Available online at www.sciencedirect.com
ScienceDirect
Social Science Research 36 (2007) 1156–1183
Social Science RESEARCH
www.elsevier.com/locate/ssresearch

Teaching executives to see social capital: Results from a field experiment ☆

Ronald S. Burt *, Don Ronchi

University of Chicago, Graduate School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA

Available online 7 November 2006

Abstract

This paper is about the benefits of teaching executives to understand the network structure of social capital. There is abundant cross-sectional evidence of performance correlated with network structure. Corroborating that evidence, we run a field experiment in which executives educated in the network structure of social capital show performance improvement relative to a control group of untrained, but otherwise equally able peers: program graduates are 36–42% more likely to receive top performance evaluations, 43–72% more likely to be promoted (an effect that builds in the 2 years following the program), and 42–74% more likely to be retained by the company. Active participation matters. The subsequent careers of executives who were quiet spectators in the program cannot be distinguished from the careers of people in the control group, peers who never attended the program.

Keywords: Social capital; Social network; Executive education; Adult learning; Training; Field experiment

☆ We are grateful for financial support from the University of Chicago GSB. Materials for the "Business Leadership Program" and "Strategic Leadership" course referenced in this paper can be downloaded at http://ChicagoGSB.edu/fac/ronald.burt/teaching. Related training can be found through links from the website of the International Network for Social Network Analysis (www.insna.org). The text has been improved in response to comments from colleagues Matt Bothner, Tracy Cox, Holly Raider, and the two reviewers for *Social Science Research*.

* Corresponding author. Fax: +1 773 702 9919.
E-mail addresses: ron.burt@chicagogsb.edu (R.S. Burt), dronchi@chicagogsb.edu (D. Ronchi).

0049-089X/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.ssresearch.2006.09.005

1. Introduction

One of the opening acts for the 21st century was venture capital discovering social networks. Tens of millions of dollars were invested in software companies purporting to make people better off through network services providing lucrative or emotionally rewarding relationships. Companies sprang up one after another like spring flowers, drawing comment from establishment media such as *Fortune* (Kirkpatrick, 2003) and the *Wall Street Journal* (Bulkeley and Wong, 2003), which in turn sent technology gossip all a-blog. From Boston's Route 128, Monster facilitated connections between employers and job seekers. Manhattan's SelectMinds facilitated connections among a company's former employees. Manhattan's theSquare facilitated connections between graduates of "quality" schools while Classmates Online, in Reston, took care of people from elsewhere. It was in Silicon Valley that the market tornado touched down: Friendster facilitated connections among registered users looking for companionship and Google's invitation-only network service, Orkut, made social life more active and stimulating by facilitating existing relations and introducing friends of friends. Data-rich Craigslist provided love and business contacts all in one screen. Targeting business, companies like LinkedIn, Tacit Knowledge Systems, and Spoke Software focused on facilitating new connections between professionals.

There is a common-sense business case to be made for these companies. Lucrative and emotionally rewarding relationships must be a good thing. Certainly, "bowling alone" is not good (Putnam, 2000). Motivational anecdotes posted on company websites describe people who used the company's service and are better off today.

Unknown is the contribution that the network service made to the people being better off. To what extent would the better-off people have been better off even if they had not used the service? How does an average user benefit from the service?

The same questions can be asked of the courses that emerged in the 1990s to train students in the career and organization-performance implications of social networks. Such questions are a slice of the broad challenge that Pfeffer and Fong (2002, p. 78) put to business education, asking for evidence of value: "What data there are suggest that business schools are not very effective: Neither possessing an MBA degree nor grades earned in courses correlate with career success, results that question the effectiveness of schools in preparing their students." More specifically, what research evidence can be cited to answer the suspicion that people skilled in the network intuitions of social capital are born, not trained? In every social-capital program we have taught, someone asks whether social capital is merely a new way to measure deep-seated personality (with extroverts building bridges across groups and introverts hiding inside a group). In fact, the correlation with personality is not at all clear (Burt, 2005, pp. 47–50) and our purpose in this paper is to show that social capital skills can be taught—to the great benefit of the executive learning the skills.

We run a field experiment that reveals positive effects, quite striking effects, of an executive education program grounded in network models of social capital. We hasten to emphasize that a variety of programs with which we are familiar could be producing the same effects. What is unique here are data with which we can estimate program effects.

Our results are from a field experiment in the sense that people given the training are observed unobtrusively in their natural environment after the training, and compared

over time to a control group of equally able, untrained peers. At the same time, our field experiment is of a limited kind in that people were not assigned at random to the treatment and control groups. People were selected for the program on the basis of their past performance. Executive education is expensive. Companies reserve it for their more promising people. There are selection biases in being selected for a senior executive program that are similar in many ways to the more familiar selection biases involved in being admitted to an elite college (Brand and Halaby, 2006). However, with access to the personnel files used to select people for the training, we were able to create a risk function that sorted the many untrained people for their probability of being selected for the program. Thus we had two groups: a group of people who received the training and a control group of people who were not trained, but at similar risk of being trained.[1]

This is a productive complement to the traditional cross-sectional research design in which performance indicators are correlated with network structure. If the studies are correct in claiming that current structure affected current performance, then enhancing the structure should improve future performance, or more generally, teaching people the network structure of social capital should result in performance improvement. That is the substance of our field experiment. We educated a few hundred executives in the network structure of social capital, worked briefly with them to apply the learning, then stood back and tracked their subsequent careers through the data in their personnel files relative to a control group of untrained but otherwise comparable peers. Performance improvements are substantial. The next step in this research direction would be to compare education programs that differ in conceptual emphasis to see how training on different dimensions of social capital are responsible for the performance improvements we report here.

2. Corporate balkanization and social capital

As venture capitalists were discovering social networks, Raytheon Company, one of America's leading electronics firms, was attacking the problem of how to coordinate across the organization silos of its acquired companies and its many product programs. In the wake of industry consolidation, operations were balkanized in the sense that programs that would do well to coordinate with one another were segregated in distant organization silos. In a nutshell, the task was to preserve efficiency and tacit knowledge within the silos, while harvesting the value of integrative work across the silos.

[1] We do not wish to draw too thick a line between field experiments and other evidence. The usual cross-sectional research design relies on post hoc statistical controls to compare people at different levels of social capital. We rely on post hoc statistical controls to create a risk function with which we define our control group. Nevertheless, our point in the text is that executive education—with its inherent purpose of increasing an X variable to affect the future value of Y—has the potential to provide more field-experiment-like control than the usual cross-sectional design. Harrison and List (2004) provide an overview of field experiments on economic questions similar to the returns to executive education analyzed here. Using Harrison and List's terminology, the results reported here are from a "natural field experiment" in that the treatment was normal company activity, subjects went back to their work environment after the treatment, and neither the subjects nor the people in the control group knew that their performance was being monitored for the experiment.

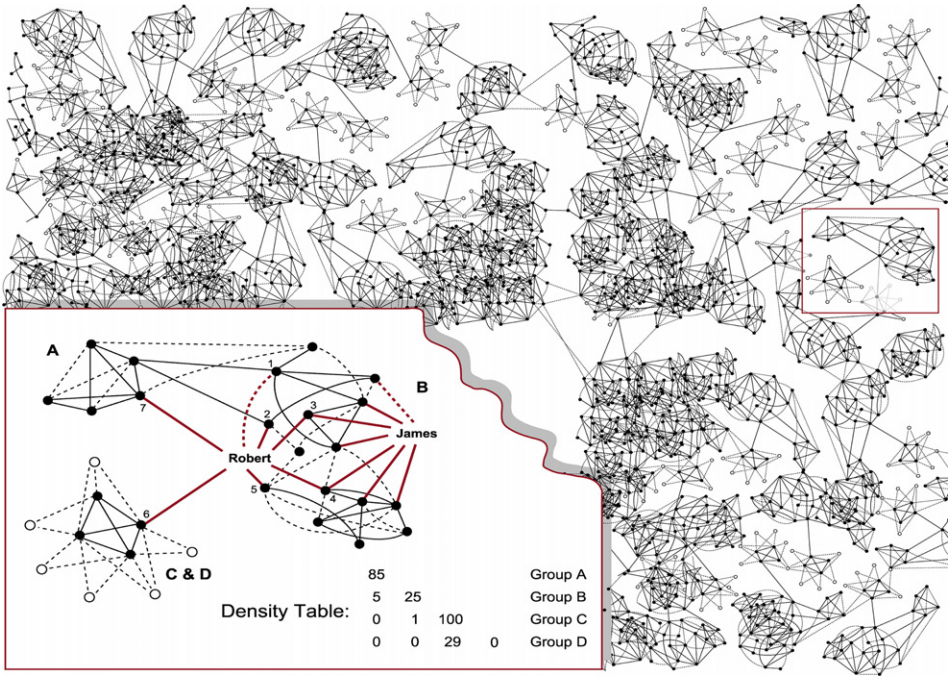

Fig. 1. The small world of organizations and markets.

That twofold task is a generic issue for contemporary organizations and the central issue addressed by the network theory of social capital.[2] Social capital exists where people have an advantage because of their location in social structure. The generic context is a social structure such as Fig. 1 (taken from Burt, 2005, p. 14). The figure is a sociogram in which lines indicate where information flows more routinely, or more clearly, between people or groups represented by dots. Solid lines indicate stronger flow. The defining feature of the social structure is clusters of dense connection linked by occasional bridge relations between clusters. As a point of reference for later discussion, a network segment is enlarged in the overlay box to highlight four clusters. Clusters A, B, and C are variably closed-network groups in the sense that relations are more dense within than beyond the group (density table shows average relations within and between groups). Cluster D (white dots in the figure) is defined by structural equivalence (density table shows that people in cluster D have stronger relations with group C than with one another). Structures of clusters connected by bridges occur in a wide variety of circumstances across levels of aggregation (Watts and Strogatz, 1998). Whether communities in a geographic region, divisions in a corporation, groups within a profession, or people in a team, people specialize within clusters and integrate via bridges across clusters.

