

The Incentive Effect of Scores: Randomized Evidence from Credit Committees *

Daniel Paravisini Antoinette Schoar
LSE MIT

July 2, 2015

Abstract

We design a randomized controlled trial to evaluate the adoption of credit scoring with a bank that uses *soft* information in small businesses lending. We find that credit scores improve the productivity of credit committees, reduce managerial involvement in the loan approval process, and increase the profitability of lending. Credit committee members' effort and output also increase when they *anticipate* the score becoming available, indicating that scores improve incentives to use existing information. Our results imply that credit scores improve the efficiency and decentralize decision-making in loan production, which has implications for the optimal organization of banks.

*We are grateful to Heski Bar-Isaac, Greg Fischer, Luis Garicano, Danielle Li, Ulrike Malmendier, Andrea Prat, Anjan Thakor, John Van Reenen, and the Arizona State University, Boston University, Columbia University, Dartmouth College - Tuck School of Business, HEC School of Management, Imperial College Business School, European Bank for Reconstruction and Development, Federal Reserve Bank of New York, Helsinki Center of Economic Research, Massachusetts Institute of Technology - Sloan School of Business, Northwestern University - Kellogg School of Management, Nova School of Business and Economics, Queen Mary School of Economics and Finance, Swiss Finance Institute - USI, University of Illinois Urbana-Champaign, University of Lausanne, University of London - Royal Holloway College, University of Namur, University of Oxford - SAID Business School, University of Toulouse, and Washington University seminar participants; and the CEPR Workshop on Incentives, Management and Organization, CEPR/AMID Development Economics Symposium, the NBER Corporate Finance Workshop, the NBER Productivity, Innovation, and Entrepreneurship Workshop, and the NBER Organizational Economics Workshop participants for helpful comments and discussions. We wish to thank BancaMia for their support, and Gustavo Caballero, Ximena Cadenas, Isabela Echeverry, Santiago Reyes and Lucia Sanchez for excellent research assistance. We thank IPA for the funding that made this study possible. Please send correspondence to Daniel Paravisini (d.paravisini@lse.ac.uk) and Antoinette Schoar (aschoar@mit.edu).

1 Introduction

The use of statistical models to analyze the determinants of borrower repayment and default behavior is standard in consumer and small business lending.¹ The trade-offs involved in summarizing complex, multidimensional information into a single summary statistic—a credit score—are well understood in theory. On the benefits side, credit scoring has a direct effect on productivity by reducing the transaction cost of analyzing information. It can have a complementary effect through reducing the cost of incentivizing the agents who evaluate loans by making the information they possess easier to monitor by the principal. On the costs side, relevant information may be lost in score-based lending. This could occur either because the information is difficult to incorporate into statistical models (*soft* information), or because reliance on the score encourages agents to only consider the information captured in the score but to ignore any other information (multitasking problems).² Thus, from a theoretical standpoint scores may affect lending quantity and quality through different mechanisms, and the sign of the effect is ambiguous.

In the present paper we assess empirically whether and through which channel does credit scoring affect lending in environments where soft information is a crucial input in the lending decision. It is typically difficult to evaluate empirically the trade-offs in adopting credit scoring technologies because adoption occurs in response to changes in the environment and is bundled with other organizational innovations, such as changes in job descriptions or compensation structures.³ Since these also affect productivity and

¹Proprietary credit scoring models are widespread in the banking, credit card, crowd-funding, and other consumer lending sectors. There is also a worldwide credit scoring industry for consumer and business lending, with Equifax, Transunion, Dun and Bradstreet, and Experian as some its most prominent players.

²Rajan, Seru & Vig (2013) argue more generally that credit scoring is subject to the *Lucas Critique*: since models predict default on the basis of historical data, they may fail to account for agents' behavioral responses to the use of scores.

³This is the case, for example, in the study by Einav, Jenkins & Levin (2013), which compares the change in lending practices by a large auto finance company before and after the adoption of automated credit scoring. The problem of endogenous and bundled adoption is not unique to the credit industry. See Milgrom & Roberts (1990) for examples in manufacturing and other industries.

incentives directly, they confound the evaluation of the net effect of the introduction of credit scores, or of whether scores are substitutes or complements with effort provision by the workers involved in the screening process.

We address these identification challenges by conducting a randomized control trial (RCT) with a for-profit bank in Colombia where we were able to change the credit evaluation process of the bank. We also were able to obtain direct measures of the inputs to the lending decision, including committee effort. The RCT approach allows us to evaluate the effect of credit scoring *ceteris paribus*, holding constant all other determinants of lending outcomes. We also design a novel treatment arm that allows us to measure the effect of the score holding constant the information set of the committee. The design, which we explain in more detail below, measures how committee behavior changes when its members anticipate the score becoming available, but before actually observing it.⁴

We conduct the research with a for-profit bank in Colombia specialized in loans to small enterprises. Since most of the prospective borrowers do not have audited financial statements, the bank's screening process begins when loan officers collect borrower information in the field (e.g., measures of sales, costs, demand volatility). This implies that most of the information is only verifiable by the loan officer who collects it. The information is then evaluated by a credit committee that includes the loan officer who collects the information and the local branch manager. When committees cannot reach a decision, they can take actions that create substantial transaction costs for the bank: they either refer the application to the regional manager who will then make the credit decision, or send the officer to collect additional information and make the decision in a second round

⁴More information may affect agents' behavior directly (the agent reacts to becoming more informed herself) and indirectly (the agent reacts to the principal becoming more informed). Our treatment design evaluates how the agents' react to the expectation of the principal becoming more informed, and is similar in spirit to the empirical strategy used in Hertzberg, Liberti & Paravisini (2011). That paper evaluates how one of the lenders to a firm reacts to the expectation that the private information it possesses about the firm will become publicly observable by the other lenders to the same firm, but before the information actually becomes public.

of evaluation. Since committee compensation depends only on the size and performance of the loan portfolio, both actions represent a potential source of agency conflicts: committee members may free-ride on regional managers by referring too many applications or engage in many rounds of wasteful information collection to build a reputation of being dilligent.

The bank developed a credit scoring model with the objective of reducing the information processing and monitoring costs resulting from the evaluation of applications.⁵ To test the impact of the scoring model, we design the trial with a control group and two treatments. In the control, group the committees make decisions using the detailed information collected by the loan officer but without observing the applicant's score. For the bank, this is the usual screening process. In the first treatment (*T1*), we provide the applicant's score to the credit committee as it evaluates the application. We use this treatment to evaluate the overall effect of scores on committee output, through information as well as incentives. In the second treatment group (*T2*), we aim to separate the information effect of the score, from its effect through incentives. For that purpose, we ask the committee to make an evaluation of the application *before* observing the value of the score (*interim* decision), but knowing that the score would be observable immediately after the interim decision has been made. Although committees could revise the interim decision after seeing the score to make a *final* decision, we find that committees almost never do so. Comparing the interim decision in *T2* and the control group allows us to evaluate whether scores affect the committee members' incentives, even when there is no change in the information they possess.

In a four-month pilot including eight urban branches, we rolled out the scoring model and randomly assigned loan applications within a committee to one of the three treatment

⁵This model was developed independently from the present research by a financial service consulting firm specializing on credit scoring models. A statistical default model was calibrated using historical application and loan repayment information and generates a statistic for the expected default probability—a credit score—for each new loan application.

arms. Randomizing within committees, after the borrower’s information is collected, assures that the treatment is not correlated with the information content of the applications (e.g., if loan officers knew which treatment arm an application was assigned to beforehand, they could adjust their information collection effort depending on the treatment arm). This design comes at the potential cost of spillover effects, if credit committees learn over time to infer the score even in applications that do not include it. This could lead to an attenuation bias over time. For this reason, we exploit the timing of the pilot and complement the experimental results with difference-in-differences estimates that compare outcomes in the eight pilot branches with matched control branches. The identification assumptions of the difference-in-differences approach are plausible since the researchers, and not the bank, selected the timing and the identity of the pilot branches.

The main findings can be summarized as follows. First, the availability of a credit score has a positive effect on committee effort (time spent per application) and output (fraction of applications where the committee reaches a decision). The increase in effort and output is concentrated in marginal—difficult to evaluate—applications (e.g., applications for larger amounts). The output gain is economically large: the fraction of non-decisions drops by more than 40% when committees observe a score. We also find evidence suggestive of learning spillovers—committee output also increases in control applications by the fourth month of the trial— which implies that the estimates represent a lower bound on the size of the true effect on output. This first set of results implies that scores and committee effort are complementary inputs in the evaluation of marginal applications, and that scores and worker effort substitute for other expensive inputs in loan production, such as manager time.

The second set of findings is obtained from the analysis of treatment $T2$. Interim output—decisions made before observing the score— also increases relative to the control group: 75% of the total output increase in $T2$ occurs before the committee observes the

score. Interim decisions are never reverted after observing the score, and the quality of the decisions (loan amount, default rate) is not different from the decisions made under treatment $T1$. These findings imply that the expectation of a score becoming public induces committees to produce more output, holding their information set constant. Most of the change in output occurs before the committee observes the score and is thus not driven by information in the score that is otherwise unknown or difficult to process by the committee.⁶ This implies that the main channel through which scores affect committee behavior in our setting is by improving its members' implicit incentives to work.

Finally, we compare pilot and control branches to evaluate the effect on borrower selection, total output, and capital allocation. Rolling out scores does not affect the pool of borrowers along observable (average credit score, requested loan amount) or unobservable (predictive power of scores in a default model) characteristics. This suggests that the introduction of scores did not have a significant impact on loan officers' incentive to collect information. Total branch output (number of loans approved, average loan size) is also unaffected. This implies that all the additional output generated by committees substitutes for output that would have been otherwise produced by regional managers or by committees after a second round of information collection.

Comparisons between pilot and control branches also show that the introduction of scores reduces the probability of default, increases the realized loan profitability, and reduces the cross-sectional dispersion of loan profitability. These results imply that credit scores improve capital allocation efficiency. This final result is important because it implies that the use of scores does not lead to a loss of information in the credit evaluation process, for example, because committee members cater to the score and ignore soft information. On the contrary, capital allocation efficiency increases in treatment branches even though

⁶The pure information gains from observing the score are bounded in our empirical setting given that a score is calculated using the same borrower-specific information that the committees observe in the application file.

the observable measures of risk, as reflected by credit scores and their ability to predict default, do not change. Thus, the introduction of scores and their effect on committee effort result in an improvement of the quality of the information effectively used in credit evaluation decisions.

