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, Revised September 2025

We are grateful to Tetyana Balyuk, Brock Blake, Greg Buchak, Kelly Cochran, Brigham Frandsen, Ryan Gilland, Khasi-marc Jamieson, Melissa Koide, Jason Kotter, Dan O'Malley, Sandip Nayak, Sarah Mazurek, Ryan Metcalf, Pradeep Mittal, David Snitkof, Abby Sleight, Sam Taussig, and Nishant Vats for support and comments. We also thank participants at the NBER Summer Institute 2025, UNC Kenan-Flagler seminar, the IMF-WIFPR Conference on Financial and Real Implications of Technologies, the BYU Mathematics Applied Analysis seminar, the Georgia Tech - Federal Reserve Bank of Atlanta Household Finance Conference, and the U.S. SBA. 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.

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, Revised September 2025
JEL No. G21, G23, G32, J14, J16

ABSTRACT

Younger entrepreneurs are disadvantaged by the common practice of relying on personal credit scores to underwrite small business loans. In causal analysis exploiting within-applicant assignment to multiple lenders, we show that younger business owners have a higher chance of loan approval when the lender incorporates recent cash flows from business checking accounts. We present a new method—Tail Analysis for Comparative Outcomes (TACO)—to assess, in a nonparametric way, who benefits from alternative models. This method is especially useful for evaluating machine learning models. We show that incorporating cash flow data into default prediction models disproportionately benefits younger entrepreneurs.

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

Sabrina T. Howell
Harvard University
Harvard Business School
and NBER
showell@hbs.edu

Mark J. Johnson
Brigham Young University
Department of Finance
markjjohnson@byu.edu

Siena Matsumoto
University of California, Berkeley
Haas School of Business
siena_matsumoto@berkeley.edu

1 Introduction

Personal and small business lending in the U.S. relies heavily on personal credit scores. Credit registries increase efficiency, but they are imperfect (Einav et al., 2013; Blattner et al., 2022). In this paper, we begin with the observation that personal credit scores 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 greater weight on the length of one's credit history than on individual repayment events. Furthermore, fair lending laws discourage the direct use of age in underwriting, but permit favorable treatment for the elderly. These features suggest the possibility of a structural disadvantage for younger entrepreneurs, which we explore in this paper.

Such a disadvantage may help explain why entrepreneurs are most commonly in their 40s when they found their firms, despite evidence that latent demand is higher among younger people (Azoulay et al., 2020; Harris, 2024). Little is known about how financial constraints vary across the lifecycle of business owners. We document that standard risk scoring models relying on personal credit scores disadvantage younger owners. First, using causal analysis, we show that younger entrepreneurs applying for small business loans are more likely to be approved when the lender makes use of timely information about repayment ability from bank statements ("cash flow data"). Second, we develop a novel method to compare predictions for subgroups across models, which is especially useful for policy and research on machine learning (ML) model interpretation. Applying this method, we show that more young entrepreneurs are deemed creditworthy when we add cash flow data to default prediction models that otherwise rely primarily on personal credit scores.

The context for our analysis is that in the 1990s, underwriting shifted from local loan officers using soft information to arm's length models that relied heavily on personal credit scores.[1] After the Financial Crisis, new fintech lenders exploited digitization and automation to incorporate cash flow information from bank statements into credit scoring models for both personal and small business lending (FinRegLab, 2019). Unlike credit scores, which emphasize history and delinquency, cash flow data highlight recent account activity and repayment capacity. These new data are now being adopted more widely in conventional lending.[2] We hypothesize that incorporating cash flow data into underwriting models could mitigate the mechanical

[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 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).

disadvantage facing younger entrepreneurs from lower or noisier personal credit scores.

Small businesses offer a useful laboratory because they depend heavily on debt (Robb and Robinson, 2014). Understanding their financing is economically important, as these firms employ about half of private-sector workers and account for nearly all firms. Existing evidence that alternative data can improve lending technologies motivates our analytical strategy. Examples include the presence of cashless payments in India and China (Ouyang, 2022; Ghosh et al., 2023), delivery app transactions in Mexico (Chioda et al., 2024), and alternative consumer characteristics such as education in the U.S. (Di Maggio et al., 2022). To our knowledge, this is the first paper to study bank account-based cash flows (i.e., inflows and outflows) in underwriting, and the first to study how entrepreneur age is relevant to the impacts of credit risk scoring or alternative data.

We use data on applications, originations, and loan performance from three fintech companies serving small businesses between 2013 and 2024. Two are lenders (who we call Lenders A and B) and one is a platform.[3] From all three companies, we obtain standard underwriting inputs such as the credit score, cash flow-related variables from bank statements ("cash flow variables"), and industry. The cash flow variables are standard among fintech lenders, capturing revenues, withdrawals, balances, volatility, and distress signals such as negative balances or insufficient funds transactions. We also obtain information they collect but which is not used in underwriting, such as the business owner's age. The platform forwards these data in an application packet to a subset of 111 participating lenders. We observe offers from the lenders and loan take-up. In total, our data includes about 1.1 million applications and 74,000 loans.

To establish the relevance of the cash flow variables, we examine how they and FICO predict loan default and approval, after controlling for the traditional inputs to underwriting models of industry, firm age, firm size, time, and state. Consistent with intuition, measures of business health and revenue—such as more inflows (i.e., credits)—are associated with lower chances of default and higher chances of approval, while measures of distress—such as overdrafts—predict higher chances of default and lower chances of approval. The relationships between cash flow variables and default are robust to controlling for the interest rate, and cash flow variables do not strongly predict interest rates.

Our first main analysis causally identifies the effect of cash flow-based underwriting on access to credit, and examines how it varies by entrepreneur age. The Platform lenders vary in whether they use cash flow data to underwrite. While applicants are not randomly assigned to lenders,

[3]These companies have requested anonymity. There are no Paycheck Protection Program loans or otherwise subsidized or mission-driven loans in the sample.

we use application fixed effects to remove potential confounders from applicant quality or lender preferences (following Khwaja and Mian (2008)). This is possible because applications are sent to multiple lenders. We classify lenders as cash flow–intensive (CFI) using a machine learning procedure: For each lender, we compare a Baseline model that relies primarily on FICO to a Cash Flow model that augments it with bank statement variables. Lenders for whom the Cash Flow model improves approval prediction are designated CFI (70% of lenders). We manually verify that CFI-classified lenders require bank statements as part of an application.

We find that assignment to a CFI lender increases the chance of approval for younger entrepreneurs, with larger effects among low-FICO applicants. For example, entrepreneurs under 40 years old have a 2.4 percentage point (pp) higher chance of approval at CFI lenders, relative to non-CFI lenders and older entrepreneurs. This effect represents 12% of the mean chance of approval, which is about 20%. There is also a negative effect on the interest rate (APR), conditional on receiving at least two offers. These effects are concentrated among low-FICO applicants, especially when comparing applicants in the middle range, who have neither poor nor excellent scores. In this region, FICO score differentials may be less informative about default, causing younger entrepreneurs with lower but not very poor FICO scores to especially benefit from cash flow underwriting.

There may be concern that CFI lenders are more likely to approve younger people because they take more risk, not because they can identify "diamonds in the rough" using alternative cash flow data. To assess this possibility, we conduct three analyses. First, we show that the effects are *larger* among especially conservative lenders, filtered either on approval rate or by the FICO score distribution of approved loans. Second, 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 while approval predicts survival, being young and assigned to a CFI lender does not predict survival, in contrast to the negative relationship we would expect if risk preferences explained the main finding. Third, we confirm the result in a context with true random assignment across loan officers within the Lender B sample.

Lenders seek to minimize default risk and typically build underwriting models in part by predicting default among past borrowers. The question about whether our results might reflect excessive risk-taking by CFI lenders is related to a broader question about whether adding cash flow variables to a FICO-based model yields a more accurate underwriting model for younger entrepreneurs. If a lender is minimizing default risk and approves more younger people when using cash flow variables, then it must be the case that more young owners have lower default risk in a model that includes cash flow variables ("Cash Flow model") than in a FICO-centered "Baseline" model. In our second main analysis, we show this formally by comparing predicted default across age groups in the Baseline and Cash Flow models.

To accomplish this, we develop a new method—"Tail Analysis for Comparative Outcomes," or TACO—to assess the benefits for sub-populations of using one model versus another. TACO compares two models' predicted default probability for each applicant. Then, for the subgroup of interest, it measures the share who are reclassified into the tail where predicted default probabilities fall the most. If lenders approve borrowers in inverse relation to predicted risk, landing in the tail of applicants with the largest fall in predicted default probability implies a higher approval likelihood. We consider a subgroup that is disproportionately pushed into this tail as benefiting from the model change. (Our notion of "benefit" refers only to this risk reclassification.)

We devised TACO 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 ML models that are otherwise difficult to interpret.

With the same ML specification that we use to classify lenders as CFI, we first show that the Cash Flow model predicts default significantly better across multiple measures and is more predictive for younger owners. We then implement TACO in three steps. First, we calculate for each borrower the difference in predicted default probability between the Cash Flow and Baseline models. Second, we identify 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 compare the share of young borrowers across the helped and hurt groups. TACO with small tails identifies those for whom the model change has the largest impact. As the tail size increases—with a maximum of 50%—TACO studies reallocation more broadly in the sample. Note that traditional average treatment effects do not capture these types of nuanced distributional effects.[4]

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

[4]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 offer limited insight if the average young owner is unlikely to receive credit.

helped group divided by the young share in the hurt group. The ratio is highest for the youngest entrepreneurs (under 35). Using the top and bottom deciles to define the tails, their TACO ratio is 1.37, meaning that there are 37% more young people in the helped group than in the hurt group. The TACO ratio declines as we raise the threshold for "young," indicating that the youngest applicants enjoy the largest benefits from cash flow underwriting. Consistent with lower FICO scores disproportionately constraining young entrepreneurs, the ratio is higher when we focus on young borrowers with lower FICO scores. For those under 40 with FICO below 740, the TACO ratio is 2.1. Yet even among high-FICO applicants, the Cash Flow model is relatively better for the young. These results are robust to using specifications in which APR is among the features used to predict default (i.e., "controlling" for APR).

The evidence suggests that younger entrepreneurs benefit from cash flow data primarily because they face a credit score disadvantage, not because they have superior cash flow metrics.[5] This likely reflects FICO being both noisier and mechanically biased downward for younger people. In related work, Blattner and Nelson (2021) show how noise in credit scores can create inefficiency in lending. Our main results support a role for both channels.

We shed further light on the channel by applying the TACO default prediction results to the Platform application data. We re-estimate the causal, within-application regressions using scores representing whether the CF model (trained out of sample on Lender A and B application data) predicts the applicant to be riskier or less risky than the FICO baseline model. We show that entrepreneurs predicted to benefit from cash flow data are much more likely to be approved at CFI lenders, whereas those predicted to be 'Most Hurt' do not show lower approval rates. These asymmetric results point to FICO bias playing a strong role. This analysis also offers comforting validation of our CFI lender measure and shows that TACO is relevant to underwriting in practice. Furthermore, the results highlight how incorporating new data in underwriting does not necessarily lead to credit reallocation. Only if the supply of capital is constrained would there would be reallocation away from, in our case, people with high FICO scores but poor cash flow metrics (a group that is disproportionately older). If capital is not constrained—either because new lenders enter or existing lenders expand credit—there would be more overall lending.

Our main contribution 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, contributing to the small but growing literature on entrepreneurship over the life cycle (Zhang and Acs, 2018; Azoulay et al., 2020). Theoretical

[5] For example, within low-FICO groups, younger and older people have similar cash flow metrics.

models of entrepreneurship predict that younger entrepreneurs are likely to face financial constraints (Evans and Jovanovic, 1989; Stiglitz and Weiss, 1981). Empirical research shows that financial constraints affect entry into entrepreneurship, including through housing 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).

