

School of Finance



University of St.Gallen

CREDIT SCORING VS. EXPERT JUDGMENT – A RANDOMIZED CONTROLLED TRIAL

THOMAS GIETZEN

WORKING PAPERS ON FINANCE NO. 2017/09

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

JUNE 2017



Credit Scoring vs. Expert Judgment A Randomized Controlled Trial

- June 2017 -

Thomas Gietzen*

Abstract

Developing financial markets experience a swift increase in the availability of borrower-information from credit information sharing systems. I study whether banks can use this information to automate credit decisions. In the wake of a randomized controlled trial, a bank in Africa introduced an automated credit decision-process based on a credit scoring technology at half of its branches, while the other half kept applying an extensive screening procedure as a base for a loan officer's expert judgment. Results show that the quality of the loan portfolio in the treatment branches did not decrease significantly, at the cost of rejecting only a 6 percentage points higher share of applications, using a much simpler procedure. An analysis of the costs and benefits of the credit scoring system strongly suggests that the bank's cost of lending decreased substantially.

Keywords: Credit Scoring, Credit Information Sharing, Credit Bureaus, Loan Officer, Automation

*KfW Development Bank, Evaluation Department, Palmengartenstrasse 5-9, 60325 Frankfurt am Main, Germany, Thomas.Gietzen@kfw.de, +49(0)69-74316239 and Swiss Institute of Banking and Finance, University of St.Gallen

I would like to thank Tobias Berg, Martin Brown and Eva Terberger. I would also like to thank seminar participants at the University of St.Gallen.

Disclaimer: Any views or opinions presented in this paper are solely those of the author and do not necessarily represent those of the company.

1 Introduction

The technologies that banks apply to screen applicants are a defining part of their business. Much of the operational costs associated with giving credit are immediately related to the screening procedures. In small business lending in developing financial markets it has been common practice for banks to screen applicants extensively, owing to the fact that there is little verifiable information from tax records or bookkeeping. To overcome the informational barrier, lenders elicit and analyze applicants' accounting data and visit their businesses on site. The operational costs of such extensive screening procedures are, however, high and are detrimental to economic development. High fixed costs relative to small loan volumes are one of the main drivers of high interest rates borrowers of small loans often have to pay (Banerjee and Duflo 2010). Against the backdrop of this prevailing technology, in recent years, there has been a large increase in the availability of hard information about borrowers from credit information sharing systems (also called credit bureaus). According to the World Bank's Doing Business Report, 54 countries had no active credit information sharing system in place in 2007, while this number has more than halved by 2015. We know from more developed financial markets that this information can be helpful in predicting borrowers' performance (Kallberg and Udell 2003; Dierkes et al. 2013; Doblas-Madrid and Minetti 2013), but there is little evidence on its usefulness for small business lenders. Yet, in developing financial markets too, the emergence of this new source of information could facilitate a shift towards automated credit decisions. From a developmental perspective, this would be most welcome. Screening borrowers more efficiently could be a key element in eliminating credit constraints, in particular for borrowers of small loans who suffer most from the informational barrier between banks and their borrowers.

By means of a randomized controlled trial (RCT) conducted by a bank in Africa, this study seeks to answer the question whether data from a credit information sharing system enables banks to shift from a screening process based on the expert judgment of a loan officer to automated credit decisions from a credit scoring system. In the wake of the experiment, the bank - a small business lender - implemented a credit scoring lending technology based on credit bureau information at half of its branches that were chosen randomly. The other half of its branches kept applying the previous, more soft information-based and more labor intensive lending technology based on loan officers' expert judgment. Identification of the effects of introducing an automated credit scoring system is further provided by the fact that subsequent to the pilot period of four months, the bank rolled out the credit scoring system to all of its branches. The bank-wide roll out allows me to check whether any differences between the branches disappear as soon as the control branches start using the new system as well.

This study is the first to evaluate a major change in the credit technology of a bank towards automated lending decisions in an experimental setup. Its results show that replacing an

extensive screening technology based on loan officers' judgment with credit scoring can lead to significant efficiency gains in small business lending: Tracking eventual borrowers' repayment behavior, I find that the introduction of the automated credit scoring system did not lead to a significant decrease in the portfolio quality at the treatment branches. I ensure that the result is not driven by systematic changes in the types of loans the credit scoring branches grant. In itself - of course - the comparable performance of the credit scoring branches could be the result of a very conservative scoring system that rejected many applications because every screening system trades off type I error (accepted and defaulting borrowers) with type II errors (rejected applicants who would not have defaulted). To measure the effect of credit scoring on the quantity of credit and possible type II errors, I compare the probability of a loan application being rejected in treatment and control branches. Indeed, the likelihood of an application being rejected increases by 6 percentage points in the branches that introduce credit scoring. In a final analysis, I weigh the costs of rejecting more loans against the benefits of decreasing the costs of screening borrowers calibrated to the empirical findings of this study. The analysis of costs and benefits strongly suggests that the bank's net profits from granting loans increase under the credit scoring system, even for highly conservative assumptions about operational cost savings.

Credit scoring means estimating the default probability of a loan application based on a model that was built using historical data. In fact, credit scoring is linked to the emergence of credit registry data and dates back to the 1950s (Mester 1997). Assigning scores based on credit bureau data to evaluate loan applications was first used in retail lending, in particular credit card lending, and has since found its way into small business lending in developed economies where scoring models usually rely on both credit bureau and firms' financial data.

The primary reason for banks to use credit scoring is to lower operational costs.¹ DeYoung, Glennon, and Nigro (2008) suggest that cost savings from credit scoring may turn a marginal loan application with a negative expected return profitable to the bank. This should increase both the volume and the riskiness of the loan portfolio, given that credit scoring works just as well as the previous way of screening borrowers. Yet, while scoring may be a cheap way to evaluate the creditworthiness of borrowers, it may or may not produce less accurate predictions about the default probability of a loan application. In the context of small business lending in developing countries, for the lack of alternatives, soft information and a loan officer's expert judgment have usually been deemed indispensable.

Even though credit scoring has been used extensively in practice, empirical evidence is scarce. That may be because studying credit scoring presupposes bank-internal, micro-level data that most banks may be unwilling to supply. In fact, most of the existing studies

¹Sometimes it may also be used to render the screening process less ambiguous by basing it on hard information only. However, Berg, Puri, and Rocholl (2016) show that even supposedly hard information at banks may be subject to manipulation by loan officers.

are based on a survey conducted by the Federal Reserve Bank of Atlanta that asked 190 U.S. banks about their use of credit scoring in 1998. These studies associate small business credit scoring with higher lending volumes to small businesses and in some cases with higher riskiness of loans granted (Frame, Srinivasan, and Woosley 2001; Frame, Woosley, and Padhi 2004). There are studies that combine the Fed Atlanta data with other bank-survey data: Allen N. Berger and Miller (2005) find a positive correlation between credit scoring and the share of small-business loans in a bank's portfolio. DeYoung, Glennon, and Nigro (2008) find that both greater distance and credit scoring are correlated with higher defaults, but the negative effect of distance does not apply to banks that use credit scoring. They suggest that credit scoring outperforms soft information lending technologies when the borrower lender distance becomes large.

