NBER WORKING PAPER SERIES

MODERNIZING ACCESS TO CREDIT FOR YOUNGER ENTREPRENEURS: FROM FICO TO CASH FLOW

Christopher M. Hair
Sabrina T. Howell
Mark J. Johnson
Siena Matsumoto

Working Paper 33367
http://www.nber.org/papers/w33367

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2025

This project was partially conducted on data gathered for a project using Federal funds under award MB22OBD8020266- T1-01 to FinRegLab from the Minority Business Development Agency, U.S. Department of Commerce. The statements, findings, conclusions, and recommendations are those of the author(s) and do not necessarily reflect the views of the Minority Business Development Agency or the U.S. Department of Commerce. Separate analysis under the grant is continuing and will be forthcoming in a separate report. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

Modernizing Access to Credit for Younger Entrepreneurs: From FICO to Cash Flow
Christopher M. Hair, Sabrina T. Howell, Mark J. Johnson, and Siena Matsumoto
NBER Working Paper No. 33367
January 2025
JEL No. G21, G23, G32, J14, J16

ABSTRACT

Younger entrepreneurs are disadvantaged by traditional loan underwriting, which relies heavily on personal credit scores. With data from three fintech companies, we show that incorporating timely information about ability to repay from business checking account statements particularly improves default prediction performance for younger business owners. We develop a novel method to compare model predictions across subgroups—Tail Analysis for Comparative Outcomes (TACO)—which finds that switching from a Baseline (FICO-driven) model to a Cash Flow-enhanced model benefits younger entrepreneurs. We confirm this in causal analysis of approval decisions, showing that access to cash flow-intensive underwriting increases approval rates for younger vs. older entrepreneurs.

Christopher M. Hair
Brigham Young University
chris_hair@byu.edu

Sabrina T. Howell
NYU Stern School of Business
KMC 9-93
44 West 4th Street
New York, NY 10012
and NBER
showell@stern.nyu.edu

Mark J. Johnson
Brigham Young University
Department of Finance
Provo, UT 84602
markjjohnson@byu.edu

Siena Matsumoto
NYU Stern School of Business
sm11914@nyu.edu

1 Introduction

Arm's length small business underwriting in the U.S. has historically relied heavily on personal credit scores. In the early 2010s, fintech lenders began incorporating data from recent bank statements, an approach referred to as "cash flow-based underwriting." Cash flows shed light on the current state of repayment ability. The key economic distinction between cash flow data and credit scores are that the former captures recency and repayment capacity, while the latter captures history and delinquency. This paper shows that since credit scores mechanically disadvantage younger entrepreneurs, incorporating cash flows benefits them by identifying the subset who are more creditworthy yet have lower credit scores.

There is little work on how financial constraints vary across the lifecycle for business owners. However, personal credit scores—which are crucial for many business loans—are strongly correlated with age, far more so than with other protected class characteristics, such as gender or race (see Figures 1 and A.1). Credit scores tend to favor older people because they reward a long history of successful debt repayment, placing more weight on history length than on any single successful repayment. Furthermore, fair lending laws focus on protecting elderly applicants and prohibit the direct use of age in lending models (subject to certain caveats). The structural disadvantage from credit scores combined with the regulatory infrastructure preventing arbitrage could inadvertently create systematic bias against younger people. These dynamics may contribute to explaining why entrepreneurs are most commonly in their 40s at the time of firm founding (Azoulay et al., 2020).

By offering information that is orthogonal to personal credit scores (e.g., FICO), cash flow data drawn from recent bank statements could mitigate the mechanical disadvantage for younger entrepreneurs. The inclusion of these new data in FICO-based underwriting models has occurred amid a broader evolution of small business underwriting. The industry shifted in the 1990s away from a model of local loan officers relying on soft information and heuristics. In its place, arm's length lending relying on personal credit scores became dominant.[1] While banks have used formal financial statements to underwrite—as in Fama (1985)—they have not generally incorporated bank account data into risk scoring models.[2] The former leader of the U.S. Small Business Administration noted: "Surprisingly, large banks such as J.P. Morgan and Wells Fargo did not historically use bank account transactions in their underwriting and in many cases did not

[1] See Black and Strahan (2002); Petersen and Rajan (2002); Carter and McNulty (2005); Berger et al. (2005b). From the industry press, see Goldstein and Lauckner (2022).

[2] See Section 2.1 for discussion.

even have easy access to the information due to the siloed nature of the bank information systems" (Mills, 2024). Following the Financial Crisis, higher capital requirements led banks to retreat from small business lending (Chen et al., 2017). Fintech lenders entered to fill the gap, exploiting new digitization and automation technologies to incorporate bank statement-based cash flow information into credit scoring models for both personal and small business lending (FinRegLab, 2019). The idea was to supplement credit scores with timely and direct measures of repayment ability. Although initially targeted to riskier market segments underserved by banks, cash flow-based underwriting is making its way into conventional lending practices.[3]

We focus on small business lending for several reasons. First, lending to small businesses is especially challenging because of information frictions and heterogeneity (Getter, 2019). This historically required substantial firm-specific attention from an individual loan officer, making it difficult to lend profitably at arm's length using hard information. Yet small firms are crucial to the economy: they employ about half of all private sector workers and make up 99% of all firms (SBA, 2023). Furthermore, small businesses rely primarily on debt, in contrast to publicly traded firms or venture capital-backed startups. Finally, the fintech lenders which introduced cash flow-based lending often targeted the small business market.

This paper examines how cash flow-based underwriting affects access to credit for younger versus older entrepreneurs. Our analysis proceeds in three parts. First, to affect access to credit, new data must help lenders assess risk. Since lenders and regulators build models using default as an outcome, we show that cash flow variables predict loan performance, and have stronger predictive power for younger owners. Second, we introduce a novel framework to evaluate the differential impacts of alternative models on subpopulations, and use it to document benefits of cash flow-based models for younger entrepreneurs. Third, we use loan approval decisions in two causal identification strategies to show a positive impact of assignment to cash flow-based underwriting on approval for younger versus older applicants.

We make use of data on applications, originations, and loan performance from three fintech companies serving small businesses. Two are lenders and one is a platform.[4] We obtain information that the lenders use in underwriting, including the credit score, bank statement-based variables, and industry, as well as information they collect but do not employ directly in credit scoring, such as the principal owner's first name, age, and zip code. For the two lenders, we observe loan performance. The platform gathers a common set of application materials, including

[3]See e.g. Crosman (2024) and Lawler (2024). Also, in 2024 the big-three credit bureaus introduced consumer scores that integrate bank statement attributes (Experian, 2024; VantageScore, 2024).

[4]These companies have requested anonymity.

FICO and cash flow variables, and forwards them to lenders. There are a total of 111 lenders who make loans via the platform. In the platform's data, we observe lender offers and whether an applicant took up an offer, but not loan performance. All together, we observe just over 900,000 applications and 36,000 loans on the platform, and 163,000 applications and 38,000 loans from the two lenders. The data are comprehensive from their respective sources and span from November 2013 to May 2024.[5]

We first examine how variables drawn directly from bank statements predict default, after controlling for the traditional inputs of FICO, industry, firm age, firm size, time, and state. We employ variables that represent a common baseline among fintech lenders, including revenues, withdrawals, balances, standard deviations of balances and credits, and distress indicators such as low or negative ending balances and insufficient funds transactions. Any lender could use these data provided there is an open banking regime, like the *de facto* system in the U.S. or the regulated version in countries such as the UK.

Consistent with intuition, regression analysis shows that measures of business health and revenue—such as more inflows (i.e., credits)—are associated with lower chances of default, while measures of distress—such as overdrafts—predict higher chances. We use a machine learning model (a random forest classifier) to more flexibly capture inputs' predictive power. We compare a Baseline model with traditional inputs to a Cash Flow model, which adds a vector of bank statement variables to the Baseline model. The Cash Flow model predicts default significantly better across multiple measures, and is more predictive for younger owners. For example, the improvement in the ROC AUC, a standard performance measure, is 2.6 pp for owners under 40 years old and 0.08 pp for owners over 40, with both differences highly significant.[6]

Next, we propose a new method—"Tail Analysis for Comparative Outcomes," or TACO—to assess the benefits for sub-populations of using one model versus another. Essentially, the method directly studies the types of borrowers who are most affected by a model change. We devised the TACO methodology for our purpose, but it is easily generalized to any population characteristic or set of models, and can be used by researchers, regulators, and practitioners to determine who benefits and who loses from switching from one model to another. For example, it permits ex-post assessment of protected class implications without requiring that the model itself observe or be segmented on protected class status. It may be especially useful when comparing machine

[5]There are no Paycheck Protection Program loans or otherwise subsidized or mission-driven loans in the sample.

[6]Since cash flow data were used in underwriting the loans, these results may underestimate their predictive power relative to a setting where they were not used. However, the cash flow variables should have explanatory power because (a) we observe much more performance data than the lenders did at the time of origination; and (b) the lenders approve borrowers with positive risk, leaving room for residual predictive power.

learning models that are otherwise difficult to interpret.

We apply TACO to compare predicted default in the Cash Flow and Baseline models for various age groups. In brief, we first calculate the change in predicted default probability for each borrower under each model. We then identify the groups of borrowers in the tails: those with the largest decrease in predicted default under the Cash Flow model relative to the Baseline model (who most benefit) and those with the largest increase (who are most hurt). Finally, we analyze how the share of young borrowers differ across the helped and hurt groups. TACO can be summarized as identifying "diamonds in the rough" that surface when switching models. As the size of the tails increases—with a maximum of 50%—the analysis studies reallocation more broadly in the sample. Traditional average treatment effects do not capture nuanced distributional effects.[7]

We show that younger entrepreneurs are disproportionately represented among those helped by the Cash Flow model. We define the TACO ratio for young owners as the young share in the helped group divided by the young share in the hurt group. A ratio of one implies no net imbalance in the effects on young people from switching models, as the share of young borrowers in the helped and hurt groups is equal. The largest estimate is for the youngest entrepreneurs (less than 35 years old). Using the top and bottom deciles to define the tails, their TACO ratio is 1.60, implying that there are 60% more young people in the helped group than in the hurt group. The TACO ratio declines as we raise the threshold for "young."

Consistent with lower FICO scores disproportionately constraining young entrepreneurs, the ratio is much higher when we focus on young borrowers with lower FICO scores. For example, the TACO ratio for low-FICO young borrowers is 7.78 (again using the 35-year old threshold). In contrast, high-FICO young borrowers are overrepresented in the hurt group, with a TACO ratio of 0.66, reflecting the Baseline model's preference for high-FICO borrowers. However, the ratio is just 0.44 for the high-FICO group overall. Therefore, even among high-FICO applicants, the Cash Flow model is relatively better for the young. Notably, the relative benefit for low-FICO young borrowers (7.78) far outweighs the relative harm for high-FICO young borrowers (0.66), underscoring the Cash Flow model's potential to identify overlooked creditworthy borrowers among the low-FICO young.

This finding—and the paper's conclusions more generally—could reflect one of two mechanisms: Younger people might have better cash flow metrics, or they might be at a FICO

[7]For example, in our case it is possible that (a) most firms with young owners are higher risk, which would be confirmed by poorer average cash flow measures; but (b) a subset of firms with young owners are creditworthy, which can be observed using cash flow metrics but not FICO. Average effects would not reveal subgroup reallocation and 8 limited insight if the average young owner is unlikely to receive credit.

disadvantage conditional on cash flow. Our data are more consistent with the latter. As noted above, older applicants have much higher FICO scores, while within low-FICO groups, younger and older people have similar cash flow metrics. Overall, our new method identifies the subgroup of younger borrowers who have low FICO but reasonably strong cash flow metrics, and are thus more creditworthy.

Our analysis of default in machine learning models has direct implications for access to credit, since lenders develop risk scoring models using default prediction exercises. In the final part of the paper, we exploit our rich approval data in two methods that causally identify the effect of cash flow-based underwriting on access to credit. The primary method is a within-application analysis of Platform approval decisions. The Platform lenders vary in whether they use cash flow variables in approving loans. While applicants are not randomly assigned to lenders, we use application fixed effects to remove potential confounders from applicant quality or lender preferences. This is possible because applications are sent to multiple lenders. We find that assignment to a cash flow-intensive (CFI) lender increases the chance of approval for younger entrepreneurs. Among low-FICO applicants, those under 40 years old have a 2.3 percentage point (13.5% of the mean) higher chance of approval at CFI lenders, relative to non-CFI lenders and older entrepreneurs. The effect among high-FICO young owners is 1.1 pp, or 4.5% of the mean. Although fintech lenders tend to serve higher-risk segments, these results do not appear to reflect more risk-taking that is correlated with more intensive use of cash flow variables.[8]

As a robustness test, we confirm the result in a context with true random assignment. One of the lenders randomly assigns applications to loan officers, allowing us to again classify some as CFI and others as non-CFI. While the sample is small and there could be other characteristics of the loan officers that lead them to prefer young applicants—though they do not observe age and such a preference would be illegal—we can assess the within-officer effect of being randomly assigned a young applicant. We find that young entrepreneurs benefit from assignment to a CFI loan officer. For example, among low-FICO applicants, young people (under 403 years old) assigned to CFI loan officers are 8.7 pp more likely to be approved (18% of the mean) relative to older people. In this analysis, we see positive but noisier results for high-FICO applicants. Overall, while all young owners benefit from exposure to cash flow-based underwriting, there is a stronger effect for those

[8]The concern is that since younger people are generally riskier, CFI lenders may approve more younger people because they are less regulated and take more risk, not because they use cash flow-based underwriting. However, we find very strong results even among extremely high-FICO applicants, such as in subsamples with FICO scores above 750. Furthermore, we collect data on business survival for a random subset of about 46,000 applications (since we do not observe loan performance on the Platform). We show that among approved applications, being young and assigned to a CFI lender has if anything a positive relationship with survival, not the negative relationship we would expect if risk preferences explained the main finding.

who have weaker conventional metrics.

In sum, the key contribution of this paper is to show that younger entrepreneurs are disadvantaged by FICO-based underwriting but benefit from cash flow-based underwriting. We offer the first study of how age is related to financial constraints, joining a small group of studies on entrepreneurship over the life cycle (Zhang and Acs, 2018; Azoulay et al., 2020). Theoretical models of entrepreneurship predict that younger entrepreneurs are likely to face financial constraints (Evans and Jovanovic, 1989; Stiglitz and Weiss, 1981). Empirical research finds that financial constraints help determine entry into entrepreneurship, including via house wealth channels (Andersen and Nielsen, 2012; Adelino et al., 2015; Kerr et al., 2022). There is also a related literature on expanding access to credit for traditionally underserved entrepreneurs (Fairlie and Robb, 2007; Karlan and Zinman, 2010; Melzer, 2011; Cook et al., 2023). Another contribution of this paper is to provide a new method for measuring who gains and who loses when switching between models, which may be useful to policymakers and researchers, especially for evaluating machine learning models.

This paper contributes to the question of how alternative inputs, machine learning, and risk-taking interact in non-bank lending. There is concern that machine learning models predict risk more flexibly, which can disadvantage protected groups by better approximating disallowed inputs such as race (Fuster et al., 2022; U.S. White House, 2023). Yet fintech lenders who rely heavily on machine learning claim to expand access to credit by identifying creditworthy borrowers who would be rejected by conventional underwriting.[9] Indeed, there is evidence that fintech lenders are more likely to extend credit to under-served minorities and distressed areas (Jagtiani and Lemieux, 2018; Erel and Liebersohn, 2020; Chernenko and Scharfstein, 2022; Cornelli et al., 2024; Howell et al., 2024). However, there is also evidence that loans to underserved groups reflect fintechs taking more risk and engaging in regulatory arbitrage (Buchak et al., 2018; Bartlett et al., 2022; Propson et al., 2024). We contribute to this debate by demonstrating that incorporating new inputs, such as cash flow data, can affect how machine learning models impact protected classes.

In this way, we also contribute to a growing literature on the use of alternative data in lending. It is not obvious that recent cash flows will be useful in underwriting. Indeed, there is evidence from personal loans that in practice fintech lenders rely heavily on credit scores (Johnson et al., 2023) and, using German consumer loan data, that relationships matter more than private checking account information (Puri et al., 2017). Yet also using German data, Berg et al. (2020) show that digital footprints such as a consumer's device type can predict payment. Further, with data from

[9]For example, the American Fintech Council claims that its members are "democratizing financial services and creating critical access to families who have been left behind." (American Fintech Council, 2024).

Indian fintech lenders and payments processors, Ghosh et al. (2023) and Rishabh (2024) show that payments data are relevant for loan outcomes. Jagtiani and Lemieux (2019) and Di Maggio et al. (2022) use data from consumer lenders LendingClub and Upstart, respectively, to assess internal, proprietary risk scores alongside other alternative inputs such as education. Outside of fintechs and banks, "Big Techs" such as Alibaba and Amazon can incorporate timely cash flows into underwriting because they observe payment flows. Evidence from China and India suggests that such data can improve default risk prediction and lead to more inclusive credit (Huang et al., 2020; Ouyang, 2022; Liu et al., 2022; Gambacorta et al., 2023, 2024). Our paper builds on existing work on alternative data in lending by studying the use of replicable bank statement features, rather than proprietary metrics, and by focusing on the U.S. market. We show that bank statement-based variables are informative and expand access to credit for young entrepreneurs.

This finding is especially relevant for open banking policy, whose purpose is largely to facilitate the transfer of bank statement information to third party financial services providers (He et al., 2023). Babina et al. (2024), Alok et al. (2024), and Yu (2024) document how open banking policies promote fintech entry and unsecured lending. The U.S. unofficially has an open banking regime because private data aggregators such as Plaid use consumer credentials to access bank accounts. The Consumer Financial Protection Bureau recently issued a rule that will require banks to provide user data to third parties at their request, which is being contested by the banking industry (Consumer Financial Protection Bureau, 2024; Forcht Bank et al., 2024). In documenting the value of cash flows from bank statements in lending for younger entrepreneurs, we contribute to the open banking debate and, more generally, to work on the role of fintech lenders in the financial system (Mills and McCarthy, 2016; Seru, 2019; Tang, 2019; Vives, 2019; Vallee and Zeng, 2019; Philippon, 2019; Balyuk et al., 2020; Gopal and Schnabl, 2022; Ben-David et al., 2021).

2 Cash Flow-Based Lending and Age in Entrepreneurship

In this section, we describe how small business underwriting has changed in recent decades. Then we explain cash flow-based underwriting. Last, we discuss entrepreneurship over the life cycle and why we expect younger entrepreneurs to benefit from cash flow-based lending.

2.1 Evolution of Small Business Underwriting

Small business lending is a crucial financial service for job creation, economic mobility, and the viability of "main streets" across the U.S. Yet it is difficult to assess the credit quality of small businesses because they are highly heterogeneous, even within narrow sectors, and often have informal accounting data (Mills, 2024). Traditionally, small and regional banks were the primary source of loans. To overcome information asymmetry, they relied on soft information obtained by a loan officer through personal contact and long-term relationships (Petersen and Rajan, 1994; Berger and Udell, 1995). This changed as banks consolidated in the 1990s and 2000s: Larger banks took on a growing share of the market and they increasingly used systematic, standardized risk scoring processes (Black and Strahan, 2002; Berger et al., 2005a; Strahan, 2017). Thus small business lending became more arm's length.

While personal credit scores, like FICO, were always a key input, they took on an even larger role as lending became more arm's length and standardized (Arora, 2023; Carter and McNulty, 2005). However, personal credit scores are likely more useful for personal than for business loans. As a result, entrepreneurs with thinner personal credit files, such as immigrants, disadvantaged minorities, and young people may face financial constraints (Lauer, 2017; Cole et al., 2018; Rodriguez, 2020; Blattner and Nelson, 2021).[10] Notably, gender differentials in credit scores have disappeared in the U.S. (McGurran, 2020).

After the Financial Crisis of 2008-9, banks faced much higher capital requirements. Led by the largest banks, they reallocated their now more limited capital to more sustainable uses. The remaining small business lending at banks was held to high underwriting standards, requiring income measures drawn from annual financial statements and tax documents.[11] Firms which lacked strong historical financials going back several years, including the smallest, youngest, and riskiest firms, were left underserved (Cole and Damm, 2020). Small business credit cards became the most common form of financing for small firms (Fed, 2024), a costly financing source typically underwritten without considering bank account information.[12]

Noticing the small business lending gap, new fintech firms entered in the early 2010s. They focused on the riskier firms that were underserved by banks. Less constrained by regulation, these new lenders targeted the riskier firms underserved by banks. But they also needed to lend profitably

[10]The credit bureau Experian notes that "42% of adults lack a conventional credit score in a range that typically grants access to credit at standard rates" (Experian, 2024).

[11]Based on conversations with industry experts, including Dan O'Malley, the Founder and CEO of Numerated.

[12]See industry discussions of credit card underwriting practices at NerdWallet (2024).

at scale while still at arm's length, motivating the introduction of two key innovations—Flexible credit risk models built on state-of-the-art machine learning algorithms, and alternative data for building more holistic applicant profiles.

2.2 Cash Flow-Based Underwriting

The most important and widely used new input among U.S. fintech small business lenders was information from recent bank statements, referred to as cash flow data. To fintech lenders, this data represented a more direct measure of a borrower's ability to repay in a timely manner; that is, whether the applicant can service the debt based on their inflows, outflows, and recent signs of distress. Bank statements offer a direct window through which all lenders can assess a small business's ability to repay. In the U.S., bank statements are ubiquitous and accessible; nearly all businesses have a checking account and accessing them does not require third party approval. Moreover, checking account statements are relatively standardized, offering the same information for what are normally highly heterogeneous and difficult-to-compare businesses. Simple inflows and outflows may also be less influenced by socioeconomic factors compared to other new inputs like education or spending habits. Finally, new technologies allow them to be systematically digitized and parsed at scale.

While banks have long used client cash flow information when making large loans, they did not systematically use transaction account data to automate credit scoring for small businesses. Instead, they typically relied on annual tax or accounting data, or exclusively on personal FICO scores. Similarly, when underwriting personal credit cards, banks even today rarely analyze customer checking account cash flows, relying instead on self-reported annual income. Conversations with practitioners indicate three reasons for the limited use of checking account information in small business lending: (a) the technology to digitize and systematically score bank statements only became available after they shifted away from small business lending; (b) it was difficult to translate bank statement variables to the profit ratios that banks required to underwrite with minimal risk; and (c) the information was siloed and not easily transferrable from one unit of the bank to another.[13]

Cash flow data were deployed at scale by fintechs in the early 2010s. They pulled data electronically from PDFs or by “screenscraping" bank websites, a process that was later digitized and commodified by Plaid, Yodlee, Ocrolus, and others. While different lenders use different

[13]Based on conversations with industry experts, including Sam Taussig (formerly Kabbage and AmEx), David Snitkof (Ocrolus), Kelly Cochran (FinRegLab), Dan O'Malley (Numerated), and Karen Mills.

features and statistics, most firms use less than ten basic variables, which encompass credits and debits to capture inflows and outflows, as well as measures drawn from text about types of transactions that suggest distress, such as insufficient funds transactions or overdrafts.[14] They are now making their way into conventional lending. For example, in May 2024, both Experian and VantageScore launched new consumer credit score products that incorporate bank statement-based cashflow attributes.[15]

2.3 Financial Constraints for Younger Entrepreneurs

In this section, we establish that entrepreneurship peaks in the 40s and 50s, consistent with our data, and explore how financial constraints may help explain the limited entrepreneurial activity among younger individuals.

Average Age of Entrepreneurs. There is a great deal of evidence that entrepreneurs, especially successful ones and those establishing incorporated firms, tend to be in their 40s and 50s. Using high quality administrative data, Azoulay et al. (2020) show that the median new business founder is 42 years old. This corresponds to our more selected data on small businesses applying for fintech loans, which occurs shortly after firm founding. In our data, the median age among applicants and borrowers is 44 and 48, respectively.

Among employer firms overall, 51% of owners are over 55, a group that comprises just 23% of the labor force. In contrast, people under 35 represent 35% of the labor force but only 6% of business owners.[16] Zhang and Acs (2018) use data from the U.S. CPS to show that the propensity to be an entrepreneur rises with age, peaking at 71. The oldest entrepreneurs are more likely to be self-employed or part-time, with incorporation peaking by an entrepreneur's early 50s. Blanchflower (2004) finds that the probability of being self-employed reaches a peak at 76 in the EU and 54 in the U.S. Other work suggesting that older workers are more likely to be entrepreneurs or self-employed includes Zissimopoulos and Karoly (2007) and Fairlie et al. (2016). Overall, there seems to be an inverted U-shape in the relationship between age and entrepreneurship, which becomes steeper

[14]Based on author conversations with industry experts, including Venkatesh Bala (Biz2Credit), David Snitkof (Ocrolus), and Sam Taussig (formerly Kabbage and AmEx).

[15]Experian explains that "Cashflow Attributes, which leverages checking and savings account information, offers a more comprehensive view of an individual's financial profile. Using traditional credit data with lender-obtained cashflow information may unlock opportunities for consumers who may not have qualified if a lender was using traditional credit data on its own." VantageScore (a competitor to FICO) also released a new consumer credit score product that incorporates information from bank accounts, targeted to thin-file consumers, and marketed as exploiting "consumer permissioned open banking data." See Experian (2024) and VantageScore (2024).

[16]See the SBA fact sheet in 2021 (Headd, 2021).

approaching the age of 50 when researchers focus on businesses with employees or significant growth (Bates, 1990; Kautonen et al., 2017; Zhao et al., 2021; Brieger et al., 2021).

Constraints for Younger People. Age is highly correlated with FICO score. First, in our data, we see in Figure 1 that FICO rises essentially linearly from below 670 for entrepreneurs under 30 to 720 for entrepreneurs over 70. There is a similar correlation in the national data; Appendix Figure A.1 shows that FICO scores steadily increase with age, peaking at 757 for people in their 80s, with the steepest rise for people in their 50s to 60s. Even people in their 90s have among the highest average scores, highlighting the relevance of age across the entire lifecycle for underwriting decisions. Therefore, younger and middle-aged borrowers may be relatively FICO-constrained, as they will compare poorly to more elderly people. Consistent with this, Avery et al. (2012) conduct an exercise in which they build credit scores from scratch using information from a big-three credit bureau. They show that the scores act as a proxy for age, leading to disparate impacts for younger people, but have no measurable disparate impact by race or gender.

The constraints imposed by FICO are likely magnified by regulations that favor the elderly. While age is a protected class under the Equal Credit Opportunity Act, lenders are permitted to employ models that at least weakly favor older people. Indeed, the first paragraph about scoring models in lending from the CFPB states:

> *"Favoring the elderly: Any system of evaluating creditworthiness may favor a credit applicant who is age 62 or older. A credit program that offers more favorable credit terms to applicants age 62 or older is also permissible."* (CFPB, 2024)

A lender who hews to this provision may use age directly if it undergoes an additional validation process, a regulatory burden that deters the lenders in our data from directly using age. Otherwise, age cannot be a decisive factor in a credit scoring model, but it is acceptable to use different models for different age groups or to use scoring methods that are highly correlated with age, so long as the result is not lower scores on average for the elderly.[17]

Are younger people more financially constrained? Agarwal et al. (2009) show that among prime borrowers, those in middle age borrow at lower interest rates than younger people, after controlling for observable risk characteristics. While entrepreneurs in practice tend to be older, there is greater entrepreneurial desire or intent among younger people, especially those under 35 (Gielnik et al., 2018; Levesque and Minniti, 2006). An industry press article put it directly, with the title: "Banks Hate Young Entrepreneurs" (James, 2019). Levesque and Minniti (2011) are among

[17] See 12 CFR 202.6(b)(2) available here: https://www.ecfr.gov/current/title-12/chapter-II/subchapter-A/part-202; and https://www.fdic.gov/bank-examinations/fair-lending-implications-credit-scoring-systems.

the only papers to discuss the possible financial constraints that might yield these patterns; they theorize that there are more potential entrepreneurs among younger people but they face capital constraints. Older entrepreneurs also tend to have a greater human, social and financial capital and can therefore pursue riskier strategies (Becker, 1962; Seibert et al., 2001; Cooper et al., 1994). Younger people have not had as much time to accumulate a credit history, especially a record of mortgage repayment, which is a key input to a strong credit score.[18]

In sum, we expect that there is latent demand for entrepreneurship among young people, who are more likely to be disadvantaged by underwriting models that rely primarily on traditional credit scores. This implies that young people would especially benefit from underwriting incorporating timely information about ability to repay.

3 Data Sources and Summary Statistics

In this section, we introduce our data (Section 3.1) and describe summary statistics (Section 3.2).

3.1 Data Sources

We use data from three fintech companies. Two are non-bank, online lenders and one is a platform. These companies have requested anonymity. All three of them exclusively serve U.S. small businesses. Furthermore, all three, like many other fintech entrants, employ cash flow variables from bank statements.

The Lenders. We obtain comprehensive loan application, approval, origination, and performance data from two lenders, which are reasonably similar in the applicant characteristics they consider and the loans they underwrite. Both source capital via debt vehicles because they do not have access to deposits. Both require a personal guarantee and take out a blanket UCC lien again business assets. Their loans are described as collateralized—secured with tangible non-revenue and non-real estate business assets—but in practice they do not verify collateral or secure the loan with specific collateral.