Comparing credit scores with bank statement variables relates to research on the efficiency of credit scores and their distributional implications (Blattner and Nelson, 2021; Bhutta et al., 2025). Bakker et al. (2025) show that default predictions for young adults are systematically biased, likely due to thinner credit files and limited financial histories. In theoretical work, Chatterjee et al. (2023) micro-found why credit scores would, in general, increase with age. They argue that asymmetric information makes the history of actions a valuable signal of type, and older people have longer histories. Our results suggest that the credit score—a single number summarizing a person's history—leaves room for timely data on repayment ability to improve access to finance for entrepreneurs who, by virtue of being young, have relatively short histories.

A second 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 ML models. This connects to research on how alternative inputs, ML, and risk-taking interact in non-bank lending. There is concern that ML 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). There is also evidence that fintech lenders take more risk and engage in regulatory arbitrage (Buchak et al., 2018; Bartlett et al., 2022; Propson et al., 2024). Yet fintech lenders who rely heavily on ML claim to expand access to credit by identifying creditworthy borrowers who would be rejected by conventional underwriting.[6] We contribute to this debate by demonstrating that incorporating new inputs changes how ML model adoption impacts protected classes.

Finally, our results contribute to ongoing policy discussions about open banking, which enables consumer bank data to flow to third-party financial services providers (He et al., 2023; Alok et al., 2024; Yu, 2024; Babina et al., 2025). While evidence from other countries shows that digital footprints or payment information can improve default prediction (Berg et al., 2020; Ghosh et al.,

[6]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). In practice, fintech lenders are more likely to extend credit to under-served minorities and distressed areas (Jagtiani and Lemieux, 2018; Erel and Liebersohn, 2022; Chernenko and Scharfstein, 2024; Cornelli et al., 2024; Howell et al., 2024).

2023; Agarwal et al., 2021; Rishabh, 2024), we provide the first large-scale evidence in the U.S. that standardized bank statement variables improve credit access for younger entrepreneurs. In documenting these benefits, we contribute both to the open banking debate and to broader work on the role of fintech in the financial system (Mills and McCarthy, 2016; Seru, 2019; Tang, 2019; Vives, 2019; Vallee and Zeng, 2019; Philippon, 2019; Balyuk et al., 2025; Gopal and Schnabl, 2022; Ben-David et al., 2025).

2 Age and Entrepreneurship

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. Substantial evidence shows that entrepreneurs—particularly successful ones and those founding 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. Zhang and Acs (2018) (U.S. CPS) and Blanchflower (2004) (U.S. and EU) find that the probability of entrepreneurship rises with age, in the 40s–50s. Other evidence shows that older workers are more likely to be self-employed or business owners (Zissimopoulos and Karoly, 2007; Kautonen et al., 2014; Fairlie et al., 2016). Among U.S. employer firms, 51% of owners are over 55, compared to just 23% of the labor force, while those under 35 account for 35% of the labor force but only 6% of owners (Headd, 2021). Overall, the relationship between age and entrepreneurship takes an inverted U-shape, steepest in the years approaching 50 (Bates, 1990; Kautonen et al., 2017; Zhao et al., 2021; Brieger et al., 2021). Our fintech loan application data are consistent with these patterns: the median applicant is in their 40s (discussed in Section 4).

Constraints for Younger Business Owners. FICO scores increase systematically with age. In our data, this pattern is illustrated in Figure 1. FICO rises almost linearly from below 670 for entrepreneurs under 30 to 720 for entrepreneurs over 70. Nationally, scores peak near 757 for people in their 80s, with the steepest gains in the 50s and 60s (Appendix Figure A.1). The highest average scores are for people in their 90s, showing how age is relevant to underwriting across the entire life cycle. Consistent with these patterns, Avery et al. (2012) show that credit scores act as a proxy for age, disadvantaging younger people while showing little disparate impact by race or gender.

The constraints imposed by FICO are reinforced 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.[7] Lenders who hew to this provision may use age directly, but only if they submit to an additional disclosure and oversight. This is a regulatory burden that deters the lenders in our data from directly using age.[8]

Are younger people more financially constrained as a result of their age? There is some data to suggest that this may be the case. While entrepreneurs in practice tend to be older, there is greater entrepreneurial desire or intent among younger people, especially those under 35 (Babson, 2022; Gielnik et al., 2018; Levesque and Minniti, 2006). A recent Harris Poll found the highest rates of interest in becoming an entrepreneur among the youngest group; with 71% of respondents aged 18–34 reported wanting to start a business, compared to 53% of those aged 45–54, with even lower rates among older groups (Harris, 2024). An industry press article put it directly, with the title: "Banks Hate Young Entrepreneurs" (James, 2019).

Few studies explicitly examine financial constraints. 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. Levesque and Minniti (2011) theorize that while younger people have more entrepreneurial potential, they face limited access to capital. In contrast, older entrepreneurs tend to have greater human, social, and financial capital, enabling them to 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 determinant to a strong credit score.[9]

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 models that incorporate timely information about repayment ability.

[7] Indeed, a guide to underwriting compliance provided by a federal agency begins with the following: *"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)

[8] 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.

[9] 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`.

3 Cash Flow-Based Lending to Small Businesses

In this section, we first describe how small business lending in the U.S. has changed in recent decades. This sets the stage for the advent of cash flow-based lending, which we use in analysis as an alternative to personal credit score-based underwriting.

Background on 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 often rely on informal accounting practices (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). As banks consolidated in the 1990s and 2000s, large banks expanded their market share and increasingly employed standardized risk scoring processes (Black and Strahan, 2002; Berger et al., 2005a; Strahan, 2017). This shift made lending more arm's length and made personal credit scores such as FICO a key input in underwriting (Arora, 2023; Carter and McNulty, 2005).

Following the Financial Crisis, higher capital requirements led banks to retreat from small business lending (Chen et al., 2017). The remaining small business lending at banks was held to high underwriting standards, requiring income measures drawn from annual financial statements and tax documents.[10] Firms which lacked strong historical financials were left underserved (Cole and Damm, 2020).[11] The retreat of larger banks from arm's length, non-credit card small business lending created a gap filled by new fintech companies, which entered the market in the early 2010s. They exploited new digitization and automation technologies to lend at arm's length.

Cash Flow-Based Underwriting Fintechs' most important and widely used new input was information from recent bank statements, referred to as cash flow data. Bank statements offer a direct measure of repayment ability; that is, whether the applicant can service debt based on inflows, outflows, and recent signs of distress. Bank statements are ubiquitous and accessible; nearly all businesses in the developed world have a checking account. Moreover, checking

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

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

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 generally use transaction account data to automate credit scoring for small businesses. 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 even have easy access to the information due to the siloed nature of the bank information systems" (Mills, 2024). Instead, banks typically relied on annual tax or accounting data, tied the loan to specific collateral, or focused exclusively on personal FICO scores. Similarly, when underwriting personal credit cards, banks even today rarely analyze customer checking accounts, relying instead on self-reported annual income. Conversations with practitioners indicate three reasons for the limited use of checking account information among traditional small business lenders: (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.[12]

Fintechs initially pulled data electronically from PDFs provided by customers or by "screen-scraping" bank websites, a process that was later digitized and commodified by Plaid, Yodlee, Ocrolus, and others. While inputs and modeling strategies vary widely across lenders, we understand that most use less than ten basic variables from bank statements, using credits and debits to capture inflows and outflows, and flagging transaction types that signal distress (e.g., insufficient funds transactions or overdrafts).[13] Cash flow data are now making their way into conventional lending. For example, in May 2024, both Experian and VantageScore launched new

[12]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.

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

consumer credit score products that incorporate bank statement-based cash flow attributes.[14]

Fintech lenders spurred the open banking movement, demanding that banks facilitate customer-permissioned sharing of financial information. Babina et al. (2025) document that by 2021, 80 countries had taken steps toward open banking, and in 2024 the U.S. established a regulatory structure set to take effect between 2026 and 2030, though it is contested. In a world with open banking and accessible AI that can interpret complex text, information from bank statements is likely to become a central input for assessing credit risk of individuals and businesses alike.

4 Data Sources and Summary Statistics

4.1 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 exclusively serve U.S. small businesses and request cash flow variables from bank statements as part of the application process.

The Platform. The Platform is a fintech company that connects prospective small business borrowers with lenders. Businesses apply through the Platform's website, which routes them to one or more of its 111 partner lenders. All applicants provide the same, standardized set of data as part of the initial application process, and all of these data are forwarded to the lenders. The lenders include fintechs and conventional banks. 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 use age or other protected class indicators. 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.

[14] 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 using "consumer permissioned open banking data." See Experian (2024) and VantageScore (2024).

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. We call one "Lender A" and the other "Lender B." Lender A provides both term loans and lines of credit.[15] Its underwriting process is highly automated and it uses a risk scoring model. Lender B only provides term loans. Lender A relies on an automated risk scoring model, whereas Lender B incorporates discretionary judgment by loan officers. Specifically, Lender B's underwriting process includes a manual component, with each case assigned to an individual underwriter Lender B also 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.[16]

Data Types and Collection. We obtained three types of data from all three companies: applications, third party underwriting inputs, and loan characteristics. Information about the applicant firm and owner is collected from the firm directly, such as owner name, address, owner age, firm founding date, and industry.[17] The owner first name is important for us to construct a gender indicator. Note that not all of this application data is used in underwriting (e.g., it would be illegal to use gender). We also observe underwriting inputs collected from third parties, specifically the FICO score from a credit bureau and bank statement information, provided by Plaid (Lender A) and Ocrolus (Lender B and the Platform).[18]

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 ultimately originated. Since we only observe performance (i.e., default) for the two Lenders and not for the platform, we collect firm survival data for a subset of the Platform applicants (introduced below). There are 104,150 applications and 18,434 loans for Lender A, from November 2013 to June 2024. For Lender B, there are a total of 58,668 applications and 25,762 loans, spanning August 2015 to May 2024. Finally, we observe 199,300 applications, 904,471 application-forwards, 191,421 offers,

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

[16] Lenders A and B as well as most lenders on the platform require a personal guarantee and take out a blanket UCC lien against business assets. Their loans are described as collateralized with non-revenue, non-real estate business assets, but in practice they do not verify collateral or secure the loan with specific collateral.

[17] 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".

[18] 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 three months of bank statements. Lender B acquires six months. Industry practice is typically three months, and rarely more than six.

and 35,821 originated loans on the Platform, spanning January 2022 to December 2023. Across all lenders, there are 262,723 unique firms.

4.2 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 observe loan offers and origination (but not performance). 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. Although we introduce specific variables below, all variables are defined in Appendix 8.

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 higher for Lender A and B applicants (46%) than for Platform applicants (21%). The average requested and originated amounts are between \$100,000 and \$150,000 across the different samples. Interest rates on originated loans average 16%. The higher rates in the approval-only columns reflect offers not taken up.[19] While some lenders on the platform provide term loans comparable to those offered by Lenders A and B, others focus on short-term loans with automatic ACH transfers, which have very high annualized interest rates.[20] Figure A.2 shows the number of applications over time for each of the three data sources. We construct an indicator variable for a non-performing loan (which we interchangeably call "default"). It 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. Lenders A and B have a 17% default rate.

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,

[19]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.

[20]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%.

where applicants with subprime, or <660, FICO scores are removed. 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. 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.[21] The second is the number of low or negative balances, which occur when the balance goes below zero or below \$1,000. 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 Lenders A and B, which reflect 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.

Entrepreneur Age. We measure owner age using the date of birth of the primary owner. The median age on the Platform is 43, close to the 42 years reported by \cite{azoulay2020age} for all new companies. in the U.S. For this reason, we focus on a threshold of 40 years to define "young" entrepreneurs, but also present results using thresholds of 35, 45, and 50.[22] Figure 2 shows that younger entrepreneurs appear to be at a disadvantage. They have a 21.5% loan approval rate, compared to 25% for older owners (Panel A), yet among borrowers they have identical default rates (Panel B).[23] Figure 3 shows that the chance of approval is strongly increasing in entrepreneur age and has a much steeper slope for younger applicants.

The disadvantage is largely tied to FICO. Younger applicants score 17 points lower on average– or 25% of a standard deviation–yet once they borrow, default rates are similar (Figure 2 Panel C). By contrast, cash flow metrics differ little by age. As one example, Figure 2 Panel D shows the number of insufficient funds transactions. We see that while approval clearly selects on this

[21] Some banks impose an insufficient funds fee and reject the transaction, while others permit the account to go negative and impose an overdraft fee.