The results in this paper contribute to several strands of the literature in finance and economics. In an influential paper on the organization of financial intermediaries, Stein (2002) argues that the adoption of technologies that *harden* information in bank lending can lead to larger, more centralized, banks. The reason is that these technologies substitute for loan officer discretion and inputs, crowding out the use of soft information. Our results highlight a countervailing effect, noted in Aghion & Tirole (1997): hardening soft information may make lower-level loan officers easier to incentivize, which might allow for greater decentralization of decision-making. We are not able to pin down the exact mechanism through which scores enhance incentives. But we are able to show that since scores are not tied to compensation, the mechanism must be of a non-contractual nature. It may be related, for example, to career concerns and the way scores interact with committee members' concern of being perceived as good evaluators.⁷

More generally, our results present causal evidence on the role of information technologies in shaping the optimal organization of production. Our empirical setting is close to the theoretical setting in Garicano (2000), where workers solve routine problems independently, but ask managers for advice on difficult ones. Our findings provide the first direct corroboration of the main prediction of the model: innovations that lower the information processing costs for workers lead to more decentralized decision-making inside the

⁷In a specific example, committee members may be reluctant to make a decision based on soft information, even if it is the right decision *ex ante*, because a negative outcome *ex post* will negatively affect their reputation. In this example scores embolden decision-makers by providing a summary of the soft information that is easy to interpret and communicate. It is important to emphasize, however, that in theory more information does not lead unambiguously to better decisions. When agents want to “conform” to the norm (Prendergast 1993), or have private information about the productivity of their actions (Prat 2005), innovations that improve transparency may reduce performance.

firm. Consistent with the delegation theory in Dessein (2002), our results also indicate that these information processing costs may, at least in part, result from agency problems between workers and the principal.⁸

Existing work attempting to identify the incentive effect of information technology adoption using administrative data has had to rely on ad hoc assumptions on whether the technology has an incentive effect or not. For example, Hubbard (2000) identifies two classes of on-board computers in the trucking industry, classifying one as *incentive-enhancing* and the other as *resource-allocation-improving*.⁹ A key methodological contribution of our paper is to provide a research design that allows evaluating the incentive effect of a technological innovation that does not rely on the researcher’s judgement.

Finally, our results are related to a literature that associates technological change with broad changes in worker productivity, skill composition, and the wage structure of industries. This work has shown that the diffusion of information technologies is accompanied by increases in productivity and the skill premium across firms and industries (Katz & Murphy 1992, Autor et al. 2003, Autor & Dorn 2013).¹⁰ For the reasons stated above, it has been difficult to identify empirically the channel through which information adoption affects productivity and its heterogeneity across worker skill levels. The results in this paper show that information technology can increase the productivity of both low and high skill individuals by inducing higher worker effort and increasing managers’ span of control, respectively.

The rest of the paper proceeds as follows. In Section 2, we provide a description of the tasks and incentives of the credit committees and the characteristics of the credit scoring

⁸In theory, decentralization implies in the long-run a lower unit cost of loan production, a larger optimal span of control, and an increase in total output. However, the empirical approach limits us to studying the short-term effects of score adoption, holding loan prices and the size of the firm fixed.

⁹Also, Baker & Hubbard (2004) assume that the introduction of an on-board computer system improves performance through better monitoring. Bloom, Garicano, Sadun & Van Reenen (2011) classify technologies into communication-enhancing and information-enhancing.

¹⁰For early surveys on information technology adoption, see Brynjolfsson & Yang (1996) and Brynjolfsson & Hitt (2000).

system. Section 3 describes the experimental design and provides descriptive statistics on the loan applications. Section 4 presents the experimental results of introducing the score on committee output and productivity, and Section 5 presents the difference-in-differences results comparing pilot to control branches. Section 6 concludes.

2 Setting

The study was implemented with BancaMia, a for-profit bank in Colombia that focuses on issuing unsecured loans to micro and small enterprises. The business model of BancaMia, similar to those of other for-profit micro and small business lenders, is based on issuing large numbers of small loans. During October 2010, the month prior to the roll-out of the pilot, the bank issued 20,119 loans totalling \$US25.9 million through its 143 branches (average size below \$US1,300). The bank makes lending decisions based on information about prospective borrowers collected first-hand by loan officers in the field. This information collection mechanism is costly but necessary, since micro and small enterprises in Colombia do not have audited financial statements or other secondary data that a bank could use for credit assessment.¹¹

The loan screening process has three unique features. First, in the absence of audited financial statements, most field-collected information is *soft* according to the definition in Stein (2002): it is difficult to verify by anyone other than the individual that collects it. For example, the loan officer may gather the sales figures for a small restaurant by counting the number of tables served during lunchtime and multiplying by the average price of a meal. Thus, it is impossible to ascertain by third parties the accuracy or veracity of the sales figure in the application. Second, the loan officer makes an active

¹¹BancaMia also offers very small loans to borrowers with high ex ante default probabilities to elicit information about their true ex post propensity to repay in a manner consistent with Rajan (1992) and Petersen & Rajan (1995). However, a very low fraction of the study's sample applications are first-time borrowers, so we ignore this learning aspect of lending in the analysis.

decision to fill out an application and bring it to the committee for consideration. Thus, applications that reach the committee do not represent the universe of potential borrowers, but only those that pass the initial screening by the officer in the field.¹² Third, officers advise prospective borrowers on how to fill out the application and may encourage riskier borrowers to request small loan amounts to improve the probability of approval. Thus, application loan amounts do not represent the borrower’s unconditional demand for credit. A consequence of prescreening and ex ante risk adjustments in requested loan amounts is that the rejection rate of loans that reach the committee is very low.

Due to the selection occurring at the loan prospecting stage, we must consider the possibility that the introduction of scores changes the officers’ incentives to gather information and select applicants in the field. Our experimental design, discussed in Section 3, randomizes at the application level once the application reaches the committee to ensure that changes in the application pool do not affect the internal validity of the results. In the results section, we also explore how the observable characteristics of the applicant pool change during the pilot week and find little evidence of changes in the prospecting effort.

2.1 Credit Committees

Loan officers upload borrower information and the bank’s management information system compiles in a single application file all the available borrower information for review by the committee: the first-hand information collected by the loan officer, the borrower’s past credit history in BancaMia, and any available third-party information (e.g., the borrower’s credit score from a private credit bureau).

The application file is reviewed by a credit committee composed of the loan officer who collects the information, the branch manager (the loan officer’s immediate superior),

¹²All the information regarding potential applicants that do not reach the committee review stage is discarded by BancaMia and is not available for this study.

and one or two additional credit specialists. The committee's main task is to reach a decision on an application (accept the application and set the terms, or reject it). The committee may make modifications to the amount and maturity of the requested loan to improve its expected profitability. For example, the committee may only approve \$US500 of a \$US1,000 requested amount if the borrower is deemed to be too unlikely to repay the latter amount. Alternatively, the committee may approve an amount larger than the application amount when it deems that the loan officer was too conservative in his recommendation to the borrower. The committee has no discretion to set the interest rate of loans, which is determined by headquarters based on the type of loan (first-time versus repeat borrower, urban or rural loan).

When the members of the committee cannot reach a decision, because they cannot agree on the terms of the loan or on whether to reject or accept the application, they have two alternative courses of action. The first is to refer the application to a regional manager, who evaluates the application and reaches a decision.¹³ Upper-level managers are more skilled and experienced than committee members and are expected to be more likely to reach the *correct* decision on difficult applications. The second option is to send the officer to collect additional information about the borrower. In this case, the committee must review the application again after the information is collected and reach a decision (or send the application to the regional manager).

BancaMia managers expressed during informal interviews that referrals and additional rounds of information collection represent a substantial cost to the bank in terms of managers' and officers' opportunity cost of time. It is difficult to quantify these costs precisely. The base fixed wage of a regional manager is four to eight times that of a loan officer, which gives a lower bound on the incremental cost of a referral to upper management.

¹³Loans above 8 million pesos go directly to the regional manager for approval. Randomization insures that this mechanical relationship between loan size and approval level is orthogonal to the scores. Also, adding the application amount as a control does not change the estimated effect of scores.

Further, since the regional manager must evaluate the application without the officer that collected the information present, she must incur additional communication costs to access any information not reflected in the application. Referrals imply additional delay costs. This is in part due to the large volume of referred applications and the time constraints of regional managers who supervise between 15 and 80 offices.

Committee member bonus compensation is an increasing function of the number, amount, and value-weighted performance of the loans issued by a branch. Performance pay to loan officers based on lending amount and loan performance is common in most types of lending institutions.¹⁴ Bonuses are calculated on the basis of loans issued, regardless of whether the decision was made by the committee or referred to the manager.¹⁵

2.2 Credit Scores

In 2010, BancaMia contracted with an external consulting firm the development of a credit risk model that uses the historical information in the credit applications to predict the repayment performance of prospective borrowers. The scoring model uses both quantitative (gender, age, location, number of years in business, frequency of late payments in past three years) and qualitative (overall knowledge of business, general sense of the level of organization, quality of information provided, quality of business location, stability and diversity of income) information in the applications. The risk model would be used to produce a credit risk score to be included in the application file. Management's stated objective is to "improve identification of the best and worst clients, decentralize the loan approval process, and reduce the labor costs involved in loan application evaluation."¹⁶ It

¹⁴The combination of bonuses based on the number of loans and value-weighted loan performance are meant to provide incentives to issue small loans to the riskiest borrowers—compensation based solely on the dollar volume of lending would discourage officers from making small loans.

¹⁵There are two potential reasons for compensating committees for loans issued. First, pay based on decisions made penalizes committees for referrals, which may lead to too many bad decisions at the committee level. Second, pay based on decisions made would eliminate the officer's incentives to monitor the borrower after the loan has been issued.

¹⁶Source: interviews with management.

is important to note that the credit score relies entirely on information that loan officers would also have had anyway, since the score relies on the information collected by loan officers.

The score is a proxy for the expected default probability of the loan. Figure 1, panel (a), plots the out-of-sample relationship between the default probability and score in the population of loans issued during October 2010 (the scoring model is calibrated using data for loans issued in 2009). A loan is considered to be in default if interest or principal payments are more than 30 days overdue at six or twelve months after the loan is issued. There is a strong positive association between credit scores and default probabilities. The tight standard error band implies that scores have a good out-of-sample predictive power for future default, and that the data in the application is informative about borrowers' repayment prospects.

There is a negative relationship between requested loan amounts and default probability in the loan population (Figure 1, panel (b)). This relationship is consistent with loan officers screening large applications by risky borrowers, or recommending risky borrowers to request smaller loan amounts. This relationship suggests that loan officers form an estimate of the borrower's default probability before bringing the application to the committee.

3 Research Design and Descriptive Statistics

We design a randomized controlled trial (RCT) with two goals. The first is to measure the causal effect of scores on committee effort and output, as well as overall loan output and performance. The second goal is to decompose the causal effect of scores into two broad mechanisms: the provision of information and, holding information constant, the provision of incentives.

In the pure information provision mechanism, scores provide a signal of borrower

quality that reduces the cost of committee deliberation. For example, the population-based score may provide a more precise signal of the creditworthiness of the borrower than the committee can gather based on their own assessment. Alternatively, the score may provide the same information that the committee can gather from the application, but at a lower cost. Also, the score may reduce the cost of communicating the information collected by the loan officer to the other committee members by providing a simple to interpret summary statistic.

The score may also affect the committee members incentives to divulge and use the information they already possess by providing a publicly observable and verifiable measure of the creditworthiness of the borrower. For example, scores may deter the committee from overstating the creditworthiness of family, friends, or acquaintances of its members. A score may also embolden the committee to make decisions in marginal cases because its members can use the score to justify the decision if the loan defaults *ex post*. On the other hand, precisely because the score may justify decisions *ex post*, it may discourage committee members from using valuable private information in the deliberation (catering to the score). Thus, not only is it an empirical question whether there is an incentive effect at all, but if there is, the direction of the effect is ambiguous.