Naturally, the cross sectional survey data studies share common weaknesses of respondents' bias and of failing to rule out reverse causality (the use of credit scoring may follow a strategic shift towards small business loans). Here, I overcome such weaknesses by studying credit scoring in a unique experimental setup. The only other micro-level study on the adoption of credit scoring in particular is Einav, Jenkins, and Levin (2013) in a very different setting to this study. Einav, Jenkins, and Levin (2013) use data from an automobile finance company which introduced credit scoring and find that the lender's profitability increased for two reasons: Credit scoring allowed the lender to screen out high-risk clients and at the same time to target low-risk clients with bigger cars and larger loans.

As the credit scoring model used in this study's experiment was based on data from a credit information sharing system, this study also adds to the growing literature on the usefulness of such data in general. In a setting of highly developed financial markets, Kallberg and Udell (2003) for the U.S. case and Dierkes et al. (2013) for the German case find that credit bureau data offers additional predictive power over firms' financial data. Doblas-Madrid and Minetti (2013) use the staggered entry of lenders to an information sharing system in the U.S. to show that access to this information reduces defaults in the bank's portfolio. My results suggest that these findings from developed markets can be transferred to small business lending in opaque markets and also help to explain why credit information sharing can resolve problems of adverse selection between banks and borrowers in general (Bos, De Haas, and Millone 2015; Liberti, Seru, and Vig 2016; Gietzen 2016). Eventually, this underpins macro-level findings that credit information sharing reduces credit risk (Jappelli and Pagano 2002) and increases the private credit to GDP ratio (Djankov, McLiesh, and Shleifer 2007).

From a developmental point of view, the implications of the finding that credit decisions in small business lending can be automated are wide ranging. With the help of high quality credit bureau information and computer-based credit scoring, costs for banks to reach out to borrowers could be reduced dramatically. This will have the strongest positive effect on

borrowers of small loans, who - due to the high fixed costs relative to loan volumes - face the most severe credit constraints.

2 Methodology


2.1 The Experiment

In the wake of the experiment the lender replaced its screening technology based on loan officers's expert judgments with a credit scoring system at half of its branches chosen randomly from a total of 28 branches. Credit scoring was applied to all borrowers at the treatment branches. Due to the low number of branches and to ensure that randomization would still work in practice, I stratified the branches prior to the random draw based on data previously obtained from the bank. I separated treatment and control branches into four groups along the median of two of the outcome variables, the Portfolio-at-Risk over 90 days (Par 90) and the total lending volume of the branch (in that order), and then randomly drew from the four buckets.² The treatment period lasted from May to end of August 2015. At the beginning of September, the bank rolled out the new credit scoring technology to all branches and I define September to end of December 2015 to be the post-treatment period. I additionally use data from a pre-treatment period, lasting from January 2014 to end of April 2015, to check for pre-treatment trends and to support the RCT assumptions. Figure 1 shows the number of loan applications in control and treatment branches during the pre-, post- and treatment period. Throughout, the number of applications is a little higher in the treatment branches.

2.2 Control Branches

In the control branches, the lender applies an extensive screening process that can be considered standard in much of microfinance and small business lending in developing financial markets. The screening procedure is centered around a site visit by the loan officer to the borrower which is why the costs of this screening procedure are particularly high for small loans. During the site visit, the loan officer inspects the physical appearance of an applicant's business and gathers information about the financial situation, all in collaboration with the applicant. This data includes the borrower's monthly income and expenses (both business and private) and also a basic balance sheet geared towards items of interest to the bank, including liquid assets (e.g. cash), fixed assets (e.g. a vehicle) and current liabilities like loans with other banks. The information about the current financial status of

²On the usefulness of stratifying according to outcome variables see Bruhn and McKenzie (2009).

**Figure 1**

Timeline and Number of Observations

the applicant culminates in financial ratios. The most important ratios are the debt-service ratio (the ratio of the future interest and principal payments divided by earnings) and the future installment to equity ratio. The financial ratios form the basis of information for the credit committee that takes the final decision on the application. The loan officer who visited the borrower is part of the committee that also consists of other loan officers, branch managers or regional managers depending on the size of the loan. A considerable amount of soft information enters the screening process. The loan officer's subjective judgment about the creditworthiness of the borrower may affect the financial data gathered during the site visit and the loan officer also influences the credit committee's decision making process.³ Importantly, even the loan officers in the control branches had access to information from the credit bureau; but in the form of a simple list of information (a standard credit report) other than a credit score that condenses all the information from the standard credit report into a prediction about the borrower's likelihood to default.

2.3 Treatment Branches

In the treatment branches, the previous lending technology was replaced by a credit scoring system. Due to a much larger degree of automation, the credit scoring is a lot less labor intensive (and therefore less expensive) than the screening technology in the control branches. The system estimates a score for each application based only on the existing information

³There is an extensive literature on the role of the loan officer in bank lending (for example Hertzberg, Liberti, and Paravisini 2010; Qian, Strahan, and Yang 2015, and the sources within).

about an applicant in the credit bureau.⁴ The credit bureau covers all loans granted by financial institutions that are regulated by the central bank. The data includes the full loan details (interest paid, loan term, loan volume, lending institution, possible arrears, collateral posted, guarantors) but also additional information like bounced checks, guarantees given, rejected or canceled applications. The system also records a certain amount of personal information like age, gender (not considered as a possible scoring variable), marital status and employment status.

Credit scores that range from 0-800 represent estimated probabilities of default, with higher scores indicating a lower risk of default. Prior to the beginning of the treatment-period, the lender had chosen two cutoff scores indicating risks that should immediately be accepted (green light) or risks that should immediately be rejected (red light) at the treatment branches. For all applications in-between (orange light) the lender decided to follow the previous loan officer-based technology (the control branches' technology). Whenever there was insufficient information about an applicant in the system (because he or she may have been a first-time borrower), the system offered no prediction and the application fell into the orange category.

The score does not, however, translate into a decision one-to-one. In addition to the score, other rules are able to trigger a decline decision by the system, for example when borrowers are too young, in default at other institutions or have more than a pre-specified number of other loans outstanding. Finally, the system's decisions are only recommendations and the loan officer, branch manager or regional manager (depending on the volume of the loan) is able to override the system's decision. The possibility of overrides of the system's decision was intentionally built into the process to avoid severe disruptions in the loan screening process due to possibly faulty system decision. All overrides of the system's decision had to be signed off by management.