[22] Figure A.3 shows histograms of age in the pooled applicant (Panel A) and originated loan (Panel B) samples.

[23] Panel A uses data from all three sources; Table A.2 shows a similar pattern in the Lender A & B data, where 35% of young applicants are approved, compared to 43% of older owners.

variable, there is little difference across age groups. Similarly, Table 2 shows that the disparity in credits is just 9% of the standard deviation, and the distress proxies are essentially the same across the age groups.[24] Table 2 also indicates that young owners are more likely to have younger firms (4.6 vs. 9 years) and smaller firms (6.3 vs 7.8 employees), but are otherwise similar across gender, socioeconomic status (determined by location), and business survival outcomes.

Finally, among low-FICO applicants (< 740), younger entrepreneurs actually show comparable or stronger cash flow profiles—for example, fewer overdrafts and MCA loans (Table A.3). This suggests that cash flow underwriting may be especially valuable in identifying creditworthy but FICO-constrained young owners.

Selection into the Data. The applicants in our data are selected 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. As mentioned above, the distribution of entrepreneur age is also representative.

5 Effect of Cash Flow Information on Credit Access

In this section, we study how exposure to cash flow-based underwriting affects access to credit. In two preliminary analyses, we first show that (a) in the full sample, cash flow variables predict loan approval decisions, indicating that lenders use them on average; and (b) the intensity of use varies substantially across individual lenders (Section 5.1). Leveraging this observed variation in lenders' reliance on cash flow information, we use a within-application design to provide causal evidence on the relationship between cash flow-based underwriting and access to credit for younger entrepreneurs. We assess whether application-forwards to lenders who more intensively use cash flow data result in a higher approval rate for younger applicants compared to older

[24]In Table A.2, we present these statistics for the Lender A and B subset. We see similar patterns across young and old, even though the applicant pool is overall lower risk. The average FICO is 718 for young owners vs. 732 for mature owners, while the cash flow variables are generally similar.

[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`.

applicants (Section 5.2). We supplement with a second causal design using data from Lender B, where applications are randomly assigned to individual loan officers who vary in their use of cash flow information (Section 5.3).

5.1 Cash-Flow Variables and Loan Approval

We first examine the degree to which cash flow variables explain loan approval decisions by constructing basic underwriting models with and without these variables. We do this for the overall sample and then for each Platform lender, forming the basis for our empirical design. The Baseline model regresses a loan approval indicator on FICO, firm size (employees), firm age, and fixed effects for industry, state, loan type, and 36 calendar quarters. The Cash Flow model is identical, but adds a vector of bank statement metrics: average credits, debits, balances, volatility, and distress flags for overdrafts and MCA withdrawals. We compare the models using both out-of-sample ROC-AUC in random forest models and OLS regressions.

The random forest machine learning algorithm that we implement (Ho, 1995), like OLS and other classical models, aims to minimize prediction error.[27] There are three steps to the estimation process. First is the selection of "hyperparameters" which act as knobs governing model behavior.[28] To select hyperparameters, we follow the standard technique in the machine learning literature of "tuning" (Probst et al., 2019). This is a Bayesian optimization procedure where hyperparameters are selected to maximize model performance (AUC-PR) across random subsamples of the data using k-folds cross-validation.[29] 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;

[27] 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.

[28] This is similar to the way an econometrician must select the bandwidth in a regression discontinuity design or the number of lagged periods in a vector autoregressive model. Examples of hyperparameters for random forests are the total number of trees in the forest, how many variables each tree considers simultaneously, and the number of samples required to split an internal node of a tree.

[29] Research in the field of machine learning presents this technique as a robust hyperparameter selection strategy. See, e.g., Kohavi et al. (1995); Snoek et al. (2012); Turner et al. (2021)

[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.

"bootstrapping" to obtain average estimates and standard errors. This estimation method is explained in more detail in Appendix C.

Across all three datasets, we find convincing evidence that on average, cash flow variables are used in underwriting. In Table 3, we show the OLS regressions in Panel A and the performance metrics from the ML models in Panel B. Moving from the Baseline to Cash Flow model improves the OLS R^2 by 9.3%. Using the more flexible machine learning model, the ROC-AUC increases from 0.803 to 0.816–a comparable improvement (4.2%) relative to a random-guess benchmark.[31] Column 2 in Panel A shows that cash flow variables predict approval in expected directions. For example, more credits to the account and higher bank balances are associated with higher chances of approval, whereas more withdrawals and higher volatility in revenues are associated with lower chances of approval. Daily pay loans and the number of days with low or negative balances also predict lower approval. There is no difference in the coefficients or R^2 when we split the sample by age (columns 3-4), consistent with the lenders not using different models for different age groups.

While these results indicate that the lenders on average use cash flow variables, there is significant heterogeneity. To quantify it, we re-estimate the random forest ML model for each lender, and examine the improvement in ROC-AUC between the Baseline and Cash Flow models. Figure A.5 depicts the distribution of ROC-AUC improvement. There is a broad range that includes negative numbers (i.e., the Cash Flow model performs worse). Some lenders may not use cash flow variables because their training samples are too small to support high-dimensional models without overfitting, they lack the sophistication to recognize their importance, or they rely on alternative technologies. Our results are consistent with the Platform's assertion that it works with both conventional and fintech lenders.

5.2 Within-Application Assignment to Cash Flow-Intensive Lenders

Having established that lenders vary in how intensively they use cash flow data in underwriting, we turn to estimating the causal effect of assignment to a relatively more cash flow intensive (CFI) lender. Note that 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. Our identification strategy addresses this challenge by exploiting within-applicant variation, following Khwaja and Mian (2008). The Platform typically forwards the same application to multiple lenders; within our analysis sample, the average and median are

[31] The calculation is (0.816-0.50)/(0.803-.50) where both models are benchmarked against an ROC-AUC of 0.5–the random-guess model.

five and three lenders, respectively, at the unique application level. By including application fixed effects in our regressions, we control for all observed and unobserved applicant risk characteristics, effectively isolating lender-side differences in underwriting practices. A lender who emphasizes cash flow metrics will estimate a different level of risk for the very same applicant than a lender who does not. Within-applicant regressions can identify whether that variation impacts approval for younger vs. older business owners.

Empirical Approach. To implement this, we identify lenders as CFI or not based on the incremental predictive power of cash flow variables for loan approval decisions. Specifically, using the random forest ML model described in Section 5.1, we train a Baseline model (FICO + controls) and a Cash Flow model (Baseline + bank statement variables) for approval decisions made by each lender with sufficient data. Lenders for whom the Cash Flow model yields a higher ROC AUC than in the Baseline model are classified as CFI.[32] Roughly 30% of the lenders have zero or negative values for AUC ROC improvement: the 70% with positive improvement are designated as CFI. Further details about the methodology for creating the cash flow-intensity measure are in Appendix C.

To ensure that the CFI lenders actually use cash flow data, we clerically researched the application and underwriting process for the subset of the lenders for which the Platform provided the lender name rather than only an ID number (53%). We found evidence online that all of these lenders use bank statements in underwriting (either via statement PDFs or integration with bank account connection platform such as Plaid), with one exception where the lender no longer offers small business loans.[33] We were also able to verify that some of the non-CFI lenders explicitly do not require bank statements, though the Platform sends all the lenders the same cash flow variables, which they may or may not use.

Summary statistics at the lender level by CFI status are in Table 4. The first four columns report statistics for all Platform lenders. The following four separate by CFI status, and the final column reports the difference in means between CFI and non-CFI lenders. Note that our within-application design means that it is not a concern if CFI and non-CFI lenders differ in the types of applicants they get or the loans they make. Offers from CFI lenders have much higher interest rates and

[32] We exclude lenders receiving under 50 application forwards or with an approval rate over 95%.

[33] Examples of evidence are statements such as the following from lender websites or application portals: “A representative will contact you and request three to four months business bank statements along with standard identification;" “Simply submit a one-page application and 3 months’ business bank statements to get funding within as little as 24 hours"; “Last, we will ask for a secure connection to your business bank account or for bank statements for the most recent 3 months;" and “bank statements are absolutely central to your loan application."

shorter maturities. Applicants are quite similar across the two groups. While those forwarded to CFI lenders have slightly lower owner FICO scores (683 vs. 697), this difference is significant only at the 10% level. We expand on this further below, where we explore whether the results reflect risk-taking, but it is worth noting that this difference is consistent with fintech lenders targeting borrowers with positive cash-flow attributes that mitigate the risk implied by their credit scores. We include statistics that confirm this pattern at the Application-Forward level in Appendix Table A.5. Summary statistics on the random forest model performance by CFI status are included in Table A.4 at the Lender, Application-Forward, and Application level. We see higher improvements in all performance metrics for CFI lenders.

We regress approval decisions for application i at time t from lender l on the interaction between the lender-level CFI_l indicator and the Young Owner_i indicator:

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

The model includes lender fixed effects (Lender_l) and application fixed effects (Application_i), which absorb all observable and unobservable attributes of the applicant that may influence approval decisions. 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. Standard errors are clustered by application.

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. The Platform forwards applications without conveying real-time information between lenders. Furthermore, lenders typically make 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.

Main Results. We present estimates of Equation 1 in the full sample in Table 5. The dependent variable in Panel A is an indicator for loan approval (i.e., an offer), which is our main outcome of interest. We use four age thresholds across the columns: 35, 40, 45, and 50. For the below-50 thresholds, we omit from the sample borrowers between that threshold and 50, so that we are not comparing very young and young owners. For example, in column 1, we compare under-35 owners with over-50 owners. Table 5 Panel A column 1 indicates that younger entrepreneurs (< 35) have a 2.5 pp higher chance of approval with CFI lenders, relative to older entrepreneurs and non-CFI lenders. This is 13% of the mean chance of approval, which is about 20%. The effect is similar at 12% of the mean using the 40 year old threshold. It effect declines slightly across the columns, to 10.5% of the mean using the 50 year old threshold (column 4). In Panel B, the dependent variable is the interest rate (APR) of the offer. The sample is restricted to applicants who receive at least two offers, so that we can compare APRs. Of course, a lower APR is better for the applicant, and this also limits the sample to a more creditworthy population. Column 1 indicates that younger entrepreneurs (< 35) enjoy a 1.8 pp lower APR from CFI lenders, which is 2.2% of the mean. The effect is similar across the columns.

In the Appendix, we report alternative specifications for Table 5. First, we include all applicants regardless of age in Table A.6. The results remain robust, albeit with lower magnitudes in the initial columns as we would expect, since the youngest are compared with the younger. Next, in Table A.7, we omit application fixed effects, instead employing industry, quarter, and lender fixed effects. The results are generally similar to those in Table 5. Last, we restrict the sample to applications forwarded to both CFI and Non-CFI lenders in Table A.8. The results are very close to the full-sample results.

Heterogeneity in FICO Score. Younger entrepreneurs with weaker or noisier FICO appear to benefit most from cash flow underwriting. While we cannot observe FICO accuracy, we test whether the causal effect of assignment to a CFI lender is larger among lower-FICO applicants. The lending industry defines score ranges from Poor-Fair (below 670) to Exceptional (above 800, which we call "Super").[34] In Panel A of Table 6, we split the sample into four groups using the industry-standard breakpoints and examine effects on approval (columns 1-4) and interest rate conditional on receiving at least two offers (columns 5-8). Using the 40-year threshold for

[34] In between Poor and Exceptional are two groups, one from 670-739 (sometimes called "Good" in the industry), and 740-799 (sometimes called "Very Good"). We call these "Low" and "High", respectively, for ease of interpretation. See, for example, here: `https://www.myfico.com/credit-education/credit-scores` or here: `https://www.experian.com/blogs/ask-experian/credit-education/score-basics/what-is-a-good-credit-score/`.

