Then the observed difference in performance between firms that hire a type l CEO and those that hire a type m CEO is

$$y_{\cdot\mathrm{l}} - y_{\cdot\mathrm{m}} = [s_{\mathrm{l}}(\lambda_{\mathrm{l}} + \Delta) + (1 - s_{\mathrm{l}})\lambda_{\mathrm{m}}] - [s_{\mathrm{m}}(\lambda_{\mathrm{m}} + \Delta) + (1 - s_{\mathrm{m}})\lambda_{\mathrm{l}}],$$

where $s_i$ is the share of CEOs who are correctly assigned to their firm types. That is, the average performance of firms led by type l CEOs is equal to the performance of type l firms when correctly matched $(\lambda_{\mathrm{l}} + \Delta)$, weighted by the share of type l CEOs who are correctly assigned $(s_{\mathrm{l}})$ plus the performance of misassigned type m firms $(\lambda_{\mathrm{m}})$ weighted by the share of type l CEOs who are wrongly assigned $(1 - s_{\mathrm{l}})$.

Simplifying and imposing the condition of proposition 2 by which all scarce-type CEOs are correctly matched in equilibrium (that is, $s_{\mathrm{l}} = 1$) yields

$$y_{\cdot\mathrm{l}} - y_{\cdot\mathrm{m}} = s_{\mathrm{m}}(\lambda_{\mathrm{l}} - \lambda_{\mathrm{m}}) + (1 - s_{\mathrm{m}})\Delta. \tag{4}$$

Equation (4) highlights two important points. First, the case in which performance differentials reflect entirely firm heterogeneity through the $(\lambda_{\mathrm{l}} - \lambda_{\mathrm{m}})$ term maps into a situation in which CEOs are horizontally differentiated and there are no matching frictions—that is, $s_{\mathrm{m}} = 1$. Second, there are two alternative mechanisms through which CEO behavior may lead to cross-sectional performance differentials across firms:

- *Horizontal differentiation in CEO behavior with matching frictions.* In this case, there is demand for both types of CEOs, but matching is imperfect, such that $0 < s_{\mathrm{m}} < 1$. Performance differentials capture the costs of the mismatches of type m CEOs $(\Delta)$, as well as firm heterogeneity.
- *Vertical differentiation in CEO behavior.* In this case, there is no demand for type m CEOs; that is, $s_{\mathrm{m}} = 0$. In this case, performance differentials reflect entirely the costs of the mismatches of type m CEOs $(\Delta)$.

In the absence of exogenous variation that would allow us to distinguish between these different mechanisms, we evaluate the plausibility of these alternatives by estimating the model and assessing which values of the parameters $s_{\mathrm{m}}$, $\Delta$, and $(\lambda_{\mathrm{l}} - \lambda_{\mathrm{m}})$ best fit the data.

## *B. Model Estimation*

The main data input of the model is firms' conditional productivity; that is, the residuals of a regression of productivity on firm characteristics as estimated in column 1 of table 3, without country fixed effects, which we model separately for reasons explained below. We denote the residual of firm $f$ run by CEO $i$ as $\hat{\varepsilon}_{if}$.[43] To obtain an empirical proxy of $x_i$, we use $\hat{x}_i = 1$ whenever $\hat{\theta}_i \geqslant 0.5$. That is, we discretize the CEO behavior index, using 0.5 as a cutoff, such that all CEOs above this threshold are classified as leaders and the rest as managers.

### 1. Nonparametric Evidence

The theoretical model suggests that, under vertical differentiation, the distribution of productivity for managers is drawn from a single distribution corresponding to inefficient matches, while the productivity for leaders is drawn from a single distribution with a higher mean. In contrast, under horizontal differentiation, the distribution of productivity for managers is a mixture of two distributions: one corresponding to inefficient matches with a lower mean and one corresponding to efficient matches with a higher mean.

