

NBER WORKING PAPER SERIES

CEO BEHAVIOR AND FIRM PERFORMANCE

Oriana Bandiera
Stephen Hansen
Andrea Prat
Raffaella Sadun

Working Paper 23248
<http://www.nber.org/papers/w23248>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138
March 2017, Revised September 2017

This project was funded by Columbia Business School, Harvard Business School and the Kauffman Foundation. We are grateful to Morten Bennedsen, Robin Burgess, Wouter Dessein, Bob Gibbons, Rebecca Henderson, Ben Hermalin, Paul Ingram, Amit Khandelwal, Nicola Limodio, Michael McMahon, Antoinette Schoar, Daniela Scur, Steve Tadelis and seminar participants at ABD Institute, Bocconi, Cattolica, Chicago, Columbia, Copenhagen Business School, Cornell, the CEPR Economics of Organization Workshop, the CEPR/IZA Labour Economics Symposium, Edinburgh, Harvard Business School, INSEAD, LSE, MIT, Munich, NBER, Oxford, Politecnico di Milano, Princeton, Science Po, SIOE, Sydney, Stanford Management Conference, Tel Aviv, Tokyo, Toronto, Uppsala, and Warwick for useful suggestions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Oriana Bandiera, Stephen Hansen, Andrea Prat, and Raffaella Sadun. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

CEO Behavior and Firm Performance
Oriana Bandiera, Stephen Hansen, Andrea Prat, and Raffaella Sadun
NBER Working Paper No. 23248
March 2017, Revised September 2017
JEL No. J22,J24,M12,O4

ABSTRACT

We measure the behavior of 1,114 CEOs in six countries parsing granular CEO diary data through an unsupervised machine learning algorithm. The algorithm uncovers two distinct behavioral types: “leaders” and “managers”. Leaders focus on multi-function, high-level meetings, while managers focus on one-to-one meetings with core functions. Firms with leader CEOs are on average more productive, and this difference arises only after the CEO is hired. The data is consistent with horizontal differentiation of CEO behavioral types, and firm-CEO matching frictions. We estimate that 17% of sample CEOs are mismatched, and that mismatches are associated with significant productivity losses.

Oriana Bandiera
London School of Economics
o.bandiera@lse.ac.uk

Stephen Hansen
University of Oxford
Department of Economics
Manor Road Building
Manor Road
Oxford OX1 3UQ
United Kingdom
stephen.hansen@economics.ox.ac.uk

Andrea Prat
Columbia Business School
3022 Broadway, Uris 624
New York, NY 10027-6902
andrea.prat@columbia.edu

Raffaella Sadun
Harvard Business School
Morgan Hall 233
Soldiers Field
Boston, MA 02163
and NBER
rsadun@hbs.edu

1 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, Bennesen et al. 2007, Kaplan et al. 2012), is that the identity of the CEO matters for firm performance. What is less known is what different CEOs do differently, and whether and how differences in CEO behavior have significant economic implications.

Scholars have approached this question in two ways. At one end of the spectrum, an influential cluster of studies starting with Mintzberg (1973) have focused on the measurement of the actual behavior of executives. They do so by “shadowing” CEOs in real time through personal observation. While shadowing exercises have revealed a wealth of information on the nature of the managerial job, they are based on small and selected samples and, as such, are difficult to generalize. 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 (Dessein and Santos (2016); Hermalin (1998, 2007)).¹

This paper develops a new methodology to bring back quantitative measurements of CEO behavior by scaling up the shadowing methods to large samples. Our aim is to advance knowledge on questions that have managerial behavior at their core, through new, large-scale and internationally comparable evidence on what CEOs actually do.

This approach involves two primary challenges: a) how to shadow a large number of CEOs, and b) how to aggregate granular information on their activities into a summary measure that has a consistent meaning across subjects. 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. 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, UK and the US. Overall, we collect data on 42,233 activities covering an average of 50 working hours per CEO. We record the same five features for each activity: its type (e.g. meeting, plant/shop-floor visits, business lunches etc.), planning horizon, number of participants involved, number of different functions, and the participants’ function (e.g. finance, marketing, clients, suppliers, etc.).² We find that CEO

¹Hermalin (1998) and Hermalin (2007) propose a rational theory of leadership, whereby the leader possesses private non-verifiable 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, less influence activities, and more communication. Bolton et al. (2013) highlights 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.

²In earlier work (Bandiera et al. 2017) 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).

behavior differs considerably along all five features. In particular, while the majority of CEOs spend most of their time in meetings, they differ in the extent to which their focus is on firms' employees vs. outsiders, and within the former, whether they mostly interact with high-level executives vs. production employees. CEOs also differ in how they organize these interactions in terms of duration, number of people involved, number of functions these people represent and planning horizon. We also show that these dimensions of time use are correlated so that, for instance, CEOs who focus on production also tend to have short, one-to-one meetings.

CEO diaries yield a wealth of information that 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, and generates a one-dimensional behavior index that represents a CEO as a convex combination of the two pure behaviors. The algorithm finds the combination of features that best differentiates between the sample CEOs. The first of the two pure behaviors is associated with more time spent with employees involved with production activities, and 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 (1999). 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.

The scalar behavior index can be used to investigate a range of questions about the causes and consequences of CEO behavior. A natural starting point is to study the correlation between CEO behavior and firm performance, which we do by merging the behavior index with firm balance sheet data. We find that leader CEOs are more likely to be found in larger and more productive firms. The correlation is economically and statistically significant: increasing the CEO behavior index by one standard deviation is associated with an increase of 7% in sales controlling for labor, capital, and other standard firm-level variables.

The correlation between CEO behavior and firm performance can be interpreted in three ways: (i) CEO behavior simply reflects firm heterogeneity correlated with performance; (ii) CEO behavior affects performance and CEOs are vertically differentiated, i.e. leader CEOs improve performance regardless of the type of firms they work for, but they are scarce; (iii) CEO behavior affects performance and CEOs are horizontally differentiated, i.e. firm performance is a function of the correct firm-CEO match, but there are CEO-firm matching frictions and some firms are run by CEOs with a firm-inappropriate behavior, thus causing a performance loss.

We use the subset of firms for which we have productivity data before and after the appointment

of the current CEO to investigate the first alternative, i.e. performance differentials across CEOs simply being driven by firm heterogeneity. This exercise reveals two results. First, firm performance pre-appointment is not correlated with CEO type: firms that hire leader CEOs have the same productivity growth as firms that hire manager CEOs before the CEO appointment. Second, the hire of a leader CEO is associated with a significant change in firm productivity relative to the pre-appointment period, which emerges gradually and increases over time. Overall, these results are in contrast with the idea that CEO behavior is merely a reflection of differential pre-appointment trends or firm-level, time-invariant differences in performance.

We then turn to the second class of alternatives—that is, that CEO behavior affects firm performance, either via vertical or horizontal differentiation in CEO behavior. In the absence of exogenous variation that would allow us to tell these alternatives apart, we develop a simple model of CEO-firm assignment that encompasses both possibilities, and estimate it 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 incorporates pure vertical differentiation—where all firms need leaders but leaders are scarce, and hence firms that end up with leaders perform better—and horizontal differentiation—where some firms need managers and others leaders, but matching frictions imply that some of the firms that need leaders end up with managers. The model estimation is consistent with horizontal differentiation of CEOs with matching frictions. More specifically, 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.

The main contribution of this study is a new method to measure CEO behavior in large samples, with an approach that can be easily replicated in a variety of contexts. To the best of our knowledge, the largest CEO shadowing exercise besides ours is still Mintzberg (1973) and it comprised five CEOs.³ Our approach to the measurement of managerial behavior can be used to address questions that have been core to the field of organizational economics, but has so far been subjected to limited, if none, empirical investigation. For example, the coordinating role of entrepreneurs has been of interest to economics since Coase (1937), and Roberts (2006) emphasizes the critical role played by leadership behavior in complementing the organizational design tasks of general managers.⁴

The paper is also related to a growing literature documenting the role of management processes on firm performance (Bloom and Van Reenen 2007 and Bloom et al. 2016). The correlation

³Other authors have performed shadowing exercises of executives below the CEO level (For instance, Kotter (1999) 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.

⁴More recently, Cai and Szeidl (2016) have shown that exogenous shifts in the interactions between an entrepreneur and his/her peers is associated with large increases in firm revenues, productivity and managerial quality.

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 et al. (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. Differently from 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.

Finally, the paper is related to the literature that studies CEO traits such as skills and personality (Kaplan et al. (2012), Kaplan and Sorensen (2016) Malmendier and Tate (2005) and Malmendier and Tate (2009)) or self-reported management styles Mullins and Schoar (2013). 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 needs to rely on third party evaluations, self reports or indirect proxies for individual preferences.

The paper is organized as follows. Section 2 describes the data and the machine learning algorithm. Section 3 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 4 interprets the correlation between CEO behavior and firm performance by estimating a simple CEO-firm assignment model. Section 5 concludes.

2 Measuring CEO Behavior

2.1 The Sample

We drew the sampling frame randomly from the set of firms classified in the manufacturing sector in the accounting database ORBIS, an extensive commercial data set produced by the company Bureau Van Dijk that contains company accounts for more than 200 million companies around the world.

The sample covers CEOs in six of the world’s ten 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 authority in charge of the organization who has executive powers and reports to the board of directors. While titles may differ across countries (e.g. Managing Director in the UK), we refer to them 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 prior to the survey.⁵ This yielded a sample of 6,527 firms in 32 two-digits SIC 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 with the opportunity to learn about their own time use with a personalized time use analysis, to be delivered after the data had been collected.⁶

Of the 6,527 firms included in the screened ORBIS sample, 1,114 (17%) participated in the survey,⁷ of which 282 are in Brazil, 115 in France, 125 in Germany, 356 in India, 87 in the UK and 149 in the US.

Table A.1 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 (ROCE) (see the Appendix for details). 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. The average tenure is 10 years, with a standard deviation of 9.55 years.⁸ 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).

2.2 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 their personal assistant (PA), or with the CEO himself (43% of the cases). We record diaries over a week that we chose based on an arbitrary ordering of firms. Enumerators collected daily information on all the activities the CEO planned

⁵We 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 were still active. We could find CEO contact details for 7,744 firms and, of these, 1,217 later resulted not to be eligible. 310 of the 1,217 could not be contacted to verify eligibility before the project ended. Among this set 1,009 were located in Brazil; 896 in Germany; 762 in France; 1,429 in India; 1,058 in the UK; 1,372 in the U.S. 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 which we could only hire one for each country. The sample construction is described in detail in the Appendix.

⁶The report was delivered two 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.

⁷This figure is at the higher end of response rates for CEO surveys, which range between 9% and 16% (Graham et al. (2013)). 1,131 CEOs agreed to participate but 17 dropped out before the end of the data collection week for personal or professional contingencies that limited our ability to reach them by phone.

⁸The 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 a family, 23% by disperse shareholders, 9% by private individuals, and 7% by private equity. Ownership data is 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 to be owned by an entity if this controls at least 25.01% of the shares; if no single entity owns at least 25.01% of the share the firm is labeled as “Dispersed shareholder”.

to undertake that day as well as those actually done.⁹ 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 shows a screenshot of the survey tool.¹⁰ The survey collects information on all the 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 duration, equivalent to 225,721 15-minute blocks, 90% of which cover work activities.¹¹ The average CEO has 202 15-minute time blocks, adding up to 50 hours per week on average.

2.3 The Data

Figure 1, Panel 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 be 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), Mintzberg (1973) and Mintzberg (1979).¹²

The richness and comparability of the time use data allows for a much more detailed description of these interactions relative to prior studies. We use as primary features of the activities their: (1) type (e.g. meeting, lunch, etc.); (2) duration (30m, 1h, etc.); (3) whether planned or unplanned; (4) number of participants; (5) functions of participants, divided between employees of the firms, which we define as “insiders” (finance, marketing, etc.), and non-employees, or “outsiders” (clients, banks, etc.). Panel B shows most of this interactive time is spent with insiders. This suggests that most CEOs chose to direct their attention primarily towards 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. Panel 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. Panel D shows most CEOs engage in planned activities with a duration of longer than one hour with

⁹70% of the CEOs 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.

¹⁰The survey tool can also be found online on www.executivetimeuse.org.

¹¹The non-work activities cover personal and family time during business hours.

¹²Mintzberg (1973), for example, documents that in a sample of five managers 70-80% of managerial time is spent communicating.

a single function. There is no marked average tendency towards 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 at the bottom quartile devote just over 40% of the time to meetings whereas those at the top quartile reach 65%; CEOs at the 3rd quartile devote over three times more time to production than their counterparts at 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 under explored phenomenon.