"Young," we find that young applicants with Poor or Low FICO scores are 2.4 pp more likely to be approved by CFI lenders, equal to about 16% of their mean approval rate (15–18%). The effect is smaller and statistically insignificant for High and Super FICO applicants, at roughly 6% of the mean. Results for APR are directionally similar but noisier: very poor-FICO applicants see interest rates about 5.5 pp lower, though this is only 5% of the relatively high baseline APR average.

In Panel B, we present results for all the age groups in triple interaction models. We focus here on the middle-range of FICO where most business applicants fall (670-799), and where we expect continuous FICO to relatively less informative. In the odd columns, we interact $\mathbb{1}(\text{Young Owner}_i)\times \mathbb{1}(\text{CFI}_l)$ with the continuous FICO score. For example, the coefficient in column 3 indicates that a single point higher FICO in this range is associated with a 0.03 pp lower chance of approval for young applicants at CFI lenders, relative to their counterpart groups. Column 3 shows that the results from Panel A columns 2-3 are significantly different. Here, we fully interact the fixed effects so that the difference in coefficients replicates the split sample. We do not find significant interaction effects for APR. While this may reflect the smaller sample, it could also suggest that young people with lower or less informative FICO scores benefit from cash flow-based lending primarily via more access to credit rather than the intensive margin decision of interest rate.

In Table A.9 we present alternative results where we split the entire FICO distribution into two groups, Low and High, and consider each of the age thresholds from 35 (columns 1-2) to 50 (columns 7-8). We report results for approval (Panel A) and APR (Panel B). Across all models, we see greater benefits among lower-FICO applicants, with the effects attenuating as we increase the age threshold.

Risk-taking as an Alternative Explanation. As noted in Section 3, fintech lenders have often targeted markets that are traditionally viewed as higher-risk. This could reflect simple regulatory arbitrage, or alternatively, the ability to identify "diamonds in the rough" using cash flow data. Note that by construction, a "diamonds in the rough" strategy requires targeting borrowers who appear risky using conventional metrics. The question is whether cash flow underwriting is associated with riskier loans in practice. We assess this in two ways.

The first test focuses on subsamples of lenders who are relatively conservative in their loan offers. In Table A.10 Panel A, we restrict the sample to lenders who approve no more than 10% of applicants, which is the average rate for non-CFI lenders when the unit of observation is the application-forward. Even though the sample is only about one quarter the size used in Table 5

Panel A, the effects are considerably *larger*. For example, using the 40 year age threshold, the effect is 4.6 pp relative to a mean of just 4.4%. In other words, among conservative lenders, assignment to a CFI lender doubles a young entrepreneur's chance of an approval. Panel B repeats the exercise using an alternative defintion of conservative lender–those for whom the 25th percentile FICO score for approved applications is above 669.[35] We see the results are similar or slightly larger than in Table 5 Panel A.

The second test examines whether CFI lenders are simply taking more risk. If this were the case, we would expect their loans to have higher default rates. Because we do not observe loan performance in the Platform, we proxy for it using business survival.. Clerical workers collected survival outcomes for a random subset of 46,000 unique applications (250,000 application-forwards and 43,000 approvals), checking whether the business remained open or appeared closed online as of September 2024.[36] Regressions predicting business survival are reported in Table 7. Column 1 shows that approval is strongly correlated with survival. However, approval by a CFI lender is no more or less predictive of survival than approval by other lenders. Next, in columns 2-5, we restrict the sample to approved application-forwards, and assess whether young people at CFI lenders are more or less likely to survive. The interaction coefficients are zero, instead of the negative coefficient we would expect if risk preferences explained the main finding. In sum, it is it not the case that the benefit to young people of assignment to a CFI lender reflects lender risk preferences. Below, we offer further confirmation using a within-lender supplementary analysis.

5.3 Random Assignment to Cash Flow-Intensive Loan Officers

We confirm our main finding using an instance of true random assignment drawn from 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.

Similar to our classification of lenders on the Platform, we identify individual loan officers at Lender B as CFI based on the predictive power of cash flow variables in their approval decisions

[35] This is the industry threshold separating Poor/Fair from Low, and is also very close to the average 25th percentile approved FICO score across lenders (which is 660 for CFI lenders, and 671 for non-CFI lenders).

[36] 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.

for applications they review. 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. Using a logit model, we estimate the incremental predictive power of cash flow variables (measured by ROC AUC improvement) for approval decisions made by each officer with sufficient observations. We use logit because there are not enough within-officer observations for an ML model to be useful. Since Lender B is generally cash-flow intensive, all officers show some improvement in ROC AUC; we classify loan officers in the top-quartile of this AUC improvement as CFI. The results are similar using other thresholds.

We present summary statistics about these data in Table A.11. Panel A reports data at the application level, while Panel B reports data at the unique loan officer level. 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.11 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. Importantly, there is no difference in default rate between CFI and non-CFI loan officers; in fact, the rate is slightly lower for CFI loan officers. Furthermore, CFI loan officers approve a somewhat lower share of applicants. Therefore, CFI loan officers do not take higher risk.

The estimation approach is as follows:

$$\begin{aligned}\mathbb{1}(\text{Approved}_i) = \beta_1 \mathbb{1}(\text{Young Owner}_i) + \beta_2 \mathbb{1}(\text{Young Owner}_i) \times \mathbb{1}(\text{CFI}_i) \\ + \text{Loan Officer}_k + \text{Industry}_i + \alpha_t + \varepsilon_{ik}\end{aligned} \tag{2}$$

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. Standard errors are clustered jointly by applicant and loan officer. 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. 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. Note also that as mentioned below, the default rates are the same across CFI and non-CFI loan officers, so there is no difference in risk preferences.

We present the results in Table 8, using the same structure as the Platform analysis.[37] Assignment to a CFI loan officer has a strong positive effect on approval, which declines as we raise the age threshold, consistent with the strongest benefits accruing to the youngest applicants. For example, the effect is 8.75 pp using the 35 year threshold (17% of the mean), and 7 pp using the 40 year old threshold (13.5% of the mean), and 3.6 pp using the 50 year old threshold (7% of the mean). Next, in Panel B, we divide the sample around median FICO.[38] Lower FICO scores explain the large effects among younger applicants. In sum, while this sample is smaller and limited to a single lender, it is comforting that we are able to essentially replicate the main finding.

6 Tail Analysis for Comparative Outcomes (TACO)

In Section 5, we provided causal evidence that younger entrepreneurs disproportionately benefit when lenders adopt underwriting models that incorporate bank statement cash-flow variables. Importantly, increased access to credit by young owners does not come at the expense of increased business failure, a proxy for default risk. In this section, we argue that cash-flow intensive lenders serve young entrepreneurs more effectively than traditional lenders because cash-flow variables are especially informative for this group. Timely bank information allows lenders to identify "diamonds in the rough" that a FICO-centric model would miss.

We introduce a novel methodology—Tail Analysis for Comparative Outcomes (TACO)——that can be broadly applied in both research and policy settings to analyze the outcome of switching models on a specific population group. Unlike methods that rely on identifying sources of exogenous variation, TACO evaluates who benefits or is harmed when new predictive variables are incorporated into a binary outcome model without an explicit counterfactual. It does this by focusing on observations in the tails of the distribution of predicted outcomes between two models, where the predictions diverge most dramatically. While we apply TACO specifically to younger versus older entrepreneurs, this method generalizes easily and can identify disparate impacts for any subgroup of interest.

We use TACO to compare traditional models to those enhanced by cash-flow variables, highlighting specifically how model predictions differ for younger and lower FICO entrepreneurs. By illuminating these subgroup-specific impacts, TACO provides practical insights that complement the causal analysis of Section 5 and establishes a useful framework for broader

[37] As above, we also adjust the sample to compare a group of young owners to over-50 owners in each column. The full sample results are similar, reported in Table A.12.

[38] We divide FICO around the median due to the smaller sample and higher FICO scores in this population.

policy evaluations of alternative underwriting models.

6.1 Cash-Flow Informativeness and Age

Why are cash-flow variables particularly useful in identifying young creditworthy borrowers? We explore this question by assessing the informativeness of these variables in predicting default separately for old and young entrepreneurs. If the Cash Flow model adds more predictive power for young owners, it would indicate that bank statement variables are particularly useful in identifying young "diamonds in the rough".

To test the differences in informativeness we turn again to our random forest models using the Baseline and the Cash Flow specifications described in 5.1, but now we use default as the outcome variable. Our focus will be on the differences in out-of-sample performance between the two models for each age group. A larger ROC-AUC improvement between the Cash Flow and Baseline model indicates that the bank statement variables are particularly important in assessing risk for that age group. We bootstrap the differences and report statistical significance in the difference in ROC-AUC between the models.

The results in Table 9 show that young owners see substantially higher gains in predictive performance when cash-flow variables are added to the Baseline model. For example, owners that are younger than 35 have an AUC of 0.608 in the Baseline model and see an increase of 0.030 (or 27.8%) with the addition of cash-flow variables. Whereas, for those older than 35 the improvement is only 0.008 (or 1.2%) with a Baseline AUC of 0.671. While the differences are statistically significant for every age group, the economic significance is substantially higher for the younger age groups regardless of where the age split occurs. We also report alternative measures of model accuracy, ROC-AUC of the precision recall curve and the H-measure, all of which lead to the same conclusions. We plot the curves in Figure A.6.

In addition to the ML results, we use OLS regressions to show in Table 10 the direction and informativeness of cash flow variables on default and APR. Panel A indicates that for the sample of originated loans, cash flow variables predict loan interest rates and default. In Panel B, we separate by age. Columns 1-2 show that there is little difference in R-squared for APR, consistent with age not being used directly in underwriting. Columns 3-4 confirm the results from Table 9, showing a much larger increase in informativeness from adding cash flow variables for younger entrepreneurs. Relative to Panel A column 3, the increase is 58% for younger owners and 36% for older owners.

As mentioned above, these results may reflect personal credit scores being less accurate in predicting default risk for younger people; research suggests that this noise—driven by data errors and “thin” credit files—is larger for minorities and low-income people (Blattner and Nelson, 2021; Federal Trade Commission, 2012; Avery et al., 2004). In the case of age, the source of noise is mechanical; as explained in Section 2, traditional credit scores are based on a history of repaying debt; while they explicitly favor a longer history, our results show for the first time that they are also more accurate when there is a longer history. A caveat is that our results pertain to small business loans, though there is no reason to think the same would not be true for personal loans.

6.2 Tail Analysis for Comparative Outcomes (TACO)

Next we introduce a novel method, TACO, which 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 binary target outcome, y, from a vector of observed characteristics X. Note that f and g could differ with respect to model types (e.g., OLS vs. ML), 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. Formally, let $z \subset X$ be the characteristic of interest and define

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

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.

TACO measures the reallocation across subgroups implied by a single lender switching underwriting models. Specifically, when z is a binary variable (e.g., "Young Owner") and the number of total loans is fixed, the TACO ratio measures reallocation of the two groups $z = 1$ and $z = 0$ across the tails. However, TACO is not strictly zero-sum for two reasons. First, some applicants with relatively small changes in predicted risk will not fall in the tails (when tails are defined as less than 50%). Most importantly, TACO does not capture general equilibrium changes in the number of lenders or their underwriting methods. To the degree that new data expands the observed pool of creditworthy borrowers, we would expect higher overall credit supply either because new lenders enter or existing lenders expand.

6.3 Implementing TACO in Our Setting

We use TACO to evaluate how switching from the Baseline to the Cash Flow (CF) model to predict default affects younger entrepreneurs. We follow the procedure in Section 5 to fit Baseline and CF models to predict default probabilities for holdout samples. Then 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. Next, 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 40 years old). We then 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 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 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.[39]

Table 11 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.[40] 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 1.37. In other words, the group that most benefits from the CF model has 37% 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 1.34, 1.21, and 1.12. We plot these results in Figure 4, 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 11, are uniformly much larger. For example, the TACO ratio for owners below the age of 40 is 2.08. This indicates that among low-FICO borrowers, young entrepreneurs are over two 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

[39]We also examine the overall distributions of predicted probabilities for the Baseline and CF models. As shown in Figure A.7, 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.

[40]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.8.

