SOCIAL NETWORKS IN THE BOARDROOM

Francis Kramarz
CREST (ENSAE)

David Thesmar
HEC Paris

Abstract
This paper provides evidence that social networks strongly affect board composition and are detrimental to corporate governance. Our empirical investigation relies on a large data set of executives and outside directors of French public firms. This data source is a matched employer–employee data set that provides detailed information on directors/CEOs as well as information about the firm employing them. We find a strong and robust correlation between the CEO's network and that of his directors. Networks of former high-ranking civil servants are the most active in shaping board composition. Our identification strategy takes into account not only firm and directors' fixed effects but also the matching of firms and director in terms of one observable and one unobservable characteristic. Turning to the direct effects of such network activity, we find that firms in which these networks are most active pay their CEOs more, are less likely to replace a CEO who underperforms, and engage in less value-creating acquisitions. These findings suggest that social networks are active in the boardroom and have detrimental effects on firms' governance. (JEL: G34, P16, L14)

1. Introduction

That social networks affect market outcomes is a well-documented fact (for early references, see Rees 1966; Granovetter 1973). This paper investigates their impact on corporate performance. Toward this end we focus on the market for nonexecutive directors, a market in which networks are important. There are two opposing views about how these networks affect corporate performance. On the one hand, current directors advise management on the difficult task of finding new directors; by

The editor in charge of this paper was Stefano DellaVigna.

Acknowledgments: For their decisive help in collecting the data we are most grateful to Denise François, Claude Jacquin, and their colleagues at INSEE. We also thank Nicolas Depetris-Chauvin and Elsa Kramarz for their research assistance. Many of the improvements since the first draft are due to the valuable comments of the editor and the referees of this journal, Roland Bénabou, Mariasunta Gianetti, Michel Gollac, Augustin Landier, Laurent Lesnard, and Ernst Maug, as well as of seminar participants from the COST conference at Uppsala, the NBER summer institute, Amsterdam University, CREST, Ente Einaudi, and the Stockholm School of Economics. Kramarz is a Research Fellow at CEPR and IZA. Thesmar is a Research Fellow at CEPR.

E-mail: kramarz@ensae.fr (Kramarz); thesmar@hec.fr (Thesmar)

Journal of the European Economic Association August 2013 11(4):780–807
 DOI: 10.1111/jeea.12021

channeling information about candidates to management, networks improve director–management match and thereby improve corporate performance (Saloner 1985). On the other hand, the supervisory role of directors means that using social networks may come at a cost: relying on executives' networks to hire their own supervisors will be detrimental to directors' independence, so supervision will be ineffective. In this view, firms will be less well managed. Overall, the economic effect of social networks in the market for directors is a priori ambiguous and can be settled only by empirical investigation.

Here we examine this question in the case of France. This paper provides evidence that CEOs' social networks strongly affect board composition and that social networks in the boardroom reduce its efficiency: (i) firms in which social networks are active are less likely to replace CEOs who underperform; (ii) socially connected CEOs tend to have higher compensation, in particular with respect to stock options; and (iii) connected CEOs make lower quality deals (as measured by changes in the firm's stock price after an acquisition announcement).

In order to examine social networks in the boardroom, we use a unique data set on CEOs and nonexecutive directors of all corporations listed on the Paris stock exchange from 1992 to 2003. France is particularly well suited for this study because its elites are highly concentrated and because (at least some of) their networks are well known, easily identified, and easy to measure. Indeed, the sociological literature documents the coexistence among French business elites of two broad and distinct networks: engineers and former high-ranking civil servants.[1] Members of these two networks are mostly recruited within graduates of two elite institutions: Ecole Polytechnique (for engineers) and Ecole Nationale d'Administration (for administrators). Firms run by CEOs from these two networks account for 12% of all firms traded on the Paris Stock Exchange (65% in asset-weighted terms). Yet it's not just that the alumni of these two schools are overrepresented among top executives; more importantly, entering one of them constitutes virtually the only way to obtain a high-level job in the civil service. Given these specific institutional features, data on social networks are relatively easy to collect via alumni directories and the French edition of *Who's Who*. In this way we gather data on the socioeconomic background, education, and career of directors and CEOs; we then match these data with accounting and financial information on the employing firms.

Our empirical investigation involves two steps. First, we provide evidence that social networks distort the labor market for nonexecutive directors. To do this we estimate, for each individual in our sample, the probability of being hired by a given firm. The key regressor in this model is the interaction between the candidate's network and the network of the firm's CEO: if both are the same, the probability of being hired should increase. This is our test of the prevalence of networks. Because we exploit the full variability and identification power provided by our matched employer–employee data, we are able to account for two important dimensions of unobserved heterogeneity

1. For references in English, see Swartz 1986; Kadushin 1995; Frank and Yasumoto (1998). References in French include Bauer and Bertin-Mourot 1997; Suleiman 1997.

that could bias our estimates of network effects. The first dimension is the inherent ability of each individual to become a director in general and to be appointed in firms that have particular observable characteristics. For instance, top-level bureaucrats may simply be more intelligent than others and thus more likely to run or supervise a large firm; hence their intelligence alone might account for their presence in a firm as CEO or directors. Our methodology allows one to account for this possibility. The second dimension is the (unobservable) firm-level propensity to hire directors and CEOs with particular observable characteristics. For instance, firms with an authoritarian corporate culture may prefer to hire older directors and CEOs, individuals from a generation in which (say) civil servants may be overrepresented. Also firms that anticipate difficulties may be more inclined to hire politically connected CEOs and directors. We give a formal proof that the data deliver enough variability to identify network effects, even in the cross section, while taking these two dimensions of unobserved heterogeneity into account.

We follow the sociological literature and define three main networks: (1) former civil servants who graduated from Ecole Nationale d'Administration, (2) former civil servants who graduated from Ecole Polytechnique, and (3) Polytechnique graduates with no history of civil service. We take all other CEOs (possibly belonging to other networks, or to none) as the reference group. We find that the probability of being hired in a given firm is greater when the individual and the CEO belong to the same network, *but only when this network is related to a past career in civil service*. We then look at hiring equations (flows), rather than employment (stock) equations, which allows us to discriminate between two effects—that of the CEO's network and that of past board composition, on each individual's probability of employment. Doing so confirms our results: civil service–related networks of CEOs still affect the recruitment policies of directors. However, the composition of the board has no significant impact on the identity of newly recruited directors. We interpret this finding as tentative evidence that it is the CEO, not the existing directors, who "shapes the board".

The second step in our analysis looks at governance in firms run by former high-ranking bureaucrats. In all these tests, we compare firms whose CEOs are former civil servants to firms whose CEOs have had full private-sector careers. This approach is based on the tendency of CEOs in civil service–related networks to hire directors from the same background. Other CEOs (in particular, those with an engineering background in the private sector) do not appear to hire from their own networks. We then look at three measures of corporate governance and ask whether firms run by network-connected CEOs (and hence with a connected board) tend to score lower on these measures. Our first measure is the sensitivity of CEO turnover to bad performance. Such sensitivity has been found to be greater in better-governed firms (for a survey, see Bebchuk and Weisbach 2010). We show that firms run by connected CEOs are less likely to replace them following bad performance.

As our second measure we examine CEO pay, which the literature has found to be higher in badly governed firms. Disclosure of management compensation was not mandatory in France before 2003, so our data are limited to a single cross-section at the end of our sample period. After controlling for size and industry, we find that

connected CEOs receive about 50% more compensation than do nonconnected CEOs. This difference is attributable in large part to the stock options that former civil servants are more likely to receive.

Third we measure the quality of acquisitions in terms of the stock price reaction at announcement. We find that acquisitions made by connected CEOs create less value for the firm. If the bidder is nonconnected, the firm's stock price typically increases by 1.7% upon announcement; thus the market expects the deal to create 1.7% additional shareholder value. This positive market reaction is in line with existing literature (Bradley and Sundaram 2006). For connected bidders, the stock price does not react at all to the announcement. The difference between these two reactions is both large and statistically significant.

The paper is organized as follows. Section 2 discusses the recent literature on the impact of social networks on corporate governance and describes our own contribution. Section 3 looks at the French elite from a historical and sociological perspective, which allows us to present how we gathered information on networks of outside directors and executives. Section 4 describes the data set and thus provides further descriptive information. Section 5 presents the statistical model and discusses identification. Section 6 examines the extent of networks, and Section 7 addresses their economic costs. Section 8 concludes.

2. Related Literature

We focus first on recent work that shows how social networks affect board composition and the hiring of CEOs. We then review papers that explain how social networks affect corporate governance and firms' outcomes. Finally, we describe the contributions of our paper with respect to these two areas of research.

A number of papers document the existence of networks. Some of them look at CEOs and directors in general. Barnea and Guedj (2008) use data on all directors and CEOs of firms in the S&P 1500 composite index between 1996 and 2004. These authors find that connected directors are more likely to obtain new directorships in the future. Liu (2008) also focuses on US directors but incorporates much more detailed information on their employment history. She finds evidence that connected CEOs are more likely to move to better jobs, in particular to firms that have a related director. Because data are readily available, another strand of the literature focuses on the mutual fund industry; Kuhnen (2009) explores the connections between US mutual funds and their subadvisors.

