CEO Behavior and Firm Performance[*]

Oriana Bandiera
London School of Economics

Stephen Hansen
University of Oxford

Andrea Prat
Columbia University

Raffaella Sadun
Harvard University

February 6, 2017

Abstract

We measure the behavior of 1,114 CEOs in Brazil, France, Germany, India, UK and US using a new methodology that combines (i) data on every activity the CEOs undertake during one work-week and (ii) a machine learning algorithm that projects these data onto one CEO behavior index. Low values of the index are associated with plant visits, and one-on-one meetings with production or suppliers, while high values correlate with meetings with high-level C-suite executives, and several functions, both from inside and outside the firm. We use these data to study the correlation between CEO behavior and firm performance within the framework of a firm-CEO assignment model. We show results consistent with significant CEO-firm matching frictions, which appear to be especially severe in lower income regions. The productivity loss generated by the mismatches is equal to 13% of the productivity gap between high and low income countries in our sample.

1 Introduction

The impact of CEOs on firm performance is at the core of many economic debates. The conventional wisdom, backed by a growing body of empirical evidence (Bertrand and Schoar 2003, Bennedsen et al. 2007, Kaplan et al. 2012) is that the identity of the CEO matters for firm performance. But what do CEOs actually do? Coase (1937) formalized the role of the CEO primarily as the coordination of activities performed by different individuals in the firm, and by relevant parties outside of it. Yet coordination can be achieved in different ways. Do different CEOs perform this role differently? And is there a "best practice", or do different circumstances call for different CEO behaviors?

In this paper we develop a new methodology to measure CEO behavior in large samples combining: (i) a survey that records each activity the CEOs undertake during one week; and (ii) a machine learning algorithm that projects the many dimensions of observed CEO behavior onto a low-dimensional behavior index. We use these data to study the correlation between CEO behavior and firm performance within the framework of a firm-CEO assignment model.

Our survey methodology is inspired by the classic study by Mintzberg (1973), who shadowed five CEOs over the course of one week. We scale up this methodology by calling the CEOs or their PAs to record the CEOs' diaries, rather than shadowing individuals directly. This approach allows us to collect comparable data on the behavior of 1,114 CEOs of manufacturing firms in six countries at different stages of development: Brazil, France, Germany, India, UK and the US. We record the same five features for each activity that involves the CEO's interactions with other people: 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.).[1]

Overall, we collect data on 42,233 activities covering an average of 50 working hours per CEO. Each of these activities can take one of the 4,253 combinations of the five features described above. We find that CEOs' behavior differs considerably along all five features. In particular, while most 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 activities. CEOs also differ in how they organize these interactions in terms of duration, number of people involved,

---

[1]In earlier work (Bandiera et al., 2013) we use 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 was led by a family CEO).

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.

To fully capture the heterogeneity and correlation structure of the data, we use an unsupervised Bayesian machine learning algorithm, the Latent Dirichlet Allocation (Blei et al. 2003) that projects the high-dimensional feature space onto a one-dimensional CEO behavior index. Low values of the index are associated with CEOs who spend more time with production, and one-on-one meetings with firm employees or suppliers. In contrast, high values of the index are associated with more time spent with C-suite executives, several participants and multiple functions from both inside and outside the firm.

While the diary data reveal that different CEOs behave differently, there is no theoretical reason to expect either type of behavior to be associated with better firm performance. Rather, the fact that such heterogeneity exists might be symptomatic of the fact that different types of firms require CEOs to behave differently. We formalize this idea in a simple model of firm-CEO assignment with, potentially, screening frictions and imperfect governance. The model is based on the assumption that firms and CEOs have heterogeneous types, and that a correct CEO-firm match results in better firm performance.

The model shows that a correlation between CEO behavior and firm performance can emerge in equilibrium for two very different reasons. First, the correlation may simply capture differences in firms' unobservable "innate" productivity levels - i.e. more productive firms may systematically hire CEOs with certain behavioral traits (or influence them to behave in a specific way). Importantly, however, a non-zero correlation may also reflect sufficiently severe screening and governance frictions. To see this, consider the scenario in which performance differences across firms are fully observable, so that the role of unobservable firm traits discuss above is irrelevant. In this case, if all CEOs are matched with the correct type of firm, the correlation between CEO behavior and performance conditional on firm observables would be zero. In contrast, if matching frictions are large, the mismatched CEOs would be associated with worse firm performance in the cross section. In other words, matching frictions are necessary to generate the variation in the data needed to identify a non-zero correlation between CEO behavior and performance.

Using firm level accounting data matched to our CEO index, we reject the null hypothesis of no correlation between CEO behavior and firm performance: high values of the CEO behavior index are significantly correlated with higher firm productivity, a key metric of firm performance (Syverson, 2011). A standard deviation increase in the CEO behavior

3

index is associated with a 0.07 log points increase in productivity, which is about 10% of the increase associated with a standard deviation increase in capital.

Focusing on the subset of firms for which we have productivity data before and after the appointment of the current CEO we establish that: (1) firms that eventually appoint high index and low index CEOs exhibits similar changes in productivity before appointment; (2) firms that hire a high index CEO experience a larger increase in productivity after the CEO appointment relative to firms that hire a low index CEO; (3) this differential effect on productivity materializes three years after the CEO is in office (even when we restrict the sample to the subset of CEOs whose behavior is measured in our survey in their first three years of tenure). These three patterns allay the concern that the correlations we observe are *solely* due to unobservable and time invariant firm traits, and are consistent with the presence of matching frictions.

Furthermore, we exploit the regional variation present in our data to show that the correlation between CEO behavior and firm performance is much stronger in lower income countries and lower income regions within countries. In light of the model, this finding can be interpreted in two ways: either firms in poorer regions are better able to shape the behavior of the CEO, or matching frictions are larger in these areas. The within-firm findings discussed above, as well as evidence pointing to the presence of significant shortage of human capital and poor quality of corporate governance in the two low income countries in our sample (Brazil and India), provide in our view further support to the matching friction interpretation.

Building on these results, in the last part of the paper we structurally estimate the model to shed light on two issues. First, is there a best practice in CEO behavior, that is: would all firms be more productive with high-index CEOs? Second, we quantify the productivity losses due to misassignement and compute the share of the productivity gap between rich and poor countries that can be attributed to CEO misassignements. Our estimates indicate that, while low-index CEOs are optimal for some of the sample firms, their supply generally overstrips demand, such that 17% of the firms end up with the "wrong" CEO. These bad matches are more frequent in poorer 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.

This study contributes a new method to measure CEO behavior in large samples and evidence on the link between CEO behavior and firm performance. The management literature provides some examples of time use analysis, but typically on much smaller samples and

for managers on lower rungs of the hierarchy.[2] In economics, our findings are complementary to the literature that studies the correlation between CEO personality traits and firm performance, rather than behavior. Kaplan et al. (2012) and Kaplan and Sorensen (2016) have detailed data on skills and personality traits of several CEOs candidates; they show the CEOs mostly differ along three dimensions: managerial talent, execution skills and interpersonal skills. Of these, only talent and execution skills correlate with firm performance but interpersonal skills increase the likelihood that the candidate is hired.[3] This is consistent with our model assumption that screening is imperfect, and that firms can end up hiring the wrong CEO. Our methodology is complementary to Mullins and Schoar (2013), who use self-reported survey questions to measure the management style and values of 800 CEOs in emerging economies. Their focus however differs from ours, as they aim to explain variation in style and values rather than the link with performance.

The paper is also related to a growing literature documenting the role of managers and management processes on firm performance (Bloom and Van Reenen (2007) and Bloom et al. (2016b)). The relationship between CEO behavior and firm performance that we identify is of the same order of magnitude as the correlation with management practices. Furthermore, for a subset of our firms we have both CEO behavior data and management scores (measured at middle managerial levels) and we are able to check that those variables are correlated, but retain independent explanatory power, thus suggesting that they might reflect two distinct channels through which managerial choices influence firm performance. Finally, we share the objective of Lippi and Schivardi (2014) to quantify the output reduction caused by distortions in the allocation of managerial talent.

The paper is organized as follows. Section 2 describes the data and the machine learning algorithm. Section 3 presents the assignment model, which informs the empirical analysis in section 4. Section 5 quantifies the share of mismatches and their consequences for productivity differentials. Section 6 concludes.

---

[2]For instance Kotter (1999) covers 15 general managers and Luthans (1988) 44 mostly middle managers. Professional surveys (e.g. McKinsey 2013) sometimes collect recall data on aggregate time use.

[3]Malmendier and Tate (2005) and Malmendier and Tate (2009) focus on overconfidence; they find that this is correlated with higher investment–cash flow sensitivity and mergers that destroy value.

2 Measuring CEO Behavior

2.1 The Sample

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-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.

Our sampling frame was randomly drawn from the set of firms classified in the manufacturing sector in the accounting database ORBIS, for a total of 6,527 eligible firms in 32 two-digit SIC industries.[4] Of these, 1,114 (17%) participated in the survey, [5] 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.[6] Finally, sample firms are very heterogeneous in size and sales values.

---

[4]This set was derived from 11,500 potential sample firms with available employment and sales data. We could find CEOs contact details for 7,744 firms and, of these, 1,217 later resulted not to be eligible. The reasons for non eligibility included recent bankruptcy or the company's not being in manufacturing. 310 of the 1,217 could not be contacted to verify eligibility before the project ended.

[5]This 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 reasons.

[6]The heterogeneity is mostly due to the distinction between family and professional CEOs, as the former have much longer tenures. In our sample 57% of the firms are owned by 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".

2.2 The Survey Tool

To measure CEOs' 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 to undertake that day as well as those actually done. [7] 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 screen-shot of the survey tool.[8] 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 reshape 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.[9] The average CEO has 202 15-minute time blocks, adding up to 50 hours per week on average.

2.3 CEO Behavior: Raw Data

For each activity we collect information on the following features: (1) type (e.g. meeting, travel, 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 or "insiders" (finance, marketing, etc.) and "outsiders" (clients, banks, etc.). We also ask whether each activity was undertaken in response to an emergency and we are able to measure the extent to which CEOs are able to plan ahead by comparing scheduled activities with the activities that eventually took place.

The CEOs in our sample largely set their agenda rather than responding to shocks. We infer this from three related facts. First, the comparison between the planned and the actual agenda shows that CEOs typically undertake all the activities scheduled for a given day - overall just under 10% of planned activities were cancelled. Moreover, only 4% of CEOs' time was devoted to dealing with emergencies.

---

[7]70% 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.

[8]The survey tool can also be found online on `www.executivetimeuse.org`.

[9]The non-work activities cover personal and family time during business hours.

