Information Sharing, Social Norms and Performance*

Marco Di Maggio
Massachusetts Institute of Technology
dimaggio@mit.edu

Marshall Van Alstyne
Boston University
mva@bu.edu

February 29, 2012

Abstract

What drives workers to seek information from their peers? And how does communication affect employee performance? Answers have proven elusive due to problems obtaining precise measures of white collar output and of observing the information individuals consume. We address these questions using an original panel data set that includes all accesses to an information-sharing platform, together with performance measures of all loan officers at a major Japanese bank. This paper makes three contributions. First, we show that skill level differences, job rotation, and differences among branches each affect the demand for information. There also exists substitution between an agent's ability and the amount of information consumed. Low skill agents benefit the most from consuming others' information. Second, restricting attention to officers who switched branches, we show that they perform on average significantly worse than before the switch, suggesting that job rotation destroys specialized human capital. We also find that an officer who shares information increases his chances of promotion rather than competes for promotion less effectively. Third, we measure the size of productivity gains based on consuming shared information. After controlling for unobserved heterogeneity over time, between branches, and among officers, a standard deviation increase in information access increases performance by roughly ten percent. By instrumenting the demand for information with the exogenous variation arising from cultural differences among branches, we are able to assess the causal effect of communication on performance.

*We thank Daron Acemoglu, Abhijit Banerjee, Tim Bresnahan, Robert Gibbons, Rocco Macchiavello, Decio Coviello, Marco Pagano, Christopher Palmer, Andrea Prat, Antoinette Schoar, Marianne Bertrand, John Van Reenen, Chad Syverson, and seminar participants at the NBER Summer Institute, the EEA, the MIT organizational economics seminar, Carnegie Mellon University, NYU Stern, and the conference on "The Economics of Intellectual Property, Software and the Internet." We would like to thank the NSF (Grant #0925004) for financial support. We also thank the outstanding editorial assistance of Emily Gallagher.

1 Introduction

A recurrent economic problem in organizations is how to use available knowledge efficiently. Information, however, is often dispersed among agents, which prevents optimal decisions if communication is absent. One important consequence is that organizations operating in markets requiring specialized knowledge can underperform. Credit, insurance, and financial markets are but a few examples. Agents might, in fact, decline a loan to a reliable client or fail to identify the appropriate risk class for a new client. Instead, when communication is available, agents can acquire information from co-workers in the organization. As Hayek (1945) pointed out: "it is a problem of the utilization of knowledge which is not given to anyone in its totality."

In the same spirit, Arrow's 1974 seminal work on the limits of organization argues that one key activity of any organization is internal communication. He has also pointed out that the trade-off between the cost of communication and the benefit of communication lies at the core of the agenda of organizations.

The effects of improved communication on productivity are in fact not obvious ex ante. On the one hand, low-skill agents may ask others for direction, thus communication might improve performance through a *learning effect*. On the other hand, high-skill agents might spend a larger fraction of their time helping others, which could negatively affect their performance, via a *substitution effect* as they swap information provision tasks for work completion tasks. Furthermore, even when production and use of information are not explicitly rewarded, a strategic motive might drive agents away from efficiency. Consider a high-skill worker facing the opportunity to share his knowledge with others. He might use this as an instrument to signal his ability in order to increase his chances of promotion. Alternatively he might prefer to hoard his knowledge to prevent a competitor from overtaking him on the ladder to promotion. Similarly, a low-skill worker might be reluctant to ask others for information required to complete his tasks out of fear of revealing his low ability. There exists then an internal job market *signaling effect* on behalf of high ability and a *adverse selection effect* on behalf of low ability workers, which influences agents' decisions. These effects are summarized in Table 1 where diagonal entries represent causal effects and off-diagonal entries represent selection effects. Finally, it is challenging to empirically disentangle the effects of communication versus those of innate ability regarding an agent's performance and demand for information.

The effect of knowledge sharing has recently attracted scholars' attention as one of main factors driving consumption and saving decisions (Moretti (2011) and Duflo

Table 1: Competing Hypotheses

Hypothesis	*Information Consumption*	*Hypothesis*	*Information Production*
↑ : Learning	People who ask questions gain skill & advice and are *more* productive.	↑ : Signaling	People who answer questions are higher skill and *more* productive.
↓ : Adv. Select.	People who ask questions are lower skill and *less* productive.	↓ : Substitution	People who answer questions are distracted and *less* productive.

and Saez (2003)), the creation of new technologies (Jaffe et al. (1993)) and adoption of technology (Foster and Rosenzweig (1995); Kremer and Miguel (2007); Conley and Udry (2010)). Our study extends this literature by empirically investigating how information sharing and communication affect workers' productivity within a large organization. Moreover, by doing so we are also able to understand what drives the workers' demand for information. We analyze these issues within the corporate division of a global bank primarily located in Japan. Starting in 2003, the bank implemented an online platform to allow employees to access information from headquarters and to share their private information with each other. Every employee gained online access to documents provided by headquarters as well as the ability to pose questions of and provide answers to other employees. Adopting this new technology allowed for more efficient communication, both vertically, between headquarters and bank branches, and horizontally among loan officers.

Information provided by bank headquarters mainly concerns legal and taxation issues, new financial instruments and services provided by the organization, and general policy guidelines as well as details about the most successful management practices to adopt. The information exchanged by loan officers is more often related to clients, for example, account management, credit worthiness, deal closure, collecting debts, and avoiding defaults. This provides us with a natural environment in which to understand how communication, and information sharing generally, affects productivity.

We exploit two unique properties of our dataset. First, we have very detailed data on the performance of each loan officer inside the bank for the two-year period 2006-2008. We observe the individual-level targets set by bank headquarters and the results achieved by each officer along a number of different dimensions, such as gross profit, loan volume and revenues. Since these measures are expressed in yen and easily quantifiable, the performance of the agents on these dimensions constitutes our *objective* performance measure. However, we also collected the branch manager's evaluations of each loan

officer's performance on softer more informal dimensions, such as his or her contributions to branch operations and customer service. These evaluations can be interpreted as *subjective* performance measures because they are not directly related to the attainment of a pre-determined target. Moreover, as reported by the bankers these evaluations try to correct for luck and to reward officers' effort.

Second, we are able to exploit a source of exogenous variation coming from an anti-corruption law that requires loan officers to switch branches every two to five years. This allows us to use an instrumental variable procedure in order to assess the causal effect of information sharing on employees' productivity. Moreover, our results are not contaminated by incentive considerations, in fact, there is no material incentive to use the new technology. Hence, if acquiring knowledge from others became an important part of the loan officer's job, it is exclusively due to the reliability and the productivity benefits of access to available information. Furthermore, the banking sector in Japan is very different from that in the United States. In particular, the incentive system implemented by this bank provides loan officers with only de minimus end-of-the-year bonuses, so incentives are almost exclusively constituted by the possibility of promotion.

This paper provides two main sets of results. First, exploiting the longitudinal dimension of our data, we are able to identify a significant positive effect of information sharing on performance. Moreover, we also highlight what is the main mechanism through which communication affects performance. We provide evidence that low-skill agents benefit the most from acquiring information from others. The magnitude of the effect is large and significant. In fact, controlling for unobserved heterogeneity over time, between branches and among officers, a standard deviation increase in information access increases performance by more than ten percent. This supports the idea that agents might under-perform because they are not aware of more successful management practices implemented elsewhere or lack the information necessary to fill the gap (Bloom et al. (2011)).

We can also conclude that asking questions and providing answers to others is mainly driven by innate ability, as captured by individual fixed effects. Intuitively, low-ability agents, in an attempt to improve their performance, are more likely to ask questions to their peers. Interestingly, we also show suggestive evidence that the likelihood of being promoted is positively associated with the information officers produce (e.g. number of answers posted) and negatively affected by how much information officers acquire (e.g. number of questions posted). This evidence, albeit not conclusive, suggests that a signaling motive might be present.

Second, we exploit the source of exogenous variation by restricting attention to agents

who switched branches. We can study how job rotation affects (1) the demand for information and (2) the officers' performance. We find that they tend to increase their demand for information immediately after switching jobs. However, this effect declines over time. That job rotation increases information demand supports the hypothesis that learning is a major factor affecting performance. Intuitively, when a loan officer is forced to switch branches, he does not possess, for example, the knowledge of the local market conditions to assess the reliability of new clients. Without the ability to communicate with more experienced officers, whom he also has trouble identifying, an officer's productivity can suffer.

Restricting our attention to "switchers," gives us the opportunity to investigate their performances in the new branches. We find strong evidence that switching negatively affects their performance and, after the switch, they perform on average significantly worse than before. This result is of independent interest as it suggests specialized human capital is destroyed when they move to a different branch. This is surprising because we are considering the same worker within the same organization, controlling for regional and branch differences. This finding contributes to the labor literature on the accumulation and destruction of specialized human capital and the effect of worker displacement.[1] As we discuss in section 6, this presents us with an opportunity to quantify and bound the costs associated with the anti-corruption law.

Finally, we further exploit the mandatory switching of loan officers across branches in order to assess the causal effect of communication and information sharing on performance. Motivated by branch variation in usage of the information-sharing platform, we construct an instrument based upon the attitudes of a branch towards the new technology. For each officer i, our instrument is the amount of information accessed in the previous branch excluding officer i. That is, if specific officer A works in a branch where problems are usually resolved within the branch and without attempting to find the solutions elsewhere in the organization, then when officer A moves to a different branch, he will tend to communicate less with other officers in other branches. In contrast, if officer A works for a branch where access to information provided by others is encouraged, he will tend to communicate more across the organization when in the new branch.

