

INFORMATION AND INCENTIVES INSIDE THE FIRM: EVIDENCE FROM LOAN OFFICER ROTATION

ANDREW HERTZBERG

JOSE MARIA LIBERTI

DANIEL PARAVISINI*

March 2007

Abstract

This paper provides evidence that a bank policy that routinely reassigned loan officers to different borrowers acts as an incentive device. We argue that the new loan officer who is assigned to the task has no reputation incentive to hide bad information. Therefore, the threat of rotation induces the incumbent to reveal bad news so as to avoid being uncovered by her successor. Using a proprietary monthly panel of internal risk rating data, we show that a three-year rotation rule induces incumbent loan officers to reveal negative information about the creditworthiness of the firms they manage. Consistent with our theory, loan officers systematically downgrade firms leading up to the three-year rule, and these downgrades are informative about the future probability of default. In line with our reputation concerns argument, we document that loan officers who fail to report bad news and are subsequently exposed by a successor, later go on to manage smaller sized lending portfolios. Finally we show that loan officer rotation impacts upon the capital allocation decisions of the bank.

*Kellogg School of Management, Northwestern University; Kellogg School of Management, Northwestern University and Columbia Business School respectively. E-mails: a-hertzberg@kellogg.northwestern.edu, j-liberti@kellogg.northwestern.edu and dp2239@columbia.edu. We wish to thank Patrick Bolton, Doug Diamond, Anil Kashyap, David Matsa, Jonah Rockoff and seminar participants at Northwestern University (Kellogg), Columbia Business School and Boston University for helpful comments and suggestions. The superb research assistance of Michael Nicssner and Ronald Chan is greatly appreciated. We specially thank Germán Blanco for providing his time and patience in answering an infinite number questions enabling us to assemble this database. All errors are our own.

I. Introduction

This paper provides evidence that a bank policy that routinely reassigns loan officers to different borrowers acts as an incentive device. We find that rotation induces loan officers to reveal negative information about the creditworthiness of the borrowers they manage. Several papers provide a theoretical account of why an organization might choose to adopt a policy of reassigning tasks among workers (Holmstrom (1982), Meyer (1994), and Prescott and Townsend (2006)). This paper provides a new explanation for this phenomenon and provides the first detailed empirical account of the way that rotation increases the incentive of a delegated monitor to reveal private information.

It has long been recognized that agency problems constrain the practice of delegated monitoring (Aghion and Tirole (1997), Stein (2002)). The agency problem stems from the fact that the monitor has access to information which the principal cannot obtain directly. We consider an environment where a loan officer performs a dual role: she is responsible for the state of the loan to which she has been assigned (active monitoring) as well as obtaining and reporting information about the state of the loan (passive monitoring). As a result of this dual role the loan officer faces a trade-off when deciding what news to report to the bank. She may hide unfavorable information about the state of a loan because it will reflect poorly on her active monitoring performance.

We argue that rotation can reduce this incentive to hide information by temporarily separating the active and passive monitoring roles. A newly assigned officer is more willing to reveal bad news about a loan because this does not reflect poorly upon her own abilities. On the contrary, she demonstrates her ability to detect bad information. As a result, an incumbent loan officer has reduced incentives to conceal bad news, since she recognizes that the information is likely to be released in the near future by her successor. When a loan is downgraded soon after rotation, the bank can infer that the previous officer was hiding bad news.

We begin by documenting a three-year loan officer rotation policy of a large multi-national commercial bank. The bank-wide policy recommends that firms are reassigned after being under the management of the same loan officer for three years. Using hand collected panel data from the bank's operations in Argentina between 1997 and 2004, we follow the monthly assignment of firms to loan officers. We find that 89% of relationships last less than three years and, conditional on having managed a firm for 34 months, a loan officer is rotated with a 58% probability within the next three months.

Using the three-year rule as an exogenous source of variation in the probability of rotation, we

assess empirically the effect of rotation on the information reporting behavior of loan officers. Loan officers collect information about the firms they oversee and assign risk ratings that are recorded internally in the bank. Using monthly data on these risk ratings we show that loan officers on average downgrade firms under their management during the four quarters leading to the three-year rule limit.

We verify that the observed patterns in risk ratings are not induced by time varying firm specific shocks. We match the firms in the bank database with a Central Bank Credit Registry that contains both the amount of loans outstanding and risk ratings assigned by all formal financial intermediaries to every borrower. We show that the risk ratings and lending outcomes of the borrowers with other financial institutions do not vary systematically around the three-year rule. Matching to other banks' ratings also allows us to show that the rating applied by the bank is systematically optimistic twelve months prior to the rule. The systematic downgrading of firms leading up to the three year rule undoes this bias.

The evidence is consistent with the hypothesis that rotation increases the incentives of a loan officer to reveal bad news about the firms under management. To further confirm this hypothesis, we show that the downgrades induced by rotation are informative about the state of the loans. Using actual loan performance data, we measure informativeness of ratings as their ability to predict future default. Using the three-year rule to instrument for rating changes, we find that the downgrades induced by rotation significantly predict future default rates.

This paper is related to a literature that stresses the role that career concerns can have in shaping asset allocation decisions. Most closely related is Rajan (1994), who argues that bankers may choose to hide poorly performing loans so as to preserve the labor market's public assessment of their reputation. As in our model, he proposes that a banker's reputation will fall if the market discovers that they have presided over poorly performing loans. Our results show that the threat of rotation leads to stronger downgrades by loan officer who have originated above average amount of loans in the course of the relationship. This finding is consistent with our theory, which suggests that a loan officer will have the strongest incentive to hide bad information when she bears a high responsibility for the current state of the loan.

To further corroborate the career concerns assumption in our empirical context, we show that the volume of assets under management of a loan officer is increasing in measures of her reputation. We construct two new indexes of reputation based on the inferences that the bank draws when a firm is

downgraded right after a loan officer rotation. Under our theory this event hurts the reputation of the previous loan officer because it indicates that she was managing a low quality loan and failed to detect and report it. Conversely, the event improves the reputation of the new loan officer because she demonstrates ability to detect bad performance and bears minimal responsibility for this performance. We show that loan officers who accumulate more positive (negative) reputational events go on to manage larger (smaller) loan portfolios.

Finally we document that loan officer rotation impacts upon the capital allocation decisions of the bank. We conjecture that on average the bank will approve more lending when the precision of the information it obtains from loan officers increases. Consistent with this hypothesis, the results indicate that the total amount of lending increases leading up to the three-year rule, at the time when rotation induces more informative ratings. Again, this result holds after controlling for firm specific increases in the demand for credit by obtaining estimates relative to the total amount of lending at all other banks. We also show evidence that this lending expansion is directed to the firms with the best repayment prospects in the cross section.

Our paper is the first to provide an empirical account of rotation within an organization. The closest empirical support comes from Osterman (1994) and (2000) who provides workplace survey evidence that job rotation is a widely used practice in many firms.¹ The accounting literature, largely motivated by the five year mandated audit partner rotation included in Section 203 of the Sarbanes-Oxley Act of 2002, has provided mixed evidence on the role that auditor rotation has on the quality of auditing services.²

A number of other papers have presented theories which argue that rotation can be used to elicit information from workers. Prescott and Townsend (2006) provide a model in which rotation can be beneficial because it allows a principal to contract with an agent who does not have an information advantage about the task to which they are assigned. Several papers argue that organizations use rotation as a credible commitment not to use the agent's private information against her in the future (Prescott and Townsend (2006), Arya and Mittendorf (2004), Hirao (1993)). Ickes and Samuelson (1987) present this argument as a solution to the *ratchet effect*, whereby a worker who anticipates

¹This evidence refers to workers being reassigned across functional areas of a firm and hence goes beyond the type of rotation that fits with that described in most of the theoretical literature. For comparison, the loan officers in our sample would not be classified as engaging in job rotation since they remain as loan officers to small and medium sized borrowers over their entire time in the sample.

²See for example Arrunada and Paz-Ares (1997), Dopuch King and Schwartz (2001), and Myer, Myer, and Omer (2003).

staying in the same task for a long period of time chooses to underperform in order to convince the principal that the task has low productivity. This leads the principal to set less demanding remuneration schemes in the future.

The explanation for rotation in this paper differs from existing theories by stressing that the new agent who is assigned to the task has no incentive to hide bad information. Moreover the threat of rotation induces the incumbent agent to reveal bad news so as to avoid being uncovered by her successor. In our story, absent rotation, the agent who leaves the task would prefer to hide bad news from the principal since this reflects poorly upon her type. By contrast, existing theories of rotation rely on the fact that the departing agents bear no consequence when they admit poor performance at the end of the assignment. This seems unlikely to be a good description of an environment where the past performance of agents impacts upon their future career, as our evidence suggests.

Our theory is closest in spirit to an explanation for job rotation discussed in Holmstrom (1982). He argues that job rotation can be used to facilitate relative performance evaluation when workers perform heterogeneous tasks of unknown productivity. We focus on the principal's threat to use information that will be released by future agents to gauge the performance of the previous agent on the job. More broadly this research is related to a larger literature of papers that study organizational forms designed to elicit information from delegated monitors. Aghion and Tirole (1997) study how the allocation of formal authority can be used to solve this problem. Stein (2002) shows how different hierarchical structures provide the best incentives for the collection and transmission of soft and hard information.

The rest of this paper proceeds as follows. Section 2 sets out our theoretical explanation for the way that rotation can be used to induce loan officers to reveal private information about the loans to which they have been assigned. We form a set of empirical predictions to test our theory. Section 3 describes the data and sets out our identification strategy. We also use this section to document our motivating fact: the routine use of loan officer rotation within the bank. Section 4 presents our key empirical results. Most importantly we show that the threat of rotation induces loan officers to release bad news about the borrowers to which they have been assigned. We also provide evidence which corroborates our supposition that the ratings that are issued by loan officers affect their future career in the bank. Section 5 documents the effect of rotation on lending which allows us to conclude that the patterns in ratings which we document do in fact affect real allocation decisions within the bank. Section 6 provides a brief conclusion and discussion.

II. Theoretical Framework

In this section we build a stylized model of delegated monitoring to demonstrate that rotation can be used to increase a monitor’s incentives to truthfully reveal her private information about the projects to which she is assigned. The particular application we have in mind is a commercial bank (the principal) that delegates the monitoring of corporate lending relationships to loan officers (the monitors). We abstract from many of the features of a lending relationship and focus on the loan officer’s role in collecting and transmitting information about the loan to her superiors. We presume that these reports are used by the bank in the process of approving the terms of lending to each client. We do not model the use of this information explicitly but rather assume that the bank prefers more information to less.

A. Set-Up

Consider a bank that must assign a loan officer to monitor a loan. We suppose that there are two periods (denoted $t = 1, 2$) and that the bank must assign one loan officer to the loan in each period.³ The bank has two officers labelled a and b who can be assigned to the loan. The bank must commit in advance to an allocation policy. The same officer can be assigned in each period $\{a, a\}$, we will refer to this as “no rotation”. Alternately, the loan can be handled by a different officer in each period $\{a, b\}$, we will refer to this as “rotation”.⁴ We identify conditions under which the bank strictly prefers rotation in the sense that it elicits more information from the officer assigned to the loan in each period. We assume that loan officers have heterogeneous skill. To capture this suppose that each officer can be either of high or low type denoted by $i \in \{h, l\}$. Alternately, a loan officer’s type can be interpreted more broadly as her willingness to collude with the borrower (although we do not model the relationship between the officer and the borrower here). Under this interpretation a low type is more willing to collude with the borrower. At the beginning of the model the loan officer and the bank are symmetrically uninformed about the type of each officer. Let $\mu \in (0, 1)$ be the common prior belief about each officer’s type.

We suppose that the loan officer’s monitoring role has two distinct components: active and passive

³We rule out the possibility of assigning more than one officer to the loan in each period. We know that this does not occur in our sample.

⁴We adopt the convention that loan officer a is assigned in the first period. Since both loan officers are ex-ante identical this comes without any loss of generality.

monitoring. Active monitoring captures the loan officer's role in setting the terms of new loans and overseeing the repayment of current outstanding loans. When these tasks are performed by a highly skilled loan officer the expected profitability of the loan is higher. We will refer to this as the state of the loan in each period, which we denote by θ^t . The state of the loan in each period can be either high or low: $\theta^t \in \{\theta_l, \theta_h\}$ where θ_h (θ_l) refers to the high (low) state. In the first period the state of the loan is θ_h with probability $p_i \in \{p_l, p_h\}$ ($p_h > p_l$) where this probability is determined by the type $i \in \{h, l\}$ of the officer assigned to the loan in the first period. The evolution of the state of the loan in the second period is influenced by the types of the officers assigned in both the first and second periods. This captures the idea that there is some persistence in the profitability of the loan which is determined by the active monitoring performed by the first loan officer. The degree of persistence in the true state of the loan is denoted by $\phi \in [0, 1]$ and will play a crucial role in our analysis of rotation. Let j denote the type of the officer assigned to the loan in the second period. The state of the loan in the second period is determined in the following way:

$$\Pr(\theta^2 = \theta_h | \theta^1) = I(\theta^1 = \theta_h) [\phi + (1 - \phi) p_j] + (1 - I(\theta^1 = \theta_h)) [(1 - \phi) p_j]$$

where $I(\cdot) \in \{0, 1\}$ is an indicator variable that takes on the value of one if the state of the loan was high in the first period.⁵ Notice that if there is no persistence in the state of the loan ($\phi = 0$) then $\Pr(\theta^2 = \theta_h) = p_j$. Conversely, if there is full persistence ($\phi = 1$) then $\Pr(\theta^2 = \theta_h | \theta^1) = I(\theta^1 = \theta_h)$, the state of the loan in the second period is unaffected by the type of the officer assigned to the loan in the second period. High persistence can be thought to represent cases where the new loan officer has only been assigned to the relationship for a short period and has not been responsible for originating much new lending. In this case the profitability of the loan primarily reflects the first officer's skill.

The true state of the loan is not directly observed by either the loan officer or the bank. The loan officer, in her passive monitoring role, acquires information about the profitability of the borrower and potential problems with the loan indirectly through soft information. We model this in the following way. In each period the officer assigned to the loan privately observes a signal s^t of the true state of the loan. The officer either detects bad news, which we denote by s_l , or detects nothing which we

⁵Note that $\Pr(\theta^2 = \theta_l | \theta^1)$ is determined analogously as $1 - \Pr(\theta^2 = \theta_h | \theta^1)$.

denote by s_h . The probability of a loan officer of type i privately observing bad news is

$$\Pr(s^t = s_l | \theta^t) = J(\theta^t = \theta_l) q_i$$

where $J(\cdot) \in \{0, 1\}$ is an indicator variable and $q_i \in \{q_h, q_l\}$ is determined by the type of the officer assigned to the loan in period t . We assume that $q_h > q_l$ so that, conditional upon the true state of a loan being low, a high type loan officer is more likely to detect bad news. Without loss of generality, we assume that if $\theta^t = \theta_h$ then the loan officer detects nothing $s^t = s_h$ with certainty.

The loan officer's only decision in each period is whether or not to report bad news to her principal. Following Stein (2002) the loan officer can harden the soft information she has observed and pass it on to her superiors. If the loan officer observes no bad news she must report nothing which we denote by r_h .⁶ Conversely if the loan officer has privately observed bad news she can decide whether or not to pass this information on. The loan officer can pass this information on by issuing the low report (denoted r_l) or she can choose to suppress this information and report nothing (denoted r_h).