To explore the raw data, we compute the average share of time that each CEO spends in different categories and represent the resulting distributions as box-and-whisker plots in Figure 1. A first notable finding from Panel A is 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) rather than working alone or traveling. This is in line with Coase's view of the CEO as someone who mostly coordinates other people's activities which cannot be easily mediated by markets. Moreover, Panel B shows most of this interactive time is spent with insiders which suggests that, at least in our sample, most CEOs chose to direct their attention primarily towards internal constituencies, rather than serving as "ambassadors" for their firms.[10] 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 a single function. There is no marked average tendency towards meeting with one or more than one person.
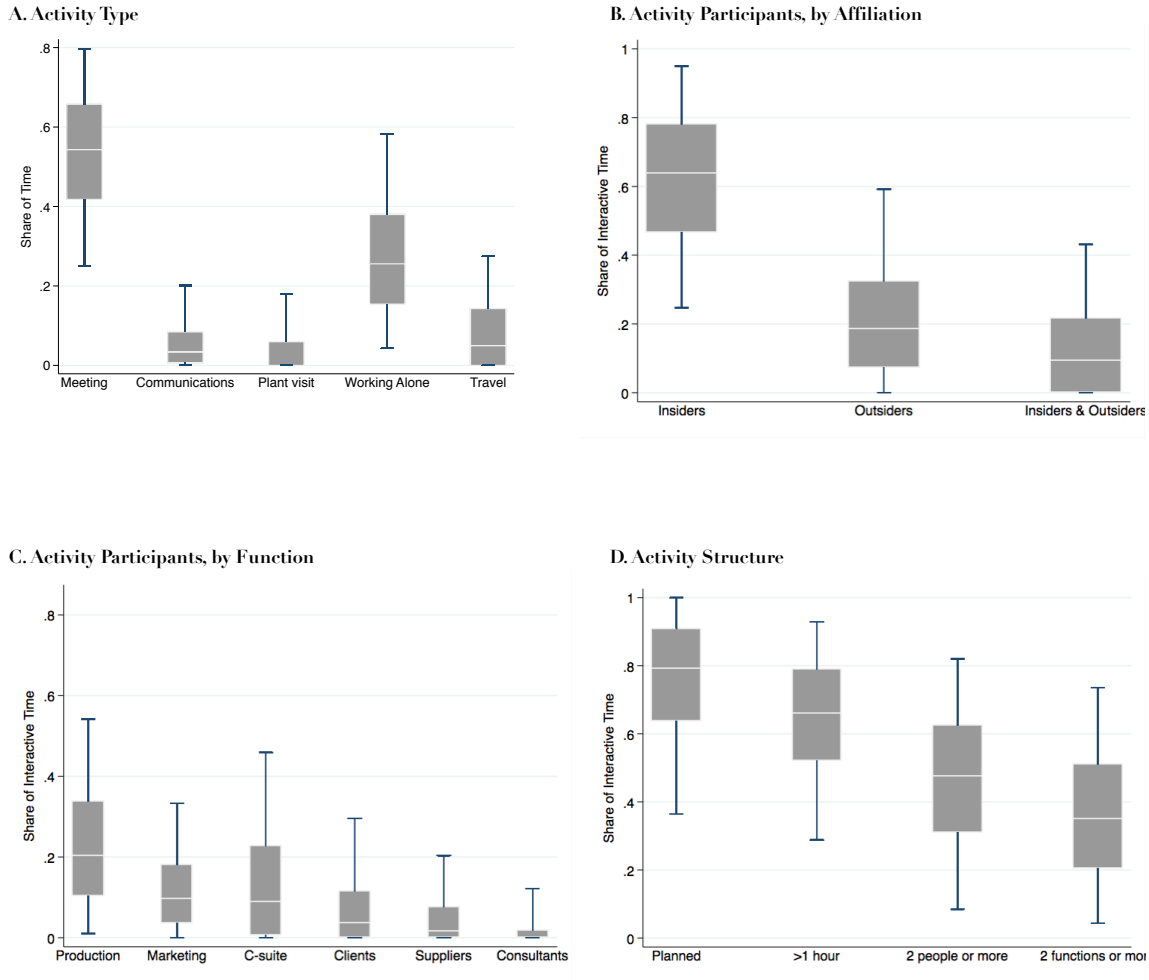
Figure 1 also reveals substantial heterogeneity underlying these average tendencies. For example, CEOs at the bottom quartile devote just over 40% of the time to meetings whereas those at the top quartile reach 65%; CEOs at the 3rd quartile devote over three times more time to production than their counterparts at the first quartile; and the interdecile ranges for time with two people or more and two functions or more are well over 50%.

Furthermore, there are 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-occurence. 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.

---

[10]See Nohria and Porter (2010) for a discussion of the role of the CEO as broker of both internal and external constituencies.

Figure 1: CEO Behavior: Raw Data

A. Activity Type

B. Activity Participants, by Affiliation

C. Activity Participants, by Function

D. Activity Structure

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.

9

Figure 2: CEO Behavior: Correlations

	Meeting	Plant visits	Communications	Planned	More 1 participant	More than 1 function	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	0.2063	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	0.1514	0.1371	0.3511	-0.3252	-0.0512	1					
Production	-0.1565	0.4114	-0.0823	-0.1157	0.0246	-0.1387	0.3455	-0.2917	-0.1092	-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	1			
Clients	-0.0945	0.0028	-0.028	0.0134	-0.1714	-0.1389	-0.4156	0.4275	0.0729	-0.1789	-0.134	-0.0455	1		
Suppliers	-0.0358	0.1089	-0.1622	-0.0381	-0.1702	-0.1703	-0.3264	0.3492	0.0384	-0.2192	0.0214	-0.0723	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	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.

2.4 The CEO Behavior Index

The most flexible way of representing the raw data is to describe each time block as a combination of the five distinct features measured in the data (type of activity, duration, planning horizon, number of participants, type of functions involved). We focus on the "interactive activities" – or combinations that involve time spent with at least one other person – which make up the bulk of CEOs' time and are most related to the coordinative role of the CEO that we seek to study. There are 4,253 unique interactive activities, but we drop those in fewer than 30 CEO's time use diaries since variation in rare activities is not informative of generalizable management behaviors. This leaves 654 combinations and 98,347 time blocks in our baseline empirical exercise.[11]

While the richness of the diary data allows us to describe CEO behavior in great detail, it makes standard econometric analysis of the relationship between CEO behavior and performance (and of the drivers of heterogeneity in behavior) challenging, as we effectively have almost as many variables (654, when we focus on the most frequent activities) as we have CEOs in our sample (1,114). At the same time, the patterns of co-occurence in time use suggest that one can view the high-dimensional raw activity data as being generated by a low-dimensional set of latent managerial behaviors. The next section discusses how we construct a scalar CEO behavior index using a widely-used machine learning algorithm.

2.4.1 Methodology

To reduce the dimensionality of the data we use the Latent Dirichlet Allocation (LDA) (Blei et al., 2003), an unsupervised machine learning algorithm for discrete data.[12] 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 linked explicitly to statistical parameters. LDA posits that the actual behavior of

---

[11]On average we drop 23% of the interactive activities in CEOs' diaries by excluding rare activities. In the appendix we explore the sensitivity of the results to dropping combinations present in fewer than 15 and 45 CEOs' time use diaries, and find little difference.

[12]An alternative approach would be to use a supervised learning algorithm that used variation in time use to directly predict firm performance. This would "force" the data to explain performance. Instead, we adopt a two-step approach in which we first identify the primary dimensions along which CEOs differ in their time use, and then examine whether variation in these dimensions explains differences in firm performance.

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 654 combinations of the five distinct features, which are potentially significantly more complex than the correlations between individual feature categories described in figure 2.

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 $\boldsymbol{\beta}^k$ over X that is common to all CEOs.[13]

In our baseline specification, we focus on the simplest possible case in which there exist only two possible pure behaviors: $\boldsymbol{\beta}^0$ and $\boldsymbol{\beta}^1$, 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.[14] We refer to the weight θ_i as the *behavior index* of CEO i.

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 $\boldsymbol{\beta}^0$ or $\boldsymbol{\beta}^1$). So, the probability that CEO i assigns to activity x_a is $\chi_a^i \equiv (1 - \theta_i)\beta_a^0 + \theta_i \beta_a^1$.

If we let $n_{i,a}$ be the number of times activity a appears in the time use of CEO i, then by independence the likelihood function for the model is simply $\prod_i \prod_a \chi_i^{n_{i,a}}$.[15] While in principle one can attempt to estimate $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ 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 $\boldsymbol{\beta}$ and θ terms. For estimating posteriors we follow the Markov Chain Monte Carlo (MCMC) approach of Griffiths and Steyvers (2004).[16] Here we discuss the estimated object

---

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

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

[15]The independence assumption of time blocks within a CEO is appropriate for our purpose to understand overall patterns 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.

[16]We set a uniform prior on θ_i–i.e. a symmetric Dirichlet with hyperparameter 1–and a symmetric Dirichlet

Figure 3: Data-Generating Process for Activities with Two Pure Behaviors
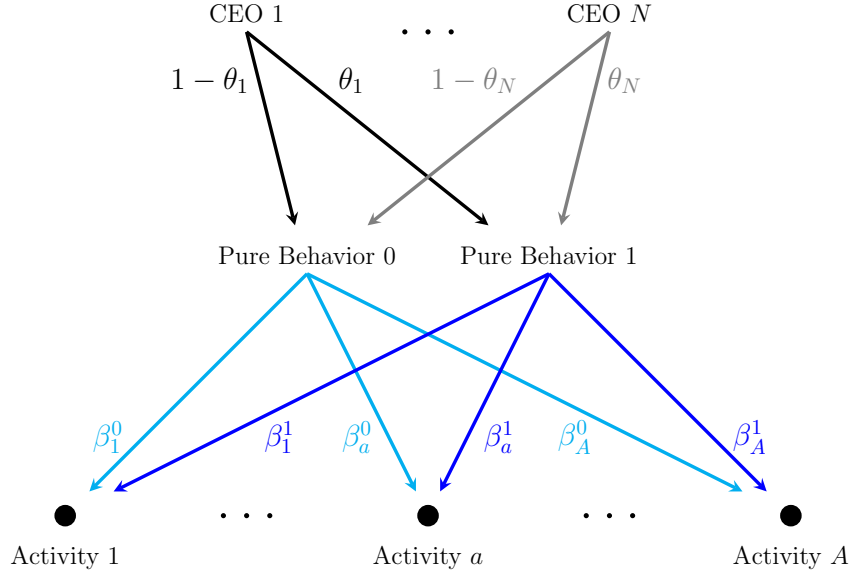
Notes: This figure provides a graphical representation of the data-generating process for the time-use data. First, CEO i chooses – independently for each individual unit of his time – one of the two pure behaviors according to a Bernoulli distribution with parameter θ_i. The observed activity for a unit of time is then drawn from the distribution over activities that the pure behavior defines.

of interest, which are the two estimated pure behaviors $\widehat{\boldsymbol{\beta}}^0$ and $\widehat{\boldsymbol{\beta}}^1$, as well as the estimated behavioral indices $\widehat{\boldsymbol{\theta}}_i$ for every CEO $i = 1, \ldots, N$.

2.4.2 Estimates

To illustrate differences in estimated pure behaviors, in Figure 4 we order the elements of X according to their estimated probability in $\widehat{\boldsymbol{\beta}}^0$ and then plot the estimated probabilities of each element of X in both behaviors. The figure shows that the combinations that are most likely in pure behavior 0 have low probability in pure behavior 1 and vice versa. Tables B.1 and B.2 list the five most common activities in each of the two behaviors.

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

---

with hyperparameter 0.1 on $\boldsymbol{\beta}^k$. This choice of hyperparameter promotes sparsity in the pure behaviors. Source code for implementation is available from https://github.com/sekhansen.