The data also shows that systematic patterns of correlation across these distributions, as we show in the heat map of Figure 2. This exercise reveals significant and intuitive patterns of co-occurrence. For example, CEOs who do more plant visits spend more time with employees working on production and suppliers. The data also shows that they tend to meet these functions one at the time, rather than in multi-functional 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, 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.

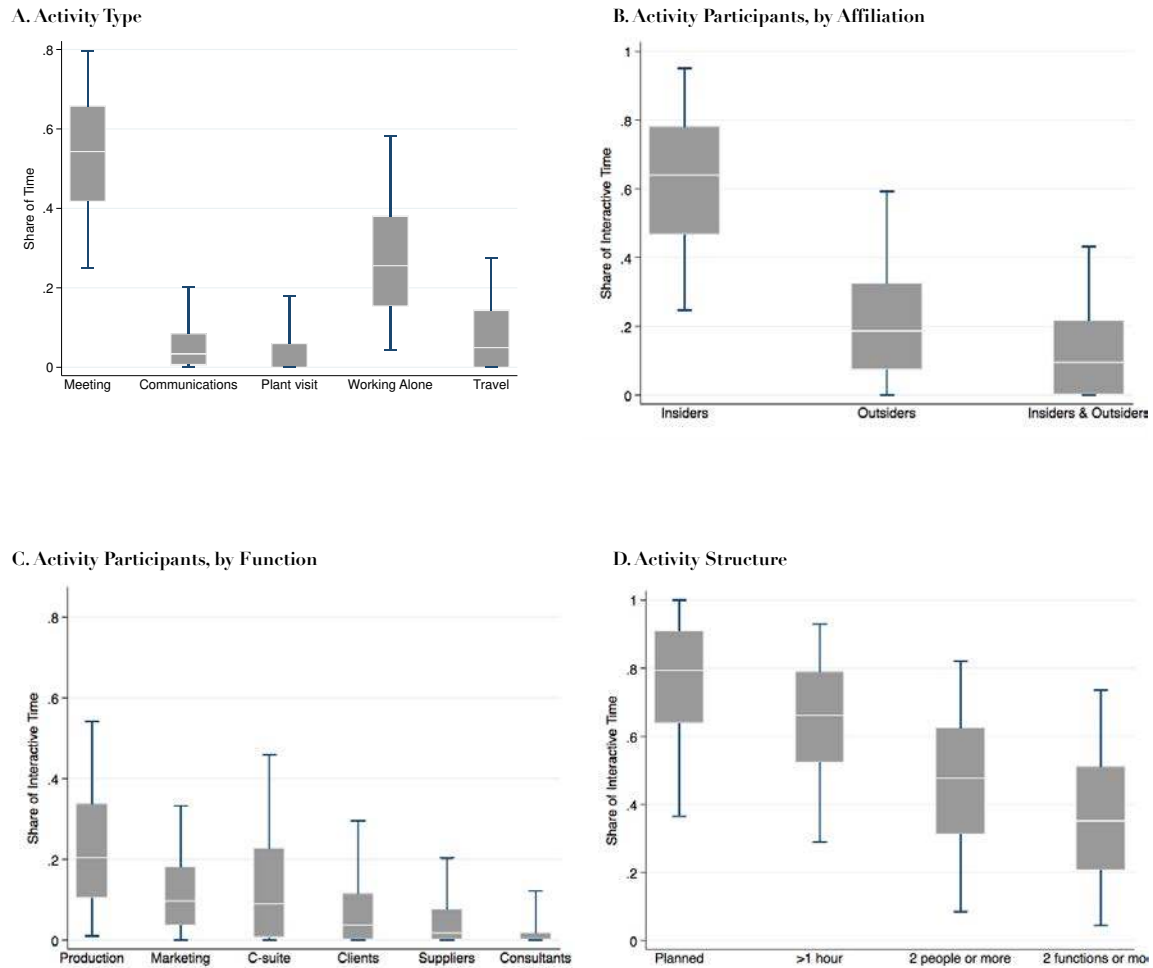
The activities also appear to largely reflect conscious planning vs. mere reactions to external contingencies. To assess this point, we asked whether each activity was undertaken in response to an emergency: only 4% of CEOs’ time was devoted to activities that were defined as emergencies. Furthermore, 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 cancelled.

2.4 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 Figure 2—that the high-dimensional raw activity data is 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.

Figure 1: CEO Behavior: Raw Data



Notes: 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 CEOs level.

Figure 2: CEO Behavior: Correlations

	Meeting	Plant visits	Communications	Planned	More 1 participant	Insiders	Outsiders	Insiders & Outsiders	C-suite	Production	Marketing	Clients	Suppliers	Consultants		
Meeting	1															
Plant visit	-0.5218	1														
Communications	-0.4673	-0.1647	1													
Planned	0.2009	-0.1169	0.0339	1												
More 1 participant	0.1056	0.0032	0.0921	0.2883	1											
More than 1 function	0.1816	-0.1736	0.1289	0.2043	0.511	1										
Insiders	-0.0486	0.0587	0.1632	-0.0941	0.0329	0.0018	1									
Outsiders	0.034	-0.057	-0.1877	0.0337	-0.1827	-0.406	-0.7052	1								
Insiders & Outsiders	0.0975	-0.0977	0.0096	0.1144	0.2122	0.5444	-0.482	-0.2224	1							
C-suite	-0.0363	-0.1394	0.2441	0.1147	0.1371	0.1371	0.3511	-0.3252	-0.2917	-0.1092	1					
Production	-0.1565	0.4114	-0.0823	-0.1157	0.0246	-0.1387	0.3455	-0.2917	-0.1092	0.0787	-0.303	1				
Marketing	0.0926	-0.1456	0.0645	-0.0228	0.0129	0.1662	0.1931	-0.2684	0.0787	-0.1882	-0.1447	-0.134	1			
Clients	-0.0945	0.0028	-0.028	0.0134	-0.1714	-0.1389	-0.4156	0.4275	0.0729	-0.1789	-0.134	-0.0655	-0.0723	1		
Suppliers	-0.0358	0.1089	-0.1622	-0.0381	-0.1702	-0.1703	-0.3264	0.3492	0.0384	-0.2192	0.0214	0.0444	-0.0723	0.0444	1	
Consultants	0.0387	-0.0483	-0.0676	-0.0182	-0.0817	-0.0251	-0.2367	0.2154	0.0931	-0.0344	-0.1429	-0.0746	-0.0606	-0.0085	-0.0606	1

Notes: Each cell reports the correlation coefficient between the variables listed in the row and column. Each variable indicates the share of time spent by CEOs in activities denoted by the specific feature (this is the same data used to generate Figure 1). Cells are color coded so that: dark (light) gray=positive (negative) correlation, reject H0: correlation=0 with p=.10 or lower, white= cannot reject H0: correlation=0.

Methodology

To reduce the dimensionality of the data we use latent Dirichlet allocation (LDA) (Blei et al., 2003), a hierarchical Bayesian factor model for discrete data.¹³ Simpler techniques like 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 which provides a complete probabilistic description of time-use patterns.¹⁴ 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 figure 2. 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 et al., 2007; Gross and Manrique-Vallier, 2014). Ours is the first application to survey data in the economics literature that we are aware of.

To be more concrete, suppose 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 $X \equiv \{x_1, \dots, x_A\}$ be the set of activities. A *pure behavior* k is a probability distribution β^k over X that is common to all CEOs.¹⁵

In our baseline specification, we focus on the simplest possible case in which there exist only two possible pure behaviors: β^0 and β^1 , and discuss alternative approaches and sensitivity of the main results using models with more than two pure behaviors in Section 4. In this simple case, the *behavior* of CEO i is given by a mixture of the two pure behaviors according to weight $\theta_i \in [0, 1]$, thus the probability that CEO i generates activity a can lie anywhere between β_a^0 and β_a^1 .¹⁶ We refer to the weight θ_i as the *behavior index* of CEO i .


¹³LDA 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 can be fragile in the presence of highly correlated covariates, which makes projecting them onto a latent space prior to regression analysis attractive.

¹⁴Tipping 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, though, our survey data is not Gaussian, so PCA lacks an obvious statistical interpretation in our context.

¹⁵Importantly, 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.

¹⁶In contrast, in a traditional clustering model, each CEO would be associated with one of the two pure behaviors, which corresponds to restricting $\theta_i \in \{0, 1\}$.

Figure 3: Data Generating Process for Activities with Two Pure Behaviors



Notes: 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 θ_i . The observed activity for a unit of time is then drawn from the distribution over activities that the pure behavior defines.

Figure 3 illustrates the LDA procedure. For each activity of CEO i , one of the two pure behaviors is drawn independently given θ_i . Then, given the pure behavior, an activity is drawn according to its associated distribution (either β^0 or β^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 $\prod_i \prod_a \chi_i^{n_{i,a}}$.¹⁷ While in principle one can attempt to estimate β and θ via direct maximum likelihood or the EM algorithm, in practice the model is intractable due to the large number of parameters that need to be estimated (and which grow linearly in the number of observations). LDA overcomes this challenge by adopting a Bayesian approach, and placing Dirichlet priors on the β and θ_i terms. For estimating posteriors we follow the Markov Chain Monte Carlo (MCMC) approach of Griffiths and Steyvers (2004).¹⁸ Here we discuss the estimated object of interest, which are the two estimated pure behaviors $\hat{\beta}^0$

¹⁷While 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 θ_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.

¹⁸We set a uniform prior on θ_i —i.e. a symmetric Dirichlet with hyperparameter 1—and a symmetric Dirichlet with hyperparameter 0.1 on β^k . This choice of hyperparameter promotes sparsity in the pure behaviors. Source code for implementation is available from <https://github.com/sekhansen>.

and $\hat{\beta}^1$, as well as the estimated behavioral indices $\hat{\theta}_i$ for every CEO $i = 1, \dots, N$.

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 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 in the main results.

Estimates

To illustrate differences in estimated pure behaviors, in Figure 4 we order the elements of X according to their estimated probability in $\hat{\beta}^0$ 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.¹⁹ To construct a formal test of whether the observed differences between pure behaviors are consistent with a model in which there 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 this simulated data to estimate the LDA model with two pure behaviors as in our baseline analysis, and find systematically less difference between pure behaviors than in our actual data (for further discussion see the Appendix).

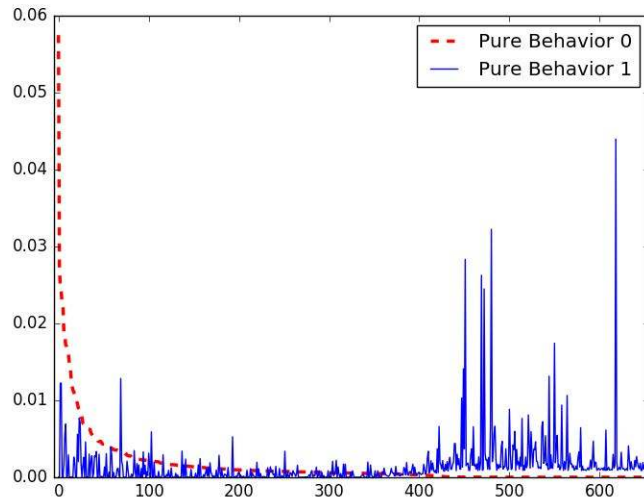
The two pure behaviors we estimate represent extremes. As discussed above, individual CEOs generate activities according to the behavioral index θ_i that gives the probability that any specific activity is drawn from pure behavior 1. Figure 5 plots both the frequency and cumulative distributions of the $\hat{\theta}_i$ 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 5 shows, though, the bulk of CEOs lies away from these extremes, where the distribution of the index is essentially uniform.

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 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 one hour; (2) are planned;

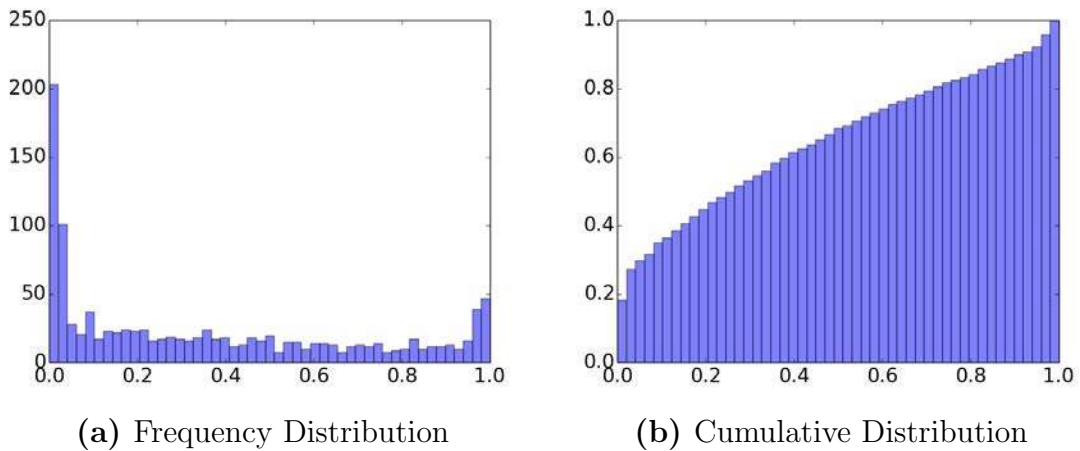
¹⁹Table B.3 displays the estimated average time that CEOs spend with the different categories in figure 1 derived from the estimated pure behaviors and CEO behavioral indices. Reassuringly, there is a tight relationship between the shares in the raw data and the estimated shares.