As an initial nonparametric test of the competing hypotheses, we plot kernel densities of firm productivity (de-meaned by country) according to CEO behavior in figure 5, both in the overall sample and broken down by income level. The low- and middle-income countries are Brazil and India, while the high-income countries are France, Germany, the United Kingdom, and the United States. The rationale for splitting the sample between high and low income levels is that we expect the level of development in a country to be negatively correlated with assignment frictions. This idea, in turn, is based on the existing evidence documenting a positive relationship between development, the supply of managerial capital, and good governance.[44]

While the pattern is somewhat masked in the full sample, the kernel densities in low-income countries (and, to some extent, in high-income countries) clearly indicate that the productivity distribution for manager-led firms can indeed be thought of as a mixture of two underlying

[43] To maintain comparability in the pooled vs. regional results that we discuss in the next section, we also limit the sample to those firms for which there is at least one observation per region, industry, and year, since these are used as controls in the estimation of the residuals. This leaves 851 observations out of 920.

[44] For example, Gennaioli et al. (2013) report wide differences in the supply of managerial/entrepreneurial human capital, using regional data for a large cross section of countries. Differences in the availability of basic managerial skills across countries and their relationship with development and firm performance are also discussed by Bloom, Sadun, and Van Reenen (2016). Furthermore, development is also likely to affect the quality of corporate governance, which affects both the selection and the dismissal of misassigned CEOs. La Porta, Lopez-De-Silanes, and Shleifer (1999) and La Porta et al. (2000) study the heterogeneity of corporate governance and ownership structures around the world. More recently, and specifically related to CEOs, Urban (2019) reports large differences in the percentage of CEOs dismissed for bad performance in public firms in Brazil and India (both 16%) vs. France (29%), Germany (40%), the United Kingdom (35%), and the United States (27%).