# Figure 4: Probabilities of Activities in Estimated Pure Behaviors
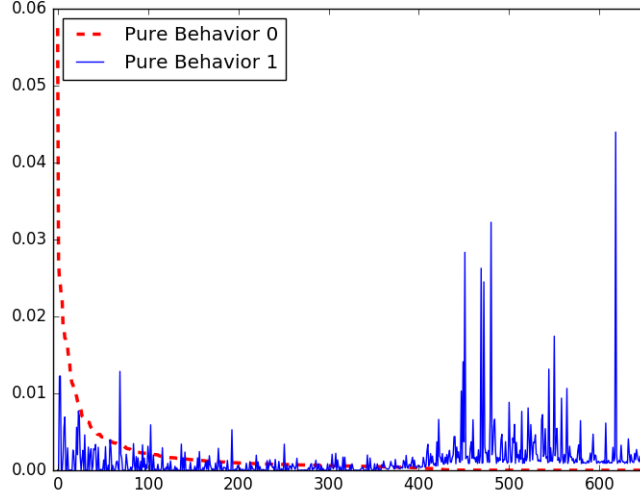
Notes: The dotted line plots the estimated probabilities of different activities in pure behavior 0, the solid line plots the estimated probabilities of different activities in pure behavior 1. The 654 different activities are ordered left to right in descending order of their estimated probability in pure behavior 0.

our baseline analysis. We perform 1,000 simulations, and store the estimated pure behaviors and behavior indices for each one. A standard distance metric for probability distributions is Hellinger distance, which lies in $[0, 1]$. The Hellinger distance between the distributions in figure 4 is 0.776. The maximum Hellinger distance in the simulated data is 0.412, while the mean is 0.327 and the minimum is 0.176. We conclude from this that the observed differences in pure behaviors are not consistent with there being no actual underlying heterogeneity in CEO time use. (For a plot of the simulated vs. actual distances, see the Appendix).

The two pure behaviors we estimate represent extremes. As discussed above, individual CEOs generate time use according to the behavioral index θ_i that gives the probability that any specific time block's feature combination is drawn from pure behavior 1. Figure 5 plots both the frequency and cumulative distributions of the $\widehat{\theta}_i$ 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, away from these extremes the distribution of the index is essentially uniform, and the bulk of CEOs draw their time use from both pure behaviors. This again highlights the value of using a model that allows CEOs to be associated with both estimated behaviors.

Figure 5: CEO Behavior and Index Distribution

(a) Frequency Distribution



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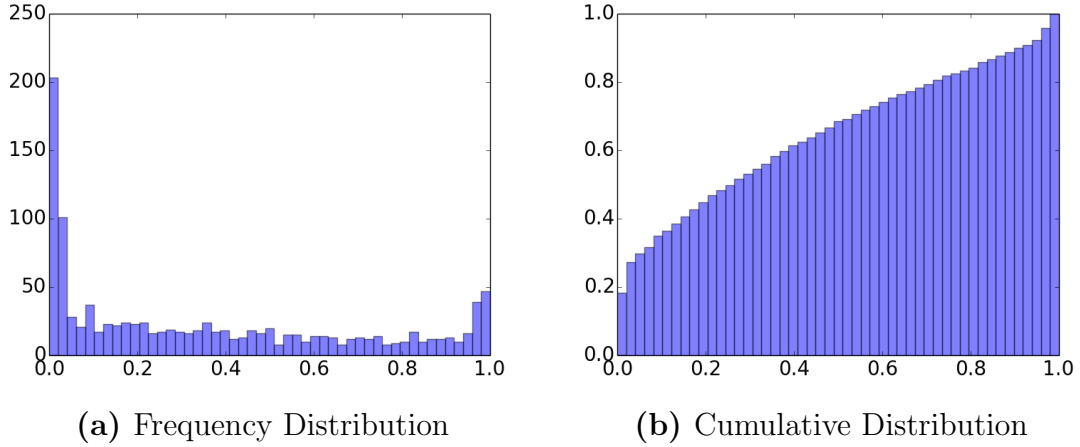
Notes: The left-hand side plot displays the number of CEOs with behavioral indices in each of 50 bins that divide the space [0,1] evenly. The right-hand side plot displays the cumulative percentage of CEOs with behavioral indices lying in these bins.

2.4.3 Interpretation of the CEO Behavior Index

Table 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. 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 category from both pure behaviors. Figure 6 displays the ratios of these marginal distributions (pure behavior 1 over pure behavior 0). A value of 1 indicates that each pure behavior generates the category with the same probability; a value of 0.5 indicates that pure behavior 1 is half as likely to generate the category; and so on.

Several striking differences emerge across pure behaviors. Pure behavior 1 is substantially more likely to engage in communications (phone calls, video conferences, etc), spend time with C-suite executives, bring together inside and outside functions, and bring together more than one function of any kind. Pure behavior 0 is more likely to devote time to plant visits, interactions with employees responsible for production, interactions with outsiders in general, and interactions with clients and suppliers in particular. Less marked differences arise for other categories. However, we have constructed simulated standard errors for the differences

15

Figure 6: Marginals

Panel C

Panel D
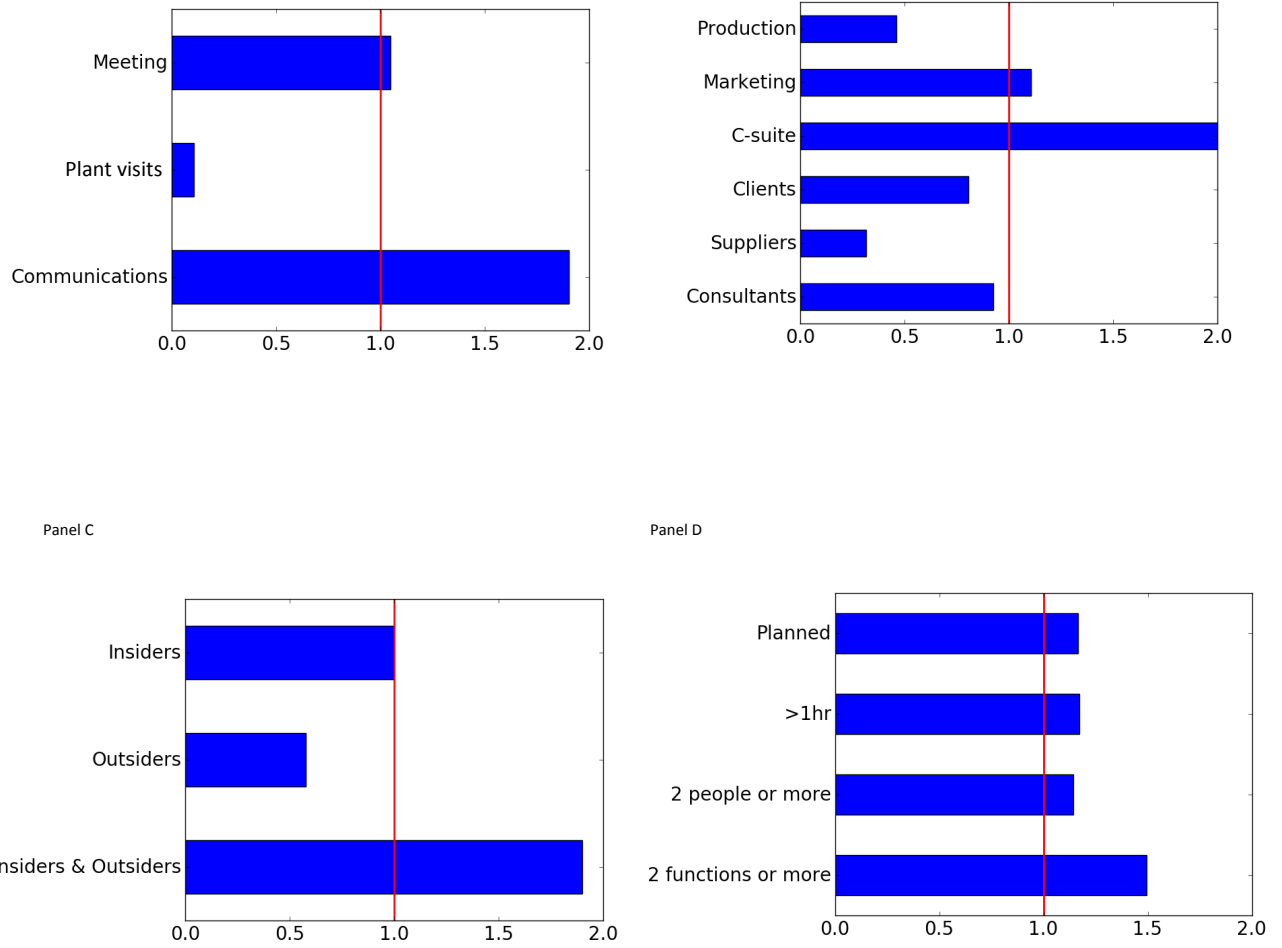
Notes: We generate these figures in two steps. First, we create marginal distributions for each behavior for each feature in activities. Then, for each category analyzed in figure 1, we report the probability of the category in behavior 1 over the probability in behavior 0. Panel D represents four separate marginal distributions. Each has two categories, so we report the ratio for only one.

in probabilities of each feature reported in the figure based on draws from the Markov chains used to estimated the reported means. All differences are highly significant except time spent with insiders. In sum, differences in the CEO behavior index mirror differences in the way CEOs coordinate the input of others: low index CEOs deal with one individual at the time, who is more likely to be of the production division; high index CEOs bring several individuals together, mostly at the top of the hierarchy.

CEOs do not seem to prefer either type of behavior. We find no correlation between the behavioral index and the CEO's satisfaction with their jobs (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.

The question of interest is whether CEO behavior is associated with differences in firm performance. A priori, there is no reason to expect either behavior to be more or less beneficial to all firms. Indeed, it is easy to imagine how different behaviors can be optimal under different circumstances. For instance, the coordinative role of CEOs may be more relevant in more complex organizations, either in terms of size of the organization or nature of activities undertaken, relative to a purely operational focus. We illustrate this point in Figure E.2 in the Appendix, in which we show that the CEO behavior index is correlated with various metrics proxying for firm size (number of employees, multinational and listed status) and organizational structure (e.g. whether the firm also employees a COO), as well as different CEO characteristics (work experience abroad, skills). These patterns of correlations illustrate a simple but crucial point: the allocation of CEOs with different behaviors across firms is correlated with firm characteristics, i.e. firms select CEOs on the basis of their own needs and observable CEO characteristics.

In the next section we present a simple theoretical framework to explicitly model the assignment of CEOs to firms, and to illustrate how this affects the interpretation of the association between CEO behavior and firm performance using cross-sectional data.

3   Modeling the Assignment of CEOs to Firms

This section develops a simple assignment model to interpret the cross-sectional correlation between CEO behavior and firm performance. The model allows for both directions of causality: firm performance can affect CEO behavior and CEO behavior can affect firm performance. The outcome of the model is the null hypothesis of zero correlation between CEO behavior and firm performance, which we test. The model specifies the conditions

under which this correlation is different from zero, and how a non-zero correlation may reflect the importance of firm level unobservables factors, and/or mismatches in the assignment of CEOs to firms.

3.1 Set-up

There are two possible behaviors that CEO i can adopt: $x_i = 0$ and $x_i = 1$. Once a CEO is hired, he decides how he is going to manage the firm that hired him. CEO i has a type $\tau_i \in \{0, 1\}$. Type 0 prefers behavior 0 to behavior 1. Namely, he incurs a cost of 0 if he selects behavior 0 and cost of $c > 0$ if he selects behavior 1. Type 1 is the converse: he incurs a cost of 0 if he selects behavior 1 and cost of c if he selects behavior 0. 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 $\tau_f \in \{0, 1\}$. The output of firm f assigned to CEO i is

$$y_{fi} = \lambda_f + \left(I_{\tau_f = x_i}\right) \Delta \tag{1}$$

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 the firm's type. 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. 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.