[2] For detailed discussion, see the book from which this brief overview is extracted (Burt, 2005). Specifically, Fig. 1 is taken from Burt (2005, p. 14), evidence on returns to brokerage is discussed in Chapter 1 of the book, and evidence on returns to closure is discussed in Chapter 3. The overview is needed because we will be talking later about the importance of experience in seeing structural holes.

Two leadership roles are highlighted, illustrated by Robert and James in Fig. 1. The two roles have long been the subject of work—for example, Schumpeter (1912) on entrepreneurial "leaders" bringing together elements from separate production spheres in which people live by routines, or Merton (1949), and Katz and Lazarsfeld (1955) on the diffusion of tastes through cosmopolitan "opinion leaders" whose relationships bridge the gaps between social worlds (see Burt, 1999), Rees (1966) on "extensive" search for information on job opportunities versus "intensive" search for information about a specific opportunity, Milgram (1967) on the "small world" phenomenon in which people at great geographic remove can communicate with one another through surprisingly few intermediaries because of bridges between social worlds (see Watts, 1999), Granovetter (1973) on the critical role that "weak ties" would play in information access and flow if bridge relations were weak rather than strong, Burt (1982, 1992) on the information access and control advantages created when relations span the "structural holes" between groups, Kotter (1990) on "leaders" motivating people in a new direction versus "managers" planning and budgeting to reduce complexity in the old direction, March (1991) on organizations "exploring" for new opportunities versus "exploiting" known revenue streams, or Padgett and Ansell (1993) on the "robust action" made possible by structural holes between groups. A theme in this work is that behavior, opinion, and information broadly conceived, are more homogeneous within than between groups. People focus on activities inside their own group, which creates holes in the information flow between groups, or more simply, structural holes.

Look at the networks around Robert and James. Each has the same number of contacts. They differ in the way their contacts are connected with one another. The seven people connected to James are densely connected with one another within cluster B. The seven people connected to Robert are not connected with one another and are drawn from separate groups A, B, and C.

Robert is better positioned for the social capital of brokerage. Given greater homogeneity within than between groups, people whose networks bridge the structural holes between groups have earlier access to a broader diversity of information and have experience in translating information across groups. They have a vision advantage in detecting and developing good ideas. They are able to see early, see more broadly, and translate information across groups. Like over-the-horizon radar in an airplane, or an MRI in a medical procedure, brokerage across the structural holes between groups provides a vision of options otherwise unseen. There is abundant and accumulating empirical evidence of the returns to brokerage—in terms of more positive performance evaluations, faster promotions, higher compensation, and more successful teams. Illustrated by Robert in Fig. 1, brokerage is about the growth advantages of introducing variation into group opinion and behavior. The social capital of brokerage is the theoretical foundation for training programs intended to enhance leadership skills in building bridges across organization silos.

James is better positioned for the social capital of closure. Where mutual contacts close the network around bridge relations, reputation pressures encourage the trust and collaboration needed to deliver the value of bridges, creating a social capital advantage defined in terms of closure across structural holes, especially at extreme levels of closure coordinating across extensive bridge relations as in a skunkworks or crisis team. Illustrated by James in Fig. 1, the social capital of closure is about the

efficiency advantages of driving variation out of group behavior or opinion. Where Robert is positioned to benefit from differences between people who vary in their behavior and opinion, James is positioned to integrate the work of people who have much in common. Where James is positioned to drive variation out of group B, Robert is positioned to introduce into group B variation from the other groups A and C with which he is familiar. The social capital of closure is the theoretical foundation for training programs (e.g., Six Sigma and Lean Production) intended to enhance skills in process efficiency.

The two forms of social capital are twice complementary. The greatest returns to informal organization occur when both forms are present, when closure within a group occurs with brokerage beyond the group. Second, brokerage and closure are complements in providing a cure for the other's failure mode. Unrestrained brokerage can create organization chaos, manifest in errors such as resources allocated to conflicting goals and units in the same organization competing against one another. Closure's reputation mechanism brings people back into alignment. It is no accident that reputation metrics provided by multi-point (360) evaluations were adopted in so many large companies at about the same time that companies were removing layers of bureaucracy, which made them more dependent on informal brokerage to integrate operations across groups. On the other hand, closure's reputation mechanism can create groupthink and rigidity. Stories echoing within the closed network amplify opinion to positive and negative extremes, making the existing structure resistant to change, deepening the structural holes that segregate groups, again especially at extreme levels of closure. Brokerage has the potential to crack closure-induced arthritis with selection and synthesis among conflicting alternatives.

3. A Business Leadership Program (BLP)

We created Raytheon's senior leadership development program, titled the "Business Leadership Program" (BLP), to teach director and vice-president executives to see the social capital in how they organize to create value. Relative to a traditional multi-week senior executive program, the BLP is a quick surgical strike that combines academic work with project applications to ongoing company operations. A cohort of two to four dozen executives spends 5 days at the University of Chicago Graduate School of Business. The gist of the program is as follows (details can be obtained from the website in the authors' note to this paper): The program begins with an academic introduction to the vision mechanism by which growth occurs when people build network bridges across the structural holes in an organization, and the reputation mechanism by which trust and efficiency occurs when people close the network around a group. Written and video cases are combined with research evidence to communicate the network mechanisms and highlight their failure modes, the agency problems that can occur with brokerage and the groupthink, stereotyping, and rigidity that can occur with closure. The network mechanisms are then used to discuss traditional questions of business and corporate strategy about market behavior, where to position programs in an organization, the choice between tight or loose couplings between programs, and how to effectively connect with customer markets. Again, written and video cases are combined with research evidence to drive home the content. The fourth and fifth days shift to application and active learning. There is a transition morning on Raytheon strategy in light of

the principles discussed, followed by a day and a half in project teams applying the principles to practical issues in Raytheon. Six weeks later, the cohort assembles for a "mid-course" day at corporate headquarters to discuss project progress and have a two-hour workshop with the CEO on Raytheon strategy and leadership. Six weeks after that, the cohort assembles for a "graduation" day at headquarters. They debrief one another and the CEO on their final proposal, and their progress on a plan to transfer the BLP project into the company for development or implementation. The BLP has shifted with strategic initiatives in Raytheon Company, and morphed in response to the interests and style of a new CEO, but two goals have remained consistent: (1) Improve Director and Vice-President alignment with Company strategy and emerging strategic initiatives. (2) Enhance market-driven collaborative connections among senior people across the Company. It was hoped that fulfilling these goals would enhance the ability of senior leaders to identify and effectively act on strategic opportunities.

4. Preliminary evaluation

Extensive research evidence supports the network theory of social capital. Research reviews, opinions, and practical advice can be found in Baker (2000), Burt (2005), Cohen and Prusak (2001), Cross and Parker (2004), and Lin (2002). The abundant evidence can be cited to justify programs, such as the BLP, designed to improve the social capital skills of executives. However, executive education is difficult to assess. Lower-level training has concrete outcomes, where the training is intended to make a known process more efficient—reduce the number of defective parts, increase the number of units shipped, or reduce the number of customer complaints. The goals of executive education programs are less concrete, less about refining what exists than creating what does not yet exist. A typical goal is to improve the informal network of collaborative relations among senior people so they are better aligned with company strategy and better able to quickly detect and pursue market opportunities consistent with the strategy. Research demonstrating the network mechanisms and substantial consequences of social capital leaves unanswered the practical question of whether executives can be taught to better see the social capital opportunities around them.

Donald Kirkpatrick proposed in 1959 a four-level typology of training evaluations that, as a recent literature review explained (Salas and Cannon-Bowers, 2001), "continues to be the most popular framework for guiding evaluations." Kirkpatrick (1994) offers detailed discussion and application, but the gist of the typology is as follows: reaction (Did participants like the program?), learning (Did participants change their attitude, knowledge, or skills?), transfer (Did participants change their behavior?), and results (Did participants become more productive?).

4.1. Level one: reaction

Level one evaluation refers to the "smilesheets" distributed at the conclusion of training programs. With respect to the BLP, the combination of advanced research, practical application to one's own work, and the opportunity to think through strategy problems with the CEO is well received by participants. Opinion of the faculty is tightly clustered on a five-point scale around a 4.94 mean. Although participants were well along in their careers (48 years of age, on average), most say they had learned the maximum "quite a bit" that would

Fig. 2. Evidence on learning a new organization.

help them be a better leader. An even larger majority would recommend the program "without reservation" to an able colleague.[3]

4.2. Level two and three: learning and transfer

Level two and three evaluations of the BLP refer to cognitive and behavioral change in participants. As illustration, Fig. 2 contains results from a recent report on people learning networks. The results are for very young people—college and day-program M.B.A. students. We do not offer the results as a baseline for executives. The results are illustrative and consistent with our experience.

Structural holes are the foundation for social capital, but they are not easily seen. Whatever the time required to learn how people are connected, there is additional time required to learn where people are not connected. For example, using DeSoto (1960) experimental design for measuring the difficulty of learning a social structure, Freeman (1992) asked college students to learn the relationships in a small network that contained a structural hole. The design consists of presenting flash cards with pairs of people displayed and asking how the people are connected. The learning metric is how many passes through the cards a subject requires before he or she has learned the connections between each pair of people. Errors occur when subjects fail to recall a relationship that existed, but Freeman reports that the most frequent error was to fill in structural holes by saying that disconnected people were connected.