We now describe the incentives of a loan officer when she is deciding whether or not to pass bad news to her superiors in the bank. We assume that the bank is unable to compensate a loan officer based directly on the report she makes. One obvious motivation for this assumption is that the reports are unverifiable either because (i) they contain soft information or (ii) they contain information which the bank does not want to reveal publicly.⁷ We know in the case of the bank that we study that loan officers are not paid directly based on the ratings they issue for the loans under their control. Absent any direct incentives the loan officer will decide what report to issue so as to maximize her reputation at the end of the second period (i.e. the bank's assessment of her type). The bank will update its belief about the type of a loan officer based on the report issued in each period. We suppose that a loan officer with a higher reputation can expect to obtain more favorable employment from the bank (higher pay, less likely to be fired, better tasks) in the future. We do not model the reward to a high reputation directly and instead simply assume that a loan officer's utility is increasing linearly⁸ in this

⁶We assume that the loan officer is unable to fabricate bad news (no cheap talk). In practice the loan officer will have to justify any claim that she makes about the state of the loan. However, less specific justification is required for reporting no news at all. From a modelling perspective this assumption is important. Without it the loan officer's report will never be informative - all loan officers will issue the same cheap talk to maximize their reputations. This assumption makes it possible for the report to be informative.

⁷Alternately the bank may not wish to pay an officer directly based on these reports so as to avoid distorting her other actions (Holmstrom Milgrom (1987)).

⁸The assumption of linearity is not important for the analysis. We would obtain similar qualitative results if we assumed a more general utility function that is increasing in reputation.

reputation. Later in the paper we provide evidence for this career concerns assumption by showing that loan officers who accumulate observable events that are good (bad) for their reputation go on to manage larger (smaller) total lending portfolios.

B. No Persistence: $\phi = 0$

Begin by supposing that there is no persistence between periods: $\phi = 0$. This breaks the connection between each period and thus we can study the loan officer's decision in either period in isolation. Focus on officer a who is assigned to the loan in the first period. Suppose that officer a has observed bad news s_l . There are two opposing forces which affect the loan officer's decision to reveal bad news by reporting r_l as opposed to concealing this information and reporting nothing r_h :

- a low state of the loan (θ_l) is informative for the loan officer being low type ($p_h > p_l$). This is the force which may lead the loan officer to hide bad news so as not to damage her reputation as an active monitor.
- conditional upon the loan performing poorly, seeing bad news is informative for the loan officer being high type ($q_h > q_l$). This is the force which may lead the loan officer to pass on bad news since it demonstrates her ability as a passive monitor.

In the static model the loan officer's decision to reveal bad news will balance these two forces. We show in the appendix that when $\phi = 0$ the unique perfect Bayesian Nash equilibrium is for the loan officer assigned to the loan in each period to conceal bad news (i.e. always report r_h) if and only if

$$\frac{1 - p_l}{1 - p_h} > \frac{q_h}{q_l}. \quad (1)$$

The intuition behind condition (1) is easily understood by noticing that when a loan officer reports bad news that the bank learns both that the loan is in the low state and that the loan officer was able to observe the low signal. This condition compares the likelihood ratios associated with each of these two pieces of information. If the state of the loan is more informative for the officer's type than their ability to detect bad news then the loan officer will choose not to reveal bad news. When parameters are such that (1) is reversed the loan officer assigned to the loan will reveal bad news in each period for any level of persistence.

For the remainder of the analysis we will assume that (1) holds, which ensures that the loan officer does not have an incentive to reveal bad news in the static model (i.e. when $\phi = 0$). We ask whether the bank can induce truthful reporting by committing to a policy of loan officer rotation.

C. Rotation and Equilibrium Reporting

We now return to the case where there is some persistence in the state of the loan ($\phi > 0$). We show in the appendix that if the same officer is assigned to the loan in each period that the unique equilibrium is for her to conceal bad news in each period. Intuitively, condition (1) ensures that officer a is unwilling to reveal bad news in the first period. Moreover in the second period revealing bad news still reflects fully upon a 's ability as an active monitor, and hence she still has no incentive to issue the low report.

Now suppose that the bank commits to a policy of rotation. Rotation effects the incentives for both loan officers to report bad news. Consider first officer b , who is assigned to the loan in the second period. She is only partially responsible for the true state of the loan (since $\phi > 0$) and receives full credit for being able to detect bad news. Thus, if persistence is sufficiently high she is willing to truthfully report bad news whenever she observes it. The threat that officer b will reveal bad news also alters a 's incentives in the first period. When a observes bad news she knows that officer b is likely to also see bad news next period and reveal that the loan is performing poorly. Faced with the threat of being discovered by her successor, a has stronger incentives to report bad news (this at least demonstrates her ability as a passive monitor). Under a policy of rotation the equilibrium reporting strategy of both agents will be to report bad news if persistence (ϕ) is sufficiently high. This logic is formalized in the following Proposition.

Proposition 1 *Assume that condition (9) holds and that the bank commits to a policy of loan officer rotation. There exists a perfect Bayesian Nash equilibrium in which both officers truthfully reveal bad news if $\phi \geq \bar{\phi}$, for some $\bar{\phi} \in (0, 1)$.*

Proof. See Appendix.⁹ ■

It follows directly that the bank strictly prefers to commit to a policy of loan officer rotation when $\phi \geq \bar{\phi}$ since this induces loan officers to reveal private information.¹⁰

⁹Condition (9) holds when q_h and q_l are sufficiently large. It ensures that there exist some parameters for which the threat of b revealing information is sufficient to induce a to also reveal bad news.

¹⁰For simplicity, we have abstracted from costs that may arise from loan officer rotation. A policy of rotation is likely

D. Empirical Implications of the Theory

In the next Section we document the widespread practice of rotating loans across officers within the bank. Moreover we show that the bank recommends to a three-year rotation rule, which means that a loan is likely to be re-assigned to a new loan officer after it has been in the hands of a loan officer for three years. Our model provides an explanation for the policy: the bank uses rotation as a way of inducing loan officers to divulge information about their clients. We test our theory in several ways. An implication of Proposition 1 is that a loan officer will be more willing to divulge bad news when the threat of rotation is high. Immediately around the point in time when a loan is reassigned the state of the loan will be predominately determined by the actions of the incumbent loan officer. Put differently, in months close to a reassignment persistence (ϕ) will be high. Thus, as the threat of rotation increases (when an assignment approaches three years), loan officers will be more likely to reveal bad news about the loans they supervise. Moreover, we predict that the informativeness of their reports should increase as the threat of rotation rises.¹¹

The model delivers other testable implications. In the proof of Proposition 1 we show that b 's incentive to reveal bad news is strictly increasing in the degree of persistence ϕ . The logic is that when ϕ is high, the state of the loan is largely unrelated to b 's type and hence it does not hurt her reputation to reveal bad news. One way to interpret ϕ is a measure of the degree of lending that has originated with the new loan officer. When the new loan officer has overseen a high degree of new lending the state of the loan will be primarily determined by that officer's type (low ϕ). Thus we expect that loan officers will have the most incentive to conceal bad news when they have supervised a high volume of origination in the relationship. Accordingly, we expect that the threat of rotation will elicit the revelation of more bad news from these loan officers. We test this in the paper by comparing the effect of the threat of rotation for relationships that have had high and low origination. Our results confirm that rotation produces the most pronounced decline in ratings in lending relationships that

to come at the cost of losing some of the information and expertise which the incumbent loan officer has accumulated during her assignment. The anticipation of being rotated may also lower her ex-ante incentives to collect information. The bank's optimal rotation policy will trade-off these costs with the benefits of inducing more accurate reporting from loan officers.

¹¹For simplicity we have limited the model to two periods. One concern regarding these empirical predictions is that truth-telling at the end of a relationship may unravel the incentive for a loan officer to conceal bad news earlier in the relationship. In a three period version of the model where a is assigned for two periods and then is replaced by b we have found parameters for which a conceals bad news at $t = 1$, reports truthfully at $t = 2$ and b reports truthfully at $t = 3$. We have omitted a formal analysis of this extension from the paper for concision. Intuitively, a 's report at $t = 2$ has less power to induce truth telling at $t = 1$ because when a reveals bad news at $t = 2$ the reputational effects are mitigated by the fact that she receives credit for having observed the bad state. This force is absent when b reports bad news at $t = 3$.

have had high origination.

The model assumes that career concerns are a key force that motivates loan officers. A loan officer will choose what information to report based on how it will affect the public assessment of her type. We have justified this by assuming that the bank will use this information in the future when renewing the terms of the loan officer's employment. To confirm this assumption we look for evidence that loan officers are in fact rewarded (punished) for reports which increase (decrease) their reputation. The model makes clear predictions about reporting patterns which will raise or lower an officer's reputation. Consider the equilibrium defined in Proposition 1. The posterior beliefs which arise in that equilibrium have the following ordering¹²:

$$\mu > \hat{\mu}_a(r_h^a, r_l^b).$$

In words, the bank's assessment of a 's type will fall if she does not report bad news, and, after rotation the new loan officer discovers and reports bad news. This lowers a 's reputation for two reasons. First, the new officer's report indicates that the state of the loan under a may well have been low. Second, it indicates that a was unable to detect this bad news. The counterpart to this is:

$$\hat{\mu}_b(r_h^a, r_l^b) > \mu.$$

An officer who reports bad news early in her assignment will increase her reputation. This occurs because early in the assignment the state of the loan is only partially related to b 's type whereas she is able to demonstrate that she was able to detect the bad signal. We show in the data that officers that accumulate bad reputation events go on to manage a smaller total volume of lending and a smaller number of lending relationships.

III. Loan Officer Rotation and the Three-Year Rule

A. Data and Stylized Facts

To document loan officer rotation we assembled a unique unbalanced monthly panel, covering the 7-year period from December 1997 to December 2004, of loan officer-firm relationships (relationships, henceforth) from a multinational bank's operations in Argentina ('The Bank, henceforth). Our sample

¹²The notation we employ here is as follows: $\hat{\mu}_a(r_h^a, r_l^b)$ is the bank's posterior belief about a 's type after a has reported r_h in the first period and b reported r_l in the second period.

includes all lending by The Bank to small and medium sized firms (defined by the internal rules of the bank as a firm with net sales below \$50 million).¹³ At each point in time there is a single loan officer responsible for monitoring and originating loans for each borrower in the sample.

The data allows us to observe which borrowers were assigned to each loan officer in every month. For each loan officer, firm, month combination, we observe the amount of the loans approved and outstanding and a set of internal credit ratings produced by the loan officer. These credit ratings reflect the repayment status of the loan, observable borrower characteristics and the assessment of the loan officer about the credit quality of the borrower. The officers assessment is based on qualitative information obtained from visits and interviews with the management of the firm.¹⁴

To obtain information on the relationships of the borrowers in the sample with other financial institutions, we name-matched this database with the records of the Argentinean Central Bank Public Credit Registry (CDSF - Central de Deudores del Sistema Financiero). The CDSF provides monthly information on the amount of loans outstanding, the amount of collateral posted and standardized credit ratings issued by every financial institution to every borrower in the sample.

Table I shows the summary statistics of the main variables. We focus our analysis on the pre 2002 sample to avoid capturing in the results the effects of the economic crisis in Argentina in 2002. The internal Bank record data indicates that the mean outstanding loan amount is \$493,000 (median \$201,000) and the default rate in the sample is 9%. The average loan size is larger and the default rate lower than the average over the entire banking system during the same period.¹⁵ There is no significant difference between these figures and the ones obtained from the CDSF data for the same sample, which highlights the accuracy of the information reported by the Central Bank. The borrowers in the sample obtain finance from multiple banking sources. The median borrower has 7 banking relationships, a total bank debt of \$1.3 million, and obtains 17% of its bank debt from The Bank.

The final sample includes 1,248 firms and 100 loan officers. During the sample period we observe a total of 4,191 non-censored firm-loan officer relationships. Slightly more than 70% of the firms in the sample are observed to have two or more distinct relationships. The average firm has 3.19 relationships

¹³By comparison, according to the Small Business Administration's size definitions, a small business in the US has less than \$6.5 million of average annual revenue for retail and service industries.

¹⁴Loan officers responsible for medium and small business lending have complete discretion in terms of their actions and duties. Particularly there is no overruling or overlapping of their actions (as in Aghion and Tirole (1997) or Stein (2002)). On the contrary, Liberti (2004) and Liberti and Mian (2006) focus *only* on the large corporate lending portfolio of this financial institution where specific loan approval rules have to be complied with.

¹⁵Using the Central Bank data Paravisini (2006) documents an average default rate of 12% and an average loan size of \$16,000.

and sees 3.04 different loan officers during the entire 7 year sample period. The difference is due to the fact that 12% of the firms encounter the same loan officer more than once in two non-consecutive relationships. Table II summarizes this information.

The average length of non-censored relationships in the sample is 22.1 months with a median of 18 months. The median firm is observed for 62 months, and the median loan officer is observed for 47 months. These numbers indicate that there is a substantial amount of relationship turnover that is driven by loan officer rotation across borrowers, and not by firm or loan officer attrition. Less than 10% of the observed relationship turnover is driven by attrition. Furthermore, the firms under management of any loan officer are rotated gradually. The median number of firms under a loan officer's responsibility that gets reassigned in any single month is 3, conditional on any reassignment. This is small when compared to the average number of firms in a loan officers portfolio (25).

B. Three-Year Rule Rotation

The Bank's internal lending policies explicitly call for loan officer rotation. The following is an excerpt from the Internal Credit Policies of The Bank (March 2000):

The maximum length of a business relationship for Account Managers (AM) is recommended to be 3 years.

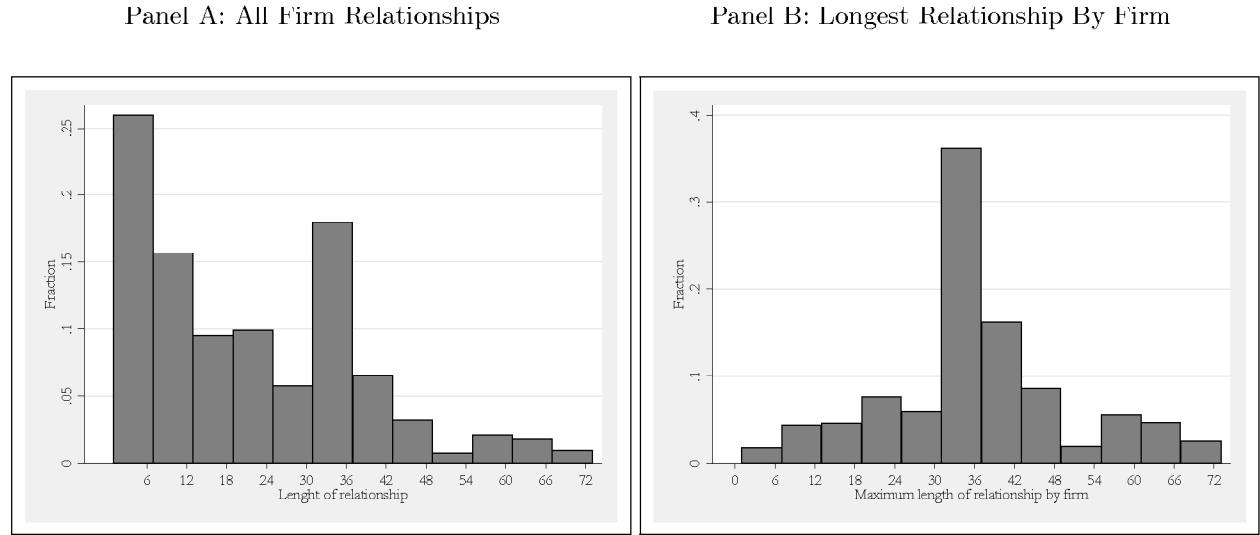
In the empirical context under consideration, account managers are loan officers, and a business relationship refers to a loan officer-firm pair assignment. The internal rules of the bank suggest a three-year limit to the duration of a relationship between loan officers and borrowers. We already documented that there is a substantial amount of rotation of loan officers and that the median relationship is below this three-year limit. However, conditional on lasting three years, the three year limit does effect the probability of rotation. To verify this, we plot the histogram of the length of non-censored relationships in Figure 1, Panel A. The histogram shows that the relationship length distribution is bimodal, with substantial mass between 1 and 6 months (25.5%) and between 31 and 36 months (19.4%). The mass of relationship lengths to the left of 36 months is consistent with the three-year rule. Furthermore, 83.3% of relationships end before 36 months.¹⁶

¹⁶The same patterns emerge when we look at the rotation and relationship length using the world-wide data for The Bank.