Within this line of research, our main contribution lies in our measure of social networks. Most of the literature just cited leverages available director data to identify personal connections more accurately—for instance, by assuming that two individuals sitting on the board "know" each other. However, we follow the work of those authors (Hwang and Kim 2009; Nguyen 2009; Braggion 2011) who instead use results from the sociological literature to identify directly the contours of social networks: we will assume, for instance, that two former civil servants are likely to know each other.

Our work is thus related also to the earlier empirical literature on economic outcomes of social networks (see among others, Bertrand, Luttmer, and Mullainathan 2000; Munshi 2003; Bayer, Ross, and Topa 2005). This literature generally relies on indirect identifying assumptions. Our network identification is more precise and direct because we are able to observe both the referee and the applicant. Being able to observe networks within the firm allows us to devise a more refined statistical model whose identification and estimation are detailed in this paper.[2] Our econometric model is a standard matched employer–employee one; in this respect it differs from techniques imported from graph theory which have become popular in the finance literature. An important advantage of our approach is that the underlying identifying assumptions are relatively transparent, and this allows us to control for much of the unobserved heterogeneity characteristic of situations—like the one we study—in which there is no clear instrument.

The second research area features papers that seek to assess the welfare impact of networks. In finance, there is evidence that social networks can be beneficial because they are the channel through which information flows. Hochberg, Ljunqvist, and Lu (2007) find that venture capital funds with parent firms that enjoy stronger network relations (measured using graph theory, as in Barnea and Guedj 2008) have better performance. Cohen, Frazzini, and Malloy (2008) find that mutual fund managers trade more on the stocks of firms whose directors they are connected with and that these trades are profitable. Thus, social networks contribute to making markets efficient: because of their trades, information is progressively impounded into market prices. In the corporate governance literature, existing work finds that social networks actually hurt corporate governance and performance. Barnea and Guedj show that connected firms pay their CEOs more, and Liu (2008) likewise finds that better-connected CEOs receive higher compensation (see also Larcker et al. 2005). Focusing on the end of the nineteenth century, Braggion (2011) finds that firms connected with Freemasonry are more leveraged and slightly less profitable. Hwang and Kim (2009) examine social ties created by a common regional origin, alma mater, military service, or industry. They show that firms with socially independent boards award lower compensation levels while exhibiting stronger pay–performance, and turnover–performance sensitivity. In a paper much like ours but written independently, Nguyen (2009) looks at the same French business elite networks. His data differ slightly: our sample period is longer and covers nearly 600 publicly listed firms per year (versus the top 250 in his paper). As we do, Nguyen finds that CEO turnover is less sensitive to bad performance when the CEO belongs to elite civil–service related networks. He also finds that connections help applicants find better jobs (in larger firms). Finally, and more specifically to the French context, Nguyen demonstrates that connected CEOs tend to lose their jobs after political events, such as the arrival of a new government. Key differences with our work are that he looks at neither CEO compensation nor acquisitions and does not develop a framework to identify and test for the presence of network effects.

2. See Kramarz and Skans (2010) for an extensive use of this framework in the context of family networks, where firms and classrooms are the two dimensions of heterogeneity.

Overall, our paper provides a relatively broad assessment of the negative effects of social networks on corporate governance. We follow the literature in examining the sensitivity of CEO turnover to firm performance, and the impact of social networks that we find is in line with the results reported in that literature. However, our paper is the first to incorporate French evidence on compensation, and our results are consistent with what US studies have found (Barnea and Guedj 2008; Liu 2008). In addition, it is (so far) the only paper to provide evidence that the social connections of firms' directors have a negative effect on the quality of acquisitions.

3. The French Business Elite

For historical and sociological reasons, France's economic elites have two distinctive features (Swartz 1986; Bauer and Bertin-Mourot 1997; see also Bourdieu 1989). First, they tend to be drawn from a handful of Grandes Ecoles, which form separated networks. Second, a large part of the contemporaneous French business elite comes from the civil service and have relatively homogeneous and standardized careers. These two features are easy to observe and will guide our empirical strategy (a fuller description is given in the working paper version's appendix).

The "tyranny of diploma" is a distinguishing feature of the French business elite (Bauer and Bertin-Mourot 1997); in other words college degrees obtained before age 25 tend to overdetermine career prospects. The French postsecondary educational system is split into two parts (Suleiman 1997). The first is the usual university system, which is free and to which access after high school graduation is guaranteed by law (hence there is no selection; in the mid-1990s, this system comprised some 1.2 million students). The second part consists of many small and elite schools attended by about 50,000 students altogether. Within this latter group, the two most prestigious schools produce a large fraction of the business elite (Swartz 1986): the Ecole Nationale d'Administration (ENA) and the Ecole Polytechnique. The former was created after World War II to supply the civil service with highly trained professionals. Ecole Polytechnique was originally founded by Napoleon to recruit and train military officers during the French Revolution; it gradually evolved into an engineering school. Nowadays, most students in each cohort choose the private sector while the best students generally opt for the civil service.

The second distinctive feature of the French business elite is the prevalence of former civil servants. The tight relation between business and the administrative world began in the aftermath of WWII, a reconstruction period largely supervised by the government. Since 1945, the best students in a given class at ENA or Polytechnique have systematically joined prestigious bureaucratic careers paths, called "Grands Corps d'Etat" (Kadushin 1995; Suleiman 1997). These career paths were fast tracks to top-level positions, as manufacturing experts (for Polytechnique), high-level administrators (ENA), and top level managers in firms (both schools). Such careers typically involved first, a few years as a civil servant, then, some time as a cabinet advisor, and finally, access to the top management of a large private or state-owned company.

4. The Data

With the foregoing two features of the French elite in mind, we describe our data sources as follows.

4.1. Data Sources

Our data set matches information on the employees—the CEO and the directors—with data on the employing firms. To construct it, we use two main data sources: (1) the DAFSA yearbook of French listed firms, which provides us with firm-level variables (including the names of the CEO and members of the board); and (2) the ENA and Polytechnique alumni directories, which yield the education history and partial information on careers of graduates. We supplement this information with the French edition of *Who's Who*, which is not exhaustive (it does not include all directors and CEOs) but does allow us to extend our coverage beyond ENA and Polytechnique graduates.

The DAFSA yearbook is an annual publication that compiles the consolidated accounts of listed companies. Yearbooks go back to the 1950s, but detailed income statements are only available from 1984 onward. We extract this information from the 1988–1993 printed issues of the yearbook and from its 1994–2003 electronic issues. We restrict ourselves to firms listed on the two main segments of the stock exchange ("premier marché" and "second marché"). Both segments have, on average about 300 firms listed each year; the first segment lists firms that are larger and whose stock is more liquid.

Along with accounting information, the DAFSA yearbook provides us with the names of the CEO, the chairman and the directors. Henceforth, we shall use the terms "nonexecutive directors" and "directors" interchangeably, since their meanings are identical in the French context. Many CEOs are also chairmen.

We retrieve personal information on the CEOs and the directors primarily from the ENA and Polytechnique alumni directories. These directories provide standard information about education but none about socioeconomic backgrounds and little about bureaucratic careers (if any). We match alumni directories with DAFSA data on CEOs and directors of public firms, using both first and last names. Because these directories are exhaustive, we capture nearly all of ENA and Polytechnique graduates in the sample of CEOs and directors (except for some individuals with very common names and surnames).

To identify other former civil servants and political advisors (i.e, other than Polytechnique and ENA graduates), we supplement this information with the 1994 and 2000 issues of *Who's Who*—a list of prominent people in politics, business, and entertainment. For each individual, the available information is well standardized and includes self-reported measures of parent's occupation, place and date of birth, marital status, number of children, education, current occupation, and past career. We use these data to construct a "former civil servant" dummy (see Section 7). The *Who's*

Who information on each individual listed in the DAFSA database as a director or CEO between 1992 and 2003 is hand-coded using first and last names. On average, some 51% of all CEOs of listed corporations are found in *Who's Who*. According to information in the 1994 and 2000 issues of *Who's Who*, this percentage declines steadily over the period under study from 60% in 1992 to 45% in 2003. The percentages are somewhat lower for directors, who are less likely to be listed in *Who's Who*. On average approximately 36% of directors are listed, and this percentage likewise declines (from 40% to 27%) over our study period.

Relying on the historical and sociological evidence just reviewed, we identify three networks in our sample: (1) ENA graduates, practically all of whom had an early career as civil servants; (2) Polytechnique graduates who started their careers as "civil service" engineers; and (3) Polytechnique graduates who started in the private sector. We now turn to a descriptive investigation of our data to establish the prevalence of these three networks among the directors and CEOs of large listed corporations.

4.2. The French Business Elite in the 1990s

A rough overview of our data confirms and updates the findings of sociologists on a much larger sample. First, the French business elite consist almost entirely of Polytechnique and ENA graduates and civil servants. Second, this pattern has become even more pronounced over the recent period for which we have data (1992–2003).