Equation 1 makes precise that the correlation between CEO behavior and firm performance can arise for two reasons: (1) Different firms have different baseline productivities and this affects what kind of CEO they look for; (2) If there are frictions on the CEO market, firms may not get the CEO they are looking for, which leads to productivity losses.

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

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

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. $w(y_{fi}) - I_{\tau_f = x_i} c$. 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 $\tau_i = x_i$, 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.

3.2   Pairing Firms and CEOs

Now that we know what happens once a CEO begins working for a firm, let us turn our attention to the assignment process. There is a mass 1 of firms. A proportion ϕ of them are of type 1, the remainder are of type 0. The pool of potential CEOs is larger than the pool of firms seeking a CEO. There is a mass $m >> 1$ of potential CEOs. Without loss of generality, assume that a proportion $\gamma \leq \phi$ of CEOs are of type 1. The remainder are of type 0. From now on, we refer to type 1 as the *scarce* CEO type and type 0 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-1 CEOs is actually more numerous than the share of type-0 firms. The model also nests the case of pure vertical differentiation, where no firm actually wants a type-0 CEO. This happens 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-0 or type-1 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.[17] This last assumption distinguishes our approach from existing theories of

---

[17]The implicit assumption is that CEOs have private information about their types, while firms' types are

manager-firm assignment, where the matching process is assumed to be frictionless, and the resulting allocation of managerial talent achieves productive efficiency (Gabaix and Landier (2008), Tervio (2008), Bandiera et al. (2015)). One 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.

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.

The following proposition makes precise the conditions under which there is no correlation between CEO behavior and firm performance. Define residual productivity as total productivity minus type-specific baseline productivity: $y_{fi} = \lambda_f$.

Proposition 1 *Firms led by CEOs who choose behavior 1 and those led by CEOs who choose behavior 0 have equal residual 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, either because of unobservable firm traits or because of mismatches. The following proposition characterizes how mismatches can occur in equilibrium.

Proposition 2 *If the screening process is sufficiently unreliable, governance is sufficiently poor, and one CEO type is sufficiently abundant,*[18] *then in equilibrium:*

---

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.

[18]Formally, this is given by the conditions: $B < \frac{c}{\Delta}$, and

$$\rho < \frac{\phi - \gamma}{\phi - \gamma\phi}.$$

- *All scarce-type CEOs are correctly matched;*

- *Some abundant-type CEOs are mismatched;*

- *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 where $\phi = 1$, that is when no firm demands type 0 CEOs, these 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 sufficient to cover all firms that prefer that CEO type ($m >> 1$), it would be possible to give all firms their preferred type and thus increase overall production.[19]

4 CEO Behavior and Firm Performance: Evidence

4.1 From Theory to Data

As described in Equation (1), the output of firm f assigned to CEO i depends on firm type and CEO behavior. We observe CEOs' behavior and firm performance but not firm type. Then the observed difference in performance between firms that hire a type 1 CEO and those that hire a type 0 CEO is:

$$y_{.1} - y_{.0} = s_1[\lambda_1 + \Delta] + (1 - s_1)\lambda_0 - [s_0[\lambda_0 + \Delta] + (1 - s_0)\lambda_1]$$

---

[19]If side transfers were feasible, this would also be a Pareto improvement as a type-1 CEO assigned to type-0 firm generates a higher bilateral surplus than a type-0 CEO matched with a type-1 firm, and the new firm-CEO pair could therefore compensate the now unemployed type-0 CEO for her job loss.

where s_i is the share of CEOs who are correctly matched, thus, for instance, the average performance of firms led by type-1 CEOs is equal to the performance of type-1 firms $(\lambda_1 + \Delta)$ weighted by the share of type-1 CEOs who are correctly matched (s_1) plus the performance of mismatched firms 0 (λ_0) weighted by the share of type-1 CEOs who are wrongly matched $(1 - s_1)$. Simplifying yields:

$$y_{.1} - y_{.0} = (s_1 + s_0 - 1)(\lambda_1 - \lambda_0) + (s_1 - s_0)\triangle \tag{2}$$

Equation (2) makes precise that the cross-sectional correlation, or the difference in means of productivity for the discrete case, captures both differences in firms' "innate" productivity levels $(\lambda_1 - \lambda_0)$ and the extent of the mismatch times its costs $(s_1 - s_0)\triangle$.

Note that a priori $(\lambda_1 - \lambda_0)$ can be positive or negative. Furthermore, if there are no frictions and the CEOs are all correctly matched, then $s_1 = s_0 = 1$ and the difference in means only captures average productivity differences between firms that hire a CEO 1 and those that hire a CEO 0. Conversely, $(s_1 - s_0) > 0$ if and only if type-1 CEOs are more likely to be correctly matched - that is, if they are more scarce - than types 0, otherwise it is negative.

4.2 Cross-Sectional Correlations

To estimate the correlation between CEO behavior and firm performance, we combine our CEO behavior data with accounting information extracted from the ORBIS database. 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.[20]

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:

$$y_{ifts} = \alpha\widehat{\theta}_i + \delta^E e_{ft} + \delta^K k_{ft} + \delta^M m_{ft} + \zeta_t + \eta_s + \varepsilon_{ifts} \tag{3}$$

where y_{ifts} is the log sales (in constant 2010 USD) of firm f, led by CEO i, in period t and sector s. To smooth out short run fluctuations and reduce measurement error in performance, y_{ifts} is average sales computed over up to the three most recent years pre-dating

---

[20]Of these: 29 did not report sales information at all; 128 were dropped due to extreme values in the productivity data, 37 had data that referred only to years in which the CEO was not in office or in office for less than one year. See the data Appendix for more details.

the survey, conditional on the CEO being in office, but the results are very similar when we use yearly data and cluster the standard errors by firm (Appendix Table E.2, column 2).[21] $\widehat{\theta}_i$ is the estimated 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.[22] We include country and year dummies throughout, as well as a set of interview noise controls listed in the notes to Table 1.

Basic productivity results Column 1, Table 1 shows the estimates of equation (3) just 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 sample of 920 firms. The estimate of α is positive and we can reject the null of zero correlation between firm labor productivity and the CEO behavior index at the 1% level.[23] 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 and is still significant at the 5% level in the subsample. The magnitude of the CEO behavior index is about 10% of the effect of a one standard deviation increase in capital.[24] In Column 3 we add materials, which further restricts the sample to 448 firms. Even in this even smaller sample the effect of standard inputs has the expected magnitude and is precisely estimated, but the magnitude and the precision of coefficient of the CEO behavior index remains unchanged. 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 (E.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.

---

[21]In practice we have 3 years for 58% of the sample, 2 years for 24% and 1 year for the rest.

[22]Since the data is averaged over three years, year dummies are set as the rounded average year for which the performance data is available.

[23]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. To this purpose, in Table E.1 we show the results of equation (3) 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. In addition, we obtain the same results when using more standard dimensionality reduction techniques such as k-means and principal components (see Table (E.2)).

[24]To 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*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).

Table 1: CEO behavior and Firm Performance

Dependent Variable	(1)	(2)	(3)	(4)	(5)	(6)
			Log(sales)			Profits/Emp
CEO behavior index	0.343***	0.227**	0.322***	0.641**	0.505**	9.987**
	(0.108)	(0.111)	(0.121)	(0.279)	(0.235)	(3.994)
log(employment)	0.889***	0.555***	0.346***	0.339**	0.804***	0.067
	(0.040)	(0.066)	(0.099)	(0.152)	(0.075)	(0.052)
log(capital)		0.387***	0.188***	0.194*		
		(0.042)	(0.056)	(0.098)		
log(materials)			0.447***	0.421***		
			(0.073)	(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.

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. (2016a)).[25] We were able to gather management data for 191 firms in our CEO sample. The CEO behavior index is indeed 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.[26] For the 156 firms for which we could match the management and CEO behavior data with accounting information, we find both variables to be independently correlated with productivity. The coefficients reported in Column (5) 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.[27] We also find a similar pattern when controlling for other potentially confounding firm and CEO characteristics, and for CEO total hours worked; Table E.2 shows that including other firm or CEO observables hardly changes the magnitude of the CEO behavior index.

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 matched with high index CEO. 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 approximatively $3,010 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

---

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

[26]See Appendix Table E.3 for details. To our knowledge, this is the first time that data on middle level management practices and information on CEO behavior is systematically analyzed. 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. 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.

[27]When 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. (2016a) in the full management sample (0.15).

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 the estimated parameters to predict the held-out data.

This approach shows that a model with eleven pure behaviors is best at prediction. However, as we discuss in the Appendix, the pure behavior with the largest correlation with productivity is actually among the most similar to pure behavior 1 used in the simple $K = 2$ model. We conclude from this exercise that – in spite of its simplicity – the model with two pure behaviors is actually able to capture many of the salient performance-related distinctions in CEO behavior.[28]

4.3   Why is CEO Behavior Correlated with Firm Performance?

The model makes it clear that the the results in Table 1 - i.e. a non-zero correlation between CEO behavior and firm performance - can be due to: (1) unobservable firm traits that determine both the need for a high index CEO and firm performance; (2) mismatches in the assignment of CEOs to firms, so that some of the firms requiring high index CEOs end up - due to an imperfect allocation mechanism - with low index CEOs. In other words, the correlation might be capturing both the effect of firm performance on CEO behavior or the

---

[28]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).

effect of CEO behavior on firm performance.

Due to the cross sectional nature of the CEO behavior data, we cannot separately identify these two channels, but we can provide some evidence on their relative importance exploiting both within and cross sectional firm heterogeneity.

4.3.1 Firm Performance Before and After the CEO Appointment

To provide evidence on the relevance of time invariant and unobservable firm-specific factors in driving the correlation between CEO behavior and firm performance, we exploit the fact that, for a subset of CEOs in our sample who have been recently appointed, we can observe firm performance before and after the CEO appointment. Therefore, while our survey measures CEO behavior just at one point in time and within a single firm, we can use this pseudo-panel to test whether firms that eventually appoint a high-index CEO (i.e., with a CEO behavior index greater than 0.5) have different productivity levels or trends relative to firms that eventually appoint a low-index CEO (CEO behavior index less than 0.5). This analysis is informative of the practical relevance of firms' unobservable traits, because it allows us to measure whether the appointment of a high-index CEO improves productivity within the same firm controlling for time invariant firm unobservables, and whether the appointment of a high-index CEO is preceded, and thus possibly driven, by a period of high growth.

To implement this approach, we restrict the sample to the 204 firms that have accounting data within a 5 year interval *both* before and after CEO appointment. To start with, Column 1 show that these firms are representative of the larger sample in terms of the correlation between the CEO behavior index and performance. The correlation is 0.258 (standard error 0.073) for firms that do not belong to the subsample, and the interaction between the CEO behavior dummy and the dummy denoting the subsample equals -0.112 and is not precisely estimated. Next, since we have data on multiple years before appointment we can test whether the parallel trend assumption, namely that treatment (here firms that eventually appoint behavior 1 CEO) and control (here firms that eventually appoint behavior 0 CEO) have the same trend before appointment. Column 2 shows that this is indeed the case.

Finally, we look at the within firm changes in productivity for firms with a behavior 1 vs. behavior 0 CEO, by estimating the following difference in differences model:

$$y_{ft} = \alpha A_t + \beta A_t \widetilde{\theta}_f + \delta^E e_{ft} + \zeta_t + \eta_f + \varepsilon_{it} \tag{4}$$

Where $t = 0$ the year the CEO is appointed, and we estimate (4) for $t \in [-5, +5]$. η_f

27

Table 2: 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 appointment of the current CEO	Before and after appointment of the current CEO (after current CEO divided in 2 subperiods)	Before and after appointment of the current CEO (after current CEO 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	572	408	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 1, column 1. The sample in columns 1-4 includes the set firms with at least one year of non missing productivity data 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. 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 build firm fixed effects. Columns 4 and 5 split the after CEO appointment period into 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.