The fundamental feature of the incentive mechanism is that it affects outcomes holding the information set of the committee constant. Thus, to isolate the incentive effect, we design a treatment that allows us to measure how the committee's actions change when its members expect the score to become available, but before observing it.

3.1 Pilot Design

We implement a four-month pilot program with an RCT design in eight urban branches. In each pilot branch, we randomly select the treatment applications for which the committee will be able to see the applicant's score. Randomization occurs in real time when

the committee begins to discuss an application but after the information of the client has been collected. This process guarantees that treatment is orthogonal to applicant characteristics and the information gathering effort by the loan officer. The committee members are informed of the group assignment at the beginning of the discussion.¹⁷

We randomize in real time at the application level to guarantee that treatment is orthogonal to applicant characteristics and the information gathering effort by the loan officer. The main disadvantage of our design is the potential for learning spillovers from treatment to control applications. In the extreme case where committee members eventually would learn perfectly the algorithm that maps borrower characteristics into scores, then there will be no difference between treatment and control application outcomes. Thus, our estimates can be interpreted as lower bounds on the treatment effect of scores.

In the control group, the committee evaluates the application without observing the score. In the first treatment group ($T1$), the committee receives the score before evaluating the application. This first treatment allows us to measure the overall effect of scores on committee effort, output, and productivity. In the second treatment group ($T2$), the committee makes an *interim*, non-binding, decision before receiving the score. The committee makes the interim decision knowing that the score will become available to the supervisor. But the committee members have no additional information relative to control applications, and they know that the interim decision can be amended after observing the score. The interim decision in treatment $T2$ allows measuring how committee behavior changes through incentives alone, holding the information set constant. The difference between the interim and the final decision allows measuring the incremental effect of the information content of the score.

Consider the two extreme cases in which the effect of the score is purely through the

¹⁷To ensure that there was no *migration* of applications between control and treatment groups, a research assistant sitting in the committee performed the randomization in real time using a random number generator on the computer and informed the committee of the outcome.

information channel or purely through the incentive channel. If the score affects behavior through its information content alone, then it will be optimal for the committee to wait until the score is revealed to make a full evaluation of the application. In this case, interim decisions will be uninformed and often revised. On the other hand, if the score changes the incentives of the committee members to use the information they already possess, then interim decisions will be informed and not frequently changed. In this case, committee output will increase before the score is revealed.

The main advantage of the design of $T2$ is that it allows measuring the incentive effect regardless of the nature of the underlying agency problem that the score mitigates (or creates) between committee members. The interim action is non-binding and unobservable to anyone outside the committee, and thus the interim action in $T2$ is more likely to affect conflicts of interest amongst members of the committee than those between the committee and the rest of the organization.

In a short training workshop before the roll-out of the scores, we provide branch managers and loan officers at the eight pilot bank branches with a general explanation of the credit risk model, the scores, the objective of the study (researching the usefulness of the score as an input to the credit evaluation process), and a detailed description of the three treatment groups and the randomization procedure. We report in Appendix Table A.1 the number of applications in each group per branch in the study sample.

3.2 Descriptive Statistics

We present descriptive statistics of the applications in the control and treatment groups in Table 1, as well as the p-values of difference of means tests between the three groups. Predetermined application characteristics—characteristics determined before the randomization takes place—are shown in Panel A. In the control group, average loan amount is US\$1,551, average maturity is 20.9 months, average risk score is 0.151, and the fraction

of first-time borrowers is 14.6%. Randomization implies that any differences in predetermined variables between the treatment and control groups are purely by chance. Table 1 corroborates that randomization created groups that are comparable in terms of predetermined characteristics, with the only significant difference in means occurring for application maturity, which is one month shorter in treatment group *T1* than in the control group. Our main specification will include pre-determined characteristics to account for chance differences between groups.¹⁸

Table 1, Panels B through E, presents the statistics for committee and loan outcomes. Some outcomes, such as the time the committee needs to take an action, are measured for all applications. Others are measured conditional on a particular action of the committee. For example, the approved loan amount is measured conditional on the committee approving the loan.

The average time spent evaluating an application in the control group, measured as the difference in the time stamp assigned by the research assistant to the beginning and end of each evaluation, is 4.68 minutes (Std. Dev. 3.28). Committees reach a decision (accept or reject a loan) in 89% of the control applications, and conditional on reaching a decision, in 0.3% of decisions the committee rejects a loan in the control group.

Conditional on loan approval, the committee approves a loan amount different than the requested one in 92% of the applications. The average ratio of approved to requested loan amount is 0.975, indicating that the mean size of approved loans differs little from the mean application amount. Nevertheless, there is substantial variance (Std. Dev. 0.419), and the average absolute value of the difference between the approved and requested amount is \$US266, or 17% of the average requested loan amount. Similar patterns can be

¹⁸The application amount and score distributions are indistinguishable between the treatment and control groups in a two-sample Kolmogorov-Smirnov test for equality of distributions, with corrected p-values of 0.81 and 0.94, respectively. We do not perform a test for maturity because the maturity distribution is highly non-normal, with many observations concentrated in whole year numbers (12, 24, and 36 months).

found for loan maturity, although the proportion of cases in which the committee modifies the loan application maturity is lower (26.2% of the applications in the control group).

The low rejection rate and frequent rate of loan size and loan maturity adjustments suggest that committees' decisions occur mostly in the intensive margin, e.g., on how much to lend as opposed to whether or not to lend. In addition, the fact that the average requested and approved loan amounts are close indicates that loan officers' evaluation of the borrowing capacity of the applicant is in line, on average, with that of the committee. In other words, loan officers do not seem to be systematically aggressive or conservative relative to the committee in their assessments of borrowers' risk.

Not all approved loans are issued: only 83.5% of the loans approved during the pilot program appear as issued in the bank's information system. The bank does not record the reason why the borrower decides not to take up the approved loan. The default rate among the issued loans—fraction of loans more than 30 days late in repayment measured six (twelve) months after the loan was issued—is 3.3% (9.5%) in the control group.

Comparing the unconditional outcomes in the treatment and control groups in Table 1, shows that, on average, committees spend more time reviewing applications in the treatment groups, although the difference is only significant for treatment $T2$ (the difference in average time between $T1$ and $T2$ is not significant). Committees were more likely to reach a decision in both treatment groups than in the control. None of the loan characteristics or outcomes conditional on approval are statistically different in the treatment and control groups, except for the average absolute value of the change in maturity.

Table 2 shows the descriptive statistics for applications in the control group conditional on the action taken by the committee—made decision, referred application to the regional manager, or sent the officer to collect additional information. On average, the applications for which the committee reaches a decision are for smaller amounts and are more likely to be submitted by first-time applicants than applications where the committee does not

reach a decision. Applications where the committee reaches a decision are no different in their credit risk from those referred to the manager (as measured by the score), but they have a smaller credit risk than those that required additional information. Committees spend less time evaluating applications where they reach decisions than when they refer an application or collect additional information. If one equates evaluation time with effort, this implies that the committee members employ a substantial amount of effort before being able to ascertain that a decision cannot be reached.

We can track final outcomes for applications when the committee did not make a decision using BancaMia's information system. This allows us to measure the disbursed amount and the default rate of loans approved by the manager, or loans approved after a second round of information collection.¹⁹ Loan outcomes differ substantially depending on the action taken by the committee. For example, loans approved by the manager default with a probability of 8.3% after 12 months and have an approved amount that is 95% of the application amount, while loans approved by the committee after a second round of information collection default with a probability of 13.3% and have an approved amount that is 148% of the application amount.

The complexity of the task of reaching a decision is unobservable by the econometrician. In theory, in the presence of task heterogeneity, committees should make decisions on the easy-to-evaluate applications and refer to the manager or collect additional information on the difficult-to-evaluate ones (Garicano 2000). The statistics in Table 2 suggest that difficult applications take more time to evaluate and that time spent evaluating an application can be used as a measure of effort. Also, the statistics suggest that applications for larger amounts, and with longer maturities, are more difficult to evaluate, while the risk of an applicant (as measured by the score) is uncorrelated with difficulty.

¹⁹Note, however, that because not all approved loans are issued, we cannot measure the fraction of the applications rejected by the manager or in a second round by the committee (rejected applications and approved but non-issued applications are confounded in the ex post data).

4 Experimental Results

We use the following reduced form equation to estimate the effect of credit scores on committee and loan outcomes:

$$Y_i = \beta \cdot Treatment_i + X_i' \cdot \eta + \varepsilon_i, \quad (1)$$

where Y_i is an outcome related to loan application i . The variable $Treatment_i$ is an indicator for whether the application is in either treatment group $T1$ or $T2$. In some specifications, we also include an indicator equal to one if the application is in treatment $T2$ to evaluate differential effects of the two treatments (we do not find any). The vector X_i contains predetermined application characteristics: applicant's credit score, application loan amount, application loan maturity, a dummy if it is the first loan application of the potential borrower, and the date of the application (in weeks).²⁰

We begin by presenting the results for outcomes that are measured unconditionally, such as the action taken by the committee or the application evaluation time. The estimated β using these outcomes measures the Average Treatment Effect (ATE) of having a score as an input to the credit evaluation process.

4.1 Committee Output

We present in Table 3 the estimated effect of introducing a score on committee output, measured as the fraction of applications in which the committee makes a decision (accepts or rejects an application). This is a good measure of committee output because the bank can significantly save on transaction costs, without reducing the quality of the loan portfolio, if the committee reduces the number of applications it sends to upper managers or those it collects information on (we discuss and measure the effect on quality later). The

²⁰Results without controls are not significantly different, see Appendix Table A.2.

point estimate is 4.6 percentage points, statistically significant at the 5% level (column 1). This implies that when scores are added as an input in the decision process, the number of cases in which committees cannot decide is reduced by 41.8% relative to the baseline proportion of 11% in the control group. The difference in the effect between $T1$ and $T2$ is positive but not significant (Table 3, column 2).

The data allows identifying two distinct margins through which scores increase committee output: 1) by reducing the need to collect additional information from applicants, and 2) by reducing the need to use upper-level manager time in evaluating loan applications. We present in Table 4 the results of estimating a multinomial logistic specification to model committee choice between approving a loan, rejecting it, collecting additional information, or sending the application to a manager in a higher hierarchical level to make the decision.²¹

Treatment has a negative and statistically significant effect on the probability of sending the application to the manager and on the probability of collecting additional information.²² To evaluate the economic significance of the effects, we report on the bottom rows of Table 4 the implied marginal effect of treatment on the probability of each choice. Observing a score decreases the probability of referring the application to the manager by 2.3 percentage points, a 48% decline reduction in the baseline probability that an application is sent to the manager in the control group. Scores reduce the probability of collecting additional information by 1.7 percentage points, a 27% decline relative to the

²¹We estimate:

$$\ln \frac{P(D_i = m)}{P(D_i = 1)} = \beta_m \cdot Score_i + X_i' \cdot \chi_m + \varepsilon_{mi}, \quad (2)$$

where D_i represents the committee choice. We use the committee's decision to approve a loan, $D_i = 1$, as the reference category. All right-hand side variables are as in equation (1). There is one predicted log odds equation for each choice relative to the reference one, e.g., there is a β_m for rejecting a loan, one for collecting more information, and one for sending the application to the manager. A positive estimate for β_m implies that committees are more likely to take action m than to approve a loan in the treatment group relative to the control group.