Figure 4: Probabilities of Activities in Estimated Pure Behaviors



Notes: The dotted line plots the estimated probabilities of different activities in pure behavior 0, the solid line plots the estimated probabilities of different activities in pure behavior 1. The 654 different activities are ordered left to right in descending order of their estimated probability in pure behavior 0.

Figure 5: CEO Behavior and Index Distribution



Notes: The left-hand side plot displays the number of CEOs with behavioral indices in each of 50 bins that divide the space $[0,1]$ evenly. The right-hand side plot displays the cumulative percentage of CEOs with behavioral indices lying in these bins.

Table 1: Most important behavioral distinctions in CEO time use data

Feature	X times less likely in Behavior 1	Feature	X times more likely in Behavior 1
Plant Visits	0.11	Communications	1.9
Just Outsiders	0.50	Outsiders + Insiders	1.90
Production	0.50	C-suite	34.00
Suppliers	0.30	Multifunction	1.50

Notes: 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.

(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 PCA analysis explains 35% of the variance in this feature space and places a positive weight on all dimensions 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 (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 relative to the ones described above.

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 1 displays the ratios of these marginal distributions (always expressed as the ratio of the probability for pure behavior 1 relative to 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 variety of functions.²⁰ Higher values of the CEO behavior index $\hat{\theta}_i$ will thus correspond to a greater intensity of these latter types of interactions.

Differences in the CEO behavior index are systematically correlated with firm, industry and country characteristics. Starting with firm characteristics, Figure D.2 in the Appendix shows that the behavior index takes higher values in larger firms (as proxied by number of employees, and the listed status dummies), and in organizational contexts in which the CEO is more likely to share managerial responsibilities with other peers (that is, when the firm is a multinational or part of a larger corporate group, and when there is a COO). The index is significantly lower in cases firms owned and managed by a family CEO. The variation across industries shows higher values of the index 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 et al. (2003)) and greater R&D intensity (defined as industry business R&D divided by industry employment from NSF data). The variation across countries—shown in Figure D.2—shows a wide distribution of the CEO behavior index across all countries, but significantly lower averages in Brazil and India.²¹ The data also shows that CEO behavior varies systematically with specific CEO characteristics, namely CEO skills (college of MBA degree) and experience abroad. Note, however, that the correlation between CEO behavior and firm characteristics (and firm size in particular) remains large significant even when we control for CEO traits. This points to the fact that observable CEO characteristics do not fully capture differences in CEO behavior.²² In contrast, we find no correlation between the behavior index and the CEOs’ job satisfaction (on a scale from 1 to 5): the average value of the CEO index for those above (below) the median level of job satisfaction is 0.42 (0.44), the p-value of the test of zero difference is 0.32.

While the labeling of the two pure behaviors is arbitrary, the distinctions between pure behavior 0 and pure behavior 1 map into 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 (1999). This defines management primarily as monitoring and implementation tasks, i.e. “setting up systems to ensure that plans are implemented precisely and efficiently.” In contrast, leadership is needed to create organizational alignment, and requires significant investment in communication

²⁰We 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.

²¹While most of the econometric analysis will exploit within country and within industry variation in the behavior index, we will return to the possible drivers and implications of these cross country differences for firm and country productivity in Section 4.2.

²²To the extent that optimal CEO behavior varies according to firms’ needs, and that boards can observe the CV of a CEO but not his actual behavior, this may lead to mismatches of CEOs to firms. We formalize a model of how these mismatches may arise in equilibrium in Section 4.

across a broad variety of constituencies.²³ This interpretation is corroborated by the fact that leadership behavior—that is, higher values of $\hat{\theta}_i$ —is more prevalent in larger and more skill intensive firms, in which CEO time may be more effective in creating alignment relative to impersonal monitoring and/or alternative contractual approaches. For example, Drucker (1967) mentions the importance of personal meetings in the management of knowledge workers, arguing that the “[...] relationships with other knowledge workers are especially time consuming. Whatever the reason—whether it is absence of or 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 superiors 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 when knowing whether he is doing a satisfactory job or not.”²⁴

From now onwards we will 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 performance, and provide a simple framework to assess the possible reasons behind the correlation.

3 CEO Behavior and Firm Performance: Evidence

An open question is whether our index of CEO behavior, built to capture the main differences between CEOs without any information on firm performance, correlates with performance. To investigate this issue, 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 firm in the CEO sample.²⁵

²³More 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. [...] Aligning invariably involves talking to many more employees than organizing does. The target population may involve not only a manager’s subordinates but also bosses, peers, staff in other parts of the organization.”

²⁴Similarly, Mintzberg (1979) emphasizes the importance of informal communication activities in the coordination of complex organizations. Mintzberg (1979) 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.”

²⁵Of these: 41 did not report sales and employment information; 64 were dropped when removing extreme values 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 prior to the survey.

3.1 Cross-Sectional Correlations

Productivity

We start by analyzing whether CEO behavior correlates with productivity, a key metric of firm performance (Syverson, 2011). We follow a simple production function approach and estimate by OLS a regression of the form:

$$y_{ifst} = \alpha \hat{\theta}_i + \delta^E e_{ft} + \delta^K k_{ft} + \delta^M m_{ft} + \zeta_t + \eta_s + \varepsilon_{ifst} \quad (1)$$

where y_{ifst} is the log sales (in constant 2010 USD) 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 logarithm of the number of firm employees and, when available, capital and materials. ζ_t and η_s are period and three digits SIC sector fixed effects, respectively.

The performance data includes up to three most recent years of accounting data pre-dating 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 (Appendix Table D.2, column 2). We include country and year dummies throughout, as well as a set of interview noise controls.²⁷ The coefficient of interest is α , which measures the correlation between log sales and the CEO behavior index. Recall that higher values of the index imply a closer similarity with the pure behavior labeled as “leader”.

Column 1, Table 2 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 α 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 one 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

²⁶We do not condition on the CEO being in office for at least three 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 3 years of accounting for 58% of the sample, 2 years for 24% and 1 year for the rest of firms.

²⁷These are a full set of dummies to denote the week in the year in which the data was collected, a reliability score assigned by the interviewer at the end of the survey week, a dummy taking value one if the data was collected through the PA of the CEO, rather than the CEO himself, and interviewer dummies. All columns 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. Since the data is averaged over three years, year dummies are set as the rounded average year for which the performance data is available.

effect of a one 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 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 0.279). 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.

Controlling for other Firm or CEO observables

Since the CEO behavior index varies according to specific firm and managerial characteristics potentially correlated with firm performance (see Figures D.1 and D.2), in Table D.2 we investigate the robustness of the results to the inclusion of these variables. We find that family CEO firms tend to be on average less productive than firms led by professional managers, while listed firms are firms with a COO tend to be on average more productive. Productivity is higher in firms where CEOs have experience outside their home country, are older and work longer hours. Including these other firm or CEO observables, however, leaves the magnitude and significance of the CEO behavior index practically unchanged.

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 using the World Management Survey (Bloom et al. 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 one standard deviation change in the management index is associated with a 0.06 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 2 using the sample of 156 firms for which we could match the management and CEO behavior data with accounting information. The coefficients

²⁸To make this comparison we multiply the coefficient of the CEO behavior index in column 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 (coefficient of 0.387 times the standard deviation of log capital of $1.88 = 0.73$).

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

³⁰See Appendix Table D.3 for details. Bender et al. (2016) 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 due to data confidentiality it excludes the highest paid individuals, who are likely to be CEOs. This is the first time that data on middle level management practices and CEO behavior are combined.

imply that a 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 the CEO behavior index is distinct from other, firm-wide, management differences.

Profits

Column 6 analyzes the correlation between CEO behavior and profits per employee. This allows us to assess whether CEOs capture all 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 one 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 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 \$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.

More than two pure CEO behaviors

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 when it comes to the correlation between CEO behavior and performance. To investigate this issue, we followed an alternative approach, in which the optimal number of pure behaviors is chosen according to a statistical criterion. To implement this approach, we estimate LDA on randomly drawn training subsets of the data, and then use

³¹When we do not control for the management (CEO) index, the coefficient on the CEO (management) index is 0.606 (0.207) significant at the 5% level. The magnitude of the coefficient on the management index is similar to the one reported by Bloom et al. (2016) in the full management sample (0.15). The correlation between CEO behavior and management practices is driven primarily by practices related to operational practices, rather than HR and people-related management practices. See Appendix Table D.3 for details.

the estimated parameters to predict the held-out data. The details of this exercise are discussed in Appendix.

This approach shows that a model with eleven pure behaviors is best at prediction. However, we also find that, among these eleven pure behaviors, the pure behavior with the largest correlation with productivity is the most dissimilar to pure behavior 0 (manager) used in the simple $K = 2$ model. That is, in spite of its simplicity, the model with two pure behaviors is able to capture the salient correlations with firm performance that are found in the more complex $K = 11$ case. This is reassuring, given the focus of the paper on the relationship between CEO behavior and firm performance.³²

Robustness using shares of time and standard dimensionality reduction techniques

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

3.2 Firm Performance Before and After the CEO Appointment

One possible mechanism underpinning the observed correlation between CEO behavior and firm performance is that firms with traits that lead to higher productivity are more likely to hire leaders. To provide evidence on this hypothesis we use firm performance in the years pre-dating the CEO appointment. We use this pseudo-panel data to ask two complementary questions. First, do differences in productivity trends *before* the CEO appointment predict the type of CEO that is eventually hired by the firm? This provides us with a rough empirical assessment of whether time varying firm heterogeneity pre-dating the appointment of the CEO is associated with differences in CEO behavior. Second, is the CEO behavior index associated with *changes* in productivity relative to the period preceding the appointment of the CEO? This exercise provides evidence on whether time invariant firm heterogeneity is the sole driver of the cross sectional performance differentials across CEOs.

³²Other applications of LDA to the time use data may well require going for a greater number of pure styles. The tradeoff between interpretability (which favors a small number of pure behaviors) and goodness-of-fit (which favors a greater number) is well known in the unsupervised learning literature. See, for example, Chang et al. (2009).

Table 2: CEO behavior and Firm Performance

Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)
	Log(sales)					Profits/Emp
CEO behavior index	0.343*** (0.108)	0.227** (0.111)	0.322*** (0.121)	0.641** (0.279)	0.505** (0.235)	10.027*** (3.456)
log(employment)	0.889*** (0.040)	0.555*** (0.066)	0.346*** (0.099)	0.339** (0.152)	0.804*** (0.075)	-0.284 (0.733)
log(capital)		0.387*** (0.042)	0.188*** (0.056)	0.194* (0.098)		
log(materials)			0.447*** (0.073)	0.421*** (0.109)		
Management					0.187** (0.074)	
Number of 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 & m	with k & m, listed	with management score	with profits, listed

Notes: *** (**) (*) denotes significance at the 1%, 5% and 10% level, respectively. 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. The sample in Columns 1 includes all firms with at least one year with both sales and employment data. Columns 2, 3 and 4 restrict the sample to firms with additional data on capital (column 2), capital and materials (columns 3 and 4). The sample in column 4 is restricted to listed firms. "Firm size" is the log of total employment in the firm. All columns include a full set of country and year dummies, three digits 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 was collected, a reliability score assigned by the interviewer at the end of the survey week, a dummy taking value one if the data was collected through the PA of the CEO, rather than the CEO himself, and interviewer dummies. All columns 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.

To implement this approach, we restrict the sample to the 204 firms that have accounting data within a five-year interval *both* before and after CEO appointment.³³ To start, Column 1 shows that these firms are representative of the larger sample in terms of the correlation between the CEO behavior index and performance. The correlation is 0.362 (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 -0.095 and is not precisely estimated.