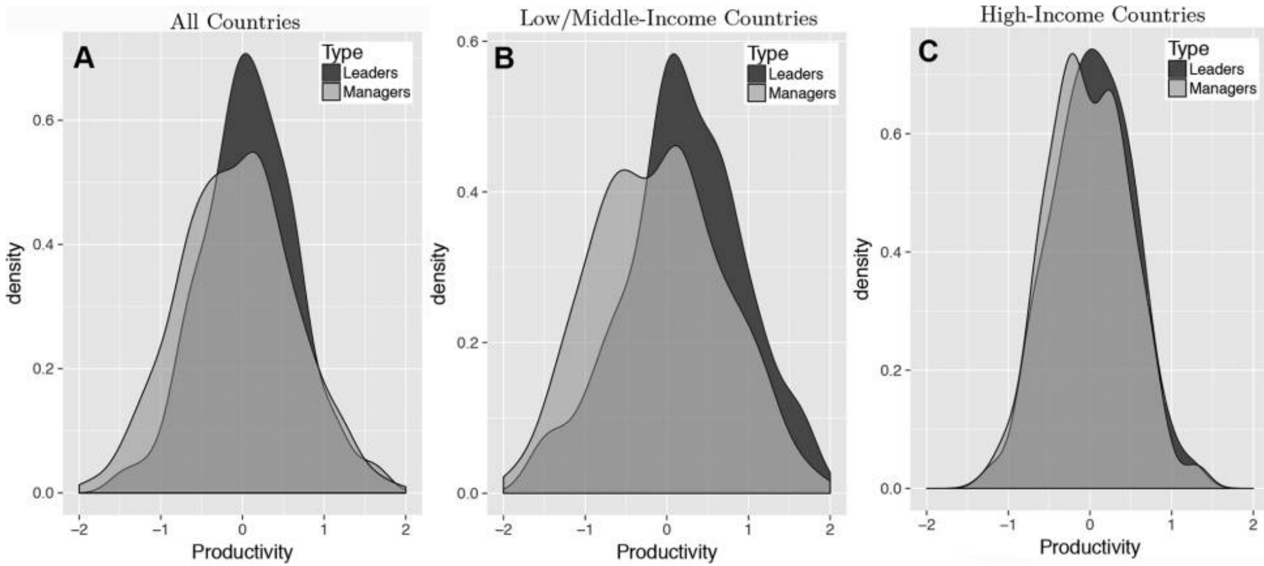


FIG. 5.—Kernel densities of productivity by CEO behavior. These figures display kernel densities of $\hat{\varepsilon}_{if}$ if de-meaned at the country level for leader-led and manager-led firms separately. *A* shows overall densities, *B* densities for Brazil and India, and *C* results for France, Germany, the United Kingdom, and the United States. A color version of this figure is available online.

distributions, the more productive of which appears to have a mean nearly identical to that of leader-led firms. This shows that the cross-sectional correlation between CEO behavior and firm performance is not driven by leaders being uniformly more productive than managers. Instead, many managers run firms that are, on average, as productive as leader-led firms. However, a substantial mass of managers also run less productive firms, which pulls down the overall average productivity of manager-led firms.

In order to explore these patterns in more detail, we now build and estimate a parametric model.

## 2. Parametric Model

In line with the theory, we adopt the statistical model $\hat{\varepsilon}_{if} = \lambda_f + (I_{\tau_f = x_i})\Delta + v_{if}$, where $\lambda_f$ is a "baseline" productivity; $\tau_f \in \{\mathrm{m}, \mathrm{l}\}$ is the firm's type; $x_i \in \{\mathrm{m}, \mathrm{l}\}$ is the CEO's behavior; and $\Delta$ is the productivity difference between firms with the "right" CEO and firms with the "wrong" CEO behavior relative to firm needs. While we treat $\hat{x}_i$ as observed data, $\tau_f$ is a random variable.

We assume that the conditions of proposition 2 hold. That is, we assume that since all type l CEOs ($\hat{x}_i = \mathrm{l}$) are correctly assigned, whenever we observe a type l we also must have $\tau_f = \mathrm{l}$. In contrast, only a share $s_{\mathrm{m}}$ of type m CEOs ($\hat{x}_i = \mathrm{m}$) are correctly assigned: when we observe a type m CEO, $\tau_f = \mathrm{m}$ with probability $s_{\mathrm{m}} \in [0, 1]$; otherwise, with probability $1 - s_{\mathrm{m}}$ the CEO is misassigned and $\tau_f = \mathrm{l}$.

As mentioned above, note that the model nests both pure vertical and pure horizontal differentiation. In the case of pure vertical differentiation, $s_{\mathrm{m}} = 0$; that is, all manager CEOs are misassigned. In the case of pure horizontal differentiation, vice versa: $s_{\mathrm{m}} = 1$; that is, all manager CEOs are assigned to firms that need their behavior. The main objective of the statistical model is to provide some evidence on which of these two scenarios is more consistent with the data.

As for the baseline productivity, we model $\lambda_f = x_{c_f, \tau_f}$, where $c_f$ denotes the country in which firm $f$ operates. This allows the model sufficient flexibility to capture that efficient and inefficient matches might have country-specific means, which figure 5 suggests is the case. We also assume that $x_{c_f, \mathrm{l}} = A + x_{c_f, \mathrm{m}}$, so that the baseline productivity of type l firms is that of type m firms plus a common constant term. This formulation allows for observed productivity differences between firms run by CEOs with different behaviors to arise from factors innate to firm types, in addition to the assignment-friction channel. Finally, we treat $v_{if}$ as a mean-zero normal random variable whose variance is both country and assignment specific: $\sigma^2_{1, c_f} (\sigma^2_{0, c_f})$ is the standard deviation of residuals in an efficient (inefficient) CEO-firm pair.