[18]FICO divides its score drivers into five categories: payment history (35%; whether you've paid past credit accounts on time), amounts owed (30%), length of credit history (15%), new credit (10%; penalizes recent new credit) and credit mix (10%; benefits from a mortgage). While FICO emphasizes that it is current because it includes recent data, it also includes historical, long-term data. See `https://www.myfico.com/credit-education/whats-in-your-credit-score` and `https://www.fico.com/blogs/fico-fact-how-current-data-my-fico-score`.

We call one “Lender A” and the other “Lender B.” Lender A provides both term loans and lines of credit.[19] Its underwriting process is highly automated and it uses a risk scoring model. Lender B only provides term loans. Lender B’s underwriting process has a manual component, with an assigned individual who serves as the underwriter in each case. Lender B screens out subprime applicants (who have a FICO score below 660) in an initial step that we do not observe, contributing to a higher approval rate in our Lender B data.

The Platform. The Platform is a fintech company that connects prospective small business borrowers with lenders. It collects a standard set of application information and routes it to lenders, which include fintechs and conventional banks. In the data we use for analysis, there are 111 unique lenders. In our baseline data for the Platform, an observation is a forward from the platform to a lender. Each application is forwarded to an average (median) of 5 (3) lenders. Lenders provide the Platform with simple criteria, such as FICO or revenue thresholds. Conditional on meeting the lender’s requirements, the Platform sends applicants to a set of lenders based on an internal matching algorithm. While we do not observe the algorithm, we know that it does not employ age or other protected class indicators. We retain only the applicants forwarded to at least one lender. The lenders then decide whether to approve the loan and if so, on what terms. We observe these offers, which we describe as “approval” decisions. Finally, the applicant decides whether to take up an offer, and we observe which, if any, loan was originated.

Data Types and Collection. We obtained three types of data from all three companies: applications, third party underwriting inputs, and loan characteristics. We describe each in turn. First, we have information about the applicant firm and owner, some of which is used in underwriting and some of which is not. This is collected from the firm directly, such as owner name, address, owner age, firm founding date, and industry.[20] The owner first name is important for us to construct a gender indicator. We also observe underwriting inputs collected from third parties, specifically the FICO score from a credit bureau and bank statement information, which comes from the third party company Plaid for Lender A and third party company Ocrolus for Lender B and the Platform. In the latter case, we employ the original JSON files containing bank statement variables that are gathered from the Ocrolus API. Lender A and the Platform acquire

[19]We control for loan type in our analysis but do not find it has any material effects on the results.

[20]Industry is in the form of NAICS codes for both Lenders A and B, and Industry categories for the Platform. These are aggregated to 17 NAICS industries appearing in our data. We group sectors “Management of Companies and Enterprises" and “Administrative and Support and Waste Management and Remediation Services" with “Professional, Scientific, and Technical Services".

three months of bank statements, while Lender B acquires six months. We understand that three months is a common industry practice and firms rarely acquire more than six months.

Finally, for all three companies, we observe whether a loan was approved, information on the offered loan terms—amount, maturity, and interest rate—and whether the loan was originated (i.e., taken up). Unfortunately, we only observe performance (i.e., default) for the two Lenders and not for the platform. There are a total of 104,150 applications and 18,434 loans for Lender A, spanning 11/19/2013 to 6/25/2024. For Lender B, there are a total of 58,668 applications and 25,762 loans, spanning 8/15/2015 to 5/20/2024. Finally, we observe 199,300 applications, 904,471 application-forwards, 191,421 offers, and 35,821 originated loans on the Platform, spanning 1/1/2022 to 12/23/2023. Across all lenders, there are 262,723 unique firms. Although we introduce specific variables below, all variables are defined in Appendix B.

3.2 Summary Statistics

We present summary statistics about the full analysis sample of applicants and borrowers in Table 1. We use the applicant data from all three companies to analyze approval decisions, and borrower data from the two lenders to analyze default. In the table, the first columns report data from the Platform, where we see loan offers and origination but not performance, and where the unit of observation is an application-forward, so each applicant appears multiple times as they are routed to multiple lenders. The second and third sets of columns report data from Lenders A and B only on applications and borrowers, respectively.

Loan Approval and Performance. The first set of variables focus on underwriting activity. In part because of Lender B's initial screen, the approval rate (and overall applicant quality) is much higher in the Lender A and B data than in the Platform data, at 46% vs. 21%. The average requested and originated amounts are between \$100,000 and \$150,000 across the different samples. We construct an indicator variable for whether the loan is "non-performing" or not. This takes a value of one if the loan is charged off, more than 60 days past due, or the borrower has received forbearance and a modified loan. For parsimony, we refer to "non-performing" as "default." We observe in the final set of columns that 19% of loans default. The average interest rate for approved loans from Lender's A and B is 17%.[21] While some lenders on the platform provide term loans comparable to those offered by Lenders A and B, others focus on short-term

[21]There are two types of loans: Lines of credit and term loans. These compose 41% and 59% of the loans, respectively. In our analysis, we combine them and include a control for loan type.

loans with automatic ACH transfers. These loans are often extended to businesses that face challenges obtaining credit elsewhere. Given the short maturities and the higher risk associated with these borrowers, the resulting annualized interest rates tend to be elevated.[22] Figure A.2 shows the number of applications over time for each of the three data sources.

FICO and Bank Statements. The average applicant to the Platform has a credit score (FICO) of 681, which qualifies as Prime (CFPB, 2020). The average applicant at Lenders A and B has a higher FICO score of 724, which reflects the initial screening that we do not observe at Lender B, where applicants with subprime, or <660, FICO scores are removed. The average borrower at Lenders A and B has a FICO score of 728.

Bank statement variables are averaged across the three months prior to application. Credits (i.e. inflows) have a mean of roughly $100,000 for applicants, and $130,000 for borrowers. They decline slightly when adjusted for new debt obligations, although this variable is missing for some applicants. Withdrawals are slightly higher than overall credits. The average balance in the account is around $35,000 for applicants, and $42,000 for borrowers. There are three variables that may signal distress. One is the number of overdrafts or insufficient funds (NSF) transactions, which occur when the balance falls below zero. Some banks impose an insufficient funds fee and reject the transaction, while others permit the account to go negative and impose an overdraft fee. The second is the number of low or negative balances, which occur when the balance goes below zero or below $1,000. Consistent with strong selection on these variables, the averages are about five times higher among applicants to Lenders A and B than among borrowers. The third is the number of daily pay loans, which are merchant cash advances (MCA) that typically have very high interest rates and are reflected in the bank statement by daily withdrawals to the MCA lender. We also observe the standard deviation of credits and balances for the Lender A and B sample, which shed light on operational volatility.

Business Characteristics. We observe key dimensions from an underwriting perspective beyond financial information. Our proxy for firm size is employment. The Platform's applicants have seven employees on average, while applicants to Lender A and B have nine. The firms operate in diverse industries, with the highest concentrations in Construction and Retail Trade (see Table A.1). The average age of applicant firms is seven years, while it is 11 years among borrowers. We define

[22] ACH loans are often quoted in factors. For example, a 1.3 factor means that the borrower must repay $1.30 for each dollar borrowed. When payments are deducted daily over a 3 month window, this would amount to an APR above %120.

young firms as those less than five years old, which divides the Platform applications in half.

Race and Gender. We consider two protected classes besides age. The first is the Black share of the zipcode-level population.[23] There is evidence that Black people, including small business owners, face distinct challenges in accessing financial services in the U.S. (Blanchflower et al., 2003; Bayer et al., 2018; Bhutta and Hizmo, 2021; Ambrose et al., 2020; Giacoletti et al., 2021; Begley and Purnanandam, 2021; Fairlie et al., 2022). Blattner and Nelson (2021) show that relative to the White majority, racial minorities' credit scores are noisier indicators of default risk. In our applicant data, the average share of the local population that is Black is 14%, in line with the national average, though the median is much lower at 6.5%. The figures are slightly lower among borrowers. The second characteristic is gender. Women have historically composed a minority of new business founders (Coleman and Robb, 2009). In our data, women-owned firms compose 29% of applicants to the platform. The share in the Lender A and B sample is 23% of borrowers. In unreported analysis, we find little difference across genders in loan approval, performance, FICO, or cash flow variables.

Entrepreneur Age. We are primarily interested in the role of owner age, which we measure using the date of birth of the primary owner or CEO. The median age of the firm owner is 44 years old, which is close to what Azoulay et al. (2020) find across all *new* companies in the U.S. (42 years old). Among borrowers, the median is 48 (Table 1). Figure A.3 shows histograms of age in the pooled applicant (Panel A) and originated loan (Panel B) samples. We define the variable Young Entrepreneur as an indicator variable that takes the value of one if the applicant's age is below 45, and zero otherwise. 56% of the applicant sample and 40% of the borrower sample are young using this definition. In analysis, we also use alternative cutoffs. Our groupings are not only natural given our data, but also align with peak entrepreneurship in the literature (see Section 2.3.)

Younger entrepreneurs appear to be at a disadvantage. Figure 2 shows that younger owners have a 22% approval rate, compared to 25% for older owners (Panel A), yet they have identical default rates (Panel B). Panel A uses data from all three sources; Table A.2 shows that this pattern replicates in the Lender A & B data, where 37% of young applicants are approved, compared to 43% of older owners. Figure 3 shows that approval is strongly increasing in entrepreneur age and

[23] We measure this using 2019 U.S. Census Bureau American Community Survey data. We use a zip code-level measure for three reasons. First, this captures a key dimension of traditional socioeconomic disadvantage, even for individuals in these areas who are not Black. Second, we do not observe self-reported race and the traditional proxies that rely on location and name are problematic (Greenwald et al., 2023). Third, policy efforts to address minority access to credit are often place-based, in part because of legal challenges to policy based on individual race.

has a much steeper slope for younger borrowers. We break the data at 45, corresponding to the median of the sample.

The disadvantage in FICO from Figure 1 also translates to an approval-default differential; we observe that young owners are at a severe FICO deficit among applicants but not among borrowers (Figure 2 Panel C). Specifically, there is a 17.3 point FICO gap between younger and older applicants, which is 26% of a standard deviation. Yet this disparity disappears for most cash flow variables. As one example, Panel D shows the number of insufficient funds transactions. We see that while approval clearly selects on this variable, there is little difference across age groups.

Complete summary statistics by age group are in Table 2, where the first set of columns include all applicants from all three sources and the second set restricts to borrowers and thus to Lenders A & B. There are no meaningful differences in gender or the Black population share across younger and older entrepreneurs. Younger entrepreneurs generally have weaker FICO and cash flow metrics.[24] However, the differences in cash flow metrics are much smaller in percentage terms or relative to the standard deviation. For example, the average FICO differential among applicants is 26% of the standard deviation. In contrast, the disparity in inflows (credits) is just 9% of the standard deviation, and the distress proxies are essentially the same across the age groups (rates of insufficient funds transactions, low or negative balances, and daily pay (MCA) loans). In sum, it seems that the age disparity in FICO could potentially be mitigated with new inputs from recent bank statements.

Descriptive analysis further suggests that cash flow information will especially benefit relatively FICO-constrained younger people. When we consider only low-FICO borrowers, younger people appear to have similar or better cash flow measures. We see in Table A.3 that the groups have nearly identical average FICO scores, unlike the full sample of Lender A and B borrowers in the right-hand columns of Table 2. Also unlike the full sample, the average of the credits variable is nearly equal across the young and old, and the group averages of other measures appear stronger for the young. For example, within the low fico population, young entrepreneurs have lower rates of insufficient funds transactions, low balances, and daily pay (MCA) loans. This suggests that there are young borrowers who have strong cash flow measures but low FICO, and thus should benefit from cash flow-based underwriting.

[24] Note that our summary statistics in Table 2 are limited to loan applicants and approved loans in which we observe owner age. In Table A.2, we repeat these applicant statistics for the Lender A and B subset. The differences are similar, even though the applicant pool is quite different. The average FICO is 718 for young owners vs. 732 for mature owners, while the cash flow variables are generally similar.

Selection into the Data. The applicants we observe in our data are a self-selected group relative to the overall population of U.S. small businesses, in that they chose to apply to an online fintech lender. However, they are observably representative along important dimensions. For example, the average small firm in the U.S. has 11.7 employees as of 2020, which is similar to 10 employees in the Lender data and 7 in the Platform data.[25] Women make up 21.7% of employer firm owners, which is very close to the figures in our data. Table A.1 compares the industry composition with the national data on employer firms with 1-19 employees from the U.S. Census Bureau.[26] The data are overall fairly representative, especially for originated loans. As mentioned above, the distribution of entrepreneur age is also representative.

4 Recent Cash Flows and Loan Performance?

In order to test whether cash flow variables from bank statements benefit younger entrepreneurs, we first need to build a general model of how they predict default. Note that lenders build risk scoring models by predicting default, and if a variable predicts default we would expect lenders to use it in underwriting. Therefore, in this initial section, we show that they predict default conditional on the other key inputs to underwriting, especially the personal credit score (FICO).

Since cash flow data were used to originate the loans in our sample, we expect to underestimate their predictive power relative to a setting where cash flow data were observed but not incorporated into underwriting. However, cash flow variables should still have explanatory power in our sample for at least two reasons. First, the fintech lenders in our sample are new entrants, lacking years of their own performance data to train their models at the time of underwriting. Our analysis, by contrast, leverages more years of realized loan performance data. Second, the business model of the lenders entails approving some borrowers with positive default risk. This means that cash flow variables, even if used optimally in screening, may have residual predictive information about loan performance.

4.1 Regression Analysis

We begin with a series of regression analyses, which tell us the direction and magnitude of predictive power. The R^2 of the regressions shed some light on informativeness. We predict

[25] See `https://advocacy.sba.gov/wp-content/uploads/2023/03/Frequently-Asked-Questions-About-Small-Business-March-2023-508c.pdf`.

[26] Available here `https://www.census.gov/programs-surveys/susb/data.html`.

default (i.e., a non-performing loan) using variants of the model in Equation 1:

$$\mathbb{1}(\text{Non-Performing}_i) = \delta\text{FICO}_i + \beta_1'\text{Cash Flow Chars}_i + \gamma_1'\text{Firm Chars}_i + \gamma_2'\text{Loan Chars}_i + \text{Industry}_i + \text{State}_i + \alpha_t + \varepsilon_i. \tag{1}$$

The coefficients of interest are the FICO and cash flow variables. To make interpretation straightforward, we standardize to z-scores, so that the coefficient represents the effect of a one standard deviation change in the independent variable on the outcome variable.[27] We include controls for characteristics that the lenders observe and that are permissible to use in underwriting under fair lending laws. These are firm age in years and size (number of employees), as well as fixed effects for the applicant's industry, state, calendar quarter of application (we observe 36 quarters), and loan type. In some models, we further include loan characteristics including the requested and originated loan amounts, maturity, and APR. We cluster standard errors by industry and quarter.

The results, using the Lenders A and B sample, are in Table 3. The first column represents a baseline, or FICO-only OLS. The second column adds the cash flow variables. Columns 3 and 4 repeat this but with several fixed effects and controls for business age and size. We repeat this exercise with a Logistic model in columns 5-6. The key takeaway from this table is that the coefficients on the cash flow variables have economically significant, robust predictive power in intuitive directions. For example, more credits to the account and higher bank balances are associated with lower chances of default, whereas more withdrawals and higher volatility in revenues and balances are associated with higher chances of default. Daily pay loans and the number of days with low or negative balances also predict higher default. Column 7 adds loan characteristics to the OLS model (APR, loan amount, and maturity). These controls slightly reduce the magnitude for some cash flow variables, which we expect if lenders are using these variables to determine loan terms. As one example of interpretation, column 7 shows that a one standard deviation increase in FICO (49 points) is associated with an 18% decline in default relative to the mean. Similarly, a one standard deviation increase in observed account credits reduces default probability by 10% of the mean (using statistics from Table 2).

The 31% difference in R^2 between columns 1 and 2 suggests added value from cash flow variables in a minimal model. After adding fixed effects and controls for business age and size, the 13% difference in R^2 between columns 3 and 4 continues to suggest added value from cash flow variables. The logit specification also suggests a better fit in the cash flow model, with a

[27]For any missing values of the cash flow variables that are not fully populated, we impute the mean and include a binary indicator to account for missingness in the model.

14% increase in Pseudo R^2 from columns 5 to 6. We verify in a robustness test that the results from Table 3 hold for first time borrowers (Table A.4). Overall, this table shows that cash flow variables have predictive power and economically interpretable directions (i.e., whether a larger value "helps" or "hurts" on average), and points to added value for predicting loan performance.

4.2 Machine Learning Analysis

Constraining the relationship to a linear form creates misspecification bias because some variables enter nonlinearly and interact with one another (Sadhwani et al., 2021; Barbaglia et al., 2023). Therefore, our main method is a machine learning (ML) model that flexibly predicts outcomes, allowing us to compare a baseline FICO-based model to one that adds cash flow variables. The ML model is also the basis for our main analysis studying benefits for younger entrepreneurs.

We use a random forest machine learning algorithm (Ho, 1995). Like OLS and other classical models, it aims to minimize prediction error.[28] There are three steps to the estimation process. First, models such as random forest require the econometrician to select "hyperparameters," which act as knobs governing model behavior.[29] To select hyperparameters, we use random subsamples of the data. In the second step, we construct training, validation, and testing samples by randomly selecting loans from the entire dataset.[30] The training dataset represents 80% of the sample and the testing dataset (or "holdout sample") 20%. Only the training dataset is used for fitting the model, while the holdout sample is used for prediction. Within the training dataset, 20% of observations are used for a validation process that fine-tunes the model and helps to reduce overfitting. The third step is to estimate the model many times—an econometric technique called "bootstrapping"—in order to obtain average estimates and standard errors. This estimation method is explained in more detail in Appendix C.

We conduct this procedure for two models: a "Baseline" model that includes FICO, firm size, firm age, and industry roughly represents a traditional underwriting model and a "Cash Flow"

[28] Like other recent research on machine learning models in finance (e.g., Fuster et al. (2022)), we focus on random forest models. In our setting, we find that random forest performs similarly to other modern machine learning methods like XGBoost but are much faster to train and considerably easier to implement, with a smaller number of hyperparameters and less tendency to overfit.

[29] This is similar to the way an econometrician must properly select the bandwidth in a regression discontinuity design or the number of lagged periods in a vector autoregressive model. While hyperparameter choices do not directly affect the predictions of the model, appropriate choices help the model to be well-behaved. Examples of hyperparameters for random forests are the maximum depth, the number of samples required to split an internal node, and regularization hyperparameters.

[30] We split the data with stratification on the outcome variable, which ensures the same proportion of each outcome class in the training dataset and testing dataset.

model adds bank statement-based variables to the Baseline model. Our preferred approach has comprehensive controls, as in columns 3-6 of Table 3. We also estimate a minimal model corresponding to columns 1-2 of Table 3. We further include economically meaningful transformations of the Cash Flow variables (see Table A.5 for a full list of features in both the Baseline and Cash Flow models).[31]

We evaluate the performance of the models using three measures, which are all reported in Table 4. The first, known as the receiver-operating characteristic (ROC) area under the curve (AUC), is widely used to evaluate binary classification models.[32] The ROC curve describes the true positive rate (TPR) and the false positive rate (FPR) at various thresholds.[33] The ROC AUC is the integral of the ROC curve, where a value of 0.5 indicates no discriminative ability (equivalent to random guessing), and a value of 1.0 indicates perfect discrimination. Typically, an AUC of 0.6 is considered desirable in information-scarce environments, whereas AUCs of 0.7 or greater are the goal in information-rich environments. In our case, the minimal CF model has a 0.642 ROC AUC, compared to 0.624 in the Baseline model. One way to measure improvement is to compare the two estimates to the 0.5 random guessing threshold (e.g. see Berg et al. (2020)). Using this method, the CF model improves performance by 15%.[34] Using the preferred specification, the increase is about 1 pp or 7.2% relative to the 0.5 threshold. A 1 percentage point improvement in AUC is considered a significant gain in the credit scoring industry.[35] Figure 4 Panel A plots the ROC curves using this preferred specification, and shows that the Cash Flow model ROC curve is shifted northwest of the Baseline model ROC curve, indicating better performance. Additionally, Figure 4 Panel B plots the ROC curves of the same preferred specification, but with the sample

[31] Simple transformations, such as scaling variables to have a mean of 0 and standard deviation of 1, allow the model to better identify complex patterns within the cash flow data, especially when relationships are non-linear or feature scaling varies. Although these transformations could be applied to OLS models, the decision-tree structure in random forests uncovers non-linear patterns more effectively without as much risk of overfitting.

[32] We note that ROC AUC is known to be less informative in situations like ours, when positive outcomes are sparse and false positive rates are naturally low. Also, Hand (2009) demonstrates that the AUC has a potentially serious deficiency: it uses different weights when computing misclassification costs for different classifiers. This means that in some settings, AUC is incoherent. In our setting, these weights are unlikely to differ significantly.

[33] For any binary classification model of default, a decision rule (threshold) is needed to determine whether a particular observation with some non-zero probability of default (e.g., 4%) should be classified as a "0" (non-default) or a "1" (default). As this threshold is lowered, the TPR approaches 1: virtually every person who eventually defaults is accurately predicted as a "1" (default). But simultaneously, the FPR approaches 1: nearly everyone is categorized as a "1," even those who do not eventually default. Conversely, increasing the threshold leads to a lower FPR (since almost every observation is predicted as a "0"), but also lowers the TPR (observations likely to default are incorrectly predicted as a "0"). The ROC curve avoids the need to select one threshold as a decision rule, and instead maps out this relationship between the TPR and the FPR for a specific model over all possible thresholds.

[34] This improvement is calculated as $\frac{(0.642-0.50)-(0.624-0.50)}{(0.624-0.50)} = 0.145$.

[35] Iyer et al. (2016), for example, find a 5.7% AUC improvement using non-standard information (characteristics of the listing text, group and friend endorsements as well as borrower choice variables such as listing duration and listing category) compared to these AUC using the Experian credit score.

split into older and younger owners.

The other measures of performance are better suited to our context where defaults are sparse. One is the Precision-Recall curve, where Precision = $\frac{TPR}{TPR+FPR}$ (the share of predicted defaults that actually default) and Recall = $\frac{TPR}{TPR+FNR}$ (the share of defaults that were predicted). The intuition is the same as in the above, where a larger area under the curve implies a better model. We observe that the Cash Flow model outperforms the Baseline model by 6% in the minimal and 5% in the preferred specification. Next, we compute the H-Measure for each model, an alternative to AUC proposed by Hand (2009).[36] Here, the Cash Flow model outperforms the Baseline model by 26% in the minimal and 14% in the preferred specification. Finally, we establish that the Cash Flow models perform significantly better (at the 1% level) across all three measures. This is performed with 100 bootstraps using the "Corrected Resampled t-Test" proposed by Nadeau and Bengio (1999). We graph the results for the ROC AUC and H-Measure in Figure A.4.

Machine learning models capture complex, nonlinear interactions, making the interpretation of individual feature effects less straightforward than in regressions. We present a graphical representation of feature importance in our model in Figure 5. For a random forest model, feature importance quantifies the contribution of each feature to the model's predictive performance.[37] We report the 10 features from previous tables with the highest feature importance for our CF Model predicting borrower default. FICO is the single most significant predictor, reflecting its well-established role in evaluating a borrower's creditworthiness. Several cash flow measures also rank highly, and the combined importance of the cash flow variables significantly exceeds that of FICO. For example, credits, balance, and their standard deviations are together more than four times as important as FICO.

In sum, this section has established that information about cash flows from bank statements is valuable in risk prediction, even in a sample where lenders have used these variables to set the interest rate and define the population of approved borrowers.

4.3 Model Informativeness by Age Group

We now examine whether the models with cash flow variables are more *informative* for younger people. This will contribute to how these new data are applied in practice. In Table 5, we repeat the

[36]We use the hmeasure package for Python to compute the H-Measure.

[37]We assess feature importance using scikit-learn's RandomForestClassifier's implementation of the "Mean Decrease in Impurity" for a specific random forest ensemble model. This attribute reflects the average decrease in impurity (entropy) attributable to each feature across all trees in the forest. The model sums the impurity reductions for each feature whenever it is used to split a node, providing a heuristic ranking of feature influence.

default prediction regression analysis found in Table 3 columns 3-4, but separate the population by age. There are two key takeaways from the table. First, the increase in explanatory power as we shift from the Baseline model (first column) to the Cash Flow (second column) model is higher for young owners, indicating that cash flow variables are particularly relevant for underwriting decisions involving this group. The increase in R^2 is about 18% for younger owners and only 8% for older owners. The second takeaway is that the coefficients on cash flow variables for young owners are generally larger in absolute magnitude than for older owners. While FICO also has a larger coefficient for younger owners, the difference is much greater for the cash flow variables. For example, one standard deviation in credits is associated with 2.6 pp lower default chance among young owners, compared to just 1.6 pp among older owners. The coefficients for FICO are 3.8 and 3.6, respectively.

We conduct several robustness tests. First, we repeat this analysis for various age cutoffs in Table A.6. Second, the tests we conducted earlier deliver similar results for the age split. In Table A.4 we verify similar relationships for first-time borrowers. Overall, the absolute value of the magnitudes of the coefficients on cash flow variables for young owners are generally larger than for old owners. This suggests that these variables are more important in underwriting decisions for young owners than for old.

The superior model performance for young owners persists in our random forest models. We conduct subsample analyses in the ML framework, comparing the performance of the Cash Flow (CF) model with the Baseline model for young owners.[38] Table 6 presents the results. For all of the age thresholds, the improvement in the model when moving from the Baseline to the CF model is larger for the younger group. Notably, the largest difference is around the lowest threshold where the ROC AUC improvement from the CF model is 2.7 pp higher. Relative to a model without any predictive power, this translates into a 22% increase.

5 Benefits for Younger Entrepreneurs

Building on the models from the previous section, we now study the relative benefits and informativeness for younger borrowers of cash flow relative to baseline models. First, we introduce a new methodology for comparing outcomes between two predictive models for

[38] We use the method described in Section 4.2 to tune hyperparameters, train the baseline and CF models, and calculate performance statistics, such as ROC AUC. Unlike in Section 4.2, we compute performance statistics using only a subsample of observations: we split the holdout dataset into firms with and without ‘young’ business owners, and compare the performance of the baseline and CF models on these subsamples.

specific subgroups of observations, which we call “Tail Analysis for Comparative Outcomes" (TACO). TACO summarizes complex distributional differences between models and provides a statistically rigorous point estimate of the net benefit (or harm) to a subgroup from adopting a new model. The method highlights observations with the largest divergences between models—those in the tails of the outcome distribution. TACO is applicable to prediction models with univariate continuous outcome variables and allows for flexible subgrouping based on characteristics of interest. Second, we apply TACO to identify which age groups benefit or are adversely affected when switching from the Baseline model to the Cash Flow (CF) model.

5.1 Tail Analysis for Comparative Outcomes (TACO)

TACO is a generalizable method for identifying subgroups most affected by switching predictive models. It is particularly relevant for comparing models with continuous outcome variables and emphasizes the tails of the outcome distribution, where prediction differences are most pronounced.

We first summarize Tail Analysis for Comparative Outcomes (TACO) in four steps:

1. **Model Estimation**: Define two prediction models, f and g, that predict the same target outcome, y. Note that f and g could differ with respect to model types (e.g., OLS vs. machine learning), specifications, parameters, or training data.
2. **Prediction Difference**: For each observation i, compute the prediction difference:

$$h_i = g(X_i) - f(X_i)$$

   where h_i represents the change in predicted outcomes when moving from model f to g.

3. **Tail Identification**: Identify the top and bottom $p\%$ of observations based on h_i:

$$T^+ = \{X : h(X) > Q_{1-p}(h)\}, \quad T^- = \{X : h(X) < Q_p(h)\},$$

   where Q_p denotes the p-th percentile of $h(X)$. Depending on context, T^+ may represent the group benefiting most (e.g., increased approval likelihood) or harmed most (e.g., higher predicted default probability).

4. **Comparison of Characteristics**: Compare population means of T^+ and T^- for a characteristic of interest (e.g., firm age or an indicator for young entrepreneurs) to assess distributional impacts. A useful summary statistic is the ratio of the characteristic means

between the groups in each tail, which we call the "TACO ratio." A ratio of one indicates that the characteristic is proportionally represented in both tails, indicating that switching models creates no net imbalance in the characteristic.

$$\textbf{TACO Ratio}_{(i)} = \frac{\frac{1}{|T^-|}\sum_{T^-} x_i}{\frac{1}{|T^+|}\sum_{T^+} x_i}$$

Predicted outcomes from different models may not always be directly comparable due to differences in calibration or scaling. For example, a score of 0.8 in one underwriting model may represent a different level of risk than the same value in another model, even if the models achieve similar overall performance metrics. TACO addresses this challenge by focusing on the relative ranking of observations, regardless of each model's prediction distribution. By identifying the tails based on percentiles rather than absolute scores, TACO is robust to differences in scaling or calibration between models, making it broadly applicable to comparisons of predictive models.