²²The coefficients on the treatment regressors β_m are significant at the 1% level in a joint test across the four choices.

baseline.

The results suggest that scores reduce the cost of decision-making by committees. Consistent with the prediction in Garicano (2000), lower decision-making costs by committees lead to fewer referrals to managers in upper levels of the hierarchy. Scores also substitute for additional costly information collection by loan officers.

4.2 Spillovers

Scores are a mapping between information in the application and default probabilities. Committee members may learn how to infer this mapping during the course of the pilot since all the borrower-specific information used to construct the score is available to the committee. After observing a number of treatment applications, committee members may develop a heuristic that allows them to infer the score even when one is not provided. If this occurs, control applications will be affected by the pilot, committee output in control applications will tend to converge to output in treatment applications, and the comparison of treatment and control applications will lead to a downward biased estimate of the effect of treatment.

To evaluate the degree of such learning by committees, we plot in Figure 3 the measure of committee output (fraction of applications in which the committee reached a decision) in treatment and control applications, by pilot month. The plots shows that committee output in treatment applications was significantly higher than in control applications during the first month of the pilot. The difference in output between treatment and control applications shrinks as the pilot progresses and becomes indistinguishable from zero by the third month of the pilot.

Thus, the dynamics of committee output in treatment and control applications strongly suggests that committee members learn within two months how to infer scores from application files. During the first month of the pilot, the raw difference between treatment

and control application output implies that treatment increased the probability of making a decision by 8 percentage points, almost double the estimate obtained in Table 3.

This discussion implies that our estimates are a very conservative measure of the true effect of scores on committee output. We expect the same downwards bias to be present in the measurement of the effect of scores on other outcomes, which implies that we need to be cautious when interpreting results that suggest there is no effect. In Section 5 we present difference-in-differences estimates that compare across branches that are less likely to be contaminated by spillovers.²³

4.3 Committee Effort

The estimated effect of introducing a score on the time it takes to evaluate and application is 0.76 minutes, statistically significant at the 1% level (Table 5, column 1). This implies that committees spend, on average, 16.2% more time per application when scores are available, measured at the mean evaluation time in the control group of 4.7 minutes. Treatment $T2$ has a larger effect on evaluation time than $T1$, but the difference is not statistically significant (column 2). Some difference is expected since committees must make two decisions in $T2$: an interim one without observing the score and then a final decision after observing the score. This result implies that the additional time it takes to make an additional decision in $T2$ is not statistically different from zero. Although we cannot use these results to distinguish how evaluation time is split between interim and final decisions and $T2$, committees rarely amend interim decisions (we discuss this in more detail below) which suggests that proportionally more time must be spent in interim decisions.

Assuming that time evaluating applications is a measure of committee effort, these

²³The potential spillover bias in the difference-in-differences setting is inter-temporal: if the committees learn to infer the score in pilot branches, then the pilot branches will continue to be "treated" even after the pilot is over.

results suggest that scores and committee effort are complements. However, this assumption requires some validation because the additional time spent evaluating the applications may be waste and not a necessary input for production. More time spent per application could also be a symptom of the committee becoming less productive, for example, if scores are a source of distraction. To shed some light on how to interpret the results, we turn to a specific prediction in Garicano (2000): lower decision costs will lead committees to tackle *more difficult* problems. If time spent is correlated with effort, we should observe that scores increase the time spent on applications that are more difficult based on ex ante measures.

We observe from the subsample of control group applications in Table 2 that committees spend more time in applications where they cannot reach a decision (difficult applications) relative to those in which they do (easy applications). If the additional time spent due to the scores shown in Table 3 is related to effort, we expect the effect to be concentrated on applications that take longer to evaluate to begin with. Put differently, scores should shift right the upper tail of the evaluation time distribution, while they should have little or no effect on the lower tail.

To evaluate this, we characterize the effect of scores on the distribution of decision time by estimating specification (1) using simultaneous quantile regressions for the 25th, 50th, and 75th quantiles of evaluation time. The results, shown in Table 3, columns 3, 5, and 7, indicate that only percentiles at or above the median are affected by the introduction of scores, and the point estimates increase monotonically with the quantile. This heterogeneous effect is consistent with the effort interpretation of the results in Table 3. In contrast, the additional effect of $T2$ relative to $T1$ on evaluation time appears to be constant across all quartiles (columns 4, 6, and 8), suggesting that any difference in evaluation time between $T2$ and $T1$ is not related to problem difficulty.

The descriptive statistics for control group applications in Table 2 also indicate ap-

plications for larger amounts are more difficult to evaluate (committees are less likely to reach a decision for larger applications). We confirm this non-parametrically in Figure 2. In control group applications, the probability of making a decision declines (panel (a)) and the time spent per application increases (panel (b)) with the application amount. Again, if time spent in applications is related to the effort put into evaluating them, we would expect the effect of scores on evaluation time to be larger specifically for those applications where scores lead to more decisions. Turning to treatment applications in Figure 2, scores indeed induce a larger increase in the probability of making a decision and on the time spent evaluating applications for larger application amounts. The coincidence of the effect of scores on output and decision time on difficult-to-evaluate (large amount) applications is consistent with the effort interpretation and with the predictions in Garicano (2000).²⁴

We have remained purposely agnostic regarding what “difficult” means in this setting. It may relate to the cognitive costs of reaching a decision, the cost of exposing the pay to a larger risk (bonuses are negatively related to the value of defaulted loans), or the reputational costs involved in reaching consensus in applications where there is more at stake. The cross-sectional patterns documented in this subsection imply that the roll-out of scores reduces the cost of reaching decisions on such difficult applications, as measured by the counterfactual time they would have required to solve in the absence of treatment, or by the requested amount.

4.4 Conditional Outcomes

The only loan outcome that we can measure unconditionally is whether or not the loan was issued. All other outcomes—e.g., amount issued, default probability—are measured

²⁴We also perform a parametric exploration of treatment heterogeneity by augmenting specification (1) with interactions between the treatment dummy and application size. These interaction terms are not statistically significant at the standard levels. This is to be expected given the non-linear patterns observed in non-parametric plots.

only conditionally on the loan being issued. For outcomes that are measured conditionally, the interpretation of β in specification (1) is complicated by the fact that it captures a combination of two effects: 1) a direct causal effect of treatment on the outcome, and 2) a selection effect driven by the effect of treatment on the conditioning variable. Given that treatment affects the probability of making a decision, the selection component may be important when measuring committee outcomes.

To illustrate the complication introduced by measuring the effect of treatment on conditional outcomes, consider the case of loan size approved by the committee. This loan size can be measured conditional on the committee approving an application, as opposed to rejecting it, sending it to the manager, or postponing the decision until more information is collected. Scores may have a direct effect on approved loan size, holding constant the set of applications. This is the Average Treatment Effect (ATE) of scores on loan amount. Scores also change the set of applications that the committee decides on. These marginal applications are likely to be different along dimensions that are correlated to loan size. Thus, treatment changes the composition of applications approved by the committee in a way that may affect average loan size even if the ATE is zero.

Disentangling the treatment and selection effects is typically difficult without an additional instrumental variable for the selection effect (see Lee 2008, for a recent discussion). Our setting provides a unique advantage to deal with this identification problem: we can evaluate outcomes of marginal committee decisions due to selection because we observe the outcome of all decisions, regardless of whether they were made by the committee during the pilot, by the committee after additional information is collected, or by the regional manager. To follow our example above, suppose that by conditioning on the sample of loans approved by the committee, we found that loan amounts are 10% larger when the score is available. And suppose that without conditioning, we found loan amounts do not change. In this example, the change in loan size conditional on committee approval

is entirely driven by selection: scores induce committees to decide on larger applications that would have otherwise been referred to the manager or decided after collecting more information. If, on the contrary, issued loan amounts increase, we can infer that the ATE of scores on approved loan amounts is positive.

We present in Table 6 the estimates of β on outcomes conditioning on the subsample of applications where the committee made the decision during the pilot (panel A), and without conditioning (panel B). Conditional on making a decision, the probability that a committee rejects an application increases by 0.9 percentage points in the presence of scores, significant at the 10% level (column 1). This estimate implies a three-fold increase in the proportion of applications rejected by the committee relative to the baseline probability of 0.3% in the control group. Moreover, assuming that all the additional rejections come from the marginal decisions, the estimate implies that committees reject 13% of the marginal cases they decide on when scores are used as an input $((0.9 - 0.3)/4.6)$. None of the other outcomes conditional on the committee approving the loan are significantly affected by scores. So even though committees are deciding on a larger proportion of marginal cases when the score is available, the average credit supply, loan maturity, or probability of default of the loans approved by the committee do not change.

Next we turn to evaluating the effect of treatment without conditioning on whether or not the committee makes a decision. We can track the ex post outcomes of all applications that led to an issued loan through the bank's information system. This includes, for example, issued loans where the application was approved by the regional manager. The effect of scores on the probability that the loan is issued is close to zero and not statistically significant (Table 6, panel 2, column 4). This implies that the addition of scores to the loan production process does not affect the overall extensive margin of lending. Thus, committees are rejecting more applications without affecting the total number of loans issued. This implies that the full increase in loan rejections by committees is due to

selection: committees are rejecting applications that would have been rejected by the manager or by the committee in a second decision in the absence of the score.

We find that treatment does not have a statistically significant effect on the average level of other outputs (loan size and default probability, columns 5 and 6). Given the potential downward bias of the estimates due to learning spillovers, we investigate again in Section 5 whether the introduction of scores affects outcomes measured at the branch level using a difference-in-differences analysis.

4.5 Interim Decisions before Observing Scores

In this subsection, we turn our attention towards evaluating the effect of treatment $T2$ on interim decisions. In treatment $T2$, the committee performs an evaluation of the application and reaches an interim conclusion before observing the score. That is, they choose an interim action with the same information set as in the control applications, except for the knowledge that the score would become observable by all committee members immediately after the action is chosen.

We estimate the OLS equation (1) with interim committee decisions as the left-hand side variable, using for estimation only the control and $T2$ applications. The right-hand side variable of interest is a dummy equal to one if application i belongs to treatment $T2$. The coefficient on this dummy measures the effect of making the score available on committee actions *before* the committee observes the score, and thus reflects the gross effect before receiving a new signal about borrower creditworthiness.

We present the results in Table 7 (odd columns). We present in the same table the estimated effect on final outcomes for $T2$, after the committee has observed the score (even columns). The effect of the score on the probability of making an interim decision is positive and significantly different from zero at the 5% confidence level (column 1). The magnitude of the estimated coefficient is 0.039: the probability that the committee makes

a decision increases by 3.9 percentage points before observing the score in $T2$. The magnitude is smaller than that of the effect on the probability of making final decision, 0.052 (column 2), but not statistically distinguishable. Committees thus make more decisions in anticipation of receiving the score, and then make even more decisions after observing it. The point estimates suggest that 75% ($.039/0.052$) of the increase in output occurs before observing the score. In addition, the expectation of receiving a score significantly reduces the probability that committees refer an application to the manager in the interim decisions (see Appendix Table A.4).