We find even stronger results than those found with the estimation of the longitudinal model. This provides further evidence in favor of the hypothesis that there exists substitution between the loan officers' ability and the amount of information to which they have access. Moreover, this suggests that in contrast to the o-ring theory

[1] Seminal papers in this strand of the literature include Hamermesh (1987), Ruhm (1991), and Jacobson and Sullivan (1993). For an early survey see Kletzer (1998).

formulated by Kremer (1993), which predicts that the productivity of each worker is increasing in the skill level of his co-workers, we find an *asymmetric* effect. While we observe that the opportunity to share information with high-skill workers increases the low-skill workers' productivity, we do not observe any negative effect on the high-skill workers' performance.

The significance of our instrument suggests that agents take into account their peers' behavior, when deciding how to cope with their daily tasks. For example, suppose loan officer A is facing a taxation issue with one of his corporate clients. Our results show that, not only he will be more inclined to ask questions if he has underperformed in the past (i.e. indicating a low-ability agent), but he will be more inclined to do so if his peers behave in the same way. That is, there exists a *complementarity* in information sharing: the higher the number of officers who consult their colleagues, the higher is the incentive for each of them to continue sharing information. This might be explained by the absence of competition and the lower risk of being identified as a low-ability agent which based on having provided information.[2] Moreover, we show that it is more likely for officer A to meet his targets and improve his performance, by solving his taxation issue with the help of others, than by taking an uninformed decision.

The paper is organized as follows. The next section places this paper in relation to existing literature. Section 2 discusses the institutional background, and describes our data. Section 3 explains our approach and the methodology we employ to estimate the effect of information sharing on performance. Section 4 presents the main results for the effect of communication on performance, promotions, and the effect of switching on information demand and productivity. Section 5 presents the results of our instrumental variable estimates. Section 6 analyzes the relevance of our results for three different issues: cost of regulation, theory of tournaments, and relational banking. Section 7 summarizes and concludes.

1.1 Related Literature

Since Marschak and Radner (1972) pioneering work, team theory has theoretically investigated issues similar to our own and the literature has developed around the idea that information flows, and not just incentives, drive agents' behaviors inside an organization. Sah and Stiglitz (1986), for example, is an early attempt to compare decision-making in different organizational forms when agents possess heterogeneous information.[3] More

[2]For an interesting overview of the economic literature on corporate culture see Hermalin (2007).

[3]Other important contributions in this literature include Radner (1993), which studies the role of networks in minimizing human limitations in information processing and Bolton and Dewatripont (1994)

Recently, Garicano (2000) presents a theoretical model of hierarchical organization of expertise, which is applicable to our setting.[4] Decisions involve problem solving and thus acquiring the relevant knowledge for each decision. There exists a trade-off between information acquisition costs and communication costs because agents can directly acquire information at a cost or elicit the relevant information from others at a higher level in the organization. The latter is costly because agents at the higher levels need to spend time solving problems faced by others. This is what we call the substitution effect. Bloom et al. (2009) employ an international data set in order to investigate the effect of information technology and communication on worker autonomy, plant manager autonomy, and span of control. We complement their analysis by focusing on the workers' performance and their demand for information.

Researchers have long investigated the differences in productivity performance between firms and plants within sectors and across countries. The magnitudes involved are striking. For example, within the 4-digit SIC industries in the U.S. manufacturing sector, the average ratio in productivity between the highest and lowest percentiles is around 1.92.[5] Most of the applied economic research has focused on documenting and explaining the main drivers of these performance differences. Existing works in different fields have linked productivity levels to a number of features of technology, demand, human capital and market structure. However, to create persistent performance differences the advantageous inner workings must be difficult to imitate, and this suggests that part of the performance variations across similar enterprises might be due to other aspects of organizations such as management practices, communication, and individual talent, which are softer and more informal than other factors identified in the literature.[6] We contribute to this debate by identifying another factor which significantly affects individual performance, that is, the possibility of acquiring knowledge from others.

Recently Bloom and Reenen (2007) and Bloom and Reenen (2010) surveyed 732 medium-sized manufacturing firms in four countries, collecting data on their management

and Van Zandt (1999) which highlight the importance of hierarchies to diminish the costs related to processing information that flows through the network of contacts.

[4] Another related paper in this strand of the literature is Garicano and Rossi-Hansberg (2006). Niehaus (2011) considers homogeneous agents, who share their knowledge with peers whenever private benefits exceed communication costs.

[5] See Syverson (2004) for a in-depth analysis of the productivity dispersion in the U.S. manufacturing sector. Abraham and White (2006) and Foster et al. (2008), among others, study the evolution of productivity over time. Bartelsman and Doms (2000) present an earlier survey on productivity, while Syverson (2010) focuses on the recent contributions in this area that attempt to explain why businesses differ in their measured productivity levels.

[6] See Gibbons and Henderson (2010) and Gibbons (2010) for surveys of the literature on performance differences across similar enterprises.

practices regarding operations, monitoring, targets, and incentives. The implementation of these practices is highly correlated with its total factor productivity. On one hand, the nature of our data limits the analysis to one bank but, on the other, it allows us to employ the same fine-grained performance measures used by headquarters to assess loan officers' performances and to avoid problems related to survey data.[7] In contrast to Bloom and Reenen (2007), our main focus is the analysis of the effects of the information generated within the organization on white collar productivity.

Our paper is also related to the strand of the literature that examines the importance of technological adoption in increasing productivity, such as Brynjolfsson et al. (2007), Faggio et al. (2010), and Aral et al. (2007). Brynjolfsson et al. (2007) document case studies where IT enhances the speed with which firms can replicate practices they find productive in one of their lines of business across the entire organization. Our paper addresses a different question, since we can look at individual-level performance measured over time and at the intensive margin in the usage of new technologies.

Finally, this paper is related to the studies of the impact of human resource management on firm performance, such as those by Ichniowski et al. (1997), Lazear (2000), Black and Lynch (2001), and Bartel et al. (2007). In particular, Ichniowski et al. (1997) find that human resource practices are observed in bundles, rather than being independently distributed; and second, different bundles are associated with substantial differences in productivity. We find evidence supporting the hypothesis that being aware of different and more productive practices within the same organization can foster increased productivity.

2 Empirical models and results

The foregoing discussion suggests that the officers' performances might be affected by several different factors such as branch characteristics and their access to information. Because the estimation strategy is affected by data availability, this section begins with a description of the institutional background and the data.

[7] Bandiera et al. (2009) employ a similarly fine-grained personnel data to study the effect of social connections and incentives on productivity, within the fruit picking division of a UK producer of soft fruit. Loan officers' behavior have been recently studied by Hertzberg et al. (2010) who show that a rotation policy that routinely reassigns loan officers to borrowers of a commercial bank affects the officers' reporting behavior. While they focus on the loan officers' moral hazard problem in communicating with the headquarters, we analyze how communication among officers might significantly improve their productivity.

2.1 The Setting

We analyze the behavior of loan officers, also called "relationship managers," in the corporate banking division of a major Japanese bank (the Bank) during the two-year period 2006-2008. Located across more than two hundred branches throughout Japan, the officers' primary tasks are to grant and manage loans to local enterprises. Their performances are assessed every six months and are measured by the percentage of the targets met during the same time. The officers' performance can be affected by two main factors. First, there exist regional differences between branches, such as the local demand for loans and the profitability of local enterprises. Second, there is some heterogeneity in loan tasks, in fact, while some officers only deal with the public administration, others need to re-structure more profitable loans or solicit loans from new clients. However, we shall take into account these sources of heterogeneity among officers.

As explained in the introduction, one of the main differences between a U.S. bank and the Japanese bank we are analyzing is the incentive system. While end-of-the-year bonuses are extensively adopted in the U.S. banking sector, the Bank rewards its loan officers by means of promotion. We observe about two hundred instances of promotion in our sample, and we can investigate the effect that information sharing has on the probability of being promoted. Moreover, in Japan seniority is by far the most important factor that increases promotion probability.

By law, the Bank implements a switching rule as a way to prevent bribery and graft among loan officers. This regulation obliges loan officers to change branches every two to five years, which allows us to disentangle the effect of an individual branch's working environment from the officer's ability on productivity. We shall show that headquarters do not relocate officers based upon their past performance.

In this environment, information sharing among loan officers has several effects. First, allowing officers to share information lets them better assess the riskiness of client enterprises, or work on more projects at the same time. Second, it might allow the low-ability workers to bridge the gap separating them from the most productive ones. Third, high-skill workers might be required to devote a larger fraction of their time helping their colleagues, which could reduce their performance. Fourth, anticipating that this can increase their chances of being promoted, officers might share their knowledge with others in order to signal their expertise in a particular field. Our main contribution is to identify and disentangle the different effects that communication has on productivity.

We now discuss the features of this work environment that allow us to assess whether information sharing shapes individual performance.

2.2 Data Characteristics

We collected data on performance, communication and information sharing from the corporate banking division of a major Japanese bank. Our primary data source is the bank's personnel records. These include all loan officers, approximately 2800 people, located across hundreds of branches in Japan. Branches vary in size and primary type of business, mainly due to location. In general, metropolitan branches have more loan officers – between 30 and 100 – and larger enterprises as customers, while those located in suburban areas have fewer officers, about 10, and smaller businesses as customers. Our data span October 2006 through September 2008. Prior to this study, only in the local branch employing individual loan officers kept their performance data. Information asymmetry made it difficult for loan officers to observe each other's productivity across branches, likely increasing dispersion of individual output. Since the Bank had a major merger in October 2005, we focus our attention on the stable, second year of the new bank to avoid having merger activity influence results via changes in officers' performances. The strengths of the data lie in their fine-grained level of detail and the possibility of tracking each officer's performance over time.