28

are firm fixed effects, $A_t = 1$ after appointment, and $\widetilde{\theta}_f = 1$ if firm f eventually appoints a CEO with a high behavioral index. The coefficient of interest is β, which measures whether firms that eventually appoint a CEO with behavior 1 experience a change in productivity after the CEO is in office. Note that, since we do not know the behavior of the previous CEO, this is a lower bound on the effect of switching from behavior 0 to behavior 1 CEOs, because at least part of these firms would have had a behavior 1 CEO before the current appointment. Column 3 shows that, after the CEO is in office, firms that appoint a behavior 1 CEO experience a relative 0.097 log points increase in productivity (standard error 0.041). Column 4 splits the post period into two sub periods: 1-2 and 3-5 years after appointment. The results suggest that the correlation materializes three years after appointment.

Taken together, the results in Table 2 rule out that the correlation is solely driven by differences in time invariant firm level unobservables and differences in pre-appointment trends. Had this been the case, we would have detected a difference between firms that eventually appoint a behavior 0 CEO and those that appoint a behavior 1 CEO also *before* their appointments. Furthermore, the delay in the productivity increase is compatible with the idea that the actions of the new CEO may take time to affect the production process - in Appendix we show a simple dynamic extension of the model developed in Section 3 which is compatible with these dynamic patterns.[29] An alternative explanation is that CEOs with longer tenure have had more time to learn the needs of the firm, and that this experience may generate the correlation between CEO behavior and performance. To shed light on this, Column 5 restricts the sample to the 102 firms whose current CEO had been in office for less than three years when we interviewed the CEO, and for which we have firm performance data three years after the appointment.[30] The fact that the results hold, and are actually stronger in this smaller sample of less experienced CEOs points to the fact that the earlier correlation was not due to learning effects.

It is important to note that, while the findings attenuate concerns related to time in-varying correlated unobservables, the data at hand does not allow us to rule out other time-varying unobservables. For example, firms may appoint CEOs of a certain type in anticipation of events that will take place three years after the appointment date (e.g. the

---

[29]The existence of significant organizational inertia within firms has been a central theme in the management literature (Cyert (1963)), and is central to a recent strand of the organizational economics literature. For example, in the model of Halac and Prat (2016), it takes time for a corporate leader to change the existing management practice and to affect the company's culture. Empirically, Bloom et al. (2016b) estimate adjustment costs in managerial capital of similar magnitude to the ones estimated for physical capital.

[30]We do not know whether the CEOs we interview are still in office after year 3. To the extent that some of them are not, the estimates are biased downwards.

CEO appointment may coincide with a change in firm strategy which may take time to materialize). However, it is worthwhile noting that this would require extraordinary board capabilities. First, the board hiring strategy must be such that they dismiss CEOs who are currently performing adequately, because they might perform poorly four years ahead, which requires exceptional abilities in both forecasting firm's needs and to remove inadequate CEOs. Second, this strategy implies that boards are willing to bear losses due to the wrong match, or CEOs are willing to behave sub-optimally for three years (33% of the average CEO tenure in our sample), rather than appoint a new CEO when he is actually needed. Furthermore, this explanation is hard to square with the cross country and cross regional evidence in the relationship between CEO behavior and productivity, which we discuss in the next section.

### 4.3.2 Variation along the Development Path

As discussed in Proposition 2, the other reason to observe a cross-section correlation between CEO behavior and firm performance is the presence of CEO-firms matching frictions. If at all relevant, we expect matching frictions to be more severe in less developed regions due to both differences in the supply of managerial capital,[31] and in the quality of corporate governance, which affects both the selection and the dismissal of wrongly matched CEOs.[32]

Consistent with this idea, Figure 7 plots productivity residuals (estimated as in Column (1), Table 1) by country income levels (Brazil and India for low/middle income countries vs. France, Germany, UK and US for high income countries) and CEO behavior, where we once more classify CEOs as behavior 1 if their CEO behavior index is ≥ 0.50, and 0 otherwise. The figure shows a marked difference between the two sets of countries. In the low/middle income countries the density of behavior 1 CEOs first-order stochastically dominates the density of the other CEOs, while in the high income countries the densities are much more similar. This is consistent with the idea that observed CEO behavior in high-income countries is much more likely to match with the firm's needs than in low/middle-income countries, i.e. that assignment frictions are substantially lower.

---

[31]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, and show the presence of a large and significant association with regional income per capita and firm level productivity. 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. (2016b).

[32]Porta et al. (1999) 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

Figure 7: Productivity Residuals Densities by Development Level

(a) Low/Middle-Income Countries **(b) High-Income Countries**
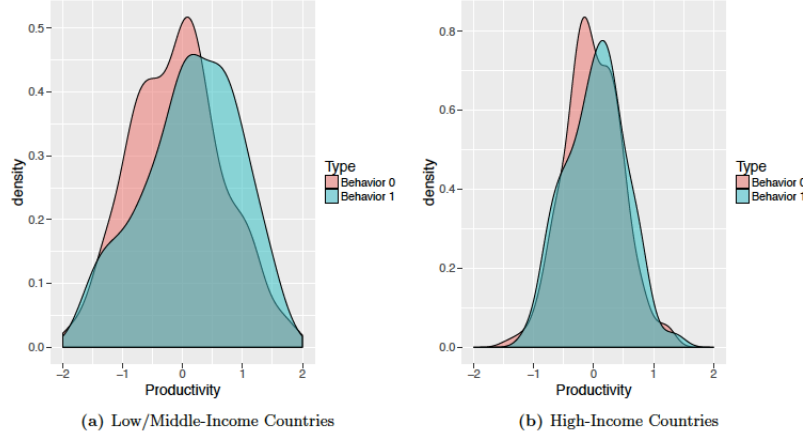
Notes: This figure plots the estimated kernel densities of the productivity residuals $\widehat{\varepsilon}_{if}$ (demeaned at the country level) by regional income and by observed CEO behavior. The low/middle income countries are Brazil and India, while the high income countries are France, Germany, the United Kingdom, and the United States.

To provide further evidence on this point, we examine whether the coefficient on the CEO behavior index in Equation (3) varies according to the level of development in the country, or in the region within a country. The results of this analysis are shown in Table 3. We start by simply reporting the baseline results shown in Column 1 in Table 1. In Column 2 we interact the continuos CEO behavior index with an indicator that equals 1 if the firm is located in a high income country (France, Germany, UK and US). The coefficient on the CEO behavior index for low/middle income countries is 0.631 and precisely estimated (standard error 0.167) in these areas, whereas it is much lower (0.631-0.619=0.012) and we cannot reject the null that it equals zero in high income countries. This pattern also holds when we use variation in GDP per capita across regions within the same country. This is shown in Column 3, where we interact the CEO behavior index interacted with the log of regional GDP level drawn from Gennaioli et al. (2013). The results show that the coefficient on the CEO behavior index declines as the level of regional GDP per capital increases; its magnitude implies that we cannot reject the null of zero correlation in the richest regions. Column 4 shows that this result holds also when we control for regional dummies to control for unobservable differences in local characteristics.[33]

(both 16%) vs. France (29%), Germany (40%), UK (35%) and US (27%).

[33]Note that if the correlation between CEO behavior and firm performance were mostly driven by firm traits, we would expect it to be stronger in higher income areas, where shareholders have in principle more control over their CEOs thanks to better governance and contract enforcement. In this light, the cross-

Table 3: Heterogeneity across Countries and Regions

Dependent variable	(1)	(2)	(3)	(4)
	Log(sales)			
log(employment)	0.889***	0.888***	0.889***	0.879***
	(0.040)	(0.041)	(0.036)	(0.041)
CEO Behavior Index	0.343***	0.631***	2.996***	3.366***
	(0.108)	(0.167)	(0.811)	(1.021)
Dummy high income country		0.375		
		(0.495)		
High income country*CEO behavior index		-0.619***		
		(0.204)		
log(regional GDP)			0.114	
			(0.103)	
log(regional GDP)*CEO Behavior index			-0.278***	-0.316***
			(0.083)	(0.105)
Number of observations (firms)	920	920	919	919
Region fixed effects	No	No	No	Yes
Cluster	Sic 2	Sic 2	Region	Region

Notes: *** (**) (*) denotes significance at the 1%, 5% and 10% level, respectively. All columns include the same controls used in Table 1, column 1. "High income country" is a dummy taking value 1 for firms located in France, Germany, UK or US. "Log regional income per capita" is in current purchasing-power-parity (PPP) dollars and is drawn from Gennaioli et al (2013). Errors clustered as noted. All columns weighted by the week representativeness score assigned by the CEO at the end of the interview week. Column (4) includes regional dummies.

Overall, the evidence presented in this section is consistent with the presence of significant differences in managerial matching frictions across countries. These findings are complementary to Alder (2016), who develops a similar idea in the context of a calibrated macro model and finds deviations from assortative matching to have large effects on output and TFP across countries. More broadly, this evidence is in line with a recent strand of the literature (see, for example, Hsieh and Klenow (2009)) which sees misallocation of factors of production - in this case, managerial inputs - as a key driver of differences in economic development across countries.

5   The Productivity Cost of Misallocation

The mismatch between firms and CEOs discussed in the previous section is consistent with a pure vertical differentiation scenario, in which behavior-0 CEOs are universally worse than behavior-1 CEOs (or $\phi = 1$ in the model), as well as with horizontal differentiation, in which both behaviors can be appropriate depending on firm circumstances ($0 < \phi < 1$).

This section builds a structural model to formally test these alternatives using cross-sectional data, and finds evidence consistent with horizontal differentiation in CEO behavior. We then use the model to quantify the productivity loss due to misassignement, and compute the share of the productivity gap between rich and poor countries that can be attributed to CEO-firms mismatches.

5.1   Structural Model

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 1. However, we omit the country fixed effects, and instead estimate them as part of firms' baseline productivities below. We denote the residual of firm f run by CEO i as $\widehat{\varepsilon}_{if}$.[34]

In line with the theory presented in Section 3, we adopt the statistical model $\widehat{\varepsilon}_{if} = \lambda_f + \mathbb{1}(\tau_f = x_i)\Delta + v_{if}$, where λ_f is a "baseline" productivity; $\tau_f \in \{0, 1\}$ is the firm's type; $x_i \in \{0, 1\}$ 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.

---

country and cross-regional evidence presented in this section suggest that the effect of mismatching may actually overweigh that driven by unobservable firm traits.

[34]To maintain comparability in the pooled vs. regional results, 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.

Throughout this section we assume the conditions of Proposition 2 hold. To obtain an empirical proxy of x_i we use $\widehat{x}_i = \mathbb{1}\left(\widehat{\theta}_i >= 0.5\right)$. While we treat $\widehat{x}_i$ as observed data, τ_f is a random variable. Since all CEOs with $\widehat{x}_i = 1$ are correctly assigned, whenever we observe $\widehat{x}_i = 1$ we also must have $\tau_f = 1$.