We first test whether productivity trends before appointment can predict the type of CEO that is eventually hired by the firm. Column 2 shows that this is not the case—in the pre-appointment period, firms that eventually appoint a leader CEO have similar productivity trends relative to 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:

$$y_{ft} = \alpha A_t + \beta A_t \hat{\theta}_i + \delta^E e_{ft} + \zeta_t + \eta_f + \varepsilon_{it} \quad (2)$$

Where t denotes whether the time period refers to the 5 years before or after the appointment of the CEO. Similarly to the results shown in Table 2, inputs and outputs are aggregated across the two different subperiods, before and after CEO appointment. η_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 β , which measures whether firms that eventually appoint 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.³⁴

Column 3 shows that the coefficient β is positive and significant (coefficient 0.130, standard error 0.057). Given this coefficient, the within firm change in productivity after the CEO appointment is -0.05, 0 and 0.07 log points for values of the CEO index that are, respectively, at the 10th, 50th and 90th percentiles of the distribution of the CEO behavior index.³⁵ In column 4 we provide more detail on the nature of the correlation between CEO behavior and performance by splitting the post period into two sub periods: 1-2 and 3-5 years after appointment. The results suggest that

³³We 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 towards 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 towards firms located in France, Germany and the UK relative to the US. This is due to the fact that accounting panel data for US private firms—of which are sample is primarily composed of—is typically less complete relative to Europe.

³⁴Note that, since we do not know the behavior of the previous CEO, this is a lower bound on the effect of switching from managers to leader CEOs, since at least part of these firms would have had already a leader CEO before the current appointment.

³⁵The overall effect turns positive for values of the CEO behavior index greater than 0.42, which corresponds to the 62nd percentile of the distribution of the index.

the correlation materializes only three 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 102 firms whose current CEO had been in office for less than three years at the time of the survey—i.e., 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 cast doubts 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.

Taken together, these results indicate that the correlation between CEO behavior and firm performance cannot be solely driven by differences in time-invariant firm level unobservables, time-varying shocks to performance pre-dating the CEO appointment, or CEOs adapting their behavior to productivity shocks. 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 three years after the CEO is appointed, this would imply that corporate boards are able to predict performance three years in advance and to replace CEOs three years before the predicted performance effects actually occur.

4 CEO Behavior and Firm Performance: Interpretation

CEO behavior can affect firm performance in two different ways, both consistent with the empirical patterns discussed above. First, CEOs may be vertically differentiated, i.e. leader CEOs are always conducive to better performance, but their supply is limited, such that firms hiring managers are penalized. Second, CEOs may be horizontally differentiated, i.e. leader CEOs may be needed in certain circumstances, while manager CEOs may be best fit in others. In this case, the performance differentials across CEOs would arise in the presence of matching frictions, i.e. if manager CEOs are more likely to be mismatched relative to leader CEOs.

We develop a simple model of CEO-firm assignment that encompasses both possibilities and use it to test which of the two is a better fit for the data.

4.1 Model

We propose a simple assignment model of CEOs with different behaviors to firms with different traits. 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.

Table 3: CEO Behavior and Firm Performance Before and After CEO Appointment

Dependent variable: log(sales)	(1)	(2)	(3)	(4)	(5)
Sample	After appointment of the current CEO	Before appointment of the current CEO	Before and after current CEO	Before and after current CEO (after divided in 2 subperiods)	Before and after appointment of the current CEO (after divided in 2 subperiods) - firms with CEO tenure <= 3 years at time of the survey
Firm appoints high index CEO	0.258*** (0.073)				
Firm is in balanced sample	0.173* (0.102)				
Firm appoints high index CEO & is in balanced sample	-0.112 (0.130)				
Trend		0.007 (0.014)			
Trend*Firm appoints high index CEO		-0.010 (0.019)			
After CEO appointment			-0.022 (0.109)		
After CEO appointment*Firm appoints high index CEO			0.097** (0.041)		
After CEO appointment (1<=t<=2)				0.124 (0.106)	0.254 (0.238)
After CEO appointment (3<=t<=5)				0.195 (0.178)	0.172 (0.382)
After CEO appointment (1<=t<=2)*Firm appoints high index CEO				0.040 (0.049)	0.083 (0.091)
After CEO appointment (3<=t<=5)*Firm appoints high index CEO				0.177*** (0.065)	0.325** (0.132)
log(Employment)	0.891*** (0.039)	0.916*** (0.091)	0.797*** (0.072)	0.763*** (0.060)	0.745*** (0.065)
Observations	2,202	684	408	572	271
Number of firms	920	204	204	204	102
Spell averages	y	n	y	y	y
Firm fixed effects	n	y	y	y	y

Notes: *** (**) (*) denotes significance at the 1%, 5% and 10% level, respectively. All columns include the same controls used in 2, column 1. The sample in columns 1-4 includes the set firms with at least one year of non missing productivity data both in the 5 year interval before and after CEO appointment; in column 5 we also exclude CEOs that 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 column 1 is aggregated as in Table 3, column 1. Column 2 uses all available yearly data within 5 years before CEO appointment. In column 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 in two sub-periods (1<=t<=2; and 3<=t<=5 years after appointment). "After CEO appointment" = 1 for the cross section computed after CEO appointment. Errors clustered by industry in column 1, by firm in column 2, and by firm and before/after CEO appointment period in columns 3 to 5.

Set-up

CEO i can adopt one of two possible behaviors : $x_i = m$ (“manager”) and $x_i = l$ (“leader”). Once a CEO is hired, he decides how he is going to manage the firm that hired him. CEO i has a type $\tau_i \in \{m, l\}$. Type m prefers behavior m to behavior l . Namely, he incurs a cost of 0 if he selects behavior m and 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 $\tau_f \in \{m, l\}$. The output of firm f assigned to CEO i is

$$y_{fi} = \lambda_f + (I_{\tau_f=x_i}) \Delta \quad (3)$$

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 λ_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 firm heterogeneity issues discussed in Section 3.2 explicitly into the model and its subsequent estimation. The second component is specific to the behavior of the CEO. Namely, if the CEO's behavior matches the firm's type, then productivity increases by a positive amount Δ . 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

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

where \bar{w} is a 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 he 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, i.e. $w(y_{fi}) - I_{\tau_i \neq x_i}c$. 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 $B > \frac{c}{\Delta}$, the CEO will adapt to the firm's desired behavior, produce an output of $\lambda_f + \Delta$, and receive a total payoff of $\bar{w} + B\Delta - c$. If instead $B < \frac{c}{\Delta}$,

the CEO will choose $x_i = \tau_i$, produce output λ_f and receive a payoff \bar{w} . We think of B as a measure of governance. A higher B aligns CEO behavior with the firm's interests.

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 ϕ 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 more numerous 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, i.e. 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 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 that submitted it. If the type of the applicant corresponds to the type of the firm, the signal has value 1. 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.³⁶ 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.³⁷

Potential CEOs maximize their expected payoff, which is equal to the probability 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 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*

³⁶The 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.

³⁷See for example Gabaix and Landier (2008), Tervio (2008), Bandiera et al. (2015). An exception in the literature is Chade and Eeckhout (2016), who present a model in which agents' characteristics are only realized after a match is formed, which leads to a positive probability of mismatch in equilibrium.

productivity if at least one of the following conditions is met: (i) Neither CEO type is sufficiently scarce; or (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 won't do it). If governance is good ($B < \frac{c}{\Delta}$), 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 only apply 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,³⁸ 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. ■

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 a scarce-type CEO. 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.

Under Proposition 2, the economy under consideration does not achieve productive efficiency. As the overall pool of scarce-type CEOs is assumed to be insufficient to cover all firms that prefer that CEO type ($P \gg 1$), it would be possible to give all firms their preferred type and thus increase overall production.³⁹

³⁸Formally, this is given by the conditions: $B < \frac{c}{\Delta}$, and $\rho < \frac{\phi - \gamma}{\phi - \gamma \phi}$.

³⁹If 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.

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:

$$y_{.l} - y_{.m} = [s_l(\lambda_l + \Delta) + (1 - s_l)\lambda_m] - [s_m(\lambda_m + \Delta) + (1 - s_m)\lambda_l]$$

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 ($\lambda_l + \Delta$), weighted by the share of type l CEOs who are correctly assigned (s_l) plus the performance of misassigned type m firms (λ_m) weighted by the share of type l CEOs who are wrongly assigned ($1 - s_l$).

Simplifying and imposing the condition of proposition 2 by which all scarce type CEOs are correctly matched in equilibrium (that is, $s_l = 1$) yields:

$$y_{.l} - y_{.m} = s_m(\lambda_l - \lambda_m) + (1 - s_m)\Delta \quad (4)$$

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

- **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 $0 < s_m < 1$. Performance differentials capture the costs of the mismatches of type m CEOs (Δ), as well as firm heterogeneity.
- **Vertical differentiation in CEO behavior:** In this case, there is no demand for type m CEOs that is, $s_m = 0$. In this case, performance differentials reflect entirely the costs of the mismatches of type m CEOs (Δ).

In 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_m , Δ and $(\lambda_l - \lambda_m)$ best fit the data.

4.2 Model Estimation

Set-Up

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, Table 2, without country fixed effects, which we model separately. We denote the residual of firm f run by CEO i as $\hat{\varepsilon}_{if}$.⁴⁰ In line with the theory, we adopt the statistical model $\hat{\varepsilon}_{if} = \lambda_f + (I_{\tau_f=x_i}) \Delta + v_{if}$, where λ_f is a “baseline” productivity; $\tau_f \in \{m, l\}$ is the firm’s type; $x_i \in \{m, l\}$ is the CEO’s behavior; and Δ is the productivity difference between firms with the “right” CEO and firms with the “wrong” CEO behavior relative to firm needs.

To obtain an empirical proxy of x_i we use $\hat{x}_i = l$ whenever $I_{\hat{\theta}_i \geq 0.5}$. That is, we discretize the CEO behavior index using 0.5 as a cutoff, such that all CEOs above this threshold are classified as type l , and the rest as type m . While we treat \hat{x}_i as observed data, τ_f is a random variable.

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

As mentioned above, note that the model nests both pure vertical and pure horizontal differentiation. In the case of pure vertical differentiation $s_m = 0$; that is, all manager CEOs are misassigned. Vice versa, in the case of pure horizontal differentiation $s_m = 1$; 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. We also assume that $x_{c_f, l} = A + x_{c_f, m}$ 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: σ_{1, c_f}^2 (σ_{0, c_f}^2) is the standard deviation of residuals in an efficient (inefficient) CEO-firm pair.

Given these observations, the likelihood function can be written as:

⁴⁰To 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.

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

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

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 (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 0.706). 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.⁴¹

The estimate of Δ 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 $\frac{\exp(0.532)-1}{\exp(0.532)}$, or around 41%.

The estimated s_m is 0.744. To test whether the data are consistent with pure vertical differentiation, we impose the restriction $s_m = 0$ in (5), which a likelihood ratio test rejects with a p-value of 0.00202. The key underlying property of the data that lets us test $s_m = 0$ is that under this restriction leader CEOs uniformly outperform manager CEOs. We can reject this in favor of a mixture model with $s_m > 0$, since we observe a large fraction of manager CEOs whose performance is similar to that of leader CEOs. Also, note that once we reject $s_m = 0$, we must necessarily reject

⁴¹Note 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 (5) beginning from the best-fit solution with $\Delta > 0$.

$s_m = 1$. In the model with $s_m = 0$ 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 $s_m = 1$ 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 $s_m = 1$ compared to the model with $s_m = 0$.

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).

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 exploit differences in the parameter estimates across high (France, Germany, UK and US) and low/middle (Brazil and India) income regions, under the assumption that the level of development in a country is negatively correlated with assignment frictions. This assumption is based on the existing evidence documenting a positive relationship between development, the supply of managerial capital and good governance.⁴²

We start from the quantification of the share of misassignments in the pooled sample. We first derive ϕ , i.e. the share of type l firms, from the market clearing condition. Overall the whole sample, we observe a share $\hat{\gamma} = 0.347$ of type l CEOs. We must then have $\phi = \hat{\gamma} + (1 - \hat{\gamma})(1 - s_m)$. 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 $1 - s_m$ of type m CEOs. Plugging in for $\hat{\gamma}$ and s_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 $0.089 (= 0.168 * \Delta)$ log points.

We then allow the s_m parameter in the likelihood function (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 4. 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 using the same steps as described above.

⁴²For 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 in Bloom et al. (2016). Furthermore, development is also likely to affect the quality of corporate governance, which affect both the selection and the dismissal of misassigned CEOs. LaPorta et al. (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 (2016) 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%), UK (35%) and US (27%).