Indeed, the data are fully consistent with the sociological and historical evidence outlined previously. Over the 1992–2003 period, ENA and Polytechnique graduates run the lion's share of French firms; many of these graduates—in particular those who are actively involved in politics—are also former civil servants. As shown in Table1, ENA and Polytechnique graduates run, on average, about 20% of the firms; this number is deceptive, however, in that their firms are typically quite large and account for some 70% of all assets traded on the stock exchange (at book value). This pattern persists even if we restrict our focus to civil servants who were political advisors: they control 6% of the firms but 52% of the assets.

Moreover, despite vigorous privatization initiatives accompanied by the deregulation of many economic sectors during the 1990s, civil servants remain prevalent among top executives of French corporations well into the twenty-first century. Figure 1 plots changes in the asset-weighted share of CEOs from various backgrounds. During the 1990s, civil servants with purely administrative background (ENA graduates) became increasingly prevalent. Also, the fraction of Polytechnique "engineers"—either from the civil service or from the private sector—declined sharply after 1999. Observe that both of these trends began with the resumption of privatizations under the right-wing government elected in 1993,[3] whereby state enterprises run by former civil servants were sold to the stock market through IPOs.

3. There was a first wave of privatizations in 1986–1988, which stopped with the re-election of President Mitterand.

TABLE 1. Firm-level summary statistics.

	Mean	S.D.	Asset-weighted mean	Mean (CEO from CS)	Mean (CEO from PS)
CEO background					
ENA graduate	0.07	0.26	0.54	0.54	0.01
Polytechnique, former civil servant	0.04	0.20	0.08	0.23	0.01
Polytechnique, always private sector	0.08	0.27	0.33	0.15	0.07
Former civil servant	0.12	0.32	0.65	1.00	1.00
Outside directors					
Total number	6.91	3.82	-	10.03	6.53
At least one ENA	0.3	0.46	0.90	0.64	0.25
At least one polytechnique, CS	0.18	0.38	0.59	0.43	0.14
At least one polytechnique, PS	0.36	0.48	0.81	0.60	0.33
Firm characteristics					
Assets (billions of euros)	5.5	45.7	-	4.4	3.3
ROA				0.05	0.06
Former SOE	0.13	0.34	0.64	0.36	0.10

Note: French public firms over the period 1992-2001. Observations: 8,014 for CEOs; 5,948 for asset-weighted and firm-level statistics; 55,409 for outside directors. Key: S.D. = standard deviation; CS = civil servant; PS = private sector; SOE = state-owned enterprise. Sources: DAFSA yearbook (director names); *Who's Who*; alumni directories.

The trend in board composition reflects changes in the (asset-weighted) share of directorships held by ENA graduates, Polytechnique graduates with a career in civil service, and Polytechnique graduates with a purely private-sector background (see also Table A.1 in the Online Appendix). These shares are high and also show a strong upward trend in the early 1990s, when privatization resumed (1993). In asset-weighted terms, between 40% and 50% of all director seats were filled by members of one of these three networks.

At the firm level, the CEO's identity seems to matter for shaping board composition. As Table 2 shows, the fraction of ENA graduates seated on the board of corporations run by ENA graduates is much higher than in other corporations. The same result holds for Polytechnique graduates with a civil service background but not for those "polytechniciens" whose entire careers are in the private sector.

This first look at the data does seem to indicate that social networks shape the composition of corporate boards. It remains unclear, however, which structural parameter is identified by this simple inspection of Table 2. Have we simply found that ENA graduates are better (and hence more sought-after) directors? Do the data show that some firms naturally attract ENA graduates as directors and CEOs—perhaps because they operate in regulated industries or because the business requires a working knowledge of the bureaucracy? Or do they suggest that ENA CEOs run larger firms with larger boards and are thus more likely to appoint directors in general, and thus in particular from ENA? To address these questions, in Section 5, we derive an empirical model from first principles. That model will enable us to interpret the descriptive results of Table 2.

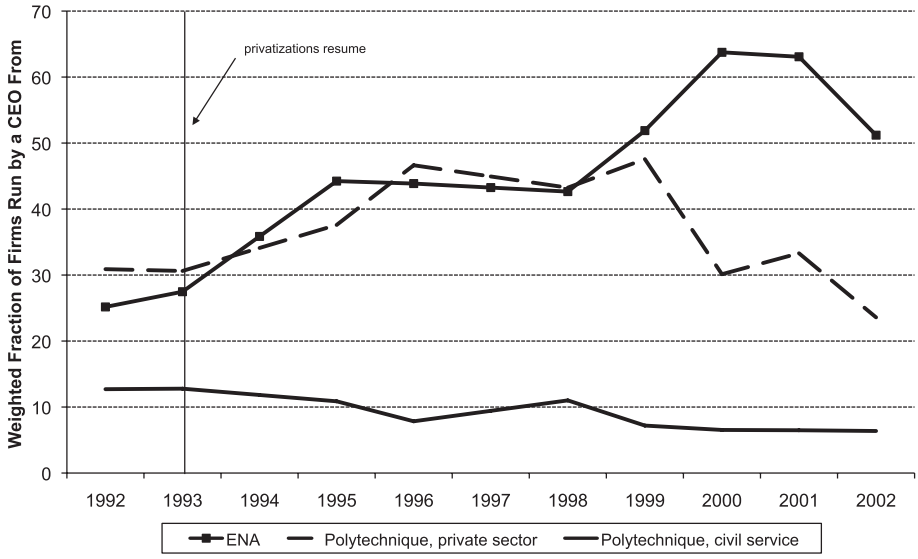

FIGURE 1. Characteristics (asset-weighted) of French CEOs.

TABLE 2. Preliminary evidence on networks: board composition as a function of CEO background.

		CEO education/career			
	All	ENA	Polytechnique civil service	Polytechnique private sector	Other
Non-weighted averages					
% of ENA graduates	0.06	0.16	0.13	0.08	0.05
% of Poly. graduates, civil servants	0.03	0.06	0.12	0.04	0.02
% of Poly. graduates, private sector	0.07	0.09	0.12	0.12	0.06
% of other	0.84	0.69	0.63	0.76	0.87
Asset-weighted averages					
% of ENA graduates	0.25	0.31	0.23	0.22	0.11
% of Poly. graduates, civil servants	0.07	0.08	0.13	0.07	0.02
% of Poly. graduates, private sector	0.12	0.14	0.13	0.10	0.09
% of other	0.56	0.47	0.51	0.61	0.77

Note: French public firms over the period 1992-2001. Observations: 8,014 for nonweighted statistics; 5,948 for asset-weighted statistics. Sources: DAFSA yearbook (director names); *Who's Who*; alumni directories.

5. Empirical Strategy

This section first proposes a model in which the impact of networks is clearly identifiable. Because the model is defined for each individual and each firm in the sample, its estimation is computationally intensive. Therefore, in a second step we describe aggregation techniques that simplify estimation and then discuss their identifying power.

5.1. The Networks Model

Consider the (matched employer–employee) panel in which individuals are indexed by i, firms by j, and time by t. We assume the existence of several (possibly overlapping) networks, which we index by k. As in Munshi (2003), we seek to establish whether belonging to the same network as the firm's CEO increases individual i's odds of being on firm j 's board:

$$E_{ijt} = \alpha_i \cdot Z_{jt} + \beta_j \cdot X_{it} + Z'_{jt} \cdot M \cdot X_{it} + \sum_{k,l} \lambda_{kl} \cdot \left(C^k_{jt} \cdot A^l_i\right) + \varepsilon_{ijt}, \tag{1}$$

where $E_{ijt} = 1$ if individual i works as a director of firm j at date t (and is equal to 0 otherwise). We use k as an index for the network; thus, $A^k_i = 1$ when individual i belongs to network k (and 0 otherwise), $C^k_{jt} = 1$ when the CEO of firm j at t belongs to network k (and 0 otherwise). The term Z_{jt} is a vector of firm-level observables and X_{it} is a vector of individual-level observables. We use α_i (resp. β_j) to denote is a vector of coefficients that differ across individuals (resp. firms).[4] Finally, M is a matrix of coefficients for the various interaction terms between variables of X_{it} and variables of Z_{jt}.

In equation (1), we measure the strength of social networks by looking at the λ_{kl} coefficients. If network effects are present then being appointed as a director in firm j should occur more frequently when the individual and the CEO share the same network. Our hypothesis can thus be written as follows:

$$\lambda_{kk} > \lambda_{kl} \quad \textit{for all } l \neq k; \qquad (\mathrm{H}_0)$$

this corresponds to evidence of network effects in the patterns of nomination.

Of course, finding directors and CEOs from the same network in the same company need not constitute evidence of networks. For instance, former civil servants tend to join (both as CEOs and directors) larger firms, firms that operate in regulated industries, or firms that depend on procurement contracts. Under an alternative interpretation, former civil servants have higher ability and large firms prefer to hire people with higher abilities. This is why equation (1) incorporate three types of controls. First, the term $\alpha_i \cdot Z_{jt}$ stands for the unobserved propensity of people α_i to serve as directors of companies with observables; Z_{jt} for example, high-IQ workers may be more likely than others to obtain seats on the boards of large firms. Second, $\beta_j \cdot X_{it}$ measures the unobserved firm propensity β_j to hire directors with observables X_{it}—as when firms with an authoritarian corporate culture prefer to hire older directors. Taken together, these two terms control for the sorting of directors and firms along one dimension that is observable and along another that is not.