Dependent Variable. Loan officers are reviewed semi-annually to assess their performance. In order to account for branch location and task differences, we control for the six main groups to which each officer may belong. These groups are: large existing account, small existing account, loan restructuring, public sector, new strategic account, and new non-strategic account. The main differences among groups depend on the different clientele. For example, officers working in the public sector group exclusively deal with public administration, while officers in the restructuring group try to renegotiate underperforming loans. Bank headquarters sets the targets for managers in these groups and to each loan officer the head of the branch assigns a score of up to 80 percentage points based on his quantitative measures and up to 20 percentage points based on qualitative measures in the categories of Table 2.

The dual sets of metrics help to increase accuracy relative to stochastic environmental shocks. As reported by bank executives, the branch managers often assign qualitative scores as a reward for major effort that did not yield results or discount low effort that did. The weighted performance of loan officers along dimensions in the first column constitutes our objective performance measure, and weighted performance along the second column constitutes our subjective performance measure. Results are robust to using either column.

Our data include all targets, objective results, and subjective scores for each loan

Table 2: Output Metrics

Quantitative	*Qualitative*
revenue	
individual loan profit	customer service
liquid deposit profit	loan reinforcement
loan volume	contribution to branch operations
reduced delinquencies	contribution to organization operations
reduced estimated losses	
bank gross profit	

officer in each group and for every branch. That is, we have all performance assessments between 2006 and 2008 for the corporate division of the Bank. We believe that the richness of our data and the fact that we need not rely on wage data to extrapolate observed performance make it highly suitable for study of productivity differences across officers.

In most of the analysis below, we focus our investigation on one dependent variable: total performance. This is the total score assigned to officers, representing the weighted sum of performance in each of the different categories and including the individual subjective score. It also represents the performance metric used by the bank.

Independent Variables. The main variables of interest capture how loan officers use the internal platform to share and gather information. We collected data on all accesses to the information platform by each officer during the period of interest. Our data include (i) the number of documents consulted by each officer, (ii) the number of questions posted, and (iii) the number of answers provided, down to the second of access and across each term.

We also have information on the number of years the officer has worked for the Bank, captured by the variable "tenure," and whether he came directly from school with no prior experience (or transferred from another bank), captured by the dummy variable "college."

2.3 Descriptives

Table 3 reports descriptive statistics for our variables of interest, and two things are worth noting. First, loan officers perform significantly differently, in fact, the mean of our main measure of performance is 52, but the standard deviation is 21. Even if we focus on the employees of the same bank, within the same region and with homogeneous tasks, we still find that their performance is heterogeneous. Second, loan officers seem

to make great use of the available information within the organization. On average they access 569 documents, post 77 questions and provide 250 answers during a six-month period. Moreover, the standard deviations of all three variables are quite high, ranging from 164 to 282. This will be relevant to interpret estimated coefficients.

Our main hypothesis is that officers can access the information generated within the Bank in order to improve their performance. In particular, we expect low performers to ask more questions and provide fewer answers.

Figure 1 shows that the kernel density of our total performance measure. Note that performance for below-the-median number of answers (above-the-median number of questions) is to the left of above-the-median number of answers (below-the-median number of questions). The loan officers who help others more often, answering their questions, on average perform better than the others. In contrast, the loan officers that ask more questions are associated with lower performance. Table 4 confirms this intuition showing the officers' performance for those who have shared information more or less than the median officer. The first column shows that there is no significant difference in performance between officers who had access to greater or fewer numbers of documents than the median. The second column, instead, shows that there exists a positive correlation between the number of answers provided and performance. The loan officers who provide a higher number of answers perform significantly better than the others. The third column confirms this result showing that those who ask more questions, above the median, perform significantly worse than the others. These results suggest that information sharing is correlated with performance and with the officers' innate ability.

In the remainder of the paper, we present formal evidence to shed light on whether these descriptive results are robust to controlling for other determinants of performance. In doing so, we make precise the underlying identifying assumptions required to interpret when this evidence is causal and present evidence in support of these identifying assumptions.

3 Information Sharing and Worker Performance

In this section we explain our approach and the methodology we employ to estimate the effect of communication and information sharing on performance.

3.1 Longitudinal Specifications

The analysis proceeds in two stages. First, we estimate the effect of information sharing on loan officers employing a longitudinal model that allows us to control for unobserved heterogeneity between branches and workers. Next, we take advantage of an anti-corruption law that requires officers to switch branches every few years to identify the causal effect of information sharing on performance.

To identify whether information sharing affects officer performance, we estimate the following panel-data regression:

$$y_{i,j,k,t} = \alpha_1 I^a_{i,j,k,t} + \alpha_2 I^g_{i,j,k,t} + \alpha_3 I^p_{i,j,k,t} + \lambda_t + \eta_k + \gamma_j + \delta_i + \beta_1 T_i + \beta_2 C_i + \varepsilon_{i,j,k,t} \quad (1)$$

where $y_{i,j,k,t}$ is officer i's log performance in branch j in group k and during term t. The main variables of interest are I^a, I^g, and I^p which capture the number of documents, questions and answers posted by officer i in branch j. The time fixed effects λ_t account for unobserved shocks that might have affected both the officers' performance and their demand for information, such as those arising during a financial crisis. The group fixed effects η_k capture permanent productivity differences across regions and tasks, such as those arising from the different clientele and heterogeneity of loan types. The branch fixed effects γ_j allows us to control for permanent productivity differences across branches, such as those arising from a more profitable location or a better head manager of the branch. Finally, individual fixed effects δ_i provides the possibility of controlling for innate ability or motivation. We also include the tenure of the loan officer T_i, when we do not include officers' fixed effects and the dummy C_i which is equal to 1 if the loan officer joined the bank right after college and equal to 0 if he had previous experience in the banking sector.

We also note that information sharing and performance are unlikely to be identically and independently distributed within a branch. We therefore adopt a conservative strategy when estimating standard errors and allow the disturbance ε_{ijkt} to be clustered by officer throughout.[8]

[8] Clustering the disturbance terms by branch leads to the standard errors on the parameters of interest being considerably smaller than those we report.

3.2 Instrumental Variable Specification

To better assess the causal effect of communication and information sharing on performance, we exploit the mandatory switching of loan officers across branches. Since there is variation in usage of the information sharing platform across branches, we can construct an instrument based upon the attitude of a given branch toward the new technology. For each loan officer i, we construct an instrument Z_{-i} which is the amount of information accessed in the previous branch excluding officer i. The choice of this instrument is motivated by the idea that if officer A worked in a branch where problems are usually resolved within the branch, without attempting to find solutions elsewhere in the organization, then even when officer A moves to a different branch, he will have been trained to communicate less with other loan officers. In contrast, if officer A worked for a branch where the access to information provided by others is encouraged, he will carry that attitude into the new branch. We construct similar instruments for each of our three endogenous variable of interests: the number of documents accessed, the number of questions posted, and the number of answers provided.

Formally, the first stage for each endogenous variable $e \in \{a, g, p\}$ is represented by:

$$I^e_{i,j,k,t} = \alpha_1 Z^a_{-i,j,k,t} + \alpha_2 Z^g_{-i,j,k,t} + \alpha_3 Z^p_{-i,j,k,t} + \lambda_t + \eta_k + \gamma_j + \beta_1 T_i + \beta_2 C_i + \eta_{i,j,k,t}$$

while the second stage is

$$y_{i,j,k,t} = \alpha_1 \widehat{I}^a_{i,j,k,t} + \alpha_2 \widehat{I}^g_{i,j,k,t} + \alpha_3 \widehat{I}^p_{i,j,k,t} + \lambda_t + \eta_k + \gamma_j + \beta_1 T_i + \beta_2 C_i + \varepsilon_{i,j,k,t}$$

where we have employed three instruments for the three endogenous variables. The validity of this instrumental-variable procedure relies on the relevance of our instruments and their exogeneity. First, we shall show that our methodology is not affected by the "weak instrument" problem, in fact, the coefficients in the first stage regressions for each endogenous variable are highly significant, and the F-Test is always above 10, the standard threshold for weak instruments. Second, we have constructed our instruments for officer i, excluding officer i from the computation of the information accessed in his branch. This should reduce the correlation between the instrument and officer i's innate ability. However, since our instruments rely on cultural variations across branches we are afraid that officer i might have contributed to the branch's culture in the past, which could bias our estimates. We address this concern by restricting attention to larger branches (with more than 50 officers), for which this possibility is, at least, less likely.

4 Baseline Results

Table 5 presents estimates of our baseline specification (1). The results show that the pattern of unconditional differences in worker performance by information sharing is robust to conditioning on a rich set of determinants of officer performance. It presents estimates for the main parameter of interest showing that the number of documents and the number of answers are positively correlated with individual performance, while the number of questions is instead negatively correlated with their performance.

Furthermore, tenure significantly affects performance, which can be the result of a learning process within the Bank. At the same time, joining the bank right after college, without any previous experience, is positively correlated with performance. This result can be interpreted as a result of the greater effort exerted by new employees.

The main concern with these results is that information sharing and the effect on performance might be driven by other factors, such as a market downturn, a greater need for information for a specific local market, or a result of a better performing branch. In order to control for all this unobserved heterogeneity, as shown by column (4), we control for time, group and branch fixed effects. Except for the effect of the number of answers, the other coefficients are still economically and statistically significant.

In column (5) we further control for the interaction of time and branch fixed effects, which shows that the results are robust to this more restrictive specification. These results suggest that when officers increase the number of documents accessed, this has a positive and significant effect on the productivity of the average worker, whereas increasing the number of questions has a negative impact on his performance. The magnitude of these effects implies that when officers increase their information access by one standard deviation it increases their performance by eleven percent. Similarly, an increase in the number of questions is associated with a reduction in performance of about five percent.