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 0.57 for the under-40 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 owners (2.08) is larger than the harm for high-FICO young owners. Also, in this high-FICO group, the Cash Flow model is relatively better for younger people; Table A.14 shows that the TACO ratio is just 0.34 for the high-FICO group overall. The CF model disadvantage for young borrowers with high FICO is smallest for the youngest groups, increasing with the age threshold.

To ensure that the results do not reflect spurious repeated sorting of winners and losers, we conduct a placebo test replacing the Baseline model with the CF model (in other words, we compare two identical CF models). The results, in Figure A.9, are statistically indistinguishable from 1.0, indicating no effect. 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.[41] 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.15, using the same format as Table 11. 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.10 and Table A.14. 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).[42] Young firms might benefit from recent cash flow data because even if they have a "thin" credit file or operating history, they may

[41]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.

[42]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).

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.14. The CF model strongly benefits low-FICO women.[43]

7 Applying TACO to Approval Decisions

In this section, we connect the two parts of our paper: the causal analysis showing how young people benefit from assignment to a cash flow-intensive (CFI) lender, and the TACO results showing how incorporating cash flow variables in standard models affects default predictions across subgroups. Here, we apply TACO's predicted improvement from the cash flow (CF) underwriting model to approval decisions among Platform lenders. Specifically, we re-estimate the causal regressions using indicators from TACO, testing whether applicants whom the CF model classifies as less risky are more likely to be approved at CFI lenders.

This exercise serves two purposes. First, it validates the CFI lender measure and shows that TACO is relevant to underwriting in practice: Entrepreneurs predicted to benefit from the CF model we use in TACO—trained on Lender A & B data—are preferred by CFI lenders out of sample in the Platform data. Second, it offers some insight into the channel behind our main results. Credit scores could disadvantage younger people in two ways. One is that they may be less precise for younger people. Consistent with this, Section 6.1 documented a larger increase in default prediction power from adding cash flow variables for younger relative to older people (see Table 9). If this precision channel were the only mechanism at play, we expect the TACO tails—"Most Hurt by CF" (highest h_i) and "Most Helped by CF" (lowest h_i)—to be symmetrically impacted in underwriting. That is, when assigned to a CFI lender, the most hurt tail should face relatively lower chances of approval, while the most helped tail should face better chances.

The other channel is mechanical bias in FICO against younger people. This is supported by the

[43] 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.

strong correlation between FICO and age (see Figures 1 and A.1), as well as by the differentially positive effect on loan approval of assignment to CFI lenders for younger people that we showed in Section 5.2 (see Table 6). In this bias channel, we expect the "Most Helped by CF" tail to benefit, but would not necessarily expect the "Most Hurt by CF" tail to suffer. Those disadvantaged by CF models may look creditworthy enough on other dimensions, like FICO, that poor cash flows do not preclude them from receiving a loan offer. In other words, if cash flow variables benefit young applicants primarily by compensating for downward bias in FICO, we expect to see asymmetry in how correcting for this bias (independent of age) affects loan approval.

Empirical Approach. For each application forwarded to at least one lender on the Platform, we compute the difference h_i in predicted default probability between the CF model (g) and the baseline model (f):

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

Both g and f are trained using data from Lender A and Lender B, as in Section 6.3. A positive h_i indicates that the CF model views the borrower as riskier than the baseline model; a negative h_i implies the opposite.

We link these h_i values to actual approval decisions using regressions that compare, within the same application, the relative likelihood of approval from CFI versus non-CFI lenders. As above, we include application fixed effects to absorb all borrower-level characteristics and lender fixed effects to capture general lender risk tolerance, so identification comes solely from within-application variation in whether the offer comes from a CFI lender. We construct indicator variables representing the TACO tails: the group with the highest h_i is "Most Hurt by CF" while the group with the lowest h_i is "Most Helped by CF". We also employ the continuous variable h_i. These are interacted with a lender-level variable indicating CFI vs. non-CFI lenders.

Note that if TACO partitions by risk and if lenders in the platform are not selecting applicants directly based on age independently of its relationship to risk, then we should find no separate impact of h_i on approval for young owners relative to old owners. Instead of age, what should matter is which tail an applicant is sorted into by TACO. We test a triple interaction specification that adds an indicator for young owners to the interaction term to act as a kind of a placebo test.

Results. The results are presented in Table 12. The first three columns use thresholds of 10%, 20%, and 30%, respectively, to define the tails ("Most Helped by CF") and ("Most Hurt by CF"). An applicant in the bottom 10% of h_i ("Most Helped by CF") is 1.7 pp (8%) more likely to be approved for a loan by a CFI lender than a non-CFI lender. The estimates using a 20% or 30% threshold for the tails are similar—1.4 pp (6%) and 1.1 pp (5%), respectively. All of these estimates are statistically significant. Across tail specifications, applicants in the "Most Hurt" group are less likely to be approved, but the coefficients are smaller and noisier, with no statistical significance. The final model in column 4 shows that higher h_i values lead to significantly lower relative approval rates from CFI lenders. Note that the coefficient on h_i would not be identified with application fixed effects; this model serves to test whether our result persists when we do not estimate within-application.[44]

These findings confirm that the borrower reclassification signal embedded in TACO not only predicts differences in default risk across models but also aligns closely with the actual approval behavior of lenders we classify as cash-flow intensive. In this way, it offers comforting validation of both the causal empirical approach and the TACO method by showing that TACO predictions can be applied out of sample to the lender approval decision. Furthermore, we find some evidence of asymmetry, in that that the inclusion of cash flow variables appears to help overcome downward bias in FICO for some applicants, but does not measurably cause CFI lenders to more often reject those for whom the reverse is true, where the applicant's FICO score is much better than their cash flow variables.

8 Conclusion

We hypothesize that younger entrepreneurs are mechanically disadvantaged by underwriters' reliance on personal credit scores. This is because these scores favor a long history of repaying debt. In the U.S., personal credit scores rise precipitously for people in their 40s and 50s, on average, which is precisely when we see the steepest increases in entrepreneurship. In this paper, we explore the connection between these facts by studying whether younger entrepreneurs benefit from cash flow-based underwriting, in which lenders incorporate recent information from business checking accounts. We focus on cash flow information because of its rising prevalence in underwriting, but we anticipate that any information that predicts default and which is orthogonal to traditional credit scores would yield similar results for young entrepreneurs.

[44] We confirm that there is no interaction effect with age in Table A.16. We report the results of a triple interaction specification where a binary variable for "Young Owner" (owner under 40 years of age) is interacted with the indicators for "CFI" and either "Most Hurt/Helped by CF" or the continuous measure h_i. Note that the application fixed effects absorb the entrepreneur's age, so any interaction effect would reflect CFI lender preference for young that varies with the benefit of cash flow data. As expected, we find no statistically significant change in approvals among young owners over older owners conditional on already being in the most helped or hurt tails.

Interest in cash flow-based underwriting has grown 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 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 (Mills and McCarthy, 2016).

We show that cash flow information helps to predict default when added to simple, standard underwriting models. 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 through increased chances of approval for credit. 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.

References

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

Agarwal, S, S Alok, P Ghosh, and S Gupta, 2021, Financial inclusion and alternate credit scoring: Role of big data and machine learning in fintech, *SSRN, Rochester, NY* .

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.

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.

Avery, Robert B, Paul S Calem, and Glenn B Canner, 2004, Credit report accuracy and access to credit, Technical report, Federal Reserve Board.

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 Bahaj, Greg Buchak, Filippo De Marco, Angus Foulis, Will Gornall, Francesco Mazzola, and Tong Yu, 2025, Customer data access and fintech entry: Early evidence from open banking, *Journal of Financial Economics* 169, 103950.

Babson, 2022, Global entrepreneurship monitor (gem) u.s. report by babson college.

Bakker, Trevor J, Stefanie DeLuca, Eric A English, James S Fogel, Nathaniel Hendren, and Daniel Herbst, 2025, Credit access in the united states, Technical report, National Bureau of Economic Research.

Balyuk, Tetyana, Allen N Berger, and John Hackney, 2025, What is fueling fintech lending? the role of banking market structure, *The Review of Corporate Finance Studies* cfae028.

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.

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

Ben-David, Itzhak, Mark J. Johnson, and René M. Stulz, 2025, Models behaving badly: The limits of data-driven lending, *Review of Finance* 29, 711–745.

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, Aurel Hizmo, and Daniel Ringo, 2025, How much does racial bias affect mortgage lending? evidence from human and algorithmic credit decisions, *The Journal of Finance* 80, 1463–1496.

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.

Blattner, Laura, Jacob Hartwig, and Scott Nelson, 2022, Information design in consumer credit markets, *American Economic Review* .

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.

Chatterjee, Satyajit, Dean Corbae, Kyle Dempsey, and José-Víctor Ríos-Rull, 2023, A quantitative theory of the credit score, *Econometrica* 91, 1803–1840.

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 Scharfstein, 2024, Racial disparities in the paycheck protection program, *Journal of Financial Economics* 160, 103911.

Chioda, Laura, Paul Gertler, Sean Higgins, and Paolina C Medina, 2024, Fintech lending to borrowers with no credit history, Technical report, 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.

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.

Einav, Liran, Mark Jenkins, and Jonathan Levin, 2013, The impact of credit scoring on consumer lending, *The RAND Journal of Economics* 44, 249–274.

Erel, Isil, and Jack Liebersohn, 2022, Can fintech reduce disparities in access to finance? evidence from the paycheck protection program, *Journal of Financial Economics* 146, 90–118.

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 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.

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.

Federal Trade Commission, 2012, Report to Congress under Section 319 of the Fair and Accurate Credit Transactions Act of 2003.

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

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.

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

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.

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.

Harris, 2024, Rising entrepreneurs: Pathways to small business formation, *Justworks and The Harris Poll* .

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.

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.

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

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, Simon Down, and Maria Minniti, 2014, Ageing and entrepreneurial preferences, *Small Business Economics* 42, 579–594.

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.

Kohavi, Ron, et al., 1995, A study of cross-validation and bootstrap for accuracy estimation and model selection, in *Ijcai*, volume 14, 1137–1145, Montreal, Canada.

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.

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

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.

Probst, Philipp, Anne-Laure Boulesteix, and Bernd Bischl, 2019, Tunability: Importance of hyperparameters of machine learning algorithms, *Journal of Machine Learning Research* 20, 1–32.

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.

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

Robb, Alicia M, and David T Robinson, 2014, The capital structure decisions of new firms, *The Review of Financial Studies* 27, 153–179.

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.

Snoek, Jasper, Hugo Larochelle, and Ryan P Adams, 2012, Practical bayesian optimization of machine learning algorithms, *Advances in neural information processing systems* 25.

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.

Turner, Ryan, David Eriksson, Michael McCourt, Juha Kiili, Eero Laaksonen, Zhen Xu, and Isabelle Guyon, 2021, Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the black-box optimization challenge 2020, in *NeurIPS 2020 Competition and Demonstration Track*, 3–26, PMLR.

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.