[3] The 4.94 mean opinion of the faculty is in response to the question: "What did you think of the instructors?" (5 very good, 4 good, 3 ok, 2 poor, 1 very poor). Participants give a 4.70 mean response to the question: "Did you learn anything that could help you be a better leader?" (5 yes, quite a bit; 4 yes; 3 not sure; 2 no; 1 no, not a thing). Participants give a 4.90 mean response to the question: "Would you recommend this program to an able colleague who asked whether he or she ought to attend?" (5 yes, without reservation; 4 probably yes; 3 not sure; 2 probably not; 1 no, absolutely not).

Janicik (1998) used DeSoto's design with older students, and added a control for the network around each student in his or her most recent or current job. Students who have more than the median number of holes in their job network are the "familiar with holes" learners in Fig. 2. The other students are "familiar with closure" in that they are at work more surrounded by connected colleagues. The first two bars in Fig. 2 show no difference between the two categories of students in the speed with which they learn a complete and consistent hierarchical structure.[4] Students familiar with closure require an average of 9.0 passes through the flash cards to learn the network. Students familiar with structural holes require 8.4 passes.

Introduce a structural hole so that two people are not connected, but are expected to be connected, and the experience difference is clear. Students familiar with holes learn the new organization about as quickly as they learned the complete and consistent hierarchy, requiring 9.6 passes through the flash cards. In contrast, students surrounded by connected colleagues do not see the structural hole until they have spent almost twice as much time on the task, requiring an average of 17.5 passes through the flash cards. If the college students in Freeman's experiment live in small, dense friendship networks, as is typical of college students, then the summary conclusion from Freeman's and Janicik's experiments is that experience matters: People in a network that contains structural holes are more likely to recognize the holes in a new network (cf. Turnbull, 1961, description of a Pygmy unable to judge the size of distant objects because life in dense forest provided no experience of viewing distant objects).

Even minor training can significantly improve the ability to see structural holes, which is the point illustrated by the last two bars in Fig. 2. Janicik and Larrick (2005, pp. 357–358) report an experiment in which college students were randomly assigned to either of two treatments, one in which subjects discuss networks that contain structural holes, or an alternative treatment in which subjects discuss closed networks in which there are no holes. The shaded bar lower than its adjacent white bar at the end of Fig. 2 shows that just talking about structural holes speeds by 34% the rate at which students learn to see the holes in a new organization.

With these results in mind, one could add to the BLP a pre- and post-program exercise testing the rate at which participants learn the informal organization of a situation. Our decision was to take such knowledge for granted—given the program content on the social capital of structural holes, and our experience with people learning network analysis (consistent with the evidence in Fig. 2). More generally, the BLP is an executive program leveraged on the domain competence of the participants. The desired program outcome is not senior leaders conforming to a prescribed profile of beliefs or behaviors. It is senior leaders growing the business by more ably identifying and effectively acting on strategic opportunities. Opportunity detection, and how participants act upon their opportunities, is up to the vision and drive of participants. Our decision to move past levels two and three, directly to level four, was informed by our experience teaching network analysis. We know

[4] Means in the text are taken from the later-published report, Janicik and Larrick (2005, pp. 352–353), which on the same pages also displays the complete, consistent hierarchy and the network containing a structural hole. The differences in Fig. 2 in rates of learning networks containing a structural hole are statistically significant, and the binary distinction between high and low familiarity with structural holes at work captures for the students the effect of a continuous measure of network closure (Janicik and Larrick, 2005, p. 352n).

that people can be taught to see network structure more accurately. Given that ability, the BLP focuses on the value produced.[5]

4.3. Level four: productivity

We use Raytheon Company's own evaluation metrics to measure program outcomes. Assuming that the annual performance evaluations reflect the quality of an employee's work, is the BLP associated with more positive evaluations? Assuming that people are promoted for merit, are BLP graduates more likely to be promoted? Assuming that people stay with the Company if they find engaging work, does the program lower the odds of a promising executive leaving the Company?

Using Company evaluation metrics is attractive, but highlights a selection bias: Executives were invited to the BLP because they appeared promising according to company metrics. Consider Fig. 3. The Company's annual performance reviews assign employees to three performance categories: Meets, Exceeds, and Far Exceeds Expectations.[6] In Fig. 3, program graduates are significantly more likely to far exceed expectations (41% of BLP graduates versus 17% of same-rank people not through the BLP, 9.38 z-score test statistic). Program graduates are significantly less likely to simply meet expectations (12% of BLP graduates versus 37% of those not through the BLP, −7.81 z-score).

It is tempting to conclude that executive performance is higher for BLP graduates, but people who were doing well were at higher risk of being invited to the BLP. This is a generic problem for evaluation: The people graduating from training programs are not a random sample of the eligible population. For example, Burt et al. (2000, pp. 139–140) report

[5] Nevertheless, in October of 2005, which is 4 years after the first BLP cohort and 4 months after the most recent, the 311 program graduates still with the company were sent an email asking about their BLP project and whether they had developed consequential relationships from the BLP. Responses came back from 59% of the graduates. With the exception of vice presidents being less likely to respond (38% of vice presidents responded versus 58% of directors and other ranks), the probability of response is uncorrelated with the risk of being invited to the BLP, the year or period in which a person attended the BLP, promotion last year, the graduate's business unit, business function, age, relative salary, last performance evaluation, having been an active participant in the BLP, or having been on a BLP team whose members believe their project was a success. Most graduates felt that the BLP had resulted in consequential relationships (90%). Beyond the "yes" responses, illustrative comments include the following: "Many relationships were established. In particular, the close relationships with my project teammates as well as the other class members. The project also gave us the opportunity to meet and talk to many other Raytheon employees I would have never met as part of my normal assignments." "It was a great group of people and it is a pleasure to continue those relationships four years later." "Not just my project team members, but others in the class. This has proven to be the most beneficial aspect of the class from my perspective. If I get a call for assistance from a member of my class, it is a high priority for me. I feel my calls for support are also treated as a high priority from class members." "It broadened my cross-functional network across Raytheon." "Yes absolutely! Both in my business and with people in other businesses around the company." And some of the few people who answered "no" seemed to be saying "yes" in a round-about way with illustrative comments such as the following: "No, but I have contacted members of the BLP when I had similar issues arise in my business to projects that were studied in the cohort." "I knew many in my BLP course. The course DID help me solidify and enhance some of those relationships." The comments show that graduates focus on relations with people in their BLP project team, then people with whom they attended the BLP (giving perhaps too little attention to relations across cohorts and into the broader organization), but the comments are consistent in attributing ongoing consequential relationships to attending the BLP.

[6] There is a fourth performance category, Needs Improvement, that lies below Meets Expectations. There are very few Needs Improvement evaluations (two in one year and two in another for the job ranks studied here), and the Needs Improvement evaluations that occur have the same correlates as Meets Expectations, so we combine the two negative evaluations into one category.

that participation in a company executive program is the only factor predicting network brokerage in a large French company. Managers who attended the program are more likely to have bridge relations across groups in the company (like Robert in Fig. 1), which is associated with higher salary relative to peers, but did the bridge relations precede, or result from, the executive program? It is reasonable to assume that managers involved in bridge relations, and therefore doing well in the company, were more likely to be selected for the executive program. Returning to Raytheon, job rank is a generic selection bias illustrated in Fig. 3. Of eligible employees, those in more senior job ranks are more likely to be invited to the BLP. The table in Fig. 3 shows Raytheon's top five job ranks: Vice President, Senior Director, Director, Senior Manager, and Manager. BLP participants are typically Vice Presidents and Directors (14.6%, 26.9%, 40.9%) with a few "high-potential" participants from the lower two ranks (11.8% and 2.8%). In contrast, the population of eligible employees is concentrated in the bottom two ranks (27.1% and 60.3%). In other words, job rank is a risk factor: higher-rank people are more likely to be invited to the BLP.

5. Control group

To estimate program effects on performance, we need to hold constant selection bias so program graduates can be compared to peers at equal risk of being invited to the program. We need a control group—a set of people not attending the BLP but otherwise similar to the people who did attend. This is a well-studied task of creating "propensity scores" that define for individuals the conditional probability of treatment assignment (Rubin, 1973; Rosenbaum and Rubin, 1985; Rubin and Thomas, 1996). The goal is to compare treated people with untreated people who had similar a priori risk of treatment.

5.1. Time periods

Time is one of the factors to hold constant. Raytheon evolved in ways that could affect program effects. Fig. 4 displays data on the 11 BLP cohorts (a group of people

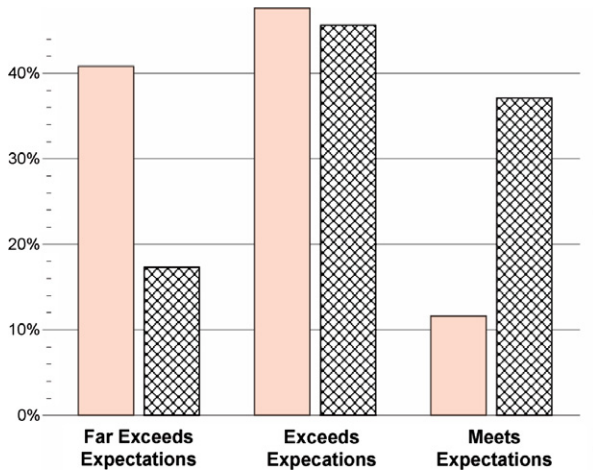

	BLP Graduates	Other Employees in Same Ranks
VP	14.6%	1.4%
Sr. Director	29.9%	2.3%
Director	40.9%	8.9%
Sr. Manager	11.8%	27.1%
Manager	2.8%	60.3%
(total)	(100%)	(100%)

Fig. 3. BLP graduates tend to be outstanding employees.