2.4 RCT Assumptions

By design, the experiment does not suffer from drop out or substitution, as all control and treatment branches stuck to their assigned lending technology all the way through the treatment period.⁵ There is no attrition as all borrowers recorded in the system are monitored by the bank and their performance is being reported back to the system.

⁴The credit scoring model was developed by the operator of the credit information sharing system. During the development process, historical data on defaulted and non-defaulted loans was fed into prediction models that help to decide which pieces of information best predict whether an applicant is likely to default on a loan or not (i.e. the choice of predictors is not motivated by theory but is a purely statistical procedure).

⁵The RCT assumptions are: 1) no dropout of treatment, i.e. units assigned to treatment actually receive it 2) no substitution of controls into other, similar treatments 3) no attrition bias 4) stable unit treatment value assumption (SUTVA), i.e. the potential outcome for each

To minimize general equilibrium effects, other lenders had not been informed of the experiment (although they may have learned about it at some point) and the new technology did not affect the bank's interest rate policy during the experiment that other lenders' could have competed with. Furthermore, the bank's share of the total assets in the central bank regulated banking system is too small (slightly less than 1%) to have caused serious repercussions in the banking market.

At the borrower level, the experiment's design does not render it susceptible to violations of the no randomization bias or the stable unit treatment value assumption. The individual default probability does not depend on the treatment status of other borrowers nearby. There was never any active communication to borrowers or applicants at the bank about a change in the credit technology and the experiment was never made public to anyone besides the bank's staff. Therefore, borrowers are unlikely to have chosen their branch of the bank deliberately according to the new technology, also because traveling to another bank branch is costly. As the information channels to the outside were strictly limited, it is also unlikely that borrowers started applying to other banks because of the experiment.

The most serious source of possible violations of the RCT assumptions in this experiment is the extra effort that treatment or control units exert because of participating in an experiment. Although the bank's employees never received active communication that an experiment was being undertaken, they surely understood that the introduction of the new technology in only half of the bank's branches was a test-situation.⁶ Loan officers may have feared to be replaced by the new technology and may have worked more to beat the system or less because their motivation dropped. It is, however, instructive to note, that the loan officers' task is to reject an application or not. Loan officers play only a minor role for the default probability of a single loan after it has been granted as there is no standardized monitoring of the loans or follow-up visits, at least not before the loans are severely delinquent. Therefore, in my setup, loan officers can affect the loan performance by granting fewer or more risky loans but not by altering borrowers' behavior. Fortunately, I am able to use data on the system's recommendation in the control branches (those are hypothetical values, i.e. they were calculated but never applied or shown to anyone) and data from the pre- and post-treatment period to investigate changes in the behavior of rejecting or accepting applications in the control branches (see Chapter 3). Furthermore, the four months length of the experiment is likely to capture a normalization of loan officers' behavior, even more so as loan officers and branch managers did not know about the four months period while the experiment was going on.

borrower is unrelated to the treatment status of others 5) no randomization bias, i.e. the applicant pool remains unchanged despite the experiment.

⁶To the extent that a non-experimental introduction of credit scoring would also prompt an increase or decrease in the motivation of loan officers, changes in the behavior of staff in the treatment branches are part of the effect I try to measure.

2.5 Estimation

I estimate the effect of credit scoring on the probability of loans to default later in their lifetime at the borrower level, using Probit models (here in latent variable formulation), for each of the different treatment periods separately.

$$Y^* = \alpha + \beta T + \delta \mathbf{Z} + \epsilon$$

The outcome variable is the Portfolio-at-Risk over 30 days (Par 30), a dummy variable that is equal to one whenever a borrower is more than 30 days late on any repayment after the start of the loan contract over the lifetime of the loan.⁷ T is an indicator variable equal to one if a borrower applied at a treatment branch and zero if the borrower applied at a control branch. Despite the full randomization and stratification, I use two additional control variables at the branch level to increase the estimations' precision (contained in the matrix $\mathbf{Z}$). I control for the distance of the respective branch to the headquarter to capture possible effects of differences in managerial control between treatment and control branches caused by their geographic location. Additionally, I control for the competitive environment of the branch and use the number of other bank branches in the vicinity of 5 km as a proxy for that.

Next, I check whether credit scoring led to a change in the type of loans that the treatment branches grant which may in turn have affected the likelihood of those loans to default. In a set of linear models (estimated by OLS), I use the interest rate, the loan volume and the loan term of the granted loans as outcome variables to see whether credit scoring systematically affects other loan features.

Any screening technology trades off type I error (defaulting borrowers) with type II errors (rejected applicants who would not have defaulted) by being more or less restrictive in terms of accepting or declining applications. To capture the second part of this trade-off, I estimate whether the credit scoring affected likelihood of a loan application to be rejected, again using Probit models.⁸ Throughout, I cluster standard errors at the branch level to account for possible dependencies among borrowers in branches, i.e. branch specific shocks.

⁷I argue that measuring the Par 30 over the lifetime of the loan is preferable to measuring it over a certain time period. Even if longer-term loans have more chances to miss repayments in absolute terms, it seems reasonable to assume that borrowers default after a certain share of their loan's time to maturity has passed and not always after e.g. four months.

⁸I estimate all models with and without the actual or hypothetical credit score to control for changes in the underlying loan application quality but present only results without (results are similar). The disadvantage of using the credit score as a control variable is that the score is linked to the actual treatment of the system's recommendation (even if not one-to-one), and picks up some of the treatment's effect because a visibly high score may change the loan officer's behavior regardless of the system's recommendation.

A standard problem of clustered RCTs with a small amount of clusters (here: branches) is the fact that the design’s power rapidly decreases in case the interdependence between observations within single clusters becomes large, no matter how high the number of observations within the single clusters may be. By nature of the experiment, the number of clusters and observations within these clusters could not be increased. However, to estimate whether this poses a real problem for the experiment’s power, I use data from 2014 to estimate the intra-cluster correlation (ICC) with regard to the Portfolio-at-Risk over 90 days (using STATA’s `xtprobit` command with no additional controls). The result suggests that the branch characteristics account for only about 3% of an individual’s propensity to be in default as measured by the Par 90, i.e. the level of intra-cluster correlation is low in the year prior to the experiment. This suggests that my tests using 28 clusters and several thousand observations for each of them are high powered.⁹

3 Data and Summary Statistics

Figure 2 shows the applications at the bank branches during the pre-treatment, treatment and post-treatment period (all following graphs sort observations into the respective month of application). In total, about 27,000 borrowers applied during the four months long treatment period and about 25,000 in the post-treatment period (that is about 1,000 borrowers on average per branch over the four months long treatment period). The ratio between the control and treatment branches remains stable over time. There is no indication of applicants deliberately shifting to either control or treatment branches. However, there is somewhat of a downward trend in the number of applications over time.