A concern with these results is that the estimation might be picking up heterogeneous effects that are unrelated to information sharing, in particular one of the main factors for which we cannot directly control: officer's ability. It is plausible that innate ability has a significant effect, which would create a spurious correlation between information sharing and performance. For example, it is likely that a loan officer who is able to close a higher number of deals and identify the most profitable ones will post fewer questions. Then, observing a negative correlation between the number of questions and the officers' performance might just be driven by heterogeneous innate ability or differences in training.

Table 6 therefore provides evidence on the effects of communication and information sharing on the productivity of the *same* worker. We exploit the longitudinal nature of our data and control for individual fixed effects. In accord with the descriptive evidence presented in the previous section, once we control for the individual fixed effects as in column (5), the number of questions is not significant anymore. However, we find an even stronger effect for the number of documents. This correlation suggests that performance could be significantly affected by the number of documents consulted by the loan officers, even controlling for individual ability. As highlighted in the introduction, we interpret this as evidence of the possibility that loan officers know more successful management practices implemented elsewhere in the bank. We instrument this explicitly in Section 5

4.1 Who Benefits the Most from Information?

To explore whether the effects of information sharing are heterogeneous across loan officers, we use quantile regression methods to estimate the conditional distribution of the log of performance of loan officer i in branch j, and group k during the term t, $y_{i,j,k,t}$, at different quantiles, θ. We therefore estimate the following specification:

$$Quant_{\theta}\left(y_{i,j,k,t}|\cdot\right) = \alpha_{1\theta} I^{a}_{i,j,k,t} + \alpha_{2\theta} I^{g}_{i,j,k,t} + \alpha_{3\theta} I^{p}_{i,j,k,t} + \beta_{\theta} X_{i,j,k,t} + \varepsilon_{i,j,k,t} \qquad (2)$$

All variables are as previously defined, and bootstrapped standard errors, based on 1000 replications, are calculated throughout. The effect of information access, gathering, and production on officers' performance at the θth conditional quantile of log performance is measured by the vector α_{θ}.

Table 7 reports the estimates of α_{θ} from the specification above at various quantiles, controlling for tenure and experience as well as time, group and branch fixed effects. Two points are of note. First, the effect of information access is zero for the top two quantiles, and is positive and significant at the bottom three quantiles. Second, the effect of information gathering as measured by the log of the number of questions posted is negative and significant for all quantiles. In particular, a ten percent increase in the number of documents predicts a performance increase of at least twenty percent, whereas the same increase in the number of questions posted predicts a decrease of more than twenty-five percent.

The data suggest that information access increases the performance of loan officers in the left tail of the productivity distribution, while it has no significant effect on officers' performance in the right tail of the distribution. These results provide evidence

that information transfers might help the low performing officers learn from the most productive officers without disrupting the higher performers. That is, the learning effect dominates the substitution effect.

This result has two implications. First, it suggests that in a distributed decision *polyarchy* (see Sah and Stiglitz (1986)), like the setting investigated here, the agents benefit from communicating with each other due to the replicability of their decisions. For example, two loan officers serving different clients might improve their productivity by sharing information, because the profitability of these loans is affected by common factors such as the credit market conditions and the available financial products. Second, since we do not observe any negative effects on high-skill agents' productivity, we can conclude that these white-collar workers do not suffer from the "information overload" problem identified by Van Zandt (2004). This is probably due to the digital platform's ability to disseminate reusable information without having high performers re-enter answers to the same questions.

4.2 Promotion Probability

Since the Bank's incentive system is mainly based upon promotion, it is natural to check whether communication among loan officers, and its effect on productivity, are reflected in promotion probability. Although the Bank does not directly incentivize loan officers to share their information with their colleagues, Bank headquarters might recognize, on the one hand, that a loan officer has developed valuable expertise in one particular field, and that he should be rewarded for answers and knowledge spillovers that benefit others. On the other hand, headquarters might also infer a lack of skill from the number of questions posted by the officer. These signaling effects should affect his probability of advancement.

Although headquarters might not directly observe the amount of information demanded and produced by each loan officer, this data can still be reflected in the promotion probability. In fact, the head manager of each branch, who has a better knowledge of the information generated within the branch, makes recommendations to headquarters regarding which loan officers deserve promotion.

Table 8 presents evidence to support these hypotheses. Specifically, we observe roughly two hundred promotions in our sample, which allows us to investigate the effect of communication on the probability of being promoted. Estimates show that, as expected given the institutional background, tenure has a positive and significant effect on promotion probability, and it is by far the greatest predictor of promotion.

Column (1) shows that the coefficients on the first lag of productivity, as captured by the officers' performance in the previous term, is positive but not significant. This means that variations in the number of questions or answers do not pick up promotion effects otherwise driven by lagged performance. Even controlling for lagged productivity, time, group and branch fixed effects, there exists a positive correlation between the number of answers and the probability of being promoted, while the correlation with the number of questions is negative. These results, even if they cannot be interpreted as causal and definitive, are consistent with our interpretation. In particular, column (5) shows that there might be an element of signaling when officers decide to share their information with others, and an element of skill-level disclosure when they ask for information. Providing others with relevant information increases, in fact, the probability of being promoted, even if it has no significant or direct effect on performance, as described in the previous section.

4.3 Effect of Switching

Up to now we have found evidence that communication and information sharing might help the low-performing officers acquire the necessary knowledge to improve their performance. If this is true, we should then expect officers to significantly increase their access to the available information when they switch branches. A loan officer might, for example, start working in a different environment, with different existing customers, and a different local credit market, which should have a significant impact on his demand for information. The switchers might then require information about the reliability of the customers and the conditions offered by competitors. We observe 618 loan officers who switched branches, as prescribed by the anti-corruption law, after two years of experience in the same branch.

As a first step, we investigate whether or not Bank headquarters relocates officers to different branches based upon their performance. On the one hand, it might be that in order to improve the productivity of a branch with below-average productivity, the Bank might find it optimal to allocate the best officers to the branches that need to improve their productivity. On the other hand, the high-performing officers might be rewarded by being allocated to the "best" branches. Figure 2 favors the latter case. It displays the slightly positive relation between the mean officers' performance before the switch and the productivity of the branch where they work after the switch. This means that the Bank does not employ the switching rule to strategically locate officers across branches to improve branch productivity.

Table 9 reports coefficient estimates on job rotation relative to document consumption – indicator variable "switch" is equal to one when a loan officer moves from one branch to another. As highlighted by columns (1) and (2) the results are consistent with the learning hypothesis. Even controlling for time, group, branch and individual fixed effects, the coefficient is positive and both statistically and economically significant. This suggests that switching might be an important determinant of the demand for information. Intuitively, the less-experienced officers would try to acquire a greater amount of information, as shown by the coefficients on tenure, negative but insignificant, and the coefficient on college, which is instead positive and significant.

However, if the demand for information is driven by a temporary need driven by the new environment, we should observe a diminishing effect of switching over time. Columns (3) and (4) investigate this issue, presenting the estimate for another indicator variable, "after switch," which equals one for all terms after the switch. The effect is still positive, but no longer significant. This confirms the hypothesis that switching has only a short-run impact on the demand for information.

Since we have identified a significant effect of switching on communication, it is now natural to investigate the level of costs associated with the application of this anti-corruption law. In particular, we can investigate if switching has a positive or negative effect on the loan officers' performance. On the one hand, switching might result in the officers exerting more effort during the first few months of the new appointment to signal their ability. On the other, loan officers might have acquired some specific knowledge about the type of firms and market conditions in the previous branch, which suggests that after the switching they would need more time to learn work practices in other environments.

Table 10 presents evidence that strongly supports the latter hypothesis. As shown by columns (3) and (4), switching has a negative and significant effect on performance. Moreover, this effect is even stronger for longer-tenured officers, as suggested by the negative coefficient on the interaction term between the indicator variable and the officers' tenure. This means that even if loan officers might tend to work more when they are forced to change branch, the overall impact on their performance is negative. As in the case of the demand for information we should expect a decreasing effect of switching on the officers' performance over time. Column (5) shows the coefficient estimates on the indicator variable that accounts for all the time after the change of branch. Although still significantly negative, its magnitude is diminished.[9]

[9]The negative effect of relocation on performance can be driven by adjustment costs borne by the loan officers, who have to adapt themselves to a different branch cultures. However, at least part of this

The last two columns (6) and (7) assure that these results are robust to the inclusion of individual fixed effects. Overall this evidence suggests that implementing a switching rule as a way to prevent bribery and corruption can impose high costs. In particular our estimates suggest that some specialized human capital is destroyed when a loan officer switches from one branch to another. In contrast to the existing literature on worker displacement (which investigates the effect of layoffs on earnings), we are able to estimate the effect of turnover directly on performance. Moreover, we have the advantage of analyzing a sample of white-collar workers that switch locations within the *same* organization, and with the same tasks. This guarantees that the negative shock to officers' performance is not driven by the relocation to a different firm or to a job that requires another set of qualifications.

One problem with the interpretation of these results might come from the possibility of endogenous adjustment of performance targets after the switch. Incoming officers might be assigned less desirable loan portfolios. Yet, if bias were present, less attractive loan portfolios would reduce performance, making the positive effects of access *a fortiori* a stronger result. Downward bias would be present if our performance measures were the absolute results achieved by the loan officers, instead, our performance measure is the score assigned to the loan officers based upon the percentage of the targets met aggregated across all relevant dimensions. Then, even an endogenous change in targets is reflected in our dependent variable without necessarily biasing our results.

5 Instrumental Variable Estimates

Up to now, the evidence presented strongly suggests the existence of a significant effect of communication on officers' performance. Specifically, there is substitution between the demand for information and the innate ability of officers. Moreover, low performers or officers who just switched to a new branch significantly increase their demand for information. We can now address a natural endogeneity problem that can arise in our context. The loan officer who is facing a market contraction, for example, can *decide* to acquire more information in order to improve his performance. Although in the previous estimates we have accounted for a variety of unobserved shocks, with the inclusion of a set of fixed effects, we now try to address this endogeneity issue in order to understand whether we can interpret the results as causal or not.