The third control, $Z'_{jt} \cdot M \cdot X_{it}$, is for matching of directors and firms along purely observable dimensions. For instance, former civil servants may tend to join

4. Because intercepts are always present in the vectors X_{it} and Z_{jt}, model (1) always includes "pure" person effects and "pure" firm effects.

the boards of former state-owned enterprises, engineers may prefer employment in more technology-intensive industries, or directors with more education may be more often found in larger firms. The elements of the M matrix control for the strength of sorting along observables in the data.

Model (1) cannot be estimated as such. Indeed, the original data, by construction, includes only observations for which $E_{ijt} = 1$. However, it is virtually impossible to generate all observations for which $E_{ijt} = 0$. Because there are, a priori, nearly 600 firms and 5,000 directors every year over a ten-year period, the sample of all (i, j, t) would amount to some 30 million observations. Hence, in Section 5.2 we derive estimable models that *require only* the knowledge of the "$E_{ijt} = 1$" observations.

5.2. The Firm-Level Model

Here we show how model (1), expressed as a match between an individual and a firm, may be aggregated as a firm-level model in which the parameters of (1) can be identified. Let us introduce a few more notations. First, let

$$n_{jt}^k = \sum_i E_{ijt} \cdot A_i^k$$

be the total number of directors sitting on firm j's board who belong to network k, where $n_{jt} > n_{jt}^k$ is the total number of directors of j, n_t^k is the total number of members in network k, and n is the total labor force.

In the following derivation, we assume for simplicity that $X_{it} = 1$—in other words, that directors do not differ in terms of observable characteristics. This is an admittedly strong assumption, but we can dispense with it in the individual-level model (see the Online Appendix). The purpose of this assumption is mostly to clarify matters (however, detailed calculations that forgo the simplifying assumption are reported in the Online Appendix). After a few manipulations, which amount to computing n_{jt}^k and n_{jt} using model (1), we show in the Online Appendix that

$$Y_{jt}^k = \left(\frac{n_{jt}^k}{n_t^k} - \frac{n_{jt}}{n_t} \right) = a_t^k \dot{Z}_{jt} + \sum_m b_t^{mk} \cdot C_{jt}^m + u_{jt}^k \tag{2}$$

with

$$b_t^{mk} = \lambda_{mk} - \sum_l \lambda_{ml} \frac{n_t^l}{n_t},$$

where Y_{jt}^k is the proportion of members of network k selected to j's board in excess of the proportion of all candidates for that board. The $a^k \cdot Z_{jt}$ term in equation (2) allows us to control for firm–director matching along firm observables and director unobserved characteristics. This control is performed by simply including the Z_{jt} firm-level controls in the linear regression of Y_{jt}^k on the CEO's network C_{jt}^m. The b_t^{mk} coefficient measures the relation between a CEO's identity and the board's composition while controlling for the fixed effects mentioned previously. These coefficients are not

exactly equal to the λ values, because any network can be present on a given firm's board simply as the result of its size in the overall population. The expected fraction of m, even in the absence of network effects, would be n^m/n. Therefore, the specific effect on k will be *underestimated* in the firm-level specification if we do not correct for this bias.

Finally, testing for the presence of networks is fairly straightforward. By comparing b_t^{kk} and b_t^{kl}, we can restate hypothesis H_0 in terms of the estimated parameters from (2),

$$b_t^{kk} > b_t^{kl} \quad \textit{for all } l \neq k. \qquad (\mathrm{H}_0')$$

That is, we look at the difference between the coefficients of C_{jt}^k in the regressions explaining (i) the proportion of members of k ending up on j's board and (ii) the proportion of members of l ending up on j's board.

Our data sources have two dimensions, firm and individual, so clearly an equivalent strategy can be derived using the individual dimension. The advantage of aggregating equation (1) at the individual level is that we can dispense with the assumption that directors are identical with respect to observables ($X_{it} = 1$). Analogously, it is convenient to assume that firms are identical ($Z_{jt} = 1$). Thus, even though we make different assumptions about the matching process of directors to firms in our derivation of the individual- and firm-level models, we view the findings that result as being complementary. This strategy is described in the Online Appendix.

5.3. Sources of Identification

It is crucial to understand why our transformations, in both the individual- and firm-level models, are able to dispose of the pure individual and firm effects even in the cross-section. The intuition is that our identification strategy is similar to the "within" transformation used in panel data analysis. To see how, let us focus on a version of equation (1) with fixed effects only ($Z_{jt} = X_{it} = 1$). For each individual i, we know in which firms this individual is a director and also in which firms she is *not* a director. This differs from typical wage models with pure individual and firm effects in employer–employee data sets (see Abowd and Kramarz 1999) because the wage paid to individual i is known only for those firms where she is employed. In our setting, all the "$E_{ij} = 0$" observations yield information on the individual effect. Because there are many such observations, the data has enough identifying power to eliminate the pure individual effect (as described in the Online Appendix). Similarly, for any firm j, all those who are not on j's board yield information about firm j's propensity to hire directors in general. Because there are many such observations, it is relatively easy to eliminate the pure firm effect by using an appropriate transformation, as described in Section 5.2.

5.4. Possible Biases

There are multiple sources of estimation biases. First, measurement error obviously could arise if our categorization of the various networks were inappropriate. Yet, even

unbiased mistakes in measuring networks would a priori attenuate the magnitude and significance of our estimates.

Second, our model controls for *observable* tendencies of firms to hire directors from particular networks—for example, if firms in regulated industries tend to hire former civil servants (the $Z'_{jt} \cdot M \cdot X_{it}$ term in equation (1)). However, our approach does not control for *unobservable* firm "tastes" for some networks—as when, for example, some firms (because of their corporate culture) have a tradition of promoting and hiring engineers rather than top-level bureaucrats. This limitation of our approach is easy to see in the individual-level model (see Section A.3 of the Online Appendix), in which director observables are allowed to vary ($X_{it} \neq 1$). Let us look at the propensity of firms to hire from particular networks; in the language of model (1), this means $X_{it} = (A_i^m)$ for some m. As shown by equation (A.3) in the Online Appendix, a linear regression will not be able to identify this effect ($c_t^k \cdot X_{it}$) separately from network effects ($d_t^{km} \cdot A_i^m$). In theory we could account for this by including a firm fixed effect in equation (2) (see our derivation in the Online Appendix). Unfortunately, there is very little turnover of ENA CEOs when they do leave their replacement CEO is almost always another ENA graduate. Clearly, introducing firm fixed effects in equation (2) would make parameters hard to identify. Hence, it is in practice difficult to identify a given firm's fixed tendency to hire (say) ENA graduates as directors and CEO *separately* from its tendency to hire such graduates because the CEO is an ENA graduate.

Third, it is impossible to control for sorting along unobservable characteristics *on both sides* (pure unobservable matching). If directors of high IQ tend to join firms with high-IQ CEOs and if IQ is correlated with Grandes Ecoles graduation, then our estimates will be upward biased. This concern is difficult to address.

6. Evidence of Networks

6.1. Estimating the Probability of Employment

In a first step, we shall abstract from matching considerations and simply posit that $X_{it} = Z_{jt} = 1$, which means that generally some firms have a greater tendency than others to appoint and some individuals have a greater tendency than others to be appointed. (We will deviate from these assumptions in Section 6.2). In most of this section we focus on firm-level aggregations of equation (1). We obtain similar results in Online Appendix Table A.2 when using individual-level aggregation, which rests on slightly different assumptions about heterogeneity.

We start by estimating the following version of (2):

$$\frac{n_{jt}^k}{n_t^k} - \frac{n_{jt}^0}{n_t^0} = a_t^k + \sum_m \underbrace{(\lambda_{mk} - \lambda_{m0})}_{c_{km}} C_{jt}^m + u_{jt}^k, \tag{3}$$

where j indexes the firm, t indexes time, and k stands for the network under scrutiny (ENA, Polytechnique with civil service, or Polytechnique without civil service). Equation (3) is obtained by subtracting equation (2) for network k from equation (2) for network 0. Thus, the difference from the previous firm-level equation is that we take one network as the reference. Now, the variable on the left-hand side is the fraction of members of network k that are employed in firm j *minus* the fraction of members of the reference network that are employed in firm j. We define the reference category as those who are members of neither ENA nor Polytechnique networks. Finally, u_{jt}^{k} is an error term and the indicator C_{jt}^{m} is set equal to 1 only when firm j's CEO belongs to network k. We are interested in the coefficients for these indicator variables ($\lambda_{mk} - \lambda_{m0}$), which have a simple structural interpretation: they measure the probability that a member of a given network k is a director of a firm run by a member of network m *minus* the probability that a member of k is a director in a firm run by a CEO who does not belong to any of the networks.

Table 3 reports estimates of (3) for all three networks of interest (ENA, Polytechnique with civil service, Polytechnique without civil service). The left panel (columns (1)–(3)) presents estimates with year dummies, and the right panel (columns (4)–(6)) presents estimates with further economic controls. These regressions are jointly estimated using the SURE method, which permits error terms of the three equations to be correlated with each other for a given firm. Indeed, if a given firm has many ENA directors, for example, then it will be less likely to have many Polytechnique graduates; hence the two equations are not totally independent. We also allow the error terms to be correlated across observations of a given firm, using the White correction method for standard errors. The last six rows of Table 3 report tests of the null hypothesis of equality of the CEO coefficients across equations.