In contrast, only a share q of CEOs with $\widehat{x}_i = 0$ is correctly assigned: when we observe $\widehat{x}_i = 0$, $\tau_f = 0$ with probability $0 \le q \le 1$; otherwise, with probability $1 - q$ the CEO is misassigned and $\tau_f = 1$. Note that the model nests both pure vertical and pure horizontal differentiation. In the case of pure vertical differentiation $q = 0$; that is, all CEOs with $\widehat{x}_i = 0$ are mis-assigned because there is no demand for their behavior. Vice versa, in the case of pure horizontal differentiation $q = 1$; that is, all CEOs with $\widehat{x}_i = 0$ are assigned to firms that demand their behavior. 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 $x_{c_f,1} = A + x_{c_f,0}$ so that the baseline productivity of type 1 firms is that of type 0 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

$$
\prod_{f \in \Theta(0)} \left\{ \begin{array}{l} \frac{q}{\sqrt{2\pi}\sigma_{H,c_f}} \exp\left[-\frac{1}{2\sigma_{H,c_f}^2} \left(\widehat{\varepsilon}_{if} - x_{c_f,0} - \Delta\right)^2 \right] + \\ \frac{1-q}{\sqrt{2\pi}\sigma_{L,c_f}} \exp\left[-\frac{1}{2\sigma_{L,c_f}^2} \left(\widehat{\varepsilon}_{if} - A - x_{c_f,0}\right)^2 \right] \end{array} \right\} \times
$$

$$
\prod_{f \in \Theta(1)} \frac{1}{\sqrt{2\pi}\sigma_{H,c_f}} \exp\left[-\frac{1}{2\sigma_{H,c_f}^2} \left(\widehat{\varepsilon}_{if} - A - x_{c_f,0} - \Delta\right)^2 \right]. \tag{5}
$$

where $\Theta(0)$ and $\Theta(1)$ are the sets of firms managed by CEOs with behaviors 0 and 1, respectively. Behavior-1 CEOs are always efficiently assigned to type-1 firms and their residuals are drawn from a normal distribution with mean $A + x_{c_f,0} + \Delta$; in contrast, firms managed by behavior-0 CEOs have their residuals drawn from a mixture of two normals, one with mean $x_{c_f,0} + \Delta$ if the assignment is efficient and another with mean $A + x_{c_f,1}$ if the assignment is inefficient. The mixing probability is simply q, the probability that CEOs with behavior 0 are assigned to type-0 firms. We use the EM algorithm to maximize (5).

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 behavior-0 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 behavior-1 CEOs. (Note that in the E-step we explicitly infer the probability that behavior 0 CEOs are efficiently assigned, which allows us to then estimate parameters in the M-step).

As is standard in structural estimation, the log likelihood is defined under the assumptions of the theoretical model, namely that $\Delta > 0$, and that type-1 CEOs are scarce and all correctly matched; 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 matched firm-CEO pairs are more productive. Of course, nothing in the structural 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 $\Delta > 0$.

The estimate of Δ is 0.532, which implies that he loss associated with an incorrect assignment 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 q is 0.744. To test whether the data are consistent with pure vertical differentiation, we impose the restriction $q = 0$ in (5), which a likelihood ratio test rejects with a p-value of 0.00202. This is an important finding since the reduced-form results cannot distinguish pure vertical differentiation from a model with firm types. The key underlying property of the data that lets us test $q = 0$ is that under this restriction behavior-1 CEOs uniformly outperform behavior-0 CEOs. We can reject this in favor of a mixture model with $q > 0$ since we observe a large fraction of behavior-0 CEOs whose performance is similar to that of behavior-1 CEOs. Also, note that once we reject $q = 0$ we must necessarily reject $q = 1$ because in both cases we fit two mean parameters, while in the former we also fit two variance parameters (per country), whereas the latter is more restricted with one variance parameter. Overall, then, 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).

5.2 Misallocation and Development

The structural model also allows us to quantify the productivity losses due to the misallocation of CEOs to firms. Following Section 4.2.2, we use the distinction between high and low/middle income regions to proxy for different intensities in matching frictions, and we quantify the productivity losses in the sample as a whole, and separately for the two subsets of countries.

To quantify the share of mismatches we first derive ϕ , i.e. the share of type-1 firms, from the market clearing condition. Overall we observe a share $\widehat{s} = 0.347$ of CEOs who adopt behavior 1. We must then have $\phi = \widehat{s} + (1 - \widehat{s})(1 - q)$. The right-hand side of this expression is the total share of CEOs assigned to type-1 firms: all behavior-1 CEOs and a portion $1 - q$ of behavior 0 CEOs. Plugging in for $\widehat{s}$ and q, we obtain $\phi = 0.514$ so that slightly over half of firms have type 1. This in turn implies that a share $\phi - \widehat{s} = 0.168$ 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 q parameter in the likelihood function (5) to vary according to whether the firm operates 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. Confirming the notable differences in figure 7, behavior 0 CEOs in low/middle income countries are efficiently assigned with probability 0.546, while the corresponding probability for behavior 0 CEOs in high income countries is 0.893. The derived parameters in the table are obtained using the same steps as described above.

One possible explanation for these different probabilities across countries is that firms in high-income countries have higher demand for coordinating rather than micro-managing behavior. Indeed, consistent with this idea, the data shows a much larger share of behavior-1 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-1 firms–are in fact very similar in both regions (if anything there is slightly higher demand for coordinating behavior in poorer countries).[35]

Instead, the main difference between regions emerging from the structural exercise is that type-1 firms in low/middle-income countries are unable to locate and hire behavior-1 CEOs. It is important to reiterate that this is not necessarily due to scarcity of type-1 CEOs in

---

[35]Also of note is that the ϕ estimates suggest that horizontal differentiation exists in both regions, since firms demand substantial shares of both behaviors. We have repeated the same chi-squared tests for restrictions on q 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.

Table 4: Estimation Results by Region

	Estimated Parameters		Derived Parameters		
	Δ	q	$\widehat{s}$	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 high-index CEOs in each region. The fourth is the value of ϕ consistent with market-clearing given q and the observed shared of high-index CEOs, while the fifth is the difference between the fourth and third, as this gives the share of type-1 firms run by behavior-0 CEOs.

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 matched type-1 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.[36]

To conclude, we use our estimates to quantify how much productivity in India and Brazil would increase if the assignment process were as efficient as in the richer countries in the sample. This implies building a counterfactual where q increases from 0.546 to 0.893, which requires the share of high-index CEOs to increase from 0.216 to 0.521 to maintain market clearing, and which yields a drop in the share of mismatches from 0.356 to 0.051. Given that the productivity difference between a good and a bad match is now estimated at a somewhat higher value of 0.667, productivity would increase by 0.203 log points.

We benchmark this magnitude against the macro differences in labor productivity across countries observed in the time interval covered by our survey and productivity data (2010-2014) 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

---

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

up to 13% of the cross country differences in labor productivity.[37]

6 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 a number of dimensions. Guided by a simple firm-CEO assignment model, we show that while there is no "best practice" in CEO behavior, that is a behavior that is optimal for all the firms, there is evidence of significant matching 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 mismatched and that mismatches 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 mismatches is equal to 13% of the labor productivity gap between firms in high and middle-low income countries in our sample.

While this paper has intentionally taken an agnostic approach to defining the relationship between the observed CEO behaviors and more general notions of leadership, an obvious next step would be to explore in more detail the precise mechanisms through which different administrative behaviors affect firm performance, and why different firms need different behaviors. The CEO behavior that according to our CEO-firm assignment model and our data is scarcer in the population of actual CEOs (and hence produces a better average performance), features multi-functional meetings and a focus on higher-level executives rather than production functions. One tentative interpretation is that a CEO that displays this pattern of behavior is closer to the notion of *general manager*, who primarily coordinates at a high level, delegates operational tasks to other executives and spends his time ensuring good communication in the top management team. On the other hand, the other pattern of behavior may be more closely associated with the role of an *operating manager*, who tends to intervene directly in operational aspects, prefers one-on-one meetings with a variety of internal and external constituents, and puts less emphasis on cross functional coordination.

To the best of our knowledge, while the general idea of leader types is present in recent papers in the economic leadership literature, the dichotomy between general and operational managers is much less prominent in this literature - within and outside economics.[38] Future

---

[37]The 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.

[38]A notable exception in this respect is Barnard (1938), who emphasizes the coordinative role of top

work could utilize information about CEO behavior to inform alternative leadership models, and in particular explore in detail the underlying CEO-firm 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 the importance of managerial interactions and team behavior (Hambrick and Mason, 1984), which are also now largely absent from our analysis. We leave these topics for further research.

---

executives in large ad complex organizations.

References

Alder, S. D. (2016). In the Wrong Hands: Complementarities, Resource Allocation, and TFP. *American Economic Journal: Macroeconomics*, 8(1):199–241.

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., Prat, A., and Sadun, R. (2013). Managing the Family Firm: Evidence from CEOs at Work. Harvard Business School Working Papers 14-044, Harvard Business School.

Barnard, C. (1938). *The Functions of the Executive*. Harvard University Press.

Bender, S., Bloom, N., Card, D., Reenen, J. V., 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 Reenen, J. V. (2016a). Management as a Technology? NBER Working Papers 22327, National Bureau of Economic Research, Inc.

Bloom, N., Sadun, R., and Van Reenen, J. (2016b). Management as a Technology? mimeograph.

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.

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.

Cyert, Richard; March, J. G. (1963). *A Behavioral Theory of the Firm.* Wiley-Blackwell.

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.

Halac, M. and Prat, A. (2016). Managerial Attention and Worker Performance. *American Economic Review*, 106(10):3104–32.

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.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly Journal of Economics*, 124(4):1403–1448.

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.

Lippi, F. and Schivardi, F. (2014). Corporate control and executive selection. *Quantitative Economics*, 5:417–456.

Luthans, F. (1988). Successful vs. Effective Real Managers. *Academy of Management Executive*, 2(2):127–132.

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 & Row., New York.

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.

Porta, R. L., Lopez-De-Silanes, F., and Shleifer, A. (1999). Corporate Ownership Around the World. *Journal of Finance*, 54(2):471–517.

Syverson, C. (2011). What Determines Productivity? *Journal of Economic Literature*, 49(2):326–65.

Tervio, M. (2008). The Difference That CEOs Make: An Assignment Model Approach. *American Economic Review*, 98(3):642–68.

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)*.

Appendix Tables and Figues

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.[39] To ensure comparability, we adopted the same protocol and retained the same project manager across both waves. The enumerators where 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

---

[39]The 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.

Figure A.1: Survey Instrument

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.[40]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. 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.

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 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

---

[40]For 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.

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***	0.695***	0.655***	0.559*
	(0.074)	(0.075)	(0.079)	(0.288)
Country=France	0.210***	0.256***	0.143	0.562**
	(0.073)	(0.074)	(0.104)	(0.221)
Country=Germany	0.115	0.194**	0.152*	0.476**
	(0.072)	(0.078)	(0.082)	(0.222)
Country=India	0.658***	0.699**	1.227***	0.672
	(0.247)	(0.272)	(0.371)	(0.425)
Country=UK	-0.178**	-0.139*	-0.153**	0.088
	(0.074)	(0.074)	(0.077)	(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.