Given these observations, the likelihood function can be written as

$$\prod_{f\in\Theta(\mathrm{m})}\left\{\frac{s_{\mathrm{m}}}{\sqrt{2\pi}\sigma_{H,c_f}}\exp\left[-\frac{1}{2\sigma^2_{H,c_f}}\left(\hat{\varepsilon}_{if}-x_{c_f,\mathrm{m}}-\Delta\right)^2\right]+\frac{1-s_{\mathrm{m}}}{\sqrt{2\pi}\sigma_{L,c_f}}\exp\left[-\frac{1}{2\sigma^2_{L,c_f}}\left(\hat{\varepsilon}_{if}-A-x_{c_f,m}\right)^2\right]\right\}$$
$$\times\prod_{f\in\Theta(\mathrm{l})}\frac{1}{\sqrt{2\pi}\sigma_{H,c_f}}\exp\left[-\frac{1}{2\sigma^2_{H,c_f}}\left(\hat{\varepsilon}_{if}-A-x_{c_f,\mathrm{m}}-\Delta\right)^2\right]. \tag{5}$$

where $\Theta(\mathrm{m})$ and $\Theta(\mathrm{l})$ are the sets of firms managed by type m and type l CEOs, respectively. Type l CEOs are always efficiently assigned to type l firms, and their residuals are drawn from a normal distribution with mean $A + x_{c_f,\mathrm{m}} + \Delta$; in contrast, firms run by type m CEOs have their residuals drawn from a mixture of two normals, one with mean $x_{c_f,\mathrm{m}} + \Delta$ if the assignment is efficient, and another with mean $A + x_{c_f,\mathrm{l}}$ if the assignment is inefficient. The mixing probability is simply $s_{\mathrm{m}}$, the probability that type m CEOs are assigned to type m firms. We use the EM algorithm to maximize equation (5).

### 3. Estimates

The $A$ parameter is estimated to be $-.026$. Since the EM algorithm does not directly yield standard errors, we formally test the restriction $A = 0$ by plugging this value into equation (5) and maximizing with respect to the other parameters. A simple likelihood ratio test then fails to reject the restriction (the associated $p$-value is .706). Intuitively, when we divide type m CEOs into two groups, one with high performance and one with low performance, the high-performing group has productivity residuals with a mean statistically indistinguishable from that of the residuals of type l CEOs.[45] This is fully consistent with the pattern observed in figure 5.

The estimate of $\Delta$ is 0.532, which implies that the loss associated with an incorrect assignment of CEOs is substantial. Given that the units of the residual are log points, the estimate implies that moving from a correct assignment to an incorrect one reduces firm productivity by $(\exp(0.532) - 1)/\exp(0.532)$, or around 41%.

The estimated $s_{\mathrm{m}}$ is 0.744. To test whether the data are consistent with pure vertical differentiation, we impose the restriction $s_{\mathrm{m}} = 0$ in equation (5), which a likelihood ratio test rejects with a $p$-value of .00202.

[45] Note that in the E-step we explicitly infer the probability that type m CEOs are efficiently assigned, which allows us to then estimate parameters in the M-step. As is standard, the log likelihood is defined under the assumptions of the theoretical model, namely, that $\Delta > 0$ and that leader CEOs are scarce and all correctly assigned; thus, while there are combinations of parameters with $A > 0$ and $\Delta = 0$ that produce the same value of the likelihood, these violate the basic assumption of the model that correctly assigned firm-CEO pairs are more productive. Of course, nothing in the statistical model rules out both $\Delta > 0$ and $A > 0$, but, importantly, we find no role for $A$ when we optimize eq. (5) beginning from the best-fit solution with $\Delta > 0$.