Table 1 presents summary statistics for individuals who applied to control and treatment branches and for eventual borrowers, by pre-treatment, treatment and post-treatment period. The borrowers’ age, gender and the requested loan amount (the first three rows) are application data. Borrowers in the sample are about 40 years old and more likely to be male, eventhough there are more female borrowers in the control branches. The borrowers run very small businesses and request loan amounts of about 760 USD (control branches) and 850 USD (treatment branches) during the treatment period.¹⁰

⁹I refrain from further ex-post power calculations as strongly suggested by e.g. Hoenig and Heisey (2001).

¹⁰For the sake of anonymity, all loan amounts in this study were translated into US Dollars. The US Dollar amounts were calculated using the same exchange rate from the local currency to USD for all loans from May 2015. I censor eleven applications in terms of the application’s loan volume at the top because a few unrealistically high loan amounts were recorded that skew the distribution.




Figure 2
Number of Applications - Treatment vs. Control

Table 1
Summary Statistics

	<u>Control Branches</u>			<u>Treatment Branches</u>		
Variable	<i>N</i>	Mean	Std. Dev	<i>N</i>	Mean	Std. Dev
Pre-Treatment	69,269 Applications			72,445 Applications		
Loan Applications / all						
Age	65,543	40.25	11.06	67,937	40.13	11.10
Gender*** (0=M, 1=F)	68,270	0.52	0.50	71,354	0.45	0.50
Amount Requested	69,266	684.51	3,322.76	72,438	668.51	2,892.74
Loan Applications / with existing information on the credit bureau						
Previous Nb. Loans***	56,548	6.07	5.07	58,233	6.28	5.56
Current Balance Outst.*	56,548	428.27	2,960.68	58,233	471.68	2,906.84
Continues on next page						

Continues on next page

Variable	<i>N</i>	Mean	Std. Dev	<i>N</i>	Mean	Std. Dev
Total Amount Ever**	56,548	2,930.06	9,761.52	58,233	3,081.24	9,560.70
Credit Score***	56,548	658.05	22.81	58,233	657.39	23.87
Loans Granted						
Annual Interest Rate***	65,517	64.60	9.54	68,347	64.74	9.36
Loan Term***	65,517	238.39	94.80	68,347	241.13	98.98
Loan Amount	65,517	553.93	1,336.82	68,347	552.88	1,388.67
PAR 30 (ever)***	65,515	16.62%	0.37	68,342	17.43%	0.38
Treatment	12,495 Applications			14,008 Applications		
Loan Applications / all						
Age	12,267	40.40	11.05	13,628	40.16	10.98
Gender*** (0=M, 1=F)	12,366	0.55	0.50	13,763	0.45	0.50
Amount Requested	12,494	762.34	4,958.33	14,007	856.78	4,758.11
Loan Applications / with existing information on the credit bureau						
Previous Nb. Loans**	10,324	6.63	5.51	11,436	6.42	5.90
Current Balance Outst.	10,324	401.07	2,327.64	11,436	463.35	2,650.21
Total Amount Ever	10,324	3,151.36	9,450.02	11,436	3,318.15	10,141.63
Credit Score*	10,324	659.00	23.19	11,436	659.65	23.92
Loans Granted						
Annual Interest Rate	11,891	64.24	9.65	12,536	64.10	9.72
Loan Term	11,891	277.02	85.57	12,536	276.50	88.29
Loan Amount	11,891	606.37	1,751.34	12,536	629.16	1,814.46
PAR 30 (ever)***	11,891	21.58%	0.41	12,536	24.83%	0.43
Post-Treatment	11,509 Applications			13,438 Applications		
Loan Applications / all						
Age	11,082	39.28	10.85	13,040	39.41	10.79
Gender*** (0=M, 1=F)	11,266	0.51	0.50	13,160	0.42	0.49
Amount Requested	11,509	1,016.99	5,685.77	13,434	1,002.04	4,203.57
Loan Applications / with existing information on the credit bureau						
Previous Nb. Loans	8,319	6.17	5.32	9,281	6.04	5.69
Current Balance Outst.	8,319	610.68	4,586.00	9,281	650.23	4,027.45
Total Amount Ever	8,319	3,908.09	12,310.25	9,281	4,154.68	12,882.68
Credit Score***	8,319	658.17	25.28	9,281	658.05	25.14

Continues on next page


Continues on next page

Variable	<i>N</i>	Mean	Std. Dev	<i>N</i>	Mean	Std. Dev
Loans Granted						
Annual Interest Rate	10,118	63.10	10.73	11,660	63.12	10.63
Loan Term	10,118	352.19	84.02	11,660	353.57	82.25
Loan Amount	10,118	757.44	1,784.62	11,660	761.11	1,865.27
PAR 30 (ever)	10,118	29.22%	0.46	11,660	36.05%	0.48

All variables referring to loan amounts are in USD. Previous Nb. Loans is the number of loans ever registered on the credit bureau for the applicant. Current Balance Outst. is the sum of all of the applicant's current open accounts registered on the credit bureau. Total Amount Ever is the total amount ever borrowed and recorded on the credit bureau. PAR 30 (ever) is a dummy variable indicating whether the borrower was ever more than 30 days late on any repayment. The number of observations for age and gender do not amount to the number of applications because sometimes the information is missing. Asterisks indicate results of *t*-tests for an equal mean of the two groups assuming unequal variances. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Rows four to seven present data for applicants who have existing information in the credit bureau, i.e. are not first-time borrowers. It is common practice in the market to take up credit repeatedly; borrowers have between six to seven previous loans recorded in the credit bureau at the time of their application. Hence, for many borrowers there is a lot of information that can be fed into a credit scoring model to forecast their performance. The average credit score is about 660 over the whole period. The data on the final borrowers (rows eight to ten) reveals that the average loan term is much shorter than a year, which partly explains the interest rates of more than 60% annually. The summary statistics reveal a trend towards larger loans over time. The average loan amount increases from about 550 USD to about 760 USD from the pre- to the post-treatment period. There is also an upward trend in the share of borrowers who are more than 30 days late (Par 30) on any repayment over the lifetime of their loan.


T-tests for the balance between the treatment and control branches reveal that, in fact, borrowers at the treatment branches are a bit more likely to default in the pre-treatment period despite using bank internal data for stratification on the Par 90 prior to the experiment. In terms of quality of the applicant pool,

**Figure 3**

(Hypothetical) System Score assigned to applications

however, during the treatment period, there are no discernible differences in the average credit score of treatment and control branches. Figure 3 underscores the successful randomization: The relative scores assigned to the applications remained stable between treatment and control branches.


Figures 4 and 5 show the system's recommendations based on the score and the other rules incorporated into the scoring model (the decisions for the pre-experiment period and the control branches are hypothetical). The decisions remain relatively stable over time. The system recommends the rejection of about 15% of applications and to accept 65% of the applications right away. About 20% of the applications are recommended for a thorough, manual screening procedure. In the post-treatment period, there was an influx of first-time applicants with no information stored in the credit bureau, which is why there is a surge of refer-decisions.