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

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

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

Hellinger Distances: Simulated (Blue) and Acutal (Red)
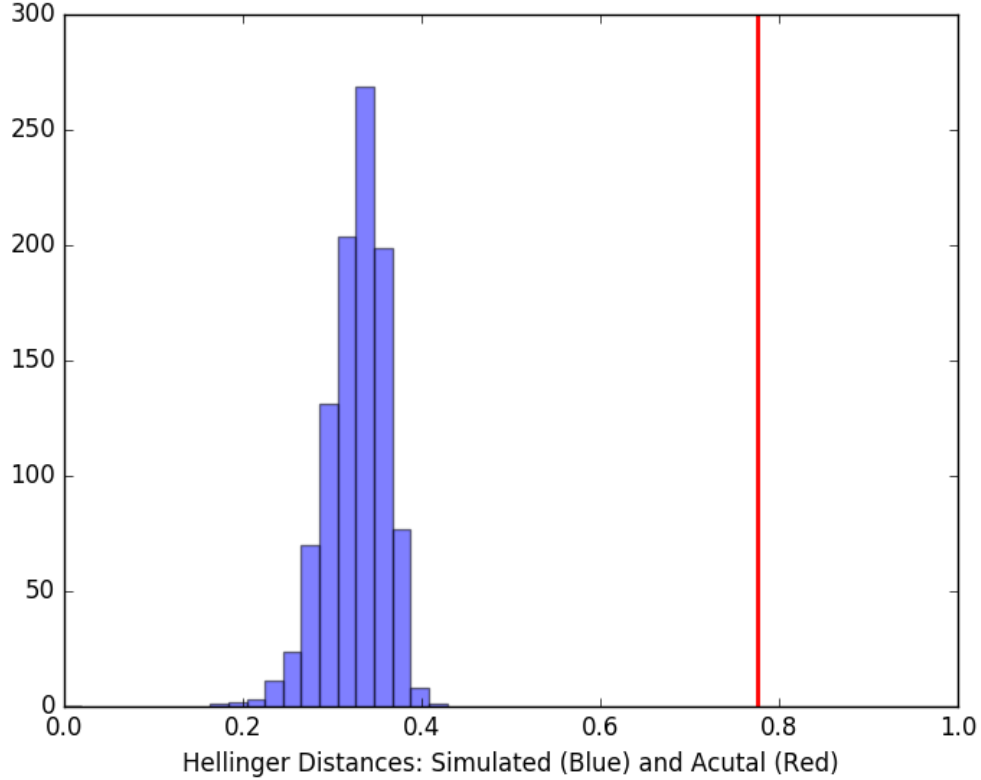
Hellinger distance between the two estimated pure behaviors. We repeat this procedure 1,000 times.

Figure B.1 plots the distribution of the Hellinger distances in the 1,000 simulations. The red line denotes the Hellinger distance we observe 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

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

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 $\bar{w} + B = 1$. 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-0 firm faces an abundant supply of type-0 CEOs. As all the applications it receives come from type-0 CEOs, the firm will simply hire the first applicant. A type-1 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 $s = 1$ and hire him.

We now consider CEOs. Suppose that all type-1 CEOs apply to type-1 firms and type-0 CEOs apply to type-1 firms with probability z and to type-0 firms with probability $1 - z$.

If a type-0 CEO applies to a type-0 firm, he will get a job if and only if his application is downloaded. The mass of type-0 firms is $1 - \phi$. The mass of type-0 CEOs applying to type-0 firms is $(1 - \gamma)(1 - z) m$. The probability the CEO is hired is

$$P_0 = \frac{1 - \phi}{(1 - \gamma)(1 - z) m}.$$

If instead a type-0 CEO applies to a type-1 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-1 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-1 firms is therefore:

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

Given that the mass of applicants to type-1 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-0 applicant passes the screening process is $1 - \rho$. Thus, the probability that a type-0 applicant is hired by a type-1 firm is

$$P_1 = \frac{(1 - \rho) \phi}{m((1 - \gamma) z (1 - \rho) + \gamma)}.$$

In the equilibrium under consideration a type-0 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-1 firm is b, the indifference condition is $P_0 = bP_1$, which yields:

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

50

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 0-types will apply to 1-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-1 CEOs always produce 1, while the average productivity of a type-0 CEO is equal to the probability that he is matched with a type 0 firm, which is

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

By replacing z, we find the average productivity of a type-0 CEO:

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

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 Dynamic Implications of the Assignment Model

We now explore the dynamic implications of our CEO-firm assignment model. Suppose that we know the behavior of the current CEO, but not the type of the firm and the behavior of the previous CEO. What can we say about the evolution of firm performance over time?

Let us assume that the conditions for Proposition 2 are satisfied. There are two types of CEOs ($i \in \{0, 1\}$) and two types of firms ($j \in \{0, 1\}$). We assume that the abundant CEO type is $i = 0$. Using a reduced form expression from the previous section, assume that the performance of a firm is $y_j + x_{ij}$, where $x_{ij} = 1$ if the firm type and the CEO type match ($i = j$) and $x_{ij} = 0$ if there is a mismatch ($i \neq j$)x_{ij}, and the term y_j indicate that the two firm types may have different baseline productivities.

Let us consider a firm whose CEO is replaced at time 0. Let x_{ij}^{old} and x_{ij}^{new} denote the match quality of the previous CEO and the current CEO, respectively. The performance of the firm at time $t < 0$ was determined uniquely by the performance of the old CEO (thus assuming that he had been in the job sufficiently long). The performance at $t \geq 0$ is given by

$$Y_t = y_j + (1 - \alpha_t)x_{ij}^{\text{old}} + \alpha_t x_{ij}^{\text{new}},$$

where α_t is increasing and s-shaped in t. Namely,$\alpha_0 = 0$,$\alpha_t' > 0$, $\lim_{t \to 0+} \alpha_t' = 0$, $\lim_{t \to \infty} \alpha_t =$

1, and $\alpha''_t > 0$ if t is low and $\alpha''_t < 0$ if t is high. As time passes, the company's performance is determined more and more by the type of the new CEO as his tenure increases. The s-shaped assumption captures the idea that the effect of a new CEO is limited in the beginning, it increases with time, but then it reaches a stable plateau.

Consider a large sample of firms. Suppose we observe the type of the current CEO, but we do not observe the type of the previous CEO, nor the type of the firm. What can we say about them?

If the current CEO belongs to the scarce type, we know for sure that the firm has type-1. The previous CEO was the scarce type too with probability π and the abundant-type with probability $1 - \pi$.[41]

Focus on performance growth, taking $t = 0$ as the baseline year: $\Delta Y_t = Y_t - Y_0$. If the current CEO belongs to the scarce type, we have

$$\Delta Y_t \left(i^{\text{new}} = 1 \right) = \begin{cases} 0 & \text{if } t < 0 \\ \left((1 - \alpha_t) E \left[x_{ij}^{\text{old}} | x_{ij}^{\text{new}} = 1 \right] + \alpha_t \right) - E \left[x_{ij}^{\text{old}} | x_{ij}^{\text{new}} = 1 \right] & \text{if } t > 0 \end{cases}$$

but note that $E \left[x_{ij}^{\text{old}} | x_{ij}^{\text{new}} = 1 \right] = \pi < 1$. Therefore,

$$\Delta Y_t \left(i^{\text{new}} = 1 \right) = \begin{cases} 0 & \text{if } t < 0 \\ \alpha_t (1 - \pi) & \text{if } t > 0 \end{cases},$$

which implies that average performance is flat before the CEO replacement and follows $\alpha_t (1 - \pi)$ thereafter.

If instead we consider a sample of firms run by abundant-type CEOs, a specular argument applies: we would observe that the average performance decreases after the current CEO is hired and follows a similarly s-shaped curve. Therefore we have:

Proposition 3 *The average performance of a sample of firms who are currently run by scarce-type (abundant-type) CEOs was flat before the new CEOs were hired and it becomes increasing (decreasing) and s-shaped thereafter.*

Figure D.1depicts the average performance of a set of firms run by scarce-type CEOs, $\Delta Y_t \left(i^{\text{new}} = 1 \right)$, under the assumption that α_t is a sigmoid function ($\alpha_t = t/\sqrt{1 + t^2}$) and $\pi = \frac{1}{2}$. The average effect of having a scarce-type CEO is positive, gradual, and s-shaped. This result implies that if we observe a set of firms run by scarce type CEOs who were all hired at the same date, we should predict that the average performance of those firms is
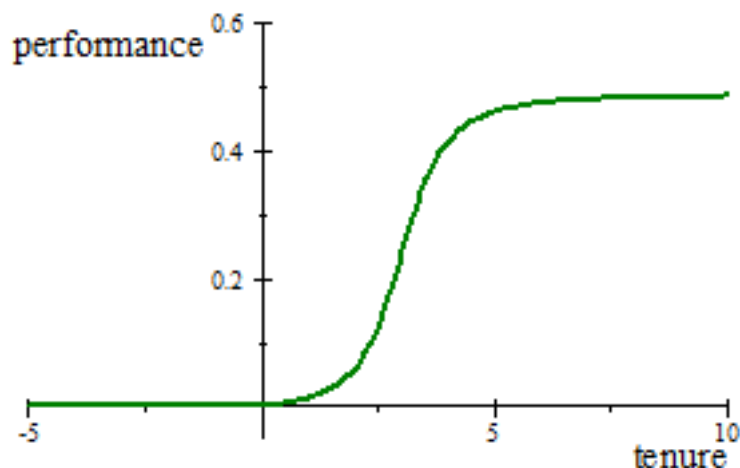
---

[41]This probability is given in equilibrium by

$$\pi = \frac{\gamma}{\gamma + (1 - \gamma) z},$$

where

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

Figure D.1: Average performance of firms run by type-1 CEOs

constant before the CEOs are hired, almost constant right after they are hired, and increasing and s-shaped afterwards.

E   Additional Results

E.1   CEO Behavior Index: Additional Descriptives

E.1.1   Variation across Countries and Industries

Figure E.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.

E.1.2   Correlation with Firm and CEO Characteristics

Panel A, Figure E.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 high-index 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

Figure E.1: CEO Behavior Index: Variation across Countries and SIC 2 industries
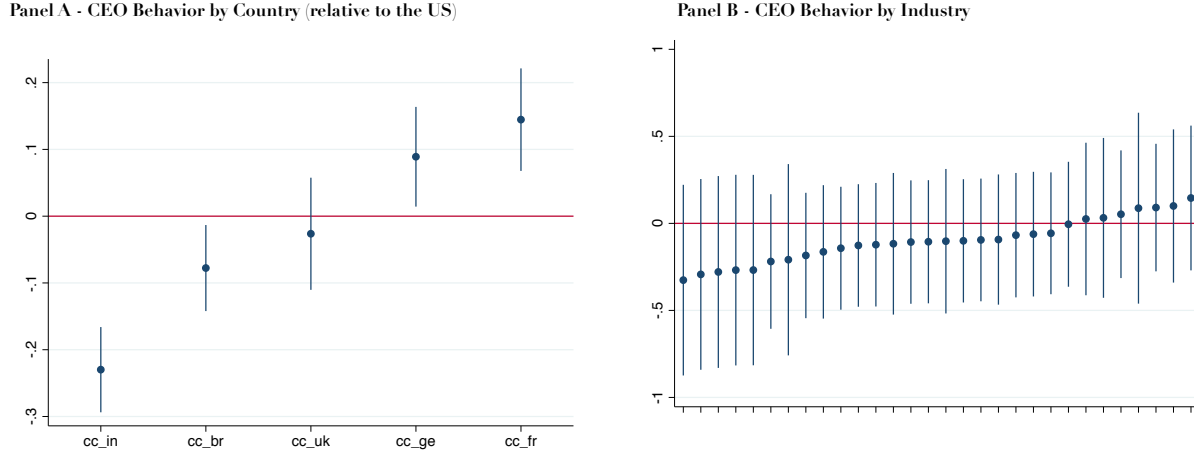
Panel A - CEO Behavior by Country (relative to the US)   **Panel B - CEO Behavior by Industry**

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

no evidence that the index is related to the age of the executive, or to whether the CEO was promoted to the role within the organization.