Conditional on making a decision, committees are also more likely to reject applications during the interim action, and before observing the score. In this case, the increase in the probability of rejection in the interim action, 1.3 percentage points (column 3), is larger than the increase in the final outcomes, 1.1 percentage points (column 4), although again, the estimates are not statistically distinguishable. Appendix Table A.3 presents in matrix form the transitions between interim and final decisions for all the applications in treatment $T2$ and shows that committees never revise an interim decision to reject an application. This implies that the decline in the point estimate on approval probability between interim and final action occurs due to an increase in number of decisions made and approved (the denominator of the fraction).

Finally, the estimated effects on approved loan amounts, on the probability of adjusting the application loan size, and on the absolute value of the adjustment, are all statistically insignificant for both interim decisions and final ones.

The main conclusion from this set of results is that the bulk of the effect of scores on committee output occurs even while holding committees' information set constant. This is consistent with an incentive effect of scores on output: scores induce committees to make more decisions. However, recall from the motivation in the introduction that the incentive effect of scores need not increase the overall amount of information used

in the decision process. Scores could lead committee members to forego useful private information that is not contained in the score. When we estimate separately the effect of scores on the default probability from treatments $T1$ and $T2$, we find the same result presented in Table 6: the point estimate is negative but not statistically different from zero (not shown). Interpreting this as evidence that scores do not deteriorate the quality of lending, the results imply that most of the relevant information contained in the scores is already known by the committee members, and scores provide incentives to use this information more effectively in the decision process. We revisit the issue of the effect on loan quality in the next section when we present the difference-in-differences results.²⁵

5 Difference-in-Differences Results

In this section, we evaluate the effect of introducing scores by comparing the pilot branches with a matched sample of control branches. There are two rationales for this analysis. The first is that it allows us to test whether the introduction of scores affected the inputs that are fed into the loan production function before the application arrives in the committee. Namely, screening effort, applicant selection, and information collection performed by the loan officer in the field. The experimental design was not designed to capture these effects, since the randomization occurs at the committee level after the selection and information collection have take place.

The second rationale is that the spillovers documented in the previous section imply that the comparison of treatment and control applications may hide important effects of the introduction of scores on lending output. In fact, the patterns observed in Figure 3 suggest that all applications in pilot branches were de facto treated during the second

²⁵In unreported regressions we find that the effect of scores on output is heterogeneous on the score of the applicant, although the results are significant only in the subsample of applications with a requested amount above the median. This is consistent with an interpretation in which scores induce effort because they make observable to the principal some of the private information in the committee.

half of the sample. A comparison between pilot and non-pilot branches is more likely to provide an unbiased estimate of the effect of scores on lending output.

This analysis compares outcomes of the pilot branches during the weeks of the trial relative to other weeks, and relative to propensity score-matched, non-pilot branches of the bank during the same weeks. This is equivalent to estimating a difference-in-differences estimator, since the timing of the introduction of the pilot and the branches in which the experiment was rolled out was determined by the researchers and not by an endogenous choice of the bank. We estimate two types of difference-in-differences specifications. The first one is at the application level:

$$Y_i = \gamma \cdot PilotWeek_i^{Branch} + \alpha_i^{Branch} + \delta_i^{Week} + \psi_i^{Branch}t + \varepsilon_i, \quad (3)$$

where Y_i is either an application characteristic (e.g., borrower score, requested amount) or a loan outcome (e.g., issued loan amount, in default after 12 months) for loan i . The specification includes a full set of branch fixed-effects (α_i^{Branch}), week fixed-effects (δ_i^{Week}), and branch specific trends ($\gamma_i^{Branch}t$). The right-hand side variable of interest is $PilotWeek_{it}^{Branch}$, a dummy equal to one if the branch where loan i was issued was a pilot branch of the randomized control trial during week t .

The second specification includes a similar set of right-hand side variables but is estimated after aggregating the data at the branch-week level:

$$Y_{jt} = \gamma' \cdot PilotWeek_{jt} + \alpha'_j + \delta'_t + \psi'_j t + \varepsilon'_{jt}, \quad (4)$$

where Y_{jt} is a branch level output (e.g., the number of loans issued, the average and standard deviation of realized loan returns) of branch j during week t .

These specifications are estimated using all the loans approved between week 41 of 2010 (four weeks before the pilot starting date) and week 26 of 2011 (four weeks after

the pilot end date). Every pilot branch is matched with a control branch selected from the population of bank branches using the nearest distance measured with a propensity score. The propensity score is estimated using the branches' number and total amount of loans approved, average approved loan size, and the average borrower score in October 2010, the month prior to the beginning of the pilot.

5.1 Loan Screening

We present the results of estimating the loan level specification (3) in Table 8. The roll-out of scores in a branch does not have a statistically significant effect on either the score of the applicants that receive loans (column 1) or on the amount of the loan requested in the application by loan recipients (column 2). The first measure captures the quality of loan applicants as measured by the risk model; the second their quality as assessed by the loan officer in the field (recall that loan officers seem to adjust downwards the application amount for riskier borrowers). Thus, these results suggest that scores do not change applicant selection by loan officers, or if it does, the change in applicant selection is fully undone by the committee.

Even if applicant selection does not change, it is possible that scores change the information collection effort by the officer. If this is the case, then the information in the applications, which is the information used by the default model to predict the score, could change. To evaluate whether this is the case, we test whether scores' predictive power of default changes during pilot weeks. In Table 8, column 3, we show the results of estimating specification (3) using a default dummy as the dependent variable. The right hand-side is augmented with the score, and the score interacted with the experiment week dummy. The coefficient on the interaction term, which represents the change in the predictive power of scores during the pilot, is not statistically significant at the conventional levels. This suggests that loan officers did not alter in a significant manner

their information collection effort in response to the introduction of scores.

5.2 Branch Output and Capital Allocation Efficiency

To evaluate the effect of score introduction on branch output we use the branch-week specification (4) and present the results in Table 8. Rolling out the scores in a branch neither affects the number of loans approved in a branch (column 4) nor the total amount of lending (column 5). These results corroborate the conclusions from the previous section: the entire increase in committee output caused by the scores substitutes for output previously generated by managers with a higher opportunity cost of time, or by the same committee after engaging in costly information collection. In other words, scores and committee effort substitute for more expensive inputs to loan production.

We also investigate whether the introduction of scores affects the default probability of issued loans. Loans issued during the pilot week have a default probability that is 1.2 percentage points lower than loans in control branches, although the measured effect is significant only at the 10% level (column 6). The point estimate implies a 12.6% decline relative to the baseline default probability in the control group of 9.5%. Note that the decline in the default probability occurs even though the average score of applicants does not change, which means that the default probability drops above and beyond what would be predicted by a default model based on ex ante borrower characteristics.

This result is important because it highlights that the source of the improvement in loan performance is not due to a better selection of applicants based on their observable ex ante characteristics. Instead, the improvement in performance may be due to additional information produced by the committee during the deliberation, or because the committee uses existing information to better tailor the loan contract characteristics to the circumstances of the borrower. Either of these mechanisms would imply that scores lead to a better allocation of capital across borrowers. A better capital allocation implies,

in turn, an increase in the expected profitability of lending, and an equalization of the expected marginal return across loans. We investigate this by evaluating the effect of rolling out scores on the average and the cross-sectional standard deviation of the realized gross return of loans issued in branch j during week t . The realized gross return of a loan is calculated as the ratio of the collections divided by the amount lent. This ratio goes up when borrowers repay a larger amount of principal and interest relative to the amount borrowed. Consistent with the decline in default rates, the roll-out of scores also increased the average realized return of lending by 1.5 percentage points (column 7). Consistent with an improvement in capital allocation across borrowers, we find that the variation of realized returns across loans issued by a branch drops during pilot weeks (column 8).

6 Conclusions

The last set of results is important because the overall effect of scores on the allocation of capital is, in theory, ambiguous. Scores may drive committee members to cater to the score and ignore soft information. We find that on the contrary, scores interact with other organizational inputs to improve repayment above and beyond what the level of the score can predict.

Put together with the experimental results, the findings imply that scoring technologies improve the overall efficiency of lending through two mechanisms: 1) they lower the cost of loan production by increasing the productivity of committee members, decentralizing the decision process and substituting for more expensive inputs, and 2) they increase the efficiency of capital allocation.

By design, the pilot trial holds constant interest rates, managers' span of control, committee members' compensation, and other dimensions of the loan production process. However, the new equilibrium level for these variables is likely to change after score adoption (interest rates may drop to reflect the lower production costs, managers may

oversee more branches, etc.). Thus, the estimated effects on output might not capture fully the general equilibrium impact of scores. Changes along these dimensions are likely to induce further improvements in productivity.

Our results also highlight that there are two potential mechanisms that drive the increase in committee productivity: pure information and incentives. We evaluate how committee decisions change in a treatment that makes the score available to all committee members soon after they have reached a decision to disentangle the two mechanisms. We find that more than half of the increase in committee output occurs as a response to the anticipation of the scores becoming available, holding the committee's information set constant. This suggests that the incentive channel plays a first-order role in the overall effect of scores on lending productivity.

These findings have interesting implications regarding the design of incentives inside organizations. Information technology based solutions that increase the ease with which the principal can monitor the actions of the agents may have first-order effects on productivity and organizational design.²⁶ In our context, the supervisors are able to observe loan officers' choices even in the absence of the score, for example, when they review the loan officers' performance and bonus payments on a quarterly basis. Thus, scores increase the *immediacy* and *ease* with which the principal can monitor the agents but not *whether* they get reviewed. Scores also affect how salient the information is to both the agent and the principal, and is thus related to the work by Cadenas & Schoar (2011), who change the frequency of incentives to help loan officers overcome procrastination issues. It is suggestive that these relatively subtle changes in how agents are monitored induce significant changes in behavior. As such, information technology solutions may represent an effective and low cost alternative to steepening or increasing monetary incentives.