5.2 Implementation of TACO in Our Setting

We use TACO to evaluate how switching from the Baseline to the CF model to predict default affects younger entrepreneurs.[39] First, we follow the procedure in Section 4.2 to fit Baseline and CF models to predict default probabilities for holdout samples. Second, we compute the prediction differences $\widehat{h(x)} = \widehat{g(x)} - \widehat{f(x)}$ and identify the top and bottom deciles (T^+ and T^-) of $\widehat{h(x)}$ across borrowers. Third, we calculate the mean values of relevant characteristics within these tails. In our case, this requires identifying the share of young people in each tail (e.g., the mean value of the indicator for being less than 45 years old). Fourth, we bootstrap these steps to estimate average effects and standard errors. Finally, we calculate the TACO ratio as the share of young people in the group most helped (T^-) divided by the share of young people in the group most hurt (T^+).

In our specific application, we evaluate whether switching from the Baseline to the CF model meaningfully changes default predictions for borrowers, particularly in the tails of the prediction distribution. If the models diverge significantly, the thresholds defining T^+ and T^- will reflect large differences in predicted default probabilities, indicating greater potential reallocation of credit. In our case, the prediction differences in the tails correspond to substantial changes in predicted

[39]Note that we cannot perform TACO on approval decisions since lenders use CF variables to make this decision. An approval model trained without CF variables is inherently the wrong model to benchmark against: we observe no baseline in which approval decisions are made without using CF variables. On the other hand, a baseline model can make predictions about ex-post default *conditional on approval* under the assumption that, once approved, default is not directly influenced by the underwriting model that was used to make the approval decision.

default rates: the T^+ threshold represents an increase of more than 4.8 percentage points, while the T^- threshold represents a decrease of 4.6 percentage points. Both thresholds deviate meaningfully from the mean default rate of 16.7 percentage points.[40]

Table 7 reports the TACO ratios and key inputs for four age indicators. We use deciles (top and bottom 10%) to define the tails (T^+ and T^-) of default prediction changes, where T^+ represents those most hurt by a switch from the Baseline to the CF model, and T^- represents those most helped.[41] The first two columns describe owner age among all observations in the bootstrap holdout samples that fall into the tails, while the next four columns disaggregate T^+ and T^-. The last column reports the TACO ratios: They are less than one if the group is adversely affected and greater than one if the group benefits. In the first group of estimates, we see that the TACO ratios far exceed one. The stars indicate that they are significantly different from one. For the owners younger than 35, the TACO ratio is 1.60. In other words, the group that most benefits from the CF model has 60% more young owners than the group that is most hurt. The TACO ratios decline monotonically as we increase the threshold for defining young owner. For thresholds of 40, 45, and 50 years old, the ratios are 1.45, 1.39, and 1.30. We plot these results in Figure 6, showing that the results are robust to alternative definitions of the tail. We further report estimates where the tail groups are based on quintiles, quartiles, or terciles (20th, 25th, and 33rd percentiles).

To examine the connection between age and FICO as a source of financial constraints, we focus on the subsample that is both young and low-FICO. The results, reported in the second set of estimates in Table 7, are uniformly much larger. For example, the TACO ratio for owners below the age of 35 is 7.78. This indicates that among low-FICO borrowers, young entrepreneurs are over seven times as likely to appear in the group benefiting most from the CF model as they are to appear in the group most hurt by the model. This represents a dramatic reallocation within this group. In contrast, for high-FICO younger entrepreneurs, shown in the bottom set of estimates, the CF model with its reduced reliance on FICO tends to hurt more than help; the TACO ratio is 0.66 for the under-35 group. This is expected because the CF model reduces reliance on FICO, which is generally less favorable for high-FICO borrowers. However, the benefit for low-FICO young people (7.8) is an order of magnitude larger than the harm for high-FICO young people (0.66). Even in this high-FICO group, the Cash Flow model is relatively better for younger people. Appendix Table A.7 shows that the TACO ratio is just 0.44 for the high-FICO group overall, while

[40]We also examine the overall distributions of predicted probabilities for the Baseline and CF models. As shown in Figure A.5, the distributions are similar, with nearly identical means (16.7%) and standard deviations (5%). This similarity further supports the appropriateness of comparing predicted probabilities of default using TACO.

[41]The underlying difference in predicted probability of default when switching from the Baseline to the CF model for young and older owners is graphed in Figure A.6.

the CF model disadvantage for young borrowers with high FICO is smallest for the youngest groups, increasing with the age threshold.

We conduct a placebo test to ensure that the results do not reflect spurious sorting of winners and losers repeatedly. Here, we replace the Baseline model used in the comparisons with the same CF model. The results, in Figure A.7, are precisely zero. Finally, we check robustness by computing an alternative statistic, the Adverse Impact Ratio (AIR), in each tail. The AIR represents a rate for the protected group divided by that rate for the comparison group. For example, courts employ AIR in the "four-fifths" rule, where the selection rate—such as for employment—of a protected group must be at least 80% that of the majority group.[42] We calculate AIR as the share of younger individuals divided by the share of older individuals. To mimic TACO, we do this only in the tails. The results are presented in Table A.8, using the same format as Table 7. A higher AIR in the benefit group relative to the overall sample means that younger business owners benefit from the CF model. Conversely, a lower AIR in the hurt group suggests that younger owners are less represented among those who are negatively impacted. Younger owners consistently have higher AIR values in the benefit group compared to the full sample average and lower AIR values in the hurt group. This pattern is more pronounced in the Low FICO sample.

As supplementary analysis, we consider other characteristics in Figure A.8 and Table A.7. As expected, there are huge benefits to the CF model for low-FICO borrowers, consistent with the CF variables offering orthogonal information about creditworthiness. We next consider business age. Like young entrepreneurs, we see that younger businesses benefit, consistent with the literature showing that their lack of operating and credit history contributes to financial constraints (Guariglia, 2008; Howell, 2017; Ma et al., 2022).[43] Young firms might benefit from recent cash flow data because even if they have a "thin" credit file or operating history, they may be creditworthy if they have earned enough in recent months to service new debt. We also consider high Black population shares (at the owner's home zipcode level) using two thresholds, the median and 75th percentile within our sample. While the ratios are positive for high-Black areas, they are not statistically significant. Finally, for female borrowers, we see that a higher proportion of women are hurt by the CF model compared to those who are helped, though the ratio is close to one. Interestingly, this reverses among low-FICO women, as shown in Table A.7. The CF model

[42]The AIR, along with the four-fifths rule, is used to assess potential discrimination in U.S. courts and by agencies such as the U.S. Department of Labor and the U.S. Equal Employment Opportunity Commission (EEOC). For example, see https://www.justice.gov/crt/fcs/T6Manual7.

[43]Research on public firms has found that younger firms are more financially constrained (Whited and Wu, 2006; Hadlock and Pierce, 2010), though the evidence is mixed (Farre-Mensa and Ljungqvist, 2016).

strongly benefits low-FICO women.[44]

6 Effect of Cash Flow Information on Credit Access

The analysis in Section 5 focused on predicting default, which should be externally valid beyond our particular sample (subject to the degree to which our population of borrowers is selected). In the final step, we leverage the richness of our data—which includes both application and approval information—to make causal claims and confirm the main finding: access to cash flow-based underwriting provides significant benefits to younger entrepreneurs. Specifically, we assess whether assignment via the Platform to a more cash flow-intensive lender differentially affects approval rates for younger versus older entrepreneurs. In a supplementary test in Section 6.2, we show that the effect is robust to using random assignment of loan officers at one of our lenders, and also document how cash flow variables predict approval.

Before using approval decisions in causal analysis, we ensure that they predict approval decisions in our data on average. This would be consistent with being used in underwriting, which we know is the case at Lenders A and B. We repeat the analysis above but use approval rather than default. Figure A.9 reports the performance measures graphically for the ML model, using the same preferred specification as the default model. Panel A shows that the ROC AUC for the CF model is 0.85, compared to 0.82 for the Baseline model. At the bottom of Table 4, we report all three performance metrics for approval. The improvement of the CF model shows that lenders are indeed relying on cash flow measures when making approval decisions. We also present regression results in Tables A.9 and A.10.[45] Below, we will show that not all lenders on the Platform appear to use cash flow variables. However, in the overall sample, enough do in order for these variables to be predictive on average.

[44] This gender disparity could reflect a number of mechanisms; given the similarity of FICO scores across the two groups, one possibility is industry sorting. Women are much more likely to be in healthcare and education services, and less likely to be in construction and transportation and warehousing. These figures are not reported, but for example, among borrowers, women represent 36% of health care and 44% of education, but only 13% of construction and 14% of transportation and warehousing.

[45] We additionally run two robustness tests using loan approval models. Under this analysis, the same broad conclusions hold as for default. See the appendix for tables that restrict observations to first-time applicants (Table A.11) and change the age cutoffs (Table A.12).

6.1 Within-Application Assignment to Cash Flow-Intensive Lenders

We are interested in the causal effect of a lender's reliance on cash flow-based underwriting on loan approval outcomes. A naive regression of approval on lender characteristics would be biased if the Platform systematically forwards applicants with certain characteristics to certain lenders based on, say, risk profile. Such selection makes it difficult to disentangle the causal impact of cash flow underwriting from differences in the applicant mix.

Our identification strategy addresses this challenge by exploiting within-applicant variation (à la Khwaja and Mian (2008)). The Platform typically forwards the same application to multiple lenders; within our analysis sample, the average and median are five and four lenders, respectively, at the unique application level. We identify those lenders for which cash flow variables are predictive of approval ("cash flow-intensive," or CFI). By including application fixed effects in our regressions, we control for all observed and unobserved applicant characteristics, effectively isolating lender-side differences in underwriting practices.

We calculate CFI for each lender using the ML model from Section 4.1, but where the outcome variable is loan approval. Specifically, we run the Baseline model (i.e., FICO plus controls) and the Cash Flow model (adding the bank statement variables) for each lender. We identify CFI lenders as those for which the AUC ROC is higher in the Cash Flow model than in the Baseline model.[46] Consistent with the Platform's assertion that it works with both conventional and fintech lenders, 38% of the lenders have zero or negative values for AUC ROC improvement. We categorize the 62% with positive improvement as CFI. Summary statistics about the assignment are in Table A.13, where we compare application-forwards, applications, and lenders across the CFI and non-CFI categorizations. While the average age of applicants is the same across the types of lenders, in general the CFI lenders receive higher-risk applicants, consistent with fintech lenders serving riskier segments relative to conventional lenders, as we noted earlier.

We regress approval decisions for application i at time t from lender l on the interaction between the lender-level CFI indicator variable ("CFI") and an indicator for whether the applicant is young ("Young Owner"):

$$\mathbb{1}(\text{Approved}_{ilt}) = \beta_1 \mathbb{1}(\text{Young Owner}_i) \times \mathbb{1}(\text{CFI}_l) + \text{Lender}_l + \text{Application}_i + \varepsilon_{ilt} \qquad (2)$$

The model includes lender fixed effects (Lender_l) and application fixed effects (Application_i).

[46] We exclude lenders receiving under 50 application forwards, with an approval rate over 95%, or with an AUC improvement over 20%.

Application fixed effects absorb all observable and unobservable attributes of the applicant that may influence approval decisions. Specifically, since applicant i is held constant, potential omitted variables related to borrower risk are differenced out. That is, the applicant's risk characteristics are constant across lenders, so we do not require firm- or owner-level controls for identification. As a result, even systematic patterns in applicant forwarding, such as consistent steering of young applicants to CFI lenders or to "tough" lenders, would not bias our estimates. Any differences in approval across lenders must arise from lender-side behavior, such as how intensively a lender uses cash flow data.

Our setting provides relatively strong guarantees of overlap, SUTVA, and conditional independence, which are also needed for causal identification. Many applicants are forwarded to at least one lender in each category, ensuring within-applicant variation in lender types. Specifically, of the 184,063 unique applications (comprising 879,889 observations), 77,124 (606,442 observations) are forwarded to both CFI and Non-CFI lenders. In addition, the Platform forwards applications without conveying real-time information between lenders and lenders generally provide offers roughly simultaneously. This leaves minimal scope for a lender to condition on the presence (or decision) of other lenders. Taken together, these assumptions imply that *conditional on application fixed effects*, variation in lender types is "as good as random" from the perspective of any single applicant-lender match.

The results are in Table 8. We split the sample around a FICO score of 700.[47] We also adjust the sample to compare younger owners to over-50 owners. For example, in column 1, we compare under-35 owners with over-50 owners.[48] We observe that low-FICO young owners benefit most from access to cash flow-based underwriting. Using the 35-year old threshold and restricting the sample to high-FICO owners, the effect of young interacted with CFI lender is 0.86 pp, or 3.6% of the mean (column 1). In the low-FICO sample, the effect is higher at 2.2 pp, or 13% of the mean (column 2). As the age threshold increases, high-FICO younger people tend to benefit more, but the coefficients are consistently larger for the low-FICO group. For example, using the 50-year old threshold, the effect is 1.3 pp for high-FICO applicants and 2.0 pp for low-FICO applicants (columns 7-8).

We report four alternative models in the Appendix. In Table A.15 we present the results without

[47] This cutoff identifies lower-score individuals while ensuring a large sample in each category. It is roughly the midpoint of the "Good" definition, see e.g. https://money.usnews.com/credit-cards/articles/credit-score-is-it-good-or-bad and https://www.equifax.com/personal/education/credit/score/articles/-/learn/what-is-a-good-credit-score/.

[48] We omit from the sample owners between 35 and 50, because in that case we are comparing very young and young owners, vitiating the effect. The full sample results are similar, and reported in Table A.14.

application fixed effects, instead employing industry, quarter, and lender fixed effects. The results are generally similar to those in Table 8. We restrict the sample to applications forwarded to both CFI and Non-CFI lenders in Table A.16. The results are very close to the full-sample results. We show a triple interaction specification in Table A.17. Last, we show the result without splitting by FICO in Table A.18 Panel A.

Do CFI Lenders Simply Take Higher Risk? It is possible that these results reflect a different lender characteristic correlated with more intensive use of cash flow variables that tends to favor younger applicants. In this case, the results would still identify benefits to young owners of access to fintech lenders who use cash flow variables, but they would have different implications for the broader use of cash flow variables by conventional lenders. The specific concern is that the results may reflect fintech lenders taking higher risk *and* being more likely to use cash flow variables. As we explained in Section 2, fintech lenders have tended to target higher-risk segments. Since younger people are generally riskier, CFI lenders may approve more younger people because they are less regulated and take more risk, not because they use cash flow-based underwriting. A first counterargument is that there are large effects among high-FICO applicants in the odd columns of Table 8. Further, we find positive effects even restricting the sample to applicants with FICO scores over 800 (Table A.18), which represents about the top 20% of FICO scores nationally.[49] Therefore, while low-FICO younger owners benefit the most, the result does not simply reflect high-risk segments.

To further explore this concern, we examine whether younger entrepreneurs assigned to CFI lenders are more likely to experience business failure, which we expect if the higher approval rate reflects higher risk. We collect data on business survival because we do not observe loan performance in the platform. Clerical workers researched outcomes for a random subset of 46,000 unique applications (corresponding to about 250,000 application-forwards and 43,000 approval decisions).[50] We first show in column 1 of Table A.19 that the overall result holds in this subsample; among low-FICO borrowers, younger people assigned to CFI lenders are 1.9 pp more likely to be approved, about 12% of the mean. Next, in column 2, we use all application-forwards and predict business survival. If CFI lenders tend to take higher risk, the interaction between approval and CFI lender should be negative. Instead, it is precisely zero. The key test is in column

[49] See `https://www.fico.com/blogs/average-u-s-ficor-score-716-indicating-improvement-consumer-credit-behaviors-despite-pandemic`.

[50] The clerical workers assessed whether the business appears to be open based on a Google search of the business name and state, with a check on the industry for common names. They examined whether the business is identified as active based on a Google sidebar. The data were collected in September, 2024.

3. Here, we restrict the sample to approved application-forwards. We regress business survival on the interaction between CFI lender and young. The interaction coefficient is positive and insignificant, instead of the negative coefficient we would expect if risk preferences explained the main finding. In sum, it does not seem to be the case that being young and approved by a CFI lender reflects lender risk preferences.

6.2 Random Assignment to Cash Flow-Intensive Loan Officers

Since the Platform analysis is within-application, non-random assignment to lenders should not affect the results. However, we confirm the finding using an instance of true random assignment that we observe in the Lender B's underwriting process. At Lender B, human loan officers make a final decision, taking into account cash flow variables and other inputs such as FICO. Lender B randomly assigns applications to loan officers. While Lender B is a cash flow-intensive lender overall, there is variation among the loan officers in their reliance on cash flow information.

We observe 60 loan officers (by first and last name), of which 15 have sufficient observations where owner age is observed to construct a measure of cash flow intensity (CFI). These loan officers considered about 11,500 applicants. Since there are a smaller number of observations, we calculate CFI for each loan officer using the logit model from Section 4.1. Since there is improvement in ROC AUC for all officers (Lender B being a CFI underwriter), we identify CFI loan officers as those with top-quartile improvement. (The results are similar using other thresholds.)

We present summary statistics about these data in Table A.20. Panel A reports data at the application level, while Panel B reports data at the unique loan officer level. Almost exactly half of observations at the application level are assigned to CFI officers. Random assignment predicts that key ex-ante cross-sectional applicant characteristics should be roughly evenly split across CFI and non-CFI loan officers, though we expect some noise due to the small number of officers. As above, the final column of Table A.20 reports t-test results for all variables. Importantly, FICO, cash flow variables, and age are the same across both types of loan officers, consistent with random assignment. CFI loan officers approve a somewhat lower share of applicants but have no significant difference in ultimate default rate. Overall, while again the sample of officers is very small, it is consistent with random assignment and thus offers a useful supplementary test.

The estimation approach is as follows:

$$\mathbb{1}(\text{Approved}_i) = \beta_1 \mathbb{1}(\text{Young Owner}_i) + \beta_2 \mathbb{1}(\text{Young Owner}_i) \times \mathbb{1}(\text{CFI}_i) \\ + \text{Loan Officer}_k + \text{Industry}_i + \alpha_t + \varepsilon_{ik} \tag{3}$$

Here, the coefficient of interest is β_2 on the interaction of younger owner and a CFI loan officer. We include loan officer, applicant industry, and application quarter (α_t) fixed effects. The loan officer and time fixed effects ensure that any variation in the composition of applicants should not bias the results. However, CFI loan officers could be different from their counterparts, raising the concern that the type of officer who is CFI also favors certain age groups. (This is analogous to the concern about lender risk preferences from Section 6.1. While we cannot exclude this possibility, we believe it is unlikely for two reasons: First, age is not directly used in underwriting; second, such a preference would be illegal. The industry controls rule out a channel where the preference comes via industry heuristics. Standard errors are clustered jointly by applicant and loan officer.

We present the results in Table 9, using the same structure as the Platform analysis.[51] As in the Lender analysis, we also adjust the sample to compare a group of young owners to over-50 owners in each column. For example, in column 1, we compare under-35 owners over-50 owners.[52] Focusing on columns 1-2, which use the 35 year old threshold, we see that young people with a low-FICO assigned to CFI loan officers are 13.3 pp more likely to be approved relative to older people, which is 28% of the mean. In contrast, the result is an insignificant 3.8 pp among high-FICO applicants. Looking across the table, we observe a strong and secular decline in the coefficient magnitude as we raise the age threshold, suggesting that when causally identified, the youngest borrowers benefit the most from assignment to CFI underwriting.

7 Conclusion

Interest in cash flow-based underwriting has risen among policymakers and practitioners in recent years, especially in the debate about Open Banking mandates. Advocates argue that incorporating timely information about ability to repay from bank statements can help democratize access to credit and foster innovation and competition in the financial services industry (e.g., Chopra (2023)). In this paper, we offer the first analysis of variables drawn directly and transparently from bank

[51] We divide FICO around the median here, because if we use Low FICO as above there are insufficient observations; the coefficient magnitude grows but loses statistical significance.

[52] The full sample results are similar, reported in Table A.21.

statements, which any lender could replicate. We compare these with FICO, the traditional central input for credit scoring. Unlike much existing work on alternative inputs in lending, we focus on the U.S., where there is an active policy debate about access to data and where there are particularly high barriers to entry for new lenders, due to the "spaghetti soup" of regulatory authorities and the market power of banks (Mills and McCarthy, 2016).

We also focus on small business lending, which is an important context for at least three reasons: Small businesses are key drivers of employment and upward mobility; information asymmetry and heterogeneity create financial constraints; and large banks have to a significant degree pulled out of the sector. All of this suggests that new, non-bank entrants who leverage cash flows to underwrite could add value. Consistent with this, we show that cash flow information helps to predict default when added to simple, standard underwriting models.

Our central thesis is a natural implication of the power of cash flow information relative to FICO. Young entrepreneurs, who are mechanically disadvantaged by FICO, should benefit from the inclusion of cash flow-based variables. We document that age is strongly correlated with FICO, which is not the case for other protected classes. We offer a novel method—Tail Analysis for Comparative Outcomes (TACO)—to compare the benefits of one model relative to another for population subgroups. We use the method to show that young people benefit from switching from a baseline, FICO-driven model to a model with cash flow information. Finally, we employ application and approval data in two quasi-experimental designs to explore implications for access to credit. We find that random assignment to more cash flow-intensive underwriting benefits younger entrepreneurs.

References

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2015, House prices, collateral, and self-employment, *Journal of Financial Economics* 117, 288–306.

Agarwal, Sumit, John C Driscoll, Xavier Gabaix, and David Laibson, 2009, The age of reason: Financial decisions over the life cycle and implications for regulation, *Brookings Papers on Economic Activity* .

Alok, Shashwat, Pulak Ghosh, Nirupama Kulkarni, and Manju Puri, 2024, Open banking and digital payments: Implications for credit access, Technical report, National Bureau of Economic Research.

Ambrose, Brent W., James N. Conklin, and Luis A. Lopez, 2020, Does borrower and broker race affect the cost of mortgage credit?, *Review of Financial Studies* 34, 790–826.

American Fintech Council, 2024, Fintech lending, Policy priorities on responsible fintech lending and access to affordable credit.

Andersen, Steffen, and Kasper Meisner Nielsen, 2012, Ability or finances as constraints on entrepreneurship? evidence from survival rates in a natural experiment, *The Review of Financial Studies* 25, 3684–3710.

Arora, Rohit, 2023, How to improve your credit score and secure a business loan, Technical report, Forbes.

Avery, Robert B, Kenneth P Brevoort, and Glenn Canner, 2012, Does credit scoring produce a disparate impact?, *Real Estate Economics* 40, S65–S114.

Azoulay, Pierre, Benjamin F Jones, J Daniel Kim, and Javier Miranda, 2020, Age and high-growth entrepreneurship, *American Economic Review: Insights* 2, 65–82.

Babina, Tania, Saleem A Bahaj, Greg Buchak, Filippo De Marco, Angus K Foulis, Will Gornall, Francesco Mazzola, and Tong Yu, 2024, Customer data access and fintech entry: Early evidence from open banking, Technical report, National Bureau of Economic Research.

Balyuk, Tetyana, Allen N Berger, and John Hackney, 2020, What is fueling fintech lending? the role of banking market structure, *The Role of Banking Market Structure (June 23, 2020)* .

Barbaglia, Luca, Sebastiano Manzan, and Elisa Tosetti, 2023, Forecasting loan default in europe with machine learning, *Journal of Financial Econometrics* 21, 569–596.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace, 2022, Consumer-lending discrimination in the fintech era, *Journal of Financial Economics* 143, 30–56.

Bates, Timothy, 1990, Entrepreneur human capital inputs and small business longevity, *The review of Economics and Statistics* 551–559.

Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross, 2018, What drives racial and ethnic differences in high-cost mortgages? the role of high-risk lenders, *Review of Financial Studies* 31, 175–205.

Becker, Gary S, 1962, Investment in human capital: A theoretical analysis, *Journal of political economy* 70, 9–49.

Begley, Taylor A., and Amiyatosh Purnanandam, 2021, Color and credit: Race, regulation, and the quality of financial services, *Journal of Financial Economics* 141, 48–65.

Ben-David, Itzhak, Mark J Johnson, and René M Stulz, 2021, Why did small business fintech lending dry up during the covid-19 crisis?, Technical report, National Bureau of Economic Research.

Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri, 2020, On the Rise of FinTechs: Credit Scoring Using Digital Footprints, *The Review of Financial Studies* 33, 2845–2897.

Berger, Allen N, W Scott Frame, and Nathan H Miller, 2005a, Credit scoring and the availability, price, and risk of small business credit, *Journal of money, credit and banking* 191–222.

Berger, Allen N, Nathan H Miller, Mitchell A Petersen, Raghuram G Rajan, and Jeremy C Stein, 2005b, Does function follow organizational form? evidence from the lending practices of large and small banks, *Journal of Financial economics* 76, 237–269.

Berger, Allen N, and Gregory F Udell, 1995, Relationship lending and lines of credit in small firm finance, *Journal of business* 351–381.

Bhutta, Neil, and Aurel Hizmo, 2021, Do minorities pay more for mortgages?, *Review of Financial Studies* 34.

Black, Sandra E, and Philip E Strahan, 2002, Entrepreneurship and bank credit availability, *The Journal of Finance* 57, 2807–2833.

Blanchflower, David G, 2004, Self-employment: More may not be better.

Blanchflower, David G., Phillip B. Levine, and David J. Zimmerman, 2003, Discrimination in the small-business credit market, *Review of Economics and Statistics* 85, 930–943.

Blattner, Laura, and Scott Nelson, 2021, How costly is noise? Data and disparities in consumer credit, *arXiv preprint arXiv:2105.07554* .

Breiman, Leo, 2001, Random forests, *Machine learning* 45, 5–32.

Brieger, Steven A, Anne Bäro, Giuseppe Criaco, and Siri A Terjesen, 2021, Entrepreneurs' age, institutions, and social value creation goals: A multi-country study, *Small Business Economics* 57, 425–453.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru, 2018, Fintech, regulatory arbitrage, and the rise of shadow banks, *Journal of financial economics* 130, 453–483.

Carter, David A, and James E McNulty, 2005, Deregulation, technological change, and the business-lending performance of large and small banks, *Journal of Banking & Finance* 29, 1113–1130.

CFPB, 2020, Borrower risk profiles, Technical report, Consumer Financial Protection Bureau.

CFPB, 2024, Regulation b (equal credit opportunity act): Official interpretations of § 1002.6, Provides official interpretations of § 1002.6 regarding rules for creditworthiness evaluations under the Equal Credit Opportunity Act.

Chen, Brian S, Samuel G Hanson, and Jeremy C Stein, 2017, The decline of big-bank lending to small business: Dynamic impacts on local credit and labor markets, Technical report, National Bureau of Economic Research.

Chernenko, Sergey, and David S. Scharfstein, 2022, Racial disparities in the Paycheck Protection Program, Working Paper 29748, National Bureau of Economic Research.

Chopra, Rohit, 2023, Prepared remarks of cfpb director rohit chopra on the proposed personal financial data rights rule, Accessed: 2024-12-10.

Cole, Rebel A, and Jason Damm, 2020, How did the financial crisis affect small-business lending in the united states?, *Journal of Financial Research* 43, 767–820.

Cole, Rebel A, Hamid Mehran, and Germana Giombini, 2018, Gender and the availability of credit to privately held firms: evidence from the surveys of small business finances, *FRB of New York staff report* .

Coleman, Susan, and Alicia Robb, 2009, A comparison of new firm financing by gender: evidence from the kauffman firm survey data, *Small Business Economics* 33, 397–411.

Consumer Financial Protection Bureau, 2024, Cfpb finalizes personal financial data rights rule to boost competition, protect privacy, and give families more choice in financial services, Press release detailing the finalization of the

Personal Financial Data Rights Rule, including its implementation phases and consumer impact.

Cook, Lisa D, Matt Marx, and Emmanuel Yimfor, 2023, Funding black high-growth startups, *Available at SSRN* .

Cooper, Arnold C, F Javier Gimeno-Gascon, and Carolyn Y Woo, 1994, Initial human and financial capital as predictors of new venture performance, *Journal of business venturing* 9, 371–395.

Cornelli, Giulio, Jon Frost, Leonardo Gambacorta, and Julapa Jagtiani, 2024, The impact of fintech lending on credit access for us small businesses, *Journal of Financial Stability* 101290.

Crosman, Penny, 2024, How banks can overcome challenges of cash flow underwriting technology, Technical report, American Banker.

Di Maggio, Marco, Dimuthu Ratnadiwakara, and Don Carmichael, 2022, Invisible primes: Fintech lending with alternative data, Technical report, National Bureau of Economic Research.

Erel, Isil, and Jack Liebersohn, 2020, Does fintech substitute for banks? evidence from the paycheck protection program, Technical report, National Bureau of Economic Research.

Evans, David S, and Boyan Jovanovic, 1989, An estimated model of entrepreneurial choice under liquidity constraints, *Journal of political economy* 97, 808–827.