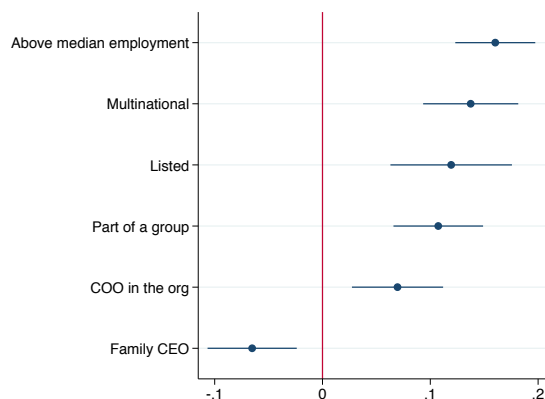## E.2   Production Function: Robustness Checks

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

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

Table E.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

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

Panel A - CEO Behavior Index and Firm Characteristics **Panel B - CEO Behavior Index and CEO Characteristics**
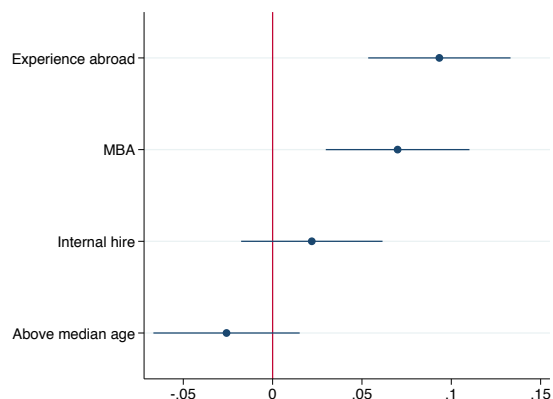
Notes: Each point represents the coefficient obtained when regressing the CEO behavior index on each dummy variable, including country and SIC2 industry dummies as additional controls.

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.

E.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 (E.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.

Notes: Each point represents the coefficient obtained when regressing
the CEO behavior index on a full set of SIC2 industry dummies,
including country dummies as additional controls.

E.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. (2013). 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.[42]

E.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 (E.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. [43] In the columns 8 and 9 of Table E.2, Panel A we show that these alternative ways of classifying CEOs do not fundamentally alter the relationship between behavior and firm performance.

E.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 (E.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

---

[42]Among 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.

[43]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 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 analysis as described above.

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. [44] Again, we find a significant relationship between the behavior index and firm performance, as shown in Table (E.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.

E.2.6 Alternative estimation techniques

Table (E.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).

E.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 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

---

[44]In 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 E.1: Production Function Results Using Shares of Time

	(1)	(2)	(4)	(5)	(3)
Dependent Variable: log(sales)					
log(employment)	0.889***	0.895***	0.893***	0.907***	0.876***
	(0.040)	(0.039)	(0.041)	(0.040)	(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 1, column 1.

Table E.2: Robustness Checks

Panel A

Experiment	(1) Baseline	(2) Firm by year accounting data, cluster at the firm level	(3) No weighting	(4) Control for hours worked	(5) Control for firm observables	(6) Control for CEO observables	(7) Discretized version (>=.5)	(8) Principal Component	(9) K-means
Dependent Variable: Log sales									
CEO behavior index	0.343***	0.345***	0.298***	0.338***	0.265**	0.309***	0.234***	0.119***	0.259***
	(0.108)	(0.090)	(0.104)	(0.105)	(0.114)	(0.108)	(0.064)	(0.035)	(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	2554	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) Baseline	(2) Exclude alla activities not present in at least 15 CEOs	(3) Exclude alla activities not present in at least 45 CEOs	(4) Include Travel and Email	(5) Mixture model	(6) Index computed on sales sample only	(7) Index computed by high and low income countries separately	(8) Olley Pakes productivity residual
Dependent Variable: Log sales								
CEO behavior index	0.343***	0.305***	0.316***	0.265***	0.133**	0.347***	0.292***	0.470***
	(0.108)	(0.110)	(0.102)	(0.096)	(0.064)	(0.102)	(0.081)	(0.105)
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 1, 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 E.3: Perplexity for Different Numbers of Pure Behaviors
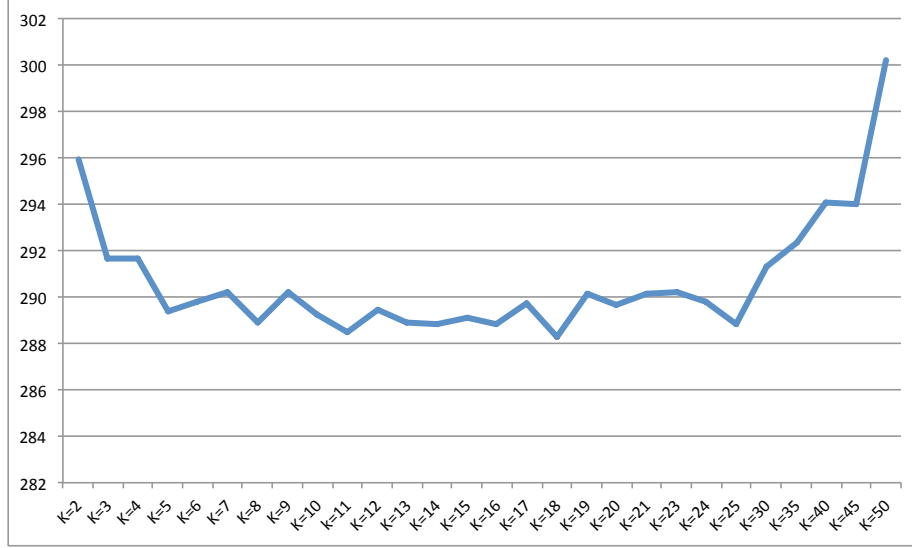
Notes: This graph plots the average perplexity computed on test data from ten randomly drawn sets of training data. The split between training data and test data is two-thirds / one-third. Lower values of perplexity indicate better goodness-of-fit. There are gaps in the values for K due to save on computation time.

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 E.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 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 1. The point estimates and confidence intervals are displayed in figure E.4.

An initial result is that an F-test for the joint significance of the variables is highly signifi-

Figure E.4: Coefficient Estimates for Model with Eleven Pure Behaviors
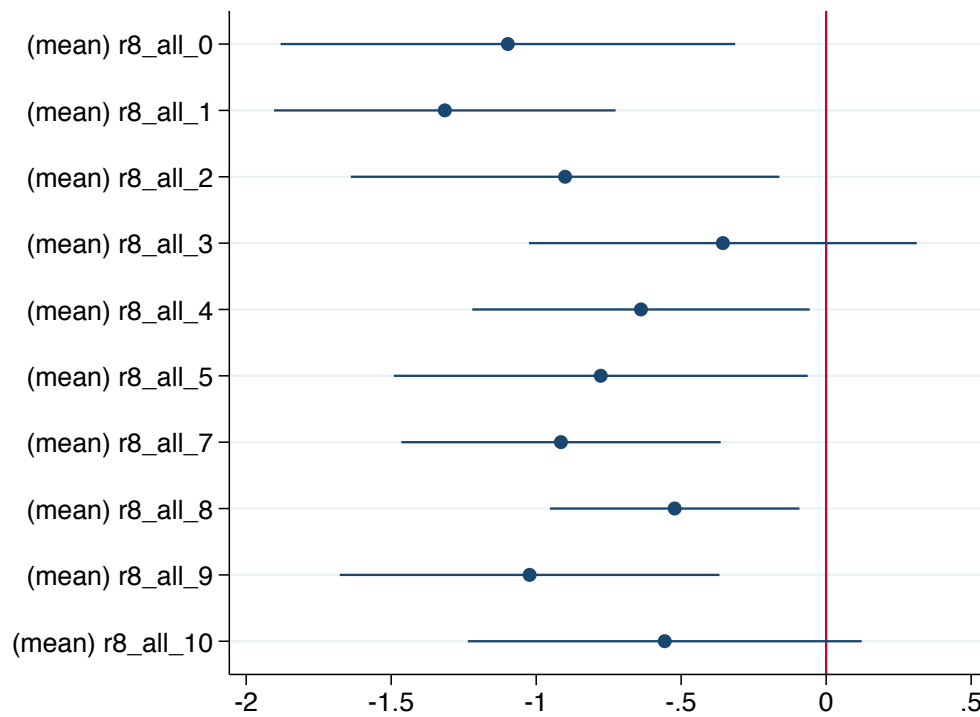
Notes: This graph plots the coefficient estimates for the productivity regression in column (1) of table 1 when we replace the scalar behavioral index with the output of the model with $K = 11$. The omitted category is the probability put on the sixth pure behavior. The dots represent point estimates, and the lines 95% confidence intervals.

Figure E.5: Hellinger Distance from Pure Behavior 0 vs. Coefficient in Productivity Regression
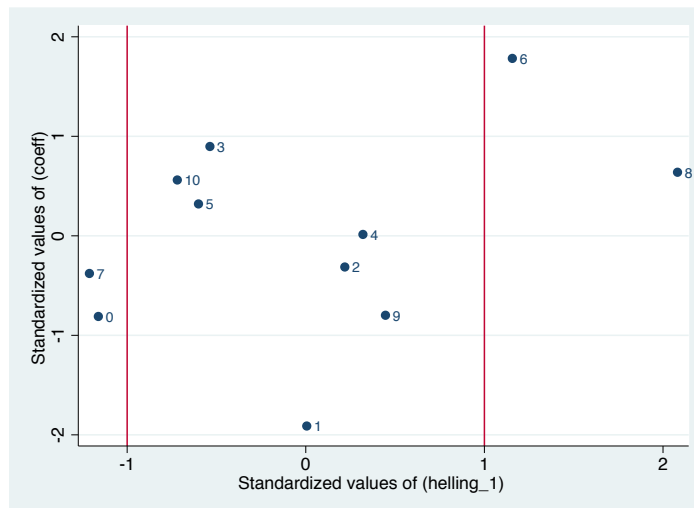
cant, which implies that the underlying heterogeneity in the probability of choosing different pure behaviors among CEOs is important for explaining differences in firm performance. In terms of individual coefficients, eight are significantly negative relative to the sixth pure behavior, while two are 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 E.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 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.

E.3   CEO Behavior Index and Management Practices

E.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).[45] This uses an interview-based evaluation tool that defines 18 basic management

---

[45]More details can be found at http://worldmanagementsurvey.org/

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 on "open" questions. For example, on the first monitoring question we start by asking the open question, "tell me how your 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?".[46] 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 we randomly sample 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. More details on th methodology can be found in Bloom et al. (2016b).

E.3.2 Management and CEO Behavior

We look at the cross sectional correlation between the management data and the CEO behavior index in Table E.3. Columns 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.

---

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

Table E.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**		0.550*
				(0.277)		(0.280)
Management (z-score)	0.054*				0.207**	0.187**
	(0.030)				(0.082)	(0.075)
Operations, Monitoring, Targets (z-score)		0.057*				
		(0.029)				
People (zscore)			0.043			
			(0.034)			
log(employment)	0.106***	0.109***	0.104***	0.848***	0.880***	0.843***
	(0.033)	(0.033)	(0.034)	(0.093)	(0.067)	(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.