The key underlying property of the data that lets us test $s_{\mathrm{m}} = 0$ is that under this restriction leader CEOs uniformly outperform manager CEOs. We can reject this in favor of a mixture model with $s_{\mathrm{m}} > 0$, since we observe a large fraction of manager CEOs whose performance is similar to that of leader CEOs. Also, note that once we reject $s_{\mathrm{m}} = 0$, we must necessarily reject $s_{\mathrm{m}} = 1$. In the model with $s_{\mathrm{m}} = 0$, we estimate separate mean parameters for managers and leaders and also separate variance parameters—these are match-quality specific, and managers are in a bad match while leaders are in a good match. By contrast, in the model with $s_{\mathrm{m}} = 1$ we fit separate mean parameters for managers and leaders but a single variance parameter, since all CEOs are in a good match. So the maximized likelihood will be lower for the model with $s_{\mathrm{m}} = 1$, compared to the model with $s_{\mathrm{m}} = 0$.

Overall, a model with heterogenous firms and assignment frictions fits the data significantly better than one without firm heterogeneity (pure vertical differentiation) or one without such frictions (pure horizontal differentiation). This formalizes the nonparametric observations above.

## 4. Quantifying the Importance of Matching Frictions for Aggregate Productivity

We now use the model to study the aggregate performance implications of CEO-firm matching frictions. To do so, we return to the differences in the parameter estimates across high- and low-/middle-income regions discussed at the beginning of the section.

We start from the quantification of the share of misassignments in the pooled sample. We first derive $\phi$, that is, the share of type l firms, from the market-clearing condition. Over the whole sample, we observe a share $\hat{\gamma} = 0.347$ of type l CEOs. We must then have $\phi = \hat{\gamma} + (1 - \hat{\gamma})(1 - s_{\mathrm{m}})$. The right-hand side of this expression is the total share of CEOs assigned to type l firms: all type l CEOs and a portion $1 - s_{\mathrm{m}}$ of type m CEOs. Plugging in for $\hat{\gamma}$ and $s_{\mathrm{m}}$, we obtain $\phi = 0.514$, so that slightly over half of firms are of type l. This in turn implies that a share $\phi - \hat{\gamma} = 0.168$ of firms are misassigned in our data, leading to an overall productivity loss of 0.089 ($=0.168 \times \Delta$) log points.

We then allow the $s_{\mathrm{m}}$ parameter in the likelihood function (eq. [5]) to vary according to whether the firm is located in a low-/middle- or high-income country. We restrict $A = 0$ in line with the results above. The estimation results are in table 6. In low-/middle-income countries, CEOs are efficiently assigned with probability 0.546, while the corresponding probability for CEOs in high-income countries is 0.893. The derived parameters in the table are obtained with the same steps described above.

One possible explanation for these different probabilities across countries is that firms in high-income countries have higher demand for type l

TABLE 6
Estimation Results by Region

| | Estimated Parameters | | Derived Parameters | | |
|---|---|---|---|---|---|
| | $\Delta$ (1) | $s_m$ (2) | $\hat{\gamma}$ (3) | $\phi$ (4) | Firms Mismatched (5) |
| Low/middle income countries | .667 | .546 | .216 | .572 | .356 |
| High income countries | .667 | .893 | .495 | .549 | .054 |

Note.—In cols. 1–2, this table displays the estimated parameters resulting from maximizing eq. (5) via the EM algorithm under the restriction that $A = 0$. Column 3 shows the observed share of leader CEOs in each region. Column 4 shows the value of $\phi$ consistent with market clearing, given $s_m$ and the observed share of leader CEOs, while col. 5 shows the difference between cols. 3 and 4, as this gives the share of type l firms run by manager CEOs.

CEOs. Indeed, consistent with this idea, the data show a much larger share of type l CEOs in high-income countries than in low-/middle-income countries (0.495 vs. 0.216). However, note that the $\phi$ parameters we extract—which capture the share of type l firms—are in fact very similar in both regions (if anything, there is slightly higher demand for type l CEOs in poorer countries).[46]

Instead, the main difference between regions emerging from the exercise is that type l firms in low-/middle-income countries are unable to locate and hire leader CEOs. It is important to reiterate that this is not necessarily due to scarcity of type l CEOs in the population per se. Rather, barriers to the allocation of talent might prevent the right individuals from entering the CEO job market. Regardless of the deeper cause, the share of inefficiently assigned type l firms in these countries is 0.356, compared to 0.054 in high-income countries. While there is still a sizable number of inefficient assignments in richer countries, the share in poorer countries is over six times as large.[47]