Experian, 2024, New experian tool empowers financial inclusion through open banking insights, Technical report, Experian Press Release, May 9.

Fairlie, Robert, Alicia Robb, and David T Robinson, 2022, Black and white: Access to capital among minority-owned start-ups, *Management Science* 68, 2377–2400.

Fairlie, Robert W, EJ Reedy, Arnobio Morelix, and Joshua Russell-Fritch, 2016, Kauffman index of startup activity: National trends 2016, *Available at SSRN 2828359* .

Fairlie, Robert W., and Alicia M. Robb, 2007, Why are black-owned businesses less successful than white-owned businesses? The role of families, inheritances, and business human capital, *Journal of Labor Economics* 25, 289–323.

Fama, Eugene F, 1985, What's different about banks?, *Journal of monetary economics* 15, 29–39.

Farre-Mensa, Joan, and Alexander Ljungqvist, 2016, Do measures of financial constraints measure financial constraints?, *The review of financial studies* 29, 271–308.

Fed, 2024, 2024 report on employer firms: Findings from the 2023 small business credit survey, Technical report, Federal Reserve Banks.

FinRegLab, 2019, The use of cash flow data in underwriting credit: Empirical research findings, Technical report, FinRegLab, Accessed: 2024-12-10.

Forcht Bank, N.A., Kentucky Bankers Association, and Bank Policy Institute, 2024, Complaint for declaratory and injunctive relief: Forcht bank, n.a., kentucky bankers association, and bank policy institute v. consumer financial protection bureau and rohit chopra, Case No. 5:24-cv-00304-DCR, United States District Court for the Eastern District of Kentucky, Lexington Division. Filed October 23, 2024.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, 2022, Predictably unequal? the effects of machine learning on credit markets, *The Journal of Finance* 77, 5–47.

Gambacorta, Leonardo, Yiping Huang, Zhenhua Li, Han Qiu, and Shu Chen, 2023, Data versus collateral, *Review of Finance* 27, 369–398.

Gambacorta, Leonardo, Yiping Huang, Han Qiu, and Jingyi Wang, 2024, How do machine learning and non-traditional data affect credit scoring? new evidence from a chinese fintech firm, *Journal of Financial Stability* 101284.

Getter, Darryl E, 2019, Small business credit markets and selected policy issues, *Report R45878, Congressional Research Service* .

Ghosh, Pulak, Boris Vallee, and Yao Zeng, 2023, Fintech lending and cashless payments, Technical report, Proceedings of Paris December 2021 Finance Meeting EUROFIDAI - ESSEC.

Giacoletti, Marco, Rawley Heimer, and Edison G. Yu, 2021, Using high-frequency evaluations to estimate discrimination: Evidence from mortgage loan officers, Working Paper 21-04, FRB of Philadelphia.

Gielnik, Michael M, Hannes Zacher, and Mo Wang, 2018, Age in the entrepreneurial process: The role of future time perspective and prior entrepreneurial experience., *Journal of Applied Psychology* 103, 1067.

Goldstein, Caroline, and Sally Lauckner, 2022, Is there a minimum credit score for business loan eligibility?, Online article, Updated June 30, 2022.

Gopal, Manasa, and Philipp Schnabl, 2022, The rise of finance companies and fintech lenders in small business lending, *The Review of Financial Studies* 35, 4859–4901.

Greenwald, Daniel, Sabrina T Howell, Cangyuan Li, and Emmanuel Yimfor, 2023, Regulatory arbitrage or random errors? implications of race prediction algorithms in fair lending analysis, *Implications of Race Prediction Algorithms in Fair Lending Analysis (April 12, 2023)* .

Guariglia, Alessandra, 2008, Internal financial constraints, external financial constraints, and investment choice: Evidence from a panel of uk firms, *Journal of banking & finance* 32, 1795–1809.

Hadlock, Charles J, and Joshua R Pierce, 2010, New evidence on measuring financial constraints: Moving beyond the kz index, *The review of financial studies* 23, 1909–1940.

Hand, David J, 2009, Measuring classifier performance: a coherent alternative to the area under the roc curve, *Machine learning* 77, 103–123.

He, Zhiguo, Jing Huang, and Jidong Zhou, 2023, Open banking: Credit market competition when borrowers own the data, *Journal of financial economics* 147, 449–474.

Headd, Brian, 2021, Small business facts: Business ownership demographics, Analysis of business ownership diversity, firm sizes, and demographic trends.

Ho, Tin Kam, 1995, Random decision forests, in *Proceedings of 3rd international conference on document analysis and recognition*, volume 1, 278–282, IEEE.

Howell, Sabrina T, 2017, Financing innovation: Evidence from r&d grants, *American economic review* 107, 1136–1164.

Howell, Sabrina T, Theresa Kuchler, David Snitkof, Johannes Stroebel, and Jun Wong, 2024, Lender automation and racial disparities in credit access, *The Journal of Finance* 79, 1457–1512.

Huang, Yiping, Longmei Zhang, Zhenhua Li, Han Qiu, Tao Sun, and Xue Wang, 2020, Fintech credit risk assessment for smes: Evidence from china, Technical report, IMF working paper.

Iyer, Rajkamal, Asim Ijaz Khwaja, Erzo F P Luttmer, and Kelly Shue, 2016, Screening Peers Softly: Inferring the Quality of Small Borrowers, *Management Science* .

Jagtiani, Julapa, and Catharine Lemieux, 2018, Do fintech lenders penetrate areas that are underserved by traditional banks?, *Journal of Economics and Business* 100, 43–54.

Jagtiani, Julapa, and Catharine Lemieux, 2019, The roles of alternative data and machine learning in fintech lending: evidence from the lendingclub consumer platform, *Financial Management* 48, 1009–1029.

James, Ethan, 2019, Banks hate young entrepreneurs, Technical report, Lendza, June 17.

Johnson, Mark J, Itzhak Ben-David, Jason Lee, and Vincent Yao, 2023, Fintech lending with lowtech pricing, Technical report, National Bureau of Economic Research.

Karlan, Dean, and Jonathan Zinman, 2010, Expanding credit access: Using randomized supply decisions to estimate the impacts, *The Review of Financial Studies* 23, 433–464.

Kautonen, Teemu, Ewald Kibler, and Maria Minniti, 2017, Late-career entrepreneurship, income and quality of life, *Journal of Business Venturing* 32, 318–333.

Kerr, Sari Pekkala, William R Kerr, and Ramana Nanda, 2022, House prices, home equity and entrepreneurship: Evidence from us census micro data, *Journal of Monetary Economics* 130, 103–119.

Khwaja, Asim Ijaz, and Atif Mian, 2008, Tracing the impact of bank liquidity shocks: Evidence from an emerging market, *American Economic Review* 98, 1413–1442.

Lauer, Josh, 2017, *Creditworthy: A history of consumer surveillance and financial identity in America* (Columbia University Press).

Lawler, Ryan, 2024, Plaid introduces cash flow data-based credit report, Technical report, Axios.

Levesque, Moren, and Maria Minniti, 2006, The effect of aging on entrepreneurial behavior, *Journal of business venturing* 21, 177–194.

Levesque, Moren, and Maria Minniti, 2011, Age matters: How demographics influence aggregate entrepreneurship, *Strategic Entrepreneurship Journal* 5, 269–284.

Liu, Lei, Guangli Lu, and Wei Xiong, 2022, *The big tech lending model* (National Bureau of Economic Research).

Ma, Song, Justin Murfin, and Ryan Pratt, 2022, Young firms, old capital, *Journal of Financial Economics* 146, 331–356.

McGurran, Brianna, 2020, Women and credit 2020: How history shaped today's credit landscape, Technical report, Experian.

Melzer, Brian T, 2011, The real costs of credit access: Evidence from the payday lending market, *The Quarterly Journal of Economics* 126, 517–555.

Mills, Karen, and Brayden McCarthy, 2016, The state of small business lending: Innovation and technology and the implications for regulation, *Harvard Business School Entrepreneurial Management Working Paper* 17–042.

Mills, Karen G, 2024, *Fintech, Small Business & the American Dream How Technology Is Transforming Lending and Shaping a New Era of Small Business Opportunity* (Palgrave Macmillan; Second Edition).

Nadeau, Claude, and Yoshua Bengio, 1999, Inference for the generalization error, *Advances in neural information processing systems* 12.

NerdWallet, 2024, Business credit card requirements: What you need to apply, Published 7 months ago; provides an overview of the necessary information and qualifications for applying for a business credit card.

Ouyang, Shumiao, 2022, Cashless payment and financial inclusion, *Working Paper* .

Petersen, Mitchell A, and Raghuram G Rajan, 1994, The benefits of lending relationships: Evidence from small business data, *The journal of finance* 49, 3–37.

Petersen, Mitchell A, and Raghuram G Rajan, 2002, Does distance still matter? the information revolution in small business lending, *The journal of Finance* 57, 2533–2570.

Philippon, Thomas, 2019, On fintech and financial inclusion, Technical report, National Bureau of Economic Research.

Propson, Drew, Emina Ajvazoska, Felipe Ferri de Camargo Paes, Stanley Mutinda, Dana Salman, Jill Lagos Shemin, Krishnamurthy Suresh, Zhifu Xie, and Bryan Zhang, 2024, The future of global fintech: Towards resilient and inclusive growth, Technical report, World Economic Forum.

Puri, Manju, Jörg Rocholl, and Sascha Steffen, 2017, What do a million observations have to say about loan defaults? opening the black box of relationships, *Journal of Financial Intermediation* 31, 1–15.

Rishabh, Kumar, 2024, Beyond the bureau: Interoperable payment data for loan screening and monitoring, *Available at SSRN 4782597* .

Rodriguez, Lorena, 2020, All data is not credit data, *Columbia Law Review* 120, 1843–1884.

Sadhwani, Apaar, Kay Giesecke, and Justin Sirignano, 2021, Deep learning for mortgage risk, *Journal of Financial Econometrics* 19, 313–368.

SBA, 2023, Frequently asked questions about small business, U.S. Small Business Administration, Office of Advocacy. Statistics based on various government sources, including Census Bureau and Bureau of Labor Statistics data.

Seibert, Scott E, Maria L Kraimer, and Robert C Liden, 2001, A social capital theory of career success, *Academy of management journal* 44, 219–237.

Seru, Amit, 2019, Regulating banks in the era of fintech shadow banks, *The Andrew Crocket Memorial Lecture, Bank for International Settlements, Basel* 30.

Stiglitz, Joseph E, and Andrew Weiss, 1981, Credit rationing in markets with imperfect information, *The American economic review* 71, 393–410.

Strahan, Philip E, 2017, Are concentrated banks better informed than diversified ones?, *Journal of Accounting and Economics* 64, 278–283.

Tang, Huan, 2019, Peer-to-peer lenders versus banks: Substitutes or complements?, *Review of Financial Studies* 32, 1900–1938.

U.S. White House, 2023, Executive order on the safe, secure, and trustworthy development and use of artificial intelligence, Technical report, White House Press Release, October 30.

Vallee, Boris, and Yao Zeng, 2019, Marketplace lending: A new banking paradigm?, *The Review of Financial Studies* 32, 1939–1982.

VantageScore, 2024, Credit bureaus + bank data = vantagescore 4 plus, Technical report, VantageScore Website.

Vives, Xavier, 2019, Digital disruption in banking, *Annual Review of Financial Economics* 11, 243–272.

Whited, Toni M, and Guojun Wu, 2006, Financial constraints risk, *The review of financial studies* 19, 531–559.

Yu, Tong, 2024, Data as collateral: Open banking for small business lending, Technical report, Imperial College Working Paper.

Zhang, Ting, and Zoltan Acs, 2018, Age and entrepreneurship: nuances from entrepreneur types and generation effects, *Small Business Economics* 51, 773–809.

Zhao, Hao, Gina O'Connor, Jihong Wu, and GT Lumpkin, 2021, Age and entrepreneurial career success: A review and a meta-analysis, *Journal of Business Venturing* 36, 106007.

Zissimopoulos, Julie M, and Lynn A Karoly, 2007, Transitions to self-employment at older ages: The role of wealth, health, health insurance and other factors, *Labour economics* 14, 269–295.

Figure 1: FICO and Age Relationship

This figure uses applicant data from Lender A, Lender B, and the Platform to show the relationship between FICO score and age (N = 1,027,837). This is a binscatter with 20 equal-sized age bins and a line of best fit.

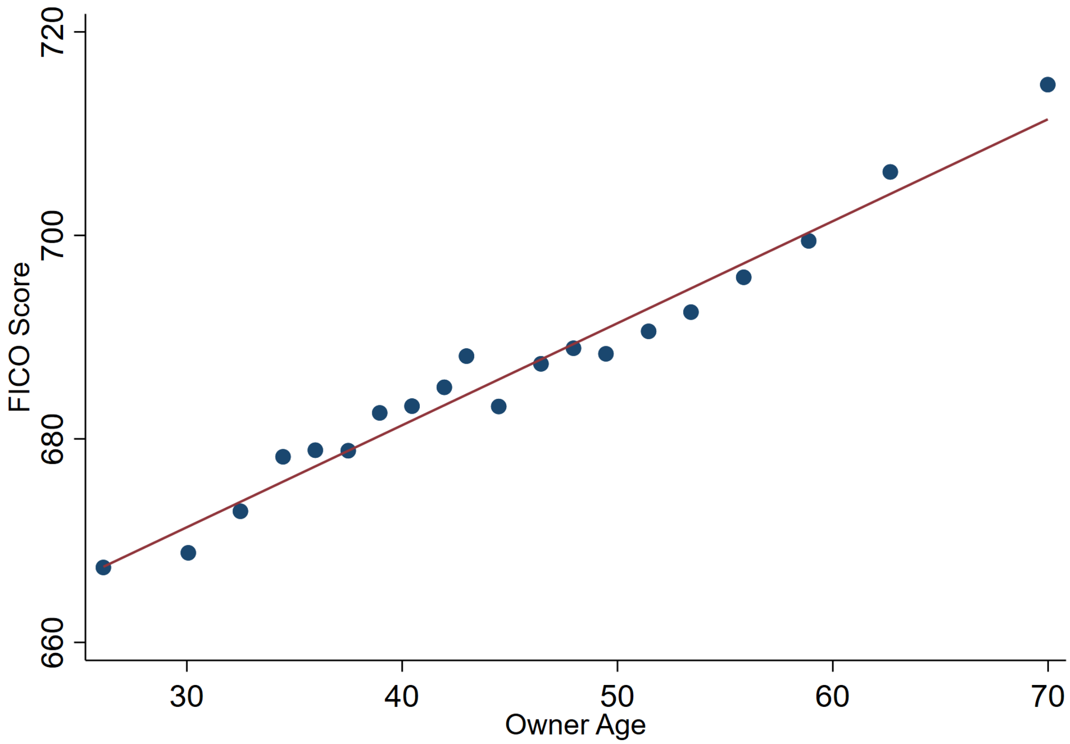

Figure 2: Variation in Loan Outcomes and Inputs by Entrepreneur Age

This figure uses applicant data from Lender A, Lender B, and the Platform to show the differences in key analysis variables for young and mature owners (N = 1,027,837). Panel A shows the shares of younger and older owners whose loan applications are approved. Panel B shows, within the sample of borrowers, the shares of younger and older owners whose loans default (N = 20,190). Panels C and D present the mean FICO score and number of insufficient funds transactions in the bank statement for four groups: younger applicants, older applicants, younger borrowers, and older borrowers. Young owners are those less than 45, while mature owners are at least 45. Each bar includes 95% confidence intervals using standard deviation.

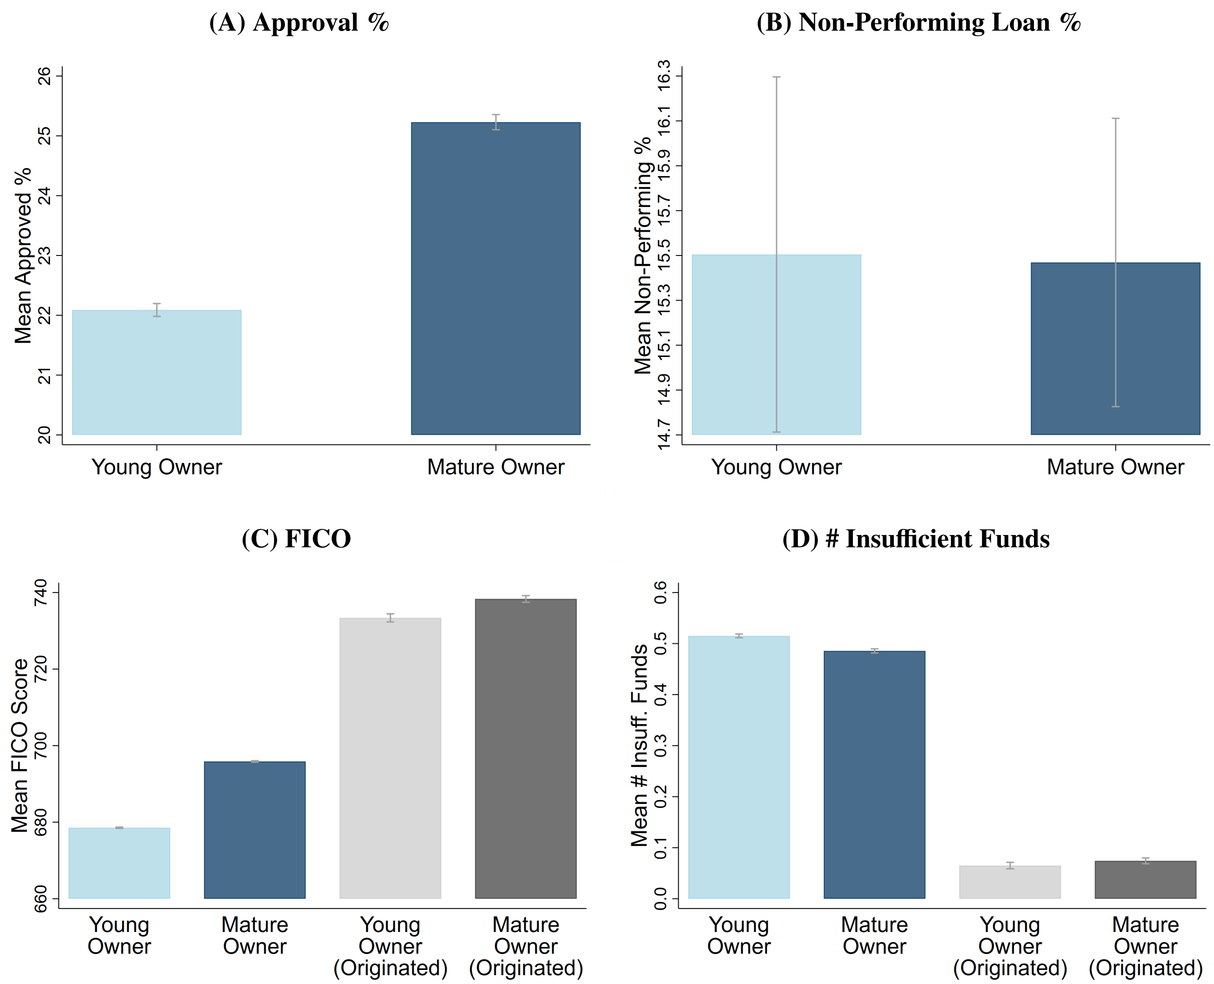

Figure 3: Approval and Owner Age

This figure uses applicant data from Lender A, Lender B, and the Platform to show the how approval percentages change with age (N = 1,027,837). This is a binscatter below and above the age 45 cutoff with 20 equal-sized age bins. The coefficients and standard errors for the owner age coefficient are reported for either side of the discontinuity.

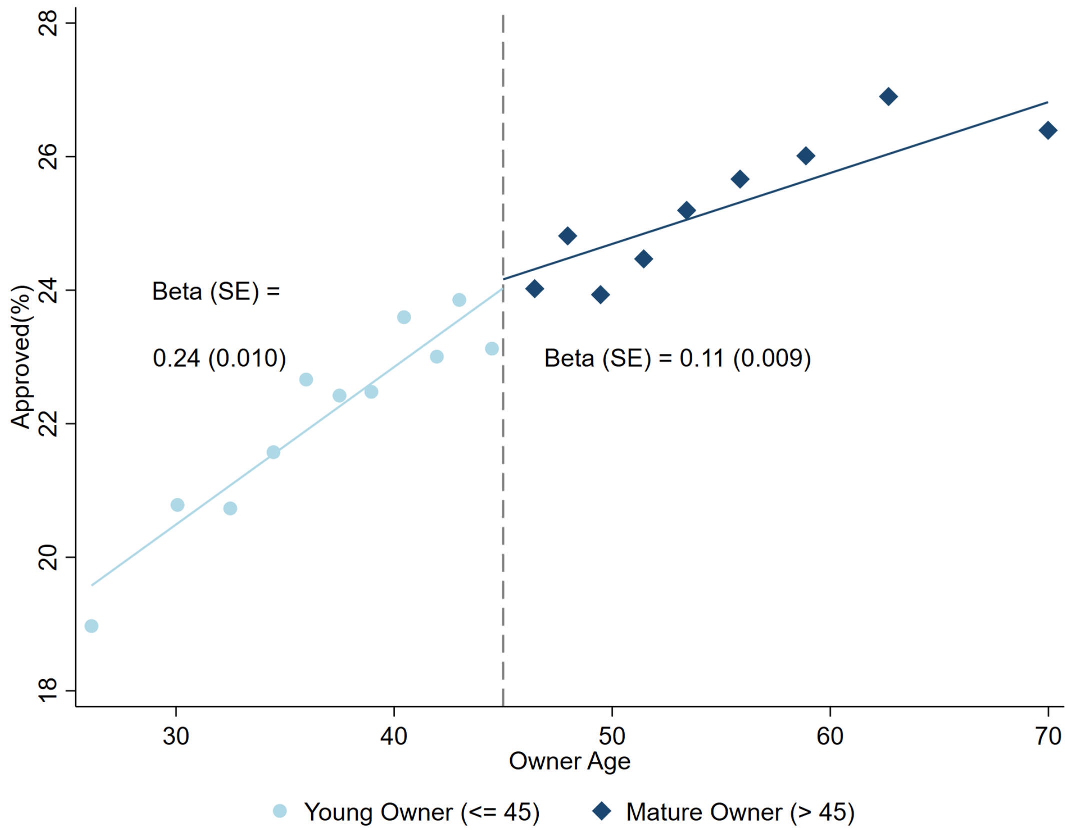

Figure 4: Performance of Loan Default Prediction Models

This figure uses data from Lender A and Lender B on originated loans (N = 38,021) to report the performance of the baseline (FICO) and Cash Flow (including bank statement variables) random forest models in predicting loan defaults on the test dataset, presented in Table 4. In Figure A, we plot the ROC curve. In Figure B and Figure C, we plot the receiver operating characteristic (ROC) curve for Young (< 45) and Mature (≥ 45) Owners individually.

(A) Overall ROC Curves

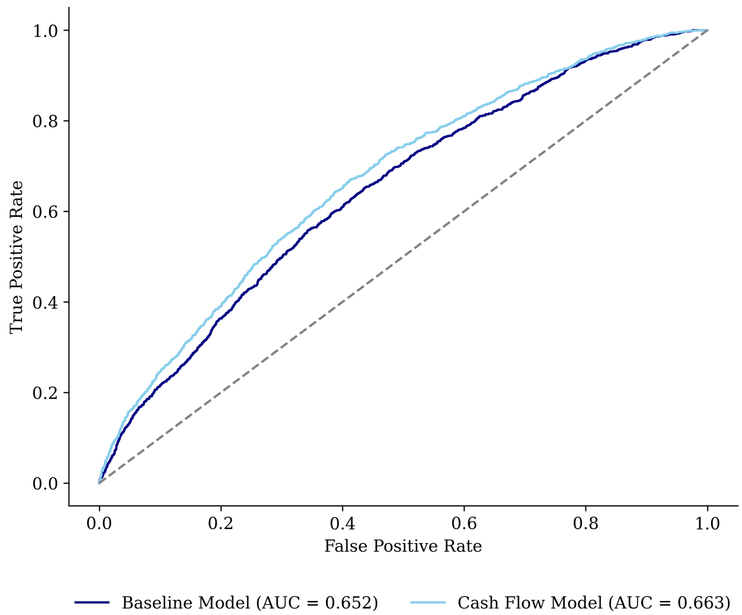

(B) Young Owners ROC Curves

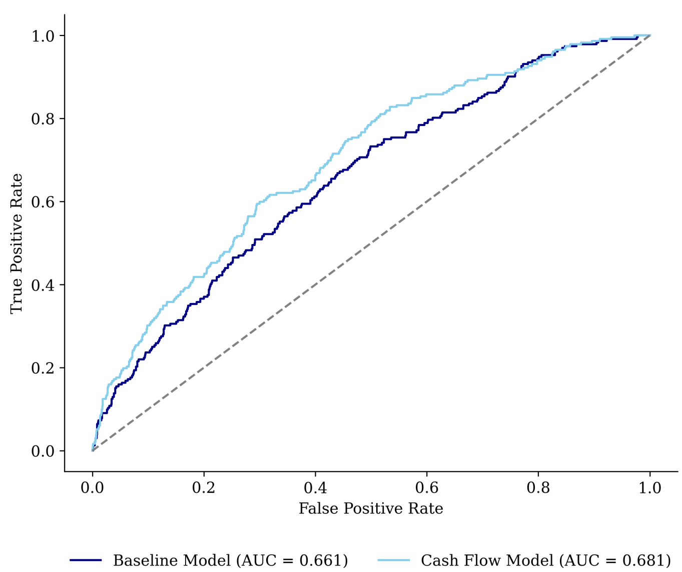

(C) Mature Owners ROC Curves

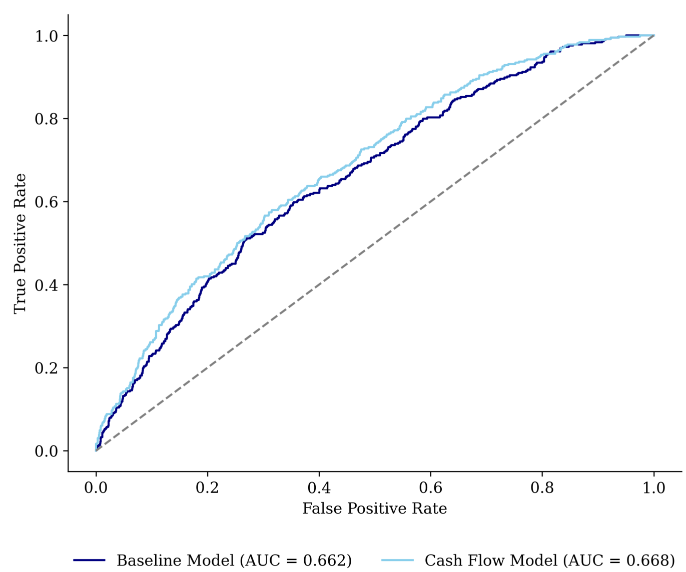

Figure 5: Feature Importance with Random Forest ML Model

This figure shows the importance of applicant features in random forest prediction models, restricting to CF and FICO features using data from Lender A and Lender B on originated loans (N = 38,021) to present the relative importance of CF and FICO in predicting loan defaults. Higher numbers indicate greater importance.