The large frequency of short relationships (less than 6 months) masks the importance of the 36 month rule. Relationships that last less than 6 months are in large part due to transitions towards longer relationships. To show this we plot in Panel B of Figure 1 the histogram of the length of the longest observed relationship per firm. Only a very small fraction of firms have relationships that are all shorter than 6 months. More importantly, 34.6% of relationships are terminated between 31 and 36 months, and 15.5% are terminated between 37 and 42 months. Overall, the longest relationship of more than half of the firms finishes within one year of the 3 year rule cut-off.

The histograms also highlight that the three-year rule is not binding. There are relationships that last up to 6 years in the data. The rule induces an increase in the unconditional probability that a firm under a loan officer's management will be allocated to another loan officer. In other words, the rule induces an increase in the ex-ante threat of rotation. We provide evidence in Section V that there is no systematic selection in the incidence of rotation at the three-year rule. In the next section we exploit the rule-induced threat of rotation to analyze the effect of rotation on loan officer behavior.

FIGURE 1
RELATIONSHIP LENGTH HISTOGRAMS



IV. Empirical Strategy

Our contention is that loan officer rotation across borrowers can affect the incentives of loan officers to report information to The Bank. Our theory in Section II argued that the anticipated threat of

rotation would induce a loan officer to reveal the bad news she had collected about a borrower. To empirically assess the effect of the threat of rotation on loan officer behavior it is necessary, first, that changes in the likelihood of rotation occur. Second, that these changes are, both, anticipated by loan officers and observable by the econometrician. The three-year rotation rule unambiguously assures these two conditions. As the previous section documented, it induces an increase (and subsequent decline) in the unconditional probability of rotation. The timing of the changes in the probability are based on a written rule, and determined by an observable relationship characteristic; its length. Our empirical strategy thus consists on using the length of a relationship and the three-year rule as a source of variation in the probability of rotation.

The crucial identification condition of this empirical strategy is that the time series variation in the probability of rotation is not correlated with shocks to firm creditworthiness or demand for credit. This condition allows us to distinguish changes in the credit rating of a firm that are induced by the effect of rotation on loan officer reporting, from credit rating changes that are induced by underlying changes in firm creditworthiness.

The three-year rule provides a source of variation in the threat of rotation that is plausibly uncorrelated with firm level shocks. The increase, and subsequent decline, of the probability of rotation as the loan officer-firm relationship age approaches 36 months, are driven by the date the relationship began. It is unlikely that a three year clock would be correlated with firm specific shocks. We now provide evidence that this is the case.

A. Exogeneity of Three-Year Rule

In order to further justify our identification strategy we show that the three-year clock is not correlated with firm specific shocks to the demand for credit or the creditworthiness of the borrowers in our sample. To do this we study how the firms' credit ratings and lending outcomes *with all lenders excluding The Bank* vary around a loan officer rotation predicted by the three-year rule. Shocks to a firm's creditworthiness or demand for credit should affect its ratings and borrowing from sources other than The Bank. In order to demonstrate that the three-year rule is uncorrelated with firm specific shocks, we show that lending outcomes with other banks do not vary systematically around officer rotations predicted by the rule.

We look at changes in lending behavior during the four quarters preceding and following an ex-

pected loan officer rotation according to the three-year rule. We do so by estimating the following specification:

$$\begin{aligned}
y_{ijt,k} = & \beta_{+1Q}^{3year} (I_{ThreeYearClock+1/+3})_{ijt,k} + \dots + \beta_{+4Q}^{3year} (I_{ThreeYearClock+10/+12})_{ijt,k} + \\
& \beta_{-1Q}^{3year} (I_{ThreeYearClock-1/-3})_{ijt,k} + \dots + \beta_{-4Q}^{3year} (I_{ThreeYearClock-10/-12})_{ijt,k} + \\
& \alpha_i + \alpha_j + \delta_{Industry \times t} + \varepsilon_{ijt}
\end{aligned} \quad (2)$$

The left hand side variable is the lending outcome of interest of firm i , managed by loan officer j of bank k at month t ($y_{ijt,k}$). The right hand side includes eight indicator variables, one for each of the four quarters of the year prior to and following a loan officer rotation for firm i as *predicted by the three-year rule*.

The predicted loan officer rotation is constructed as follows. Suppose that loan officer j is assigned to firm i at month t . A clock starts counting the number of months since t , and predicts that firm i 's loan officer is more likely to change between $t + 34$ and $t + 36$, if loan officer j is still assigned to i at $t + 33$. If firm i is assigned a different loan officer than j at any time between t and $t + 33$, the clock is reset to zero and starts counting again at the date of the last change.¹⁷

The quarter dummies on the right hand side of (2) are then assigned relative to the predicted quarter of loan officer rotation. In other words, the dummy for the first quarter after (before) a predicted change is set to one between months $t + 37$ and $t + 39$ ($t + 31$ and $t + 33$) when a loan officer change occurred at t and no change occurred until $t + 33$. The coefficients on the quarter indicators represent the average of the outcome of interest during each of the four quarters preceding (β_{-1Q}^{3year} through β_{-4Q}^{3year}) and following (β_{+1Q}^{3year} through β_{+4Q}^{3year}) a rule induced increase in the probability of loan officer rotation.¹⁸¹⁹

We also include firm and loan officer fixed effects to control for unobserved cross sectional het-

¹⁷The use of a rule induced clock to identify exogenous increases in the probability of rotation is similar to the empirical strategy employed by Cole (2007) and Khemani (2004).

¹⁸Note that a firm must be observed in the data for at least three years before an increase in the probability of rotation can be predicted using this procedure. Given that the sample period begins on December 1997, the analysis in this section takes into consideration increases in the probability of rotation induced by the 3 year rule after December 2000. The findings of previous subsection are quantitatively and qualitatively similar when estimated using the sub-sample of loan officer changes that occur after that date, and are thus still useful for comparison with the results presented here.

¹⁹With this procedure we are estimating the local effect of turnover for a the selected sample of relationships that reach the three year cutoff. We documented a substantial amount of turnover before the three year cutoff. This early turnover seems to be associated with default. Thus, firms that reach the three year cutoff are likely to have a lower credit risk than the average firm in the sample.

erogeneity (α_i, α_j). Finally, we include a full set of industry-month interactions to control for time series variation in lending outcomes that is driven by shocks to the demand for credit ($\delta_{Industry} \times t$). The fixed effects specification ensures that the averages are measured as deviations from firm and loan officer means. The industry-month dummies imply that we measure deviations relative to the average firm in the same industry. All the standard errors in this specification (and in every specification hereafter unless otherwise noted) are heteroskedasticity-robust and estimated allowing for clustering at the firm level.

FIGURE 2

FIRM LENDING OUTCOMES AROUND ROTATION, WITH ALL LENDERS EXCLUDING THE BANK

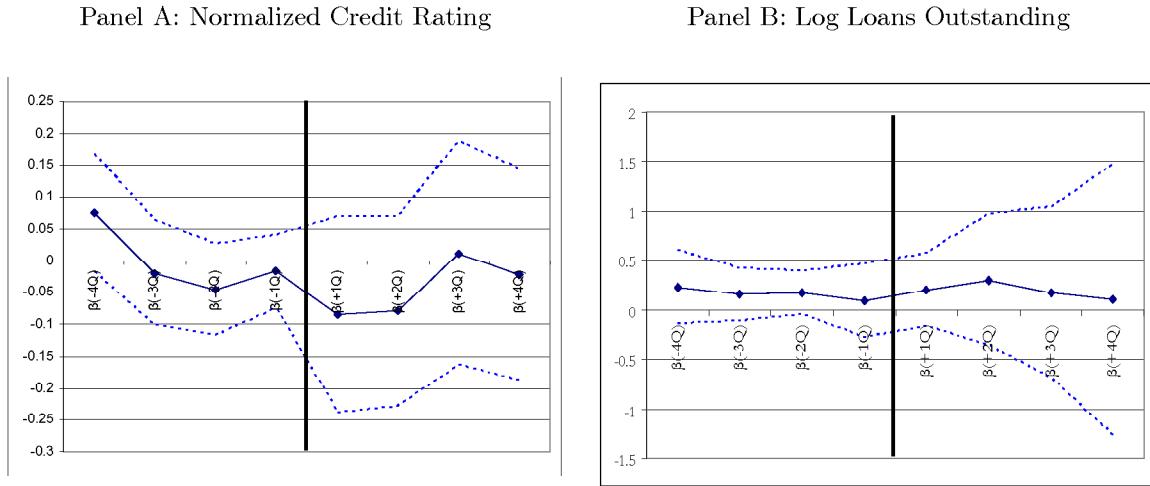


Figure 2 plots the quarter effects obtained using two left hand side variables. Panel A shows the results using the normalized weighted credit rating issued by all banks excluding The Bank. If The Bank is indexed by k , then the weighted risk rating is given by:

$$wrr_{ijt,-k} \equiv \frac{\sum_{n \neq k} (rr_{ijt,n}) (debt_{ijt,n})}{\sum_{n \neq k} (debt_{ijt,n})} \quad (3)$$

where n indexes all the banks in the financial system. Panel B shows the results using the log of the total amount of outstanding debt with all banks excluding The Bank ($\ln(\sum_{n \neq k} debt_{ijt,n})$). The two plots are similar in the lack of obvious or significant pattern in lending outcomes around the predicted

rotation date. These plots corroborate the identifying assumption. The timing of the loan officer rotation induced by the three-year rule is uncorrelated with either the credit quality or the demand for credit of the firms in the sample. In the next section we will show how the credit ratings assigned by loan officers within The Bank vary around the three-year rule. The results in this section insure that these patterns are not driven time-varying firm-specific shocks. Thus we can interpret our results in the next section to be caused by the ex-ante threat of rotation itself.

V. Rotation, Credit Ratings and Information

A. The Effect of Rotation on Credit Ratings

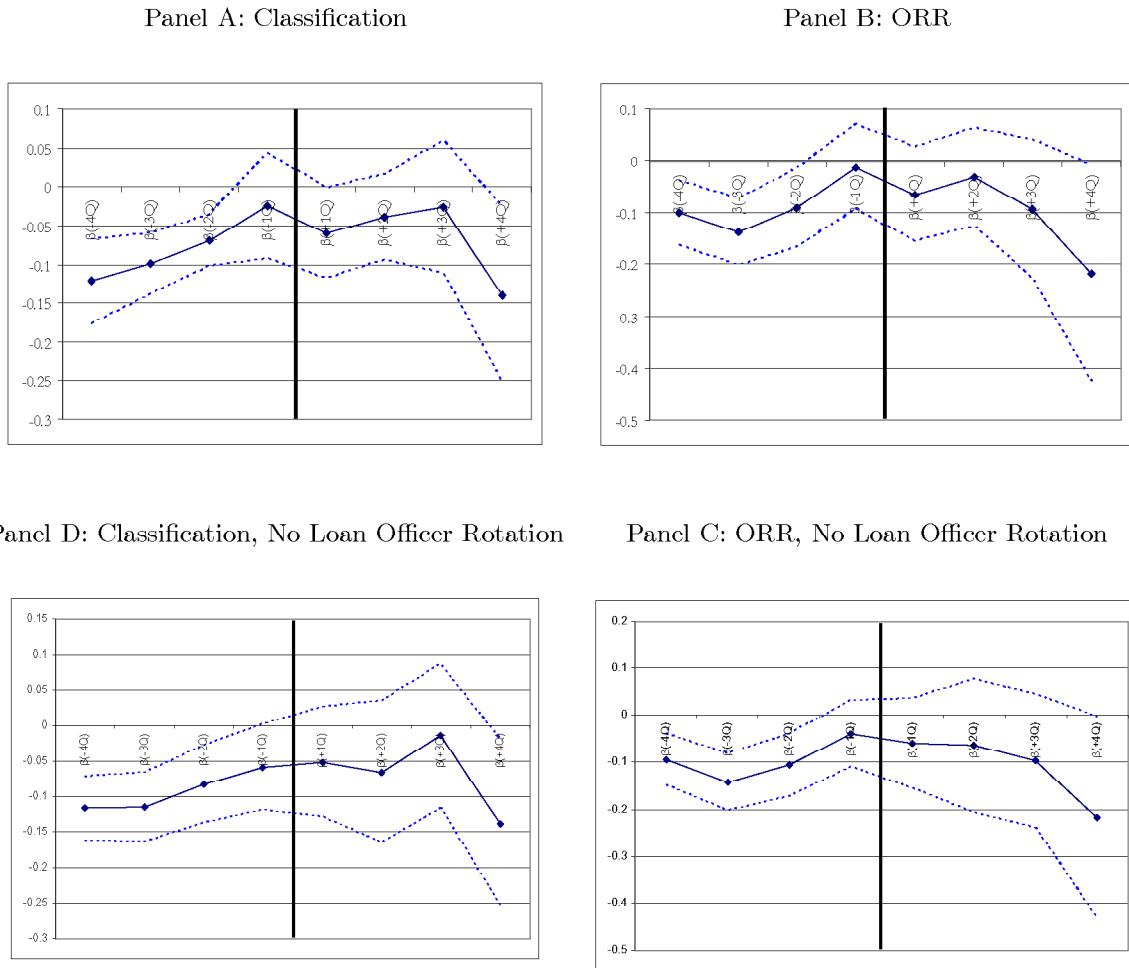
We have shown so far that there is a rule induced increase in the probability of loan officer rotation when a relationship approaches three years. We also showed that this increase is uncorrelated with firm specific shocks after controlling for firm, loan officer and industry-month dummies. This implies that any systematic change in the observed internal credit ratings of The Bank around the three year-rule can be interpreted as a consequence of rotation, after controlling for those fixed effects. In this section we show that such patterns exist, and that they are consistent with the testable implications of the model discussed in the theoretical framework.