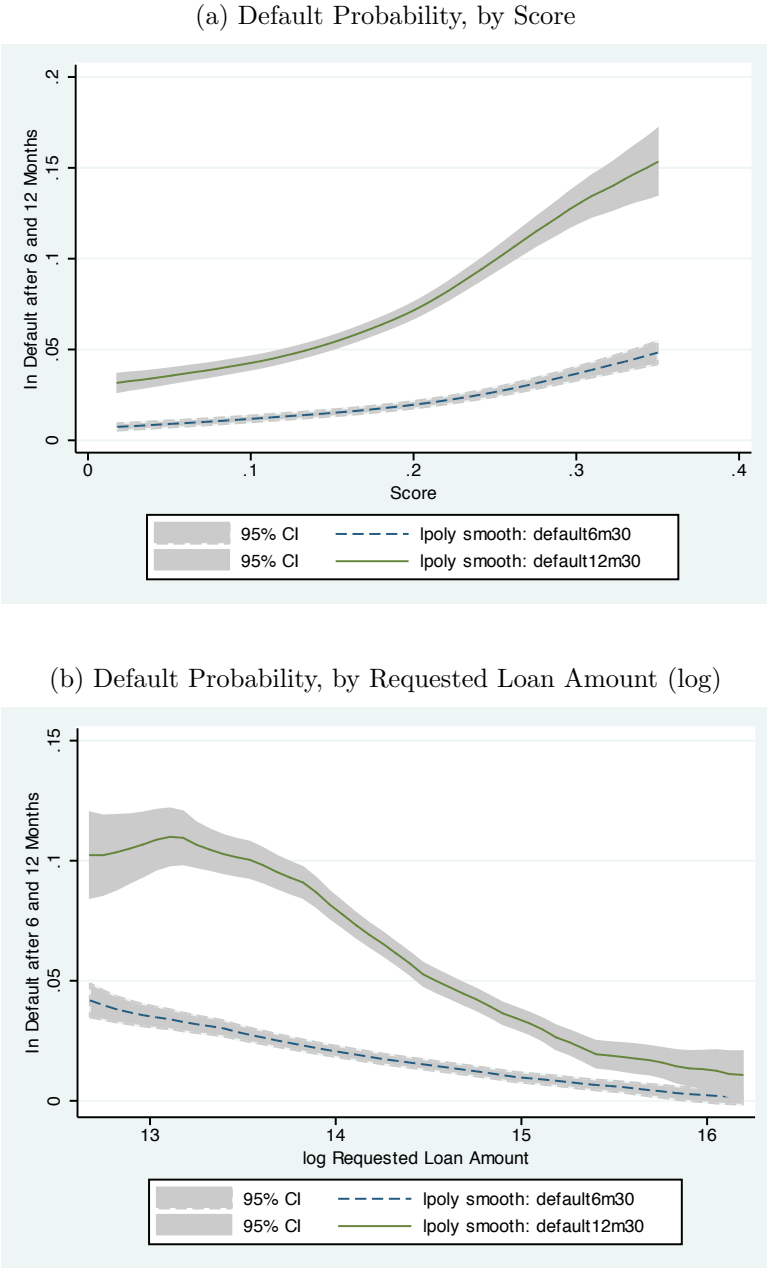
²⁶Note, in contrast, that our results have nothing to say about information technology based solutions that completely substitute committee judgement. Berg, Puri & Rocholl (2013) show evidence that piece-rate incentives and a lending rule solely based on a credit scoring model lead to loan officers distorting the information they enter into the computer terminal.

References

- Aghion, P. & Tirole, J. (1997), ‘Formal and Real Authority in Organizations’, *Journal of Political Economy* **105**(1), 1–29.
- Autor, D. & Dorn, D. (2013), ‘The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market’, *The American Economic Review* **103**(5), 1553–1597.
- Autor, D., Levy, F. & Murnane, R. (2003), ‘The Skill Content of Recent Technological Change: An Empirical Exploration’, *Quarterly Journal of Economics* **118**(4), 1279–1334.
- Baker, G. & Hubbard, T. (2004), ‘Contractibility and Asset Ownership: On-Board Computers and Governance in U.S. Trucking’, *Quarterly Journal of Economics* **119**(4), 1443–1479.
- Berg, T., Puri, M. & Rocholl, J. (2013), ‘Loan Officer Incentives and the Limits of Hard Information’, *NBER Working Paper* . 19051.
- Bloom, N., Garicano, L., Sadun, R. & Van Reenen, J. (2011), The Distinct Effects of Information Technology and Communication Technology on Firm Organization. Available at <http://cep.lse.ac.uk/pubs/download/dp0927.pdf>.
- Brynjolfsson, E. & Hitt, L. (2000), ‘Beyond Computation: Information Technology, Organizational Transformation and Business Performance’, *The Journal of Economic Perspectives* **14**(4), 23–48.
- Brynjolfsson, E. & Yang, S. (1996), Information Technology and Productivity: A Review of the Literature, in M. Zelkowitz, ed., ‘Advances in Computers’, Vol. 43, Elsevier, pp. 179–214.
- Cadenas, X. & Schoar, A. (2011), ‘Fighting Procrastination within the Firm’, *NBER Working Paper* . 16944.
- Dessein, W. (2002), ‘Authority and Communication in Organizations’, *The Review of Economic Studies* **69**(4), 811–838.
- Einav, L., Jenkins, M. & Levin, J. (2013), ‘The Impact of Information Technology on Consumer Lending’, *The RAND Journal of Economics* **44**(2), 249–274.
- Garicano, L. (2000), ‘Hierarchies and the organization of knowledge in production’, *Journal of Political Economy* **108**(5), 874–904.
URL: <http://www.jstor.org/stable/10.1086/317671>
- Hertzberg, A., Liberti, J. M. & Paravisini, D. (2011), ‘Public Information and Coordination: Evidence from a Credit Registry Expansion’, *The Journal of Finance* **66**(2), 379–412.

- Hubbard, T. (2000), ‘The Demand for Monitoring Technologies: The Case of Trucking’, *Quarterly Journal of Economics* **115**(2), 533–560.
- Katz, L. & Murphy, K. (1992), ‘Changes in Relative Wages, 1963-1987: Supply and Demand Factors’, *Quarterly Journal of Economics* **107**(1), 35–78.
- Lee, D. (2008), ‘Training, wages, and sample selection: Estimating sharp bounds on treatment effects’, *Review of Financial Studies* **76**(3), 1071–1102.
- Milgrom, P. & Roberts, J. (1990), ‘The Economics of Modern Manufacturing: Technology, Strategy, and Organization’, *American Economic Review* **80**(3), 511–528.
- Petersen, M. & Rajan, R. (1995), ‘The Effect of Credit Market Competition on Lending Relationships’, *Quarterly Journal of Economics* **110**(2), 407–443.
- Prat, A. (2005), ‘The Wrong Kind of Transparency’, *The American Economic Review* **95**(3), 862–877.
- Prendergast, C. (1993), ‘A Theory of Yes Men’, *The American Economic Review* **83**(4), 757–770.
- Rajan, R. (1992), ‘Insiders and Outsiders: The Choice between Informed and Arm’s-Length Debt’, *The Journal of Finance* **47**(4), 1367–1400.
- Rajan, U., Seru, A. & Vig, V. (2013), ‘The Failure of Models that Predict Failure: Distance, Incentives and Defaults’, *Journal of Financial Economics* **Forthcoming**.
- Stein, J. (2002), ‘Information production and capital allocation: Decentralized versus hierarchical firms’, *The Journal of Finance* **57**(5), 1891–1921.
URL: <http://dx.doi.org/10.1111/0022-1082.00483>

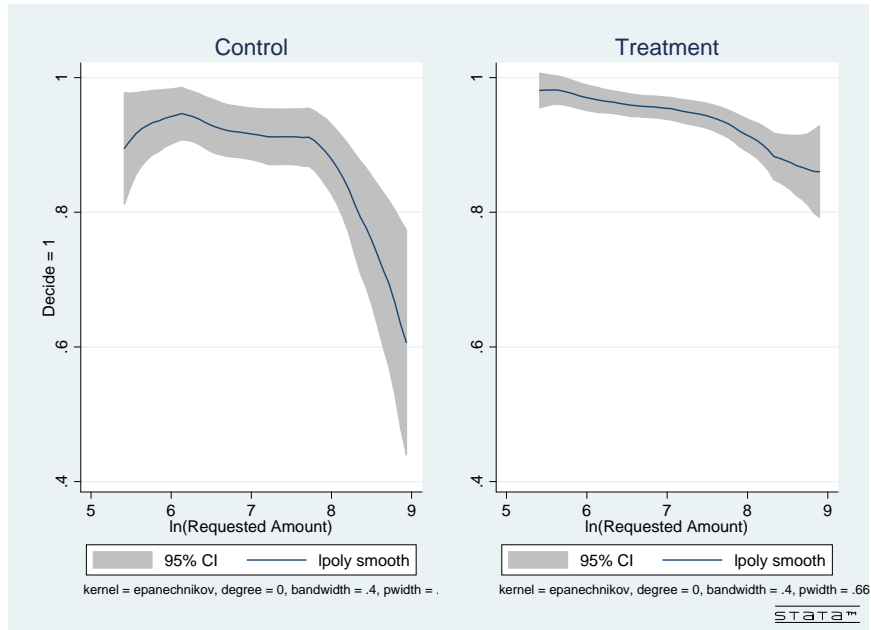
Figure 1: Population Relationships between Default Probability and Credit Scores/Requested Loan Amount



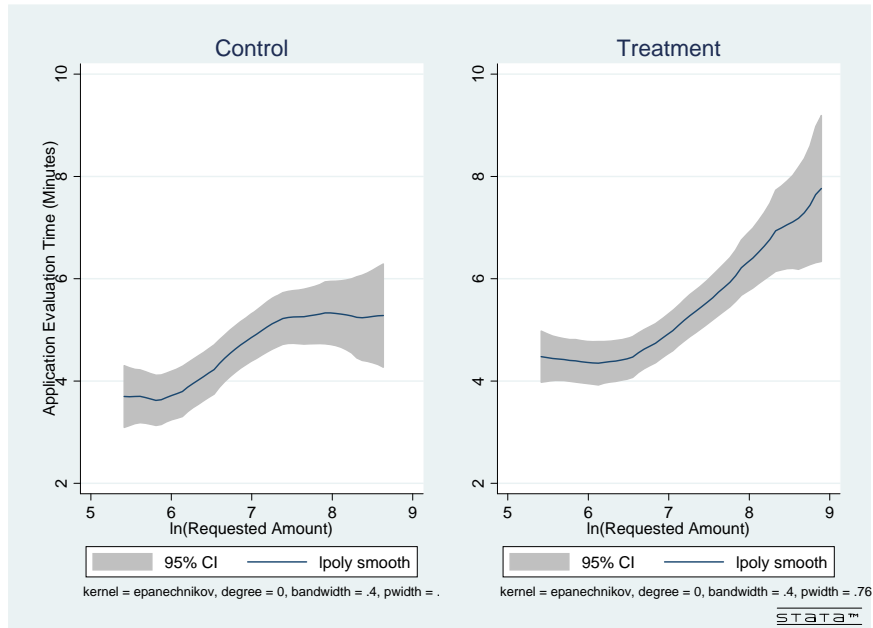
Non-parametric relationship between 6-month and 12-month default probabilities and (a) credit score, (b) requested loan amount, estimated on the sample of *all* loans approved by BancaMia during October 2010, one month before the roll-out of the randomized pilot program.

Figure 2: Probability of Decision and Evaluation Time, by Application Amount

(a) Probability that Committee Makes Decision

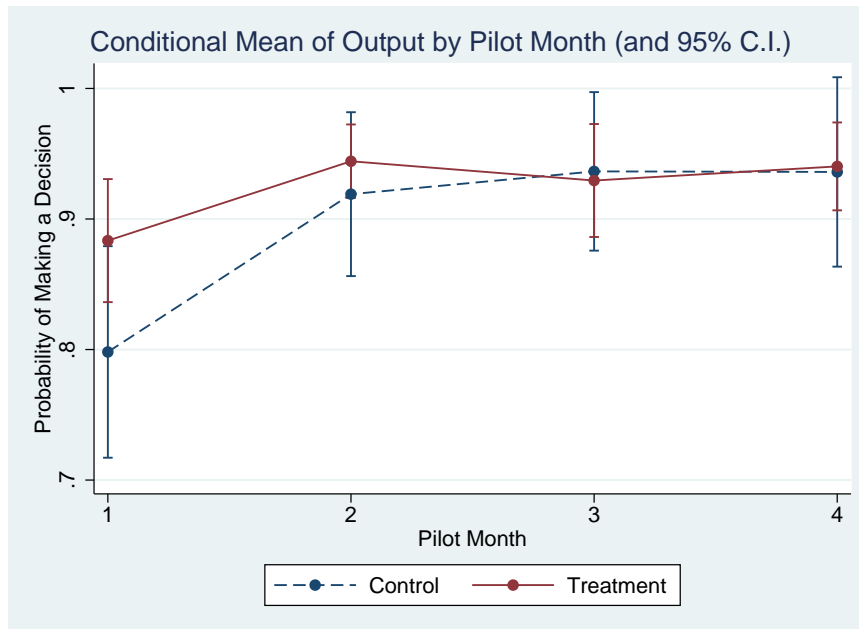


(b) Application Evaluation Time



Non-parametric relationship of (a) probability that committee makes a decision on an application (approve or reject) and (b) evaluation time, with application amount.