Since we do not have data before adoption of the new technology, which allowed for

is captured by the inclusion of the branch fixed effects, which should control for unobserved cultural heterogeneity.

information to be shared, we are not able to run a natural experiment to understand the effect of information sharing on officers' performance. Nevertheless, we are able to address this issue by employing the anti-corruption law as a source of exogenous variation. For each officer i who switched from branch j at time t to branch j' at time $t+1$, we use the number of documents (as well as questions and answers) accessed in the branch j as an instrument for the number of documents (and the number of questions and answers) that officer i had access to in branch j' at time $t+1$. That is, we exploit the variation in branch attitudes toward information sharing to determine the effect of communication on performance. Then, we shall use the cross-sectional variation among switchers to identify the effect of communication on performance.

Table 11 presents the first stage estimate for each one of the endogenous variables. The first column reports the coefficient estimates of our instrument for the number of documents, which shows that both the coefficient (positive and statistically significant at the one percent level) and the F-test (above 10) strongly suggest that our instruments are not weak. Columns (2) and (3) present the first stage instruments for the number of questions and the number of answers. As for the documents, our instruments seem to significantly affect the demand for information. Intuitively, for all three variables of interest, tenure has a negative and significant effect, which confirms that even restricting attention only to the switchers, the more experienced people demand less information. Given the small sample of switchers we are not able to control for branch fixed effects, but we include both time and group fixed effects.

Table 12 presents the ordinary least-square estimates restricted to the switchers sample and the instrumental-variable estimates. The coefficient on the number of documents is negative and not significant for all the OLS estimates while positive and highly significant for the IV estimates. This can be interpreted as evidence in favor of our substitution hypothesis, that is, high-ability officers tend to seek out less information than their colleagues. The magnitude is also interesting, in fact; the most conservative specification in column (6) suggests an effect of about ten percent on performance. This means that incentivizing the usage of the information produced by others within the same organization might actually result in a significant improvement in productivity.

Interestingly, Black and Lynch (1996) found that a 10% rise in average education, roughly one year of schooling, led to an 8.5% productivity increase in manufacturing and a 12.7% increase in non-manufacturing. The ten percent gain we find in banking therefore appears comparable to just under one year of education.

Our second variable of interest, the number of questions, has a negative effect on performance in both the OLS and IV estimates. This is consistent with the previous

results, and suggests that even the exogenous variation in the number of questions negatively affects performance. The magnitude is higher for our IV estimates than in the OLS results, ranging from five percent to almost twenty percent. This confirms the substitution between officers' ability and the number of questions posted.

Finally, in contrast to the panel estimates presented above, the number of answers has a significant, positive effect on performance. These estimates show that the exogenous variation in the number of answers has an impact on performance. However, the coefficients on both the number of questions and answers should be interpreted carefully because, based on the panel analysis of the previous section, we know that these might not be robust to the inclusion of individual fixed effects.

The main concern with our instrument is that it might fail the exogeneity restriction. In particular, one might imagine that even if we do not include officer i's demand for information in the construction of our instrument, he might have had an effect on the cultural attitude of the branch regarding information sharing. However, this effect should be more pronounced for small branches than for larger branches.

Table 13 shows that even when we restrict attention to the subsample of branches with more than 50 loan officers, we find the same results. That is, the number of documents consulted and the number of answers provided positively affect performance, while the impact of the number of questions is negative. As expected given the lower number of observations, the estimates are significant only at the five- and ten-percent level in the most conservative specification shown in column (6). However, both the magnitudes and the signs are consistent with the previous results.

6 Discussion

We devote the next section to implications of our empirical findings for three strands of literature. In the first section, we consider theoretical and empirical results from tournament research that bear on information sharing incentives and thus group productivity. We then analyze human capital inefficiencies generated by the mandatory-transfer regulation. Finally, we interpret our findings in the context of existing studies on relational banking.

6.1 Information Sharing and Tournaments

If information sharing affects productivity, an organizationsl design question arises as to how promotion incentives interact with sharing incentives. Tournament theory, as

modeled by Lazear and Rosen (1981), models promotions as a *relative* game, that is, prizes depend on relative rather than absolute performance. The compensation at one level of the firm, in addition to motivating individuals at that level, motivates those at lower levels.

These basic ideas of tournament theory have been extended in numerous ways.[10] In particular, Dye (1984) and Lazear (1989) consider how the potential for collusion, sabotage, or other forms of non-cooperative behavior counter the incentive value generated by promotions and tournaments. Chan (1996) suggests that handicapping insiders in the tournament, with respect to external hires, can help mitigate the possibility of influence activity or sabotage. Then, the existing theoretical literature has recognized the costs and inefficiencies generated by the implementation of a promotion-based incentive system. Prendergast (1999) discusses the possibility that incentives from promotion methods give rise to dysfunctional behavioral responses and that companies adjust management compensation to address some of the negative responses from promotion incentives.

Empirically, many studies have confirmed these predictions from tournament theory. Using a survey of Australian firms, for example, Drago and Garvey (1998) show that individuals are less helpful and work harder when promotion incentives are strong. This seems to suggest that workers incentivized via promotions are less willing to cooperate with each other, because the effort to help others may reduce their own probability of being promoted.

In contrast to these observations, we find that loan officers intensively cooperate with each other, sharing their knowledge and in doing so they indirectly improve their colleagues' performance. This result seems to suggest that it is possible to reconcile competitive promotion incentives with cooperation in knowledge-sharing. We attribute this novel result to the relative weight placed on tenure in Japanese banking and to the dual role of posting questions and providing answers. On the one hand, sharing information with competitors can increase their chances of promotion via increased performance. On the other hand, helping others signals skill and is recognized and rewarded by the organization. This explains the strong correlations presented in Table 8.

6.2 The Cost of Anti-Corruption Regulation

Corruption is recognized as a driving factor in persistent poverty in less-developed countries but also a source of inefficiency and rent extraction in developed countries. The

[10] See Gibbons and Waldman (1999) for a survey of this literature.

World Bank ranks the fight against corruption as a top priority for poverty reduction (WorldBank 2006). Much theoretical work, since Becker and Stigler (1974) and Tirole (1986), focused on understanding the incentives and the constraints within corruptible bureaucracies.[11]

Existing evidence, however, on anti-corruption policies shows that corruption is fought and defeated with very simple tools. For example, Klitgaard (1991) describes successful cases of corruption elimination, such as in the Hong Kong Police Force and the Singapore Excise Department. The main factors were better monitoring and replacing individual bad actors. Similarly, Olken (2007) analyzes a randomized field experiment on reducing corruption in Indonesia suggesting that traditional top-down monitoring can play an important role in reducing corruption, even in a highly corrupt environment.

This gives rise to a more fundamental question: if these levers for eliminating corruption are within the choice set of governments, why are they not more often implemented? A possible answer is suggested by Acemoglu and Verdier (2000). They identify a trade-off between market failures and government failures. That is, since preventing all corruption is excessively costly, the second-best intervention may involve tolerating a certain fraction of bureaucratic corruption. Then, government failures may indicate an unavoidable price of dealing with market failures.

In our setting we do not directly observe corruption, but we are able to quantify the cost associated with the anti-corruption regulation. If the implementation of the rotation rule has been efficient, the social cost of enduring corruption should be higher than the banking inefficiencies generated by the remedial regulation. Then, we can infer the costs of corruption by analyzing the productivity loss associated with the officers' relocation. As a first step, we collected loan officers' performance in dollar terms. On average each officer generated almost four million dollars in bank gross profit every six months. Then, as is shown by Table 10, and holding other factors constant, switching reduces performance by more than 10%. This translates to a reduction in profits of 939,200 dollars a year for each officer transferred.

From this, we can conclude that in two years the Bank passed up more than 200 million dollars due to the adoption of this anti-corruption law. Since our data span only two years, we cannot analyze potential long term productivity gains that might arise from exposure to multiple branches. Yet, if corruption regulation has been optimally implemented, this estimate would constitute a reasonable bound on the amount of corruption avoided due to regulatory intervention.

[11] See Banerjee and Mullainathan (2009) for a recent survey of this literature.

6.3 Relational Banking

Boot (2000) defines relationship banking as the provision of financial services by a financial intermediary on the basis of long-term investment in obtaining firm-specific information through multiple interactions with diverse financial services. Banks are interested in relationship-based banking mainly for two reasons. First, the cost of information gathering is reduced by learning through repeated transactions. Second, financial contracts are typically incomplete: banks and customers can build commitment and reputation through repeated transactions across services. This banking model has been the predominant one in Japan.[12]

The possibility for a bank to build a long-term relationship with a client enables collection of soft information that is otherwise unavailable. The bond markets and the rating agencies collect financial disclosures, accounting reports, and default histories which can be considered hard information. Banks collect information on the client's ability and his honesty, which cannot be easily communicated to others Petersen and Rajan (1994).

Stein (2002) argues that larger, more hierarchical banks, where the decision maker is further from the information collector, are more likely to rely on hard information, because such organizations are expected to be less efficient at making relationship loans. Information in a large bank is potentially collected by one individual or group and a decision made by another. Thus the decisions must be made on information that is easy to transmit across physical or organizational distances. Consistent with this intuition, Berger et al. (2005) find that larger banks are more likely to lend to more customers at a greater distance and communicate with the borrower more impersonally, i.e. by mail or phone as opposed to face-to-face.