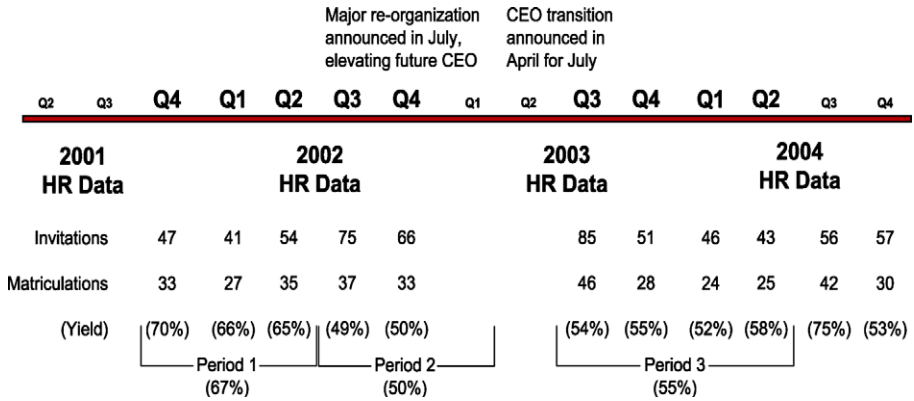

Fig. 4. BLP Time Line. A large Q indicates quarter containing a BLP cohort included in this analysis. The nine cohorts involved 508 invitations to 361 employees, yielding 288 participants.

taking the program together) from the time that the program was launched in the fourth quarter of 2001 through the end of 2004. Our data on employee job rank, division, geographic location, supervisor, and so on are updated for September of each year (indicated by "HR Data" in Fig. 4), after the annual performance evaluations. We estimate program effects for the first nine cohorts (distinguished by large Q quarters in Fig. 4), tracking employees through the 2004 evaluation cycle.

Three periods are distinguished in Fig. 4: one quiet, one turbulent, and a period of settling down. The first three cohorts go through the BLP under a continuing CEO and a continuing organization structure. These are the cohorts in which invitations are most often accepted.[7]

Enter a touch of turbulence. The 2003 Q3 and Q4 cohorts go through the BLP just after a major re-organization is launched that elevates the next CEO, Bill Swanson, to a position of President. In the following April, an orderly succession is announced with CEO Dan Burnham moving to Chairman when Bill Swanson becomes the new CEO at the end of June. Senior positions removed and re-defined through this period make the end of 2002 and beginning of 2003 a stressful period for the eligible population. Anticipating employee reservations about being away from the office during this period, the 2002 Q3 BLP is delayed and 50% more invitations are sent out. The yield on invitations is the lowest of all cohorts (−1.98 z-score test statistic). Following the June announcement of the new CEO, BLPs in 2003 are delayed until the third quarter. Anticipating residue employee concerns, a large number of invitations are sent out, but concerns turned out to have subsided judging from the many acceptances (46 in the Q3 cohort).

The third period, following the summer of 2003, involves settling down in the new organization structure. The BLP becomes an element in routine business operations, with

[7] There are no statistically significant differences between the first nine cohorts in the tendency for BLP invitations to be accepted (9.92 χ^2, 8 df, $P = .27$), but the acceptance rate is significantly higher for the first three cohorts as a group (2.87 z-score test statistic).

invitations tied to the Company's annual performance reviews. The program graduation project briefings become a forum in which the new CEO shares intuitions and arguments about the Company with his senior people.

5.2. Risk of invitation

Table 1 is an inventory of factors predicting invitation to the BLP. The two columns of effect test statistics in Table 1 come from two logit models, one predicting who was invited to the BLP and the second predicting who accepts their invitation. Effects are estimated across three panels of data corresponding to the three periods in Fig. 4: The first panel uses data from the 2001 HR review on people in the top five job ranks (excluding direct reports to the CEO because they are not eligible for the BLP) to predict invitation to any of the first three BLPs. Risk factors are recorded before the period in which risk is estimated. For example, job rank for the first period is the rank an employee held before invitations went out for the first BLP. The second panel uses data from the 2002 HR review (excluding direct reports to the CEO and anyone who already graduated from the program) to predict invitation to either of the two BLPs in the second period. The third panel uses data from the 2003 HR review (excluding direct reports to the CEO and anyone who already graduated from the program) to predict invitation to any of the four BLPs in the third period. Across the three periods, there are a total of 27,884 observations for the estimation in Table 1.[8]

The "Job Rank" rows in Table 1 corroborate Fig. 3: Vice Presidents and Senior Directors are likely to be invited to the BLP (7.18 and 8.87 z-scores). People in the two Manager ranks are unlikely to be invited (-10.90 and -12.22 z-scores).

Some variables are noteworthy because they are not risk factors. To improve integration across functions, for example, people from all kinds of work are invited. Differences between kinds of work in Table 1 are negligible with one exception: The 359 people who run key programs are targeted (3.20 z-score for Lead Program Managers).[9]

Certain demographic categories are at higher risk. Younger people are targeted for development (-2.99 z-score means younger age is at higher risk of invitation). Asians and African-Americans are at higher risk of invitation (5.58 and 2.30 z-scores).

To improve the Company's chances of getting a return on employee time in the program, people doing well are at greater risk of invitation. Recalling the performance categories in Fig. 3, employees invited to the BLP are likely to far exceed expectations

[8] Three observations of 8552 employees at risk of invitation in all 3 years, two observations of 640 employees at risk of invitation in two of the three years, and one observation of 945 employees at risk of invitation in one of the three years.

[9] Ten kinds of work are distinguished in Table 1 based on the department to which each employee is assigned: business development, three kinds of engineering, production, two kinds of program management, finance, human resources, and other administrative work. "Lead Program Managers" are the 359 people in charge of the company's most significant programs. "Lead Engineers" are the 150 people in charge of primary engineering organizations. These two kinds of key people were identified by the head of Raytheon's Engineering Institute, Charlie Case, using organization charts, his decades with the company, and his network of engineering colleagues across the company.

(5.93 z-score) and unlikely to simply meet expectations (−2.46 z-score). People with graduate degrees are slightly more likely to be invited (2.31 z-score). People with salaries above their peers are more likely to be invited (3.55 z-score).[10] People in the Company-sponsored Six Sigma efficiency programs are more likely to be invited. The lowest level of expertise—attending a short course and becoming a Specialist—does not trigger an invitation (1.05 z-score). The higher-level expertise involving several months of course work and project certification does (3.91 z-score).[11]

More than any other factor in Table 1, doing well in the Strategic Leadership Program (SLP) is a trigger for invitation (13.77 z-score). Launched a year before the BLP, the SLP is a 2-day mobile executive-education program at the Company's local sites. The SLP introduced employees through video, lecture, and cases to the social capital of brokerage, closure, and reputation to help middle managers think strategically about their role in growing the business. Three or four SLPs occur each month, reaching slightly more than a thousand employees a year. After each SLP, participants are sorted by the instructor into three categories (active leader, active participant, or observer), which is source of nominees for the BLP. The SLP performance data are scored 3, 2, 1 for Table 1, with a 0 given to employees who had not attended an SLP. SLPs are sponsored, but not mandated, by the Company. SLPs occur where local leaders request them. In other words, employees are not at equal risk of participating in an SLP. Strong SLP performance in Table 1 indicates two things: an able employee, *and* a local leader who supported training to improve employee strategic thinking. Both factors affect invitation to the BLP since local business leaders are consulted about nominations to the BLP.

The last three factors predicting invitations in Table 1 adjust for an administrative shift in the risk function. When we predict invitations with a separate logit model for each of the three periods, Table 1 effects recur in each period—with one exception: In the third period, after invitations to the BLP had been integrated into the annual performance review process, the BLP is used as a socialization mechanism so recent hires are more likely to be invited. In the first 2 years of the program, during 2001 and 2002, two percent of the people just hired into senior rank were invited to the BLP. In the third period, the percentage increases to 12 percent. In Table 1, new hires are not at higher risk of invitation to the BLP before 2003 (−.05 z-score). In the third period,

[10] Annual salaries for all 27,884 employees in Table 1 are regressed across four job-rank variables (job rank and job-rank squared for engineers separate from non-engineers), age, age in years past 55 (based on an inspection of residuals from age alone), the two education variables in Table 1 (bachelor and graduate degree), a time adjustment for an employee's salary increasing from 2001 to 2002, and another adjustment for the increase in 2003 (.69 squared multiple correlation). With adjustment for autocorrelation within an employee across the three panels, all of the effects are statistically significant at well beyond a .001 level of confidence, except the negligible difference between salaries in 2001 and 2002. Relative salary in Table 1 is the studentized residual from the prediction. Scores measure the extent to which an employee's salary is above his or her peers defined by job rank, age, and education.

[11] Of employees in the top five job ranks, 3,417 are Six Sigma Specialists (of whom 4.2% were invited to the BLP), 269 graduated from the several-month course work to become an Expert (of whom 6.7% were invited to the BLP), and 226 went all the way to become a certified Expert by completing a suitable Six Sigma project (of whom 10.2% were invited to the BLP). The Expert variable in Table 1 is scored 2 for a certified Expert, 1 for an un-certified Expert, and 0 for everyone else.

invitations are likely for new hires (4.11 z-score).[12] The shift to new hires carries no evaluative element. It is not good or bad. It is simply another risk factor to take into account in defining a control group. The control group for a BLP cohort in the third period should contain more new hires than the control group for cohort in the first two periods.

5.3. Risk of matriculation

The BLP is not a mandatory program. Matriculation is a variable yield on the number of people invited (tabulated in Fig. 4). Many invitations are re-invitations to people unable to accept an earlier invitation, typically because of a critical customer or leadership-team meeting back home. Of 361 people invited to the BLP by the middle of 2004, 259 had been invited once (of whom 85% attend), 76 had been invited twice (of whom 70% attend), 16 had been invited three times (of whom 69% attend), and 10 people had been invited four or more times (of whom 50% attend).