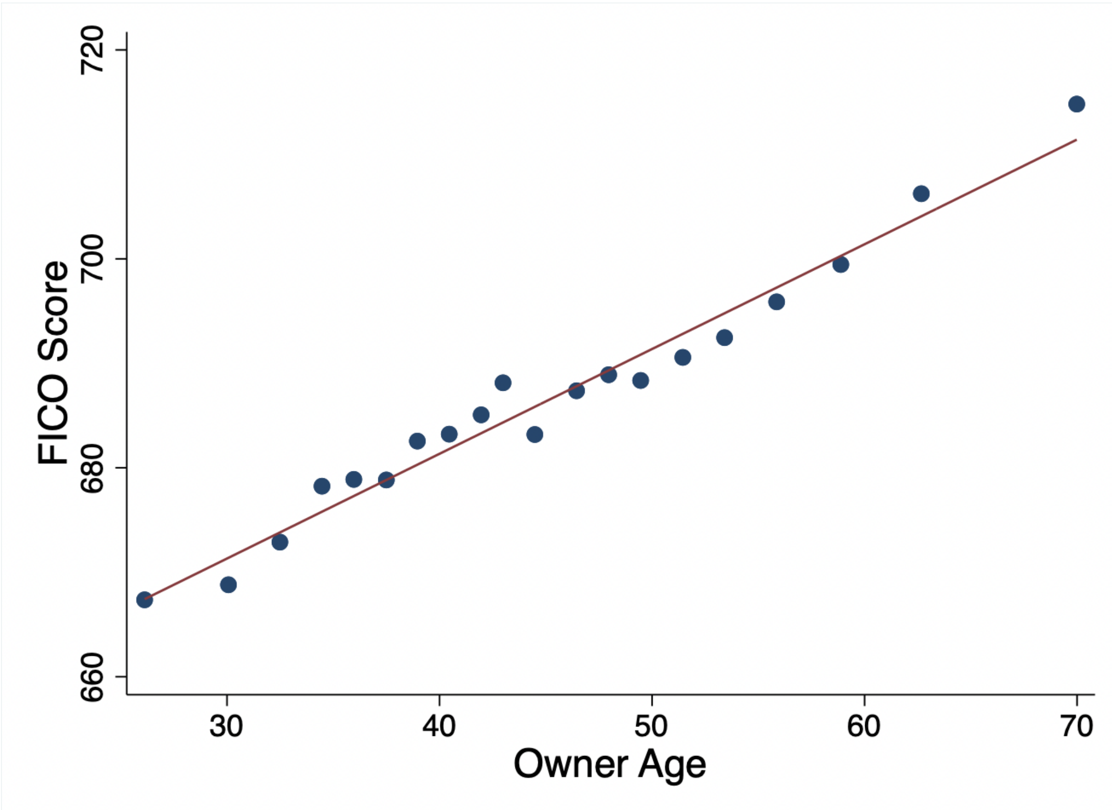

Figure 2: Variation in Loan Outcomes and Inputs by Entrepreneur Age

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 40, while mature owners are at least 40. Each bar includes 95% confidence intervals using standard deviation.

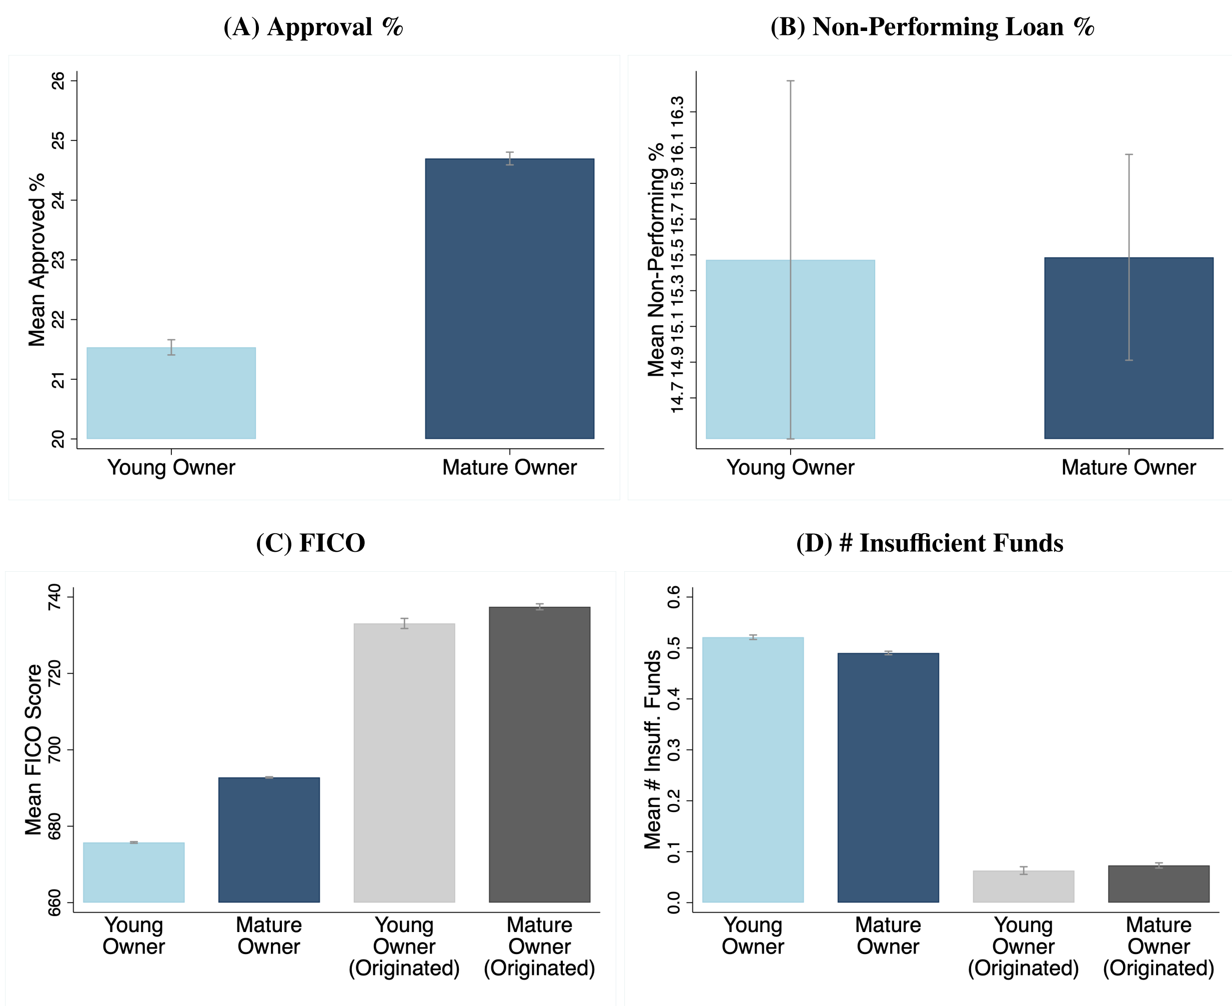

Figure 3: Approval and Owner Age

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 40 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.

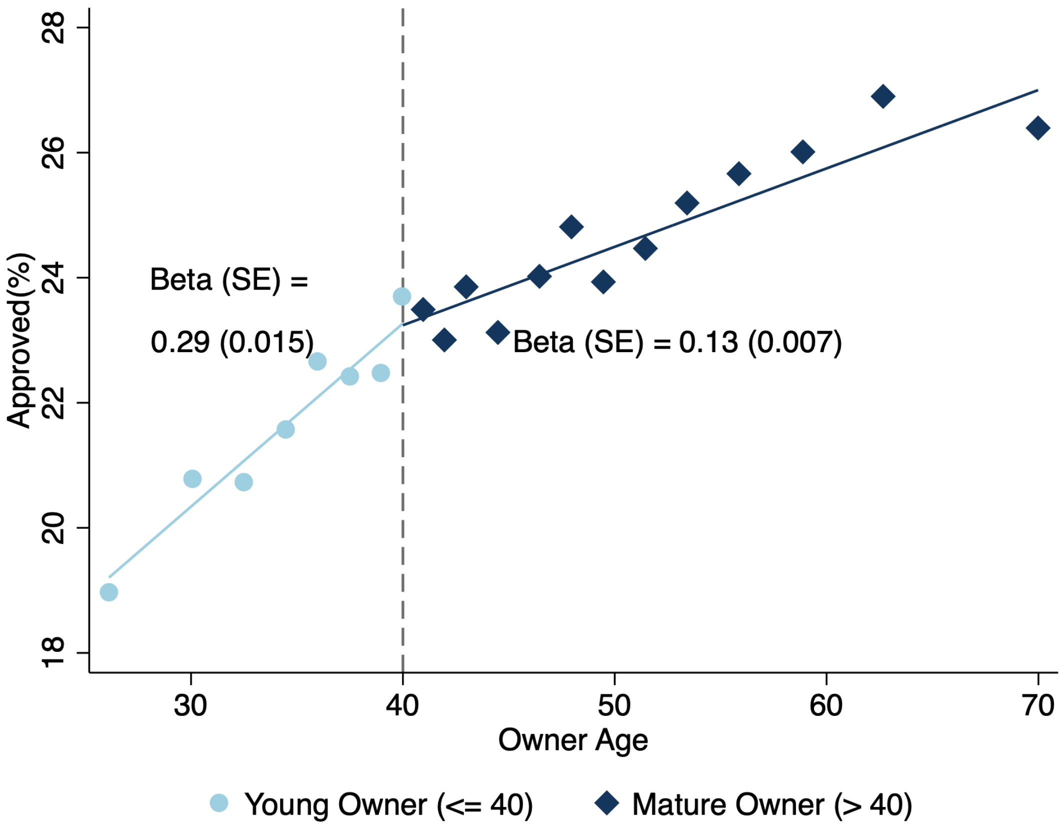

Figure 4: Comparison of Benefit from Cash Flow Model vs. Baseline Model by Age Group

This figure plots the Tail Analysis for Comparative Outcomes (TACO) ratios for young owners from the first set of estimates in Table 11 (see table notes and Section 6 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.

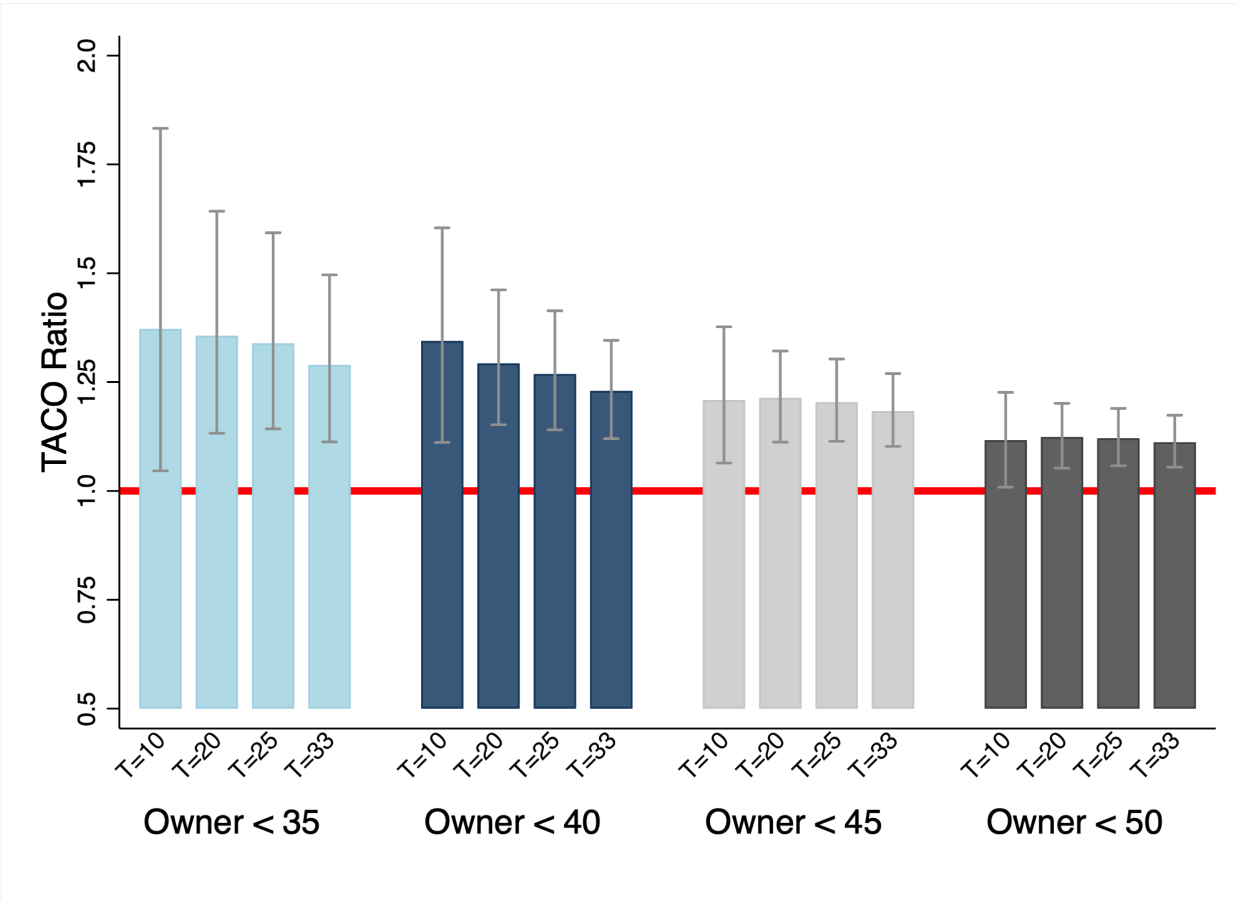

Table 1: Summary Statistics

This table contains summary statistics about loan applications and originated loans from the three sources of data (Platform company and Lenders A & B). The Platform statistics are at the application-forward level so that an application often enters the sample more than once since it is being evaluated by multiple lenders.