Our paper contributes to this literature in two ways. First, we show that information sharing among relationship managers allows a large bank, like the one analyzed here, to build long-term relationships with clients by reducing the costs of communicating the soft information collected. The platform used by the Bank allows officers to store the relevant information about diverse clients and to effectively communicate this to their colleagues. Second, we highlight the dark side of relationship banking, the possibility that loan officers could be bribed by clients. The bank has adopted the two-year-switching rule to prevent the relationship between clients and officers degenerating into suboptimal decisions from the point of view of the organization and society.

[12] See Hoshi et al. (1990) and Hoshi et al. (1991) for an empirical analysis of the the role of banks in Japan.

7 Concluding Remarks

To address the question of whether access to information produced and gathered within the same organization affects the performance of information workers, we examined two years of micro data from a major Japanese bank. Data include all accesses to an information-sharing platform, objective and subjective performance measures, and all promotions and job rotations among more than 2,800 loan officers. Exogenous legal requirements, aimed at curbing corruption by compulsory office rotation, permit analysis of loan officer performance in different settings.

We find that a standard deviation increase in the number of shared documents predicts an 11% rise in output, in specifications with time, location and individual fixed effects. This is comparable to just under one year of education among non-manufacturing workers Black and Lynch (1996). Quantile regression estimates suggest that communication and information sharing greatly benefits the low-performance officers. In fact, questions are more salient among workers of lesser ability and, when productivity gains exist, they appear strongly on the left tail of the distribution, at the 10^{th} and 25^{th} percentiles, but do not appear on the right tail, at the 75^{th} and 90^{th} percentiles.

We observe over 200 instances of job promotion. Tenure within the bank is by far the strongest predictor, while lagged productivity is not significant. Loan officer answers are associated with faster promotion while loan officer questions are associated with slower promotion. These findings are consistent with a signaling hypothesis on behalf of high-ability workers, who distinguish themselves from low-ability workers by providing answers in order to demonstrate their expertise and signal higher ability.

We also observe 618 instances of loan officers switching branches as a result of an anti-corruption law. This exogenous shock provides an opportunity to observe the same knowledge worker in different contexts. With this instrument, and controlling for unobserved heterogeneity over time and branch, a standard deviation increase in shared document consumption boosts productivity by at least 10%. These results appear to be causal.

The difference between OLS and IV specifications provides evidence of the substitution hypothesis: high-ability officers demand less information than low-ability officers, while low-ability officers can compensate for low independent performance by consuming information provided by others. We also find that switching jobs significantly reduces overall performance, possibly indicating destruction of job specific human capital. The anti-corruption law should therefore avoid economic losses from graft of at least this magnitude. Interestingly, officers of all abilities increase their demand for information

on switching jobs.

A range of potential extensions is left for future research. It is important to understand how information sharing and communication are related to the incentive system in place. One could address this question by developing a similar analysis on micro data from a major U.S. bank, which should clarify if end-of-the-year bonuses and an incentive system heavily based upon performance affects the information shared between white collar workers. Moreover, it would be interesting to understand whether it is possible to increase the flow of information within the organization by explicitly relating information consumption to monetary incentives. Finally, we are also interested in the robustness of our results with respect to the possibility of anonymously seeking and providing information. On one hand, this could reduce embarrassment costs in requesting information and increase the provision of novel but controversial ideas. On the other hand, the quality of information could decrease due to a lower signaling and reputation-building effects. This would enable tests of how incentives and reputation interact with organizational theories of the firm.

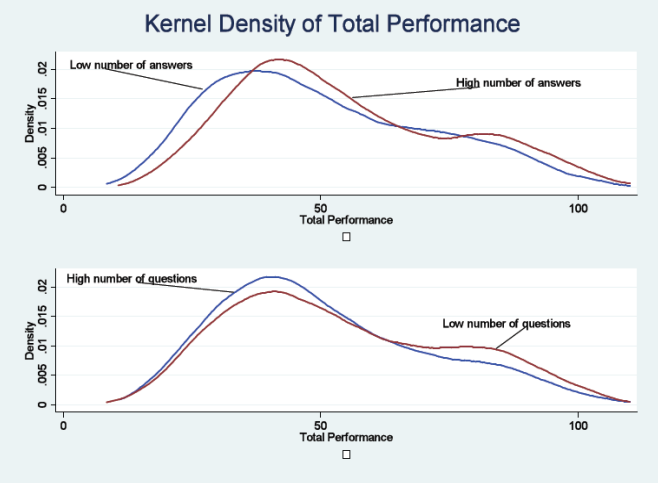

Figure 1: The plot displays the kernel density estimation of Total Performance for managers that provided an above (below) the median number of questions and answers. The density estimates are calculated using an Epanechnikov kernel.

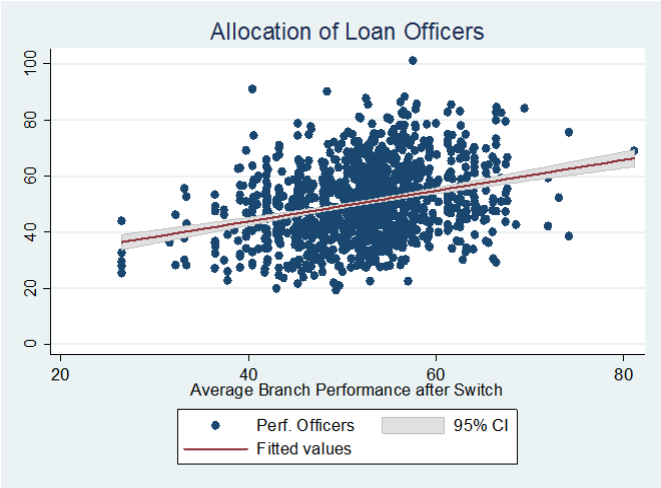

Figure 2: The plot displays the relationship between the officers' average performance in the term before the relocation to another branch, and the average productivity of the branch to where they are relocated. The fitted values and the 95% confidence interval are reported.

Table 3: Summary Statistics

	Mean	Median	Min	Max	Std. Dev.
Total Performance	52	48	8	110	21
Number Documents	569	522	102	1618	282
Number Questions	77	45	4	585	123
Number Answers	250	238	4	813	164
Tenure	10	11	0	32	5

Notes: entries are the summary statistics for our main variables of interest. "Total Performance" is the sum of the objective and the subjective performance measures. "Tenure" is the number of years managers worked for the Bank. On average each manager obtains a score of 52 out of 100, downloads 569 documents, post 77 questions, and provides 250 answers over a six-month period. Overall there are 2451 manager-branch-term observations.

Table 4: Managers Performance by Information Sharing

	Number Documents	Number Answers	Number Questions
Below the Median	52.603	50.356	53.919
	(0.291)	(0.294)	(0.308)
Above the Median	52.025	54.438	50.727
	(0.306)	(0.306)	(0.287)
Difference	0.577	4.085***	-3.192***
	(0.422)	(0.420)	(0.421)

Notes: *** denotes significance at 1%, ** at 5%, and * at 10%. The standard errors clustered by manager are reported in parenthesis. Performance is measured as the total score assigned to a manager in a given branch.

Table 5: Panel Model – Fixed Effects Estimates

Log(Tot. Performance)	(1)	(2)	(3)	(4)	(5)
Log(Number Documents)	0.0493***	0.0562***	0.0226**	0.0300***	0.0216**
	(0.010)	(0.011)	(0.010)	(0.011)	(0.011)
Log(Number Questions)	-0.0505***	-0.0335***	-0.0246***	-0.0217***	-0.0195***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
Log(Number Answers)	0.0508***	0.00848	0.00533	0.00337	0.00475
	(0.003)	(0.006)	(0.006)	(0.006)	(0.007)
Log(Tenure)	0.0632***	0.0674***	0.105***	0.100***	0.1000***
	(0.009)	(0.009)	(0.008)	(0.008)	(0.008)
College	0.108***	0.113***	0.105***	0.0924***	0.0891***
	(0.028)	(0.028)	(0.026)	(0.026)	(0.025)
Time Fixed Effects		YES	YES	YES	YES
Group Fixed Effects			YES	YES	YES
Branch Fixed Effects				YES	YES
Time * Branch Fixed Effects					YES
Observations	9,805	9,805	9,805	9,805	9,805
R-squared	0.0629	0.0582	0.1567	0.3049	0.467

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. The dependent variable in all columns is the log of total performance. "College" is a dummy variable equal to one if the manager joined the Bank directly after college. The time period is 2006-2008. The estimation method in all columns is OLS. Standard errors in brackets under coefficients in all columns are clustered by individual (i.e. robust to heteroskedasticity and autocorrelation of unknown form). Columns include a full set of time, group and branch fixed effects. As additional controls Column (5) include time dummies interacted with a branch dummies.

Table 6: Panel Model – Individual FE Estimates

Log(Tot. Performance)	(1)	(2)	(3)	(4)	(5)
Log(Number Documents)	0.0480***	0.0430***	0.0467***	0.0292*	0.0374**
	(0.010)	(0.016)	(0.017)	(0.017)	(0.018)
Log(Number Questions)	-0.0543***	-0.0408***	-0.0176**	-0.0153**	-0.0104
	(0.005)	(0.007)	(0.007)	(0.007)	(0.007)
Log(Number Answers)	0.0511***	0.0519***	-0.00604	-0.00668	-0.0062
	(0.003)	(0.003)	(0.007)	(0.007)	(0.007)
Managers Fixed Effects		YES	YES	YES	YES
Time Fixed Effects			YES	YES	YES
Group Fixed Effects				YES	YES
Branch Fixed Effects					YES
Observations	9,806	9,806	9,806	9,806	9,806
R-squared	0.032	0.048	0.059	0.081	0.157

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. The dependent variable in all columns is the log of total performance. The time period is 2006-2008. The estimation method in all columns is OLS. Standard errors in brackets under coefficients in all columns are clustered by individual (i.e. robust to heteroskedasticity and autocorrelation of unknown form). Columns include a full set of time, group and branch fixed effects. As additional controls Column (2)-(5) include managers fixed effects.