Table 4: Estimation Results by Region

	Estimated Parameters		Derived Parameters		
	Δ	s_m	$\hat{\gamma}$	ϕ	% firms mismatched
low/middle income	0.667	0.546	0.216	0.572	0.356
high income	0.667	0.893	0.495	0.549	0.054

Notes: In its first two columns, this table displays the estimated parameters resulting from maximizing (5) using the EM algorithm under the restriction that $A = 0$. The third column is the observed share of leader CEOs in each region. The fourth is the value of ϕ consistent with market-clearing given s_m and the observed share of leader CEOs, while the fifth is the difference between the fourth and third, as this gives the share of type l firms run by manager CEOs.

One possible explanation for these different probabilities across countries is that firms in high-income countries have higher demand for type l CEOs. Indeed, consistent with this idea, the data shows a much larger share of type l CEOs in high-income countries relative to low/middle-income countries (0.495 vs. 0.216). However, note that the ϕ 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).⁴³

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.⁴⁴

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

⁴³We have repeated the same chi-squared tests for restrictions on s_m as described above for each region separately. While the power of the tests is lower due to reduced sample size, we are able to reject pure vertical and horizontal differentiation at a 10% significance level in both regions.

⁴⁴Our 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.

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 (2010-2014) using the Penn World Table data v.9 (Feenstra 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.⁴⁵

5 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 which distinguishes between “managers” –i.e. CEOs that are primarily involved with production-related activities– and leaders –i.e. CEOs that 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 under explored 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 utilize our data and methodology to inform new leadership models, which 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 the current paper. 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 (2015)), but have so far been largely unexplored from an empirical perspective. We leave these topics for further research.

⁴⁵The average labor productivity for high (low/middle) income countries in our sample is 11.4 (9.83). These values are calculated using data on output-side real GDP at chained PPPs and the total number of persons engaged from the Penn World Tables.

References

- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.
- Bandiera, O., Guiso, L., Prat, A., and Sadun, R. (2012). What Do CEOs Do? CEP Discussion Papers dp1145, Centre for Economic Performance, LSE.
- Bandiera, O., Guiso, L., Prat, A., and Sadun, R. (2015). Matching Firms, Managers, and Incentives. *Journal of Labor Economics*, 33(3):623 – 681.
- Bandiera, O., Lemos, R., Prat, A., and Sadun, R. (2017). Managing the Family Firm: Evidence from CEOs at Work. *Review of Financial Studies*.
- Bender, S., Bloom, N., Card, D., Van Reenen, J., and Wolter, S. (2016). Management Practices, Workforce Selection and Productivity. NBER Working Papers 22101, National Bureau of Economic Research, Inc.
- Bennedsen, M., Nielsen, K. M., Perez-Gonzalez, F., and Wolfenzon, D. (2007). Inside the Family Firm: The Role of Families in Succession Decisions and Performance. *The Quarterly Journal of Economics*, 122(2):647–691.
- Bertrand, M. and Schoar, A. (2003). Managing with Style: The Effect of Managers on Firm Policies. *The Quarterly Journal of Economics*, 118(4):1169–1208.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022.
- Bloom, N., Sadun, R., and Van Reenen, J. (2016). Management as a Technology? NBER Working Papers 22327, National Bureau of Economic Research, Inc.
- Bloom, N. and Van Reenen, J. (2007). Measuring and Explaining Management Practices Across Firms and Countries. *The Quarterly Journal of Economics*, 122(4):1351–1408.
- Bolton, P., Brunnermeier, M. K., and Veldkamp, L. (2013). Leadership, Coordination, and Corporate Culture. *The Review of Economic Studies*, 80(2):512–537.
- Cai, J. and Szeidl, A. (2016). Interfirm Relationships and Business Performance. NBER Working Papers 22951, National Bureau of Economic Research, Inc.
- Chade, H. and Eeckhout, J. (2016). Stochastic Sorting. mimeograph.
- Chang, J., Gerrish, S., Boyd-Graber, J. L., and Blei, D. M. (2009). Reading Tea Leaves: How Humans Interpret Topic Models. In *Advances in Neural Information Processing Systems*.

- Coase, R. (1937). The nature of the firm. *Economica*, 4, New Series(16):386–405.
- Deming, D. J. (2015). The Growing Importance of Social Skills in the Labor Market. NBER Working Papers 21473, National Bureau of Economic Research, Inc.
- Dessein, W. and Santos, T. (2016). Managerial Style and Attention. mimeograph.
- Drucker, P. (1967). *The Effective Executive*. New York : Harper and Row.
- Erosheva, E. A., Fienberg, S. E., and Joutard, C. (2007). Describing Disability through Individual-Level Mixture Models for Multivariate Binary Data. *The Annals of Applied Statistics*, 1(2):502–537.
- Feenstra, Robert C., R. I. and Timmer, M. P. (2015). The next generation of the penn world table. *American Economic Review*, 10(105):3150–3182.
- Gabaix, X. and Landier, A. (2008). Why has CEO Pay Increased So Much? *The Quarterly Journal of Economics*, 123(1):49–100.
- Gennaioli, N., Porta, R. L., de Silanes, F. L., and Shleifer, A. (2013). Human Capital and Regional Development. *The Quarterly Journal of Economics*, 128(1):105–164.
- Graham, J. R., Harvey, C. R., and Puri, M. (2013). Managerial attitudes and corporate actions. *Journal of Financial Economics*, 109(1):103–121.
- Griffiths, T. L. and Steyvers, M. (2004). Finding Scientific Topics. *Proceedings of the National Academy of Sciences*, 101(Suppl. 1):5228–5235.
- Gross, J. H. and Manrique-Vallier, D. (2014). A Mixed Membership Approach to the Assessment of Political Ideology from Survey Responses. In Airoldi, E. M., Blei, D., Erosheva, E. A., and Fienberg, S. E., editors, *Handbook of Mixed Membership Models and Its Applications*. CRC Press.
- Hambrick, D. C. and Mason, P. A. (1984). Upper Echelons: The Organization as a Reflection of Its Top Managers. *The Academy of Management Review*, 9(2):193–206.
- Hermalin, B. E. (1998). Toward an Economic Theory of Leadership: Leading by Example. *American Economic Review*, 88(5):1188–1206.
- Hermalin, B. E. (2007). Leading for the Long Term. *Journal of Economic Behavior & Organization*, 62(1):1–19.
- Hoffman, M. and Tadelis, S. (2017). How do managers matter? evidence from performance metrics and employee surveys in a firm. *mimeo, University of Toronto*.

- Kaplan, S. N., Klebanov, M. M., and Sorensen, M. (2012). Which CEO Characteristics and Abilities Matter? *The Journal of Finance*, 67(3):973–1007.
- Kaplan, S. N. and Sorensen, M. (2016). Are CEOs Different? Characteristics of Top Managers. mimeograph.
- Kotter, J. P. (1999). *John Kotter on What Leaders Really Do*. Harvard Business School Press, Boston.
- La Porta, R., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. (2000). Investor protection and corporate governance. *Journal of Financial Economics*, 58(1-2):3–27.
- LaPorta, R., Lopez-De-Silanes, F., and Shleifer, A. (1999). Corporate Ownership Around the World. *Journal of Finance*, 54(2):471–517.
- Lazear, E. P., Shaw, K. L., and Stanton, C. T. (2015). The value of bosses. *Journal of Labor Economics*, 33(4):823 – 861.
- Lippi, F. and Schivardi, F. (2014). Corporate control and executive selection. *Quantitative Economics*, 5:417–456.
- Malmendier, U. and Tate, G. (2005). CEO Overconfidence and Corporate Investment. *Journal of Finance*, 60(6):2661–2700.
- Malmendier, U. and Tate, G. (2009). Superstar CEOs. *The Quarterly Journal of Economics*, 124(4):1593–1638.
- Mintzberg, H. (1973). *The Nature of Managerial Work*. Harper and Row., New York.
- Mintzberg, H. (1979). *The Structuring or Organizations*. Englewood Cliffs, N.J.: Prentice-Hall.
- Mullins, W. and Schoar, A. (2013). How do CEOs see their Role? Management Philosophy and Styles in Family and Non-Family Firms. NBER Working Papers 19395, National Bureau of Economic Research, Inc.
- Roberts, J. (2006). The Modern Firm: Organizational Design for Performance and Growth John Roberts (Oxford: Oxford University Press, 2004); ISBN 0-19-829376-3, 318 pages. *Business Ethics Quarterly*, 16(01):109–109.
- Tervio, M. (2008). The Difference That CEOs Make: An Assignment Model Approach. *American Economic Review*, 98(3):642–68.
- Tipping, M. E. and Bishop, C. M. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society B*, 61(3):611–622.

Urban, D. (2016). The Effects of Culture on CEO Power: International Evidence from Executive Turnover The Effects of Culture on CEO Power: International Evidence from Executive Turnover. *mimeo, Technische Universität München (TUM)*.

Van den Steen, E. (2010). On the origin of shared beliefs (and corporate culture). *RAND Journal of Economics*, 41(4):617–648.

APPENDIX TABLES AND FIGURES - NOT FOR PUBLICATION

A Data Appendix

A.1 The Time Use Survey

The time use survey took place in two stages: in the Spring of 2011 a team of 15 analysts based in Mumbai and led by one of our project managers collected data on India, while the rest of the countries were covered in a second survey wave in the Spring of 2013 by a team of 40 enumerators based at the London School of Economics.⁴⁶ To ensure comparability, we adopted the same protocol and retained the same project manager across both waves. The enumerators were typically graduate students (often MBAs) recruited specifically for this project. All enumerators were subject to a common intensive training on the survey methodology for three days at the beginning of the project, plus weekly team progress reviews and one to one conversations with their supervisors to discuss possible uncertainties with respect to the classification of the time use data. Each interview was checked off at the end of the week by one supervisor, who would make sure that the data was complete in every field, and that the enumerator had codified all the activities according to the survey protocol. Each enumerator ran on average 30 interviews.

Each enumerator was allocated a random list of about 120 companies, and was in charge of calling up the numbers of his or her list to convince the CEO to participate in the survey, and to collect the time use data in the week allocated to the CEO. One project manager, five full time supervisors and one additional manager working on a part time basis led the survey team. We actively monitored and coached the enumerators throughout the project, which intensified their persistence in chasing the CEOs and getting them to participate. We also offered the CEOs a personalized analysis of their use of time (which was sent to them in January 2012 to the Indian CEOs and in June 2014 to the rest of the countries) to give them the ability to monitor their time allocation, and compare it with peers in the industry.

The survey instrument is available at www.executivetimeusesurvey.org. A screenshot of the blank instrument is shown in Figure A.1.

A.2 Sampling Frame

The sampling frame was drawn from ORBIS, an extensive commercial data set that contains company accounts for several millions of companies around the world. Our sampling criteria were as follows. First, we restricted the sample to manufacturing and additionally kept firms that were classified as “active” in the year prior to the survey (2010 in India and 2012 for the other countries) and with available recent accounting data.⁴⁷ These conditions restricted our sample to 11,500 firms. Second, we further restricted the sample to companies for which we could find CEOs contact details.

⁴⁶The data collection methodology discussed in this section is an evolution of the approach followed in Bandiera et al. (2012) to collect data on the diary of 100 Italian CEOs. While the data collection of the Italian data was outsourced to a private firm, the data collection described in this paper was internally managed from beginning to end. Due to this basic methodological difference and other changes introduced after the Italian data was collected (e.g. the vector of features used to characterize every activity) we decided not to combine the two samples.

⁴⁷For the Indian sample, we also restricted the sample to firms headquartered in the fifteen main Indian states. This excluded firms located in Assam, Bihar, Chandigarh, Chhattisgarh, Dadra, Daman and Diu, Goa, Himachal Pradesh, Jammu and Kashmir, Jharkhand, Orissa and Uttarakhand, each of which accounts for less than 3% of Indian GDP.