A-1. Rating Changes Around Rule-Induced Rotation

We begin by showing graphically how the credit ratings assigned by loan officers behave around the three-year rule using specification (2), augmented to control for the weighted risk rating assigned by the other lenders of firm i as defined in (3). The quarter dummies are estimated using two different internal risk ratings of The Bank as dependent variables, *Classification* and *ORR*. These risk ratings are assigned by loan officers as measures of the perceived probability of default of the firm. Loan officers have a significant amount of discretion in assigning these risk ratings. *Classification* is a number between 1 and 5 and its choice is entirely under the discretion of the loan officer. *ORR* is obtained from a formula that uses as inputs both information from the financial statements of the firm and unverifiable assessments made by the loan officer, e.g., quality and trustworthiness of the management. Both risk rating measures are normalized in the sample to have zero mean and unit variance to ease the interpretation of the results. Figure 3 plots the estimated coefficients of the quarter dummies from

specification (2).²⁰ The quarter effects represent the average deviation of the risk rating obtained by a firm from The Bank's loan officer, relative to the mean risk rating assigned to the firm by all other banks.

FIGURE 3
CHANGES IN INTERNAL RISK RATINGS (NORMALIZED)
INDUCED BY THE THREE-YEAR ROTATION RULE



The plotted average risk ratings are increasing during the four quarters preceding the high probability of rotation period. This is consistent with loan officers systematically downgrading firms as the threat of rotation increases due to the three-year rule. Furthermore, estimated internal risk ratings are negative and significant three and four quarters before the high rotation period, and they converge to

²⁰We have omitted the tables for this regression for conciseness. Refer to Table V for details of the parametric counterpart to this regression (as explained below).

zero as the high threat of rotation period approaches. Our interpretation of these patterns is as follows. When the threat of rotation is low (twelve months before the three year rule) loan officers conceal bad information about a loan. It is for this reason that the risk ratings they assign are systematically lower (i.e. more optimistic) than those assigned to the same firm by all other banks that it borrows from. As the threat of rotation increases loan officers undo the bias from the ratings they assign and are more willing to reveal bad news.

After the high rotation period, ratings remain constant and statistically indistinguishable from zero for 3 quarters. Ratings decline again and become significantly negative (firms are upgraded) four quarters after the high rotation period when the threat of rotation has declined. Again, the results indicate that loan officers systematically upgrade the credit rating of firms when the probability of rotation declines.

Our interpretation of Panel A and B of Figure 3 is that it is the increase in the *ex-ante* probability of rotation which induces a loan officer to reveal bad news. To confirm this interpretation we repeat the estimations with all the quarter dummies interacted with a dummy equal to one if the relationship between firm i and loan officer j is not turned over in three years. The plotted coefficients in Panels C and D of Figure 3 represent the average risk ratings of the firms that remained with the same loan officer throughout the entire high rotation period. The observed patterns are the same as those described for Panels A and B. Risk ratings increase before the high rotation period, then remain stable for three quarters and finally fall.

These results indicate, first, the observed increase in ratings before the high rotation period is not exclusive to those firms whose loan officers change. This allows us to rule out the possibility that firm downgrades are inducing loan officer rotations around the three year period. And second, firms experience an upgrade in their credit ratings after the high rotation period *even if their loan officer did not change*. This means that the ratings change pattern observed after the high rotation period is not due to new, uninformed loan officers taking over the management of the relationship with the firm.

A-2. Parametric Tests of Rating Changes

The results in Figure 3 suggest that ratings increase prior to rotation and decline afterward. We test for a change in this trend around the rule-induced rotation period by estimating the following

specification:

$$\begin{aligned}
y_{ijt,k} = & \eta_{SB} (\tilde{t}_{ijt} \times Before3YR_{ijt}) + \eta_{SA} (\tilde{t}_{ijt} \times After3YR_{ijt}) + \\
& \eta_{LB} (Before3YR_{ijt}) + \eta_{LA} (After3YR_{ijt}) + \\
& \varphi y_{ijt,-k} + \phi_i t + \alpha_i + \alpha_{jt} + \delta_{Industry} \times t + \varepsilon_{ijt}
\end{aligned} \tag{4}$$

The right hand side variable is, as before, the normalized risk rating (Classification, ORR) of firm i , assigned by loan officer j , in month t and bank k (The Bank). $Before3YR_{ijt}$ and $After3YR_{ijt}$ are two dummy variables that turn to one during the year preceding and following an expected increase in the probability of rotation according to the 3-year-rule. \tilde{t}_{ijt} is a firm specific time trend, normalized to zero at the time of the predicted increase in the probability of rotation. Thus, the interaction coefficients η_{SB} and η_{SA} , measure the rate of change of the outcome of interest before and after the rule induced rotation quarter. According to the patterns in Figure 3, we expect $\eta_{SB} > \eta_{SA}$. The control variables remain as before: firm, loan officer and industry times month dummies, and the average normalized risk rating assigned to firm i by all other lenders.

The estimated parameters are shown in Table III and the results of Figure 3 are confirmed. There is a significant decline in the rate of change of ratings around the rule-induced rotation in all specifications. Moreover, the rate of change of the internal risk rating measures (columns 1 and 3) is positive before the high rotation quarter and negative after. The magnitudes of the coefficients indicate that the change in the rate of change of risk ratings was substantial. Risk ratings increased by 12% (13%) of a standard deviation during the year preceding the rule induced rotation, and declined by 13% (32%) of a standard deviation during the year following the rule induced rotation when risk ratings are measured using classification (ORR).

The results put together corroborate that the observed rating patterns are driven by the effect that the probability of rotation has on loan officer behavior. Loan officer rotation affects risk rating bias: ratings change systematically in response to the threat of rotation in a way uncorrelated with underlying firm characteristics. We explore further the direction of the bias at the end of this section when we look at how the rating changes reflect the probability of default of a firm.

A-3. Relationship to Loan Origination

Our theoretical framework suggested that loan officers have incentives to hide bad information about borrowers because this reflects poorly on the loan officer's ability to originate and manage successful lending relationships (active monitoring). If this hypothesis is true, we expect that a loan officer will have stronger incentives to produce biased rating reports for a firm when she has originated more loans for it. A testable implication of this hypothesis is that the threat of rotation should induce starker adjustments of ratings in those loan officer-firm relationships characterized by above average origination.

We verify this testable implication by looking at how the previously described risk rating patterns vary according to the amount of origination in the relationship. First we calculate the origination rate of every loan officer-firm relationship as the average percentage monthly change in the amount of debt outstanding. Then we classify relationships by creating a dummy $HighOrigination_{ijt}$ that turns to one for every firm-loan officer-month cell that belongs to a relationship that has an origination rate above the median. Finally, specification (4) is estimated including an interaction between the trend in rating changes before the high rotation period ($\tilde{t}_{ijt} \times Before3YR_{ijt}$) and $HighOrigination_{ijt}$.

Columns 2 and 4 of Table III show the estimated coefficients using both Classification and ORR as dependent variables. The interaction between the $HighOrigination_{ijt}$ dummy and other dependent variables is included in the regression but the coefficients are omitted from the table for conciseness. The coefficient on the interaction term of interest represents the difference between the rate of change of ratings before the high rotation period amongst the high origination and the low origination loan officers. The estimated coefficient in both specifications is positive and significant, indicating that high origination loan officers tend to downgrade firms more as the high rotation period approaches, relative to the low origination loan officers. In fact, the coefficient on the slope becomes statistically insignificant when the interaction term is included, which means that most of the observable downgrade is due to the high origination loan officers.

These results confirm our theory that, absent rotation, loan officers have the strongest incentive to conceal bad news when the state of the loan is most informative for their type. Our results demonstrate that it is the subgroup of loan officers with high origination who underestimate the credit risk of their clients twelve months away from the high rotation period. Moreover, rotation has the strongest impact on the ratings issued by these loan officers since they are the subgroup that has been

previously concealing bad news.

B. Information Content of Risk Ratings

Up to this point we have documented how the *level* of ratings vary in response to the ex-ante threat of rotation. We now demonstrate that ratings do in fact contain more information as the probability of rotation increases. To study the effect of loan officer rotation on the information content of risk ratings, we explore how changes in risk ratings induced by the three-year rule are related to the *actual* probability of default. Using the three year rule as a source of variation in ratings, we estimate the predictive power of internal risk ratings on the probability of default of loans. In other words, we will use the time to, and after, a predicted change according to the three-year rule ($\tilde{t}_{ijt} \times Before3YR_{ijt}$, $\tilde{t}_{ijt} \times After3YR_{ijt}$) as instruments for changes in risk ratings. The results presented so far provide the sufficient and necessary conditions for the validity of these instruments. First, they are significantly correlated with rating changes, as documented by Table III. Second, they are uncorrelated with other firm characteristics affecting risk ratings (Figure 2).

To measure the predictive power of the bank's internal risk ratings on the probability of default of a firm we estimate the following linear default probability model:

$$Default_{ijt+12,k} = \gamma (irr_{ijt,k}) + \varphi (wrr_{ijt,-k}) + \phi_i t + \alpha_i + \alpha_{jt} + \delta_{Industry \times t} + \omega_{it} \quad (5)$$

The left hand side variable is a dummy equal to one if firm i is in default at $t + 12$. The right hand side variable of interest, irr , is the internal risk rating assigned by loan officer j (of bank k , or The Bank) to firm i at month t . As in the previous subsection, we will use two different internal risk ratings of the bank (*Classification, ORR*). The control variables include the weighted average of the risk rating of firm i with other banks at t (from 3), firm fixed effects, loan officer fixed effects, and a set of industry times month dummies.

The parameter of interest, γ , represents the predictive power of changes in the internal risk rating of k on the probability of default, above and beyond the risk rating received by the firm from other banks. Our interest resides on the predictive power of the risk rating around loan officer changes. Thus, the parameters are estimated restricting the sample to the 12 months preceding and following any change in loan officer.

Column 1 of Table IV shows the estimation results using OLS. The estimated γ^{OLS} is 0.106 (se

0.013) which implies that, on average, an increase in one standard deviation in the risk rating is associated with an increase in 10 percentage points of the probability of default. Column 2 (Column 3) shows the estimated γ^{OLS} over the subsample of 12 months before (after) a loan officer change. The estimated parameters suggest there is a strong and positive association between changes in internal risk ratings and changes in the probability of default both before and after a change in loan officer. The predictive power of risk ratings set by the new loan officers during the year following rotation is lower: a one standard deviation increase change in the risk rating is associated with a 5 percentage point increase in the probability of default.

Columns 4 through 6 of Table IV show the estimated parameter of interest using 2SLS, where $Before3YR_{ijt}$ and $After3YR_{ijt}$ and their interaction with \tilde{t}_{ijt} are used as instruments for the internal risk rating ($rr_{ijt,k}$). γ^{2SLS} represents the predictive power internal risk ratings on the probability of default, when the risk ratings changes are induced by the 3-year-rule. It also estimates the predictive power of ratings on a selected sample of firms: those that reach a three year relationship with their loan officer. These firms have had a longer than average relationship with their current loan officers and are likely to be of better than average credit quality (early rotations are on average associated with poor firm performance).

Column 4 shows the estimated γ^{2SLS} using the variation in internal risk ratings during the full two years surrounding a predicted increase in rotation probability. The point estimate is positive but insignificant. On average, changes in internal risk ratings induced by the 3-year-rule are not informative about the probability of default of the firm. A different picture appears when the γ^{2SLS} is estimated separately over the before (column 5) and after (column 6) high rotation periods. The estimated γ^{2SLS} over the two subsamples are 1.07 (se 0.447) and -1.02 (se 0.786) respectively. Changes in the internal risk rating that occur before the predicted increase in rotation probability are positively correlated and with the probability of default of the firm. The point estimate suggests that the average increase in the standardized risk rating before the 3-year-rule (0.13 from Figure 5) is associated with a 13 percentage point increase in the probability of default. The standard error associated with the estimated parameter is high: the 95% confidence interval ranges from 2.5 to 25 percentage points. Conversely, risk rating changes that occur after the 3-year-rule induced rotation are negatively correlated with the probability of default, although not significantly so.

The results are consistent with our hypothesis that the threat of rotation induces loan officers

to reveal bad news about the firm. When the threat of rotation increases, firms are systematically downgraded, and these downgrades are informative about the future probability of default of the firm. Once the probability of rotation declines, firms are systematically upgraded and these upgrades do not contain information about the future probability of default of the firm. In fact, the negative point estimates in the information content estimation (column 6) suggest that these upgrades reflect an optimistic bias by the loan officers.

As before, we look for whether there is any cross sectional heterogeneity in these patterns across firms that did and did not experience a loan officer change during the high rotation period. We repeat the 2SLS estimation using the 3-year-rule instruments interacted with a dummy equal to one if firm i did (did not) change loan officers around the rule-induced high rotation quarter. The estimated predictive power of ratings are presented in columns 7 through 9 (10 through 12) of Table 6. The observed patterns confirm the hypothesis that it is the anticipated threat of rotation which leads to an increase in the informativeness of ratings. Changes in ratings that occur before the rule-induced high rotation period are informative about the future probability of default for both groups of firms.

Changes in ratings after the high rotation period are not significantly correlated with changes in the probability of default for either group. This confirms that the lack of predictive power of the ratings after the high rotation period are not due to new and uninformed loan officers. In fact the data suggest the opposite: after the high rotation period, the point estimate of the group that experienced a loan officer change is positive, and the point estimate for the no-change group is negative. The coefficients are not different from zero, but they are significantly different from each other. After the three year rule the changes in ratings of the no-change group appear to be systematically less informative than those assigned by new loan officers. This suggests that new loan officers release more information early in their tenure than the incumbent loan officer would choose to reveal. This is consistent with our argument that it is the threat coming from the additional information released by new loan officers that compels incumbent loan officers to reveal bad information prior to rotation.

VI. Reputation and Career Concerns

One of the key forces in our theory of the effects of rotation is that loan officers care about their reputation. Bad news reflect poorly on the ability of the loan officer, and thus loan officers have an incentive to hide it. On the other hand, when a new loan officer is assigned to a relationship she

has minimal responsibility for the current state of the loan and revealing bad news demonstrates her ability as a monitor. In this section we document that the career prospects of a loan officer are affected by the reputation effects of the reports they make to the bank in a way that is consistent with our theory.

We build indexes of loan officer reputation based on the predictions of the model (see Section 2). The bank’s assessment of the ability of a loan officer deteriorates if she does not report bad news and then, after rotation, the new loan officer discovers and reports bad news. This demonstrates (i) that the state of the loan under the previous officer was low and (ii) that she was unable to detect this. Thus, we construct an index of bad reputation by counting the number of events, up to time t , in which the firm under assignment to loan officer j is downgraded by the new loan officer during the 6 months after a rotation ($NBad_{jt}$). We exclude from the count the observations when the firm under management is already classified as very likely to default or in default to prevent mechanical changes in ratings after firms are allocated to risk managers. In other words, we do not take into consideration the post-rotation downgrades when the initial risk ratings of the firm have $Classification \geq 2$, or $ORR \geq 8$. We also scale this count by the number of rotations a loan officer has experienced up to time t , so the bad reputation index of loan officer j at time t is given by $\frac{NBad_{jt}}{NRot_{jt}}$. Scaling by the number of rotations assures that a loan officer is not assigned mechanically a higher index when it goes through more rotations when, for example, she handles a larger portfolio. The bad reputation index is the proportion of all j ’s loan assignments that were downgraded within six months of handing them to another officer.

The bank’s assessment of the loan officer’s ability improves when she detects and reports bad news after a rotation. We count the number of times a loan officer i has downgraded a firm during the six months after a rotation up to time t ($NGood_{jt}$). Then we scale it by the number of rotations experienced by loan officer j up to time t to obtain a measure of good reputation. As before, the cases where the firm was rated as having a high probability of default before rotation are not included in the count.