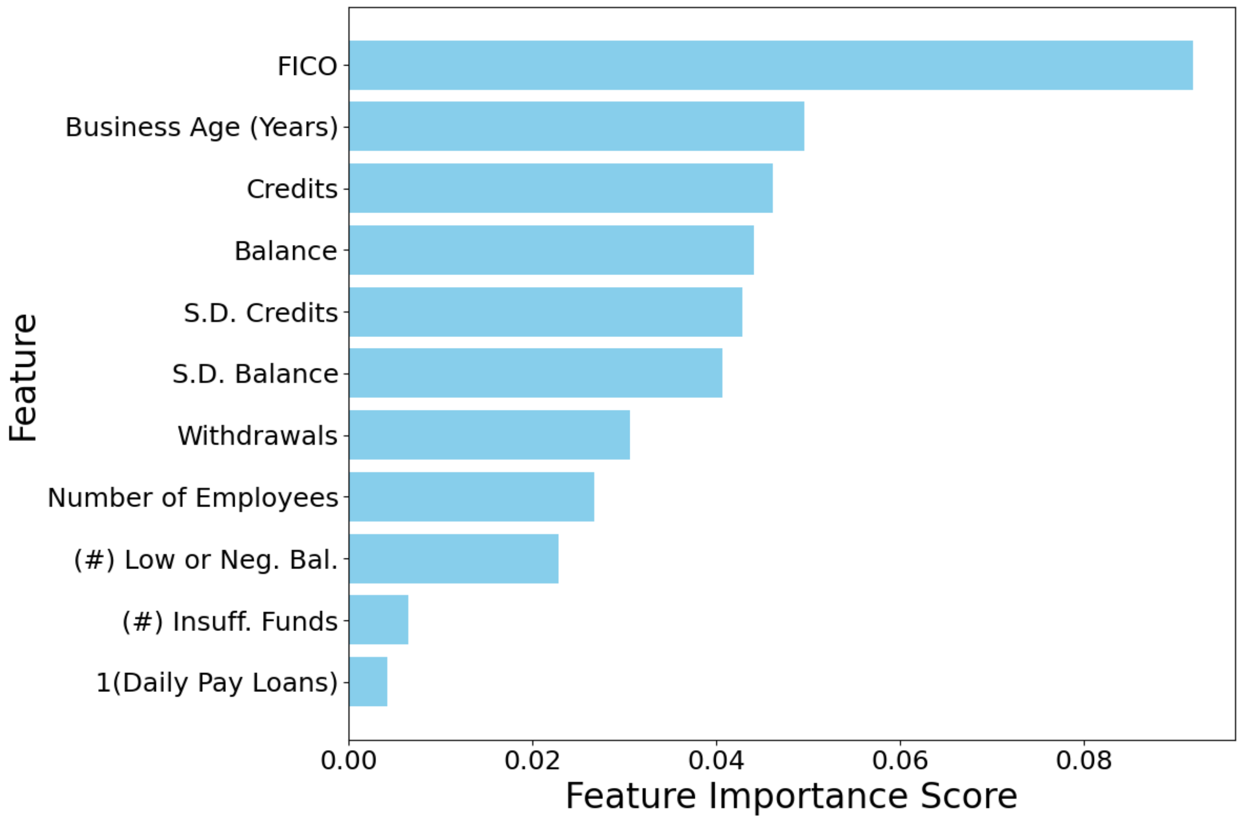

Figure 6: Comparison of Benefit from Cash Flow Model vs. Baseline Model by Age Group

This figure plots the Tail Analysis for Comparative Outcomes (TACO) ratios for young owners from the first set of estimates in Table 7 (see table notes and Section 5 for details), using 10%, 20%, 25%, and 33% respectively to define the tails. This analysis is bootstrapped 1,000 times and reports 95% confidence intervals around the TACO ratio.

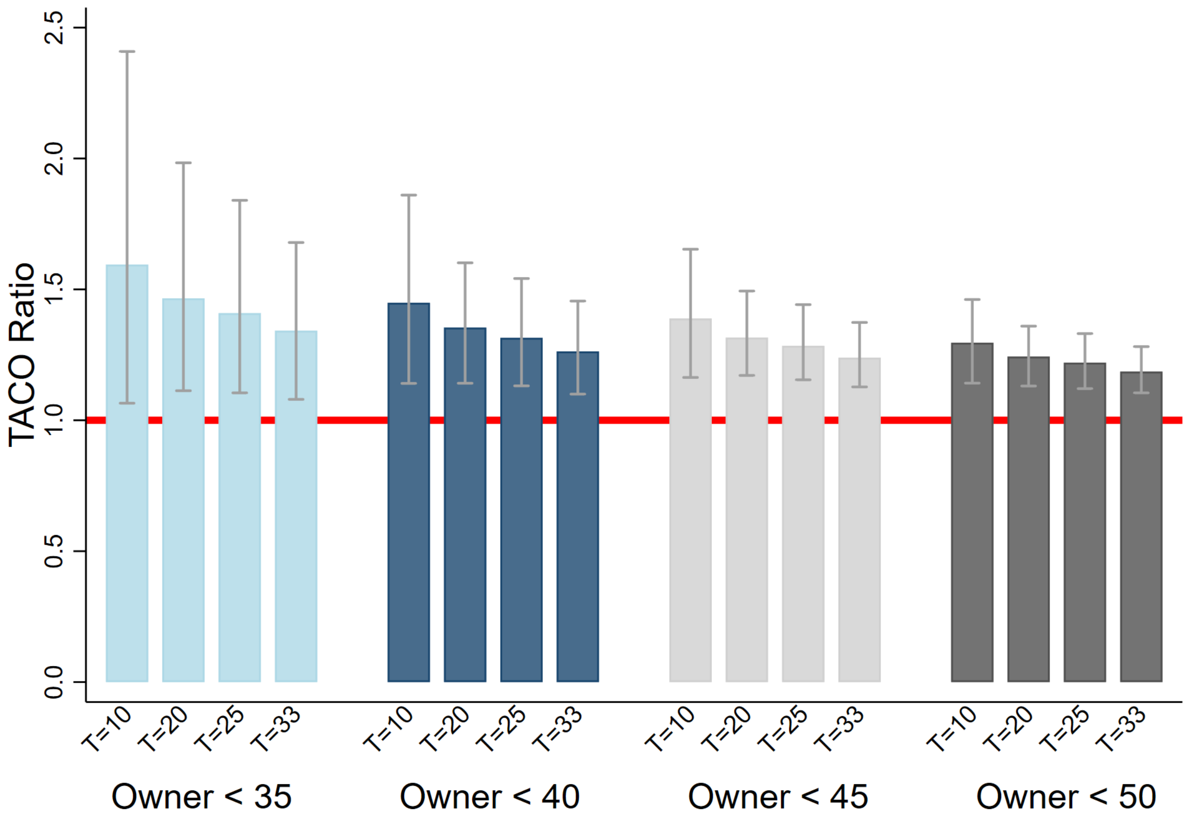

Table 1: Summary Statistics

This table contains summary statistics about loan applications and originated loans from the three sources of data (Platform company and Lenders A & B).

	Applications (Platform)				Applications (Lender A & B)				Borrowers (Lender A & B)			
	N	Mean	Median	SD	N	Mean	Median	SD	N	Mean	Median	SD
Loan Variables:												
Approved (%)	904,471	21			162,818	46			38,021	100		
Requested Loan Amount (Th$)	904,001	151	100	193	162,816	104	75	100	38,021	109	80	98
APR (%)	146,779	88	80	57	69,017	19	18	6.44	36,252	16	15	5.11
Originated (%)	904,471	3.96			162,818	27			38,021	100		
Non-Performing Loan (%)									38,021	17		
Originated Loan Amount (Th$)									38,021	115	78	104
Loan Maturity (Years)									38,021	3.15	3.00	1.41
Credit Score & Cash Flow (Bank Statement) Variables:												
FICO	904,471	681	679	66	162,818	724	725	60	38,021	728	726	49
Credits (Th$)	904,471	99	47	153	162,818	109	52	165	38,021	131	73	170
Balance (Th$)	904,471	33	11	64	162,818	35	13	63	38,021	42	20	64
(#) Insuff. Funds	904,471	0.53	0.00	1.43	162,818	0.27	0.00	1.05	38,021	0.05	0.00	0.25
(#) Low or Neg. Bal.	904,471	0.56	0.00	1.39	162,818	1.76	0.00	3.40	38,021	0.42	0.00	1.06
Withdrawals (Th$)	904,469	99	48	154	70,208	150	73	210	20,353	168	93	209
Credits (less new debt) (Th$)	904,471	83	39	129	58,100	165	90	200	19,370	169	97	195
$\mathbb{1}$(Daily Pay Loan)	903,228	0.35			131,949	0.14			23,950	0.06		
S.D. Credits (Th$)					162,818	12	6.42	14	38,021	13	7.63	14
S.D. Balance (Th$)					162,818	5.18	2.55	6.21	38,021	5.86	3.39	6.15
Borrower Age:												
Owner Age	904,189	44	43	11	123,648	48	47	12	20,190	49	48	11
Young Owner (< 35)	904,189	0.23			123,648	0.14			20,190	0.12		
Young Owner (< 40)	904,189	0.40			123,648	0.28			20,190	0.25		
Young Owner (< 45)	904,189	0.58			123,648	0.43			20,190	0.40		
Young Owner (< 50)	904,189	0.72			123,648	0.58			20,190	0.55		
Other Borrower Characteristics:												
Female	881,418	0.29			160,727	0.23			37,820	0.23		
Business Age (Years)	904,471	6.99	4.83	7.00	162,818	10	7.00	7.65	38,021	11	7.71	7.64
Young Firm (< 5)	904,471	0.52			162,818	0.32			38,021	0.26		
Pct Black Pop (%)	904,471	14	6.50	18	162,818	13	5.60	17	38,021	11	5.30	16
High Pct Black Pop (> 6%)	904,471	0.52			162,818	0.48			38,021	0.46		
Number of Employees	904,471	7.09	4.00	10	162,818	8.88	5.00	12	38,021	10	6.00	13

Table 2: Summary Statistics by Owner Age

This table compares business owners who are young (<45 years old) with those who are more mature (≥ 45). The Loan Applicants sample uses data from Lender A, Lender B, and the Platform (N = 1,067,289). The Originated Loans sample uses data from Lender A and Lender B (N = 38,021).

	Loan Applicants								Originated Loans							
	Young Owner (≤ 45)				Mature Owner (> 45)				Young Owner (≤ 45)				Mature Owner (> 45)			
	N	Mean	Median	SD	N	Mean	Median	SD	N	Mean	Median	SD	N	Mean	Median	SD
Loan Variables:																
Approved (%)	575,767	22			452,070	25			8,030	100			12,160	100		
Requested Loan Amount (Th$)	575,459	138	77	181	451,916	149	100	188	8,030	78	65	62	12,160	78	70	61
APR (%)	95,477	83	75	60	87,024	65	55	54	7,253	15	15	5.10	11,168	15	15	4.72
Originated (%)	575,767	5.29			452,070	6.82			8,030	100			12,160	100		
Non-Performing Loan (%)									8,030	15.5			12,160	15.5		
Originated Loan Amount (Th$)									8,030	74	59	58	12,160	77	60	59
Loan Maturity (Years)									8,030	2.48	2.08	1.20	12,160	2.51	2.08	1.21
Credit Score & Cash Flow (Bank Statement) Variables:																
FICO	575,767	679	677	66	452,070	696	694	67	8,030	733	731	48	12,160	738	735	49
Credits (Th$)	575,767	92	44	145	452,070	105	51	160	8,030	106	62	136	12,160	114	68	143
Balance (Th$)	575,767	30	10	60	452,070	36	12	67	8,030	37	17	59	12,160	37	17	58
(#) Insuff. Funds	575,767	0.52	0.00	1.41	452,070	0.49	0.00	1.39	8,030	0.06	0.00	0.29	12,160	0.07	0.00	0.32
(#) Low or Neg. Bal.	575,767	0.76	0.00	1.91	452,070	0.74	0.00	1.93	8,030	0.48	0.00	1.19	12,160	0.45	0.00	1.13
Withdrawals (Th$)	535,475	93	45	147	400,099	108	52	163	1,081	125	59	193	1,658	126	65	180
Credits (less new debt) (Th$)	529,566	78	37	124	393,902	91	44	139	656	169	89	203	1,100	157	90	180
$\mathbb{1}$(Daily Pay Loan)	574,018	0.33			449,610	0.33			7,907	0.04			11,912	0.05		
S.D. Credits (Th$)	53,221	11	5.43	13	70,427	12	6.34	14	8,030	13	7.49	13	12,160	13	7.89	14
S.D. Balance (Th$)	53,221	4.48	2.01	5.86	70,427	4.95	2.37	6.12	8,030	5.49	3.09	5.99	12,160	5.49	3.13	5.97
Borrower Characteristics:																
Owner Age	575,767	37	37	5.49	452,070	55	53	7.25	8,030	38	38	5.00	12,160	56	55	7.51
Female	559,201	0.28			443,855	0.28			8,006	0.23			12,135	0.23		
Business Age (Years)	575,767	5.22	3.92	4.83	452,070	10	7.00	8.55	8,030	7.29	7.00	5.90	12,160	13	15	7.96
Young Firm (< 5)	575,767	0.62			452,070	0.35			8,030	0.45			12,160	0.21		
Pct Black Pop (%)	575,767	14	6.90	19	452,070	13	5.80	18	8,030	12	5.90	16	12,160	12	5.40	17
High Pct Black Pop (> 6%)	575,767	0.54			452,070	0.49			8,030	0.49			12,160	0.47		
Number of Employees	575,767	6.57	4.00	10	452,070	8.07	4.00	11	8,030	8.50	4.00	11	12,160	10	7.00	12

Table 3: Predicting Default in Regressions

This table uses data from Lender A and Lender B (N = 38,021) on originated loans and their performance to show how credit score, cash flow, and borrower characteristics predict default. The dependent variable is an indicator for the loan being non-performing. The model is logit in columns 5 and 6, and is OLS in all other columns. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Is Non-Performing (%)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
FICO	-4.68***	-4.29***	-4.37***	-4.11***	-0.33***	-0.32***	-3.10***
	(0.21)	(0.22)	(0.21)	(0.22)	(0.02)	(0.02)	(0.22)
Credits (Th$)		-3.04***		-2.72***		-0.23***	-1.73***
		(0.36)		(0.36)		(0.03)	(0.37)
Withdrawals (Th$)		1.85***		1.14***		0.10***	-3.06***
		(0.33)		(0.33)		(0.03)	(0.36)
Balance (Th$)		-1.21***		-1.23***		-0.11***	-0.75**
		(0.33)		(0.32)		(0.03)	(0.32)
$\mathbb{1}$(Daily Pay Loan)		1.15***		0.97***		0.06***	0.77***
		(0.25)		(0.24)		(0.01)	(0.25)
(#) Low or Neg. Bal.		1.53***		1.58***		0.10***	1.71***
		(0.23)		(0.22)		(0.01)	(0.23)
(#) Insuff. Funds		0.06		0.05		0.01	0.14
		(0.21)		(0.21)		(0.01)	(0.21)
S.D. Credits (Th$)		1.85***		2.16***		0.17***	1.96***
		(0.34)		(0.33)		(0.03)	(0.33)
S.D. Balance (Th$)		1.12***		1.29***		0.10***	0.78**
		(0.35)		(0.33)		(0.03)	(0.34)
Business Age (Years)			-0.17***	-0.17***	-0.01***	-0.01***	-0.16***
			(0.03)	(0.03)	(0.00)	(0.00)	(0.03)
Number of Employees			-0.07***	-0.05***	-0.01***	-0.00***	-0.04**
			(0.01)	(0.02)	(0.00)	(0.00)	(0.02)
Requested Loan Amount							0.94
							(0.60)
Originated Loan Amount							4.52***
							(0.69)
Loan Maturity (Years)							2.72***
							(0.17)
APR (%)							0.82***
							(0.06)
Observations	38,021	38,021	38,019	38,019	37,995	37,995	36,252
Industry FE	No	No	Yes	Yes	Yes	Yes	Yes
State FE	No	No	Yes	Yes	Yes	Yes	Yes
Quarter FE	No	No	Yes	Yes	Yes	Yes	Yes
Loan Type FE	No	No	Yes	Yes	Yes	Yes	Yes
R-squared	0.016	0.021	0.038	0.043			0.073
Pseudo R-squared					0.0446	0.0507	
Y-mean	16.72	16.72	16.72	16.72	16.73	16.73	17.30

Table 4: Performance Comparison of Baseline and Cash Flow Machine Learning Models

This table presents our performance evaluation of the Baseline and Cash Flow random forest models for predicting loan default and approval. For variables in each model, see Table A.5. The Minimal and Preferred Specification for default use data from Lender A and Lender B (N = 38,021) on originated loans where loan performance is available. The Preferred Specification for approval uses data from Lender A, Lender B, and the Platform (N = 1,067,289) on loan applications. Performance metrics are calculated as the mean of 100 bootstrap iterations. Definitions of the performance metrics *ROC AUC*, *AUC-PR* and *H-measure* are provided in section 4.2; larger numbers indicate better predictive performance. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

	ROC AUC	**AUC-PR**	**H-Measure**
Minimal Specification (Borrower Default):			
FICO Model	0.624	0.244	0.057
Cash Flow Model	0.642	0.259	0.072
Difference	0.018***	0.015***	0.015***
Preferred Specification (Borrower Default):			
FICO Model	0.652	0.266	0.081
Cash Flow Model	0.663	0.279	0.092
Difference	0.011***	0.013***	0.011***
Preferred Specification (Loan Approval):			
FICO Model	0.820	0.635	0.324
Cash Flow Model	0.850	0.685	0.388
Difference	0.030***	0.050***	0.064***

Table 5: Predicting Default in Regressions by Age Group

This table uses data from Lender A and Lender B (N = 38,021) to show how credit score, cash flow, and borrower characteristics predict default conditional on origination. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Is Non-Performing (%)			
	Owner (< 45)		Owner (> 45)	
	(1)	(2)	(3)	(4)
FICO	-4.07***	-3.78***	-3.82***	-3.59***
	(0.41)	(0.41)	(0.32)	(0.33)
Credits (Th$)		-2.61***		-1.60***
		(0.66)		(0.57)
Withdrawals (Th$)		4.27***		2.87***
		(1.27)		(0.92)
Balance (Th$)		-2.04***		-1.15**
		(0.53)		(0.48)
𝟙(Daily Pay Loan)		0.75**		0.73**
		(0.37)		(0.32)
(#) Low or Neg. Bal.		1.81***		1.26***
		(0.46)		(0.41)
(#) Insuff. Funds		-0.20		-0.14
		(0.38)		(0.31)
S.D. Credits (Th$)		1.85***		1.59***
		(0.61)		(0.48)
S.D. Balance (Th$)		1.39**		0.62
		(0.56)		(0.55)
Business Age (Years)	-0.17***	-0.15**	-0.17***	-0.17***
	(0.06)	(0.06)	(0.04)	(0.04)
Number of Employees	-0.06*	-0.04	-0.08***	-0.07**
	(0.03)	(0.04)	(0.03)	(0.03)
Observations	8,027	8,027	12,158	12,158
Industry FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Loan Type FE	Yes	Yes	Yes	Yes
R-squared	0.051	0.060	0.048	0.052
Y-mean	15.51	15.51	15.47	15.47

Table 6: Machine Learning Model Performance by Age Group

This table presents our performance evaluation of the Baseline and Cash Flow random forest models for predicting loan default for different owner age groups based on the birth date of the primary owner or CEO. We present results from the Preferred Specification, for variables in this model, see Table A.5. This table uses data from Lender A and Lender B (N = 38,021) on originated loans where loan performance is available. Performance metrics are calculated as the mean of 100 bootstrap iterations. Definitions of the performance metrics *ROC AUC*, *AUC-PR* and *H-Measure* are provided in Section 4.2; larger numbers indicate better predictive performance. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

	ROC AUC	**AUC-PR**	**H-Measure**
Under 35:			
FICO Model	0.622	0.238	0.087
Cash Flow Model	0.649	0.272	0.119
Difference	0.027***	0.034***	0.032***
Over 35:			
FICO Model	0.666	0.264	0.100
Cash Flow Model	0.677	0.279	0.112
Difference	0.011***	0.015***	0.012***
Under 40:			
FICO Model	0.651	0.247	0.097
Cash Flow Model	0.677	0.273	0.125
Difference	0.026***	0.026***	0.028***
Over 40:			
FICO Model	0.664	0.266	0.100
Cash Flow Model	0.672	0.280	0.110
Difference	0.008***	0.014***	0.009***
Under 45:			
FICO Model	0.661	0.256	0.100
Cash Flow Model	0.681	0.279	0.123
Difference	0.020***	0.024***	0.023***
Over 45:			
FICO Model	0.662	0.265	0.100
Cash Flow Model	0.668	0.278	0.109
Difference	0.007***	0.012***	0.008***
Under 50:			
FICO Model	0.658	0.259	0.095
Cash Flow Model	0.673	0.278	0.112
Difference	0.015***	0.019***	0.017***
Over 50:			
FICO Model	0.666	0.265	0.106
Cash Flow Model	0.674	0.279	0.117
Difference	0.009***	0.014***	0.010***

Table 7: TACO Results on Benefit of Cash Flow vs Baseline Model by Age Group

This table shows the results from implementing Tail Analysis for Comparative Outcomes (TACO, see Section 5). We compare two random forest models to predict default: a Baseline model (containing FICO, firm size, firm age, and industry among others) and the Cash Flow (CF) model, which adds bank statement variables to the Baseline model. The table uses data from Lender A and Lender B (N = 38,021) on originated loans. The observation counts represent the sum across 1,000 bootstrap holdout samples. The first two columns ("Tails") show the group's share in the decile tails population of bootstrapped sample observations. The next two columns restrict to the 10% of each bootstrap sample with the highest increase in default chance between the Baseline and the CF model, who are thus adversely affected by switching from the Baseline to the CF model. The mean shows the share of young owners in this group, which can be compared to the "Tails" mean column. The next two columns show the same metric for the bottom 10% (the group that most benefits from switching to the CF model). The last column shows the ratio between the "Hurt" and "Benefit" means, which we call the TACO ratio. A ratio of one implies no implication of switching models, a ratio less than one implies that the group is adversely affected, and a ratio greater than one implies that the group benefits. We calculate standard errors for the TACO ratio using the percentile bootstrap.

	Tails		Top 10% Default Increases w/ CF Model (Hurt)		Bottom 10% Default Increases w/ CF Model (Benefit)		
	N	Mean	N	Mean	N	Mean	TACO Ratio
Full Sample:							
Young Owner (< 35)	736,507	0.130	384,203	0.101	352,304	0.161	1.595**
Young Owner (< 40)	736,507	0.270	384,203	0.222	352,304	0.322	1.449***
Young Owner (< 45)	736,507	0.418	384,203	0.352	352,304	0.489	1.389***
Young Owner (< 50)	736,507	0.578	384,203	0.506	352,304	0.656	1.296***
Low FICO (< 700):							
Young Owner (< 35)	736,507	0.056	384,203	0.013	352,304	0.103	7.778***
Young Owner (< 40)	736,507	0.110	384,203	0.029	352,304	0.199	6.926***
Young Owner (< 45)	736,507	0.169	384,203	0.044	352,304	0.305	6.894***
Young Owner (< 50)	736,507	0.226	384,203	0.062	352,304	0.404	6.496***
High FICO (≥ 700):							
Young Owner (< 35)	736,507	0.073	384,203	0.088	352,304	0.058	0.660
Young Owner (< 40)	736,507	0.160	384,203	0.193	352,304	0.123	0.633**
Young Owner (< 45)	736,507	0.249	384,203	0.308	352,304	0.185	0.599***
Young Owner (< 50)	736,507	0.352	384,203	0.444	352,304	0.251	0.566***

Table 8: Effect of Assignment to Cash Flow-Intensive Lender (Within-Application)

This table uses data from the Platform to test whether cash flow-intensive lenders are more likely to approve young entrepreneurs (N = 879,889). The level of observation is the application-lender, so an applicant may appear multiple times as the application is forwarded to multiple lenders. Application fixed effects mean that covariates such as entrepreneur age, industry, and time are not identified. Young is defined according to the column header and only applicants that fall below that threshold or are 50 years or older are included in the regressions. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood for young applicants forwarded to cash flow-intensive lenders relative to applicants over 50 and those forwarded to non-cash flow-intensive lenders. In odd (even) columns, the sample is restricted to applicants with a FICO score above (below) 700. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)							
	Young (< 35)		Young (< 40)		Young (< 45)		Young (< 50)	
FICO:	High	Low	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1	0.86*	2.16***	1.11***	2.32***	1.26***	2.13***	1.34***	1.99***
	(0.47)	(0.38)	(0.40)	(0.33)	(0.37)	(0.31)	(0.35)	(0.30)
Observations	159,582	258,687	209,323	355,366	261,661	448,378	306,295	522,477
Application FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.320	0.316	0.322	0.318	0.323	0.316	0.323	0.316
Y-mean	24.01	16.81	24.16	16.94	24.42	16.99	24.53	17.08

Table 9: Effect of Random Assignment to Cash Flow-Intensive Loan Officer

This table uses data from Lender B to test whether young entrepreneurs are more likely to have their loan application approved when randomly assigned to a cash flow-intensive Loan Officer (N = 11,535). The level of observation is an application. Young is defined according to the column header and only applicants that fall below that threshold or are 50 years or older are included in the regressions. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood for young applicants assigned to cash flow-intensive loan officers relative to applicants over 50 and those forwarded to non-cash flow-intensive loan officers. In odd (even) columns, the sample is restricted to applicants with a FICO score above (below) the median. Standard errors are clustered by quarter and approver. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)							
	Young (< 35)		Young (< 40)		Young (< 45)		Young (< 50)	
FICO:	High	Low	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1	3.80	13.34**	4.69	8.70**	3.94	5.39*	3.29	3.28
	(5.35)	(5.60)	(3.35)	(3.75)	(2.94)	(3.05)	(2.49)	(2.75)
Young=1	-6.74*	-9.47**	-5.82***	-5.79**	-4.55**	-4.57*	-3.51**	-3.27
	(4.02)	(4.03)	(2.22)	(2.72)	(2.01)	(2.31)	(1.64)	(2.11)
Observations	3,319	3,231	4,023	3,947	4,884	4,842	5,782	5,753
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.075	0.067	0.070	0.065	0.070	0.062	0.068	0.059
Y-mean	55.80	48.28	55.65	48.39	55.63	48.14	55.76	48.20

Appendix (For Online Publication)

A Supplemental Tables and Figures

Figure A.1: Relationship between FICO and Demographics Nationally

This figure uses National data from Experian as of 2019 Q2 on the average FICO for each reported age.

Credit Scores by Age		
Age Range	Average FICO® Score Q2 2015	Average FICO® Score Q2 2019
23 to 29*	660	660
30 to 39	652	672
40 to 49	667	683
50 to 59	688	703
60 to 69	718	733
70 to 79	745	754
80 to 89	755	757
90 to 99	754	753

Source: Experian. *In Q2 2015, data was only available for ages 23 to 29; we compared to the same age range in Q2 2019.

Figure A.2: Applications by Source

This figure shows the number of applications over time for each of the three data sources. Panel A uses data from Lender A (N = 104,150), Panel B uses data from Lender B (N = 58,668), and Panel C uses data from the Platform (N = 904,471).

(A) Lender A

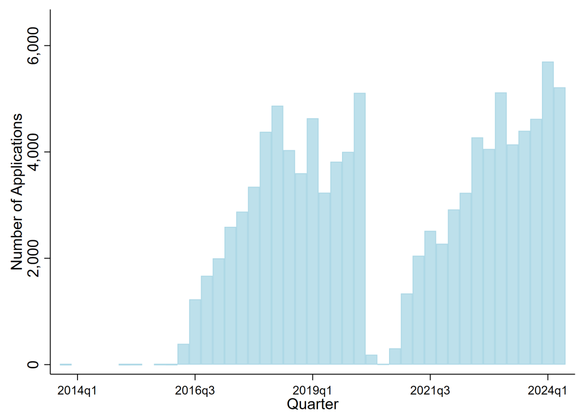

(B) Lender B

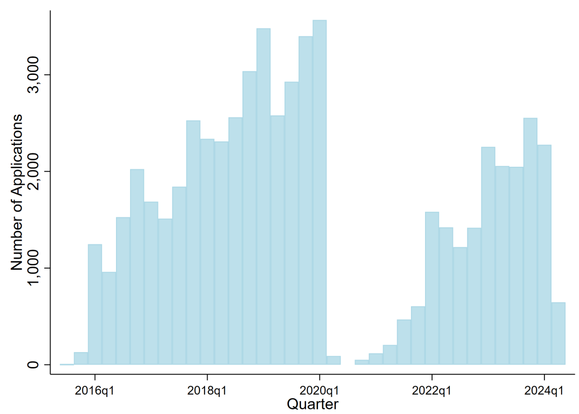

(C) Platform

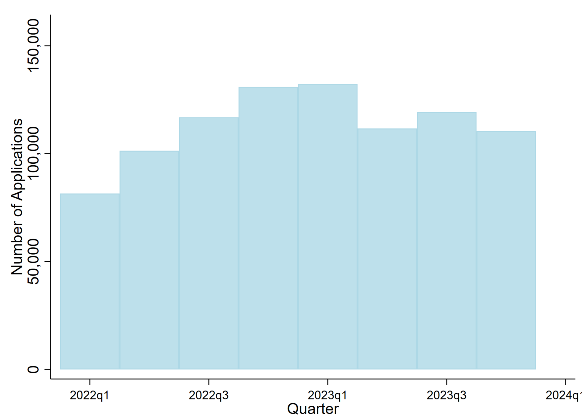

Figure A.3: Entrepreneur Age Distributions

This figure shows the distribution of age among owners in the data (each firm is identified as having one primary CEO or owner). Panel A uses data from Lender A, Lender B, and the Platform (N = 1,027,837) on loan applications without missing age. Panel B uses data from Lender A and Lender B (N = 20,190) on originated loans without missing age.