Figure 3: Committee Output by Pilot Month



Non-parametric relationship of (a) probability that committee makes a decision on an application (approve or reject) and (b) evaluation time, with application amount.

Table 1: Descriptive Statistics by Randomized Subsample

	(1)		(2)		(3)		(4)		
	Control		Treatment T1		Treatment T2		p-value		
	(n = 335)		(n = 563)		(n = 523)		(1) = (2)	(1) = (3)	(2) = (3)
	Mean	SD	Mean	SD	Mean	SD			
Panel A. Ex Ante Application Characteristics									
Application Amount (USD)	1,551.5	1,321.4	1,571.7	1,405.6	1,532.3	1,256.8	0.832	0.852	0.649
Application Maturity (Months)	20.9	9.8	22.0	10.4	22.1	10.4	0.109	0.086	0.880
Credit Risk Score	0.151	0.068	0.155	0.074	0.158	0.080	0.459	0.201	0.512
First Application (Dummy)	0.146	0.354	0.147	0.355	0.159	0.366	0.962	0.631	0.616
Panel B. Committee Outcomes									
Evaluation Time (minutes)	4.68	3.28	5.13	5.24	5.43	5.34	0.156	0.021	0.353
Committee Approves/Rejects (Dummy)	0.890	0.314	0.931	0.254	0.950	0.218	0.032	0.001	0.221
Panel C. Committee Outcomes, Conditional on Reaching decision									
Loan Approved (Dummy)	0.997	0.058	0.987	0.11	0.984	0.13	0.161	0.100	0.717
Panel D. Committee Outcomes, Conditional on Approval									
Amount Approved \neq Application (Dummy)	0.698	0.460	0.737	0.441	0.692	0.462	0.230	0.849	0.108
Approved Amount/Application Amount	0.979	0.435	0.975	0.318	0.950	0.293	0.905	0.271	0.187
Approved Approved - Application Amount	266.4	478.8	249.8	484.3	245.6	486.0	0.635	0.557	0.892
Maturity Approved \neq Application (Dummy)	0.262	0.440	0.278	0.449	0.307	0.462	0.609	0.174	0.314
Approved Maturity/Application Maturity	0.985	0.290	1.000	0.264	0.983	0.371	0.471	0.922	0.404
Approved Maturity - Application Maturity	2.3	4.7	2.4	5.0	3.2	6.0	0.616	0.023	0.032
Loan Issued (Dummy)	0.835	0.372	0.855	0.353	0.840	0.367	0.447	0.840	0.524
Panel E. Final Outcomes, Conditional on Loan Issued									
Disbursed Amount/Application Amount	0.959	0.382	0.965	0.297	0.974	0.549	0.828	0.702	0.755
In Default after 6 Months (Dummy)	0.036	0.188	0.039	0.193	0.037	0.190	0.881	0.954	0.917
In Default after 12 Months (Dummy)	0.095	0.293	0.088	0.283	0.089	0.284	0.757	0.781	0.976
Defaulted Amount (6 months)	27.26	166.22	26.43	147.97	35.04	193.26	0.947	0.604	0.476
Defaulted Amount (12 months)	62.67	238.07	71.84	257.73	74.11	265.39	0.650	0.584	0.902

Table 2: Descriptive Statistics by Committee Action, Control Group Applications

	Decide (n = 298) (1)		Send Up (n = 16) (2)		More Information (n = 21) (3)	
	mean	sd	mean	sd	mean	sd
Panel A. Ex Ante Application Characteristics						
Application Amount (US\$)	1,443	1,170	2,480	2,126	2,476	1,994
Application Maturity (Months)	20.3	9.3	26.3	12.2	25.1	13.3
Credit Risk Score	0.152	0.069	0.155	0.060	0.138	0.046
First Loan (Dummy)	0.154		0.125		0.048	
Panel B. Outcomes						
Time to decision (minutes)	4.608	3.188	5.438	3.405	5.105	4.508
Loan Issued (Dummy)	0.832		0.750		0.714	
Amount Approved \neq Application (Dummy)	0.924		1.000		1.000	
Approved Amount/Application Amount	0.945	0.272	0.950	0.227	1.486	1.807
Approved - Application Amount	287.4	499.1	262.9	309.8	1477.0	2153.0
In Default after 6 Months (Dummy)	0.028		0.000		0.200	
In Default after 12 Months (Dummy)	0.093		0.083		0.133	
Defaulted Amount (after 6 Months)	25.8	164.0	0.0	0.0	121.4	321.2
Defaulted Amount (after 12 Months)	66.4	246.6	28.6	99.2	0.0	0.0

Comparison of application characteristics where the officer reaches a decision—approves or rejects application—(column

1), those where the officer refers the application to the regional manager (column 2), and those where the committee decides to collect additional information (column 3).

Table 3: Average Treatment Effect of Scores on Committee Output

Estimation: Dependent Variable:	OLS	
	Committee Decides	
	(1)	(2)
Treatment (T1 and T2)	0.046** (0.018)	0.038** (0.019)
Treatment (T2)		0.016 (0.013)
ln(Application Amount)	-0.032*** (0.012)	-0.032*** (0.012)
ln(Application Maturity)	-0.026 (0.018)	-0.026 (0.018)
Credit Risk Score	-0.105 (0.111)	-0.107 (0.111)
First Application	0.009 (0.018)	0.009 (0.018)
Trend	Yes	Yes
Observations	1,406	1,406
R-squared	0.041	0.042

OLS estimates of the effect of treatment on committee output measured as the probability that committee reaches decision, with robust standard errors in parenthesis. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Table 4: Costly Information and Referrals

Committee Choice	Approves (Omitted)	Rejects	More Information	Send to Manager
	(1)	(2)	(3)	(4)
Treatment (T1 and T2)		1.3236 (1.049)	-0.5439* (0.305)	-0.9038*** (0.344)
ln(Application Amount)		0.0851 (0.466)	0.7971*** (0.243)	0.1790 (0.308)
ln(Application Maturity)		-0.2176 (0.731)	-0.1358 (0.386)	1.7474*** (0.570)
Credit Risk Score		4.8843** (2.121)	0.2146 (2.001)	2.8701* (1.537)
First Application		0.4442 (0.673)	-0.6193 (0.487)	0.2574 (0.419)
Trend		Yes	Yes	Yes
Observations	1405			
Pseudo R-squared	0.0875			
Fraction in Control Subsample	0.8866	0.0030	0.0627	0.0478
Marginal Effects:				
Treatment	0.0281 (0.0166)	0.0124 (0.0100)	-0.0174* (0.0104)	-0.0231** (0.0094)

Multinomial Logistic Regression estimates of the effect of treatment on final committee actions: make a decision on an application (approve or reject), postpone until the loan officer collects additional information, or send the application to the manager (referrals). The first action, make a decision, is the omitted category. The bottom rows present the proportion of each action in the control group and the estimated marginal effect of treatment on the probability that the committee takes an action. Robust standard errors in parenthesis. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Table 5: Average Treatment Effect of Scores on Committee Effort

Dependent Variable:	OLS		LAD (Quantile Regression)							
	Evaluation Time		Evaluation time							
	(1)	(2)	25th %ile	50th %ile	75th %ile	(3)	(4)	(5)	(6)	(7)
Treatment (T1 and T2)	0.760*** (0.229)	0.624** (0.273)	0.183 (0.192)	0.123 (0.194)	0.426*** (0.157)	0.383** (0.193)	0.663*** (0.231)	0.662*** (0.254)		
Treatment (T2)		0.283 (0.321)		0.057 (0.170)		0.133 (0.171)		0.113 (0.249)		
ln(Application Amount)	0.751*** (0.226)	0.747*** (0.226)	0.084 (0.145)	0.053 (0.140)	0.239* (0.133)	0.238 (0.145)	0.664*** (0.183)	0.617*** (0.185)		
ln(Application Maturity)	0.653** (0.332)	0.656** (0.330)	0.676*** (0.222)	0.692*** (0.222)	0.680*** (0.208)	0.670*** (0.225)	0.227 (0.308)	0.265 (0.303)		
Credit Risk Score	-1.28 (1.457)	-1.301 (1.458)	-1.181 (0.752)	-1.175 (0.818)	-1.970** (0.945)	-1.464 (1.077)	-1.761 (1.564)	-1.971 (1.558)		
First Application	0.695* (0.387)	0.692* (0.388)	0.408** (0.187)	0.426** (0.194)	0.541*** (0.186)	0.571*** (0.204)	0.764 (0.509)	0.757 (0.535)		
Trend	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes		
Observations	1,397	1,397	1,397	1,397	1,397	1,397	1,397	1,397		
R-squared	0.048	0.049								

OLS estimates of the effect of treatment on application evaluation time in minutes, with robust standard errors in parenthesis (columns 1 and 2). LAD estimates of the effect of treatment on evaluation time, with bootstrapped standard errors (500 repetitions) estimated via simultaneous quantile regressions in parenthesis (columns 3 through 8). ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Table 6: Scores and Conditional Committee Outcomes – OLS

Dependent Variable:	Committee Approves Dum. (1)	Approved Amount (log USD) (2)	Approved Maturity (log Months) (3)	Loan Issued Dum. (4)	Issued Amount (log USD) (5)	In Default after 12 months (6)
Panel A. Subsample: applications where the committee made the decision during the pilot						
Treatment (T1 and T2)	-0.0092* (0.005)	-0.0001 (0.020)	0.0282 (0.029)	0.0051 (0.025)	0.0114 (0.024)	-0.0093 (0.020)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Conditional of Committee Approving	No	Yes	Yes	Yes	Yes	Yes
Conditional on Loan Issued	No	No	No	No	Yes	Yes
Observations	1,315	1,314	1,314	1,302	1,001	1,001
R-squared	0.007	0.840	0.010	0.011	0.778	0.043
Panel B. All applications						
Treatment (T1 and T2)				0.0020 (0.024)	0.0107 (0.018)	-0.0158 (0.020)
Controls				Yes	Yes	Yes
Conditional on Loan Issued				No	Yes	Yes
Observations				1,406	1,048	1,048
R-squared				0.007	0.773	0.043