To conclude, we use our estimates to quantify how much productivity in low-income countries would increase if the assignment process were as efficient as it is in the richer countries in the sample. This implies building a counterfactual where $s_m$ increases from 0.546 to 0.893, which requires the share of leader CEOs to increase from 0.216 to 0.521 to maintain market clearing and which yields a drop in the share of misassigned firms from

[46] We have repeated the same $\chi^2$ tests for restrictions on $s_m$ described above for each region separately. While the power of the tests is lower because of reduced sample size, we are able to reject pure vertical and horizontal differentiation at a 10% significance level in both regions.

[47] Our findings provide a counterpoint to Chade and Eeckhout (2016), who estimate the degree of mismatch in the US CEO labor market, using wage data. First, while they find substantial mismatch, based on the deviation of the observed wage distribution, from what a model with perfect matching on observables would predict, our estimates that explicitly incorporate heterogeneity in CEO behavior indicate little mismatch in high-income countries. Second, they argue that nearly all match productivity differences arise from firm rather than CEO characteristics, whereas we find an important role for CEO heterogeneity.

0.356 to 0.051. Given that the productivity difference $\Delta$ is now estimated at a somewhat higher value of 0.667, productivity would increase by 0.203 log points.

We benchmark this magnitude against the macro differences in labor productivity across countries observed in the time interval covered by our survey and productivity data (2010–14), using the Penn World Table data, v.9 (Feenstra, Inklaar, and Timmer 2015). The average differences in log labor productivity between the two subsets of countries is 1.560. Therefore, improving the allocation of CEOs to firms in low-/middle-income countries could account for up to 13% of the cross-country differences in labor productivity.[48]

## VI. Conclusions

This paper combines a new survey methodology with a machine learning algorithm to measure the behavior of CEOs in large samples. We show that CEOs differ in their behavior along several dimensions and that the data can be reduced to a summary CEO index that distinguishes between "managers"—CEOs who are primarily involved with production-related activities—and "leaders"—CEOs who are primarily involved in communication and coordination activities.

Guided by a simple firm-CEO assignment model, we show that there is no "best practice" in CEO behavior—that is, a behavior that is optimal for all the firms—rather, there is evidence of horizontal differentiation in CEO behavior and significant frictions in the assignment of CEOs to firms. In our sample of manufacturing firms across six countries, we estimate that 17% of firm-CEO pairs are misassigned and that misassignments are found in all regions but are more frequent in emerging economies. The consequences for productivity are large: the implied productivity loss due to differential misassignment is equal to 13% of the labor productivity gap between firms in high- and middle-/low-income countries in our sample.

This paper shows that an underexplored dimension of managerial activity—that is, how CEOs spend their time—is both heterogeneous across managers and firms and correlated with firm performance. Future work could adopt our data and methodology to inform new leadership models that incorporate more explicitly the drivers and consequences of differences in CEO behavior and, in particular, explore the underlying firm-CEO matching function, which is not dealt with explicitly in this paper.

[48] The average labor productivity for high-income (low-/middle-income) countries in our sample is 11.4 (9.83). These values are calculated with data on output-side real GDP at chained purchasing power parities and the total number of persons engaged from the Penn World Tables.

Furthermore, a possible next step of this research would be to extend the data collection to the diaries of multiple managerial figures beyond the CEO. This approach would allow us to further explore whether and how managerial interactions and team behavior vary across firms and correlate with firm performance (Hambrick and Mason 1984). These aspects of managerial behavior, which are now largely absent from our analysis, are considered to be increasingly important in the labor market (Deming 2017) but have so far been largely unexplored from an empirical perspective. Finally, it would be fascinating to explore the relationship between CEO behavior and other personality traits, such as the ones considered in Kaplan, Klebanov, and Sorensen (2012) and Kaplan and Sorensen (2016). We leave these topics for further research.