**Figure 4**

System Recommendation Treatment (Yes, Referral, No)

Table 2 offers insight into the behavior of bank staff in response to the new system and loan officers' adherence to the system's decision (hypothetically for the control branches), i.e. incidences of overrides. Remarkably, prior to the introduction of credit scoring, many applications that the system would have recommended to decline were in fact being approved by the bank. During the pre-treatment period, only about 13% of applications the system recommended to reject were actually rejected, both in the control and treatment branches. On the other hand, the bank accepted almost all applications the system recommended to accept during the pre-treatment period (about 97%).

When the experiment started, the share of applications that were recommended by the system for a decline but were eventually approved remained stable at the control branches, indicating no significant change in behavior by personnel at the control branches. At the treatment branches, however,

**Figure 5**

Hypothetical System Recommendation Control (Yes, Referral, No)

the number of rejected applications that had also been recommended to be declined increased significantly during the treatment period (to 37%), though the share of seemingly non-compliant decisions remains high.¹¹

In the post-treatment period, compliance with the system in the control branches jumps up to the treatment level of the treatment branches when they start using the credit scoring system (37% for rejected applications). The compliance in the treatment branches increases somewhat further after the treatment phase (46% for rejected applications). Interestingly, when branches receive the treatment they start to be more conservative in terms of risk-taking even for applications that the system recommended and start to reject more of those

¹¹Many overrides were related to applications declined by the system due to falsely recorded outstanding amounts at the bank. In most of these cases the loan officer was able to verify that the old loan had in fact been paid back fully and then overrode the decision.

Table 2
Compliance

Bank Decision / System Decision	Control Branches			Treatment Branches		
	Granted	Rejected	Total	Granted	Rejected	Total
Pre-Treatment						
Reject	8,579 (88.03%)	1,167 (11.97%)	9,746	8,740 (87.03%)	1,302 (12.97%)	10,042
Accept	42,913 (97.07%)	1,295 (2.93%)	44,208	43,688 (96.96%)	1,372 (3.04%)	45,060
Total	51,492	2,462	53,954	52,428	2,674	55,102
Treatment						
Reject	1,533 (88.72%)	195 (11.28%)	1,728	1,096 (63.17%)	639 (36.83%)	1,735
Accept	7,940 (97.39%)	213 (2.61%)	8,153	8,594 (94.63%)	488 (5.37%)	9,082
Total	9,473	408	9,881	9,690	1,127	10,817
Post-Treatment						
Reject	948 (62.57%)	567 (37.43%)	1,515	872 (54.23%)	736 (45.77%)	1,608
Accept	5,850 (92.02%)	507 (7.98%)	6,357	6,439 (90.56%)	671 (9.43%)	7,110
Total	6,798	1,074	7,872	7,311	1,407	8,718

Rows indicate the system's recommendation and column's indicate the bank's decision.

too. Altogether, this shows that the credit scoring system was set up to, and actually led to a somewhat more conservative loan policy by the bank.¹²

¹²Behavioral aspects also play a role. Interviews on the ground revealed that some of the non-compliant decisions during the experiment were caused by a lack of trust on the part of the loan officers in the decision the computer had taken. The fact that loan officers are being paid based on both volume and loan performance contributed to this behavior. As the system urged loan officers to reject more loans, concerns about their own portfolio's volume led them to grant more loans than the system was recommending to reject, even if their overrides had to be signed off by the management.

4 Results

4.1 The Effects of Credit Scoring

Figure 6 visualizes the share of borrowers who are ever more than 30 days late on repayment (Par 30 ever) over the lifetime of their loans, sorted by the respective months the loan was taken up in during the pre-treatment, treatment and post-treatment period.




Figure 6

Share of Defaults - Treatment vs. Control (sorted by month of application)

By visual inspection, the introduction of credit scoring did not increase the likelihood of a borrower to default. Only during the post-treatment phase, when all branches started using credit scoring, did the treatment branches grant loans with a lower credit quality. That means, regardless of the experiment, there is noticeable deterioration of the credit quality towards the end of the post-treatment period. This decline, however, is unrelated to the introduc-

tion of credit scoring. Firstly, the credit quality increased again significantly for loans granted after January, when the bank was still using the same credit scoring system. Secondly, Figure A4.1 in the appendix substantiates the claim that the deterioration of quality was unrelated to the introduction of credit scoring but rather the result of the bank trying to shift to more long-term loans in their portfolio towards the end of the year.




Figure 7
Share of Declined Applications

While credit scoring did not visibly affect the credit quality of the bank's portfolio, the introduction of the new system led to an increase in the likelihood of an application to be declined (Figure 7). During the first month of the treatment phase, the share of rejected applications at the treatment branches jumps up. This jump was clearly related to the new system: Just as the bank rolled out the system to all branches in September, the control branches rapidly caught up in rejecting more applications during the post-treatment

period. The share of declined applications in the control branches increased by almost 8% between August and September.

Consequently, the number of granted loans in the treatment branches decreased after the start of the treatment period. While, during the pre-treatment phase, the number of loans in the treatment branches steadily hovered above the control branches, the number of granted loans is almost similar in the months of July and August (see Figure 8). As soon as the control branches started to use credit scoring themselves, total number of loans in the treatment branches again surpassed the control branches.




Figure 8

Number of Loans Granted - Treatment vs. Control

Table 3 presents results for the main regression at the individual level that uses a dummy for ever being more than 30 days late as the outcome variable.¹³ I

¹³The results are robust to variations in the outcome variable. Using the loans' repayment status in the very last month of their time to maturity does not change the results of the default status regressions.

include proxies for competition at the branch level (neighboring bank branches) and managerial control of the branch (distance to the headquarter). There are no significant differences in loan quality between the treatment and control branches for loans granted during the pre-treatment and the treatment period (results are marginal effects). Given that, during the pre-treatment phase, loans at treatment branches were also more likely to be late on repayments (see Table 1), deducting that pre-treatment trend would only reinforce the result that credit scoring did not increase the default rate. For loans granted during the post-treatment period, when all branches used credit scoring, the treatment branches feature a higher default rate but the overall setup of the experiment suggests that this difference is part of a trend rather than related to credit scoring.

Table 3

Default Probability (Par 30 ever) - All Granted Loans

	(1) Pre-Treatment	(2) Treatment	(3) Post-Treatment
Treatment Branch	0.004 (0.21)	0.032 (1.06)	0.064* (1.97)
Log(Competing Branches (5km))	0.019 (1.88)	0.025 (1.49)	0.038 (1.91)
Log(Distance to HQ)	0.022** (2.74)	0.040*** (3.30)	0.046** (2.93)
<i>N</i>	109,602	22,240	20,263
Mean Dependent Variable	0.199	0.249	0.339

Estimates are the results of Probit regressions. Results are marginal effects. The model was estimated including a constant. *t*-statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Next, I present results for the regressions checking for systematic differences in terms of the characteristics of loans granted at treatment and control branches during the treatment period (Table 4). None of the three loan features, the interest rate, the loan term and the loan volume are significantly affected by the fact that the bank started to use credit scoring. This is in line with the intended use of the credit scoring system. The system was designed to help branch personnel decide about rejecting or accepting a loan. At the same

time, the loan characteristics were supposed to be determined the way they had been determined before, either as the result of a pre-determined growth path of a borrower's loan volume or in negotiation between borrower, loan officer and/or branch manager. The result precludes concerns that results on the loan portfolio quality and the likelihood of rejecting an application could be driven by systematic changes in the types of loans the bank branches were granting.