The model predicts a negative (positive) correlation between the bad (good) reputation index and career outcomes of the loan officer. We focus on measures of assets under management of the loan officer as proxies for career outcomes. This choice hinges on the idea that the bank will reallocate the assets under the control of different loan officers as it updates its beliefs about their ability to manage

lending relationships and monitor the creditworthiness of borrowers (see Berk and Green (2004)). Two measures of assets under management are used: the total amount of loans outstanding and the number of firms under management of a loan officer at any given time t .

Recalling the discussion in Section II, the model also allows making predictions about the relative magnitude of the correlation between the two reputation indexes and career outcomes. When the successor of a loan officer downgrades a firm after a rotation, the bank learns two things about the incumbent loan officer: that she did a poor job in managing the relationship (poor active monitoring ability) and that she failed to detect that there was bad news to report (low passive monitoring ability). On the other hand, when a loan officer downgrades a loan after a rotation, the bank learns about the high ability of the new loan officer to detect bad loans, but nothing about her active monitoring abilities. Thus, we expect bad reputation events to have a higher impact, in absolute value, on career outcomes of a loan officer than good reputation events. The bad reputation index is a more informative signal about the underlying ability of the loan officer.

We test these predictions of the model by collapsing the data at the loan officer-month level and estimating the parameters of the following specification:

$$\ln(Y_{jt}) = \theta_{Bad} \frac{NBad_{jt-6}}{NRot_{jt-6}} + \theta_{Good} \frac{NBad_{jt-6}}{NRot_{jt-6}} + \gamma X_{jt} + \alpha_j + \alpha_t + v_{it} \quad (6)$$

The left hand side variable is a measure of the career prospects of loan officer j at time t . The variables on the right hand side are the good and bad reputation indexes (times 100 to ease interpretation of point estimates), loan officer fixed effects and month dummies. The reputation indexes are lagged 6 months to allow for a response time between changes in reputation and the reassignment of assets. All specifications also include as a control the average risk rating assigned to the firms under management of loan officer j by all other banks (using Central Bank data), weighted by the amount of loans outstanding of each firm. This control is meant to account for observable features of the loan portfolio of the officer that may also affect career outcomes. All the standard errors are estimated allowing for clustering at the loan officer level.

Estimating (6) directly using OLS may produce a simultaneity bias. The bank may endogenously choose to rotate officers whose portfolios it wants to decrease. We use a standard IV procedure to isolate changes in reputation that come from exogenous variation in rotation. As instruments for the good and bad reputation indexes, we use measures of good and bad reputation based solely on the

rotations that are induced by the three-year rule. Thus, the bad reputation instrument has in the numerator the number of events, up to time t , in which the firm under assignment to loan officer j is downgraded by the new loan officer during the 6 months after a *predicted rotation using the three-year rule*. Similarly, the good reputation instrument has in the numerator number of times a loan officer i has downgraded a firm during the six months after a rule-predicted rotation up to time t . The first stage regressions are shown in Table V of the Appendix and indicate that the bad (good) reputation instrument are significant at the 1% level after controlling for the weighted risk rating, loan officer fixed effects and year dummies.

The 2SLS estimation of the parameters of specification (6) can be interpreted as the effect on the future assets under management of a loan officer, when the fraction of bad (good) reputation events increases by one percentage point. The OLS and 2SLS estimations of the parameters are reported in Table VI. The effect of our bad reputation measure is consistently negative across all measures of assets under management and estimation method. The effect of our good reputation measure is consistently positive, but only statistically significant in the OLS specifications. The 2SLS specifications (columns 3 and 4) obtain a point estimate of -0.32 when loans outstanding are used as the outcome. This implies that a one percentage point increase in the bad reputation index can be associated with a decline of 32% in the outstanding loans (15% of the firms) under management. The large impact of bad reputation events on the career prospects is consistent with the fact that the distribution of the bad reputation index is heavily skewed, with 62% of the observations in the sample having an index of zero (no bad reputation events).

These results provide strong evidence for our assumption that The Bank will use the ratings assigned by loan officers to update its belief about their type. Consistent with our model, handing over a loan which is subsequently downgraded has an economically significant impact upon the officer's future career. The effect of this event on an officer's reputation is particularly pronounced because it reflects poorly on her ability as both a passive and active monitor. The results confirm our argument that the threat of being discovered by a successor provides a strong incentive for an incumbent loan officer to reveal bad news herself.

Our results show that *the same downgrade* which hurts the reputation of the previous loan officer improves or does not hurt the reputation of the new loan officer. This report allows her to demonstrate her skill as a passive monitor. Since there is only one force driving this update we expect the

reputational effect for the new loan officer to be smaller in magnitude than the effect for the replaced officer. Our estimates confirm this. The fact that the new loan officer is not punished confirms our explanation of why rotation works. By temporarily separating active and passive monitoring rotation removes the disincentive of the new loan officer to reveal bad news about the loan.

Our results documenting the reputation effects of the loan officer's rating decisions add to the empirical literature on the role of career concerns in delegated portfolio management (see for example Chevalier and Ellison (1999)). Our results lend support to Rajan (1994) who supposes that loan officers may hide poorly performing loans in order to maintain their public reputation.

VII. Real Effects of Rotation: Lending Outcomes

So far we have focused on the incentive effects of rotation on the ratings decisions of loan officers. We have shown that the threat of rotation reduces the incentives of loan officers to systematically underestimate the credit risk of borrowers under their supervision. As a result, rotation increases the accuracy with which the bank can infer the probability of default of a borrower from the internal risk ratings. We conjecture in this section that such an increase in ability to assess borrower creditworthiness must affect lending outcomes.

In many theoretical settings total lending increases when investors (lenders) have more precise signals about the quality of the investments or the ability of the borrowers to repay. When investors are asymmetrically informed about the investment prospects of a firm, adverse selection leads to the use of lower external finance (Leland and Pyle (1977), Myers and Majluf (1984)) and credit rationing (Stiglitz and Weiss (1981)). External financing and lending increase as the investors become better informed about the investment prospects of the firm.

In the empirical context of the paper, we expect the total amount of lending to increase as the informativeness of the credit ratings increases. In other words, we expect the total amount of lending to increase as the threat of rotation increases. Moreover, we expect that the effect of the improvement of information on lending will vary in the cross section of firms. In particular, downgraded (upgraded) borrowers should observe a decline (increase) in the amount of credit. In what follows, we corroborate these patterns in the data.

A. Total Lending Patterns

We first look at how total lending by The Bank changes during the four quarters preceding and following an expected loan officer rotation according to the three-year rule. We use specifications (2) and (4) with two modifications. First, we use as a dependent variable the log of the amount of debt outstanding of firm i with The Bank (k), under the management of loan officer j at month t ($\ln(debt_{ijt,k})$). Second, we add as a control variable the log of the total amount of debt of firm i at time t with all other banks in the financial system, obtained from the Central Bank data ($\ln(\sum_{n \neq k} debt_{ijt,n})$). As before, we include firm and loan officer fixed effects to control for unobserved cross sectional heterogeneity, and a full set of industry-month interactions to control for time series variation in lending outcomes that is driven by shocks to the industry wide demand for credit.

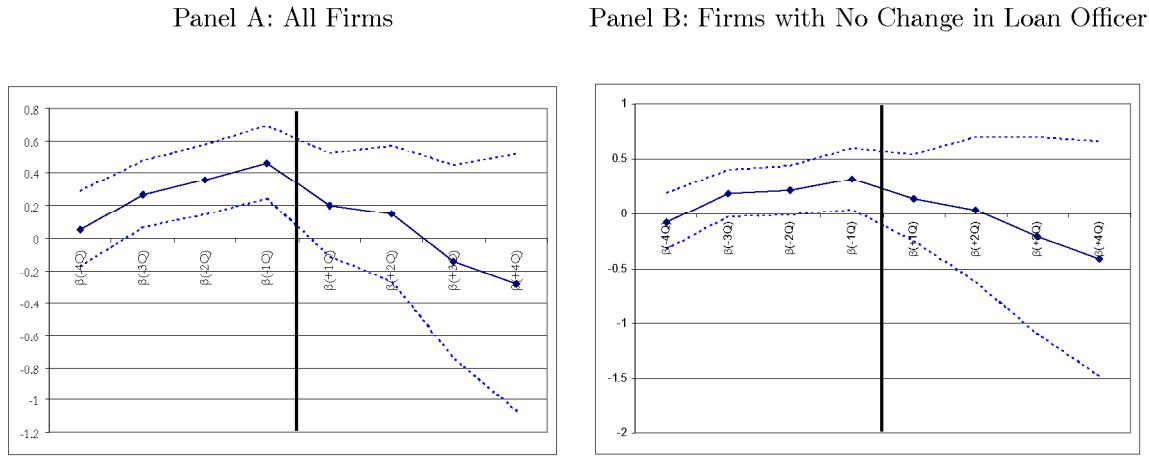
Panel A of figure 4 plots the estimated quarter effects, which represent the average lending during each of the four quarters preceding and following a predicted loan officer rotation, *relative to the amount of borrowing of the same firm with all other banks*. So a point estimate of zero for a quarter effect means that the fraction of debt outstanding from The Bank relative to other lenders is at its average in the sample. The plot shows that the quarter effects increase from zero to around 0.4 during the four quarters preceding the loan officer rotation. Given that the average firm in the sample obtains 27% of its bank debt from The Bank, these estimates imply that the fraction of debt from The Bank increases by 10 percentage points during the year leading up to the high rotation period. This is precisely the period during which the information content of credit ratings increased, as documented in Table IV.

Panel A also shows that after the high rotation period the quarter effect estimates declines and becomes statistically indistinguishable from zero. Even though the standard errors are large and all the point estimates of the quarter effects are not statistically different from zero, we corroborate using specification (4) that the point estimate of the slope is negative and significant (estimates reported in column 5 of Table III). These results indicate that the fraction of lending from the Bank declines to the sample average once the high rotation period is over.

In Panel B of Figure 4 we show the quarter estimates when interacted with an indicator variable that turns to one if the loan officer of the firm *was not* rotated at 36 months. This confirms that the observed patterns are there regardless of whether the loan officer rotation happened or not. As in our discussion of Figure 3 on credit rating patterns (Section V), this evidence suggests it is unlikely that

the increase in lending was responsible for loan officer rotation (reverse causality argument). It also indicates that the lending decline after the high rotation period is not driven by the fact that new and uninformed loan officer is managing the relationship. Overall, the findings are consistent with our argument that the ex-ante threat of rotation induced the changes in lending behavior.

FIGURE 4
LOANS OUTSTANDING RELATIVE TO OTHER BANK DEBT
AROUND THREE-YEAR RULE



These results provide a way of corroborating that the precision of all information which passes from the loan officer to the bank increases in response to the threat of rotation. This rules out the possibility that the ratings and informativeness patterns merely document a substitution in the way that loan officers report information to the bank around the high rotation period. If there was only a switch in the channel through which information was transmitted to the bank then we would not see any variation in the bank's capital allocation decisions around rotation.

B. Lending and Ratings Correlation

Total lending patterns mask an important feature of the effect of information in lending. Not all firms are expected to have access to larger amounts of credit when the bank has better information about their creditworthiness. The firms that are revealed to be of poor credit quality by the new information are likely to see either their availability of credit reduced, or the price of debt increase.

Either way, the amount of debt of downgraded firms is expected to drop when bad information is revealed. On the contrary, the expansion in total lending should be driven by the firms that were upgraded, or that were pooled in the same ratings than firms that were subsequently downgraded.

These cross-sectional patterns are verified in the data by looking at the within-firm correlation between debt outstanding and risk ratings. The time series correlation between the risk rating and the amount of debt should be negative for each individual firm in the sample if the above hypothesis is true. We estimate the within-firm correlation during the four quarters preceding the high rotation period predicted by the three-year rule using the following specification:

$$\begin{aligned} \ln(Debt)_{ijt,k} - \ln(Debt)_{ijt-1,k} &= \rho [irr_{ijt,k} - irr_{ijt-1,k}] + \\ &\quad \tau \left[\ln\left(\sum_{n \neq k} debt_{ijt,n}\right) - \ln\left(\sum_{n \neq k} debt_{ijt-1,n}\right) \right] + \\ &\quad [\alpha_{jt} - \alpha_{jt-1}] + \delta_{Industry \times t} + \varepsilon_{it} \end{aligned} \quad (7)$$

The variable on the left hand side is the log of debt outstanding of firm i with The Bank (k) at time t . The right hand side variable of interest is the measure of internal risk rating, irr . We control for the log of debt of firm i with all other banks, loan officer fixed effects and industry-time dummies as in other specifications. The first-differenced specification emphasizes that the parameter of interest ρ , represents the within-firm partial correlation between changes in debt and changes in the internal risk rating. This partial correlation nets out the effect of changes in firm specific shocks (e.g. demand for credit), unobserved loan officer heterogeneity and industry wide shocks. We expect $\rho < 0$: firms that receive a higher risk rating (credit rating downgraded) borrow less, all else being equal.

Table VII shows the estimated partial correlations in the sub-sample where the expected rotation probability increased. That is, during the four quarters prior to a predicted loan officer rotation as predicted by the three-year rule. This is the period sub-sample where we documented an increase in total lending (Figure 4, Panel A), in the average risk rating (Figure 3, Panels A and B) and in the predictive power of risk ratings (Table IV). The estimated correlation is negative and significant, using both measures of internal risk ratings (*Classification* and *ORR*).

The patterns in the data are consistent with the conjecture that banks are using the improved information to allocate credit optimally. More total lending is happening when the risk ratings are

more informative about the probability of default of borrowers. Although this lending expansion happens amidst a general increase in risk ratings, the lending is being allocated to the firms with better repayment prospects in the cross section.²¹

VIII. Conclusion

In this paper we have documented a bank policy that routinely reassigns loan officers to different borrowers. We have shown that the *ex-ante* threat of rotation induces loan officers to truthfully reveal bad information to the bank. Our empirical results allow us to evaluate how well the various theoretical accounts of rotation which exist in the literature can explain the phenomenon in this setting. Several papers have argued that rotation can induce agents to reveal private information at the end of their assignment by providing the principal a way of committing not to use this information against them in the future (Prescott and Townsend (2006), Arya and Mittendorf (2004), Hirao (1993), and Ickes and Samuelson (1987)). This is consistent with our finding that loan officers release additional information about their performance when they assess the probability of rotation to be high. However these arguments rely on the assumption that the principal is able and willing to commit not to use this information in future interactions. Instead we have argued that the bank will use this information to learn about the skill of the loan officer. As a consequence, the loan officer will consider the reputation effects of disclosing bad news at the end of a relationship. We have provided evidence to support this: a loan officer's future career within the bank is materially affected by the reputation consequences of the information she passes to the bank.

Prescott and Townsend (2006) argue that the cost of providing incentives to an agent will become increasingly costly as the agent acquires an informational advantage over the principal. Rotation can solve this problem by allowing the principal to contract with an agent who knows less. By this account we should expect to see the agency problem between the loan officer and the bank become more pronounced as time passes. On the contrary our results imply that agency problems are reduced as a loan officer approaches the end of her assignment. In particular we show that the officer assigns ratings with less bias and that these ratings are more informative for the state of the loan. This is consistent with our theory that it is the threat of the information which will be released after rotation

²¹The negative relation between ratings and lending in the data allows us to rule out the possibility that the downgrades we document around rotation are mechanically driven by an increase in lending. To the contrary, it is the firms with lower risk ratings who receive the increase in lending leading up to the three-year rule.

that provides additional incentives for the incumbent loan officer to reveal bad news to the bank.