We first comment on the left panel, which incorporates only year indicators. For civil servants, the coefficient for CEO's identity is always strong and economically significant; the probability of being director in a firm is increased on average by 0.5–1 percentage points when the CEO belongs to one of the two civil service–related networks (graduates from ENA or Polytechnique). This is a sizable increase when one considers that, with 600 firms, the probability of being employed in a particular firm averages only 0.2%.

Second, these results do not necessarily constitute strong evidence of network importance per se, since we are only comparing members of three networks to "mostly unconnected" directors. We thus test our hypotheses (H) and (H') more directly by checking whether, for a given director, the probability of being employed in a firm run by a CEO of the *same* network is significantly higher. In other words, we ask in equation (3) whether $c_{kk} > c_{km}$ for all m. These tests are reported in the last six rows of Table 3. The results indicate that the most important networks are former ENA graduates and former Polytechnique graduates with a civil service career, *not* Polytechnique graduates who went directly to the private sector. These results strongly support the intuitions of Kadushin (1995) and Frank and Yasumoto (1998): it is networks of former civil servants, not networks of private-sector engineers, that matter the most in this context.

TABLE 3. Econometric evidence on networks: effect of CEO background on director current employment.

	Firm level model					
Among currently employed directors, fraction of:	(1) ENA	(2) Polytechnique civil service	(3) Polytechnique private sector	(4) ENA	(5) Polytechnique civil service	(6) Polytechnique private sector
CEO is ENA grad.	0.62***	0.33***	0.13***	0.48***	0.35***	0.12*
	(0.08)	(0.08)	(0.04)	(0.10)	(0.11)	(0.06)
CEO is Polytechnique grad.	0.50***	0.97***	0.25***	0.42***	0.95***	0.19***
& former civil servant	(0.11)	(0.15)	(0.05)	(0.12)	(0.18)	(0.06)
CEO is Polytechnique grad.	0.21***	0.16**	0.16***	0.10*	0.09	0.17***
& always private sector	(0.06)	(0.06)	(0.03)	(0.06)	(0.06)	(0.04)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Former SOE dummy	No	No	No	Yes	Yes	Yes
Past-year firm ROA	No	No	No	Yes	Yes	Yes
Industry dummies	No	No	No	Yes	Yes	Yes
Observations		8,035			5,219	
Test ENA(1) = ENA(2)		0.00***			0.35	
Test ENA(1) = ENA(3)		0.00***			0.00***	
Test Poly, CS(2) = Poly, CS(1)		0.00***			0.01***	
Test Poly, CS(2) = Poly, CS(3)		0.00***			0.00***	
Test Poly, PS(3) = Poly, PS(1)		0.50			0.35	
Test Poly, PS(3) = Poly, PS(2)		0.97			0.36	

Notes: Reported values are SURE estimates (standard errors in parentheses). Residuals may be correlated across equations and observations of the same firm. All explanatory variables are lagged by one year. Sources: DAFSA yearbook (accounting variables); 1994 and 2000 *Who's Who* (directors' education); alumni directories (CEOs).
*Significant at 10%; **significant at 5%; ***significant at 1%.

To confirm the results obtained in Table 3, we use the individual-level model to run similar regressions and report the results in Online Appendix Table A.2, which has the same structure as Table 3. Given our assumptions that $X_{it} = Z_{jt} = 1$, results should be identical to the firm-level model (3) provided model (1) is not misspecified. In (1), the dependent variable is the fraction of seats held by individual i (at date t) that correspond to firms run by CEOs of network k. As it turns out, the same orders of magnitude and the same test statistics are obtained with this alternative way of collapsing the data. The only difference that emerges using this model is that ENA directors are just as likely to sit on boards of firms run by ENA CEOs as they are to sit on boards of firms run by Polytechnique civil servants. This suggests that the different civil service–related networks have links with each other, a pattern that we will find again in subsequent analyses.[5]

6.2. When Directors and CEOs Sort on Other Dimensions

Now, we assess the biases arising from possibility of directors sorting with firms according to observable or unobservable characteristics. We start by re-estimating our firm-level regressions while including as in equation (2) observable firm characteristics: a dummy set equal to 1 for former state-owned enterprises, industry dummies, and the firm's past profitability (in terms of return on assets, ROA, lagged by one year). This approach allows us to account for these observables being important to directors endowed with particular unobservable characteristics that might be correlated with networks. This we do, in columns (4)–(6) of Table 3, for each of our three focal networks. In fact, these controls do not have much of an effect on our estimates. The only change is that now firms run by ENA graduates are as likely to hire former civil-servants from ENA as from Polytechnique. This does not alter our general conclusion that civil servant networks are active whereas networks of Polytechnique graduates without bureaucratic careers are not. Hence, accounting for other possible sorting processes, which could overlap network effects, does not affect our results either quantitatively or qualitatively.

In the last three columns of Online Appendix Table A.2, we use individual-level regressions to control for director characteristics (age and years of education)—rather than firm-level regressions on firm-level characteristics as in Table 3. The results obtained are similar to those reported for regressions with year indicators only.

6.3. Estimating the Probability of Appointment

An important question raised by the previous regression results is whether the CEO's identity matters or whether instead it simply proxies for the board's identity. Suppose that the CEO holds no real power in appointments and that all such power in these

5. We also looked at the difference between the largest firms(within the premier marché) and the smallest (within the second marché). We find that premier-marché firms are where most of the action takes place, although some civil service–related networks appear to be operating in the second marché.

matters rests with the board of directors. Then, the board will appoint CEOs that are similar to the set of directors, in which case the causal relation would be reversed. Although this would still constitute evidence that social networks interfere with the labor market, the direction of the relation matters for corporate governance. Indeed, if the board turns out to be chosen by the firm's CEO then directors' ability to monitor management on behalf of the shareholders might be severely impaired.

We address this issue in two steps. First, we re-estimate model (1) by looking at *appointments* rather than employment. Under this new interpretation, $E_{ijt} = 1$ when i is appointed by firm j at date t. We use firm-level aggregation and thus correlate the CEO's identity with the firm's hiring policy, which yields a more stringent test of social interactions.[6] Second, we check for whether the CEO's identity in these appointment regressions is a proxy for initial board composition by including in the regression the past number of directors on the board from either network. This amounts to running the following modified version of (3):

$$\frac{n^k_{jt}}{n^k_t} - \frac{n^0_{jt}}{n^0_t} = a^k_t + b^k_{jt} + \sum_m c_{km} \cdot C^m_{jt} + \sum_m c'_{km}(\#A^m_{jt}) + u^k_{jt}.$$

Here the variable the left-hand side is now the share of network-k members newly hired by firm j *minus* the share of directors newly hired by j and $\#A^m_{jt}$ denotes the fraction of network-m members who have already been appointed to firm j's board. Note that such a regression could *not* be estimated using employment instead of appointment (as in our previous specifications) because of the co-called reflection problem (Manski 1993). That is, if A and B are similar and are both on the same board, then it is difficult to know whether A's presence is due to B's or vice versa. By introducing some dynamics, this methodology makes a "Granger causality" argument: it is A who matters if A was on the board *prior to* B.

The results of these firm-level regressions for our three selected networks are presented in Table 4. Estimation of all three equations is made jointly using the SURE methodology, and we allow for flexible correlation across observations of a same firm using the White correction. As before, industry and year indicators are included. To avoid spurious correlations, explanatory variables are lagged one year. In the table, columns (1)–(3) are based on the equivalent of equation (3) (i.e., assuming $c'_{km} = 0$); columns (4)–(6) add the controls for past board composition.

The regression results from columns (1)–(3) confirm our previous findings: education (ENA and Polytechnique versus the rest) and career (civil service versus private sector) networks affect the allocation of directors to firms, even when we analyze nominations (rather than employment). Results reported in columns (4)–(6) support the notion that CEO identity, not board composition, explains the selection of directors' appointments. First, even though including the board composition variables reduces slightly the difference between coefficients for CEO identity (compare test

6. We also ran individual-level regressions using appointments instead of employment and obtained strongly similar results (not reported here).

TABLE 4. Econometric evidence on networks: effect of CEO background on directors' *appointment*.