Table 4
Other Loan Options

	(1) Interest Rate	(2) Loan Term	(3) Amount (USD)
Treatment Branch	0.335 (0.50)	-3.450 (-0.55)	-23.107 (-0.42)
Log(Competing Branches (5km))	0.563 (0.69)	-7.878 (-1.57)	-40.775 (-0.59)
Log(Distance to HQ)	0.096 (0.19)	-6.297 (-1.93)	-45.952 (-1.10)
Constant	62.038*** (13.44)	327.527*** (11.56)	920.680* (2.37)
<i>N</i>	23,627	23,627	23,627
Mean Dependent Variable	64.24	276.00	593.86

Estimates are the results of OLS regressions. *t*-statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Finally, the regression results for the probability of a loan to be rejected in Table 5 confirm the visual impression from before: The credit scoring system led to a more conservative screening process in comparison to the soft information based technology. During the treatment phase, the likelihood of an application to be rejected is about 6 percentage points higher in a treatment branch. The roll out of the system to all branches helps lend credibility to a truly causal interpretation of the finding: After the control branches start using credit scoring, the difference between the control and treatment branches disappears.

Table 5
Rejection Probability - All Granted Loans

	(1) Pre-Treatment	(2) Treatment	(3) Post-Treatment
Treatment Branch	0.002 (0.58)	0.059*** (5.12)	0.013 (1.26)
Log(Competing Branches (5km))	0.005 (1.58)	0.017** (3.10)	0.009 (1.33)
Log(Distance to HQ)	0.002 (1.04)	0.014*** (3.41)	0.005 (1.06)
<i>N</i>	136,967	25,615	24,029
Mean Dependent Variable	0.055	0.078	0.126

Estimates are the results of Probit regressions. Results are marginal effects. The model was estimated including a constant. *t*-statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.2 Type II Errors

An increase in the rejection rate combined with a stable portfolio quality implies an increase in the total number of rejected applications that would not have defaulted (type II errors). Rejected applications that would not have defaulted have negative economic effects because it is reasonable to assume that those loans would have financed a worthy investment. It is, however, difficult to assess the extent of type II errors because the outcome of loan applications that the bank rejected is obviously unknown. Loans that were recommended for a decline by the system, but nevertheless accepted by the bank are also a poor proxy for the unknown default rate of declined applications. After all, some of these applications had been rejected by the system mistakenly (see Chapter 3) and the loan officer corrected the decision which is why they are systematically different from applications that were actually declined by the bank.

Another way of thinking about the extent of type II errors is to compare the default rate between loan applications in treatment and control branches that had been recommended for a decline by the system but were granted by the bank (see Table 6). In the treatment branches, more of the applications recom-

Table 6

Par 30 (ever) - Recommended for Decline

System Decision - Bank Decision	Control Branches		Treatment Branches	
	Number	Par 30 (ever)	Number	Par 30 (ever)
Decline - Granted	1,533	27.27%	1,096	27.10%
Decline - Rejected	195	n.a.	639	n.a.
Total	1,728		1,735	

mended for a decline are actually declined by the bank. If the bank is well able to sort out the most risky applications, the default rate among the remaining applications recommended for a decline - but granted - should decrease. This is, however, not the case: Among the applications that are recommended for a decline but granted by the bank, there is no discernible difference in terms of loan quality between treatment and control branches (27.10% vs. 27.27%). Under a stable quality of the pool of applications (see Figure 3), the credit quality among the remaining applications that were actually declined by the bank must be similar between treatment and control branches as well. It follows that the bank makes a considerable number of type II errors when starting to reject more applications by using the credit scoring system (even if many of those loans default).

It is instructive to try and understand which applicants are most likely to be hit by the increasing rejection rate at the bank (other than those targeted by the credit scoring system per se, like overly indebted borrowers). Therefore, Figures A4.2 and A4.3 in the appendix show the probability of a loan application to be rejected plotted against the loan volume applied for (sorted into equally large deciles). The figures show that the increasing rejection rate at the treatment branches hits applications of all volumes similarly. The credit scoring approach does not seem to discriminate between borrowers of small and large loans in particular.¹⁴

¹⁴The bank rejects many more very large applications that do not fit into the bank's target group anyway.

4.3 Full Compliance

Naturally, one is also inclined to ask which result the bank would have achieved under perfect compliance to the system's recommendation, to see what the system would really be capable of doing. This question is, however, difficult to answer for at least two reasons. Firstly, some non-compliant decisions by the branch staff were triggered by loans that were falsely recorded as still outstanding (see Chapter 3). Those overrides - that I do not observe - must not be understood as decisions against, but as part of the system, since the possibility to correct such recommendations was intentionally built into the credit scoring model. Secondly, in the treatment period about 5% of all applications the system recommended to accept were declined by the bank. The outcome for those loans is unobserved which is why I cannot calculate what the outcome would have been, had the loan officers complied with the recommendation. To get an idea of whether the bank could have done better by increasing compliance with the credit scoring system, I must rely on loans the bank actually granted. Table 7 compares loans granted by the bank that were either recommended for a decline by the system or recommended to be accepted (i.e. complying vs. non-complying loans) with regard to their likelihood to ever be more than 30 days late on any repayment. Results show that - even if the bank's staff had good reasons to override the system in each and every case - loans that were accepted against the system's recommendation are significantly more likely to default. The difference vanishes in the post-treatment period when there are less of such overrides in total. The bank also started to grant larger loans during the post-treatment period, which hints at the fact that the credit scoring model may have been less able to gauge the credit quality among those applicants of larger loans than among the initial target group of borrowers of very small loans it had been built for.

This is no proof, however, that full compliance with the credit scoring system would have been better for the bank. After all, fully complying with the system's recommendation would have also meant to reject more loan applications. For example, in the treatment branches during the treatment period,

Table 7

Default Probability - Complier vs. Non-Complier

	(1) Pre-Treatment	(2) Treatment	(3) Post-Treatment
Non-Complier Dummy	0.069*** (9.18)	0.054* (2.38)	0.018 (0.70)
Log(Distance to HQ)	0.027*** (3.90)	0.051*** (3.30)	0.051* (2.44)
Log(Competing Branches (5km))	0.016 (1.86)	0.031 (1.53)	0.028 (1.29)
<i>N</i>	42,155	8,644	6,587
Mean Dependent Variable	0.197	0.264	0.345

Estimates are the results of Probit regressions. Results are marginal effects. The model was estimated including a constant. *t*-statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

the number of applications that would have additionally been rejected by the scoring system under full compliance is 637.