Three points are established by the results in the second column of Table 1. All but one of the effects is statistically negligible. In other words, there is no systematic tendency for people from any rank, region, kind of work, or personal background to accept or reject an invitation to the BLP. Attendance is random after invitation, so the risk factors for attending a BLP are whatever factors define the risk of being invited.

The one exception is performance in the Strategic Leadership Program (SLP). People who had done well in the SLP are more likely to accept their invitation to the BLP. This highlights the point in the previous section about local leaders influencing invitations to the SLP. Where local leaders bring the SLP to employees, employees are more likely to accept an invitation to the BLP.

Third, repeated invitations do not affect the odds of attendance (−1.49 z-score in Table 1). Ignoring other factors, people invited multiple times are less likely to accept their invitation (−3.52 logit z-score). But people in higher rank have more difficulty getting away for the program. The average Manager in the BLP received one invitation. The average Vice President had received two invitations. The number of invitations to an employee increased systematically with the employee's job rank (4.13 t-test). Job rank is not to be confused with performance. People declining an invitation tend to be in higher-rank jobs, but performance evaluations of them are neither higher (1.33 z-score) nor lower (−.64 z-score) than people who attended when first invited. When job rank and the other factors in Table 1 are held constant, attendance is as likely for a first-time invitation as it is for a second- or third-time invitation (−1.51 z-score).

[12] The "new hire" effect is gone within 3 years. The variable "New Hire" in Table 1 equals 3 for someone hired in the current year, 2 if hired last year, 1 if hired 2 years ago, and 0 if hired three or more years ago. This variable is updated over time so a person hired in 2000 has a score of 2 on "New Hire" in the 2001 panel, 1 on the variable in the 2002 panel, and 0 on the variable in the 2003 panel. "Period 3" in Table 1 equals 1 for observations at risk of invitation in the third period, 2003 and after. "New Hire in Period 3" is the product of "New Hire" and the dummy variable "Period 3." The effect of "New Hire" is the effect of being new during the 2 years preceding the third period (when "Period 3" equals 0), and the effect of "New Hire in Period 3" is the effect of being new during the third period.

Table 1
Risk factors for invitation and attendance

	z-Score invitation	*z*-Score attend
Job rank		
Vice president	7.18	−1.38
Senior director	8.87	−.92
Director	—	—
Senior manager	−10.90	.53
Manager	−12.22	—
Kind of work		
Engineering job category	.63	.87
Business development	1.90	−.09
Lead engineer	−.02	1.62
Engineering fellow	.27	−1.49
Other engineer	−.70	−.22
Production	1.55	.59
Lead program manager	3.20	.38
Other program manager	1.17	.13
Finance	.67	.07
Human resources	—	—
Other administrative work	1.79	.15
Geography		
Arizona	1.72	−.66
California	−.82	.25
New England	—	—
Texas	.53	1.81
Other location	1.76	−.13
Personal attributes		
Age (in years)	−2.99	−.78
Age past 55 (years)	−3.36	−.13
Woman	1.82	.88
African-American	5.58	.85
Asian	2.30	−1.08
Hispanic	.78	.20
Bachelor's degree	−.43	−.79
Graduate degree	2.31	.25
Personal performance		
Meets expectations	−2.46	−1.76
Exceeds expectations	—	—
Far exceeds expectations	5.93	−1.66
Relative salary	3.55	−.21
Six Sigma specialist	1.05	.86
Six Sigma expert	3.91	1.97
Performance in the SLP	13.77	3.89
New hire (years)	−.05	−.58
Period 3 (after mid-2003)	6.81	1.78
New hire in period 3	4.11	1.43
Multiple invitations to BLP	—	−1.51

Note: These are *z*-score test statistics for two logit models, one predicting who was invited and the second predicting who accepts their invitation. Estimation is across three annual panels of data on employees in the top five job ranks (27,884 observations, adjusted for autocorrelation within employee across panels). Reference categories are indicated by "—" in place of a statistic. The "invitation" column lists *z*-scores for the row factor predicting invitation to the BLP (1144.41 χ^2, 35 df, $P \ll .001$). The "attend" column lists *z*-scores for the row factor predicting who among the invited people attended the BLP (74.65 χ^2, 35 df, $P < .001$). The job rank of "Manager" is not included in the matriculation prediction because everyone invited from that rank accepted his or her first invitation.

5.4. Control group peers

For each employee not yet attending the BLP, each year, we compute from the first logit equation in Table 1 the probability that the employee will be invited to the BLP. This is the employee's risk of being invited. The distribution of risk scores is displayed in Table 2. Eight people in the first period had a .85 or higher probability of being invited to the BLP. Of the eight people, seven were invited and four attended. Across the rows in Table 2, program

Table 2
Mean risk scores far participants and controls

Risk of invitation	People at risk	Attended BLP	Mean risk score	Control group		
				Invited, did not attend	Random others	Mean risk score
Period 1						
.9	8	4	.924	3	1	.913
.8	8	3	.836	2	3	.835
.7	12	8	.732	2	2	.772
.6	13	8	.648	1	4	.663
.5	35	15	.531	2	18	.536
.4	25	4	.466	2	6	.461
.3	60	11	.328	4	18	.331
.2	107	17	.248	7	27	.235
.1	280	9	.155	5	13	.156
.0	8744	16	.032	3	29	.046
Total	9292	95	.387	31	121	.317
Period 2						
.9	4	3	.926	0	1	.921
.8	5	2	.836	1	2	.852
.7	3	2	.749	1	0	.774
.6	4	3	.645	0	1	.671
.5	22	6	.526	4	8	.532
.4	20	5	.455	2	8	.453
.3	47	13	.334	5	21	.337
.2	82	10	.248	5	15	.247
.1	253	7	.158	7	7	.157
.0	8737	19	.029	10	28	.042
Total	9177	70	.311	35	91	.265
Period 3						
.9	11	7	.930	0	4	.935
.8	11	6	.841	2	3	.857
.7	17	8	.740	2	7	.739
.6	22	6	.637	1	11	.647
.5	39	9	.535	3	15	.543
.4	58	13	.428	4	22	.445
.3	79	14	.343	9	19	.348
.2	143	20	.235	10	30	.236
.1	362	24	.138	12	36	.139
.0	8673	16	.048	8	24	.047
Total	9415	123	.368	51	171	.320

participants are clearly concentrated at higher levels of risk, but there are also people drawn from low-risk levels reflecting special interests and initiatives in the Company.

The control group is a set of non-BLP employees matched to BLP participants on their risk of being invited to the program for each period. The control group includes two kinds of people, distinguished by column in Table 2. First, the control group for a period contains everyone who was invited during the period but did not attend. Regardless of their background, these invited non-attendees were people at high risk of attending the program in that they had already cleared the nomination process. Additionally, the control group for a period contains a stratified sample of uninvited people. Uninvited people are drawn in sufficient numbers to have two control-group employees for each program participant at each level of risk in Table 2—up to the limit of available people in the risk level (e.g., there are not sufficient people in the .9 risk level during any of the periods to double the number of program participants). Dividing risk into 10 levels as sampling strata, versus nine or eleven levels, is an arbitrary decision intended to ensure representation in the control group of the risk distribution in the population.[13] Having a control group twice the number of participants is an arbitrary decision intended to provide multiple control observations to offset random sampling for the control groups. Table 2 shows similar average risk scores for participants and controls at each level of risk within each period (.009 average row difference across the 30 risk-level rows in Table 2). The larger employee population available to provide controls at lower risk levels means that employees in the control group on average have a slightly lower risk of invitation (e.g., .317 for non-attending controls in the first period versus .387 for attendees). We hold constant employee risk-of-invitation in estimating program effects.

6. Program effects on current operations

Fig. 5 displays three program effects on Company operations at the end of the time line in Fig. 4. The three program effects are higher performance evaluations, more likely

[13] Drawing additional controls at random from uninvited employees who occupied job ranks at risk of invitation would produce too many low-risk controls since there were many more employees low on the risk factors in Table 1. We stratified the population to ensure representation of high- and low-risk employees. The 10 categories in Table 2 are sufficient to capture risk differences significant for invitation to the program. We make that statement based on the following test: Regress the continuous measure of risk over two variables, a dummy variable distinguishing invited employees, and a 10-interval variable distinguishing the 10 risk levels in Table 2. Across the 27,884 observations in Table 1, invited people are at higher risk of invitation than the uninvited (8.27 *t*-test for the higher risk for invited people, 243.88 *t*-test for the 10-level control; with standard errors adjusted for autocorrelation within an employee across the three panels). The fact that invitations have a statistically significant association with risk despite our control for the 10 risk levels in Table 2 tells us that the strata, on average, are too broad to treat the risk of invitation as random within strata. However, the problem is concentrated at the lowest level of risk, among employees with less than a .05 probability of invitation. For the 26,154 observations at the lowest risk level in Table 2, risk scores are significantly higher for invited employees (14.28 *t*-test). In contrast, invitations are independent of risk within the nine non-zero risk levels (.59 *t*-test for invited people having a higher risk of invitation, 213.54 for the risk-level control variable). In other words, the 10 risk levels in Table 2 are sufficient as sampling strata to capture the risk variation significantly associated with invitation to the program—except among the most unlikely participants. So, we draw people at random to fill the control group for each period except in the lowest risk level, where people are drawn in descending order of risk (highest risk person first, next highest second, and so on). This decision does not affect the most likely program participants because, as shown in Table 2, few people are available at high levels of risk so everyone is either a participant or in the control group.