Other explanations of job reassignment in the literature include Meyer (1994) and Ortega (2001) who argue that organizations may use a policy of worker rotation to increase their ability to learn about workers and tasks. By setting aside agency concerns, these theories cannot account for the fact that loan officers assign biased ratings to their clients and that rotation induces them to undo this. Our theory can be thought to build upon the fundamental assumption of these papers: that firms use performance to learn about the skills of their workers. Our paper documents the agency conflict that this learning can create and shows that rotation is one way to mitigate its effects.

We would expect to see rotation in an environment where workers are heterogenous and the skills required to perform specific tasks changes over time. Job reassignment would occur as the outcome of an ongoing matching process (Mortensen and Pissarides (1994)). While it is likely that some relationship turnover does arise due to matching, this theory cannot account for the results we have documented. Matching theories predict that a loan officer's skill should decline toward the end of a match, hence creating the need for reassignment. To the contrary, our results show that on average loan officers provide better information just before the high rotation period.

Tirole (1986) suggests that rotation may be used to prevent collusion between a delegated monitor and the agent she is monitoring. This can be thought of as a reinterpretation of the explanation for rotation that we have presented in this paper. Instead of skill, the state of the loan may be determined by the officer's willingness to collude with a borrower. A loan officer will be less willing to sustain collusion as the threat of rotation approaches for fear of being discovered by her successor.

IX. Mathematical Appendix

A. Equilibrium with No Persistence ($\phi = 0$)

Since $\phi = 0$ we can study a single period in isolation. Suppose that in equilibrium agent a truthfully reports bad news. The corresponding updated beliefs about the agent's type associated with a high and a low report can be found using Bayes' rule:

$$\hat{\mu}_a(r_l^a) = \frac{\mu}{\mu + (1 - \mu) \frac{(1 - p_l) q_l}{(1 - p_h) q_h}}, \text{ and } \hat{\mu}_a(r_h^a) = \frac{\mu}{\mu + (1 - \mu) \frac{1 - (1 - p_l) q_l}{1 - (1 - p_h) q_h}}.$$

Given these beliefs the agent will optimally choose to reveal bad news if and only if $\hat{\mu}_a(r_l^a) \geq \hat{\mu}_a(r_h^a)$. Simple manipulation shows this holds if and only if $\frac{1 - p_l}{1 - p_h} \leq \frac{q_h}{q_l}$. It can only be an equilibrium for the agent to reveal bad news when this condition holds.

Next suppose that we are in an equilibrium in which the agent conceals bad news. If the agent reports the high signal (as they always will in the conjectured equilibrium) this is completely uninformative and hence their reputation will remain at μ . However if the agent observes s_l and chooses to deviate and report the low signal then the belief about her type will be updated to $\hat{\mu}_a(r_l)$. The agent will optimally choose to conceal bad news (i.e. $\mu \geq \hat{\mu}_a(r_l^a)$) if and only if $\frac{1-p_l}{1-p_h} \geq \frac{q_h}{q_l}$. Thus it follows that when (1) holds that the unique perfect Bayesian Nash equilibrium is for the agent to conceal bad news. This establishes the claim in the paper.

B. Equilibrium with No Rotation

To show that the unique perfect Bayesian Nash equilibrium for a to conceal bad news in each period we need to rule out the possibility of any equilibrium where information is revealed in one or both periods (Steps 1-4) and show that given the equilibrium beliefs it is optimal for a to conceal bad news in both periods (Step 5).

B-1. Step 1: Ruling Out Truth in First or Second Period

We begin by showing that it cannot be a Bayesian equilibrium for a to truthfully reveal bad news in one period and not the other. We show this by contradiction. First suppose that it was a Bayesian equilibrium for the agent to reveal bad news in the first period and not in the second. Since no information is released in the second period this would be identical to the case where $\phi = 0$. Our assumption (1) ensures that the agent will optimally choose not to reveal information. Thus this cannot be a perfect Bayesian Nash equilibrium. A symmetric argument rules out the conjectured equilibrium where a conceals in the first period and reveals bad news in the second.

B-2. Step 2: Ruling Out Truth in Both Periods

Consider an equilibrium in which the agent chooses to reveal bad news in both periods. In order for this to be an equilibrium it must be that agent a is always willing to reveal bad news in the second period. This requires that

$$\hat{\mu}_a(r_h^a, r_l^a) \geq \hat{\mu}_a(r_h^a, r_h^a) \text{ and } \hat{\mu}_a(r_l^a, r_l^a) \geq \hat{\mu}_a(r_l^a, r_h^a). \quad (8)$$

By the law of iterated expectations the updated belief after the first period is equal to the (conditional) expectation of the belief at the end of the second period:

$$\begin{aligned} \hat{\mu}_a(r_h^a) &= \Pr(s_h^2|s_h^1)\hat{\mu}_a(r_h^a, r_h^a) + \Pr(s_l^2|s_h^1)\hat{\mu}_a(r_h^a, r_l^a) \\ \hat{\mu}_a(r_l^a) &= \Pr(s_h^2|s_l^1)\hat{\mu}_a(r_l^a, r_h^a) + \Pr(s_l^2|s_l^1)\hat{\mu}_a(r_l^a, r_l^a). \end{aligned}$$

We know that (1) ensures $\hat{\mu}_a(r_h^a) > \hat{\mu}_a(r_l^a)$. Using these facts lets ask whether agent a will optimally choose to reveal bad news after having seen s_l in the first period. If she reports bad news her expected reputation will be $\hat{\mu}_a(r_l^a)$ (by the law of iterated expectations). If agent a deviates from the equilibrium and suppresses the bad news then her expected reputation will be

$$\hat{\mu}_a^{Dev}(r_h^a) = \Pr(s_h^2|s_l^1)\hat{\mu}_a(r_h^a, r_h^a) + \Pr(s_l^2|s_l^1)\hat{\mu}_a(r_h^a, r_l^a).$$

Note that (8) ensures that the agent will optimally reveal bad news in the second period. Next, note that $\Pr(s_l^2|s_l^1) \geq \Pr(s_l^2|s_h^1)$ (intuitively this follows immediately from the fact that there is some auto-correlation in the state). Notice that

$$\hat{\mu}_a^{Dev}(r_h^a) - \hat{\mu}_a(r_h^a) = \hat{\mu}_a(r_h^a, r_h^a) [\Pr(s_h^2|s_l^1) - \Pr(s_h^2|s_h^1)] + \hat{\mu}_a(r_h^a, r_l^a) [\Pr(s_l^2|s_l^1) - \Pr(s_l^2|s_h^1)]$$

which can be re-written using the fact that $\Pr(s_h^2|s_l^1) = 1 - \Pr(s_h^2|s_l^1)$ and $\Pr(s_h^2|s_h^1) = 1 - \Pr(s_l^2|s_h^1)$

$$\hat{\mu}_a^{Dev}(r_h^a) - \hat{\mu}_a(r_h^a) = [\hat{\mu}_a(r_h^a, r_l^a) - \hat{\mu}_a(r_h^a, r_h^a)] [\Pr(s_l^2|s_l^1) - \Pr(s_l^2|s_h^1)] \geq 0$$

where the inequality comes directly from the fact that $\Pr(s_l^2|s_l^1) \geq \Pr(s_l^2|s_h^1)$ and (8). This establishes that

$$\hat{\mu}_a^{Dev}(r_h^a) \geq \hat{\mu}_a(r_h^a) \geq \hat{\mu}_a(r_l^a)$$

which implies that the agent will prefer to deviate from the equilibrium and hide bad news in the first period. Therefore it cannot be an equilibrium that the agent truthfully reports bad news in both periods.

B-3. Step 3: Contingent Strategy I

Consider the following candidate equilibrium: report truthfully in the first period and report truthfully in the second period if and only if r_h^1 was issued in the first period. Suppose now that the agent observes bad news in the first period but considers deviating and concealing this information. If the agent reports r_l^1 (as per the candidate equilibrium) her expected reputation will be $\hat{\mu}_a(r_l^a)$. Note that for it to be optimal for the agent to report truthfully in the second period after having revealed good news in the first period it must be that

$$\hat{\mu}_a(r_h^a, r_l^a) \geq \hat{\mu}_a(r_h^a, r_h^a).$$

So now suppose agent a observes s_l^1 and considers deviating by concealing this information. As per the equilibrium the agent will optimally choose to reveal bad news in the second period. Hence her expected reputation is

$$\hat{\mu}_a^{Dev}(r_h^a) = \Pr(s_h^2|s_l^1)\hat{\mu}_a(r_h^a, r_h^a) + \Pr(s_l^2|s_l^1)\hat{\mu}_a(r_h^a, r_l^a).$$

We have already showed in Case 2 that $\hat{\mu}_a^{Dev}(r_h^a) > \hat{\mu}_a(r_l^a)$ and hence it follows that the agent will choose to deviate in the first period and conceal bad news. Hence this cannot be an equilibrium.

B-4. Step 4: Contingent Strategy II

Consider the following candidate equilibrium: report truthfully in the first period and report truthfully in the second period if and only if r_l^1 was issued in the first period. Suppose now that the agent observes bad news in the first period but considers deviating and concealing this information. By the law of iterated expectations, if the agent reveals r_l^1 at $t = 1$ (as per the candidate equilibrium) her expected reputation will be $\hat{\mu}_a(r_l^a)$. However, if the agent deviates and releases r_h^1 at $t = 1$ and then follows the equilibrium strategy of concealing in the second period then her reputation will be $\hat{\mu}_a(r_h^a)$. We know that (1) ensures $\hat{\mu}_a(r_h^a) > \hat{\mu}_a(r_l^a)$ and hence this cannot be an equilibrium since the agent will choose to deviate and conceal bad news in the first period.

B-5. Step 5: Conceal in Both Periods

Consider an equilibrium where it is believed that the agent will not reveal bad news. If the agent does not deviate from this equilibrium then nothing will be learnt about her type and her final reputation will be unaltered (i.e. will remain μ).

Suppose that the agent has reported r_h at $t = 1$ (either truthfully or because she was concealing bad news). If the agent observes the low signal in the second period and decides to reveal this her updated belief will be $\hat{\mu}_a(r_l^a)$. Condition (1) ensures that $\mu > \hat{\mu}_a(r_l^a)$ and hence the agent has no incentive to deviate from the conjectured equilibrium by revealing bad news in the second period.

Finally, we can show that the agent does not have an incentive to deviate from the conjectured equilibrium and report bad news in the first period. To see this notice that if she does report bad news that the belief about her type will be updated (after $t = 1$) to $\hat{\mu}_a(r_l^a)$. By the law of iterated expectations this must be equal to her expected belief regardless of what is reported at $t = 2$. By (1) we know that $\hat{\mu}_a(r_l^a) < \mu$ and hence there is no incentive for the agent to deviate and report bad news in the first period.

C. Equilibrium With Rotation

We propose an equilibrium with rotation where both agents truthful reveal bad news and suppose that beliefs about each agent's type are formed using Bayes' rule presuming that agents act according to this equilibrium. We now look for conditions which ensure that both agents will optimally choose not to deviate from this proposed equilibrium. First we consider b 's incentives to truthfully report at $t = 2$ and then consider a 's incentive to report truthfully at $t = 1$.

C-1. Statement of Equilibrium Posterior Beliefs About b 's Type

For agent b there are four possible updated posterior beliefs about her reputation. These are found using Bayes' rule:

$$\begin{aligned}\hat{\mu}_b(r_l^a, r_l^b) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{(\phi + (1 - \phi)(1 - p_l))q_l}{(\phi + (1 - \phi)(1 - p_h))q_h} \right]}, \\ \hat{\mu}_b(r_l^a, r_h^b) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - \phi)p_l + [\phi + (1 - \phi)(1 - p_l)](1 - q_l)}{(1 - \phi)p_h + [\phi + (1 - \phi)(1 - p_h)](1 - q_h)} \right]}, \\ \hat{\mu}_b(r_h^a, r_l^b) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{q_l[(1 - \phi)(1 - p_l)(\mu p_h + (1 - \mu)p_l) + (1 - (1 - \phi)p_l)(\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l))]}{q_h[(1 - \phi)(1 - p_h)(\mu p_h + (1 - \mu)p_l) + (1 - (1 - \phi)p_h)(\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l))]} \right]}, \text{ and} \\ \hat{\mu}_b(r_h^a, r_h^b) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{(\mu p_h + (1 - \mu)p_l)(1 - (1 - \phi)(1 - p_l)q_l) + (\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l))[1 - (1 - (1 - \phi)p_l)q_l]}{(\mu p_h + (1 - \mu)p_l)(1 - (1 - \phi)(1 - p_h)q_h) + (\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l))[1 - (1 - (1 - \phi)p_h)q_h]} \right]}.\end{aligned}$$

C-2. Evaluating Beliefs for $\phi = 0$ and $\phi = 1$

(1) $\hat{\mu}_b(r_l^a, r_l^b)$. For different levels of persistence:

$$\hat{\mu}_b(r_l^a, r_l^b; \phi = 1) = \frac{\mu}{\mu + (1 - \mu) \frac{q_l}{q_h}} \text{ and } \hat{\mu}_b(r_l^a, r_l^b; \phi = 0) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - p_l)q_l}{(1 - p_h)q_h} \right]}.$$

Rearranging condition (1) yields that $\frac{(1 - p_l)q_l}{(1 - p_h)q_h} > 1$ and therefore it must be that

$$\hat{\mu}_b(r_l, r_l; \phi = 1) > \mu > \hat{\mu}_b(r_l, r_l; \phi = 0).$$

(2) $\hat{\mu}_b(r_l^a, r_h^b)$. For different levels of persistence:

$$\hat{\mu}_b(r_l^a, r_h^b; \phi = 1) = \frac{\mu}{\mu + (1 - \mu) \frac{(1 - q_l)}{(1 - q_h)}} \text{ and } \hat{\mu}_b(r_l^a, r_h^b; \phi = 0) = \frac{\mu}{\mu + (1 - \mu) \frac{1 - (1 - p_l)q_l}{1 - (1 - p_h)q_h}}.$$

Notice that

$$\hat{\mu}_b(r_l^a, r_h^b; \phi = 0) > \mu > \hat{\mu}_b(r_l^a, r_h^b; \phi = 1)$$

where the first inequality follows from the fact that (1) implies

$$\frac{1 - (1 - p_l) q_l}{1 - (1 - p_h) q_h} = \frac{\frac{1}{(1 - p_h) q_h} - \frac{(1 - p_l) q_l}{(1 - p_h) q_h}}{\frac{1}{(1 - p_h) q_h} - 1} < 1.$$

(3) $\hat{\mu}_b(r_h^a, r_l^b)$. For different levels of persistence:

$$\hat{\mu}_b(r_h^a, r_l^b; \phi = 1) = \frac{\mu}{\mu + (1 - \mu) \frac{q_l}{q_h}} \text{ and } \hat{\mu}_b(r_h^a, r_l^b; \phi = 0) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - p_l) q_l}{(1 - p_h) q_h} \right]}.$$

Notice that

$$\hat{\mu}_b(r_h^a, r_l^b; \phi = 1) > \mu > \hat{\mu}_b(r_h^a, r_l^b; \phi = 0)$$

where the second inequality comes from (1).