4.4 Costs and Benefits of Credit Scoring

The credit scoring system trades off between type I errors (accepted loan applications that default) and type II errors (rejected loan applications that would not have defaulted). To judge whether the gains from using credit scoring surpass the losses, it is instructive to weigh the costs and benefits in terms of these errors, taking into account changes in operational costs.

Here, I compare profits without credit scoring with profits after the introduction of credit scoring for each dollar invested into the bank's loan portfolio.¹⁵ Profits from granting loans without credit scoring are the loan volume LV multiplied by the economic profit π (the economic profit takes into account the opportunity costs of lending and is the surplus profit from granting loans over alternative investments). After the introduction of credit scoring at the treatment branches, economic profits accrue only on the part of the loans that

¹⁵The analysis is an adaption of the cost-benefit analysis in Berg (2015). To see how the analysis translates into a calculation for the total loan portfolio by the bank under the assumption of uniform loans, the formula can be multiplied by the number of loans.

have not been rejected by the credit scoring system $LV * (1 - \Delta Rr)$, with ΔRr being the increase in the rejection rate, and the rest of the money is assumed to be invested at opportunity costs. Profits are decreased (increased) by a rise (fall) in the default rate weighted by the loss given default ($\Delta Def * LGD$).

The bank pays a certain amount of money for each credit score provided by the operator of the credit information sharing system. The operational costs of obtaining the credit score accrue on every application ($ScoreCosts$). At the same time, credit scoring decreases operational costs by simplifying the screening process (see Chapter 2.3). I assume that the change in operational costs for screening borrowers $\Delta ScreenCosts$ accrues on every dollar invested in the loan portfolio, i.e. screening costs are increasing in the loan volume. Finally, credit scoring is profitable for the bank if:

$$LV * \pi < \\ LV * (1 - \Delta Rr) * (\pi - \Delta Def * LGD) - ScoreCosts - LV * \Delta ScreenCosts$$

This comparison can be expressed in terms of costs from 1) foregone profit on loans due to an increase in the rejection rate and 2) an increase in operational costs from paying for the credit score; and benefits from 3) a decrease in the default rate and 4) a decrease in operational costs via a simplification of the screening process.

$$LV * \pi * \Delta Rr + ScoreCosts < \\ -LV * \Delta Def * LGD * (1 - \Delta Rr) - LV * \Delta ScreenCosts$$

Based on the empirical findings, I calibrate the increase in the rejection rate at 6%. I also employ the finding that credit scoring did not change the default rate ($\Delta Def = 0$), which is why an estimate for the loss given default is not needed. The fixed costs of obtaining the credit score for each application from the external provider $ScoreCosts$ is about 2 USD. To distribute the costs among the number of loans actually granted, I multiply the costs for the score by the ratio of the number of applications to the number of loans granted by the bank. I use data on all branches from the post-treatment period during

which the bank rejects 12.7% of all applications. Hence, the costs for the score broken down to each loan granted amount to 2.29 USD. Furthermore, I set the bank's economic profit on a given dollar invested into the loan portfolio to 4%, 7% and 10%. In lack of a concrete measure of exactly how much credit scoring reduced operational costs at the bank, I need to rely on estimates. Rosenberg et al. (2013) show that average total operational costs for small business lenders and microfinance banks in Africa are about 19% of the loan volume.


Here, I build into the analysis the fact that costs of screening borrowers per USD of loan volume are generally larger for smaller loans. This is because some operational costs related to screening borrowers that are reduced by the credit scoring system are not variable, e.g. visiting a borrower of a large loan will cost almost the same as it does for a borrower of a very small loan. This is also why borrowers of small loans suffer most from high interest rates triggered by operational costs (Banerjee and Duflo 2010). Based on estimates in Research Insight (2014) and in line with the estimates in Rosenberg et al. (2013), I assume linearly declining operational costs based on costs of 30% of the loan volume for the smallest loans of 50 USD, and 15% of the loan volume for a loan of 2,000 USD.¹⁶ Finally, I assume that operational costs decrease by a highly conservative 10% due to the use of credit scoring. I rely on this very moderate estimate for the cost savings by credit scoring because an analysis of the time between application and loan approval at the bank in the experiment suggests that for the current introduction of credit scoring, most efficiency gains arose on the immediate rejection of applications rather than the immediate approval. For the loans that the system recommended to accept, in many cases, the loan officers would not process them faster than before, especially when the loan officers doubted the system's decision.

¹⁶This is still a conservative estimate of operational costs. On average, according to Rosenberg et al. (2013), operational costs constitute more than half of the total costs related to a loan which would likely be much more than 19% for this bank where the average interest rate on a loan is more than 60%. For the worldwide portfolio of institutions, the average interest yield in 2011 was 26.9%; broken down into 2.6% profits, 3.6% loan losses, 7.8% financial expenses and 14% operational expenses, all in terms of loan volume (Rosenberg et al. 2013).

Results in Figure 9 show that credit scoring is profitable even for the smallest loans at the bank. The bank reaches the break even point of using credit scoring already for a loan volume of just below 100 USD. However, an ever smaller amount of operational costs relative to the loan size means that the benefits from reducing screening costs diminish as loans become larger (the benefits are hump-shaped). Meanwhile, the total foregone profit from larger rejected applications increases. The bank would lose a lot of money on high-volume applications that the credit scoring system is more likely to reject. In this setting, credit scoring is most lucrative (in absolute, not relative terms) for loans of about 1,800 USD. Eventually, the net result from using credit scoring even becomes negative for large loans. For such high-volume borrowers, the bank is better off screening each loan application extensively and thereby being able to accept more of them, as the stakes of declining high-volume applications are too high. The bank in this study, however, has only very few of such high-volume applications. Less than 5% of all application are larger than 3,000 USD. That is why, weighing the costs and benefits, the credit scoring system clearly increases the bank's profits. Given a relatively small fee for the provision of the score, the similar performance of eventual borrowers at the cost of an increase in the rejection rate of 6% suggests that even for very conservative estimates of operational costs reductions, credit scoring pays off for the bank.

5 Conclusion

This study demonstrates that a recent surge in the availability of data from credit bureaus facilitates more automated and hard information-based credit technologies. I present results from a first-of-its-kind randomized controlled trial that was carried out by a small business lender in an African financial market. In the wake of the experiment, the lender introduced a credit scoring lending technology at a randomly selected half of its 28 branches for a four months long treatment period, while the other branches kept on applying an extensive screening technology based on loan officers' expert judgment. The

**Figure 9**

Cost minus Benefits of Credit Scoring - Non-Linear Operational Costs

fact that the bank rolled out the new credit scoring system to all branches after the initial treatment period, supports the causal interpretation of the findings.