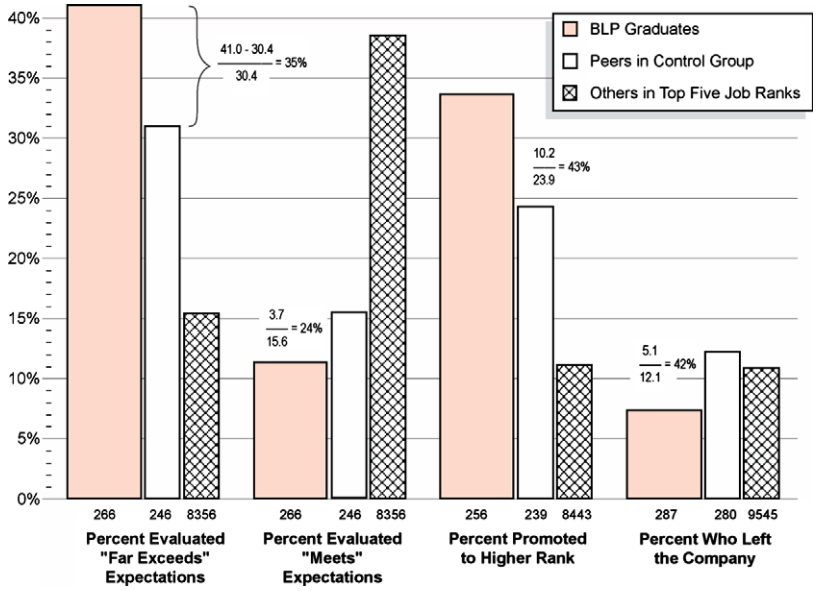

Fig. 5. BLP effects on current senior people.

promotion, and improved retention. Effects are aggregated over time in the sense that we compare, at the end of 2004, BLP graduates, peers in the control group, and other employees in the senior ranks.

The first bars in Fig. 5 show the BLP effect on performance evaluations. The BLP graduates (shaded bars) and the controls (white bars) are more likely than the general population (cross-hatch bars) to far exceed expectations (respectively, 41% and 30% versus 15%). At the other extreme, BLP graduates and the controls are less likely to receive the minimum evaluation of meeting expectations (12% and 16% versus 39% for the cross-hatch bar). BLP graduates and their peers in the control group are clearly outstanding within the broader employee population in the top five job ranks, but BLP graduates are significantly ahead of their peers. With respect to far exceeding expectations, BLP graduates are a 35% improvement over their peers in the control group.[14]

The third panel in Fig. 5 shows people moving to leadership positions after the company reorganization late in 2002. Bars indicate the percentage of employees promoted to a higher job rank in 2003 or 2004. Given the time order, this is a program correlate rather than a program effect, but the association shows the relative tendency for people in the three categories to be sought for leadership positions. We look at time-dependent effects below. Here, the BLP graduates and their peers in the control group again stand out from the broader employee population: 34% of BLP graduates are promoted and 24% of their peers in the control group are promoted—versus 11% of other people in the top five job

[14] The performance-evaluation differences are statistically significant between BLP, control group, and others in the same ranks (220.35 χ^2, 4 df, $P < .001$). Putting aside the broader population, differences between the BLP graduates and controls are significant (6.59 χ^2, 2 df, $P < .05$), and continue to be statistically significant in an ordinal logit model holding constant the summary risk measure in Table 1 (2.21 z-score BLP graduates, $P < .05$).

ranks. Again, BLP graduates are significantly ahead of their peers in the control group, with 43% higher odds of promotion.[15]

The fourth panel in Fig. 5 shows the percentage of employees who decided to leave the Company before the end of the time period. The departures could be for work elsewhere or retirement, either being a decision to leave the Company for the promise of other activity. About one in ten people leave the Company, but the odds are much lower for BLP graduates. The BLP effect on executive retention is substantial (42%) and statistically significant.[16]

7. The importance of active participation

Beginning with the first session in 2001, the faculty sorted program graduates into five categories of participation: stars (led others in successfully engaging the program material), solid (the person's questions and opinions demonstrated comprehension of the program material), good (able but undistinguished participation), weak (the person's questions and opinions demonstrated that they had trouble thinking strategically), and mute (the person neither asked questions nor offered opinion). By the middle of 2004, the 288 program graduates are distributed across the categories as follows: 16% stars, 13% solid, 28% good, 26% weak, and 17% mute.

The five categories reduce to two when we analyze participation categories with respect to post-program evaluations, promotions, and departures. Stars and solid participants are similar to one another in doing well after the program. These are the "active" participants in Fig. 6 (dark shaded bars). Good, weak, and mute participants are similar in showing minor program benefits. These are the "other" participants in Fig. 6 (light shaded bars).

The tabulations in Fig. 6 are the same as in Fig. 5, except BLP graduates are now distinguished for their "active" or "other" participation in the program. White bars again describe the control group. Cross-hatch bars again describe people in the top five job ranks other than the controls and BLP graduates.

Program benefits are concentrated in the "active" participants. This point has two parts. First, the "active" participants get more benefit from the program. They are much less likely to receive a "meets" evaluation. They are much more likely to be promoted, and much less likely to leave the Company. Second, the less active participants show little more program benefit than the control group, people who never attended the program. This is apparent in Fig. 6 from the light-shaded bars being closer in height to the white bars (control group) than the dark-shaded bars ("active" participants).

Top performance evaluations least fit the pattern. In the first cluster of Fig. 6 bars, "active" participants are slightly more likely than the "other" participants to far exceed expectations. All graduates of the BLP, however, are likely to receive a "far exceeds"

[15] Differences between the odds of promotion are statistically significant (157.96 χ^2, 2 df, $P < .001$). The higher probability of promotion for BLP graduates relative to peers in the control group is statistically significant (6.35 χ^2, 1 df, $P < .05$), and continues to be significant in a logit model holding constant the summary measure in Table 1 aggregating the risk factors predicting invitation to the BLP (2.16 z-score, $P < .05$).

[16] The logit z-score test statistic for the lower probability of a BLP graduate leaving the Company is -2.01 ($P < .05$), which is about the same (-2.10 z-score) if we add a control for the summary measure in Table 1 aggregating the risk factors predicting invitation to the BLP. There is no statistically significant difference between exit rates for the control group versus other non-BLP people in the top five ranks (.71 z-score test statistic).

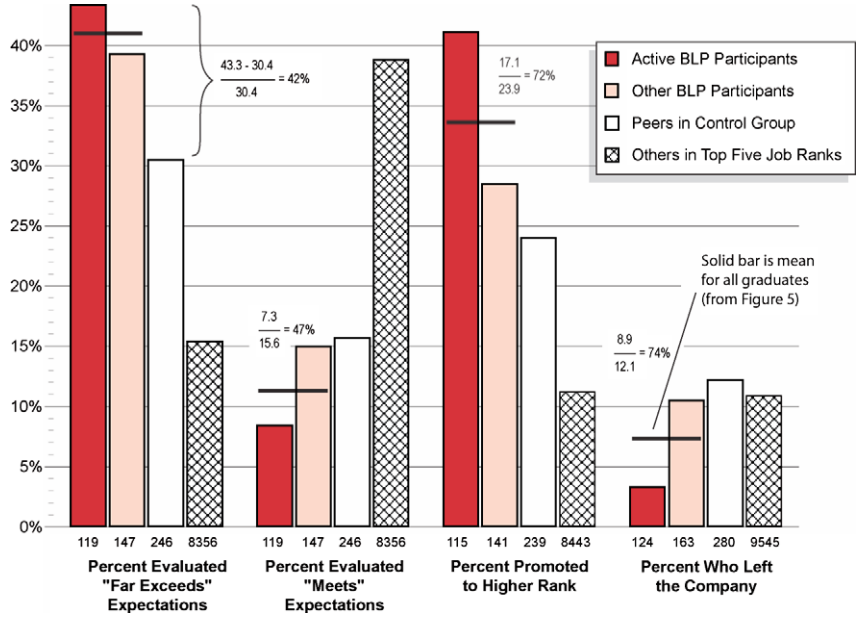

Fig. 6. BLP effects on current senior people, distinguishing active participants.

evaluation relative to the control group. We attribute this effect to the signal value of attending the Company's senior leadership program. It is not the invitation. People invited who did not attend the program were no more likely to receive a "far exceeds" evaluation than the other people in the control group (.40 χ^2, 1 df, $P \sim .53$). The signal value comes from going away to attend the program. Whether the person was active or passive in the program, participants were more likely to receive a high job evaluation when they returned to their job.

Having flagged the evidence of signal, however, we emphasize the main point illustrated in Fig. 6: Attendance alone did not protect program participants from poor evaluations, or increase their chances of promotion, or decrease their odds of leaving the firm; all of which were contingent on "active" participation in the program.

Active participation is also correlated with benefit from the program project. The program involves participants working for 12 weeks on a team project presented to the CEO. Graduates were asked in October of 2005 to describe what became of their BLP project (see footnote 5). They had three response options: (1) The project was implemented with beneficial result, (2) the project fizzled in the sense that the team met after the program but lost momentum or implementation did not yield beneficial results, or (3) nothing concrete happened in the sense that there was no concrete action or the respondent had no post-program contact with the team.

Fig. 7 displays consistency and contradiction between individual opinion and the average opinion of the individual's project team. Most people on a team reported the same outcome: their project was implemented, it fizzled after a while, or it never got off the ground. These are the 132 people in the middle of Fig. 7.