Figure A.1: Survey Instrument

ACTUAL AGENDA	
Tuesday	
On Tuesday, at what time did the Executive START working? Please consider all work-related activities (e.g. calls from home, breakfast meetings). 09:30 AM	
On Tuesday, at what time did the Executive FINISH working? Please consider all work-related activities (e.g. calls from home, dinner meetings). 09:15 PM	
Please enter all activities lasting more than 15 minutes for Tuesday. You can report up to 15 activities if necessary.	
Activity 1: <input type="text" value="Preparing daily schedule/HQ/alone"/> Activity 2: <input type="text" value="Checking MIS from Finance dept./HQ/alone"/> Activity 3: <input type="text" value="meeting / HQ/ consultant"/> Activity 4: <input type="text" value="Emails/ HQ/ alone"/> Activity 5: <input type="text" value="Phonecall/ HQ/ Deputy CFO"/> Activity 6: <input type="text" value="Emails/ HQ/ alone"/> Activity 7: <input type="text" value="Lunch/ HQ/ Executives"/> Activity 8: <input type="text" value="Meeting/ HQ/ Business Head (Drill)"/> Activity 9: <input type="text" value="Phonecall/HQ/Marketing Head"/> Activity 10: <input type="text" value="Phonecall/ HQ/ Customer"/> Activity 11: <input type="text" value="Increment Meeting/ HQ/HR Head"/> Activity 12: <input type="text" value="Meeting for grading people/ HQ/ Finance Head"/> Activity 13: <input type="text" value="Phonecall / HQ / Manufacturing Head"/> Activity 14: <input type="text" value="Emails/ HQ/ alone"/> Activity 15: <input type="text" value="Phonecall/HQ/ Marketing Head (South & west)"/>	Start Time: 09:30 AM End Time: 10:00 AM Start Time: 10:00 AM End Time: 10:30 AM Start Time: 10:30 AM End Time: 12:00 PM Start Time: 12:00 PM End Time: 12:30 PM Start Time: 12:30 PM End Time: 01:15 PM Start Time: 01:15 PM End Time: 01:30 PM Start Time: 01:30 PM End Time: 02:30 PM Start Time: 02:30 PM End Time: 02:45 PM Start Time: 02:45 PM End Time: 03:15 PM Start Time: 03:15 PM End Time: 03:30 PM Start Time: 03:30 PM End Time: 04:00 PM Start Time: 04:00 PM End Time: 04:30 PM Start Time: 04:30 PM End Time: 06:00 PM Start Time: 06:00 PM End Time: 07:00 PM Start Time: 07:00 PM End Time: 07:45 PM
Checked by supervisor? Jaidev	

Activity 1:	Type <input type="text"/>	Who participated in the activity, excluding the Executive? (check all that apply)	
	When was the activity scheduled in agenda? <input type="text"/>	<input type="checkbox"/> People employed by firm INSIDERS	<input type="checkbox"/> People not employed by firm OUTSIDERS
Start Time <input type="text"/>	If unscheduled, was the activity undertaken due to an emergency? <input type="text"/>	What type of INSIDERS participated in the activity? (i.e. people employed by the firm)	What type of OUTSIDERS participated in the activity? (i.e. people NOT employed by the firm)
End Time <input type="text"/>	Did the activity take place inside the firm and/or HQ? <input type="text"/>	<input type="checkbox"/> Finance <input type="checkbox"/> Marketing/Communication <input type="checkbox"/> Production/Logistics <input type="checkbox"/> Strategy <input type="checkbox"/> Human Resources <input type="checkbox"/> Business Unit Directors <input type="checkbox"/> Others	<input type="checkbox"/> Clients <input type="checkbox"/> Suppliers <input type="checkbox"/> Banks <input type="checkbox"/> Investors <input type="checkbox"/> Lawyers <input type="checkbox"/> Management <input type="checkbox"/> Consultants <input type="checkbox"/> Politicians <input type="checkbox"/> Government Officials <input type="checkbox"/> Journalists <input type="checkbox"/> Unions <input type="checkbox"/> Competitors <input type="checkbox"/> Others
	Where did the activity take place, relative to HQ? <input type="text"/>	If "Others", specify: <input type="text"/>	
	How many people were present at the activity, excluding the Executive? <input type="text"/>	If "Others", specify: <input type="text"/>	

Table A.1: Selection Analysis

	(1)	(2)	(3)	(4)
Sample	All	All	All	All
Dependent Variable: Dummy=1 if CEO participated				
Country=Brazil	0.677*** (0.074)	0.695*** (0.075)	0.655*** (0.079)	0.559* (0.288)
Country=France	0.210*** (0.073)	0.256*** (0.074)	0.143 (0.104)	0.562** (0.221)
Country=Germany	0.115 (0.072)	0.194** (0.078)	0.152* (0.082)	0.476** (0.222)
Country=India	0.658*** (0.247)	0.699** (0.272)	1.227*** (0.371)	0.672 (0.425)
Country=UK	-0.178** (0.074)	-0.139* (0.074)	-0.153** (0.077)	0.088 (0.218)
Ln(Sales)		-0.071*** (0.011)		
ln(Sales/Employees)			-0.018 (0.030)	
ROCE				0.000 (0.001)
Number of firms	6256	5993	4090	3492

Notes: *significant at 10%; ** significant at 5%; *** significant at 1%. All columns estimated by probit (marginal effects reported, robust standard errors under coefficient). The dependent variable in all columns is a dummy=1 if the CEO participated in the survey. The selection regression is run on the latest available year of accounting data. All columns include 2 digits SIC industry dummies.

To gather contact information we hired a team of research assistants based in Mumbai, London and Boston who verified the CEOs names and found their phone numbers and emails. This restricted the sample to 7,744 firms. Of these, 907 later resulted not to be eligible for the interviews upon the first telephonic contact (the reasons for non eligibility included recent bankruptcy or the company not being in manufacturing), and 310 were never contacted because the project ended before this was possible. The final number of eligible companies was thus 6,527, with median yearly sales of \$53,000,000. Of these, we were able to secure an interview with 1,131 CEOs, although 17 CEOs dropped out before the end of the data collection week for personal reasons and were thus removed from the sample before the analysis was conducted.

The selection analysis in Table A.1 shows that firms in the final sample have on average slightly lower log sales relative to the sampling frame (coefficient 0.071, standard error 0.011). However, we do not find any significant selection effect on performance variables, such as labor productivity (sales over employees) and return on capital employed (ROCE).

Table A.2 presents the basic summary statistics of the sample.

Table A.2: Summary Statistics

Variable	Mean	Median	Standard Deviation	Observations
A. CEOs Traits				
CEO age	50.93	52.00	8.45	1107
CEO gender	0.96	1.00	0.19	1114
CEO has college degree	0.92	1.00	0.27	1114
CEO has MBA	0.55	1.00	0.50	1114
CEO tenure in post	10.29	7.00	9.55	1110
B. Firms Traits				
Employment	1,275.47	300.00	6,497.72	1114
Sales ('000 \$)	222,033.90	35,340.49	1,526,261.00	920
Capital ('000 \$)	79,436.72	10,029.00	488,953.60	618
Materials ('000 \$)	157,287.10	25,560.02	1,396,475.00	448
Profits per employee ('000 \$)	8.62	2.55	14.87	386

Notes: Variables in Panel A and B are drawn from our survey and ORBIS, respectively.

Table B.1: Five Most Common Activities in Pure Behavior 0

Type	Planned	Duration	Size	Functions	Prob. in β^0	Prob. in β^1
Meeting	Yes	Long	Large	Production	0.057	0.000
Meeting	Yes	Long	Small	Clients	0.027	0.000
Meeting	Yes	Long	Small	Production	0.025	0.012
Meeting	Yes	Long	Large	Marketing	0.024	0.012
Meeting	Yes	Long	Large	Marketing/Production	0.023	0.000

B Further Results from LDA Model

B.1 Most common activities in each pure behavior

These tables display the most common activities in each pure behavior. In the duration category, long refers to an activity's lasting longer than one hour; in the size category, small refers to an activity's involving just one other person, while large refers to its involving more than one person. Regarding functions, groupcom refers to members of the firm's commercial group, and associations are trade association meetings.

B.2 Significance of Differences in Pure Behaviors

A natural question is whether the difference in pure behaviors is significant. To explore this, we adopt the following approach. First, we generate a dataset of activities based upon a model in which there are no underlying differences among CEOs. Specifically, we take the empirical distribution of

Table B.2: Five Most Common Activities in Pure Behavior 1

Type	Planned	Duration	Size	Functions	Prob. in β^0	Prob. in β^1
Meeting	Yes	Long	Large	C-suite	0.057	0.000
Meeting	Yes	Long	Large	Others	0.027	0.000
Meeting	Yes	Long	Large	Associations	0.025	0.012
Meeting	Yes	Long	Large	Marketing/Clients	0.024	0.012
Meeting	Yes	Long	Large	Board	0.023	0.000

the 654 activities that enter the LDA analysis and for each time unit draw an activity independently from it. This corresponds to a model in which there is a single pure behavior from which all CEOs draw their observed activities. We then estimate the same parameters on this simulated data as we do on the actual data, and compute the Hellinger distance between the two estimated pure behaviors. We repeat this procedure 1,000 times.

Figure B.1 plots the distribution of the Hellinger distances in the 1,000 simulations. The red line denotes the Hellinger distance we observe in the actual data. In no simulation does the Hellinger distance between two behaviors exceed that we observe in the actual data: the maximum simulated distance is 0.412 whereas in the actual data the distance is 0.776. We therefore conclude that it is highly unlikely that our observed data is consistent with a model in which all CEOs adopt a single pure behavior.

B.3 Estimated time shares

We also report the raw and estimated time shares in the baseline sample in table B.3. The raw shares are simply the shares of time that the average CEO is observed to spend in different categories. These differ slightly from those displayed in figure 1 since we only compute averages on the subset of activities that include non-rare feature combinations. The estimated shares are the fraction of time each behavior spends in each category, weighted by the average value of the CEO behavior index. In general there is a very close relationship between the raw and estimated shares. The largest deviations occur for time with outsiders and with insiders and outsiders together. However these are derived from the probabilities each behavior places on different combinations of individual functions rather than a feature explicitly included in the algorithm.

C Proof of Proposition 2

We verify that the situation described in the proposition corresponds to a Bayesian equilibrium. To simplify notation re-normalize all variables so that $\Delta = 1$.

First note, that if $B > 1$, all CEOs will choose the behavior that is optimal for the firm that hires them. This means that CEO behavior only depends on firm type. Therefore, in what follows we assume that governance is sufficiently poor, so $B < c$.

In that case, when a CEO is hired, her utility is $\bar{w} + B$ if she works for a firm of the same type and \bar{w} if she works for a firm of a different type. To simplify notation, further normalize $\bar{w} + B = 1$.

Figure B.1: Distribution of Hellinger Distances in Simulated Data




Table B.3: Raw and Estimated Time Shares

	Raw	Estimated
Meeting	0.803	0.801
Communications	0.068	0.06
Site Visit	0.06	0.062
Insiders	0.657	0.653
Outsiders	0.235	0.175
Insiders & Outsiders	0.108	0.171
Production	0.35	0.355
Marketing	0.206	0.208
C-suite	0.115	0.122
Clients	0.103	0.104
Suppliers	0.064	0.068
Consultants	0.026	0.026
Planned	0.764	0.782
>1 Hour	0.657	0.687
2 People or More	0.553	0.573
2 Functions or More	0.273	0.262

Notes: This table compares the observed share of time that CEOs spend on average in different activities against that estimated by LDA. To obtain the latter, we obtain the average time spent on each activity as $\frac{\sum_i \hat{\theta}_i \hat{\beta}_1 + (1 - \hat{\theta}_i) \hat{\beta}_0}{N}$.

Hence the utility of a correctly matched CEO is one and the utility of a mismatched CEO is

$$b \equiv \frac{\bar{w}}{\bar{w} + B}.$$

Note that b is a measure of the quality of governance, with $b = 1$, being the worst level of governance.

A type m firm faces an abundant supply of type m CEOs. As all the applications it receives come from type m CEOs, the firm will simply hire the first applicant. A type l firm instead may receive applications from both CEO types. If k is sufficiently low, the optimal policy consists in waiting for the first candidate with $s = l$ and hire him.

We now consider CEOs. Suppose that all leader CEOs apply to type l firms and manager CEOs apply to type l firms with probability z and to type m firms with probability $1 - z$.

If a manager CEO applies to a type m firm, he will get a job if and only if his application is downloaded. The mass of type m firms is $1 - \phi$. The mass of manager CEOs applying to type m firms is $(1 - \gamma)(1 - z)m$. The probability the CEO is hired is

$$P_m = \frac{1 - \phi}{(1 - \gamma)(1 - z)m}.$$

If instead a manager CEO applies to a type l firm, he will get a job if and only if his application is considered and the firm does not detect deception. Computing the first probability requires an additional step, because some firms consider more than one application before they find an application which passes the screening process.