	Firm-level regressions					
Among currently employed directors fraction of:	(1) ENA	(2) Polytechnique civil service	(3) Polytechnique private sector	(4) ENA	(5) Polytechnique civil service	(6) Polytechnique private sector
CEO is ENA grad.	0.13***	0.06***	0.03***	0.09***	0.04***	0.02**
	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.01)
CEO is Polytechnique grad.	0.10***	0.23***	0.03***	0.05**	0.18***	0.02
& former civil servant	(0.02)	(0.04)	(0.01)	(0.02)	(0.04)	(0.01) .
CEO is Polytechnique grad	0.04***	0.05***	0.05***	0.02	0.03**	0.04***
& always private sector	(0.02)	(0.02)	(0.01)	(0.01)	(0.02)	(0.01)
% of ENA directors (-1)	–	–	-	0.35***	0.12***	0.10***
				(0.04)	(0.04)	(0.03)
% of Poly. former CS	-	-	-	0.17***	0.36***	0.02
directors (-1)				(0.05)	(0.11)	(0.03)
% of Poly.. always PS	-	-	-	0.09***	0.03	0.07***
directors (-1)				(0.03)	(0.03)	(0.02)
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations		6,759		6,757		
Test ENA(1) = ENA(2)		0.01***		0.00***		
Test ENA(1) = ENA(3)		0.01***		0.00***		
Test Poly. CS(2) = Poly. CS(1)		0.00***		0.00***		
Test Poly. CS(2) = Poly. CS(3)		0.00***		0.00***		
Test Poly. PS(3) = Poly. PS(1)		0.72		0.18		
Test Poly. PS(3) = Poly. PS(2)		0.99		0.87		

Notes: Reported values are SURE estimates (standard errors in parentheses). Residuals may be correlated across equations and observations of the same firm. All explanatory variables are lagged by one year. Key: CS = civil servant; PS = private sector. Sources: DAFSA yearbook (accounting variables); 1994 and 2000 Who's Who (directors' education); alumni directories (CEOs).
**Significant at 5%; ***significant at 1%.

values for the first regression with those for the second), all c'_{km} coefficients for board composition remain significant and strongly positive, and all tests yield results that are virtually identical to those presented in Table 3. In addition, we now have similar results for boards: boards dominated by former civil servants tend to recruit new directors from the networks (i.e., Polytechnique or ENA) to which they belong.

7. Networks and Corporate Governance

Our findings so far suggest that networks of former high-ranking civil servants are particularly active in shaping board composition. When the CEO is a former civil servant (whether a graduate from Ecole Polytechnique or ENA), the fraction of directors from the same background (both in stock and in flow) is larger.

In principle, such arrangements may arise for two distinct reasons. In well-governed firms, CEOs may use their own social networks to find directors whose advice and monitoring will be more effective. When corporate governance is poor, CEOs may use their networks to hire friendly (or simply passive) directors who will rubber-stamp their decisions. Hence, the presence of social networks in the boardroom could signify either good or bad governance.

In order to shed light on this issue, this section examines the quality of corporate governance of firms run by former high-ranking civil servants. We do this by using three indicators that the literature has found to be correlated with (or good indicators of) governance: the sensitivity of CEO turnover to performance, CEO compensation, and M&A quality.

7.1. Sensitivity of CEO Turnover to Performance

We first use turnover to performance sensitivity as a measure of corporate governance. Weisbach (1988) shows that, when a firm underperforms, its CEO is more likely to leave when the board of directors is independent. To test for the existence of this dynamic, we run—separately for connected (i.e., with an early career in civil service) and nonconnected (i.e., with purely private-sector careers) CEOs—the following logistic regression:

$$T_{jt+1} = \alpha + \beta \cdot \text{PERF}_{jt} + \delta \cdot \text{controls}_{jt} + \varepsilon_{jt}, \tag{4}$$

where T_{jt+1} is a dummy variable set equal to 1 if the CEO loses his job over the next year (between t and $t + 1$). We then compare the turnover-to-performance sensitivity coefficient β for both categories of CEOs and test for equality. If social networks impair governance, then β should be less negative for connected CEOs. Like prior papers in this literature, we do not observe dismissals and so must consider all types of turnover; in an attempt to exclude voluntary retirement and reduce measurement error, we restrict ourselves to the sample of CEOs aged less than 65. The term $PERF_{jt}$ is an industry-adjusted measure of corporate performance (in terms of return on assets and cumulative stock returns). Since the dependent variable is binomial, we run logistic

TABLE 5. Turnover-to-performance sensitivity of connected CEOs.

	Losing CEO position in forthcoming year		
	Former civil service	Private sector	Difference
Panel A: No controls			
Industry-adjusted ROA	−2.61	−8.18***	5.56**
	(2.51)	(1.45)	(2.87)
Observations	498	1,793	0.96
Industry-adjusted stock returns	−1.21	−2.17***	
	(0.86)	(0.51)	(1.02)
Observations	346	860	
Panel B: With controls			
Industry-adjusted ROA	−1.22	−8.91***	7.69**
	(3.22)	(1.58)	(3.59)
Observations	461	1,768	1.44
Industry-adjusted stock returns	−0.90	−2.34***	
	(0.96)	(0.56)	(1.12)
Observations	302	774	

Notes: Reported values are logit estimates (standard errors in parentheses). The tested sample consists of all firms run by a CEO aged less than 65 for all years after 1991. Panel A regresses the CEO's not returning in the next year on industry-adjusted measures of annual firm performance (return on assets and annual stock returns); there are fewer observations for stock returns because we match accounting and returns data. Panel B adds controls to this regression: log of assets as well as industry and year dummies. Column (1) estimates the two models on the subsample of former civil servants, and column (2) restricts the sample to CEOs who never were civil servants. Column (3) tests for equality of the coefficients reported in columns (1) and (2); for this, we (i) regress future turnover on all RHS variables interacted with a civil service dummy and (ii) report the coeffcient for the interaction between civil service and firm performance. In all regressions, error terms are clustered at the firm level.
**Significant at 5%; ***significant at 1%.

regressions and allow error terms ε_{jt} to be correlated in a flexible fashion across observations of a given firm.

Results (with standard errors in parentheses) are reported in Table 5 and support the hypothesis that social networks in the board room reduce the effectiveness of governance. In Panel A, we include no controls; in Panel B we control for firm size (log of book assets) as well as for industry and year fixed effects. Column (1) reports the estimate of β in the sample of firms run by former civil servants, and column (2) does the same on the (smaller) sample of firms run by CEOs with the alternative background. Overall, turnover appears less sensitive to bad performance for former civil servants, since β is smaller in absolute value; this remains true whether or not we include the controls and also for both performance measures. The difference is economically large: when performance is measured by ROA, β increases from 2 (former civil servants) to 8 (private sector). The coefficient is statistically insignificant for former civil servants but is strongly significant for private-sector CEOs, although this might be due to the smaller sample of firms run by former civil servants.

In column (3) we perform a statistical test: we re-estimate model (4) on the whole sample, interacting all right-hand-side variables with the civil servant dummy; we then report the coefficient for profit interacted with this dummy, which is exactly equal to

the difference between the estimated values for β reported in columns (1) and (2). The difference in turnover-to-performance sensitivities is statistically significant (at the 5% level) when performance is measured by ROA, but the difference is insignificant when we use stock returns. Overall, we find evidence that connected CEOs are less likely to depart when the company they run underperforms. This is in line with results from Nguyen (2009), who uses the largest firms of our sample.[7]

7.2. CEO Compensation

In a cross-section of US firms, the level of CEO compensation has been found to correlate strongly with poor corporate governance (see Bebchuk and Weisbach (2010) and the references therein). A priori, a high level of compensation may mean that shareholders have a strong need to incentivize the CEO; under this "optimal contracting" view, a high level of compensation simply reflects agency rents appropriated by CEOs but willingly granted by shareholders. Under the "CEO power" view, shareholders are too weak to fight the CEO's demands. The existing literature finds evidence consistent with this second view: compensation is higher when there is no large shareholder, when directors are "busy" (in the sense that they sit on the boards of many other companies), and when the firm's charter includes anti-takeover provisions.

In our French setting, if civil service–related social networks are detrimental to the quality of firm governance then the pay of former civil servants should be higher. We shall test this hypothesis by using hand-collected data on CEO compensation, which we use to regress the log of CEO compensation on a "former civil service" dummy while controlling for firm size and industry.

Data collection imposes severe limitations on our research design. First, through most of our sample period, French listed firms were not required to disclose CEO compensation in their annual reports; in 2002, less than 5% of them willingly chose to do so. Starting in 2003, however, the New Economic Regulation Act (passed in 2001) made it mandatory for listed firms to disclose not only CEO compensation (salary and bonus) but also stock options granted (number, date of grant, strike price, maturity, vesting period). We therefore focus our analysis on 2003, retrieving annual reports from the securities regulator's web site.[8] From a subsample of 555 firms in our dataset, we find annual reports for only 224 firms—though all of them include the value of CEO pay. Out of these 224, some 178 provided a breakdown of total cash compensation into bonus and fixed salary. Of these latter firms, 75 report CEO stock options; however, we are able to compute the Black–Scholes value for only 52 of these firms because in many cases stock return data are missing.[9] Hence the variable "total

7. Indeed, in regressions suggested by a referee (not reported), we find that most of the action is in the premier marché.

8. Autorité des Marchés Financiers (AMF; http://www.amf-france.org/).

9. To value these options, we (i) compute the annual volatility using daily returns over the twelve months prior to the grant and (ii) take the stock price on the last day of the last month prior to the grant. We assume a risk-free rate of 4%.