Table 7: Quantile Regression Estimates

Log(Total Performance)	10th	25th	50th	75th	90th
Log(Number Documents)	0.0391**	0.0409***	0.0250*	0.002	-0.012
	(0.017)	(0.012)	(0.014)	(0.013)	(0.010)
Log(Number Questions)	-0.0207***	-0.0251***	-0.0372***	-0.0361***	-0.0164***
	(0.007)	(0.008)	(0.007)	(0.006)	(0.006)
Log(Number Answers)	0.004	0.008	0.013	0.0209**	0.006
	(0.010)	(0.012)	(0.012)	(0.009)	(0.007)
Observations	9,805	9,805	9,805	9,805	9,805

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. The dependent variable in all columns is the log of total performance. All specifications control for time, group, and branch fixed effects.

Table 8: Effect on Promotion Probability

Promotion (Probit)	(1)	(2)	(3)	(4)	(5)
Number Questions		-0.00172**	-0.00174**	-0.00154**	-0.00245**
		(0.001)	(0.001)	(0.001)	(0.001)
Number Answers		0.00106***	0.00107***	0.00106***	0.00153***
		(2.88E-04)	(2.87E-04)	(0.001)	(0.001)
Lag Productivity	0.00067	0.00194	0.00177	0.00152	0.00146
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Tenure	0.096***	0.107***	0.107***	0.108***	0.148***
	(0.007)	(0.010)	(0.009)	(0.009)	(0.016)
College	0.1608	-0.227	-0.229	-0.22	-0.294
	(0.294)	(0.254)	(0.254)	(0.258)	(0.267)
Time Fixed Effects			YES	YES	YES
Group Fixed Effects				YES	YES
Branch Fixed Effects					YES
Observations	6,971	6,971	6,971	6,971	6,971

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. All columns estimated by probit maximum likelihood. The dependent variable in all columns is a dummy variable equal to one if the loan officer is promoted. All specifications include "Number Documents" which is never significant. "Lag Productivity" is the managers' total performance recorded in the previous term. We include all the observations for which we have data on the managers' position inside the Bank. All columns include "Tenure" as control variable. "College" is a dummy variable equal to one if the manager joined the Bank directly after college. Additional controls include time, group, and branch fixed effects.

Table 9: Effect of Switching on Information Access

Log(Number Documents)	(1)	(2)	(3)	(4)
Switch	0.0199**	0.0218**		
	(0.009)	(0.009)		
After Switch			0.0125	0.0150*
			(0.008)	(0.008)
Log(Tenure)	-0.0218		-0.0217	
	(0.014)		(0.014)	
College	0.278***		0.279***	
	(0.042)		(0.042)	
Time Fixed Effects	YES	YES	YES	YES
Group Fixed Effects	YES	YES	YES	YES
Branch Fixed Effects	YES	YES	YES	YES
Individual Fixed Effects		YES		YES
R-squared	0.3342	0.347	0.334	0.346
Observations	10,055	10,055	10,055	10,055

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. The dependent variable in all columns is the log number of documents downloaded in a six-month period. The estimation method in all columns is OLS. Standard errors in brackets under coefficients in all columns are clustered by individual (i.e. robust to heteroskedasticity and autocorrelation of unknown form). Columns include a full set of time, group and branch fixed effects. As additional controls Column (2) and (4) include individual fixed effects. "Switch" is a dummy variable equal to one only in the first term after the manager's relocation to another branch. "After Switch" equals one for all the terms after the transfer. "College" is a dummy variable equal to one if the manager joined the Bank directly after college.

Table 10: Effect of Switching on Performance

Log(Total Performance)	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Switch	-0.111***	-0.0317	-0.118***	-0.0461*		-0.108***	
	(0.012)	(0.027)	(0.012)	(0.027)		(0.013)	
Switch*Tenure		-0.00778***		-0.00726***			
		(0.002)		(0.002)			
After Switch					-0.0889***		-0.0915***
					(0.010)		(0.012)
Log(Number Documents)	0.0470***	0.0469***	0.0287***	0.0283***	0.0274**	0.0387**	0.0348**
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.018)	(0.018)
Log(Number Questions)	-0.0515***	-0.0516***	-0.0210***	-0.0210***	-0.0205***	-0.00851	-0.0078
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.007)	(0.007)
Log(Number Answers)	0.0569***	0.0571***	0.0059	0.00618	0.00476	-0.00394	-0.0052
	(0.003)	(0.003)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
Log(Tenure)	0.0636***	0.0700***	0.0994***	0.105***	0.0994***		
	(0.009)	(0.009)	(0.008)	(0.008)	(0.008)		
College	0.106***	0.104***	0.0915***	0.0907***	0.0897***		
	(0.027)	(0.028)	(0.026)	(0.026)	(0.026)		
Time Fixed Effects			YES	YES	YES	YES	YES
Group Fixed Effects			YES	YES	YES	YES	YES
Branch Fixed Effects			YES	YES	YES	YES	YES
Manager Fixed Effects						YES	YES
Observations	9,805	9,805	9,805	9,805	9,805	9,805	9,805
R-squared	0.0698	0.0701	0.3078	0.308	0.3054	0.166	0.165

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. The dependent variable in all columns is the log of total performance. The estimation method in all columns is OLS. Standard errors in brackets under coefficients in all columns are clustered by individual (i.e. robust to heteroskedasticity and autocorrelation of unknown form). Columns include a full set of time, group and branch fixed effects. As additional controls Column (6) and (7) include individual fixed effects. "Switch" is a dummy variable equal to one only in the first term after the manager's relocation to another branch. An interaction term between "Switch" and "Tenure" is included in columns (2) and (4). "After Switch" equals one for all the terms after the transfer. "College" is a dummy variable equal to one if the manager joined the Bank directly after college.

Table 11: First Stages

	Documents	Answers	Questions
Documents Prev Branch	7.398***	-0.654*	-0.214
	(0.744)	(0.362)	(0.248)
Answers Prev Branch	-14.439***	4.190***	-1.637**
	(2.282)	(1.112)	(0.757)
Questions Prev Branch	8.420***	2.271**	7.437***
	(2.157)	(1.052)	(0.718)
Tenure	-10.419***	-4.132***	-2.688***
	(2.622)	(1.278)	(0.873)
College	54.021	19.242	-2.409
	(69.580)	(33.915)	(23.161)
F-Test	50.289	16.774	49.343
Time Fixed effects	YES	YES	YES
Group Fixed Effects	YES	YES	YES
Observations	618	618	618

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. The dependent variables are the number of documents, answers and questions in a six-month period after the relocation. Standard errors in brackets under coefficients in all columns are clustered by individual (i.e. robust to heteroskedasticity and autocorrelation of unknown form). Columns include a full set of time, and group fixed effects. "Documents Prev Branch" is the average number of documents consulted within the branch before the relocation. Similarly for "Answers Prev Branch" and "Questions Prev Branch". "College" is a dummy variable equal to one if the manager joined the Bank directly after college. Overall we observe in our sample 618 managers switching branch.

Table 12: IV Estimates

Total Performance	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV
Number Documents	-0.00123	0.0197**	-0.00138	0.0225***	-0.00347	0.0194**
	(0.003)	(0.008)	(0.003)	(0.008)	(0.003)	(0.008)
Number Questions	-0.0306***	-0.133**	-0.0287***	-0.104***	-0.0169**	-0.0908***
	(0.008)	(0.061)	(0.008)	(0.034)	(0.008)	(0.033)
Number Answers	0.0175***	0.112	0.0148**	0.0727**	0.0114*	0.0674**
	(0.007)	(0.080)	(0.007)	(0.036)	(0.007)	(0.034)
Tenure	0.387**	0.871**	0.385**	0.689***	0.603***	0.819***
	(0.172)	(0.382)	(0.174)	(0.242)	(0.183)	(0.223)
College	-3.736	-7.538	-3.714	-6.748	-3.414	-6.015
	(4.809)	(6.436)	(4.806)	(5.74)	(4.839)	(5.544)
Time Fixed Effects			YES	YES	YES	YES
Group Fixed Effects					YES	YES
Observations	618	618	618	618	618	618

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. In all columns the dependent variable is the manager's total performance. "College" is a dummy variable equal to one if the manager joined the Bank directly after college. Standard errors in brackets under coefficients in all columns are clustered by individual (i.e. robust to heteroskedasticity and autocorrelation of unknown form). Columns include a full set of time, and group fixed effects. The estimation method in columns (1), (3), and (5) is OLS. Columns (2), (4), and (6) are estimated using 2SLS. In both cases, we restrict attention to our subsample of 618 managers switching branch.

Notice that, albeit significant, we do *not* interpret causally the effect of the "number of questions" and the "number of answers" on performance. In fact, as shown by the results in table 6, these effects would disappear if individual fixed effects were added to the specification.