There were 31 people, to the left in Fig. 7, whose reports were more negative than their teammates. For example, the individual reported that the project never got off the ground

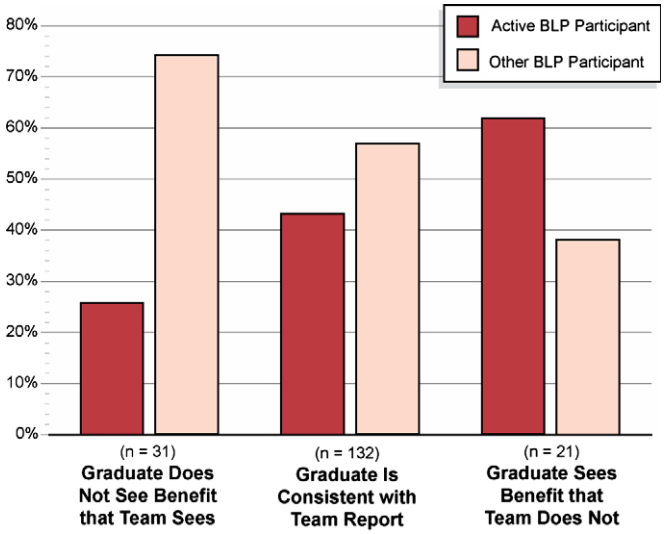

Fig. 7. Individual and team opinion of BLP projects.

while the other team members reported that the project was implemented with beneficial result. The light-shaded bar to the left in Fig. 7 shows that 74% of these people who did not see the project value that their team reported had been passive participants in the program. Respondent comments are informative here. After giving an opinion on the project outcome, respondents were asked to explain. Most people in the left pair of bars in Fig. 7 did not explain their opinion, but those who did speak of losing contact with the team: "I had no contact with the project after the BLP, nor do I know what happened." "I do not know what happened to the project." "Other people had a vested interest in the project, so they know what happened." "Actions on the project were carried on by individual members of the team, but not the team."

There were 21 people, to the right in Fig. 7, whose reports were more positive than their teammates. For example, the individual reported that his or her project was implemented with beneficial result while the other team members reported that the project never got off the ground. The dark-shaded bar to the right in Fig. 7 shows that 62% of these people who saw project value that their team did not had been active participants in the program. The shift in Fig. 7 is statistically significant from passive participants offering more negative reports than teammates to active participants offering more positive reports than teammates, but here again, personal explanations are useful.[17] People reporting project benefits despite negative reports from their team speak of implementing something on their own that proved to be beneficial: "The techniques addressed in the project contributed to Raytheon winning program (name withheld)." "The project work helped my operations

[17] In an ordinal logit model predicting the three horizontal categories in Fig. 7 from the distinction between active versus other program participants, active participants are more likely to report project value despite negative teammate reports and less likely to report project failure when teammates are positive (2.62 z-score, $P < .01$). The association remains after holding constant job rank, business function, business unit, the summary risk-of-invitation variable in Table 1, age, relative salary, job performance, promotion, the year or period of attending the BLP, and the average team opinion of the BLP project outcome.

council establish significant utilization goals by business." "The project resulted in some enterprise campaign funding and an ongoing effort." These positive comments do not contradict a negative team outcome. Rather, the comments reveal an individual who found an opportunity outside the team to generate benefit from the team project.

The importance of active participation raises the question of individual differences. Active participation could be a proxy for some unmeasured quality that distinguishes superior employees. If we knew what made a program participant "active," and could measure "active" in the control group, it could turn out that "active" people in the control group are just as likely as "active" program graduates to do well in the company. There is evidence to support the possibility. Recall the summary risk-of-invitation variable in Table 1 predicting the probability of an employee being invited to the program. The employees more at risk of invitation were in more senior job ranks, better paid than their peers, receiving higher job evaluations, et cetera. It turns out that "active" participants were at higher risk of invitation to the program. The average risk-of-invitation score for "active" participants is .37 versus .21 for the "other" participants (5.56 t-test, $P < .001$). In other words, individual differences before the program are correlated with differences in activity during the program.

However, when we hold constant the Table 1 individual differences captured by the summary risk-of-invitation measure, the effects visually apparent in Fig. 6 remain statistically significant. An ordinal logit model predicting the three levels of performance evaluation, with the summary measure held constant, shows that "active" participants received significantly higher performance evaluations than their peers in the control group (2.35 z-score, $P < .05$) while "other" participants are indistinguishable from peers in the control group (1.44 z-score). A logit model predicting promotion shows that "active" participants are more likely to be promoted than peers in the control group (2.89 z-score, $P < .05$) while the "other" participants are indistinguishable from the control group (.95 z-score). A logit model predicting employee departures shows that "active" participants are significantly less likely to leave than peers in the control group (−2.60 z-score, $P < .05$) while "other" participants are indistinguishable from the control group (−.59 z-score).

8. Time-dependent program effect on promotion

How long before program effects are apparent? Fig. 8 shows promotions by when they occurred, before and after the program.[18] Test statistics at the bottom of Fig. 8 show that the promotion-chances of the program graduates are indistinguishable from the control group in the year before attending program, and during the year they attend the program.

[18] We aggregate across data panels to create the time line in Fig. 8. The three periods of BLPs distinguished in Fig. 4 precede and follow some number of data panels. The BLPs run in the first period are grounded in the 2001 data panel and precede three panels, one in 2002, 2003, then 2004. The BLPs run in the third period follow the 2001 and 2002 data panels, are grounded in the 2003 data panel, and precede the 2004 data panel. We have one observation on each person in a BLP cohort and observations on peers drawn for the BLP cohort's control group. For example, Table 2 shows 95 graduates from the BLP during the first period and a control group of 152 peers not through the BLP. The time line in Fig. 8 was created by expressing the data on each cohort relative to the year they went through the program (or the year they were drawn for a control group). Period-one participants who were promoted in 2003 from their job rank in 2002, for example, were promoted 2 years after the program. They are tabulated in the "$t + 2$" column of Fig. 8. With respect to the volume and heterogeneity of data, we have the greatest statistical power at the center of the distribution.

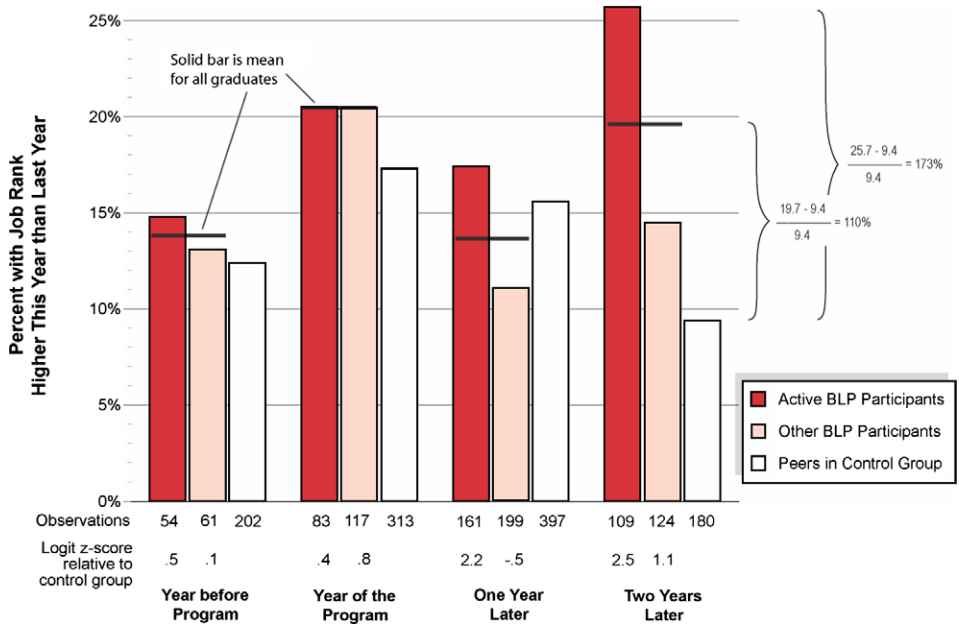

Fig. 8. Time-dependent BLP effect on promotion.

Here again, we run statistical tests holding constant the individual differences that put an employee at risk of invitation to the program.[19]

One year after the program, graduates on average are indistinguishable from the control group; however, those who had been "active" participants are more likely than the control group to be promoted (2.2 *z*-score test statistic) and are well ahead of the "other" program participants, the people who were inactive during the program.

Two years later, program graduates on average are more likely than the control group to be promoted (2.0 *z*-score), but more striking is the now 173% higher odds of promotion for graduates who had been "active" participants (2.5 *z*-score). We conclude that, holding constant individual differences coming into the program, the program enhances the promotion-potential of participants, especially those who are "active" participants.

[19] The *z*-score test statistics in Fig. 8 come from logit models predicting promotion in each time interval from the two levels of program participation ("active" versus "other") relative to the control group, holding constant the summary measure in Table 1 that aggregates the risk factors predicting invitation to the BLP, and adjusting standard errors down for autocorrelation between observations of the same people at different times as a program participant or member of the control group (some people at high risk of invitation to the BLP were controls for the early BLP cohorts). We also ran tests for associations with promotion indicated by people remaining at the same job rank but working for a new boss. There are no statistically significant differences from the control group before or after the program. We also ran tests for associations with promotion indicated by people remaining at the same job rank but doing their work in a new city. Again, there are no statistically significant differences from the control group before or after the program.

9. Conclusion

Performance improves for executives educated in the network structure of social capital. Panel data on executives in a large organization—some of whom had participated in a Business Leadership Program (BLP) grounded in principles of social capital—gave us a rare opportunity to estimate social-capital education effects. Relative to a control group of untrained, but otherwise similarly-able peers, program graduates are 36–42% more likely to receive top performance evaluations, 43–72% more likely to be promoted, and 42–74% more likely to be retained by the Company. We conclude that the program of instruction has substantial positive effects for the participants and the Company. Executive performance is significantly higher for BLP graduates relative to high-performing peers at equal risk of attending the program.