TABLE 6. Compensation of connected CEOs.

	log(1+salary)		log(1+bonus)		log(1+S.O.)		Log (TC)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Former	0.72***	0.14	5.44***	0.99	4.76***	2.32*	1.45***	0.38*
Civil service	(0.18)	(0.16)	(1.19)	(1.39)	(1.16)	(1.35)	(0.26)	(0.21)
log(assets)	-	0.24***	-	1.55***	-	1.08***	-	0.38***
		(0.03)		(0.28)		(0.26)		(0.04)
Fraction held by	-	-0.01**	-	0.00	-	−0.02	-	−0.01***
largest shareholder (%)		(0.00)		(0.02)		(0.02)		(0.00)
log(1 + age of firm)	-	0.09	-	0.55	-	0.60	-	0.16*
		(0.07)		(0.96)		(0.60)		(0.09)
Industry fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
Observations	180	147	178	145	201	164	200	163

Notes: Reported values are from ordinary least-squares (OLS) regressions (standard errors in parentheses). The leW-hand side (LHS) is the log of CEO compensation as reported in the 2003 annual reports available from the securities regulator's (AMF) web site. Columns (1) and (2) use the log of fixed compensation, and columns (3) and (4) use the log of bonus. Columns (5) and (6) use the log of (1 + stock options); because 149 firms report some compensation in the form of bonus or salary but no stock options, for them we set those options to zero (a reasonnable assumption, since the law mandates disclosure of all stock option grants). Columns (7) and (8) use the log of total compensation (i.e., salary + bonus + stock option grants) as the LHS variable. The odd-numbered columns report regression results without controls; each even-numbered column includes four controls (firm size, industry fixed effects, percentage held by dominant shareholder, age of the firm since creation). Stock options are valued using the Black Scholes formula with the strike price reported in the firm's annual report, the stock price at the end of the last month preceeding the grant, and the annualized stock price volatility of daily returns over the twelve months preceeding the grant. Key: S.O. = stock options; TC = total compensation.
*Significant at 10%; **significant at 5%; ***significant at 1%.

compensation", which includes all three types of payments, is missing for firms that report option grants but for which we cannot compute the Black–Scholes value.

We provide regression results for compensation and its components in Table 6. There are two salient features. First, former civil servants receive much higher levels of compensation. In columns (1), (4), (7), and (10) we make a raw comparison between the two types of CEOs (no controls in the regressions). For former civil servants, the salary is twice as large ($e^{0.7}$) and the bonus is 2.5 times larger ($e^{0.9}$); their average option grant is about 100 times larger than grants to non–former bureaucrats. This factor is due primarily to the much greater likelihood of former civil servants being granted any stock options: about 30% of them receive this form of compensation whereas only 9% of other CEOs do. All these differences are strongly statistically significant. Overall, the total compensation of a former civil servant is 4.5 times ($e^{1.5}$) larger than that of top executives from alternative backgrounds.

Second, this compensation discrepancy is partly (but not entirely) explained by the fact that connected CEOs run larger firms. In columns (2), (5), (8), and (11) of Table 6 we control for size (log of book assets) and industry dummies. This shrinks the excess salary of connected CEOs to almost zero. Although the difference in bonuses remains large (150%), it is rendered insignificant by the size control. The stock options differential is much larger ($e^{2.2} - 1 = 800\%$)—and is significant at the 10% level. As noted previously, this difference is largely explained by the higher probability that ex–civil servants receive stock options. After combining the factors (and controlling

for firm size), we find that—even though larger firms pay better—connected CEOs receive overall compensation about 50% ($e^{0.4} - 1$) higher than that of nonconnected CEOs. The strong explanatory power of controls raises the concern that our remaining results (columns (8) and (11)) are fragile. To strengthen the analysis, we add controls for governance (which is known to be correlated with CEO compensation) in columns (3), (6), (9), and (12) (these controls include age of the firm since creation and fraction of stocks held by the largest shareholder). Our results are unaffected: former civil servants tend to receive higher compensation, especially in the form of stock options.

One possible interpretation for our results is that former civil servants receive more performance-related compensation and enjoy lower agency rents (see Jensen and Murphy 1990). However, there is evidence in the compensation literature that agency rents would have to be implausibly large to account for the amount of stock options granted that is observed in the data (Hall and Liebman 1998; Bertrand and Mullainathan 2001). Therefore, stock options seem more consistent with the CEO power hypothesis than with shareholders designing optimal contracts (Bebchuk and Weisbach 2010).

7.3. M&A Quality

Here we use the quality of acquisitions as an indirect measure of firm governance. Following the finance literature, we proxy the quality of acquisitions via the acquiring firm's stock price reaction to the transaction's announcement. According to the efficient market hypothesis, this reaction measures the present value (net of acquisition costs) of the deal to the acquirer's shareholders. Low-quality acquisitions are generally viewed as a waste of the acquiring firm's free cash flows (Lang, Stulz, and Walking 1991). An acquirer whose stock price reacts badly to a deal announcement tends to score low on standard corporate governance indices (Masulis, Wang, and Xie 2007).

We obtain acquisition data from Securities Data Company (SDC). We retrieve from this database all "completed" acquisitions initiated between 1992 and 2003 by French firms that were either directly listed or listed through their ultimate parent. We then further restrict the sample to acquisitions whose transaction value (in millions of US dollars) is reported in the data and for which the fraction of shares held by the acquirer after the transaction is at least 50%. We end up with 1,469 deals. We then manually merge by company the resulting transaction data with our main dataset. This process leaves us with 1,103 acquisitions for which we have the transaction value, the target and acquirer's four-digit SIC codes (from SDC), and the acquirer's accounting information and CEO background. Finally, we match the resulting data set with stock return data to yield 961 deals for which we can compute the acquirer's announcement returns.

Using this metric, we find that acquisitions made by connected CEOs create less value for the firm. We calculate cumulative market-adjusted returns, starting five days before the announcement and lasting until five days afterwards. These announcement reactions are plotted in Figure 2. Upon announcement, the stock price of acquirers

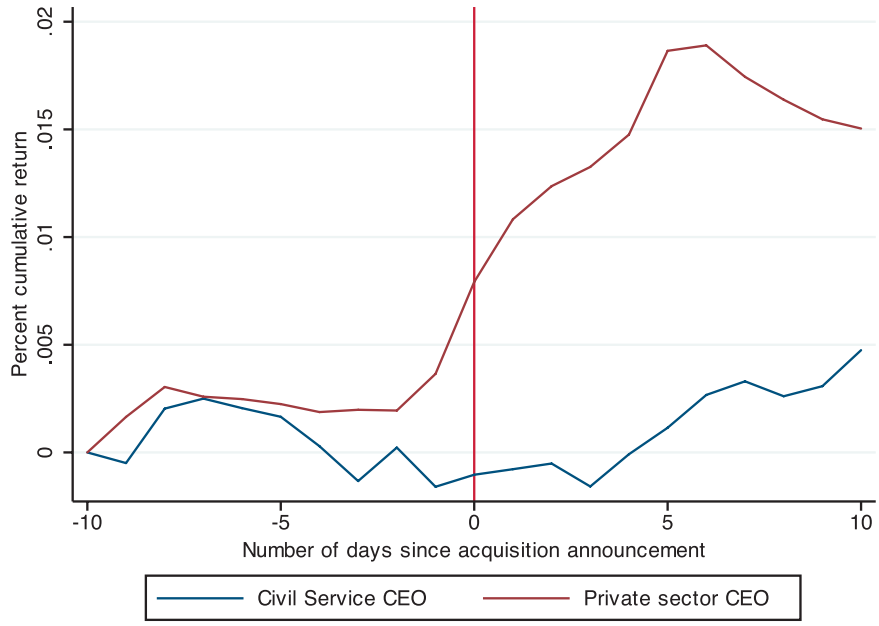

FIGURE 2. Stock price reaction (by acquirer type) to acquisition announcement.

run by nonconnected CEOs increases by 1.5%—a reaction that is consistent with the evidence from US data.[10] When the firm is run by a former civil servant, however, we find no such positive reaction. Hence, the market views acquisitions made by former civil servants as less value creating.

We report formal statistical tests in Table 7. Panel A is limited to the cumulative excess returns regressed on the civil servant dummy (simple mean return comparison); Panel B controls for 18 industry indicators as well as for acquirer and target sizes, which have been shown to affect the quality of acquisitions. Columns (1)–(4) report at the cumulative price change over different time spans surrounding the acquisition. Columns (1) and (2) reveal that pre-announcement price movements do not differ significantly across CEO background until one day *before* the announcement. In other words, there is no evidence of more insider trading in firms run by CEOs with a civil service background. The difference in price movements becomes one percentage point (or more) one day *after* the announcement; this difference is persistent and remains statistically significant even when we control for acquirer size, deal size, and industry. These numbers and tests confirm the intuition derived from Figure 2: connected CEOs seem to make acquisitions of significantly lower quality, given that they create less shareholder value than do acquisitions made by non connected CEOs. This finding is consistent with connected firms being less well governed.

10. Using a sample similar to ours (SDC, both public and private targets), Bradley and Sundaram (2006) also find a positive announcement return of 1.5%.

TABLE 7. Quality of acquisitions made by connected CEOs.