The probability that a type l firm application is accepted if it is considered is:

$$H = \frac{(1 - \gamma)z(1 - \rho) + \gamma}{(1 - \gamma)z + \gamma}.$$

The mass of applications that are downloaded by type l firms is therefore:

$$\phi(1 + (1 - H) + (1 - H)^2 + \dots) = \phi \frac{1}{H}.$$

Given that the mass of applicants to type l firms is $m((1 - \gamma)z + \gamma)$, the probability that an application is considered is

$$\frac{\phi}{m(\gamma + (1 - \gamma)z)H} = \frac{\phi}{m((1 - \gamma)z(1 - \rho) + \gamma)}$$

The probability that a type m CEO applicant passes the screening process is $1 - \rho$. Thus, the probability that a type m CEO applicant is hired by a type l firm is

$$P_l = \frac{(1 - \rho)\phi}{m((1 - \gamma)z(1 - \rho) + \gamma)}.$$

In the equilibrium under consideration a type m CEO must be indifferent between applying to the two types of firms. As the benefit of being hired by a same-type firm is one, while the benefit of being hired by a type l firm is b , the indifference condition is $P_m = bP_l$, which yields:

$$\frac{1 - \phi}{(1 - \gamma)(1 - z)} = \frac{(1 - \rho)\phi b}{((1 - \gamma)z(1 - \rho) + \gamma)},$$

yielding

$$z = \frac{(1 - \gamma)(1 - \rho)\phi b - (1 - \phi)\gamma}{(1 - \phi + \phi b)(1 - \gamma)(1 - \rho)}.$$

The solution of z will be positive – meaning that some type m CEOs will apply to type l firms – if

$$\rho < 1 - \frac{(1 - \phi)\gamma}{(1 - \gamma)\phi b},$$

which is satisfied as long as ρ is not too high, b is not too low, and γ is sufficiently smaller than ϕ . For instance, the combination of $\rho = 0$, $b = 1$, and $\phi > \gamma$ would work.

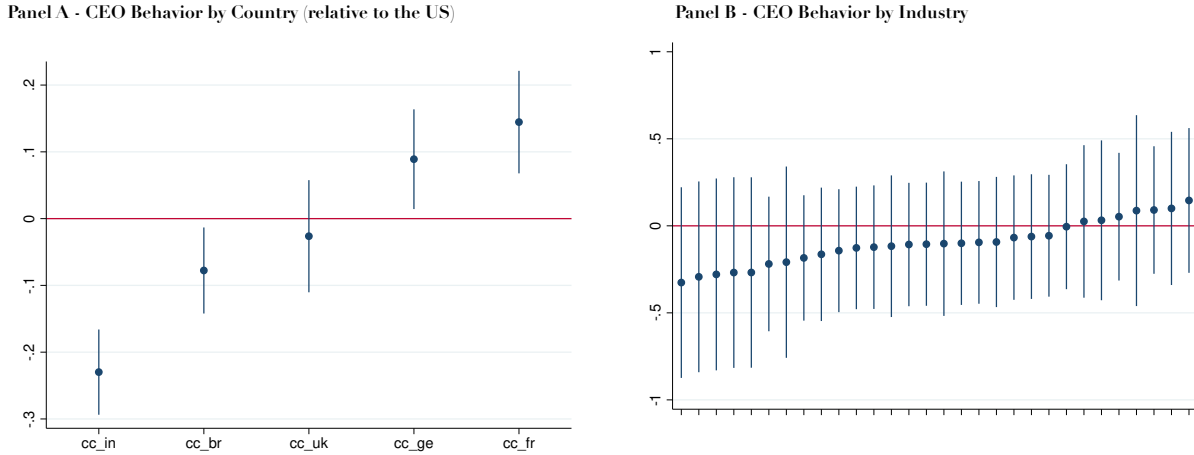
Type l CEOs always produce 1, while the average productivity of a type m CEO is equal to the probability that he is matched with a type m firm, which is

$$\frac{1 - z}{1 - z + z(1 - \rho)}.$$

By replacing z , we find the average productivity of a type m CEO:

$$\frac{(1 - \phi)((1 - \gamma)(1 - \rho) + \gamma)}{(1 - \phi)(1 - \gamma)(1 - \rho) + (1 - \phi)\gamma + ((1 - \gamma)(1 - \rho)\phi b - (1 - \phi)\gamma)(1 - \rho)},$$

Figure D.1: CEO Behavior Index: Variation across Countries and SIC 2 industries



Notes: Each point represents the coefficient obtained when regressing the CEO behavior index on country and two digits SIC fixed effects.

which is smaller than one whenever $\rho < 1$.

Finally, note that the difference between the profit (including CEO compensation) of a correctly matched firm and an incorrectly matched one is $1 - B$.

D Additional Results

D.1 CEO Behavior Index: Additional Descriptives

D.1.1 Variation across Countries and Industries

Figure D.1 shows the point estimates and confidence intervals of the regression of the CEO behavior index on, respectively, country (using the US as relative country benchmark) and SIC 2 industry dummies.

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

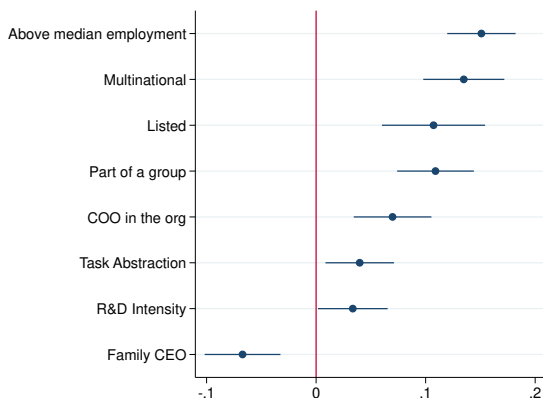
D.1.2 Correlation with Firm and CEO Characteristics

Panel A, Figure D.2 reports the correlation between CEO behavior and firm/CEO traits controlling for country and industry fixed effects. Larger firms, multinationals, listed firms and firms that have a COO are all more likely to hire a leader CEO.


The index is also correlated with specific CEO characteristics, as shown in Panel B. It is significantly larger for CEOs who report having had a study or work experience outside their home country, or to have attained an MBA degree or equivalent. In contrast, there is no evidence that

Figure D.2: CEO Behavior Index, Firm and CEO Characteristics

Panel A - CEO Behavior Index and Firm Characteristics



Panel B - CEO Behavior Index and CEO Characteristics



Notes: Each point represents the coefficient obtained when regressing the CEO behavior index on each dummy variable, including country and SIC2 industry dummies as additional controls.

the index is related to the age of the executive, or to whether the CEO was promoted to the role within the organization.

D.2 Production Function: Robustness Checks

We have examined the robustness of the basic results discussed in Table 2. The robustness checks are summarized in Tables D.1 and D.2. In each table, Column 1 simply reports the baseline results of Table 2, column 1.

D.2.1 Using shares of time instead of the CEO Behavior Index

Table D.1 shows the basic production function results when we use the share of time spent by CEOs in activities with different features rather than the CEO index. Starting with activity type, Column 2 shows that there is a negative and precisely estimated correlation between the time spent in plant visits and performance, while the correlation with time spent in communications is positive but not precisely estimated (all relative to time spent in meetings). Column 3 shows that among participants, firm performance is higher when CEOs devote more time to insiders together with outsiders as opposed to outsiders or insiders alone. Moving to specific functions, Column 3 shows that performance is negatively correlated with the time spent with production and clients and positively correlated with time spent with C-suite executives and marketing. Column 5 shows that performance is positively correlated with planning and multi-functional and multi-participant interactions but not with meeting duration. Taken together, the results suggests that most of the features for which CEOs with different indexes behave similarly (meetings, insiders, group size) are not correlated with performance. The sole exception is the share of planned time, which is positively correlated with performance but not with the index. Moreover, all the differences captured by our index (site vs communication, outsiders alone vs with insiders, production and clients vs. C-suite, single function vs multifunction interactions) are individually correlated with firm performance.

D.2.2 Alternative specification choices

We examined whether the results varied when we used annual accounting data, instead of the averaged version employed in the baseline regressions. Table (D.2), Panel A, column 2 shows that the baseline results are not sensitive to this choice. In column 3 we show that the unweighted regressions deliver a very similar coefficient on the CEO behavior index relative to the baseline results, which are weighted by the representativeness of the week as rated by the CEO at the end of the data collection week.

D.2.3 Controlling for CEO and firm characteristics

We investigated whether the coefficient of the CEO behavior index in the baseline result could capture the effect of other CEO or firm observables, which could be at the same time correlated with CEO behavior and firm performance. In column 4 we show that the coefficient on the index actually increases when we control for the overall number of hours worked by the CEO during the survey week, a proxy for effort which was extensively analyzed in Bandiera et al. (2017). In columns 5 and 6 we include a set of firm dummies to denote whether the firm is a multinational, part of a group, owned and run by a family CEO, listed on a public exchange and has a COO in the organizational chart, and CEO characteristics (dummies to capture whether the CEO holds an MBA degree or equivalent, has studied or worked abroad, is male, was promoted internally and age). While these additional variables are for the most part insignificant, the coefficient on the CEO behavior index remains large and statistically significant.⁴⁸

D.2.4 Alternative ways of expressing the CEO behavior index, including alternative dimensionality reduction techniques

We experimented with different ways of expressing the CEO behavior index.

First, we used a discretized version of the index ($=1$ if the index is ≥ 0.5), as shown in Table (D.2), Panel A, column 7. We also examined alternative dimensionality reduction approaches, namely PCA and k-means analysis, on the key marginals that emerge from LDA as being significantly different across behavior types. For each CEO, we counted the number of engagements that: (1) last longer than one 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 PCA analysis explains 36% of the variance in this feature space and places a positive weight on all dimensions 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 (4) and lower values for all others). Hence the patterns identified using simpler methods validate the key differences from LDA with two pure behaviors.⁴⁹ In the columns 8 and 9 of Table D.2, Panel A we show that these alternative ways of classifying CEOs do not fundamentally alter the relationship between behavior and firm performance.

⁴⁸Among the firm variables, the only significant ones are the family CEO dummy (negative) and the COO dummy (positive). Among the CEO variables, the only significant ones are the log of CEO age and the dummy to capture experience abroad, both positive.

⁴⁹Note 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 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 relative to the ones described above.

D.2.5 Activity selection

In the baseline analysis, we define a rare activity as one not present in the time use of at least 30 CEOs. When we drop these activities from the analysis, we discard 23% of interactive activities on average across CEOs. One potential concern is that the choice of rare activities itself is a component of behavior that we do not capture with the behavior index. To address this, we construct a behavior index based on dropping activities not present in the time use of at least 15 and, alternatively, 45 CEOs. The results are presented in Table (D.2), Panel B, columns 2 and 3. The results are essentially identical as for the baseline index.

In the baseline results, we build the index only on the basis of interactive activities, excluding traveling. Column 4 shows that we would obtain very similar results if we were to include travel in the set.

LDA is a mixed-membership model that allows CEOs to mix their time between two pure behaviors. An alternative model is a simpler mixture model in which each CEO is associated exclusively to one behavior. We have estimated a multinomial mixture model via the EM algorithm, and derived an alternative behavior index as the probability that a CEO draws activities from behavior 1.⁵⁰ Again, we find a significant relationship between the behavior index and firm performance, as shown in Table (D.2), Panel B, column 5.

The behavior index in the main paper is based on all 1,114 CEOs in our time use survey, but we have sales data for 920. We therefore also construct the index based on the subset of CEOs for which sales data is available, but as column 6 shows this does not change the coefficient.

A final concern is that the differences we capture in the behavior index arise solely from cross-region variation in time use, and that within-region variation is not related to firm performance. We therefore construct a behavior index for CEOs in low/middle-income countries based solely on time use observed in these countries, and likewise for CEOs in the high-income countries. Column 7 shows the results on firm performance, and we again find a significant relationship.

D.2.6 Alternative estimation techniques

Table (D.2), Panel B, column 8 shows the results when we regress we regress the Olley Pakes estimator of productivity on the CEO behavior index. Given the need to rely on panel data for capital, this restricts the sample to 562 firms. As a comparison, the OLS estimate of the CEO behavior index on the same sample is 0.244 (standard error 0.107).

D.2.7 Choosing number of pure behaviors with out-of-sample prediction

As discussed in the main text, we choose two pure behaviors primarily for interpretability, but an alternative is to choose the number of pure behaviors K based on a statistical criterion. We adopt perhaps the most popular approach—cross-validation—in which K is chosen based on the ability of the model to predict out-of-sample observations. We first randomly draw two-thirds of our sample of CEOs as training data, and fit an LDA model for various values of K beginning from $K = 2$. Then we take the estimated parameters and compute the goodness-of-fit for the test data (the held-out one-third of CEOs) using *perplexity*, a standard measure in the machine learning

⁵⁰In the mixture model, each CEO draws all of his/her activities from a single pure behavior, but the econometrician is unsure which behavior this is. The E-step in the EM algorithm provides a probability distribution over cluster assignments, and we use the probability of being assigned to cluster 1 as the behavior index.