## References

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Q.J.E.* 118 (4): 1279–333.

Bandiera, Oriana, Luigi Guiso, Andrea Prat, and Raffaella Sadun. 2015. "Matching Firms, Managers, and Incentives." *J. Labor Econ.* 33 (3): 623–81.

Bandiera, Oriana, Renata Lemos, Andrea Prat, and Raffaella Sadun. 2018. "Managing the Family Firm: Evidence from CEOs at Work." *Rev. Financial Studies* 31 (5): 1605–53.

Bender, Stefan, Nicholas Bloom, David Card, John Van Reenen, and Stefanie Wolter. 2018. "Management Practices, Workforce Selection, and Productivity." *J. Labor Econ.* 36 (S1): S371–S409.

Bennedsen, Morten, Kasper Meisner Nielsen, Francisco Pérez-González, and Daniel Wolfenzon. 2007. "Inside the Family Firm: The Role of Families in Succession Decisions and Performance." *Q.J.E.* 122 (2): 647–91.

Bertrand, Marianne, and Antoinette Schoar. 2003. "Managing with Style: The Effect of Managers on Firm Policies." *Q.J.E.* 118 (4): 1169–208.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *J. Machine Learning Res.* 3 (January): 993–1022.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2016. "Management as a Technology?" Working Paper no. 22327 (June), NBER, Cambridge, MA.

Bloom, Nicholas, and John Van Reenen. 2007. "Measuring and Explaining Management Practices Across Firms and Countries." *Q.J.E.* 122 (4): 1351–408.

Bolton, Patrick, Markus K. Brunnermeier, and Laura Veldkamp. 2013. "Leadership, Coordination, and Corporate Culture." *Rev. Econ. Studies* 80 (2): 512–57.

Cai, Jing, and Adam Szeidl. 2018. "Interfirm Relationships and Business Performance." *Q.J.E.* 133 (3): 1229–82.

Chade, Hector, and Jan Eeckhout. 2016. "Stochastic Sorting." Manuscript (September).

Coase, R. H. 1937. "The Nature of the Firm." *Economica*, new series, 4 (16): 386–405.

Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *Q.J.E.* 132 (4): 1593–640.

Dessein, Wouter, and Tano Santos. 2016. "Managerial Style and Attention." Manuscript, Columbia Univ.