(A) Full Sample

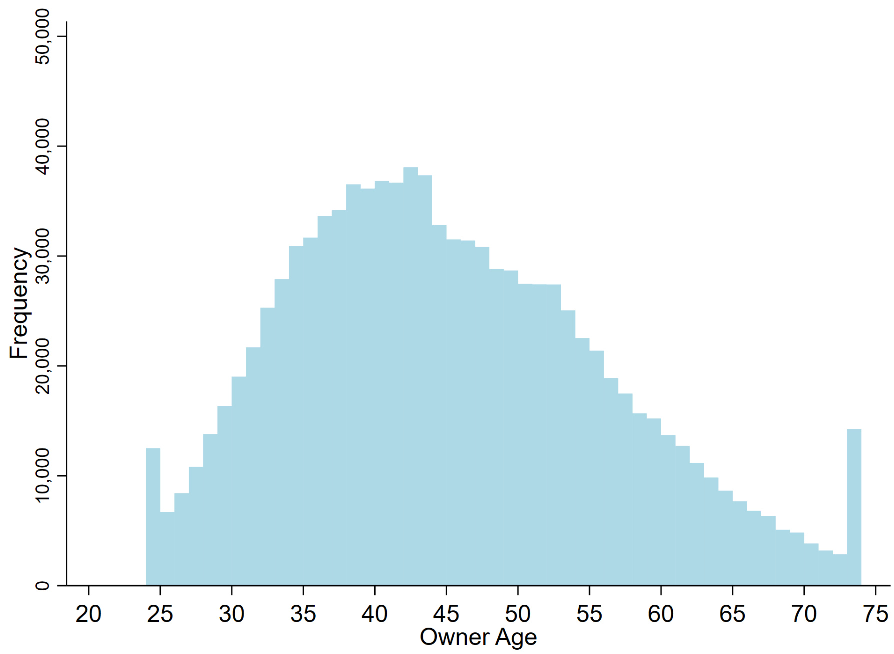

(B) Lender A and B Sample

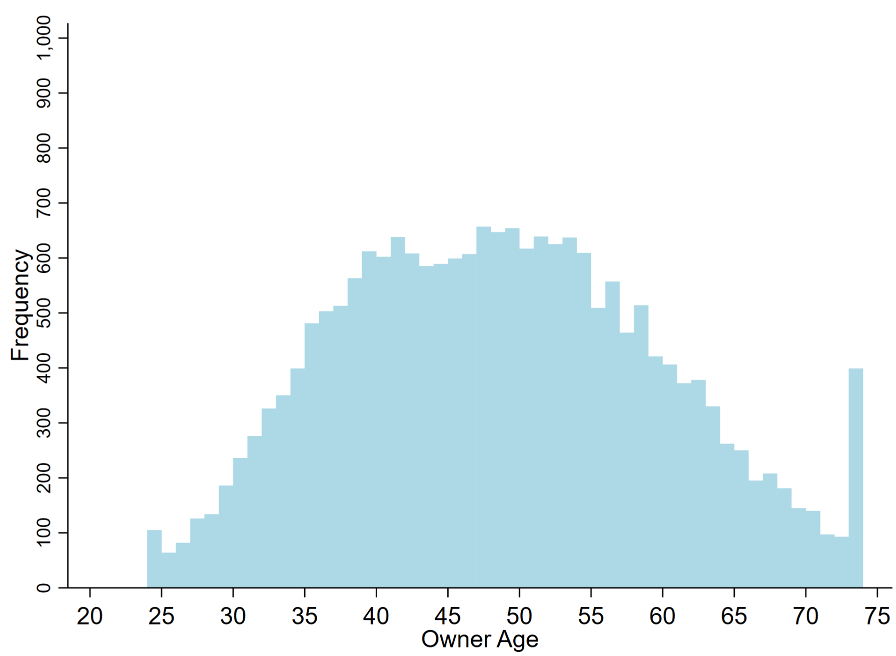

Figure A.4: Bootstrap Estimates of Differences in Predicting Default

This figure presents differences in performance metrics between the Baseline and Cash Flow random forest models for predicting loan default. We present results from the Preferred Specification; for variables in this model, see Table A.5. This table uses data from Lender A and Lender B (N = 38,021) on originated loans where loan performance is available. Paired performance metrics are calculated for 1,000 bootstrapped iterations and plotted below. Definitions of the performance metrics *ROC AUC* and *H-Measure* are provided in S+ ection 4.2; larger numbers indicate better predictive performance. In the figure, positive numbers (bars above the red line) represent that the CF model has better performance.

(A) Differences in ROC AUC

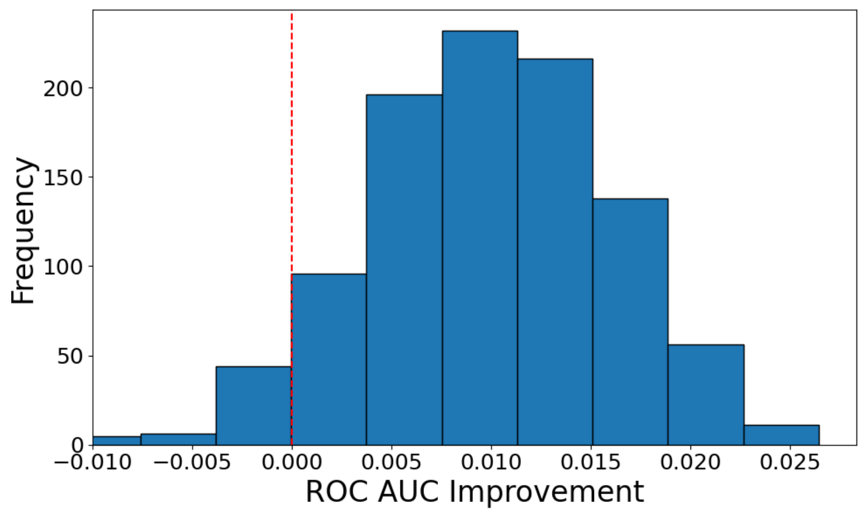

(B) Differences in H-Measure

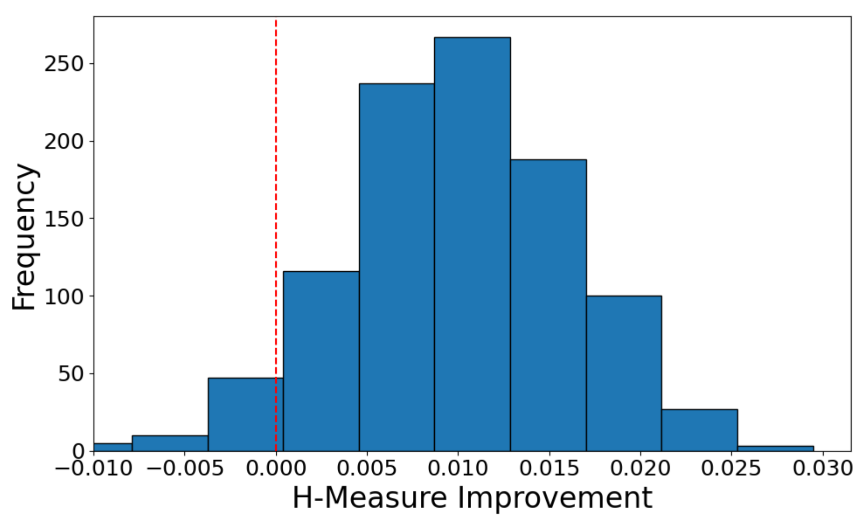

Figure A.5: Histogram of Baseline Model and CF Model Prediction Probabilities

This figure uses data from Lender A and Lender B on originated loans (N = 38,021) to show the results of a TACO analysis to predict loan default. It reports a histogram of predicted default probabilities under the Baseline and CF random forest models for each observation in the holdout sample. This figure combines results from all 1,000 bootstrap iterations.

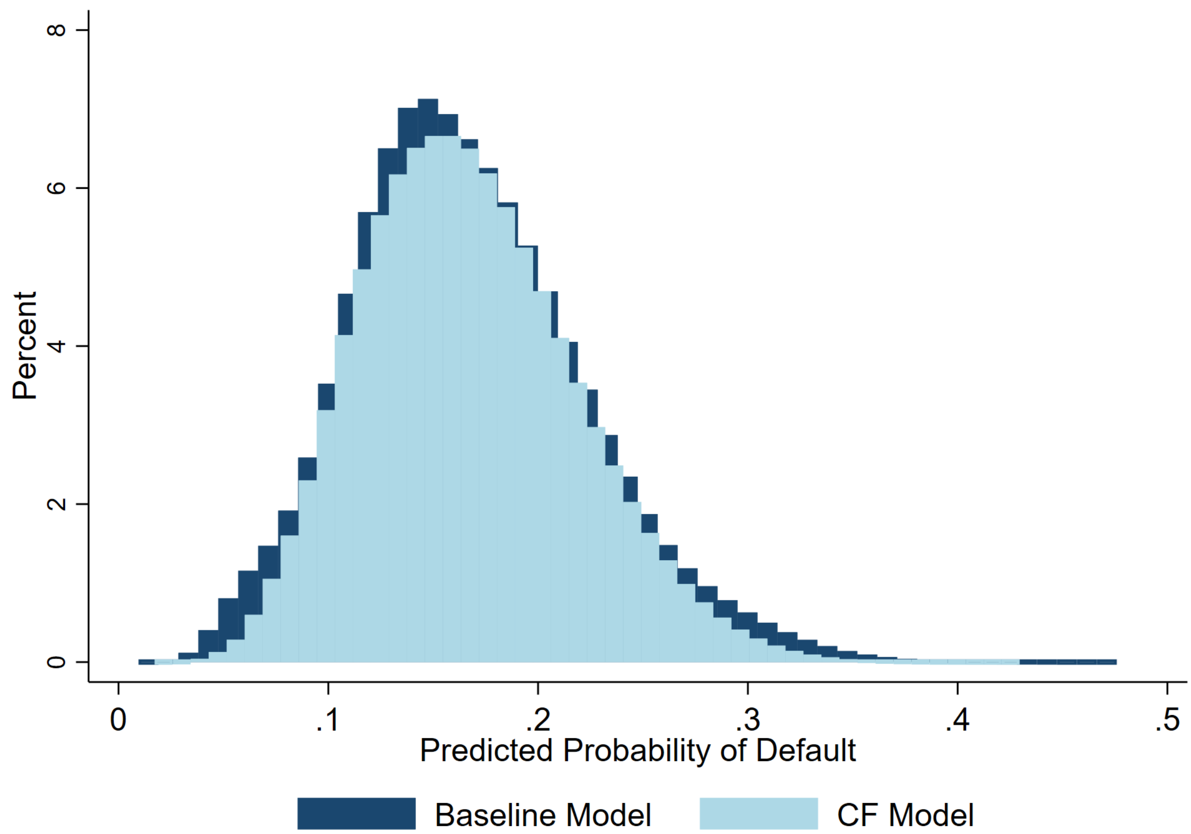

Figure A.6: Default Difference Distributions

This figure uses data from Lender A and Lender B on originated loans (N = 38,021). It reports a histogram of the difference in predicted default probabilities ($h_i = g(X_i) - f(X_i)$) under the Baseline and CF random forest models for each observation in the holdout sample. This difference is defined in more detail in Section 5. This figure combines results from all 1,000 bootstrap iterations. It plots observations by owner age at each bin with red lines at the 10th and 90th percentile values.

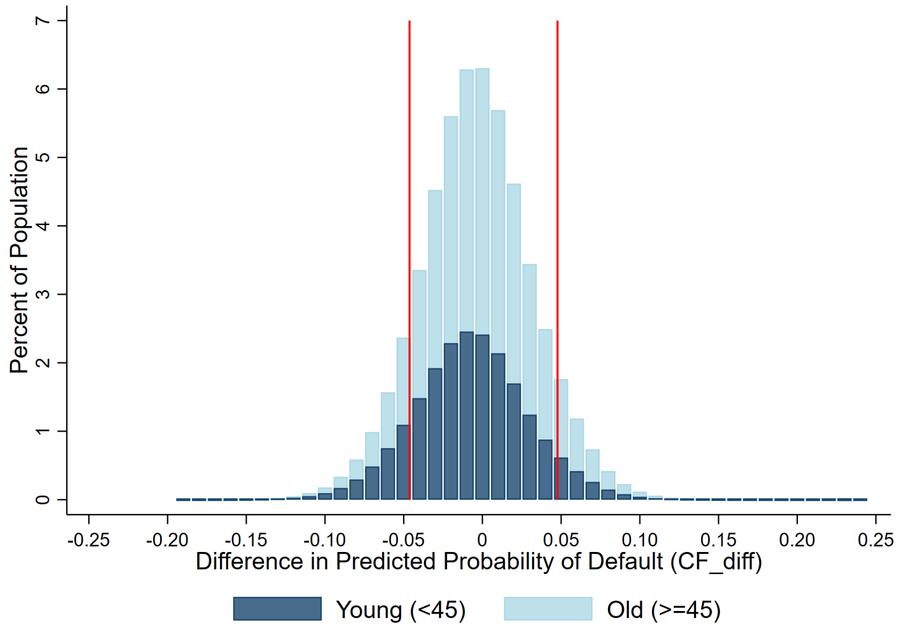

Figure A.7: Placebo Test of TACO Methodology by Age Group

This figure uses data from Lender A and Lender B on originated loans (N = 38,021) to show the results of the Tail Analysis for Comparative Outcomes (TACO) methodology (see Section 5 for details). This is similar to Figure 6 but instead–as a placebo–we compare a CF model to an otherwise identical CF model, varying only the randomized split between the training and testing (holdout) datasets across each of the 100 bootstrap iterations. We use the Preferred Model specification and a 10% threshold to define the tails, and plot the TACO ratios and 95% confidence intervals here. The expected result under this placebo test is a TACO ratio of 1.0: the characteristics of the group with the largest reduction in predicted default probability should be identical to the group with the largest increase in predicted default probability.

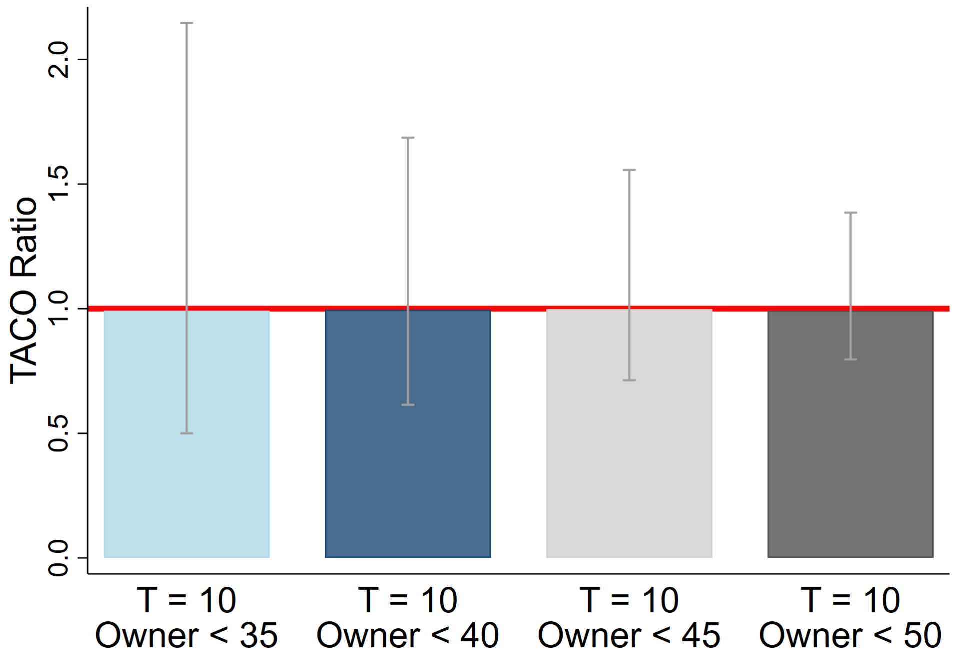

Figure A.8: Comparison of Benefit from Cash Flow Model vs. Baseline Model for Other Characteristics

This figure uses data from Lender A and Lender B on originated loans (N = 38,021) to show the results of the Tail Analysis for Comparative Outcomes (TACO) methodology (see Section 5 for details) for alternative applicant and firm characteristics. This analysis is bootstrapped 1,000 times and reports 95% confidence intervals for the TACO ratios.

(A) Low FICO

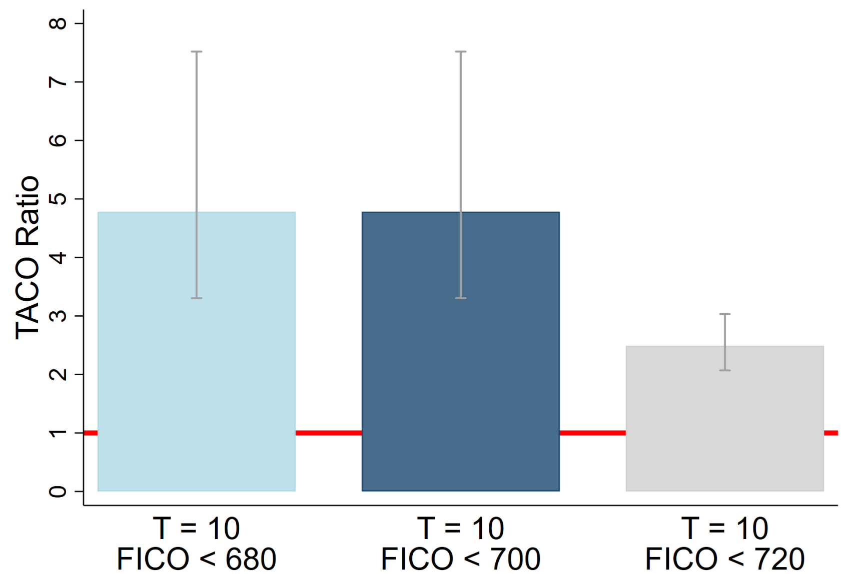

(B) High Local Black Share and Young Firms

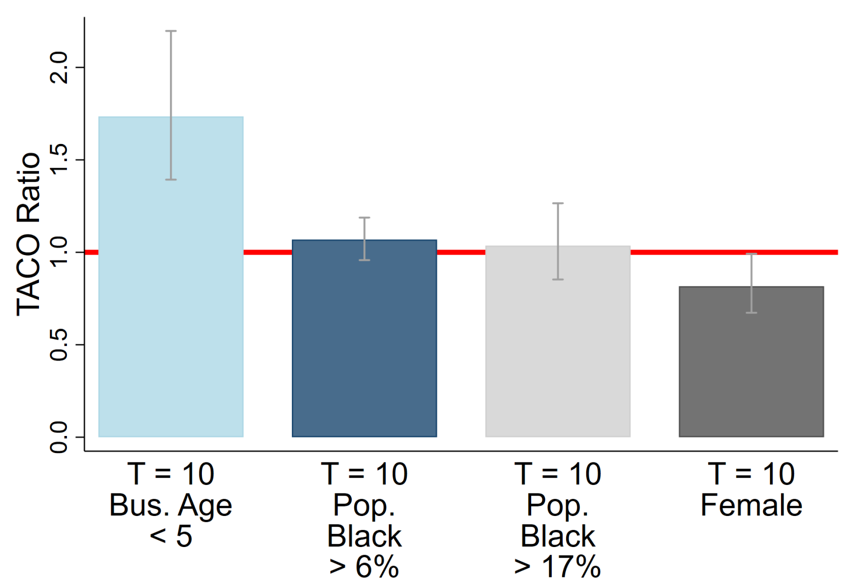

Figure A.9: Performance of Loan Approval Prediction Models

This figure uses data from Lender A, Lender B, and the Platform (N = 1,067,289) to report performance of the baseline (FICO) and Cash Flow (including bank statement variables) random forest models in predicting loan approvals on the test dataset, presented in Table 4. In Figure A, we plot the ROC curve. In Figure B and Figure C, we plot the ROC curve for Young (< 45) and Mature (≥ 45) Owners individually.

(A) Overall ROC Curves

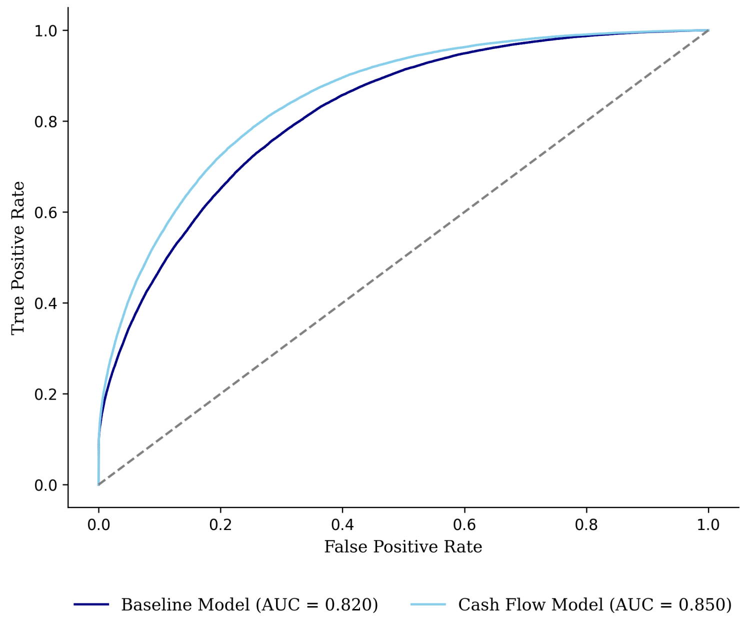

(B) Young Owners ROC Curves

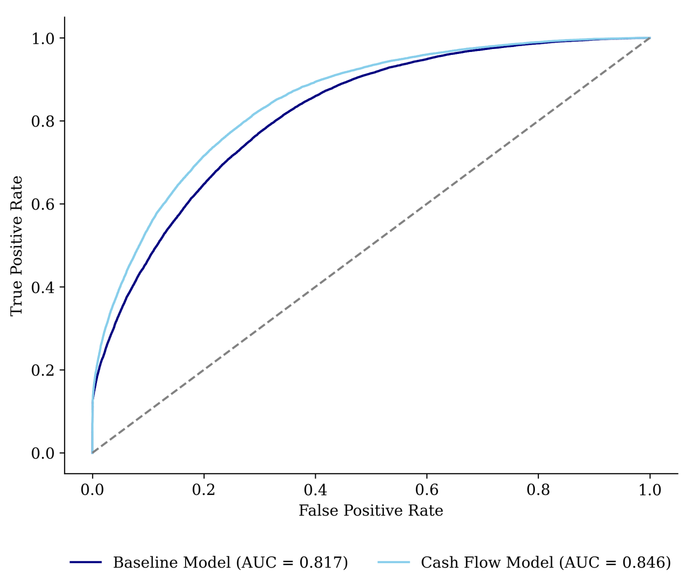

(C) Mature Owners ROC Curves

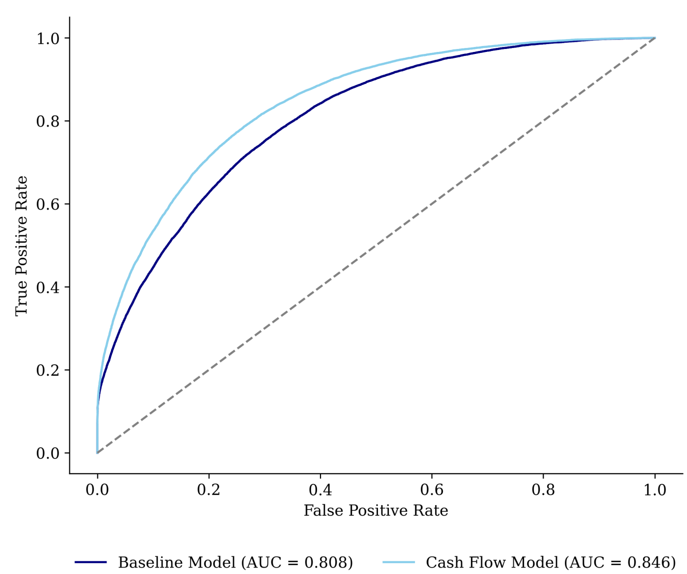

Table A.1: Industry Distribution

This table describes the industry distribution among small businesses nationwide (first column), loan applicants in our data (second set of columns) and borrowers in our data (third set of columns). We also report the share in each industry with young owners (under 45).

		Loan Applicants		**Originated Loans**	
Industry	% National	% Applications	% < 45	% Loans	% < 45
Accomodation and Food Services	7.8	10.7	54.3	8.5	40.1
Agriculture, Forestry, Fishing and Hunting	0.4	1.0	56.6	0.7	34.5
Arts, Entertainment, and Recreation	2.2	3.3	61.4	2.2	45.8
Construction	12.3	20.9	58.6	16.2	42.4
Educational Services	1.4	1.9	50.4	1.2	40.0
Finance and Insurance	4.1	2.1	51.1	1.4	40.6
Health Care and Social Assistance	10.5	8.9	45.4	13.9	37.6
Information	1.3	4.0	57.5	1.3	41.5
Manufacturing	3.3	5.0	48.1	5.0	29.9
Mining, Quarrying, and Oil and Gas Extraction	0.3	0.2	54.2	0.1	37.0
Other Services (except Public Administration)	12.1	0.2	38.8	0.8	43.2
Professional, Scientific, and Technical Services	20.0	3.9	43.6	22.5	38.9
Real Estate and Rental and Leasing	5.8	2.9	53.6	2.3	34.9
Retail Trade	10.6	16.1	60.2	13.3	42.7
Transportation and Warehousing	3.2	14.8	63.1	7.2	44.8
Utilities	0.1	0.6	58.9	0.1	21.7
Wholesale Trade	4.5	3.4	48.1	3.5	32.7

Table A.2: Summary Statistics by Age about Applicants to Lenders A & B

This table prepares summary statistics for Lender A and Lender B (N = 162,818) among loan applicants. It compares business owners who are young (<45 years old) with those who are more mature (≥ 45).

	Young Owner (≤ 45)		Mature Owner (> 45)		
	N	Mean	N	Mean	Difference
Loan Variables:					
Approved (%)	53,221	37	70,427	43	-6.551***
Requested Loan Amount (Th$)	53,221	84	70,425	88	-4.549***
APR (%)	13,540	19	22,221	18	0.814***
Originated (%)	53,221	18	70,427	22	-4.264***
Non-Performing Loan (%)	8,030	16	12,160	15	0.036
Originated Loan Amount (Th$)	9,682	83	15,815	89	-5.297***
Loan Maturity (Years)	9,682	2.79	15,815	2.95	-0.153***
Credit Score & Cash Flow (Bank Statement) Variables:					
FICO	53,221	718	70,427	732	-14.290***
Credits (Th$)	53,221	84	70,427	98	-14.138***
Balance (Th$)	53,221	28	70,427	33	-4.318***
(#) Insuff. Funds	53,221	0.36	70,427	0.30	0.061***
(#) Low or Neg. Bal.	53,221	2.48	70,427	1.92	0.561***
Withdrawals (Th$)	12,931	100	18,456	118	-18.444***
Credits (less new debt) (Th$)	7,020	143	12,259	144	-0.717
$\mathbb{1}$(Daily Pay Loan)	52,173	0.14	68,509	0.14	0.000
S.D. Credits (Th$)	53,221	11	70,427	12	-1.180***
S.D. Balance (Th$)	53,221	4.48	70,427	4.95	-0.467***
Borrower Characteristics:					
Owner Age	53,221	37	70,427	56	-19.114***
Female	52,254	0.25	69,661	0.23	0.011***
Business Age (Years)	53,221	6.53	70,427	12	-5.664***
Young Firm (< 5)	53,221	0.50	70,427	0.23	0.273***
Pct Black Pop (%)	53,221	14	70,427	13	1.042***
High Pct Black Pop (> 6%)	53,221	0.51	70,427	0.47	0.034***
Number of Employees	53,221	7.51	70,427	8.85	-1.348***

Table A.3: Summary Statistics on Low-FICO Originated Loans (Lender A & B)

This figure uses data from Lender A and Lender B on originated loans for business owners who are young (< 45 years old) with those who are more mature (≥ 45) limited to owners with a FICO score at or below 700 (N = 5,220).