	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		
Number of Employees	904,471	7.09	4.00	10	162,818	8.88	5.00	12	38,021	10	6.00	13
Business Survival	250,695	0.35										
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		

Table 2: Summary Statistics by Owner Age

This table compares business owners who are young (<40 years old) with those who are more mature (≥ 40). 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 (≤ 40)				Mature Owner (> 40)				Young Owner (≤ 40)				Mature Owner (> 40)			
	N	Mean	Median	SD	N	Mean	Median	SD	N	Mean	Median	SD	N	Mean	Median	SD
Loan Variables:																
Approved (%)	399,103	22			628,734	25			5,009	100			15,181	100		
Requested Loan Amount (Th$)	398,915	134	75	178	628,460	149	100	188	5,009	77	60	60	15,181	79	70	62
APR (%)	63,931	86	78	61	118,570	68	58	55	4,532	15	15	4.96	13,889	15	15	4.84
Originated (%)	399,103	5.06			628,734	6.54			5,009	100			15,181	100		
Non-Performing Loan (%)									5,009	15.5			15,181	15.5		
Originated Loan Amount (Th$)									5,009	72	59	56	15,181	77	60	59
Loan Maturity (Years)									5,009	2.45	2.08	1.19	15,181	2.51	2.08	1.21
Credit Score & Cash Flow (Bank Statement) Variables:																
FICO	399,103	676	675	66	628,734	693	691	67	5,009	733	731	48	15,181	737	734	49
Credits (Th$)	399,103	89	42	142	628,734	103	50	157	5,009	103	58	136	15,181	113	68	142
Balance (Th$)	399,103	29	9	59	628,734	35	12	66	5,009	36	17	59	15,181	37	17	58
(#) Insuff. Funds	399,103	0.52	0.00	1.42	628,734	0.49	0.00	1.40	5,009	0.06	0.00	0.27	15,181	0.07	0.00	0.32
(#) Low or Neg. Bal.	399,103	0.78	0.00	1.94	628,734	0.73	0.00	1.90	5,009	0.48	0.00	1.19	15,181	0.46	0.00	1.14
Withdrawals (Th$)	372,660	90	43	145	562,914	106	51	161	677	118	57	183	2,062	128	64	186
Credits (less new debt) (Th$)	368,556	76	36	123	554,912	89	43	136	398	161	88	194	1,358	161	90	188
$\mathbb{1}$(Daily Pay Loan)	397,970	0.32			625,658	0.34			4,932	0.04			14,887	0.05		
S.D. Credits (Th$)	34,609	10	5.10	13	89,039	12	6.28	14	5,009	13	7.36	13	15,181	13	7.86	14
S.D. Balance (Th$)	34,609	4.26	1.85	5.72	89,039	4.94	2.35	6.11	5,009	5.40	3.02	5.93	15,181	5.52	3.15	5.99
Borrower Characteristics:																
Owner Age	399,103	34	35	4.35	628,734	52	50	8.27	5,009	35	35	3.94	15,181	53	52	8.63
Female	387,110	0.27			615,946	0.28			4,993	0.23			15,148	0.23		
Business Age (Years)	399,103	4.62	3.42	4.30	628,734	9	6.33	7.99	5,009	6.40	4.69	5.41	15,181	12	9	7.83
Young Firm (< 5)	399,103	0.67			628,734	0.39			5,009	0.51			15,181	0.24		
Number of Employees	399,103	6.30	4.00	9	628,734	7.82	4.00	11	5,009	8.38	4.00	12	15,181	10	7.00	12
Business Survival	95,083	0.34			155,494	0.36										
Pct Black Pop (%)	399,103	15	7.10	19	628,734	13	6.00	18	5,009	12	5.90	16	15,181	12	5.50	16
Per Capita Income	398,280	36	33	16	627,373	38	34	18	5,006	38	34	18	15,163	39	35	18

Table 3: Predicting Loan Approval in OLS and ML

This table shows how credit score, cash flow, and borrower characteristics predict loan approval using data from Lender A, Lender B, and the Platform (N = 1,067,289). Panel A presents results using an OLS model. Columns 1-2 shows results for the full population and columns 3-4 show the results split by owner age. 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. Panel B presents our performance evaluation of the Baseline and Cash Flow random forest models for predicting loan approval. Additional details on the model are found in Appendix 8. Performance metrics are calculated as the mean of 1,000 bootstrap iterations. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: OLS

Dependent Variable:	Approved (%)			
	All		Owner (< 40)	Owner (> 40)
	(1)	(2)	(3)	(4)
FICO	7.14***	5.61***	5.74***	5.46***
	(0.50)	(0.33)	(0.40)	(0.30)
Credits (Th$)		4.40***	3.14**	4.75***
		(0.81)	(1.09)	(0.88)
Withdrawals (Th$)		-4.58***	-3.12**	-5.02***
		(0.86)	(1.21)	(1.00)
Balance (Th$)		0.45*	0.17	0.52*
		(0.24)	(0.21)	(0.25)
$\mathbb{1}$(Daily Pay Loan)		-2.39***	-1.96***	-2.67***
		(0.40)	(0.34)	(0.45)
(#) Low or Neg. Bal.		-5.43***	-4.55***	-5.77***
		(0.58)	(0.44)	(0.60)
(#) Insuff. Funds		-0.12	-0.26	-0.02
		(0.17)	(0.17)	(0.18)
S.D. Credits (Th$)		-1.03***	-0.60	-0.89**
		(0.23)	(0.40)	(0.31)
S.D. Balance (Th$)		0.57**	0.84**	0.77***
		(0.21)	(0.31)	(0.23)
Business Age (Years)	0.20***	0.21***	0.35***	0.17***
	(0.02)	(0.02)	(0.04)	(0.02)
Number of Employees	0.05***	0.04**	0.06***	0.04**
	(0.01)	(0.02)	(0.02)	(0.02)
Observations	1,067,272	1,067,272	399,091	628,718
Industry FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.205	0.224	0.212	0.206
R-squared Improv.		9.4		
Y-mean	24.93	24.93	21.53	24.70

Panel B: Machine Learning

	ROC AUC	**AUC-PR**	**H-Measure**
FICO Model	0.803	0.615	0.295
Cash-Flow Model	0.816	0.639	0.323
Difference	0.013***	0.023***	0.027***

Table 4: Summary Statistics by Cash Flow-Intensive Lender Status

This table reports summary statistics from the Platform (N = 879,889) comparing lenders categorized by their intensity of cash flow-based underwriting (CFI). CFI lenders are those for which including bank statement-based cash flow variables significantly enhances the predictive accuracy of loan approval models relative to the Baseline FICO-driven model. The classification relies on improvements in model performance metrics such as ROC-AUC from a random forest classifier. We compare lender-level data by CFI intensity using a two-sample t-test. P-values are based on a two-sample t-test of the FICO and Cash Flow model performance metrics. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

	All				CF Intensive Lender		Not CF Intensive Lender		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Loan Variables:									
Approved (%)	76	21.2	20	16	53	21.4	23	20.7	0.71
# Forwards	76	6.53	7.14	3.32	53	7.33	23	4.69	2.64***
Requested Loan Amount (Th$)	76	171	159	75	53	177	23	158	19
APR (%)	74	59.0	53	45	52	69.3	22	34.7	35***
Originated (%)	76	4.19	3.22	4.60	53	3.87	23	4.90	-1.03
Originated Loan Amount (Th$)	75	84	62	78	52	89	23	72	18
Loan Maturity (Years)	75	2.11	1.00	2.44	52	1.66	23	3.12	-1.46**
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	76	687	686	33	53	683	23	697	-13.65*
Credits (Th$)	76	120	114	65	53	124	23	111	13
Balance (Th$)	76	39.8	37	20	53	40.4	23	38.5	1.88
(#) Insuff. Funds	76	0.61	0.54	0.34	53	0.59	23	0.65	-0.07
(#) Low or Neg. Bal.	76	0.64	0.55	0.39	53	0.62	23	0.71	-0.10
Withdrawals (Th$)	76	121	115	65	53	125	23	111	13
Credits (less new debt) (Th$)	76	100	96	55	53	104	23	92	11
$\mathbb{1}$(Daily Pay Loan)	76	0.38	0.38	0.11	53	0.39	23	0.35	0.04
Borrower Age:									
Owner Age	76	44.4	44	1.99	53	44.4	23	44.3	0.15
Young Owner (< 35)	76	0.23	0.22	0.07	53	0.22	23	0.23	-0.00
Young Owner (< 40)	76	0.41	0.40	0.08	53	0.40	23	0.41	-0.01
Young Owner (< 45)	76	0.58	0.57	0.07	53	0.58	23	0.58	-0.01
Young Owner (< 50)	76	0.72	0.72	0.05	53	0.72	23	0.72	-0.00
Other Borrower Characteristics:									
Female	76	0.29	0.28	0.06	53	0.29	23	0.28	0.00
Business Age (Years)	76	7.08	7.07	1.61	53	7.14	23	6.93	0.22
Young Firm (< 5)	76	0.51	0.52	0.13	53	0.51	23	0.53	-0.02
Number of Employees	76	8.07	7.58	3.24	53	8.28	23	7.57	0.71
Business Survival	72	0.37	0.41	0.14	51	0.39	21	0.33	0.07*
Pct Black Pop (%)	76	13.3	13	2.79	53	13.4	23	13.3	0.07
High Pct Black Pop (> 6%)	76	0.50	0.51	0.06	53	0.51	23	0.50	0.01

Table 5: Within-Application Effect of Assignment to Cash Flow-Intensive Lender

This table uses data from the Platform to test the effect of assignment to a cash flow-intensive lender on approval and offered interest rate (APR) for young entrepreneurs. 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 dependent variable in Panel A is approval and in Panel B it is the APR. The coefficient on the interaction term "Young=1 × CFI=1" represents the difference in approval likelihood or APR 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: Effect on Loan Approval

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	2.53***	2.36***	2.19***	2.08***
	(0.31)	(0.27)	(0.25)	(0.25)
Observations	418,454	564,934	710,357	829,167
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: Effect on the Interest Rate

Dependent Variable:	APR (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	-1.79**	-2.03***	-1.84***	-1.44***
	(0.84)	(0.59)	(0.50)	(0.46)
Observations	51,934	70,435	89,285	105,129
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.878	0.878	0.879	0.878
Y-mean	81.16	82.31	82.20	81.78

Table 6: Within-Application Effect of Assignment to Cash Flow-Intensive Lender by FICO Range

This table uses data from the Platform to test the effect of assignment to a cash flow-intensive lender on approval and percent interest rate (APR) for young entrepreneurs. The level of observation is the application-lender, so an applicant may appear multiple times as the application is forwarded to multiple lenders. In Panel A columns 1-4, the outcome is approval, which is an indicator multiplied by 100, so the coefficient represents a percentage point change in the chance of approval. In Panel A columns 5-8, the outcome is the APR within loan offers, conditional on the applicant receiving at least two loan offers. The FICO ranges reflect industry standard thresholds, which we call Poor (<670), Low (670-739), High (740-799), and Super (>799). In Panel B, we present triple interactions with the continuous FICO score, and then with Low FICO as defined in Panel A, restricting to the sample with Low and High FICO scores. Lender FE are interacted with Low FICO so that the difference between coefficients is the same as in Panel A. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Split Sample