OLS regressions of conditional outcomes on treatment status. Panel 1 shows estimates on the subsample of applications where the committee reached a decision (approves or rejects) during the pilot. Panel 1 shows estimates on all applications, including those where the decision was made by the regional manager and those where the committee made a decision later. Columns (2) through (4) are estimated on the subsample of applications where the committee approved an application. Columns 5 and 6 further conditions on issued loans. Robust standard errors in parenthesis. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Table 7: Information versus Incentives: Effect on Interim and Final Actions in $T2$ – OLS

Outcome:	Committee Decides		Committee Approves		ln(Approved Amount)	
	Dummy		Dummy			
Choice:	Interim	Final	Interim	Final	Interim	Final
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment T2	0.0388** (0.019)	0.0524*** (0.018)	-0.0136** (0.007)	-0.0113* (0.006)	-0.0023 (0.023)	-0.0064 (0.022)
ln(Application Amount)	-0.0257 (0.016)	-0.0378** (0.015)	0.0013 (0.007)	0.0012 (0.007)	0.746*** (0.022)	0.776*** (0.021)
ln(Application Maturity)	-0.0572** (0.024)	-0.0352 (0.023)	-0.0015 (0.015)	-0.0012 (0.015)	0.280*** (0.036)	0.237*** (0.035)
Credit Risk Score	-0.0925 (0.132)	-0.1607 (0.132)	-0.1877 (0.129)	-0.1632 (0.129)	-0.821*** (0.155)	-0.762*** (0.152)
First Application	0.0402* (0.021)	0.0318 (0.021)	-0.0034 (0.010)	-0.0052 (0.010)	-0.0197 (0.031)	-0.0152 (0.031)
Trend	Yes	Yes	Yes	Yes	Yes	Yes
Observations	851	851	787	793	783	793
R-squared	0.050	0.056	0.023	0.019	0.838	0.845

OLS estimates the treatment effect on interim committee outcomes before observing the score (odd columns) and on final outcomes after observing the score (even columns). Robust standard errors in parenthesis. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Table 8: Branch-Level Effects - Difference-in-Differences Estimates

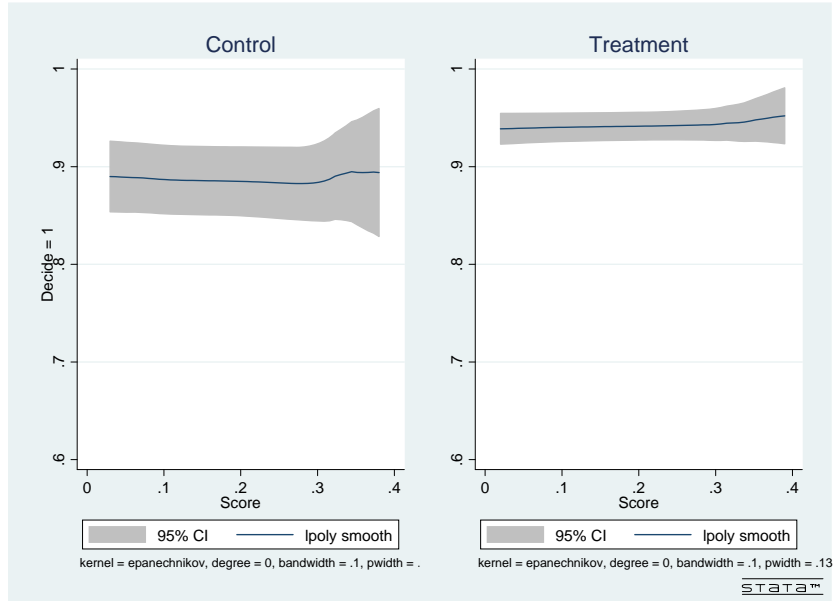
Unit of Observation:	Loan			Branch-Week				
	Score	ln(Application Amount)	In Default after 12 Months	ln(Number of Loans)	ln(Sum Loan Amount)	Fraction of Loans that Defaults	Realized Return Average	Std. Dev.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Experiment Week	-0.0014 (0.002)	-0.0194 (0.020)	-0.0071 (0.013)	-0.0273 (0.044)	0.0074 (0.047)	-0.0121* (0.006)	0.0152* (0.008)	-0.0392** (0.015)
Risk Score			0.6949*** (0.039)					
Experiment Week x Risk Score			-0.0812 (0.082)					
Branch Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Branch Trends	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	18,270	18,270	18,270	525	525	525	525	525
R-squared	0.044	0.028	0.016	0.754	0.699	0.283	0.284	0.17

OLS regression of loan characteristics on a dummy equal to one if the application was evaluated during a week in which the randomized pilot study was taking place in the branch. Sample contains only approved loans from the eight pilot branches and eight propensity score-matched branches (branch matching based on number and total amount of loans approved, average approved loan size, and borrower score measured in October 2010). The sample period is from week 41 of 2010 to week 26 of 2011 (four weeks before and after the pilot program began and ended). Columns 1 through 3 are estimated at the loan level, and 3 through 9 at the branch-week level. Robust standard errors clustered at the branch level in parenthesis. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

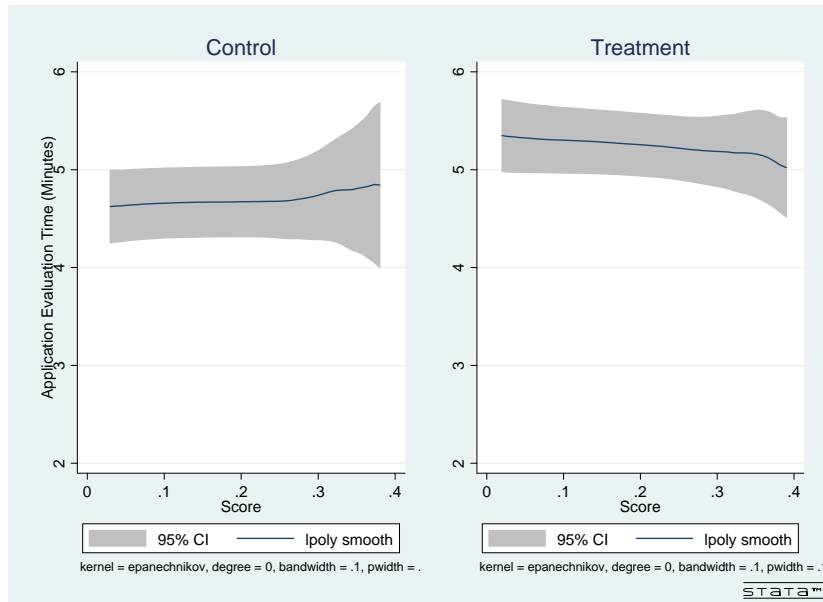
A Appendix

Figure A.1: Probability of Decision and Evaluation Time, by Score

(a) Probability that Committee Makes Decision



(b) Application Evaluation Time



Non-parametric relationship of (a) probability that committee makes a decision on an application (approve or reject) and (b) evaluation time, with application score.

Table A.1: Study Sample: Number of Applications per branch and per Treatment Status

	Control	T1	T2	Total
Branch #:				
1	44	67	62	173
2	89	153	132	374
3	26	51	66	143
4	69	88	87	244
5	18	28	27	73
6	22	26	14	62
7	20	45	38	103
8	47	105	98	250
Total	335	563	524	1,422

Control: the committee makes a decision without observing the score. *T1*: the borrower's score is made available at the beginning of the application evaluation. *T2*: the committee makes an interim decision before the score is made available, and allowed to revise the decision after observing the score.

Table A.2: Effect of Scores on Committee Output, No Controls

Sample Conditioning: Dependent Variable:	None		Committee Decides	Committee Approved		Loan Issued	
	Evaluation Time (1)	Committee Decides (2)	Committee Approves (3)	ln(Approved Amount) (4)	Loan Issued (5)	ln(Issued Amount) (6)	Defaults (7)
Score Dummy	0.5962** (0.242)	0.0506*** (0.019)	-0.0113** (0.005)	0.0281 (0.050)	0.0182 (0.028)	0.0541 (0.056)	0.0099 (0.014)
Observations	1,412	1,421	1,319	1,315	1,303	1,001	1,001
R-squared	0.003	0.007	0.002	0.000	0.000	0.001	0.000

OLS estimates of the effect of treatment on committee and loan outcomes. Columns (1) and (2) are estimated on all applications, columns (3) and (4) on the subsample of applications where the committee reached a decision, column (5) on the subsample of approved applications, and columns (6) and (7) are estimated on the subsample of issued loans. Robust standard errors in parenthesis. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.

Table A.3: Interim and Final Decisions in Treatment $T2$

	Final Decision (after Observing Score):				Total
	Accept Application	Reject Application	Obtain More Information	Send Decision to Manager	
Interim Decision:					
Accept	482	0	0	1	483
Reject	0	8	0	0	8
Obtain More Information	0	0	20	0	20
Send Decision to Manager	7	0	0	5	12
Total	489	8	20	6	523

Each observation in the matrix represents the two sequential decisions made by a committee regarding the *same* application in treatment $T2$. Interim decisions (rows) are the decisions made before observing the score and final decisions (columns) are the revised decisions after observing the score.

Table A.4: Information versus Incentives: Effect on Interim and Final Actions in T2 – ML

Estimation Action:	Interim Outcomes				Final Outcomes			
	Approve (Omitted) (1)	Reject (2)	More Information (3)	Send to Manager (4)	Approve (Omitted) (5)	Reject (6)	More Information (7)	Send to Manager (8)
Treatment T2		1.5208 (1.071)	-0.4868 (0.347)	-0.8152** (0.392)		1.4762 (1.058)	-0.5112 (0.346)	-1.5197*** (0.483)
ln(Application Amount)		-0.1783 (0.698)	0.8793** (0.359)	-0.1355 (0.453)		-0.1558 (0.701)	0.9058** (0.358)	0.5286 (0.501)
ln(Application Maturity)		0.1960 (1.352)	0.0517 (0.575)	1.8806*** (0.652)		0.1588 (1.342)	0.0159 (0.568)	1.0678* (0.618)
Credit Risk Score		6.3225** (2.507)	1.1175 (1.738)	1.7954 (2.708)		6.4498** (2.566)	1.3050 (1.774)	4.2580 (2.806)
First Application		0.2935 (0.642)	-0.9911 (0.644)	-0.5589 (0.656)		0.3135 (0.644)	-0.9733 (0.643)	-0.4016 (0.766)
Trend		Yes	Yes	Yes		Yes	Yes	Yes
Observations	850				850			
R-squared								
Pseudo R-squared	0.0975				0.114			
Fraction in Control Subsample	0.8866	0.0030	0.0627	0.0478	0.8866	0.0030	0.0627	0.0478
Marginal Effects:								
Treatment T2	0.0279 (0.0201)	0.0140 (0.0104)	-0.0172 (0.0131)	-0.0247* (0.0128)	0.0392* (0.0204)	0.0138 (0.0102)	-0.0170 (0.0129)	-0.0360*** (0.0136)

Multinomial Logistic Regression estimates of the effect of treatment on interim committee decisions made before observing the score (columns 1 through 4) and on final decisions made after observing the score (columns 5 through 8). The bottom rows present the proportion of each action in the control group and the estimated marginal effect of treatment on the probability that the committee takes an action. Robust standard errors in parenthesis. ***, **, and * indicate significance at the 1%, 5%, and 10% levels.