	All				Young Owner (≤ 45)		Mature Owner (> 45)		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	5,220	676	680	19	2,219	675	3,001	677	-1.62***
Credits (Th$)	5,220	95	59	115	2,219	93	3,001	97	-4.40
Balance (Th$)	5,220	27	13	45	2,219	28	3,001	25	2.83**
(#) Insuff. Funds	5,220	0.09	0.00	0.36	2,219	0.09	3,001	0.10	-0.01
(#) Low or Neg. Bal.	5,220	0.64	0.00	1.35	2,219	0.62	3,001	0.65	-0.03
Withdrawals (Th$)	581	91	47	125	255	94	326	89	4.28
Credits (less new debt) (Th$)	331	129	80	137	130	141	201	121	20
$\mathbb{1}$(Daily Pay Loan)	5,136	0.05	0.00	0.21	2,185	0.04	2,951	0.05	-0.02***
S.D. Credits (Th$)	5,220	12	7.09	13	2,219	12	3,001	12	-0.05
S.D. Balance (Th$)	5,220	4.73	2.66	5.43	2,219	4.81	3,001	4.68	0.13

Table A.4: Default for First Time Applicants

This table uses data from Lender A and Lender B on originated loans among first time applicants (N = 32,226) and their performance to show how credit score, cash flow, and borrower characteristics predict default. The dependent variable is an indicator for the loan being non-performing. The model is logit in column 3 and OLS in all other columns. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Is Non-Performing (%)					
					Owner (< 45)	Owner (> 45)
	(1)	(2)	(3)	(4)	(5)	(6)
FICO	-4.42***	-4.13***	-0.32***	-3.13***	-3.92***	-3.31***
	(0.22)	(0.29)	(0.02)	(0.28)	(0.57)	(0.30)
Credits (Th$)		-2.81***	-0.24***	-1.91***	-3.06***	-1.71*
		(0.65)	(0.03)	(0.55)	(0.70)	(0.85)
Withdrawals (Th$)		1.43**	0.13***	-2.69***	4.16*	2.58
		(0.57)	(0.03)	(0.66)	(2.26)	(1.86)
Balance (Th$)		-1.57***	-0.14***	-1.13**	-1.92***	-1.37***
		(0.43)	(0.03)	(0.40)	(0.48)	(0.42)
$\mathbb{1}$(Daily Pay Loan)		0.94***	0.06***	0.73**	0.56	0.72**
		(0.29)	(0.01)	(0.27)	(0.33)	(0.26)
(#) Low or Neg. Bal.		1.47***	0.09***	1.59***	1.82***	0.90*
		(0.32)	(0.01)	(0.36)	(0.51)	(0.52)
(#) Insuff. Funds		0.05	0.01	0.12	-0.10	-0.15
		(0.42)	(0.01)	(0.40)	(0.58)	(0.38)
S.D. Credits (Th$)		2.32***	0.18***	1.99***	2.15**	1.85***
		(0.53)	(0.03)	(0.50)	(0.76)	(0.52)
S.D. Balance (Th$)		1.29***	0.10***	0.86***	0.94	0.66*
		(0.19)	(0.03)	(0.22)	(0.62)	(0.36)
Requested Loan Amount				0.62		
				(0.65)		
Originated Loan Amount				4.75***		
				(1.22)		
Loan Maturity (Years)				2.56***		
				(0.32)		
APR (%)				0.84***		
				(0.09)		
Business Age (Years)	-0.18***	-0.18***	-0.01***	-0.17***	-0.13*	-0.20***
	(0.03)	(0.04)	(0.00)	(0.04)	(0.07)	(0.04)
Number of Employees	-0.06***	-0.04**	-0.00**	-0.03	-0.03	-0.05
	(0.01)	(0.02)	(0.00)	(0.02)	(0.04)	(0.04)
Observations	32,224	32,224	32,205	30,658	7,165	10,560
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Loan Type FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.040	0.045		0.073	0.062	0.050
Pseudo R-squared			0.0531			
Y-mean	16.70	16.70	16.71	17.29	15.45	15.10

Table A.5: Approval and Default ML Features

This table lists the features included in the Machine Learning analysis (Table 4 and 7). The Minimal Specification consists of the features listed under Baseline Minimal Model. The Preferred Specification includes all of the features in the Minimal Specification column plus the additional columns listed. The Cash Flow Model consists of all of the Baseline Features plus the additional features listed under Cash Flow Model. Note that this includes transformations of the cash flow variables included in the OLS regressions.

Baseline Minimal Model	**Baseline Preferred Model**
FICO Score	Industry
Business Age (Years)	Quarter Number (1-4)
Number of Employees	Late Quarter (After Median)
Lender ID	Region (NE, Midwest, South, West)
Requested Loan Amount (Log)	Loan Type
State	
Cash Flow Model	
Credits (Log)	Balance (Log)
Withdrawals (Log)	(#) Insuff. Funds
(#) Low or Neg. Balance	$\mathbf{1}$(Daily Pay Loan)
S.D. Credits	S.D. Balance
Missing Withdrawals	Missing Daily Pay Loans
Credits (less new debt) (Log)	Missing Credits (less new debt)
Debits to Credits Ratio	Balance $\times$ Credits
(#) Insuff. Funds $\times$ (#) Low or Neg. Balance	Low Credit Utilization
Coeff. Variation Balance	Coeff. Variation Credits
$\mathbb{1}$(Daily Pay Loan) to Balance Ratio	Never Low or Neg. Balance
Never Insuff. Funds	(#) Insuff. Funds > 5
Balance Volatility Ratio	Credits to Balance Ratio

Table A.6: Default with Alternate Age Cutoffs

This table uses data from Lender A and Lender B (N = 38,021) on originated loans and their performance to show how credit score, cash flow, and borrower characteristics predict default. The dependent variable is an indicator for the loan being non-performing. The model is OLS in all columns. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Is Non-Performing (%)					
	Owner (< 35)	Owner (> 35)	Owner (< 40)	Owner (> 40)	Owner (< 50)	Owner (> 50)
	(1)	(2)	(3)	(4)	(5)	(6)
FICO	-4.25***	-3.62***	-3.81***	-3.60***	-3.63***	-3.71***
	(0.80)	(0.29)	(0.43)	(0.36)	(0.47)	(0.36)
Credits (Th$)	-2.11	-1.93***	-2.54**	-1.78**	-2.19**	-1.64**
	(1.39)	(0.62)	(0.96)	(0.70)	(0.79)	(0.63)
Withdrawals (Th$)	4.82*	3.16*	4.21	3.16*	4.53**	2.02
	(2.65)	(1.72)	(2.77)	(1.68)	(2.05)	(1.34)
Balance (Th$)	-3.72**	-1.32***	-2.65***	-1.20**	-1.84***	-1.15***
	(1.49)	(0.41)	(0.72)	(0.44)	(0.54)	(0.39)
$\mathbb{1}$(Daily Pay Loan)	0.70	0.70**	0.27	0.82**	0.88*	0.50
	(0.65)	(0.32)	(0.47)	(0.36)	(0.48)	(0.33)
(#) Low or Neg. Bal.	2.10***	1.37***	2.04***	1.32***	1.50***	1.46***
	(0.51)	(0.34)	(0.50)	(0.39)	(0.39)	(0.46)
(#) Insuff. Funds	-0.10	-0.16	-0.33	-0.12	0.00	-0.31
	(0.60)	(0.45)	(0.58)	(0.42)	(0.42)	(0.50)
S.D. Credits (Th$)	2.10*	1.62***	2.12*	1.52**	1.65**	1.65**
	(1.19)	(0.47)	(1.11)	(0.56)	(0.74)	(0.61)
S.D. Balance (Th$)	2.29**	0.80**	1.36**	0.83**	1.12*	0.65
	(0.95)	(0.29)	(0.51)	(0.36)	(0.56)	(0.62)
Business Age (Years)	-0.10	-0.15***	-0.06	-0.18***	-0.15***	-0.17***
	(0.12)	(0.04)	(0.11)	(0.04)	(0.05)	(0.04)
Number of Employees	0.04	-0.07**	-0.02	-0.07**	-0.04	-0.08**
	(0.05)	(0.03)	(0.05)	(0.03)	(0.03)	(0.04)
Observations	2,317	17,865	5,006	15,178	11,192	8,994
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Loan Type FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.071	0.053	0.066	0.052	0.054	0.058
Y-mean	15.80	15.45	15.48	15.49	15.70	15.22

Table A.7: Benefit of Cash Flow vs Baseline Model for Other Characteristics

This table shows the results from implementing Tail Analysis for Comparative Outcomes (TACO, see Section 5). We compare two random forest models: a Baseline model (containing FICO, firm size, firm age, and industry among others) and the Cash Flow (CF) model, which adds bank statement variables to the Baseline model. The table uses data from Lender A and Lender B (N = 38,021) on originated loans. The observation counts represent the sum across 1000 bootstrap holdout samples. The first two columns ("Tails") show the group's share in the decile tails population of bootstrapped sample observations. The next two columns restrict to the 10% of each bootstrap sample with the highest increase in default chance between the Baseline and the CF model, who are thus adversely affected by switching from the Baseline to the CF model. The mean shows the share of a given demographic in this group, which can be compared to the first column. The next two columns show the same metric for the bottom 10% (the group that most benefits from switching to the CF model). The last column shows the ratio between the "Hurt" and "Benefit" means, which we call the TACO ratio. A ratio of one implies no implication of switching models, a ratio less than one implies that the group is adversely affected, and a ratio greater than one implies that the group benefits. We calculate standard errors for the TACO ratio using the percentile bootstrap.

	Tails		Top 10% Default Increases w/ CF Model (Hurt)		Bottom 10% Default Increases w/ CF Model (Benefit)		
	N	Mean	N	Mean	N	Mean	TACO Ratio
FICO Split:							
Low FICO (< 700):	1,522,000	0.405	761,000	0.175	761,000	0.635	3.626***
High FICO (≥ 700)	1,522,000	0.595	761,000	0.825	761,000	0.365	0.442***
Full Sample:							
Young Firm (≤ 5)	1,522,000	0.302	761,000	0.221	761,000	0.383	1.735***
High Pct Black Pop (> 6%)	1,522,000	0.473	761,000	0.457	761,000	0.488	1.069
High Pct Black Pop (> 17%)	1,522,000	0.210	761,000	0.206	761,000	0.214	1.036
Female	1,513,552	0.218	755,681	0.240	757,871	0.196	0.816**
Low FICO (< 700):							
Young Firm (≤ 5)	1,522,000	0.140	761,000	0.033	761,000	0.247	7.573***
High Pct Black Pop (> 6%)	1,522,000	0.200	761,000	0.080	761,000	0.321	4.004***
High Pct Black Pop (> 17%)	1,522,000	0.090	761,000	0.038	761,000	0.142	3.784***
Female	1,513,552	0.088	755,681	0.046	757,871	0.131	2.882***
High FICO (≥ 700):							
Young Firm (≤ 5)	1,522,000	0.162	761,000	0.188	761,000	0.136	0.722*
High Pct Black Pop (> 6%)	1,522,000	0.272	761,000	0.377	761,000	0.168	0.445***
High Pct Black Pop (> 17%)	1,522,000	0.120	761,000	0.169	761,000	0.071	0.422***
Female	1,513,552	0.130	755,681	0.195	757,871	0.065	0.334***

Table A.8: Adverse Impact Ratio Results on Benefit of Cash Flow vs Baseline Model by Age Group

This table compare two random forest models: a Baseline model (containing FICO, firm size, firm age, and industry among others) and the Cash Flow (CF) model, which adds bank statement variables to the Baseline model. The table uses data from Lender A and Lender B (N = 38,021) on originated loans. The observation counts represent the sum across 1000 bootstrap holdout samples. The first column AIR represents the ratio of the relevant group to others in the full population of bootstrapped sample observations. The next two columns ("Tails") show the group's share in the tails population of bootstrapped sample observations. The next three columns restrict to the 10% of each bootstrap sample with the highest increase in default chance between the Baseline and the CF model, who are thus adversely affected by switching from the Baseline to the CF model. The mean shows the share of young owners in this group, which can be compared to the first column. The next three columns show the same metric for the bottom 10% (the group that most benefits from switching to the CF model). We calculate standard errors for the AIRs using the percentile bootstrap.

Panel A: By Owner Age

		Tails		Top 10% Default Increases w/ CF Model (Hurt)			Bottom 10% Default Increases w/ CF Model (Benefit)		
	AIR	N	Mean	N	Mean	AIR	N	Mean	AIR
Full Sample:									
Young Owner (< 35)	0.130	736,507	0.130	384,203	0.101	0.112	352,304	0.161	0.192**
Young Owner (< 40)	0.330	736,507	0.270	384,203	0.222	0.286	352,304	0.322	0.475***
Young Owner (< 45)	0.661	736,507	0.418	384,203	0.352	0.544*	352,304	0.489	0.958***
Young Owner (< 50)	1.245	736,507	0.578	384,203	0.506	1.024*	352,304	0.656	1.904***
Low FICO (< 700):									
Young Owner (< 35)	0.034	736,507	0.056	384,203	0.013	0.013***	352,304	0.103	0.115***
Young Owner (< 40)	0.072	736,507	0.110	384,203	0.029	0.030***	352,304	0.199	0.249***
Young Owner (< 45)	0.120	736,507	0.169	384,203	0.044	0.046***	352,304	0.305	0.438***
Young Owner (< 50)	0.172	736,507	0.226	384,203	0.062	0.066***	352,304	0.404	0.679***
High FICO (≥ 700):									
Young Owner (< 35)	0.089	736,507	0.073	384,203	0.088	0.096	352,304	0.058	0.061
Young Owner (< 40)	0.221	736,507	0.160	384,203	0.193	0.240	352,304	0.123	0.140**
Young Owner (< 45)	0.410	736,507	0.249	384,203	0.308	0.445	352,304	0.185	0.226***
Young Owner (< 50)	0.689	736,507	0.352	384,203	0.444	0.797	352,304	0.251	0.335***

Panel B: By Other Demographics

		Tails		Top 10% Default Increases w/ CF Model (Hurt)			Bottom 10% Default Increases w/ CF Model (Benefit)		
	AIR	N	Mean	N	Mean	AIR	N	Mean	AIR
FICO Split:									
Low FICO (< 700):	0.425	1,522,000	0.405	761,000	0.175	0.212***	761,000	0.635	1.742***
High FICO (≥ 700)	2.351	1,522,000	0.595	761,000	0.825	4.708***	761,000	0.365	0.574***
Full Sample:									
High Pct Black Pop (> 6%)	0.844	1,522,000	0.473	761,000	0.457	0.842	761,000	0.488	0.955*
High Pct Black Pop (> 17%)	0.251	1,522,000	0.210	761,000	0.206	0.260	761,000	0.214	0.272
Young Firm (≤ 5)	0.351	1,522,000	0.302	761,000	0.221	0.284**	761,000	0.383	0.622***
Female	0.291	1,513,552	0.218	755,681	0.240	0.317	757,871	0.196	0.244**
Low FICO (< 700):									
High Pct Black Pop (> 6%)	0.166	1,522,000	0.200	761,000	0.080	0.087***	761,000	0.321	0.472***
High Pct Black Pop (> 17%)	0.069	1,522,000	0.090	761,000	0.038	0.039***	761,000	0.142	0.166***
Young Firm (≤ 5)	0.091	1,522,000	0.140	761,000	0.033	0.034***	761,000	0.247	0.329***
Female	0.076	1,513,552	0.088	755,681	0.046	0.048***	757,871	0.131	0.151***
High FICO (≥ 700):									
High Pct Black Pop (> 6%)	0.461	1,522,000	0.272	761,000	0.377	0.605***	761,000	0.168	0.202***
High Pct Black Pop (> 17%)	0.157	1,522,000	0.120	761,000	0.169	0.203**	761,000	0.071	0.077***
Young Firm (≤ 5)	0.214	1,522,000	0.162	761,000	0.188	0.232	761,000	0.136	0.157*
Female	0.184	1,513,552	0.130	755,681	0.195	0.242***	757,871	0.065	0.070***

Table A.9: Predicting Approval with Cash Flow Variables

This table uses data from Lender A and Lender B (N = 162,818) to show how credit score, cash flow, and borrower characteristics predict loan approval (receiving an offer from a lender). The level of observation is the loan-borrower, so an applicant may appear multiple times as they apply to a single or multiple lender. The dependent variable is an indicator for approval. The model is logit in columns 5 and 6 and OLS in all other columns. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
FICO	11.61***	4.79***	11.99***	6.04***	0.62***	0.39***	6.81***
	(0.78)	(0.60)	(0.69)	(0.46)	(0.01)	(0.01)	(0.44)
Credits (Th$)		5.65***		6.64***		0.36***	7.55***
		(0.90)		(0.61)		(0.01)	(0.59)
Withdrawals (Th$)		-0.70		-8.05***		-0.41***	-6.81***
		(0.96)		(0.62)		(0.01)	(0.62)
Balance (Th$)		4.73***		5.73***		0.20***	6.13***
		(0.54)		(0.51)		(0.01)	(0.52)
𝟙(Daily Pay Loan)		-14.76***		-13.43***		-0.81***	-13.07***
		(0.57)		(0.67)		(0.02)	(0.66)
(#) Low or Neg. Bal.		-8.62***		-6.94***		-0.71***	-7.24***
		(0.38)		(0.41)		(0.01)	(0.40)
(#) Insuff. Funds		-1.08***		-1.80***		-0.69***	-1.60***
		(0.27)		(0.32)		(0.03)	(0.33)
S.D. Credits (Th$)		-1.74***		-1.23***		-0.05***	-1.06***
		(0.28)		(0.20)		(0.00)	(0.21)
S.D. Balance (Th$)		-1.15***		-1.19***		-0.07***	-1.09***
		(0.23)		(0.19)		(0.00)	(0.20)
Business Age (Years)			0.32***	0.27***	0.02***	0.02***	0.28***
			(0.03)	(0.03)	(0.00)	(0.00)	(0.03)
Number of Employees			0.06*	-0.02	0.00***	-0.00	0.00
			(0.03)	(0.02)	(0.00)	(0.00)	(0.02)
Requested Loan Amount							-7.67***
							(0.49)
Observations	162,818	162,818	162,812	162,812	162,800	162,800	162,812
Industry FE	No	No	Yes	Yes	Yes	Yes	Yes
State FE	No	No	Yes	Yes	Yes	Yes	Yes
Quarter FE	No	No	Yes	Yes	Yes	Yes	Yes
Lender FE	No	No	Yes	Yes	Yes	Yes	Yes
R-squared	0.044	0.203	0.151	0.258			0.267
Pseudo R-squared					0.1205	0.2496	
Y-mean	45.85	45.85	45.85	45.85	45.86	45.86	45.85

Table A.10: Approval in the Platform Sample

This table uses data from the Platform (N = 904,471) to show how credit score, cash flow, and borrower characteristics predict loan approval (receiving an offer from a lender). The level of observation is the loan-borrower, so an applicant may appear multiple times as they apply to a single or multiple lender. The dependent variable is an indicator for approval. The model is logit in column 3 and OLS in all other columns. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)					
					Owner (< 45)	Owner (> 45)
	(1)	(2)	(3)	(4)	(5)	(6)
FICO	6.29***	5.54***	0.39***	5.52***	5.72***	5.31***
	(0.05)	(0.06)	(0.00)	(0.06)	(0.07)	(0.09)
Credits (Th$)		-4.93***	-0.33***	-4.96***	-5.47***	-4.30***
		(0.38)	(0.03)	(0.38)	(0.52)	(0.55)
Withdrawals (Th$)		4.97***	0.34***	4.95***	5.61***	4.18***
		(0.38)	(0.03)	(0.38)	(0.52)	(0.55)
Balance (Th$)		-0.02	-0.02***	-0.04	-0.04	-0.00
		(0.07)	(0.00)	(0.07)	(0.10)	(0.11)
$\mathbb{1}$(Daily Pay Loan)		-1.68***	-0.13***	-1.67***	-1.62***	-1.78***
		(0.05)	(0.00)	(0.05)	(0.07)	(0.08)
(#) Low or Neg. Bal.		-2.81***	-0.36***	-2.80***	-2.84***	-2.76***
		(0.06)	(0.01)	(0.06)	(0.07)	(0.10)
(#) Insuff. Funds		-0.81***	-0.06***	-0.81***	-0.83***	-0.80***
		(0.05)	(0.00)	(0.05)	(0.06)	(0.08)
Business Age (Years)	0.16***	0.19***	0.01***	0.19***	0.30***	0.14***
	(0.01)	(0.01)	(0.00)	(0.01)	(0.02)	(0.01)
Number of Employees	0.04***	0.06***	0.00***	0.06***	0.06***	0.06***
	(0.01)	(0.01)	(0.00)	(0.01)	(0.01)	(0.01)
Requested Loan Amount				0.15***		
				(0.06)		
Observations	904,455	904,455	899,801	904,455	522,538	381,629
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.180	0.186		0.186	0.196	0.176
Pseudo R-squared			0.190			
Y-mean	21.16	21.16	21.27	21.16	20.61	21.92

Table A.11: Approval in the Originated Sample for First Time Applicants

This table uses data from Lender A and Lender B among first time applicants (N = 140,330) to show how credit score, cash flow, and borrower characteristics predict loan approval (receiving an offer from a lender). The analysis is limited to first time applicants and the level of observation is the loan-borrower, so an applicant may appear multiple times if they apply to multiple lenders. The dependent variable is an indicator for approval. The model is logit in column 3 and OLS in all other columns. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)					
					Owner (< 45)	Owner (> 45)
	(1)	(2)	(3)	(4)	(5)	(6)
FICO	12.20***	6.07***	0.40***	6.76***	6.03***	6.04***
	(0.69)	(0.47)	(0.01)	(0.45)	(0.44)	(0.49)
Credits (Th$)		6.48***	0.36***	7.37***	6.40***	5.74***
		(0.61)	(0.02)	(0.60)	(0.67)	(0.71)
Withdrawals (Th$)		-8.11***	-0.41***	-6.93***	-10.46***	-10.27***
		(0.63)	(0.02)	(0.62)	(0.80)	(0.93)
Balance (Th$)		5.89***	0.20***	6.23***	7.70***	6.38***
		(0.57)	(0.01)	(0.57)	(0.59)	(0.66)
$\mathbb{1}$(Daily Pay Loan)		-13.26***	-0.80***	-12.91***	-12.95***	-14.56***
		(0.64)	(0.02)	(0.64)	(0.81)	(0.75)
(#) Low or Neg. Bal.		-7.00***	-0.71***	-7.29***	-6.14***	-7.62***
		(0.40)	(0.01)	(0.39)	(0.35)	(0.43)
(#) Insuff. Funds		-1.79***	-0.63***	-1.58***	-0.29	-1.83***
		(0.31)	(0.03)	(0.32)	(0.32)	(0.22)
S.D. Credits (Th$)		-1.20***	-0.05***	-1.03***	-1.22***	-0.86***
		(0.22)	(0.01)	(0.23)	(0.25)	(0.25)
S.D. Balance (Th$)		-1.19***	-0.07***	-1.09***	-1.17***	-1.06***
		(0.21)	(0.01)	(0.21)	(0.17)	(0.24)
Business Age (Years)	0.27***	0.22***	0.01***	0.23***	0.39***	0.17***
	(0.03)	(0.02)	(0.00)	(0.02)	(0.03)	(0.03)
Number of Employees	0.07**	-0.02	-0.00	0.00	-0.05**	-0.01
	(0.03)	(0.02)	(0.00)	(0.02)	(0.02)	(0.02)
Requested Loan Amount				-7.16***		
				(0.49)		
Observations	140,324	140,324	140,312	140,324	47,161	61,045
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.147	0.257		0.265	0.285	0.275
Pseudo R-squared			0.2496			
Y-mean	45.17	45.17	45.17	45.17	36.66	42.47

Table A.12: Approval in the Originated Sample with Alternate Age Cutoffs

This table uses data from Lender A and Lender B (N = 162,818) to show how credit score, cash flow, and borrower characteristics predict loan approval (receiving an offer from a lender). The level of observation is the loan-borrower, so an applicant may appear multiple times as they apply to a single or multiple lender. The dependent variable is an indicator for approval. The model is OLS in all columns. All bank variables and FICO score are standardized to z-scores and can be interpreted as the change in the dependent variable from 1 standard deviation of change. # Low or Neg. Bal. is the number of low or negative ending balances across the statements. # Insuff. Funds is the number of insufficient funds transactions. Missing values are replaced with median values. Standard errors are clustered by industry and quarter. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)					
	Owner (< 35)	Owner (> 35)	Owner (< 40)	Owner (> 40)	Owner (< 50)	Owner (> 50)
	(1)	(2)	(3)	(4)	(5)	(6)
FICO	6.19***	5.94***	6.26***	5.89***	6.03***	6.04***
	(0.43)	(0.47)	(0.44)	(0.47)	(0.43)	(0.52)
Credits (Th$)	5.71***	6.40***	6.31***	6.31***	6.63***	5.94***
	(1.07)	(0.73)	(0.86)	(0.71)	(0.83)	(0.65)
Withdrawals (Th$)	-10.31***	-10.18***	-10.10***	-10.21***	-10.30***	-10.10***
	(1.47)	(0.82)	(0.88)	(0.83)	(0.73)	(0.91)
Balance (Th$)	8.06***	6.49***	7.32***	6.47***	6.96***	6.31***
	(0.79)	(0.51)	(0.70)	(0.55)	(0.60)	(0.63)
$\mathbb{1}$(Daily Pay Loan)	-12.14***	-14.38***	-12.98***	-14.48***	-13.67***	-14.64***
	(0.79)	(0.83)	(0.80)	(0.83)	(0.83)	(0.80)
(#) Low or Neg. Bal.	-5.41***	-7.18***	-5.88***	-7.33***	-6.37***	-7.63***
	(0.31)	(0.42)	(0.33)	(0.43)	(0.38)	(0.43)
(#) Insuff. Funds	0.58***	-1.42***	0.26	-1.70***	-0.33	-2.13***
	(0.16)	(0.26)	(0.24)	(0.29)	(0.32)	(0.27)
S.D. Credits (Th$)	-1.06**	-1.06***	-1.15***	-1.02***	-1.16***	-0.93***
	(0.40)	(0.23)	(0.32)	(0.23)	(0.23)	(0.26)
S.D. Balance (Th$)	-1.54***	-1.07***	-1.25***	-1.09***	-1.10***	-1.16***
	(0.41)	(0.18)	(0.21)	(0.20)	(0.18)	(0.23)
Business Age (Years)	0.23**	0.24***	0.37***	0.22***	0.35***	0.20***
	(0.08)	(0.02)	(0.04)	(0.02)	(0.04)	(0.04)
Number of Employees	0.04	-0.03	-0.03	-0.02	-0.05*	0.00
	(0.05)	(0.02)	(0.04)	(0.02)	(0.03)	(0.02)
Observations	17,625	106,016	34,605	89,036	71,509	52,134
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.285	0.279	0.288	0.277	0.285	0.276
Y-mean	31.42	41.81	34.58	42.57	37.97	43.57

Table A.13: Summary Statistics by Cash Flow-Intensive Lender Status

This table reports summary statistics by Lender (Platform) cash flow-intensity (CFI) (N = 879,889). Panel A compares application-forward level data by CFI Lender using a two-sample t-test. Panel B compares application level data by CFI Lender using a two-sample t-test. Panel C compares lender level data by CFI intensity using a two-sample t-test. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Application-Forward Level Statistics

	All				CF Intensive Lender		Not CF Intensive Lender		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Loan Variables:									
Approved (%)	879,889	19.7	0.00	40	729,293	21.2	150,596	12.2	8.98***
# Forwards	879,889	8.60	8.00	5.27	729,293	8.51	150,596	9.02	-0.50***
Business Open on Google	43,889	0.40	0.00	0.49	38,669	0.40	5,220	0.42	-0.02***
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	879,889	681	679	66	729,293	677	150,596	700	-22.86***
Credits (Th$)	879,889	100	48	154	729,293	98	150,596	111	-13.50***
Balance (Th$)	879,889	33.4	11	64	729,293	32.2	150,596	39.5	-7.37***
(#) Insuff. Funds	879,889	0.52	0.00	1.41	729,293	0.54	150,596	0.41	0.14***
(#) Low or Neg. Bal.	879,889	0.53	0.00	1.32	729,293	0.56	150,596	0.42	0.14***
Withdrawals (Th$)	879,889	101	49	155	729,293	99	150,596	113	-13.91***
Credits (less new debt) (Th$)	879,889	84	40	130	729,293	82	150,596	94	-11.56***
$\mathbb{1}$(Daily Pay Loan)	879,889	0.36	0.00	0.48	729,293	0.36	150,596	0.35	0.00***
Borrower Age:									
Owner Age	879,889	44.4	43	11	729,293	44.2	150,596	45.4	-1.23***
Young Owner (< 35)	879,889	0.22	0.00	0.42	729,293	0.23	150,596	0.20	0.04***
Young Owner (< 40)	879,889	0.40	0.00	0.49	729,293	0.41	150,596	0.36	0.05***
Young Owner (< 45)	879,889	0.58	1.00	0.49	729,293	0.58	150,596	0.54	0.04***
Young Owner (< 50)	879,889	0.72	1.00	0.45	729,293	0.73	150,596	0.69	0.04***

Panel B: Unique Application Level Statistics

	All				CF Intensive Lender		Not CF Intensive Lender		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Loan Variables:									
Approved (%)	261,187	16.5	0.00	28	169,391	18.7	91,796	12.4	6.31***
# Forwards	261,187	5.69	5.00	4.59	169,391	5.10	91,796	6.79	-1.69***
Business Open on Google	21,650	0.38	0.00	0.49	17,039	0.37	4,611	0.42	-0.05***
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	261,187	681	679	68	169,391	673	91,796	694	-20.89***
Credits (Th$)	261,187	89	40	145	169,391	85	91,796	96	-10.47***
Balance (Th$)	261,187	30.7	9.43	62	169,391	29.0	91,796	33.8	-4.81***
(#) Insuff. Funds	261,187	0.60	0.00	1.57	169,391	0.67	91,796	0.47	0.20***
(#) Low or Neg. Bal.	261,187	0.63	0.00	1.54	169,391	0.70	91,796	0.49	0.21***
Withdrawals (Th$)	261,187	89	41	146	169,391	86	91,796	96	-10.74***
Credits (less new debt) (Th$)	261,187	74	33	123	169,391	70	91,796	80	-9.24***
$\mathbb{1}$(Daily Pay Loan)	261,187	0.35	0.00	0.48	169,391	0.36	91,796	0.35	0.01***
Borrower Age:									
Owner Age	261,187	44.2	43	11	169,391	43.9	91,796	44.8	-0.83***
Young Owner (< 35)	261,187	0.23	0.00	0.42	169,391	0.24	91,796	0.22	0.03***
Young Owner (< 40)	261,187	0.41	0.00	0.49	169,391	0.42	91,796	0.39	0.03***
Young Owner (< 45)	261,187	0.58	1.00	0.49	169,391	0.59	91,796	0.56	0.03***
Young Owner (< 50)	261,187	0.72	1.00	0.45	169,391	0.73	91,796	0.71	0.02***