These are noteworthy performance improvements, but their statistical significance is less striking. Effects are just over twice their standard error. That is statistically significant, but not by a large margin. In other words, performance did not improve equally for all graduates. In fact, active participation turned out to be critical to program effects. The subsequent careers of executives who were quiet spectators in the program cannot be distinguished from the careers of people in the control group, people who never attended the program.

These findings raise questions regarding the causal mechanisms linking the principles of social capital with improved performance on the one hand, and the relationship between active participation and outcomes on the other. Both questions carry important implications for both the practice of executive education and future research on its efficacy. In a recent interview, former CEO and chairman of CitiCorp and interim CEO and chairman of the New York Stock Exchange, John Reed, is quoted as saying, "when I look at the young people who come from business schools to work in the bank, they don't know anything that is any different from people 30 years ago. They haven't learned anything new," (Augier, 2006, p. 91). Reed's comment is rooted in substantive content and it may be that social capital represents an incremental addition to the substance of business curricula. The science of social capital has an interdisciplinary heritage and, as such, provides a comprehensive framework for analyzing the issues that contemporary executives face, issues that are inevitably interdisciplinary in nature. Furthermore, the brokerage and closure mechanisms that form the foundation of social capital (Burt, 2005) provide a wealth of practical insight on how to address the two fundamental challenges facing senior executives: how to exploit existing capabilities while at the same time paying adequate attention to the exploration of new capabilities essential for future growth and competitive advantage (March, 1991). As we suggest below, future research that systematically manipulates content can help answer this question.

With respect to the question of why active participation is so crucial to program effects, we speculate that the *experience* of the program was different for active versus passive participants. The importance of active engagement in the learning process has long been a touchstone of the field of "adult" learning (Knowles, 1984, 1990; Knox, 1978; McMullan and Cahoon, 1979), and much of expert performance can be attributed to deliberate practice (Ericsson et al., 1993). Active participation in the BLP often takes the form of practicing framing and re-framing arguments designed simultaneously to both create value through a strategic action and buy-in from stakeholders critical to its success. Active participants had many more cycles of such practice compared to their more inert counterparts.

The evidence displayed in Fig. 7 also suggests that active participants remained more engaged with projects initiated in the BLP subsequent to the program. Active participants offered fewer negative reports on the outcomes of their projects and more positive reports than their counterparts. Comments offered along with the project ratings suggest that active participants may have found alternative ways to derive value from a concept originating in their team project but not coming to fruition through the project.

Whatever the causal mechanism underlying the participation effect turns out to be, our impression is that merely communicating the effect has a dramatic impact on participation levels. In BLP programs conducted after this research was completed, we present the basic findings at the beginning of the program. Aware of a connection between participation and outcomes, everyone now participates at one time or another in the class discussion. We cannot say that the quality of discussion is improved, but there is more engagement (which could be the more important factor for returns to the company from running the program).

Experiments with program content come readily to mind for future research, but treating participation as a dependent variable is more problematic. We assembled a control group using "propensity scores" for a risk-of-invitation variable constructed from HR risk factors available to the people who nominated employees for invitation to the BLP, but "active participation" was not one of the risk factors. A practical implication of the analysis is that the risk factors should be expanded to include whatever is responsible for "active participation" in the program.

The passive participants have their own unique, if expensive, analytical value. Passive participants went through the motions of attending a senior executive program. They listened to the lectures and case discussions, ate meals with the active participants, traveled to company headquarters, were present when teammates presented to the CEO, and so on. They did not actively engage the material. In other words, the passive participants were in a kind of placebo program, and the negligible returns to passive participation are an indication that a placebo does not generate program effects. This is interesting speculation, but no more than that. The fact is that we cannot know for sure that we do not have a Hawthorne effect here. It is equally clear that there are substantial effects from the program as a complex treatment, so it is now interesting to test for component and Hawthorne effects in future research.

Apart from their practical implications, the results illustrate one way that executive education can be a productive research site. If social capital affects executive performance, then an event that enhances social-capital skills should enhance performance. We ran a field experiment to learn whether the BLP was such an event. Given the strong evidence of program effect on performance, the next step is replication with other programs and analysis of how variations in program content amplify or erode the performance effect. For example, would a program without the project work have the same effect on performance? Would more attention to the social capital of brokerage increase the program effect on performance? Returns to programs that differ in conceptual emphasis can be compared to see how training on different dimensions of social capital yields performance improvements. The bad news is the time required for such research. Fig. 8 shows modest program effect within a year, strong effect in 2 years. The good news is the felt reality of the evidence. This is not evidence of effects estimated from students role-playing managerial behavior in a campus laboratory. This is not evidence of effects estimated from a cross-section of executives using statistical controls to make inferences about

causal mechanisms. This is a field experiment with executives in vivo: Executives learn about the network structure of social capital in the program, then we watch the executives over time, along with a matched sample of their peers as controls, to see whether program-graduate performance improves relative to the controls. It does.

References

Augier, M., 2006. Making management matter: an interview with John Reed. Academy of Management Learning and Education 5, 84–100.

Baker, W.E., 2000. Achieving Success Through Social Capital. Jossey-Bass, San Francisco.

Brand, J.E., Halaby, C.N., 2006. Regression and matching estimates of the effects of elite college attendance on educational and career achievement. Social Science Research 35, 749–770.

Bulkeley, W.M., Wong, W., 2003. Six degres of exploitation? New programs help companies 'mine' relationships for key business prospects. Wall Street Journal (August).

Burt, R.S., 1982. Toward a Structural Theory of Action. Academic Press, New York.

Burt, R.S., 1992. Structural Holes. Harvard University Press, Cambridge.

Burt, R.S., 1999. The social capital of opinion leaders. Annals 566, 37–54.

Burt, R.S., 2005. Brokerage and Closure. Oxford University Press, New York.

Burt, R.S., Hogarth, R.M., Michaud, C., 2000. The social capital of French and American managers. Organization Science 11, 123–147.

Cohen, D., Prusak, L., 2001. In Good Company. Harvard Business School Press, Boston.

Cross, R., Parker, A., 2004. The Hidden Power of Social Networks. Harvard Business School Press, Boston.

DeSoto, C.B., 1960. Learning a social structure. Journal of Abnormal and Social Psychology 60, 417–421.

Ericsson, K.A., Krampe, R.T., Tesch-Romer, C., 1993. The role of deliberate practice in the acquisition of expert performance. Psychological Bulletin 100, 363–406.

Freeman, L.C., 1992. Filling in the blanks: a theory of cognitive categories and the structure of social affiliation. Social Psychology Quarterly 55, 118–127.

Granovetter, M.S., 1973. The strength of weak ties. American Journal of Sociology 78, 1360–1380.

Harrison, G.W., List, J.A., 2004. Field experiments. Journal of Economic Literature 42, 1009–1055.

Janicik, G.A., 1998. Social Expertise in Social Networks: Examining the Learning of Relations. Dissertation. University of Chicago, Graduate School of Business.

Janicik, G.A., Larrick, R.P., 2005. Social network schemas and the learning of incomplete networks. Journal of Personality and Social Psychology 88, 348–364.

Katz, E., Lazarsfeld, P.F., 1955. Personal Influence. Free Press, New York.

Kirkpatrick, D., 2003. I get by with a little help from my friends of friends of friends: ties created by social-networking sites have people excited. Fortune, Tuesday September 30.

Kirkpatrick, D.L., 1994. Evaluating Training Programs. Berrett-Koehler, San Francisco.

Knowles, M., 1984. Andragogy in Action. Jossey Bass, San Francisco.

Knowles, M., 1990. The Adult Learner. Gulf Publishing, Houston.

Knox, A.B., 1978. Adult Development and Learning. Jossey Bass, San Francisco.

Kotter, J.P., 1990. What leaders really do. Harvard Business Review 68, 103–111.

Lin, N., 2002. Social Capital. Cambridge University Press, New York.

March, J.G., 1991. Exploration and exploitation in organization learning. Organization Science 2, 71–87.

McMullan, W.E., Cahoon, A., 1979. Integrating abstract conceptualizing with experiential learning. Academy of Management Review 4, 453–458.

Merton, R.K., 1968 [1949]. Patterns of influence: local and cosmopolitan influentials. In: Merton, R.K. (Ed.), Social Theory and Social Structure. Free Press, New York, pp. 441–474.

Milgram, S., 1967. The small world problem. Psychology Today 1, 61–67.

Padgett, J.F., Ansell, C.K., 1993. Robust action and the rise of the medici, 1400–1434. American Journal of Sociology 98, 1259–1319.

Pfeffer, J., Fong, C.T., 2002. The end of business schools? Less success than meets the eye. Academy of Management Learning and Education 1, 78–95.

Putnam, R.D., 2000. Bowling Alone. Simon and Schuster, New York.

Rees, A., 1966. Information networks in labor markets. American Economic Review 56, 559–566.

Rosenbaum, P.R., Rubin, D.B., 1985. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. American Statistician 39, 33–38.
Rubin, D.B., 1973. Matching to remove bias in observational studies. Biometrics 29, 159–183.
Rubin, D.B., Thomas, N., 1996. Matching using estimated propensity scores: relating theory to practice. Biometrics 52, 249–264.
Salas, E., Cannon-Bowers, J.A., 2001. The science of training: a decade of progress. Annual Review of Psychology 52, 471–499.
Schumpeter, J.A., 1934 [1912]. The Theory of Economic Development (Redvers Opie, Trans.). Harvard University Press, Cambridge.
Turnbull, C.M., 1961. Some observations regarding the experiences and behavior of the BaMbuti Pygmies. American Journal of Psychology 74, 304–308.
Watts, D.J., 1999. Networks, dynamics, and the small-world phenomenon. American Journal of Sociology 105, 493–527.
Watts, D.J., Strogatz, S.H., 1998. Collective dynamics of "Small-World" networks. Nature 393, 440–442.