(4) $\hat{\mu}_b(r_h^a, r_h^b)$. For different levels of persistence:

$$\begin{aligned} \hat{\mu}_b(r_h^a, r_h^b; \phi = 1) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{(\mu p_h + (1 - \mu) p_l) + (\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l))(1 - q_h)}{(\mu p_h + (1 - \mu) p_l) + (\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l))(1 - q_h)} \right]} \\ \hat{\mu}_b(r_h^a, r_h^b; \phi = 0) &= \frac{\mu}{\mu + (1 - \mu) \frac{1 - (1 - p_l) q_l}{1 - (1 - p_h) q_h}} \end{aligned}$$

Notice that

$$\hat{\mu}_b(r_h, r_h; \phi = 0) > \mu > \hat{\mu}_b(r_h, r_h; \phi = 1)$$

where the second inequality follows by noticing that the fraction in the denominator of $\hat{\mu}_b(r_h^a, r_h^b; \phi = 1)$ has all terms in the numerator and denominator the same apart from $(1 - q_l)$ and $(1 - q_h)$.

C-3. Sufficient Conditions to Ensure b Reports Truthfully

In order for b to always truthfully reveal bad news at $t = 2$ it must be the case that

$$\Delta \hat{\mu}_b(r_h^a, \cdot) \equiv \hat{\mu}_b(r_h^a, r_l^b) - \hat{\mu}_b(r_h^a, r_h^b) \geq 0, \text{ and}$$

$$\Delta \hat{\mu}_b(r_l^a, \cdot) \equiv \hat{\mu}_b(r_l^a, r_l^b) - \hat{\mu}_b(r_l^a, r_h^b) \geq 0.$$

Both $\Delta \hat{\mu}_b(r_h^a, \cdot)$ and $\Delta \hat{\mu}_b(r_l^a, \cdot)$ are continuous in ϕ .²²

Next, observe that

$$\begin{aligned} \Delta \hat{\mu}_b(r_h^a, \cdot; \phi = 1) &= \hat{\mu}_b(r_h^a, r_l^b; \phi = 1) - \hat{\mu}_b(r_h^a, r_h^b; \phi = 1) > 0 \\ \Delta \hat{\mu}_b(r_l^a, \cdot; \phi = 1) &= \hat{\mu}_b(r_l^a, r_l^b; \phi = 1) - \hat{\mu}_b(r_l^a, r_h^b; \phi = 1) > 0 \end{aligned}$$

where the inequalities follow from the orderings established above. Similarly, observe that

$$\begin{aligned} \Delta \hat{\mu}_b(r_h^a, \cdot; \phi = 0) &= \hat{\mu}_b(r_h^a, r_l^b; \phi = 0) - \hat{\mu}_b(r_h^a, r_h^b; \phi = 0) < 0 \\ \Delta \hat{\mu}_b(r_l^a, \cdot; \phi = 0) &= \hat{\mu}_b(r_l^a, r_l^b; \phi = 0) - \hat{\mu}_b(r_l^a, r_h^b; \phi = 0) < 0. \end{aligned}$$

²²It is straightforward to show that (1) implies that $\frac{\partial \Delta \hat{\mu}_b(r_h^a, \cdot)}{\partial \phi} > 0$ and $\frac{\partial \Delta \hat{\mu}_b(r_l^a, \cdot)}{\partial \phi} > 0$.

It follows immediately that there exists a $\bar{\phi}_b \in (0, 1)$ such that agent b will always optimally choose to reveal bad news if $\phi \geq \bar{\phi}_b$.

C-4. Statement of Equilibrium Posterior Beliefs About a 's Type

For agent a there are four possible updated posterior beliefs about her reputation. These are found using Bayes' rule:

$$\begin{aligned}\hat{\mu}_a(r_l^a, r_l^b) &= \hat{\mu}_a(r_l^a, r_h^b) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - p_l) q_l}{(1 - p_h) q_h} \right]} \\ \hat{\mu}_a(r_h^a, r_l^b) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{p_l(1 - \phi)[\mu(1 - p_h)q_h + (1 - \mu)(1 - p_l)q_l] + (1 - p_l)(1 - q_l)[\mu(\phi + (1 - \phi)(1 - p_h))q_h + (1 - \mu)(\phi + (1 - \phi)(1 - p_l))q_l]}{p_h(1 - \phi)[\mu(1 - p_h)q_h + (1 - \mu)(1 - p_l)q_l] + (1 - p_h)(1 - q_h)[\mu(\phi + (1 - \phi)(1 - p_h))q_h + (1 - \mu)(\phi + (1 - \phi)(1 - p_l))q_l]} \right]} \\ \hat{\mu}_a(r_h^a, r_h^b) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{p_l[1 - (1 - \phi)[\mu(1 - p_h)q_h + (1 - \mu)(1 - p_l)q_l]] + (1 - p_l)(1 - q_l)[1 - \mu(1 - (1 - \phi)p_h)q_h - (1 - \mu)(1 - (1 - \phi)p_l)q_l]}{p_h[1 - (1 - \phi)[\mu(1 - p_h)q_h + (1 - \mu)(1 - p_l)q_l]] + (1 - p_h)(1 - q_h)[1 - \mu(1 - (1 - \phi)p_h)q_h - (1 - \mu)(1 - (1 - \phi)p_l)q_l]} \right]}\end{aligned}$$

C-5. Sufficient Conditions to Ensure a Reports Truthfully

Agent a will choose to reveal bad news in the first period, r_l^1 , if and only if this raises her expected reputation. Suppose agent a has observed s_l^1 . By the law of iterated expectations, her expected reputation if she reports r_l^1 is

$$E_1(\hat{\mu}_a; r_l^a, s_l^1) = \hat{\mu}_a(r_l^a) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - p_l) q_l}{(1 - p_h) q_h} \right]}.$$

Conversely if she conceals the bad news and reports r_h^1 her expected reputation is

$$E_1(\hat{\mu}_a; r_h^a, s_l^1) = \hat{\mu}_a(r_h^a, r_l^b) + \left(\hat{\mu}_a(r_h^a, r_l^b) - \hat{\mu}_a(r_h^a, r_h^b) \right) [\mu(1 - (1 - \phi)p_h)q_h + (1 - \mu)(1 - (1 - \phi)p_l)q_l]$$

So agent a will choose to reveal bad news if and only if

$$\Delta\hat{\mu}_a = E_1(\hat{\mu}_a; r_l^a, s_l^1) - E_1(\hat{\mu}_a; r_h^a, s_l^1) \geq 0$$

$\Delta\hat{\mu}_a$ is continuous in $\phi \in [0, 1]$. To show that the incentive for agent a to report truthfully is increasing in the level of persistence compare $\Delta\hat{\mu}_a(\phi = 1)$ and $\Delta\hat{\mu}_a(\phi = 0)$:

$$\begin{aligned}\Delta\hat{\mu}_a(\phi = 1) - \Delta\hat{\mu}_a(\phi = 0) &= \hat{\mu}_a(r_h, \cdot; \phi = 0) - [\mu q_h + (1 - \mu) q_l] \hat{\mu}_a(r_h, r_l; \phi = 1) \\ &\quad - [1 - \mu q_h - (1 - \mu) q_l] \hat{\mu}_a(r_h, r_h; \phi = 1)\end{aligned}$$

where

$$\hat{\mu}_a(r_h^a, \cdot; \phi = 0) \equiv \hat{\mu}_a(r_h^a, r_l^b; \phi = 0) = \hat{\mu}_a(r_h^a, r_h^b; \phi = 0) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{1 - (1 - \phi)q_l}{1 - (1 - \phi)q_h} \right]}$$

We can write

$$\begin{aligned}\Delta\hat{\mu}_a(\phi = 1) - \Delta\hat{\mu}_a(\phi = 0) &= \hat{\mu}_a(r_h^a, \cdot; \phi = 0) - [\mu q_h + (1 - \mu) q_l] \hat{\mu}_a(r_h^a, r_l^b; \phi = 1) \\ &\quad - [1 - \mu q_h - (1 - \mu) q_l] \hat{\mu}_a(r_h^a, r_h^b; \phi = 1)\end{aligned}$$

Note that for any level of persistence that in a truth telling equilibrium $\hat{\mu}_a(r_h, \cdot; \phi = 0)$ is the updated belief about a 's type when she reports r_h^1 . By the law of iterated expectations this must be equal to the expected belief conditional upon agent b 's report. Thus, if $\phi = 1$ then it must be that

$$\begin{aligned}\hat{\mu}_a(r_h, \cdot; \phi = 0) &= \Pr(s_l^2 | s_h^1; \phi = 1) \hat{\mu}_a(r_h^a, r_l^b; \phi = 1) \\ &\quad + \Pr(s_h^2 | s_h^1; \phi = 1) \hat{\mu}_a(r_h^a, r_h^b; \phi = 1)\end{aligned}$$

where

$$\Pr(s_l^2 | s_h^1; \phi = 1) = \frac{[\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l)][\mu q_h + (1 - \mu)q_l]}{[\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l)] + [\mu p_h + (1 - \mu)p_l]}$$

and

$$\Pr(s_h^2 | s_h^1; \phi = 1) = 1 - \Pr(s_l^2 | s_h^1; \phi = 1).$$

This allows us to write

$$\Delta\hat{\mu}_a(\phi = 1) - \Delta\hat{\mu}_a(\phi = 0) = [\Pr(s_l^2 | s_h^1; \phi = 1) - [\mu q_h + (1 - \mu)q_l]] [\hat{\mu}_a(r_h^a, r_l^b; \phi = 1) - \hat{\mu}_a(r_h^a, r_h^b; \phi = 1)]$$

Next, observe that (1) ensures that

$$\hat{\mu}_a(r_h, r_h; \phi = 1) > \mu > \hat{\mu}_a(r_h, r_l; \phi = 1).$$

So to show

$$\Delta\hat{\mu}_a(\phi = 1) - \Delta\hat{\mu}_a(\phi = 0) > 0$$

it is sufficient to show that

$$\Pr(s_l^2 | s_h^1; \phi = 1) < \mu q_h + (1 - \mu)q_l$$

which can be written as

$$1 > \frac{[\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l)]}{[\mu(1 - p_h)(1 - q_h) + (1 - \mu)(1 - p_l)(1 - q_l)] + [\mu p_h + (1 - \mu)p_l]}$$

which must hold since $[\mu p_h + (1 - \mu)p_l]$ is positive. This establishes that $\Delta\hat{\mu}_a(\phi = 1) > \Delta\hat{\mu}_a(\phi = 0)$.

If there is zero persistence ($\phi = 0$) then:

$$\Delta\hat{\mu}_a(\phi = 0) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - p_l)q_l}{(1 - p_h)q_h} \right]} - \frac{\mu}{\mu + (1 - \mu) \left[\frac{1 - (1 - p_l)q_l}{1 - (1 - p_h)q_h} \right]} < 0$$

where the inequality follows directly from (1). So with zero persistence agent a strictly prefers to conceal bad news. Now suppose that there is full persistence ($\phi = 1$)

$$\begin{aligned}\Delta\hat{\mu}_a(\phi = 1) &= \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - p_l)q_l}{(1 - p_h)q_h} \right]} - [\mu q_h + (1 - \mu)q_l] \frac{\mu}{\mu + (1 - \mu) \left[\frac{(1 - p_l)(1 - q_l)}{(1 - p_h)(1 - q_h)} \right]} \\ &\quad - [1 - \mu q_h - (1 - \mu)q_l] \frac{\mu}{\mu + (1 - \mu) \left[\frac{p_l + (1 - p_l)(1 - q_l)[1 - \mu q_h - (1 - \mu)q_l]}{p_h + (1 - p_h)(1 - q_h)[1 - \mu q_h - (1 - \mu)q_l]} \right]}\end{aligned}$$

This can be either positive or negative. We will assume that parameters are such that

$$\Delta\hat{\mu}_a(\phi = 1) > 0 \tag{9}$$

to ensure that it is possible to induce truth-telling for some values of ϕ . This holds when q_h and q_l are sufficiently large. Since $\Delta\hat{\mu}_a$ is continuous in $\phi \in [0, 1]$ then this ensures that there exists a $\bar{\phi}_a \in (0, 1)$ such that agent a will always optimally choose to reveal bad news in the model with rotation if $\phi \geq \bar{\phi}_a$. Therefore set $\bar{\phi} = \max \{\bar{\phi}_a, \bar{\phi}_b\} \in (0, 1)$ and it follows that if $\phi \geq \bar{\phi}$ then it will be a perfect Bayesian Nash equilibrium for both agent a and b to always report bad news in the model with rotation. This establishes Proposition 1.

References

Aghion, P., and J. Tirole, 1997, Formal and Real Authority in Organizations, *The Journal of Political Economy*, 105, 1-29.

Arrunada, B., and C. Paz-Ares, 1997, Mandatory Rotation of Company Auditors: A Critical Examination, *International Review of Law and Economics*, 17, 31-61.

Arya, A., and B. Mittendorf, 2004, Using Job Rotation to Extract Employee Information, *Journal of Law, Economics, & Organization*, 20, 400-414.

Berk, J., and R. C. Green, 2004, Mutual Fund Flows and Performance in Rational Markets, *Journal of Political Economy*, 112, 1269-1295.

Chevalier, J., and G. Ellison, 1999, Career Concerns and Mutual Fund Managers, *Quarterly Journal of Economics*, 114, 389-432.

Colc, S., 2007, Fixing Market Failures or Fixing Elections? Elections, Banks and Agricultural Lending in India, mimeo, Harvard Business School.

Dopuch, N., R. King, and R. Schwartz, 2001, An Experimental Investigation of Retention and Rotation Requirements, *Journal of Accounting Research*, 39, 93-117.

Hirao, Y., 1993, Task Assignment and Agency Structures, *Journal of Economics and Management Strategy*, 2, 299-323.

Holmstrom, B., 1982, Moral Hazard in Teams, *Bell Journal of Economics*, 13, 324-340.

Holmstrom, B., and P. Milgrom, 1987, Aggregation and Linearity in the Provision of Intertemporal Incentives, *Econometrica*, 55, 303-28.

Ickes, B. W., and L. Samuelson, 1987, Job Transfers and Incentives in Complex Organizations: Thwarting the Ratchet Effect, *RAND Journal of Economics*, 18, 275-286.

Khemani, S., 2004, Political Cycles in a Developing Economy: Effect of Elections in the Indian States, *Journal of Development Economics*, 73, 125-54.

Leland, H., and D. Pyle, 1977, Informational Asymmetries, Financial Structure, and Financial Intermediation, *Journal of Finance*, 32, 371 - 387.

Liberti, J., 2004, Initiative, Incentives and Soft Information. How Does Delegation Impact the Role of Bank Relationship Managers?, mimeo, Kellogg School of Management.

Liberti, J., and A. Mian, 2006, Estimating The Effect Of Hierarchies On Information Use, mimeo, Chicago GSB.

Meyer, M., 1994, The Dynamics of Learning with Team Production: Implications for Task Assignment, *Quarterly Journal of Economics*, 109, 1157–1184.

Mortensen, D. T., and C. A. Pissarides, 1994, Job Creation and Job Destruction in the Theory of Unemployment, *The Review of Economic Studies*, 61, 397-415.

Myers, J., L. Myers, and T. Omer, 2003, Exploring the Term of the Auditor-Client Relationship and the Quality of Earnings: A Case for Mandatory Auditor Rotation?, *Accounting Review*, 78, 779-799.