The results show that credit scoring produced screening results for loan applications of similar precision to the extensive screening technology based on loan officers' expert judgment. I do not observe a significant increase in the likelihood of a loan being late on repayment when it was granted using the credit scoring based lending technology. There is also no evidence that treatment branches started to grant different types of loans that would make treatment and control loans incomparable. Meanwhile, the application of credit scoring rendered the lender's credit technology somewhat more conservative and an individual application's probability of being rejected increased by 6 percentage points. Yet, an analysis of the costs and benefits strongly suggests that even

for conservative assumptions about the cost savings from using credit scoring, the bank's net benefit from introducing the system are positive.

Being an increasingly efficient way of screening applications credit scoring may commence a shift away from extensive and more expensive lending technologies that have dominated small business lending in developing financial markets. As the availability and quality of credit bureau data steadily increases, credit scoring in small business lending may prove to be an effective way to widely decrease operational costs. From a developmental point of view, the results suggest new opportunities to improve small business borrowers' access to credit, despite the fact that the introduction of credit scoring led to an increase in the share of loan applications rejected by the bank. While in the short-run there is less credit to borrowers, credit scoring frees up some of the bank's resources because the scoring model can do the task of screening applications more efficiently than the extensive screening procedure can. In the long-run, the bank may use such resources for the purpose of prospecting new customers. In a competitive banking market, eventually the usage of a credit scoring must lead to more credit as the marginal application turns profitable to the bank. Therefore, the results in this study are not only good news for the bank. Borrowers of small loans who face the most severe credit constraints due to high operational costs could ultimately be the winners of such a move towards more automated small business lending. Meanwhile, the role of policymakers will be to ensure that credit scoring systems do not discriminate against certain borrower groups in an unduly manner, denying access to the credit market to some.

6 Appendix






Figure A4.1

Correlation - Par 30 (ever) and Loan Term (sorted by month of application)

**Figure A4.2**

Declined Applications by Application Volume - Treatment

**Figure A4.3**

Declined Applications by Application Volume - Pre-Treatment

References

- Allen N. Berger, W. Scott Frame and Nathan H. Miller (2005). “Credit Scoring and the Availability, Price, and Risk of Small Business Credit”. In: *Journal of Money, Credit and Banking* 37.2, pp. 191–222.
- Banerjee, Abhijit V. and Esther Duflo (2010). “Giving Credit Where It Is Due”. In: *Journal of Economic Perspectives* 24.3, pp. 61–80.
- Berg, Tobias (2015). “Playing the Devil’s Advocate: The Causal Effect of Risk Management on Loan Quality”. In: *The Review of Financial Studies* 28.12, pp. 3367–3406.
- Berg, Tobias, Manju Puri, and Jörg Rocholl (2016). “Loan Officer Incentives, Internal Ratings, and Default Rates”. In: *AFA 2013 San Diego Meetings Paper*.
- Bos, Jaap W.B., Ralph De Haas, and Matteo Millone (2015). “Show Me Yours and I’ll Show You Mine: Sharing Borrower Information in a Competitive Credit Market”. In: *CentER Discussion Paper Series* No. 2015-027.
- Bruhn, Miriam and David McKenzie (2009). “In Pursuit of Balance: Randomization in Practice in Development Field Experiments”. In: *American Economic Journal: Applied Economics* 1.4, pp. 200–232.
- DeYoung, Robert, Dennis Glennon, and Peter Nigro (2008). “Borrower-lender distance, credit scoring, and loan performance: Evidence from informational-opaque small business borrowers”. In: *Journal of Financial Intermediation* 17.1, pp. 113–143.
- Dierkes, Maik et al. (2013). “Business credit information sharing and default risk of private firms”. In: *Journal of Banking & Finance* 37.8, pp. 2867–2878.
- Djankov, Simeon, Caralee McLiesh, and Andrei Shleifer (2007). “Private credit in 129 countries”. In: *Journal of Financial Economics* 84.2, pp. 299–329.

- Doblas-Madrid, Antonio and Raoul Minetti (2013). “Sharing information in the credit market: Contract-level evidence from U.S. firms”. In: *Journal of Financial Economics* 109.1, pp. 198–223.
- Einav, Liran, Mark Jenkins, and Jonathan Levin (2013). “The impact of credit scoring on consumer lending”. In: *The RAND Journal of Economics* 44.2, pp. 249–274.
- Frame, W. Scott, Aruna Srinivasan, and Lynn Woosley (2001). “The Effect of Credit Scoring on Small-Business Lending”. In: *Journal of Money, Credit and Banking* 33.3, pp. 813–825.
- Frame, W. Scott, Lynn Woosley, and Michael Padhi (2004). “Credit Scoring and the Availability of Small Business Credit in Low- and Moderate-Income Areas”. In: *The Financial Review* 39.1, pp. 35–54.
- Gietzen, Thomas (2016). “Credit Information Sharing and Hold-up”. In: *University of St. Gallen, School of Finance Research Paper* No. 2016-12.
- Hertzberg, Andrew, Jose Maria Liberti, and Daniel Paravisini (2010). “Information and Incentives Inside the Firm: Evidence from Loan Officer Rotation”. In: *The Journal of Finance* 65.3, pp. 795–828.
- Hoenig, John M. and Dennis M. Heisey (2001). “The Abuse of Power: The Pervasive Fallacy of Power Calculations for Data Analysis”. In: *The American Statistician* 55.1, pp. 19–24.
- Jappelli, Tullio and Marco Pagano (2002). “Information sharing, lending and defaults: Cross-country evidence”. In: *Journal of Banking & Finance* 26.10, pp. 2017–2045.
- Kallberg, Jarl G. and Gregory F. Udell (2003). “The value of private sector business credit information sharing: The US case”. In: *Journal of Banking & Finance* 27.3, pp. 449–469.
- Liberti, Jose M., Amit Seru, and Vikrant Vig (2016). “Information, Credit, and Organization”. In: *Institute for Monetary and Financial Stability - Working Paper Series* No. 2015-97.

Mester, Loretta (1997). “What’s the point of credit scoring?” In: *Business Review*, pp. 3–16.

Qian, Jun, Phillip E. Strahan, and Zhishu Yang (2015). “The Impact of Incentives and Communication Costs on Information Production and Use: Evidence from Bank Lending”. In: *The Journal of Finance* 70.4, pp. 1457–1493.

Research Insight, responsAbility (2014). *Efficiency is the key to lower interest rates in microfinance*. responsAbility. URL: <http://www.responsability.com>.

Rosenberg, Richard et al. (2013). *Microcredit Interest Rates and Their Determinants: 2004 - 2011*. CGAP. URL: <http://www.cgap.org/publications/microcredit-interest-rates-and-their-determinants>.