Panel C: Lender Level Statistics

	All				CF Intensive Lender		Not CF Intensive Lender		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Loan Variables:									
Approved (%)	73	21.1	19	15	45	21.3	28	20.7	0.62
# Forwards	73	6.67	7.33	3.30	45	7.59	28	5.19	2.40***
Business Open on Google	69	0.38	0.42	0.13	44	0.39	25	0.38	0.01
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	73	687	686	33	45	679	28	699	-19.99**
Credits (Th$)	73	120	114	64	45	125	28	112	13
Balance (Th$)	73	39.8	36	20	45	39.6	28	40.1	-0.55
(#) Insuff. Funds	73	0.60	0.52	0.33	45	0.62	28	0.58	0.04
(#) Low or Neg. Bal.	73	0.64	0.55	0.38	45	0.67	28	0.59	0.08
Withdrawals (Th$)	73	120	114	64	45	125	28	113	12
Credits (less new debt) (Th$)	73	100	95	54	45	104	28	93	11
$\mathbb{1}$(Daily Pay Loan)	73	0.38	0.38	0.11	45	0.40	28	0.35	0.05*
Borrower Age:									
Owner Age	73	44.4	44	1.99	45	44.3	28	44.6	-0.29
Young Owner (< 35)	73	0.23	0.22	0.07	45	0.23	28	0.22	0.00
Young Owner (< 40)	73	0.41	0.40	0.08	45	0.41	28	0.40	0.01
Young Owner (< 45)	73	0.58	0.57	0.07	45	0.58	28	0.57	0.01
Young Owner (< 50)	73	0.72	0.72	0.05	45	0.72	28	0.72	0.01

Table A.14: Effect of Assignment to Cash Flow-Intensive Lender (Within-Application), All Applicant Ages

This table uses data from the Platform to test whether cash flow-intensive lenders are more likely to approve young entrepreneurs (N = 879,889). Young is defined according to the column headers. The interaction term “Young=1 × CFI=1" represents the difference in approval likelihood for young applicants forwarded to cash flow-intensive lenders relative to older applicants and those forwarded to non-cash flow-intensive lenders. In odd (even) columns, the sample is restricted to applicants with a FICO score above (below) 700. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)							
	Young (< 35)		Young (< 40)		Young (< 45)		Young (< 50)	
FICO:	High	Low	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1	-0.12	0.89***	0.33	1.45***	0.77**	1.64***	1.34***	1.99***
	(0.41)	(0.31)	(0.35)	(0.26)	(0.33)	(0.27)	(0.35)	(0.30)
Observations	306,295	522,477	306,295	522,477	306,295	522,477	306,295	522,477
Application FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.322	0.316	0.322	0.316	0.322	0.316	0.323	0.316
Y-mean	24.53	17.08	24.53	17.08	24.53	17.08	24.53	17.08

Table A.15: Effect of Assignment to Cash Flow Intensive Lender on Platform Without Application Fixed Effects

This table uses data from the Platform on lenders that can be classified as being Cash Flow Intensive (N = 879,889) to test whether cash flow-intensive lenders approve traditionally constrained individuals. The level of observation is the application-lender, so an applicant may appear multiple times as the application is forwarded to multiple lenders. This is the same specification as Table 8 but without application fixed effects and instead including Industry, Quarter, and Lender fixed effects. Young is defined according to the column header and only applicants that fall below that threshold or are 50 years or older are included in the regressions. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood for young applicants forwarded to cash flow-intensive lenders relative to applicants over 50 and those forwarded to non-cash flow-intensive lenders. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	0.71***	0.83***	0.80***	0.79***
	(0.26)	(0.23)	(0.22)	(0.21)
Young=1	-4.06***	-3.21***	-2.71***	-2.37***
	(0.23)	(0.20)	(0.19)	(0.18)
Observations	444,412	599,988	754,161	879,889
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.091	0.091	0.090	0.090
Y-mean	19.36	19.43	19.54	19.65

Table A.16: Effect of Assignment to Cash Flow-Intensive Lender (Within-Application), Limited to Applications Sent to Both Types of Lender

This table uses data from the Platform to test whether cash flow-intensive lenders are more likely to approve young entrepreneurs. This table is similar to Table 8 but limited to applications sent to both cash flow-intensive and not cash flow-intensive lenders (N = 606,442). The level of observation is the application-lender, so an applicant may appear multiple times as the application is forwarded to multiple lenders. Young is defined according to the column header and only applicants that fall below that threshold or are 50 years or older are included in the regressions. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood for young applicants forwarded to cash flow-intensive lenders relative to applicants over 50 and those forwarded to non-cash flow-intensive lenders. In odd (even) columns, the sample is restricted to applicants with a FICO score above (below) 700. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)							
	Young (< 35)		Young (< 40)		Young (< 45)		Young (< 50)	
FICO:	High	Low	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1	0.90*	2.17***	1.14***	2.33***	1.28***	2.14***	1.36***	1.99***
	(0.47)	(0.38)	(0.40)	(0.33)	(0.37)	(0.31)	(0.35)	(0.30)
Observations	136,009	169,662	177,913	232,020	222,955	294,086	261,956	344,486
Application FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.310	0.289	0.311	0.290	0.313	0.289	0.312	0.289
Y-mean	24.30	17.67	24.45	17.78	24.70	17.81	24.82	17.91

Table A.17: Effect of Assignment to Cash Flow-Intensive Lender (Within-Application), Triple Interaction

This table uses data from the Platform to test whether cash flow-intensive lenders are more likely to approve young entrepreneurs (N = 879,889). The level of observation is the application-lender, so an applicant may appear multiple times as the application is forwarded to multiple lenders. Young is defined according to the column header and only applicants that fall below that threshold or are 50 years or older are included in the regressions. The interaction term "Young=1 × CFI=1 × Low FICO=1" represents the difference in approval likelihood for young, low FICO applicants forwarded to cash flow-intensive lenders relative to applicants over 50, high FICO applicants, and those forwarded to non-cash flow-intensive lenders. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1 × Low FICO=1	0.93	0.89*	0.58	0.39
	(0.60)	(0.52)	(0.48)	(0.47)
Young=1 × CFI=1	1.00**	1.25***	1.40***	1.46***
	(0.47)	(0.41)	(0.37)	(0.36)
CFI=1 × Low FICO=1	-5.28***	-5.38***	-5.44***	-5.42***
	(0.40)	(0.39)	(0.39)	(0.39)
Observations	418,269	564,689	710,039	828,772
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.315	0.316	0.316	0.316
Y-mean	19.56	19.61	19.73	19.83

Table A.18: Effect of Assignment to Cash Flow-Intensive Lender (Within-Application), Full Samples

This table uses data from the Platform to test whether cash flow-intensive lenders are more likely to approve young entrepreneurs (N = 879,889). The level of observation is the application-lender, so an applicant may appear multiple times as the application is forwarded to multiple lenders. Young is defined according to the column header and only applicants that fall below that threshold or are 50 years or older are included in the regressions. Panel A contains all observations and Panel B is limited to applications with a FICO score over 800.The interaction term “Young=1 × CFI=1" represents the difference in approval likelihood for young applicants forwarded to cash flow-intensive lenders relative to applicants over 50 and those forwarded to non-cash flow-intensive lenders. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: No FICO Split

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	1.00***	1.28***	1.28***	1.27***
	(0.30)	(0.26)	(0.24)	(0.23)
Observations	418,269	564,689	710,039	828,772
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.314	0.316	0.316	0.315
Y-mean	19.56	19.61	19.73	19.83

Panel B: Very High FICO Sample (> 800)

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	1.96	1.75	2.64**	2.63***
	(1.75)	(1.28)	(1.09)	(1.00)
Observations	18,563	23,987	30,098	35,165
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.352	0.358	0.361	0.358
Y-mean	28.52	28.46	28.70	28.70

Table A.19: Role of Risk in Cash Flow-Intensive Lender Approval

This table examines whether CFI lenders tend to approve riskier businesses, as measured by their ultimate survival as of September, 2024. We use data from the Platform combined with data collected from Google on business status. The level of observation is the application-lender, so an applicant may appear multiple times as the application is forwarded to multiple lenders. The sample in column 1 is limited to low FICO applicants. The sample in column 3 is limited to approved application-forwards. Approved is an indicator for being approved multiplied by 100. Quarter FE control for the application quarter. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved	Business Open on Google	
	Low FICO		Approved
	(1)	(2)	(3)
CFI=1 × Young Owner (< 45)=1	1.89***		0.02
	(0.55)		(0.01)
CFI=1 × Approved		-0.00	
		(0.01)	
CFI=1		-0.02***	-0.03***
		(0.00)	(0.01)
Approved		0.06***	
		(0.01)	
Young Owner (< 45)=1			-0.03*
			(0.02)
Observations	116,834	240,362	43,889
Application FE	Yes	No	No
Lender FE	Yes	No	No
Quarter FE	No	Yes	Yes
R-squared	0.293	0.003	0.002
Y-mean	15.78	0.36	0.40

Table A.20: Summary Statistics by Assignment to Cash Flow-Intensive Loan Officer

This table reports summary statistics by Loan Officer (Lender B) cash flow-intensity (CFI) (N = 11,535). Panel A compares application level data by CFI Loan Officer using a two-sample t-test. Panel B compares loan officer level data by CFI intensity using a two-sample t-test.

Panel A: Application Level Statistics

	All				CF Intensive Loan Officer		Not CF Intensive Loan Officer		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Loan Variables:									
Approved (%)	11,535	52.0	100	50	5,776	49.9	5,759	54.1	-4.19***
Non-Performing Loan (%)	1,222	25.0	0.00	43	606	24.4	616	25.6	-1.23
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	11,535	744	747	42	5,776	744	5,759	745	-0.85
Credits (Th$)	11,535	163	80	222	5,776	160	5,759	165	-4.58
Balance (Th$)	11,535	68.4	32	93	5,776	69.6	5,759	67.2	2.39
(#) Insuff. Funds	11,535	0.03	0.00	0.29	5,776	0.03	5,759	0.03	0.00
(#) Low or Neg. Bal.	11,535	0.35	0.00	0.96	5,776	0.35	5,759	0.34	0.01
Withdrawals (Th$)	11,535	178	94	224	5,776	177	5,759	179	-1.67
Credits (less new debt) (Th$)	11,535	159	85	197	5,776	158	5,759	160	-1.54
$\mathbb{1}$(Daily Pay Loan)	11,535	0.02	0.00	0.15	5,776	0.02	5,759	0.02	-0.00
S.D. Credits (Th$)	11,535	16.7	10	16	5,776	16.8	5,759	16.6	0.14
S.D. Balance (Th$)	11,535	8.42	5.35	7.49	5,776	8.50	5,759	8.35	0.15
Borrower Age:									
Owner Age	11,535	50.3	50	11	5,776	50.3	5,759	50.4	-0.16
Young Owner (< 35)	11,535	0.09	0.00	0.28	5,776	0.09	5,759	0.08	0.01
Young Owner (< 40)	11,535	0.21	0.00	0.41	5,776	0.21	5,759	0.21	0.00
Young Owner (< 45)	11,535	0.36	0.00	0.48	5,776	0.36	5,759	0.36	0.00
Young Owner (< 50)	11,535	0.52	1.00	0.50	5,776	0.52	5,759	0.52	0.00

Panel B: Loan Officer Level Statistics

	All				CF Intensive Loan Officer		Not CF Intensive Loan Officer		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Loan Variables:									
Approved (%)	15	56.1	56	12	5	52.9	10	57.6	-4.75
Non-Performing Loan (%)	15	20.6	20	13	5	24.1	10	18.9	5.23
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	15	740	740	8.47	5	739	10	741	-1.80
Credits (Th$)	15	177	154	72	5	162	10	184	-22.36
Balance (Th$)	15	65.9	62	23	5	62.3	10	67.7	-5.36
(#) Insuff. Funds	15	0.05	0.03	0.06	5	0.04	10	0.06	-0.02
(#) Low or Neg. Bal.	15	0.32	0.34	0.12	5	0.35	10	0.31	0.04
Withdrawals (Th$)	15	189	182	68	5	179	10	194	-14.75
Credits (less new debt) (Th$)	15	173	155	64	5	165	10	177	-12.03
$\mathbb{1}$(Daily Pay Loan)	15	0.04	0.02	0.03	5	0.05	10	0.03	0.02
S.D. Credits (Th$)	15	16.7	16	4.47	5	16.3	10	17.0	-0.70
S.D. Balance (Th$)	15	8.26	7.98	2.13	5	8.09	10	8.35	-0.26
Borrower Age:									
Owner Age	15	49.7	50	1.42	5	49.7	10	49.7	-0.07
Young Owner (< 35)	15	0.10	0.09	0.02	5	0.10	10	0.10	0.00
Young Owner (< 40)	15	0.23	0.22	0.04	5	0.23	10	0.22	0.01
Young Owner (< 45)	15	0.38	0.38	0.05	5	0.38	10	0.38	-0.00
Young Owner (< 50)	15	0.55	0.53	0.08	5	0.54	10	0.55	-0.01

Table A.21: Effect of Random Assignment to Cash Flow-Intensive Loan Officer, All Applicant Ages

This table uses data from Lender B to test whether young entrepreneurs are more likely to have their loan application approved when randomly assigned to a cash flow-intensive Loan Officer (N = 11,535). The level of observation is an application. Young is defined according to the column header. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood for young applicants assigned to cash flow-intensive loan officers relative to older applicants and those assigned to non-cash flow-intensive loan officers. In odd (even) columns, the sample is restricted to applicants with a FICO score above (below) the median. Standard errors are clustered by quarter and approver. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Approved (%)							
	Young (< 35)		Young (< 40)		Young (< 45)		Young (< 50)	
FICO:	High	Low	High	Low	High	Low	High	Low
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1	2.11	12.80**	4.14	8.93***	3.45	5.87**	3.29	3.28
	(5.49)	(4.91)	(3.09)	(3.24)	(2.85)	(2.71)	(2.49)	(2.75)
Young=1	-5.35	-8.05**	-5.24**	-4.90**	-4.53**	-4.32**	-3.51**	-3.27
	(4.22)	(3.46)	(2.06)	(2.27)	(1.96)	(2.02)	(1.64)	(2.11)
Observations	5,782	5,753	5,782	5,753	5,782	5,753	5,782	5,753
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.067	0.060	0.068	0.060	0.068	0.060	0.068	0.059
Y-mean	55.76	48.20	55.76	48.20	55.76	48.20	55.76	48.20

B Variable Definitions

B.1 Lender A (Application Level)

Loan Variables:

- Approved (%): An indicator variable that takes the value of one if a loan is approved and zero otherwise (for easy interpretation, multiplied by 100).
- Requested Loan Amount (Th$): The loan amount requested at the time of application (in thousands of dollars).
- Interest Rate (%, APR): The interest rate of the loan (in percentage points).
- Originated (%): An indicator variable that takes the value of one if a loan is originated and zero otherwise (for easy interpretation, multiplied by 100).
- Non-Performing Loan (%): An indicator variable that takes the value of one if a loan is in forbearance or default and zero otherwise (for easy interpretation, multiplied by 100).
- Originated Loan Amount (Th$): The loan amount (in thousands of dollars).
- Loan Maturity (Years): The loan-maturity (in years).

Credit Score & Cash Flow (Bank Statement) Variables:

- Applicant's Credit Score (FICO): The credit score of the borrower (330-850).
- Credits (Th$): The average monthly business revenue, implied from bank statements over the last year (in thousands of dollars).
- Balance (Th$): The average monthly balance in the bank account of the borrower (in thousands of dollars).
- (#) Insuff. Funds: The average monthly number of insufficient balance fees.
- (#) Low or Neg. Balance: The average monthly number of cases where the borrower's account balance is either low ($< \$1,000$) or negative.

- Withdrawals (Th$): The average monthly amount of debit transactions for applicants over the last six months (in thousands of dollars).
- $\mathbb{1}$(Daily Pay Loan): An indicator that takes the value of one if an applicant has daily pay loans and zero otherwise.
- S.D. Credits (Th$): The average monthly standard deviation of deposits of the borrower's account over the past three months (in thousands of dollars).
- S.D. Balance (Th$): The average monthly standard deviation of bank account balances of the borrower over the past three months (in thousands of dollars).

Borrower Characteristics:

- Owner Age (Years): The owner's age at the time of the loan origination (in years).
- Female: An indicator variable that takes the value of one if the borrower is female and zero otherwise.
- Business Age (Years): The age of the borrower's business at the time of the loan origination (in years).
- Young Firm: An indicator variable that takes the value of one if the loan's borrowing firm age is below 5 years (the 50th percentile in business age across all loans), and zero otherwise.
- Pct Black Pop: The fraction of the black population in the borrower's zip code. We collect the total population by race (Asian, Black, Native, Pacific Islander, White, Other) at the zip code level. The data is from the 2020 American Community Survey (ACS) 5-year tables. Since ACS data uses zip code tabulation areas (a Census-specific analogue to zip code), we use a crosswalk provided by Uniform Data System to map ACS data to our zip codes.
- High Pct Black Pop: An indicator variable that takes the value of one if the fraction of the black population in the borrower's zip code is above 6% or 17% (relatively the 50th and 75th percentile in Pct Black Pop across all loans), and zero otherwise.
- Number of Employees: The number of employees in a firm.

B.2 Lender B (Application Level)

Loan Variables:

- Approved (%): An indicator variable that takes the value of one if a loan is approved and zero otherwise (for easy interpretation, multiplied by 100).
- Requested Loan Amount (Th$): The loan amount requested at the time of application (in thousands of dollars).
- Interest Rate (%, APR): The interest rate of the loan (in percentage points).
- Originated (%): An indicator variable that takes the value of one if a loan is originated and zero otherwise (for easy interpretation, multiplied by 100).
- Non-Performing Loan (%): An indicator variable that takes the value of one if a loan is in forbearance or default and zero otherwise (for easy interpretation, multiplied by 100).
- Originated Loan Amount (Th$): The loan amount (in thousands of dollars).
- Loan Maturity (Years): The loan-maturity (in years).

Credit Score & Cash Flow (Bank Statement) Variables:

- Applicant's Credit Score (FICO): The credit score of the borrower (450-850).
- Credits (Th$): The average monthly credits in the bank account of the borrower over the last six months (in thousands of dollars).
- Balance (Th$): The average monthly balance in the bank account of the borrower (in thousands of dollars).
- (#) Insuff. Funds: The average monthly number of insufficient balance fees over the past six months.
- (#) Low or Neg. Balance: The average monthly number of cases where the borrower's account balance is either low ($< \$1,000$) or negative over the past six months.
- Withdrawals (Th$): The average monthly amount of debit transactions for applicants (in thousands of dollars).

- Credits (less new debt)(Th$): The average monthly business revenue adjusted for loans, implied from bank statements over the last year (in thousands of dollars).
- $\mathbb{1}$(Daily Pay Loan): An indicator that takes the value of one if an applicant has daily pay loans in the last six months and zero otherwise.
- S.D. Credits (Th$): The average monthly standard deviation of deposits of the borrower's account over the past six months (in thousands of dollars).
- S.D. Balance (Th$): The average monthly standard deviation of bank account balances of the borrower over the past six months (in thousands of dollars).

Borrower Characteristics:

- Owner Age (Years): The owner's age at the time of the loan origination (in years).
- Female: An indicator variable that takes the value of one if the borrower is female and zero otherwise.
- Business Age (Years): The age of the borrower's business at the time of the loan origination (in years).
- Young Firm: An indicator variable that takes the value of one if the loan's borrowing firm age is below 5 years (the 50th percentile in business age across all loans), and zero otherwise.
- Pct Black Pop: The fraction of the black population in the borrower's zip code. We collect the total population by race (Asian, Black, Native, Pacific Islander, White, Other) at the zip code level. The data is from the 2020 American Community Survey (ACS) 5-year tables. Since ACS data uses zip code tabulation areas (a Census-specific analogue to zip code), we use a crosswalk provided by Uniform Data System to map ACS data to our zip codes.
- High Pct Black Pop: An indicator variable that takes the value of one if the fraction of the black population in the borrower's zip code is above 6% or 17% (relatively the 50th and 75th percentile in Pct Black Pop across all loans), and zero otherwise.
- Number of Employees: The number of employees in a firm.

Loan Officer Variables

- CF Intensive: We identify CFI loan officers as those with top-quartile improvement in AUC ROC from the Baseline to the Cash Flow Logistic model.

B.3 The Platform (Application-Forward Level)

Loan Variables:

- Approved (%): An indicator variable that takes the value of one if a loan is approved and zero otherwise (for easy interpretation, multiplied by 100).
- Requested Loan Amount (Th$): The loan amount requested at the time of application (in thousands of dollars).
- Interest Rate (%, APR): The interest rate of the loan (in percentage points).
- Originated (%): An indicator variable that takes the value of one if a loan is originated and zero otherwise (for easy interpretation, multiplied by 100).

Credit Score & Cash Flow (Bank Statement) Variables:

- Applicant's Credit Score (FICO): The credit score of the borrower (400-850).
- Credits (Th$): The average monthly credits in the bank account of the borrower over the last three months (in thousands of dollars).
- Balance (Th$): The average monthly balance in the bank account of the borrower over the last three months (in thousands of dollars).
- (#) Insuff. Funds: The average monthly number of insufficient balance fees over the last three months.
- (#) Low or Neg. Balance: The average monthly number of cases where the borrower's account balance is either low ($< \$1,000$) or negative over the last three months.
- Withdrawals (Th$): The average monthly amount of debit transactions for applicants over the last three months (in thousands of dollars).

- Credits (less new debt)(Th$): The average monthly business revenue adjusted for loans, implied from bank statements over the last three months (in thousands of dollars).
- $\mathbb{1}$(Daily Pay Loan): An indicator that takes the value of one if an applicant has daily pay loans in the last three months and zero otherwise.

Borrower Characteristics:

- Owner Age (Years): The owner's age at the time of the loan origination (in years).
- Female: An indicator variable that takes the value of one if the borrower is female and zero otherwise.
- Business Age (Years): The age of the borrower's business at the time of the loan origination (in years).
- Young Firm: An indicator variable that takes the value of one if the loan's borrowing firm age is below 5 years (the 50th percentile in business age across all loans), and zero otherwise.
- Pct Black Pop: The fraction of the black population in the borrower's zip code. We collect the total population by race (Asian, Black, Native, Pacific Islander, White, Other) at the zip code level. The data is from the 2020 American Community Survey (ACS) 5-year tables. Since ACS data uses zip code tabulation areas (a Census-specific analogue to zip code), we use a crosswalk provided by Uniform Data System to map ACS data to our zip codes.
- High Pct Black Pop: An indicator variable that takes the value of one if the fraction of the black population in the borrower's zip code is above 6% or 17% (relatively the 50th and 75th percentile in Pct Black Pop across all loans), and zero otherwise.
- Number of Employees: The number of employees in a firm.

Lender Variables

- CF Intensive: We identify CFI lenders as those for which the AUC ROC is higher in the Cash Flow ML model than in the Baseline ML model. We exclude lenders receiving under 50 application forwards, with an approval rate over 95%, or with an AUC improvement over 20%.

B.3 Additional ML Variables

Credit Score & Cash Flow (Bank Statement) Variables:

- Debits to Credits Ratio: The ratio of the log of variables *Withdrawals (Th$)* over *Credits (Th$)*.
- Balance × Credits: An interaction term between variables *Balance (Th$)* and *Credits (Th$)*.
- (#) Insuff. Funds × (#) Low or Neg. Balance: An interaction term between variables *(#) Insuff. Funds* and *(#) Low or Neg. Balance*.
- Low Credit Utilization: An indicator variable that takes the value of 1 if the variable *Credits (Th$)* is below the median, and zero otherwise.
- Coeff. Variation Balance: The ratio of the log of variables *S.D. Balance (Th$)* over *Balance (Th$)*.
- Coeff. Variation Credits: The ratio of the log of variables *S.D. Credits (Th$)* over *Credits (Th$)*.
- $\mathbb{1}$(Daily Pay Loan) to Balance Ratio: The ratio of the log of variables $\mathbb{1}$*(Daily Pay Loan)* over *Balance (Th$)*.
- Never Low or Neg. Balance: An indicator variable that that takes the value of 1 if *(#) Low or Neg. Balance* is equal to zero, and zero otherwise.
- Never Insuff. Funds: An indicator variable that that takes the value of 1 if *(#) Insuff. Funds* is equal to zero, and zero otherwise.
- (#) Insuff. Funds > 5: An indicator variable that that takes the value of 1 if *(#) Insuff. Funds* is greater than 5, and zero otherwise.
- Balance Volatility Ratio: The ratio of the log of variables *S.D. Balance (Th$)* over *Credits (Th$)*.
- Credits to Balance Ratio: The ratio of the log of variables *Credits (Th$)* over *Balance (Th$)*.

Borrower Characteristics:

- Quarter Number: The quarter number (1-4) of the loan application date.
- Late Quarter: The loan application is after the median loan application date in the combined data (post 10/1/2022).
- Region: The borrower's state is mapped to the following regions: Midwest, Northeast, South, West, and Other. This classification is based on the US Census designation by state.

C Machine Learning Methods

A random forest classifier is a machine learning method that uses a series of decision trees to classify or predict outcomes. Each tree is trained on a randomly selected subset of the data (bagging) and a random subset of features, ensuring diversity across the ensemble. The classifier determines its final prediction through majority voting in classification tasks. This aggregation process reduces over fitting and improves generalization compared to individual decision trees. Since the development of the method (Ho, 1995; Breiman, 2001), random forests have been increasingly applied to complex predictive tasks in the economics and finance fields.

In this study, the primary objective is to develop classifications and predictions of target values, and to measure how these predictions change across different model specifications.[53] Random forests are well-suited for such tasks, as they not only offer predictive accuracy but also provide insights into the relative importance of input features. Feature importance is calculated directly from the trained model's attributes, ranking features based on their contributions to reducing impurity across all trees in the forest. This analysis helps evaluate how the relevance of different variables shifts between models, offering valuable insights into the underlying data-generating process.

Hyperparameter Tuning: We tune the hyperparameters for the random forest models using an optimization routine with a TPE Sampler that maximizes model AUC-PR through iterative Bayesian sampling. The hyperparameter tuning executes 100 trials of stratified K-fold cross-validation. Our implementation uses Optuna, a hyperparameter optimization framework for Python. These optimized hyperparameters are selected once for each model specification, and then used for all subsequent estimation and prediction.

[53]The analysis utilized a custom framework built in Python, with dependencies including scikit-learn, numpy, pandas, optuna and other libraries for data manipulation and visualization.

In the first stage, the training dataset is used with 5-fold cross-validation to evaluate model performance. 5-fold cross-validation splits the data into five parts, with each part serving as the validation set once while the remaining four parts are used for training, repeating this process for each fold and preparing a mean AUC-PR score.[54] This cross-validation procedure is repeated 100 times, each with a different set of hyperparameters chosen through Bayesian optimization by Optuna. After 100 iterations, Optuna identifies the best hyperparameter configuration based on the highest mean AUC-PR score achieved across all cross-validation folds. This entire process is done once for the baseline model and once for the cash flow model, resulting in two sets of optimized hyperparameters.

- Estimators: specifies the number of trees that the model will train. A higher number of trees produces a more complex model at the risk of overfitting.
- Maximum Tree Depth: defines the maximum depth of each individual tree in the model. A higher depth produces a model able to measure more complex patterns at the risk of overfitting.
- Minimum Samples to Split: The minimum number of samples required to split an internal node.
- Feature Sample by Tree: the fraction of features (columns) that will be randomly selected to grow each tree. This prevents overfit by ensuring the model does not rely too heavily on any particular feature or combination of features.
- Maximum Leaf Nodes: Maximum number of leaf nodes in the trees.
- Minimum Child Weight: the minimum sum of instance weights (hessian) needed in a child node. A node is split only if the resulting child nodes have at least a certain amount of"weight."

Bootstrapped Model Evaluation: In the second stage, the training dataset is randomly split into a training sample (80%) and validation sample (20%) for each of 100 bootstrapped iterations.[55] The training sample is used to fit the model, while the validation sample is used to assess performance and check for early stopping. This setup provides average performance metrics that are robust to sampling and reduce the chances of overfit. Each model's performance

[54] We tune our hyperparameters to maximize AUC-PR because AUC-PR can be more informative and stable for datasets like ours with unbalanced outcome classes. Choosing to maximize ROC AUC yields almost identical results.

[55] We stratify splits by the outcome class to maintain consistent proportions of the outcomes in the validation sample.

is then evaluated on the holdout (testing) dataset, and mean performance metrics are collected and standard errors are calculated. To evaluate the differences between the baseline and cash flow models, paired t-tests are conducted on the differences in performance.