Table D.1: Production Function Results Using Shares of Time

	(1)	(2)	(4)	(5)	(3)
Dependent Variable: log(sales)					
log(employment)	0.889*** (0.040)	0.895*** (0.039)	0.893*** (0.041)	0.907*** (0.040)	0.876*** (0.040)
CEO behavior index	0.343*** (0.108)				
Share of time spent in Communications		0.066 (0.253)			
Share of time spent in Plant visits (site)		-1.168*** (0.364)			
Share of time spent with Insiders only			0.375** (0.187)		
Share of time spent with Insiders and Outsiders together			-0.166 (0.166)		
Share of time spent with Production			0.055 (0.175)		
Share of time spent with Marketing			0.494*** (0.164)		
Share of time spent with C-suite managers				0.247 (0.187)	
Share of time spent with Clients				0.353 (0.237)	
Share of time spent with Suppliers					-0.661*** (0.235)
Share of time spent with Consultants					0.459 (0.299)
Share of time spent in Planned activities					0.373 (0.239)
Share of time spent in Interactions > 1hr					-0.804** (0.318)
Share of time spent in Interactions with more than 2 people					-0.462 (0.603)
Share of time spent in interactions with more than 2 functions					0.281 (0.649)
Number of observations (firms)	920	920	920	920	920
Observations used to compute means	2,202	2,202	2,202	2,202	2,202

Notes: *** (**) (*) denotes significance at the 1%, 5% and 10% level, respectively. All columns include the same controls used in 2, column 1.


Table D.2: Robustness Checks

Panel A									
Experiment	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Baseline	Firm by year accounting data, cluster at the firm level	No weighting	Control for hours worked	Control for firm observables	Control for CEO observables	Discretized version ($\geq .5$)	Principal Component	K-means
Dependent Variable: Log sales									
CEO behavior index	0.343*** (0.108)	0.275*** (0.091)	0.298*** (0.104)	0.338*** (0.105)	0.265** (0.114)	0.309*** (0.108)	0.234*** (0.064)	0.109*** (0.030)	0.259*** (0.074)
log(CEO hours worked)				0.297* (0.159)					
MNE (dummy)					0.089 (0.078)				
Part of a Group (dummy)					0.045 (0.084)				
Family CEO (dummy)					-0.215** (0.093)				
Listed (dummy)					0.143* (0.085)				
COO in the org (dummy)					0.143* (0.079)				
CEO has MBA (dummy)						-0.048 (0.072)			
CEO has Experience abroad (dummy)						0.189** (0.080)			
CEO age						0.354* (0.181)			
CEO is male (dummy)						-0.126 (0.132)			
CEO is an internal promotion (dummy)						0.063 (0.057)			
Number of observations (firms)	920	2,202	920	920	920	920	920	920	920
Observations used to compute means	2,202	2,202	2,202	2,202	2,202	2,202	2,202	2,202	2,202

Panel B								
Experiment	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Baseline	Exclude all activities not present in at least 15 CEOs	Exclude all activities not present in at least 45 CEOs	Include Travel and Email	Mixture model	Index computed on sales sample only	Index computed by high and low income countries separately	Olley Pakes productivity residual
Dependent Variable: Log sales								
CEO behavior index	0.343*** (0.108)	0.305*** (0.110)	0.316*** (0.102)	0.265*** (0.096)	0.133** (0.064)	0.347*** (0.102)	0.292*** (0.081)	0.472*** (0.106)
Number of observations (firms)	920	920	920	920	920	920	920	562
Observations used to compute means	2,202	2,202	2,202	2,202	2,202	2,202	2,202	1,431

Notes: *** (**) (*) denotes significance at the 1%, 5% and 10% level, respectively. All columns include the same controls used in 2, column 1. **Panel A:** Column 2 uses yearly accounting data instead of firm level aggregates (always based a max on an interval including 5 years per firm, during the CEO tenure in office). Column 3 shows unweighted results. Column 4 includes as additional control the log of total hours worked by the CEO during the week. Column 5 includes as additional controls a set of firm level characteristics (MNE status, part of a group, family CEO, listed firm dummies). Column 6 includes as additional controls a set of CEO characteristics (MBA, study or work experience abroad, gender and internal promotion dummies and log age). Column 7 uses the discretized version of the CEO behavior index ($=1$ if the index is ≥ 0.5). Column 8 uses an index derived using the first principal component from PCA. Column 9 derives the index from a k-means clustering approach. **Panel B:** Column 2 uses LDA excluding all activities that are not present in at least 15 CEO diaries, and column 3 does the same using 45 diaries as a threshold. Column 4 builds the index using a Mixture Model. Column 5 computes the index with the LDA method, but only using the activities of CEOs working in firms included in the production function sample. Column 6 applies the LDA approach differently by high and low/middle income countries. Column 7 uses the CEO behavior index built by high and low income country separately. Column 8 shows the results obtained when we regress the Olley Pakes estimator of productivity on the CEO behavior index.

Figure D.3: Perplexity for Different Numbers of Pure Behaviors



Notes: This graph plots the average perplexity computed on test data from ten randomly drawn sets of training data. The split between training data and test data is two-thirds / one-third. Lower values of perplexity indicate better goodness-of-fit. There are gaps in the values for K due to save on computation time.

literature given by


$$\exp \left[- \frac{\sum_{i=1}^N \sum_{a=1}^A n_{i,a} \log \left(\sum_{k=1}^K \theta_{i,k} \beta_a^k \right)}{\sum_{i=1}^N T_i} \right]$$

where $n_{i,a}$ is the total number of times activity a appears in the time use of CEO i ; $\theta_{i,k}$ is the probability CEO i adopts pure behavior k ; β_a^k is the probability that pure behavior k generates activity a ; and T_i is the total number of time units observed for CEO i . Here the relevant population of CEOs is the test sample. We use the estimated value of β_a^k from the LDA estimation on the training data, and a uniform distribution for $\theta_{i,k}$ to compute perplexity. We repeat this procedure ten times, each time randomly drawing the training data. Figure D.3 reports the average perplexity computed on the test data across these ten draws. Lower values indicate better goodness-of-fit.

As we increase the number of pure behaviors from $K = 2$, we can indeed better fit time-use patterns, as can be seen from the decreasing perplexity. Naturally, the most parsimonious model does not account for all the underlying correlations in the high-dimensional feature space. At the same time, the improvement in fit levels off fairly quickly, and the average perplexity stays essentially flat from $K = 5$ through $K = 25$ before subsequently increasing. This increase is due to the fact that high values of K capture correlations specific to the training data that do not generalize to test data.

For illustrative purposes, we choose to analyze $K = 11$, where average perplexity achieves a local minimum, although $K = 18$ corresponds to the global minimum (the difference is merely 0.25 and the interpretation difficulties for $K = 11$ will a fortiori become more severe for $K = 18$). Rather than describing behavior with a single index, the $K = 11$ model yields a ten-dimensional vector

Figure D.4: Coefficient Estimates for Model with Eleven Pure Behaviors




Notes: This graph plots the coefficient estimates for the productivity regression in column (1) of table 2 when we replace the scalar behavioral index with the output of the model with $K = 11$. The omitted category is the probability put on the sixth pure behavior. The dots represent point estimates, and the lines 95% confidence intervals.

that describes CEO behavior (we omit the probability associated with the sixth pure behavior). We use this in place of the behavioral index in the regression model of column (1) of table 2. The point estimates and confidence intervals are displayed in figure D.4.

An initial result is that an F-test for the joint significance of the variables is highly significant, which implies that the underlying heterogeneity in the probability of choosing different pure behaviors among CEOs is important for explaining differences in firm performance. In terms of individual coefficients, eight are significantly negative relative to the sixth pure behavior, while two are not significantly different.

To gain more insight into what differences these capture, we first compare each pure behavior in the $K = 11$ model to pure behavior 0 in the baseline model by computing their Hellinger distance. This is a standard metric in the information theory literature, and lies in the $[0, 1]$ interval. We then transform the distances by computing their z-values, and also standardize the estimated coefficients in the productivity regression (treating the coefficient on pure behavior six as zero). Figure D.5 displays a scatterplot of these two series. There is clearly a positive correlation between distance from pure behavior 0 in the baseline model and a more positive association with productivity. Moreover, the behavior closest (furthest) from behavior 0 are among those least (most) associated

Figure D.5: Hellinger Distance from Pure Behavior 0 vs. Coefficient in Productivity Regression



with high firm performance. In this way, the behavioral differences in the $K = 11$ model appear to capture many of the important differences that also emerge from a more parsimonious model. At the same time, interpreting the content of each separate behavior is difficult, which serves to highlight the choice of $K = 2$ on the grounds of simplicity.

D.3 CEO Behavior Index and Management Practices

D.3.1 Management Data

We were able to match the CEO behavior index with information on management practices for 191 firms in our sample. The data are drawn from the World Management Survey (WMS).⁵¹ This uses an interview-based evaluation tool that defines 18 basic management practices and scores them from one (“worst practice”) to five (“best practice”) on a scoring grid. This evaluation tool was first developed by an international consulting firm, and scores these practices in three broad areas. First, *Monitoring*: how well do companies track what goes on inside their firms, and use this for continuous improvement? Second, *Target setting*: do companies set the right targets, track outcomes, and take appropriate action if the two are inconsistent? Third, *Incentives/people management*: are companies promoting and rewarding employees based on performance, and systematically trying to hire and retain their best employees? The survey was targeted at plant managers, who are senior enough to have an overview of management practices but not so senior as to be detached from day-to-day operations.

The data is collected through interviews with production plant managers using a “double-blind” technique. One part of this technique is that managers are not told in advance they are being scored or shown the scoring grid. They are only told they are being “interviewed about management practices for a piece of work”. The other side of the double blind technique is that the interviewers do not know anything about the performance of the firm. To survey is based

⁵¹More details can be found at <http://worldmanagementsurvey.org/>

on “open” questions. For example, on the first monitoring question we start by asking the open question, “tell me how you monitor your production process”, rather than closed questions such as “Do you monitor your production daily? [yes/no]”. We continue with open questions focused on actual practices and examples until the interviewer can make an accurate assessment of the firm’s practices. For example, the second question on that performance tracking dimension is, “What kinds of measures would you use to track performance?” and the third is “If I walked around your factory, could I tell how each person was performing?”.⁵² The other side of the double-blind technique is that interviewers are not told anything about the firm’s performance in advance. They are only provided with the company name, telephone number, and industry. Since the WMS randomly samples medium-sized manufacturing firms (employing between 50 and 5,000 workers) who are not usually reported in the business press, the interviewers will generally have not heard of these firms before, so they should have few preconceptions.

D.3.2 Management and CEO Behavior

We look at the cross sectional correlation between the management data and the CEO behavior index in Table D.3. Column 1 shows that the two variables are positively correlated (all regressions include log employment, country dummies and a set of noise controls). Columns 2 and 3 show that the correlation is stronger for the operational subcomponents of the management score, while they are positive but insignificant for the questions in the survey measuring people management processes. In columns 4-6 we investigate the relationship between firm performance, management and CEO behavior. This shows that the two indices are positively and independently correlated with firm productivity.

⁵²The full list of questions for the grid can be found at <http://worldmanagementsurvey.org/wp-content/images/2010/09/Manufacturing-Survey-Instrument.pdf>.

Table D.3: CEO Behavior Index and Management Practices

Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)
	CEO behavior index			Log(sales)		
CEO behavior index				0.606** (0.277)		0.550* (0.280)
Management (z-score)	0.054* (0.030)				0.207** (0.082)	0.187** (0.075)
Operations, Monitoring, Targets (z-score)		0.057* (0.029)				
People (zscore)			0.043 (0.034)			
log(employment)	0.106*** (0.033)	0.109*** (0.033)	0.104*** (0.034)	0.848*** (0.093)	0.880*** (0.067)	0.843*** (0.075)
Number of firms	191	191	191	142	142	142

Notes: *** (**) (*) denotes significance at the 1%, 5% and 10% level, respectively. Columns (1) to (3) include country dummies. Columns (4) to (6) include also year dummies. "Management" is the standardized value of the Bloom and Van Reenen (2007) management score, "Operations, Monitoring and Targets" and "People" are subcomponents of the main management score. Noise controls are a reliability score assigned by the interviewer at the end of the survey week and a dummy taking value one if the data was collected through the PA of the CEO, rather than the CEO himself, as well a variable capturing the reliability of the management score (as assessed by the interviewer) and the duration of the management interview. In columns (4) to (6) we include at most 5 years of data for each firm and build a simple average across output and all inputs over this period. Industry controls are 1 digit SIC dummies. All columns weighted by the week representativeness score assigned by the CEO at the end of the interview week. Errors clustered at the 2 digit SIC level.