Table 13: IV Estimates: Large Branches

Total Performance	(1) OLS	(2) IV	(3) OLS	(4) IV	(5) OLS	(6) IV
Number Documents	0.007	0.026	0.006	0.0129*	0.004	0.0156**
	(0.005)	(0.042)	(0.005)	(0.007)	(0.005)	(0.007)
Number Questions	-0.0455***	0.224	-0.0384***	-0.0422**	-0.0199**	-0.0377*
	(0.009)	(0.683)	(0.009)	(0.019)	(0.010)	(0.021)
Number Answers	0.0316***	-0.387	0.019	0.020	0.011	0.0392*
	(0.011)	(1.049)	(0.013)	(0.025)	(0.013)	(0.021)
Tenure	0.243	-0.039	0.211	0.269	0.603**	0.695**
	(0.272)	(1.100)	(0.278)	(0.269)	(0.282)	(0.286)
College	-25.25***	-13.650	-22.65***	-24.11***	-19.94***	-23.44***
	(7.129)	(33.150)	(6.421)	(6.669)	(7.420)	(8.608)
Time Fixed Effects			YES	YES	YES	YES
Group Fixed Effects					YES	YES
Observations	240	240	240	240	240	240

Notes: * significant at 10%; ** significant at 5%; *** significant at 1%. In all columns the dependent variable is the manager's total performance. Large branches are those with more than 50 employees. "College" is a dummy variable equal to one if the manager joined the Bank directly after college. Standard errors in brackets under coefficients in all columns are clustered by individual (i.e. robust to heteroskedasticity and autocorrelation of unknown form). Columns include a full set of time, and group fixed effects. The estimation method in columns (1), (3), and (5) is OLS. Columns (2), (4), and (6) are estimated using 2SLS. In both cases, we restrict attention to our subsample of 240 managers switching branch.

Notice that, albeit significant, we do *not* interpret causally the effect of the "number of questions" on performance. In fact, as shown by the results in table 6, this effect would disappear if individual fixed effects were added to the specification.

References

Abraham, A. and K. White (2006, July). The dynamics of plant-level productivity in u.s. manufacturing. Working Papers 06-20, Center for Economic Studies, U.S. Census Bureau.

Acemoglu, D. and T. Verdier (2000, March). The choice between market failures and corruption. *American Economic Review 90*(1), 194–211.

Aral, S., E. Brynjolfsson, and M. V. Alstyne (2007, June). Information, technology and information worker productivity: Task level evidence. NBER Working Papers 13172, National Bureau of Economic Research, Inc.

Arrow, K. (1974). *The limits of organization*. Norton New York.

Bandiera, O., I. Barankay, and I. Rasul (2009, 07). Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica 77*(4), 1047–1094.

Banerjee, Abhijit, R. H. and S. Mullainathan (2009). Corruption. *mimeo*.

Bartel, A., C. Ichniowski, and K. Shaw (2007, November). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The Quarterly Journal of Economics 122*(4), 1721–1758.

Bartelsman, E. J. and M. Doms (2000, September). Understanding productivity: Lessons from longitudinal microdata. *Journal of Economic Literature 38*(3), 569–594.

Becker, G. S. and G. J. Stigler (1974, January). Law enforcement, malfeasance, and compensation of enforcers. *Journal of Legal Studies 3*(1), 1–18.

Berger, A. N., N. H. Miller, M. A. Petersen, R. G. Rajan, and J. C. Stein (2005, May). Does function follow organizational form? evidence from the lending practices of large and small banks. *Journal of Financial Economics 76*(2), 237–269.

Black, S. and L. Lynch (1996). Human-capital investments and productivity. *The American Economic Review 86*(2), 263–267.

Black, S. E. and L. M. Lynch (2001, August). How to compete: The impact of workplace practices and information technology on productivity. *The Review of Economics and Statistics 83*(3), 434–445.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2011). Does management matter? evidence from india. Technical report, National Bureau of Economic Research.

Bloom, N., L. Garicano, R. Sadun, and J. V. Reenen (2009, May). The distinct effects of information technology and communication technology on firm organization. CEP Discussion Papers dp0927, Centre for Economic Performance, LSE.

Bloom, N. and J. V. Reenen (2007, November). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics 122*(4), 1351–1408.

Bloom, N. and J. V. Reenen (2010, Winter). Why do management practices differ across firms and countries? *Journal of Economic Perspectives 24*(1), 203–24.

Bolton, P. and M. Dewatripont (1994, November). The firm as a communication network. *The Quarterly Journal of Economics 109*(4), 809–39.

Boot, A. W. A. (2000, January). Relationship banking: What do we know? *Journal of Financial Intermediation 9*(1), 7–25.

Brynjolfsson, E., A. McAfee, M. Sorell, and F. Zhu (2007). Scale without mass: business process replication and industry dynamics. *Proceedings* (Nov).

Chan, W. (1996, October). External recruitment versus internal promotion. *Journal of Labor Economics 14*(4), 555–70.

Conley, T. and C. Udry (2010). Learning about a new technology: Pineapple in ghana. *The American Economic Review 100*(1), 35–69.

Drago, R. and G. T. Garvey (1998, January). Incentives for helping on the job: Theory and evidence. *Journal of Labor Economics 16*(1), 1–25.

Duflo, E. and E. Saez (2003). The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quarterly Journal of Economics 118*(August), 815–842.

Dye, R. A. (1984, January). The trouble with tournaments. *Economic Inquiry 22*(1), 147–49.

Faggio, G., K. G. Salvanes, and J. V. Reenen (2010, December). The evolution of inequality in productivity and wages: panel data evidence. *Industrial and Corporate Change 19*(6), 1919–1951.

Foster, A. and M. Rosenzweig (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 1176–1209.

Foster, L., J. Haltiwanger, and C. Syverson (2008, March). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review 98*(1), 394–425.

Garicano, L. (2000, October). Hierarchies and the organization of knowledge in production. *Journal of Political Economy 108*(5), 874–904.

Garicano, L. and E. Rossi-Hansberg (2006, November). Organization and inequality in a knowledge economy. *The Quarterly Journal of Economics 121*(4), 1383–1435.

Gibbons, R. (2010). Inside organizations: Pricing, politics, and path dependence. *Annual Review of Economics 2*.

Gibbons, R. and R. Henderson (2010). Relational Contracts and the Origins of Organizational Capabilities. *Handbook of Orgaizational Economics forthcoming*.

Gibbons, R. and M. Waldman (1999). Careers in organizations: Theory and evidence. *Handbook of labor economics 3*, 2373–2437.

Hamermesh, D. (1987). The costs of worker displacement. *The Quarterly Journal of Economics 102*(1), 51–75.

Hayek, F. (1945). The use of knowledge in society. *The American Economic Review*, 519–530.

Hermalin, B. (2007). Leadership and corporate culture. *Draft of chapter for forthcoming Handbook of Organiza*.

Hertzberg, A., J. M. Liberti, and D. Paravisini (2010, 06). Information and incentives inside the firm: Evidence from loan officer rotation. *Journal of Finance 65*(3), 795–828.

Hoshi, T., A. Kashyap, and D. Scharfstein (1990). The role of banks in reducing the costs of financial distress in japan. *Journal of financial economics 27*(1), 67–88.

Hoshi, T., A. Kashyap, and D. Scharfstein (1991). Corporate structure, liquidity, and investment: Evidence from japanese industrial groups. *Quarterly Journal of Business and Economics 106*(1).

Ichniowski, C., K. Shaw, and G. Prennushi (1997). The effects of human resource management practices on productivity: A study of steel finishing lines. *The American Economic Review 87*(3), 291–313.

Jacobson, L.S., L. R. and D. Sullivan (1993). Earnings losses of displaced workers. *The American Economic Review*, 685–709.

Jaffe, A., M. Trajtenberg, and R. Henderson (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *the Quarterly journal of Economics 108*(3), 577.

Kletzer, L. (1998). Job displacement. *The Journal of Economic Perspectives 12*(1), 115–136.

Klitgaard, R. (1991). *Controlling corruption*. Univ of California Pr.

Kremer, M. (1993). The o-ring theory of economic development. *The Quarterly Journal of Economics 108*(3), 551.

Kremer, M. and E. Miguel (2007). The illusion of sustainability. *Quarterly Journal of Economics 122*(3).

Lazear, E. (1989). Pay equality and industrial politics. *Journal of Political Economy 97*(3), 561–580.

Lazear, E. P. (2000, December). Performance pay and productivity. *American Economic Review 90*(5), 1346–1361.

Lazear, E. P. and S. Rosen (1981, October). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy 89*(5), 841–64.

Marschak, J. and R. Radner (1972). *The Economic Theory of Teams*.

Moretti, E. (2011). Social learning and peer effects in consumption: Evidence from movie sales. *The Review of Economic Studies 78*(1), 356.

Niehaus, P. (2011). Filtered social learning. *Journal of Political Economy 119*(4), 686–720.

Olken, B. (2007). Monitoring Corruption: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy 115*, 200–249.

Petersen, M. and R. Rajan (1994). The benefits of lending relationships: Evidence from small business data. *Journal of Finance 49*(1), 3–37.

Prendergast, C. (1999). The provision of incentives in firms. *Journal of economic literature 37*(1), 7–63.

Radner, R. (1993, September). The organization of decentralized information processing. *Econometrica 61*(5), 1109–46.

Ruhm, C. (1991). Are workers permanently scarred by job displacements. *The American Economic Review 81*(1), 319–324.

Sah, R. K. and J. E. Stiglitz (1986, September). The architecture of economic systems: Hierarchies and polyarchies. *American Economic Review 76*(4), 716–27.

Stein, J. (2002). Information production and capital allocation: Decentralized versus hierarchical firms. *The Journal of Finance 57*(5), 1891–1921.

Syverson, C. (2004). Product substitutability and productivity dispersion. *Review of Economics and Statistics 86*(2), 534–550.

Syverson, C. (2010). What Determines Productivity? *NBER Working Paper*.

Tirole, J. (1986, Fall). Hierarchies and bureaucracies: On the role of collusion in organizations. *Journal of Law, Economics and Organization 2*(2), 181–214.

Van Zandt, T. (1999, July). Real-time decentralized information processing as a model of organizations with boundedly rational agents. *Review of Economic Studies 66*(3), 633–58.

Van Zandt, T. (2004, Autumn). Information overload in a network of targeted communication. *RAND Journal of Economics 35*(3), 542–560.

WorldBank (2006). Annual report 2006. *Discussion paper, The World Bank Group*.