Drucker, Peter. 1967. *The Effective Executive.* New York: Harper & Row.
Erosheva, Elena A., Stephen E. Fienberg, and Cyrille Joutard. 2007. "Describing Disability through Individual-Level Mixture Models for Multivariate Binary Data." *Ann. Appl. Statis.* 1 (2): 502–37.
Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2015. "The Next Generation of the Penn World Table." *A.E.R.* 105 (10): 3150–82.
Gabaix, Xavier, and Augustin Landier. 2008. "Why Has CEO Pay Increased So Much?" *Q.J.E.* 123 (1): 49–100.
Gennaioli, Nicola, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer. 2013. "Human Capital and Regional Development." *Q.J.E.* 128 (1): 105–64.
Graham, John R., Campbell R. Harvey, and Manju Puri. 2013. "Managerial Attitudes and Corporate Actions." *J. Financial Econ.* 109 (1): 103–21.
Griffiths, Thomas L., and Mark Steyvers. 2004. "Finding Scientific Topics." *Proc. Nat. Acad. Sci. USA* 101 (Suppl. 1): 5228–35.
Gross, Justin H., and Daniel Manrique-Vallier. 2014. "A Mixed Membership Approach to the Assessment of Political Ideology from Survey Responses." In *Handbook of Mixed Membership Models and Their Applications*, edited by Edoardo M. Airoldi, David M. Blei, Elena A. Erosheva, and Stephen E. Fienberg, 119–39. Boca Raton, FL: Chapman & Hall/CRC.
Hambrick, Donald C., and Phyllis A. Mason. 1984. "Upper Echelons: The Organization as a Reflection of Its Top Managers." *Acad. Management Rev.* 9 (2): 193–206.
*Harvard Business Review.* 2011. "Idea in Practice." In *HBR's* 10 *Must Reads on Leadership*, 40. Cambridge, MA: Harvard Bus. School.
Hermalin, Benjamin E. 1998. "Toward an Economic Theory of Leadership: Leading by Example." *A.E.R.* 88 (5): 1188–206.
———. 2007. "Leading for the Long Term." *J. Econ. Behavior and Org.* 62 (1): 1–19.
Hoffman, Mitchell, and Steven Tadelis. 2017. "How Do Managers Matter? Evidence from Performance Metrics and Employee Surveys in a Firm." Manuscript, Univ. Toronto.
Kaplan, Steven N., Mark M. Klebanov, and Morten Sorensen. 2012. "Which CEO Characteristics and Abilities Matter?" *J. Finance* 67 (3): 973–1007.
Kaplan, Steven N., and Morten Sorensen. 2016. "Are CEOs Different? Characteristics of Top Managers." Manuscript.
Kotter, John P. 1990. "What Leaders Really Do." *Harvard Bus. Rev.* 68 (3): 103–11.
La Porta, Rafael, Florencio Lopez-De-Silanes, and Andrei Shleifer. 1999. "Corporate Ownership Around the World." *J. Finance* 54 (2): 471–517.
La Porta, Rafael, Florencio Lopez-de Silanes, Andrei Shleifer, and Robert Vishny. 2000. "Investor Protection and Corporate Governance." *J. Financial Econ.* 58 (1–2): 3–27.
Lazear, Edward P., Kathryn L. Shaw, and Christopher T. Stanton. 2015. "The Value of Bosses." *J. Labor Econ.* 33 (4): 823–61.
Lippi, Francesco, and Fabiano Schivardi. 2014. "Corporate Control and Executive Selection." *Quantitative Econ.* 5 (2): 417–56.
Malmendier, Ulrike, and Geoffrey Tate. 2005. "CEO Overconfidence and Corporate Investment." *J. Finance* 60 (6): 2661–700.
———. 2009. "Superstar CEOs." *Q.J.E.* 124 (4): 1593–638.
Mintzberg, Henry. 1973. *The Nature of Managerial Work.* New York: Harper & Row.
———. 1979. *The Structuring of Organizations.* Englewood Cliffs, NJ: Prentice-Hall.
Mullins, William, and Antoinette Schoar. 2016. "How Do CEOs See Their Role? Management Philosophy and Styles in Family and Non-Family Firms." *J. Financial Econ.* 119 (1): 24–43.

Porter, Michael E., and Nitin Nohria. 2018. “What Do CEOs Actually Do?” *Harvard Bus. Rev.* 2018 (July–August).
Roberts, John. 2004. *The Modern Firm: Organizational Design for Performance and Growth.* Oxford: Oxford Univ. Press.
Syverson, Chad. 2011. “What Determines Productivity?” *J. Econ. Literature* 49 (2): 326–65.
Tervio, Marko. 2008. “The Difference That CEOs Make: An Assignment Model Approach.” *A.E.R.* 98 (3): 642–68.
Tipping, Michael E., and Christopher M. Bishop. 1999. “Probabilistic Principal Component Analysis.” *J. Royal Statis. Soc. B* 61 (3): 611–22.
Urban, Daniel. 2019. “The Effects of Culture on CEO Power: International Evidence from Executive Turnover.” *J. Banking and Finance* 104:50–69.
Van den Steen, Eric. 2010. “On the Origin of Shared Beliefs (and Corporate Culture).” *RAND J. Econ.* 41 (4): 617–48.