Myers, S., and N. Majluf, 1984, Corporate Financing and Investment Decisions when Firms have Information that Investors do not Have, *Journal of Financial Economics*, 13, 187-222.

Ortega, J., 2001, Job Rotation as a Learning Mechanism, *Management Science*, 47, 1361-70.

Osterman, P., 1994, How Common is Workplace Transformation and who Adopts it? , *Industrial and Labor Relations Review*, 47, 173–188.

Osterman, P., 2000, Work Reorganization in an Era of Restructuring: Trends in Diffusion and Effects on Employee Welfare, *Industrial and Labor Relations Review*, 53, 179–196.

Paravisini, D., 2006, Local Bank Financial Constraints and Firm Access to External Finance, mimeo, Columbia GSB.

Prescott, E. S., and R. Townsend, 2006, Private Information and Intertemporal Job Assignments, *Review of Economic Studies*, 73, 531-548.

Rajan, R. G., 1994, Why Bank Credit Policies Fluctuate: A Theory and Some Evidence, *Quarterly Journal of Economics*, 109, 399-441.

Stein, J., 2002, Information Production and Capital Allocation: Decentralized Vs. Hierarchical Firms, *Journal of Finance*, 57, 1891-1921.

Stiglitz, J., and A. Weiss, 1981, Credit Rationing in Markets with Imperfect Information, *American Economic Review*, 71, 393-410.

Tirole, J., 1986, Hierarchies and Bureaucracies: On the Role of Collusion in Organizations, *Journal of Law, Economics, and Organization*, 2, 181-214.

TABLE I
SELECTED SUMMARY STATISTICS

Variable*	Mean	Median	SD	Min	Max	No of Obs.
LEVEL OF BORROWING (\$000)						
Approved Loan Amount	1,111	628	2,972	0	285,564	27,243
Outstanding Amount	493	201	1,273	0	72,205	27,018
Outstanding Reported by CB	513	226	936	0	34,922	22,659
Total Bank Debt Reported by CB	2,941	1,336	4,882	0	83,139	22,659
Debt Bank/Total Debt	0.27	0.17	0.27	0	1	22,659
INTERNAL BANK RATINGS/DEFAULT						
Internal Risk Rating (ORR)	18.67	17.00	4.06	0	29	29,371
Internal Risk Rating (Classification)	1.54	1.00	1.11	1	5	30,358
Default (based on Classification)	0.09	0.00	0.28	0	1	34,239
Default Reported by CB	0.08	0.00	0.26	0	1	34,239
Default with Other Banks Reported by CB	0.09	0.00	0.29	0	1	34,239
Number of Other Lending Relationships	7.52	7.00	4.08	1	34	22,659
Length of Relationship (Months) **	17.49	7.00	29.07	0	254	32,665

* See variable definitions in the data appendix. All statistics are calculated over the small business lending, pre-2002 subsample.

** The length of relationship measures the average over all months. The average firm in the sample is observed for 67 months.

TABLE II
DESCRIPTIVE STATISTICS ON LOAN OFFICER ROTATION

Variable*	Mean	Median	SD	Min	Max
LOAN OFFICER STATISTICS					
Number of Firms in Loan Officer Portfolio	25.57	10.0	36.14	1	221
Length of Loan Officer-Firm Relationship	22.1	18.0	18.04	1	84
FIRM STATISTICS					
Number of Relationships per Firm	3.04	3.00	1.29	1.00	7.00
Number of Different Loan Officers per Firm	3.19	3.00	1.43	1.00	9.00
%Firms Repeat Loan Officer (Non-Consecutive)	28%	0.0	0.5	0.0	1.0

*Statistics based on a sample of 1,248 firms, 100 loan officers and 4,191 non-censored firm-loan officer relationships

TABLE III
PARAMETRIC REDUCED FORM - EFFECT OF LOAN OFFICER ROTATION ON OUTCOMES

Dependent Variable	Internal Risk Rating				Lending ln(O/S Amount _t)	
	Classification _t		Risk Rating _t			
	(1)	(2)	(3)	(4)		
Months to predicted Loan Officer Δ	0.010*** (0.003)	0.001 (0.005)	0.011*** (0.003)	0.004 (0.004)	0.020* (0.011)	
Months since predicted Loan Officer Δ	-0.011 (0.011)	-0.01 (0.011)	-0.027** (0.013)	-0.027** (0.013)	-0.084* (0.048)	
Pre-Predicted Change Dummy	0.023 (0.039)	-0.043 (0.075)	0.027 (0.029)	-0.024 (0.051)	0.484*** (0.131)	
Post-Predicted Change Dummy	0.009 (0.059)	0.009 (0.060)	0.056 (0.058)	0.054 (0.058)	0.207 (0.274)	
Months to predicted Loan Officer Δ x High Origination Dummy			0.014** (0.006)		0.009* (0.005)	
Pre-Predicted Change Dummy x High Origination Dummy		0.105 (0.081)		0.074 (0.058)		
High Origination Dummy			-0.074*** (0.027)		0.014 (0.021)	
Loan Officer FE	Yes	Yes	Yes	Yes	Yes	
Firm FE	Yes	Yes	Yes	Yes	Yes	
Credit Information Central Bank	Yes	Yes	Yes	Yes	Yes	
Observations	21,400	21,400	21,389	21,389	20,272	
R-Sq	0.79	0.79	0.73	0.73	0.20	

Estimates of the coefficients of specification (4) in the paper:

$$y_{ijt,k} = \eta_{SB}(\tilde{t}_{ijt} \times \text{Before3 YR}_{ijt}) + \eta_{SA}(\tilde{t}_{ijt} \times \text{After3 YR}_{ijt}) + \eta_{LB}(\text{Before3 YR}_{ijt}) + \eta_{LA}(\text{After3 YR}_{ijt}) + \phi y_{ijt,-k} + \phi_i t + \alpha_i + \alpha_{jt} + \delta_{Industry \times t} + \varepsilon_{ijt}$$

The left hand side variable is a loan outcome of firm i , managed by loan officer j at month t . Columns (1) through (4) use internal measures of risk ratings and column 5 uses the log of the loan outstanding with bank k as the dependent variable. Predicted change refers to the month at which the loan officer of firm i is predicted to change according to the three-year rule. The prediction is based on the last officer rotation observed for firm i in the data. Months to (since) predicted Loan Officer Change refers to an event clock that counts the months before (after) a predicted change. Pre-Predicted (Post) Change Dummy is a dummy equal to one during the 12 months prior (after) to a predicted loan officer rotation. Credit Information Central Bank indicates that the specification includes a control variable that accounts for the outcome of interest with all other banks in the financial system using Central Bank data. High Origination Dummy is a dummy that turns to one for every firm-loan officer-month cell that belongs to a relationship that has an origination rate above the median.

*, ** and *** indicate that the point estimate is statistically significant at the 10, 5 and 1 percent levels respectively. All significant estimates are in bold typeface. The standard errors are presented in parentheses and are estimated accounting for heteroskedasticity and clustering at the firm level

TABLE IV
INFORMATION CONTENT OF INTERNAL RISK RATINGS

Estimation Method:	OLS						2SLS						2SLS					
	Loan Officer Δ	Before Loan Officer Δ	After Loan Officer Δ	Loan Officer Δ	Before Loan Officer Δ	After Loan Officer Δ	Before/After Loan Officer Δ											
Sample period:																		
Default at t+12	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)						
Risk Rating (ORR Standardized)	0.106*** (0.013)	0.091*** (0.014)	0.048*** (0.016)	0.354 (0.489)	1.066*** (0.447)	-1.026 (0.786)	1.263** (0.469)	1.246** (0.409)	1.053 (0.989)	-0.304 (0.349)	0.969** (0.336)	-0.912 (0.733)						
Industry x Month Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes						
Loan Officer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes						
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes						
Risk Rating Central Bank	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes						
Observations	16,454	12,938	8,885	16,454	12,938	8,885	16,454	12,938	8,885	16,454	12,938	8,885						
R-Sq	0.72	0.80	0.63															

Estimates of the coefficients of the linear probability model in specification (5) in the paper:

$$\text{Default}_{ijt+12,k} = \gamma(\text{irr}_{ijt,k}) + \phi(\text{wrr}_{ijt,-k}) - \phi(\text{wrr}_{ijt,k}) + \alpha_i + \alpha_{jt} + \delta_{Industryjt} + \sigma_{it}$$

The left hand side variable is a dummy equal to one if firm i , managed by loan officer j , is in detail: at month t . Risk Rating is an internal risk rating of The Bank (ORR). Predicted change refers to the month at which the loan officer of firm i is predicted to change according to the three-year rule. The prediction is based on the last officer rotation observed for firm i in the data. Before (after) Loan Officer Change refers to the subsample of 12 months before (after) a predicted loan officer change. The 2SLS estimates use Time to (after) Predicted Loan Officer Change as an instrument for changes in risk ratings. Thus, the specification in Table III can be interpreted as the first-stage regression. Specifications 7-9 (10-12) use the instrument interacted with a dummy equal to one if the loan officer of firm i changed (did not change) at three years. Risk Rating Central Bank indicates that the specification includes a control variable that accounts for the risk rating assigned to firm i by all other banks in the financial system at time t , weighted by their outstanding loans using Central Bank information.

*, ** and *** indicate that the point estimate is statistically significant at the 10, 5 and 1 percent levels respectively. All significant estimates are in bold typeface. The standard errors are presented in parentheses and are estimated accounting for heteroskedasticity and clustering at the firm level.

TABLE V
FIRST STAGE - REPUTATION AND LOAN OFFICER'S ASSETS UNDER MANAGEMENT

Dependent variable:	Bad Reputation Index		Good Reputation Index
	(1)	(2)	
Bad Reputation Instrument	1.927*** (0.687)		0.058 (0.267)
Good Reputation Instrument	-0.025 (0.737)		2.936*** (0.614)
Weighted Risk Rating	0.641 (1.256)		4.609 (3.266)
Loan Officer FE	Yes		Yes
Year Dummies	Yes		Yes
Observations	1,242		1,242
R-Sq	0.50		0.81

Estimates of the coefficients of the first stage of specification (6) in the paper. Column 1 (2) presents the estimates of the bad (good) reputation index on the bad and god reputation instruments, loan officer fixed effects, year dummies and the loan officers portfolio weighted risk rating. The index (instrument) of bad reputation is the fraction of the number of events, up to time t , in which the firm under assignment to loan officer j is downgraded by the new loan officer during the 6 months after a rotation (a predicted rotation), scaled by the number of rotations (predicted rotations) a loan officer has experienced up to time t . The index (instrument) of good reputation is the number of times a loan officer i has downgraded a firm during the six months after a rotation (predicted rotation) up to time t , scaled by the number of rotations (predicted rotations) experienced by loan officer j up to time t .

*, ** and *** indicate that the point estimate is statistically significant at the 10, 5 and 1 percent levels respectively. All significant estimates are in bold typeface. The standard errors are presented in parentheses and are estimated accounting for heteroskedasticity and clustering at the loan officer level

TABLE VI
REPUTATION AND LOAN OFFICER'S ASSETS UNDER MANAGEMENT

Estimation Method:	OLS		2SLS	
	ln(O/S)	ln(# firms)	ln(O/S)	ln(# firms)
	(1)	(2)	(3)	(4)
Bad Reputation Index (x100) _{t-6}	-0.154* (0.086)	-0.070*** (0.011)	-0.325* (0.114)	-0.157*** (0.032)
Good Reputation Index (x100) _{t-6}	0.060** (0.026)	0.032*** (0.025)	0.06 (0.164)	0.02 (0.037)
Weighted Risk Rating _{t-6}	0.903 (1.230)	-0.1 (0.325)	1.111 (0.789)	0.158 (0.247)
Loan Officer FE	Yes	Yes	Yes	Yes
Year Dummies	Yes	Yes	Yes	Yes
Observations	1,039	1,039	803	866
R-Sq	0.70	0.89		

Estimates of the coefficients of specification (6) in the paper:

$$\ln(Y_{jt}) = \theta_{Bad} \frac{NBad_{jt-6}}{NRot_{jt-6}} + \theta_{Good} \frac{NBad_{jt-6}}{NRot_{jt-6}} + \gamma X_{jt} + \alpha_j + \alpha_t + \nu_{it}$$

The left hand side variable is the log of a measure of assets under management of loan officer j at time t . Three measures of assets under management are used: the total amount of loans approved, the total amount of loans outstanding and the number of firms under management of a loan officer at any given time t . The index (instrument) of bad reputation is the fraction of the number of events, up to time t , in which the firm under assignment to loan officer j is downgraded by the new loan officer during the 6 months after a rotation (a predicted rotation), scaled by the number of rotations (predicted rotations) a loan officer has experienced up to time t . The index (instrument) of good reputation is the number of times a loan officer i has downgraded a firm during the six months after a rotation (predicted rotation) up to time t , scaled by the number of rotations (predicted rotations) experienced by loan officer j up to time t . The reputation indexes (and instruments) are lagged 6 months. All specifications include the average risk rating of the firms under management, weighted by the amount of loans outstanding of each firm.

*, ** and *** indicate that the point estimate is statistically significant at the 10, 5 and 1 percent levels respectively. All significant estimates are in bold typeface. The standard errors are presented in parentheses and are estimated accounting for heteroskedasticity and clustering at the loan officer level

TABLE VII
WITHIN-FIRM CORRELATION BETWEEN LENDING AND RISK RATINGS

Dependent variable:	$\Delta \ln(O/S \text{ Amount})$	
	(1)	(2)
Δ Risk Rating (ORR Standardized)	-0.114* (0.068)	
Δ Risk Rating (Classification Standardized)		-0.474** (0.204)
$\Delta \ln(\text{Debt other banks})$	0.012 (0.010)	-0.017 (0.017)
Industry x Month Dummies	Yes	Yes
Loan Officer FE	Yes	Yes
Firm FE	Yes	Yes
Credit Information Central Bank	Yes	Yes
Observations	4,063	4,063
R-Sq	0.07	0.06

Estimates of the coefficients of specification (7) in the paper:

$$\ln(\text{Debt})_{ijt,k} - \ln(\text{Debt})_{ijt-1,k} = \rho[\text{irr}_{ijt,k} - \text{irr}_{ijt-1,k}] + \tau \left[\ln\left(\sum_{n \neq k} \text{debt}_{ijt,n}\right) - \ln\left(\sum_{n \neq k} \text{debt}_{ijt-1,n}\right) \right] + [\alpha_{jt} - \alpha_{jt-1}] + \delta_{Industry \times t} + \varepsilon_{it}$$

The variable on the left hand side is the log of debt outstanding of firm i with The Bank (k) at time t. The right hand side variable of interest is the measure of internal risk rating (see definitions of ORR and Classification in the data Appendix). Controls include the log of debt of firm i with all other banks at time t, loan officer fixed effects and industry-time dummies. The first-differenced specification emphasizes that the parameter of interest ρ , represents the within-firm partial correlation between changes in debt and changes in the internal risk rating. The parameters are estimated during the sub sample period defined by the four quarters preceding the high rotation period predicted by the three year rule

*, ** and *** indicate that the point estimate is statistically significant at the 10, 5 and 1 percent levels respectively. All significant estimates are in bold typeface. The standard errors are presented in parentheses and are estimated accounting for heteroskedasticity and clustering at the firm level