	R(-10,−5)	R(-5,-1)	R(-1,1)	R(0,5)
Panel A: No control				
Former civil service	−0.04 (0.30)	−0.46 (0.38)	−0.66*** (0.24)	−0.76** (0.32)
Observations	939	939	936	926
Panel B: With controls				
Former civil servant	0.45	−0.48	−0.82**	−1.14***
	(0.48)	(0.47)	(0.37)	(0.41)
log(deal size)	0.02	0.10	0.10	0.09
	(0.10)	(0.11)	(0.10)	(0.12)
log(acquirer size)	-0.01	0.19	0.06	0.33**
Year dummies	(0.10)	(0.17)	(0.11)	(0.16)
	Yes	Yes	Yes	Yes
Industry dummies	Yes	Yes	Yes	Yes
Observations	716	716	715	709

Notes: In all regressions, the dependent variable is the cumulative market-adjusted return around the announcement of an acquisition. Column (1) uses the return between ten and five days before the announcement, column (2) uses the return between five days and one day prior to the announcement, column (3) uses the return between one day before and one day after the announcement, and column (4) uses the return from the announcement day to five days after. We restrict the sample to acquisitions whose value is reported in SDC. Panel A reports OLS regression results of the cumulative return on the public service dummy with no other controls. Panel B includes several controls: log of deal size (millions of euros), a year-of-deal dummy, 18 industry dummies, and log of the acquirer's total assets (millions of euros). All error terms are clustered at the firm level (*standard errors* in parentheses).
**Significant at 5%; ***significant at 1%.

TABLE 8. Acquisitions by connected CEOs.

	At $t+1$			
	Number of # acquisitions		log(1 + value of all acquisitions)	
	(1)	(2)	(3)	(4)
Former civil	1.45***	−0.06	0.82***	0.25**
service log(assets)	(0.18)	(0.15)	(0.15)	(0.12)
	-	0.60***	-	0.27***
		(0.03)		(0.03)
Year fixed effects	Yes	Yes	Yes	Yes
Industry fixed effects	No	Yes	No	Yes
Observations	7,291	6,094	7,291	6,094

Notes: We restrict the sample to aacquisitions whose value is reported in SDC. In all regressions, the LHS measures M&A activity. Columns (1) and (2) incorporate a dummy variable set equal to 1 only if the firm makes an acquisition in year $t + 1$; columns (3) and (4) use the log of the sum of deal values. Columns (1) and (3) (Poisson regressions) report results without controls (except for year fixed effects); columns (2) and (4) (OLS regressions) also include industry fixed effects and firm size (as measured by log of total assets). All error terms are clustered at the firm level (standard errors in parentheses).
**Significant at 5%; ***significant at 1%.

There is weaker evidence that connected CEOs engage in more of these non–value creating acquisitions. In Table 8 we examine the frequency and amount of such acquisitions. In columns (1) and (2), the dependent variable is the number of acquisitions (and so we run Poisson regressions). The coefficient indicates that the average annual number of acquisitions is 15% higher in firms run by connected CEOs.

Yet when we control for the larger size of these firms and for year and industry dummies, the effect vanishes. Columns 3 and 4 focus on the overall cost of these acquisitions. After controlling for differences in composition by size, year, and industry, we find that the overall annual cost of acquisitions in connected firms is nearly 26% higher. This number is statistically significant at the 5% level.

8. Conclusion: Social Networks and Corporate Performance

This paper has demonstrated that social networks do indeed appear to shape board composition. We used French data because the history and sociology of the French business elite make it fairly easy to establish whether a given CEO or director belongs to a particular network. We have used new data and new techniques to identify the existence of networks. As it turns out, the network of former bureaucrats is the most active in determining board composition when we control for both director and firm characteristics. This phenomenon has direct implications for the sociology of the French elite, for the economics of networks, and for corporate governance. For the French elite and the role of social capital, we see that the market is far from eliminating the effects of networks. With respect to the economics of networks, the econometric techniques developed here are particularly well suited to the study of several questions of economic interest (e.g., the impact of networks within firms). As regards corporate governance, we show that firms whose directors and CEOs have a past career in the civil service are less likely to change CEOs in response to poor performance, that connected CEOs are paid more than their nonconnected counterparts, and that connected CEOs make acquisitions that are larger and of less value (as rated by the market). These findings suggest that social networks have multiple effects, most of which are detrimental to good governance.

References

Abowd, John M. and Francis Kramarz (1999). "The Analysis of Labor Markets Using Matched Employer–Employee Data." In *Handbook of Labor Economics*, Vol. 3, edited by O. Ashenfelter and D. Card. Elsevier, pp. 2629–2710.

Barnea, Amir and Ilan Guedj (2008). "*Director Networks and Firm Governance*." Working paper.

Bauer, Michel and B. Bertin-Mourot (1997). "La tyrannie du diplôme initial et la circulation des élites: La stabilité du modèle français." In *Le recrutement des élites en Europe*, edited by Mendras and Suleiman. La Découverte, Paris.

Bayer, Patrick, Stephen Ross, and Giorgio Topa (2005). "*Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes*." NBER Working Paper No. 11019.

Bebchuk, Lucian and Michael Weisbach (2010). "The State of Corporate Governance Research." *The Review of Financial Studies*, 23, 939–961.

Bertrand, Marianne, Erzo Luttmer, and Sendhil Mullainathan (2000). "Network Effects and Welfare Culture." *Quarterly Journal of Economics*, CXV, 1019–1056.

Bertrand, Marianne and Sendhil Mullainathan (2001). "Are CEOs Rewarded for Luck? The Ones without Principals Are." *Quarterly Journal of Economics*, 116, 901–932.

Bourdieu, Pierre (1989). *La Noblesse d'Etat: Grandes Ecoles et Esprit de Corps*. Editions de Minuit, Paris.

Bradley, Michael and Anant Sundaram (2006). "*Acquisitions and Performance: A Re-assessment of the Evidence*." Working paper, Duke University.

Braggion, Fabio (2011). "Managers and (Secret) Social Networks: the Influence of the Freemasonry on Firm Performance." *Journal of the European Economic Association*, 9, 1053–1081.
Cohen, Lauren, Andrea Frazzini, and Christopher Malloy (2008). "The Small World of Investing: Board Connections and Mutual Fund Returns." *Journal of Political Economy*, 116, 951–979.
Frank, Kenneth and Jeffrey Yasumoto (1998). "Linking Action to Social Structure in the System: Social Capital Within and Between Subgroups." *American Journal of Sociology*, 104, 642–686.
Granovetter, Mark (1973). "The Strength of Weak Ties." *American Journal of Sociology*, 78, 1360–1380.
Hall, Brian and Jeffrey Liebman (1998). "Are CEOs Really Paid like Bureaucrats?" *Quarterly Journal of Economics*, 63, 653–691.
Hochberg, Yael, Alexander Ljungqvist, and Yang Lu (2007). "Whom You Know Matters: Venture Capital Networks and Investment Performance." *Journal of Finance*, 62, 251–301.
Hwang, Byoung-Hyoun and Seoyoung Kim (2009). "It Pays to Have Friends." *Journal of Financial Economics*, 93, 138–158.
Jensen, Michael and Kevin Murphy (1990). "Performance Pay and Top Management Incentives." *Journal of Political Economy*, 98, 225–264.
Kadushin, Charles (1995). "Friendship Among the French Financial Elite." *American Sociological Review*, 60, 202–221.
Kramarz, Francis and Oskar Skans (2010). "With a Little Help From My . . . Parents: Family Networks and Youth Labor Market Entry." CREST Working paper.
Kuhnen, Camelia (2009). "Business Networks, Corporate Governance and Contracting in the Mutual Fund Industry." *Journal of Finance*, 64, 2185–2220.
Lang, Larry, René Stulz, and Ralph Walking (1991). "A Test of the Free Cash Flow Hypothesis: The Case of Bidder Returns." *Journal of Financial Economics*, 29, 315–336.
Larcker, David, Scott Richardson, Andrew Seary, and Irem Tuna (2005). "*Back Door Links Between Directors and Executive Compensation*." Working paper.
Liu, Yun (2008). "Employment Networks and the CEO Labor Market." Working paper.
Manski, Charles (1993). "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, 60, 531–542.
Masulis, Ronald, Cong Wang, and Fei Xie (2007). "Corporate Governance and Acquirer Returns." *Journal of Finance*, 62, 1851–1890.
Munshi, Kaivan (2003). "Networks in the Modern Economy: Mexican Migrants in the US Labor Market." *Quarterly Journal of Economics*, 549–599.
Nguyen, Bang Dang (2009). "Does The Rodolex Matter? Corporate Elite's Small World and the Effectiveness of Boards of Directors." Working paper, Chinese University of Hong Kong.
Rees, Albert (1966). "Information Networks in the Labor Market." *American Economic Review*, 56(1–2), 559–566.
Saloner, Garth (1985). "Old Boys Networks as Screening Mechanisms." *Journal of Labor Economics*, 3, 255–267.
Suleiman, Ezra (1997)."Les élites de l'administration et de la politique dans la France de la V° république: Homogénéité, Puissance, Permanence." In *Le recrutement des élites en Europe*, edited by Mendras and Suleiman. La Découverte, Paris.
Swartz, David (1986). "French Corporate Leadership: A Class Based Technocracy." *Research in Political Sociology*, 2, 49–79.
Weisbach, Michael (1988). "Outside Directors and CEO Turnover." *Journal of Financial Economics*, 20, 431–460.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Online Appendix For: "Social Networks in the Boardroom"