Dependent Variable:	Approved (%)				APR (%)			
FICO Range:	Poor	Low	High	Super	Poor	Low	High	Super
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young Owner (< 40)=1 × CFI=1	2.44***	3.91***	1.46**	1.98	-5.52***	-1.66*	-0.99	-2.20
	(0.40)	(0.43)	(0.62)	(1.39)	(1.62)	(0.86)	(0.88)	(1.71)
Observations	251,610	208,532	104,792	24,799	22,159	30,002	18,259	4,771
Application FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.330	0.304	0.337	0.357	0.849	0.868	0.874	0.865
Y-mean	15.55	21.47	25.68	28.42	108.25	76.94	59.69	53.85

Panel B: Triple Interaction within Low-High FICO Range

Dependent Variable:	Approved (%)							
Young Owner (Young) Defined As:	< 35		< 40		< 45		< 50	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1 × FICO	-0.03**		-0.03***		-0.02*		-0.01	
	(0.01)		(0.01)		(0.01)		(0.01)	
Young=1 × CFI=1 × Low FICO=1		2.13**		2.61***		1.61**		1.35*
		(0.95)		(0.82)		(0.76)		(0.73)
Observations	216,991	216,991	288,520	288,520	362,228	362,228	424,268	424,268
Application FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.312	0.312	0.313	0.313	0.314	0.314	0.314	0.313
Y-mean	22.283	22.283	22.400	22.400	22.606	22.606	22.719	22.719

Table 7: Role of Risk in Cash Flow-Intensive Lender Approval

This table uses data from the Platform to examine whether CFI lenders tend to approve riskier businesses. We proxy for risk with survival as of September, 2024, which is available for a random subset of applicants. In column 1, the level of observation is an application-forward. We do not use application fixed effects because the outcome of business survival varies at the applicant level. In columns 2-5, the sample is restricted to approved application-forwards. Quarter FE control for the time of application. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Dependent Variable:	Business Survival				
Application Sample:	All	Approved (%)			
		Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)	(5)
CFI=1 × Approved	-0.01				
	(0.01)				
Young=1 × CFI=1		0.00	0.00	0.01	0.02
		(0.03)	(0.02)	(0.02)	(0.02)
Approved	0.07^{***}				
	(0.01)				
Young Owner		0.01	0.00	-0.01	-0.01
		(0.03)	(0.02)	(0.02)	(0.02)
Observations	240,532	43,919	43,919	43,919	43,919
Lender FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.009	0.008	0.008	0.008	0.008
Y-mean	0.358	0.403	0.403	0.403	0.403

Table 8: 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. Panel A uses the full sample and Panel B splits by FICO score. In odd (even) columns, the sample is restricted to applicants with a FICO score above (below) the median. 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. Standard errors are clustered by quarter and approver. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Effect on Loan Approval

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	8.75***	6.98***	4.97**	3.63*
	(3.10)	(2.43)	(2.05)	(1.83)
Young=1	-8.53***	-6.14***	-4.78***	-3.63**
	(2.26)	(1.71)	(1.55)	(1.39)
Observations	6,550	7,970	9,726	11,535
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Loan Officer FE	Yes	Yes	Yes	Yes
R-squared	0.060	0.057	0.056	0.054
Y-mean	52.09	52.06	51.90	51.99

Panel B: Effect on Loan Approval by FICO Score

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

Table 9: Machine Learning Model Performance

This table presents our performance evaluation of the Baseline and Cash Flow random forest models for predicting loan default overall and 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.17. 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 1,000 bootstrap iterations. Definitions of the performance metrics *ROC AUC*, *AUC-PR* and *H-Measure* are provided in Section 8; 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**
Overall			
FICO Model	0.642	0.255	0.071
Cash Flow Model	0.657	0.272	0.086
Difference	0.015***	0.017***	0.016***
Under 35			
FICO Model	0.608	0.220	0.074
Cash Flow Model	0.638	0.254	0.106
Difference	0.030***	0.034***	0.033***
Over 35			
FICO Model	0.671	0.264	0.104
Cash Flow Model	0.679	0.261	0.108
Difference	0.008***	-0.003***	0.004***
Under 40			
FICO Model	0.633	0.219	0.076
Cash Flow Model	0.648	0.229	0.090
Difference	0.015***	0.011***	0.014***
Over 40			
FICO Model	0.663	0.263	0.098
Cash Flow Model	0.663	0.269	0.101
Difference	-0.000	0.005***	0.004***
Under 45			
FICO Model	0.651	0.239	0.087
Cash Flow Model	0.652	0.232	0.087
Difference	0.001***	-0.006***	0.000
Over 45			
FICO Model	0.662	0.270	0.102
Cash Flow Model	0.659	0.267	0.098
Difference	-0.003***	-0.003***	-0.004***
Under 50			
FICO Model	0.641	0.243	0.079
Cash Flow Model	0.662	0.258	0.098
Difference	0.020***	0.015***	0.019***
Over 50			
FICO Model	0.642	0.236	0.083
Cash Flow Model	0.651	0.245	0.090
Difference	0.009***	0.008***	0.007***

Table 10: Predicting Loan Outcomes in Regressions

This table shows how credit score, cash flow, and borrower characteristics can predict loan outcomes of APR and default. Panel A shows results for the full population, Panel B shows the results split by owner age. Columns 1-2 use data on originated loans from Lender A, Lender B, and the Platform (N = 76,367) to predict APR. Columns 3-5 use data on originated loans from Lender A and Lender B (N = 38,021) to 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. Standard errors are clustered by industry and quarter. R-squared improvement in Panel A represents the percent improvement in R-squared from the baseline model (columns 1 and 3). ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: All Observations

Dependent Variable:	APR (%)		Is Non-Performing (%)		
	(1)	(2)	(3)	(4)	(5)
FICO	-6.83***	-6.07***	-4.37***	-4.11***	-3.10***
	(0.18)	(0.18)	(0.21)	(0.22)	(0.22)
Credits (Th$)		0.49		-2.72***	-1.73***
		(0.92)		(0.36)	(0.37)
Withdrawals (Th$)		-3.90***		1.14***	-3.06***
		(0.94)		(0.33)	(0.36)
Balance (Th$)		-0.20		-1.23***	-0.75**
		(0.18)		(0.32)	(0.32)
𝟙(Daily Pay Loan)		-0.84***		0.97***	0.77***
		(0.15)		(0.24)	(0.25)
(#) Low or Neg. Bal.		3.14***		1.58***	1.71***
		(0.36)		(0.22)	(0.23)
(#) Insuff. Funds		0.38*		0.05	0.14
		(0.21)		(0.21)	(0.21)
S.D. Credits (Th$)				2.16***	1.96***
				(0.33)	(0.33)
S.D. Balance (Th$)				1.29***	0.78**
				(0.33)	(0.34)
Business Age (Years)	-0.46***	-0.44***	-0.17***	-0.17***	-0.16***
	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Number of Employees	-0.16***	-0.01	-0.07***	-0.05***	-0.04**
	(0.01)	(0.02)	(0.01)	(0.02)	(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	33,925	33,925	38,019	38,019	36,252
Industry FE	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	No	No	No
Loan Type FE	No	No	Yes	Yes	Yes
R-squared	0.830	0.833	0.038	0.043	0.073
R-squared Improv.		0.4		13.8	91.4
Y-mean	81.62	81.62	16.72	16.72	17.30

Panel B: Sample Split by Age Group

Dependent Variable:	APR (%)		Is Non-Performing (%)	
	Owner (< 40)	Owner (> 40)	Owner (< 40)	Owner (> 40)
	(1)	(2)	(3)	(4)
FICO	-6.07***	-6.03***	-3.81***	-3.60***
	(0.32)	(0.22)	(0.52)	(0.30)
Credits (Th$)	0.51	0.32	-2.54***	-1.78***
	(1.77)	(1.07)	(0.87)	(0.51)
Withdrawals (Th$)	-5.40***	-2.88***	4.21***	3.16***
	(1.78)	(1.09)	(1.50)	(0.95)
Balance (Th$)	0.28	-0.39**	-2.65***	-1.20***
	(0.35)	(0.20)	(0.73)	(0.43)
$\mathbb{1}$(Daily Pay Loan)	-0.82***	-0.75***	0.27	0.82***
	(0.26)	(0.18)	(0.46)	(0.28)
(#) Low or Neg. Bal.	3.31***	3.11***	2.04***	1.32***
	(0.63)	(0.44)	(0.62)	(0.37)
(#) Insuff. Funds	0.49	0.35	-0.33	-0.12
	(0.34)	(0.26)	(0.54)	(0.28)
S.D. Credits (Th$)			2.12**	1.52***
			(0.82)	(0.43)
S.D. Balance (Th$)			1.36*	0.83*
			(0.75)	(0.46)
Business Age (Years)	-0.67***	-0.35***	-0.06	-0.18***
	(0.06)	(0.02)	(0.09)	(0.04)
Number of Employees	-0.03	-0.01	-0.02	-0.07***
	(0.03)	(0.02)	(0.04)	(0.03)
Observations	13,603	20,308	5,006	15,178
Industry FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	No	No
Loan Type FE	No	No	Yes	Yes
R-squared	0.829	0.835	0.066	0.052
Y-mean	91.03	75.33	15.48	15.49

Table 11: 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 6). 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. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

	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)	709,342	0.131	346,996	0.110	362,346	0.150	1.372**
Young Owner (< 40)	709,342	0.265	346,996	0.226	362,346	0.303	1.344***
Young Owner (< 45)	709,342	0.417	346,996	0.377	362,346	0.456	1.209***
Young Owner (< 50)	709,342	0.567	346,996	0.535	362,346	0.598	1.117**
Low FICO (< 740):							
Young Owner (< 35)	709,342	0.098	346,996	0.070	362,346	0.124	1.767***
Young Owner (< 40)	709,342	0.179	346,996	0.116	362,346	0.240	2.078***
Young Owner (< 45)	709,342	0.282	346,996	0.191	362,346	0.368	1.923***
Young Owner (< 50)	709,342	0.382	346,996	0.273	362,346	0.487	1.782***
High FICO (≥ 740):							
Young Owner (< 35)	709,342	0.033	346,996	0.039	362,346	0.026	0.666**
Young Owner (< 40)	709,342	0.086	346,996	0.110	362,346	0.063	0.571***
Young Owner (< 45)	709,342	0.136	346,996	0.186	362,346	0.088	0.473***
Young Owner (< 50)	709,342	0.185	346,996	0.262	362,346	0.111	0.424***

Table 12: Applying TACO to Approval Decisions

This table examines the effects of the predicted default probabilities generated by the Cash Flow (CF) and Baseline models on applicant approval. The CF and Baseline models are trained using data from Lender A and Lender B, following the methodology in Section 6.3. The trained models are used to make out-of-sample default probability predictions on all applicants to the Platform forwarded to at least one lender. The difference $h_i = g(X_i) - f(X_i)$ is then used in an OLS regression to predict applicant approval, where g and f represent the CF-based underwriting model and the Baseline model, respectively. Here, "Most Hurt by CF" ("Most Helped by CF") indicates the top (bottom) tail of h_i by percentile within bootstrap iteration. Interaction rows report the difference for CFI lenders using the CFI classification from Section 5. Columns 1–3 include lender and application fixed effects. Column 4 reports a linear specification in h_i without application fixed effects but with controls equivalent to the Baseline model. Standard errors are based on the bootstrap distribution across 100 iterations; ***, **, * denote p-values below 0.01, 0.05, 0.10, respectively.

Dependent Variable:	Approved (%)			
	10% Tails (1)	20% Tails (2)	30% Tails (3)	Continuous h_i (4)
Most Hurt by CF × (CFI = 1)	-1.17	-0.90	-0.75	
	(1.10)	(0.64)	(0.48)	
Most Helped by CF × (CFI = 1)	1.73***	1.42***	1.11***	
	(0.53)	(0.44)	(0.36)	
Continuous h_i				-26.62***
				(4.37)
Continuous h_i × (CFI = 1)				-36.65***
				(4.58)
Observations	880,747	880,747	880,747	1,005,544
Lender FE	Yes	Yes	Yes	Yes
Application FE	Yes	Yes	Yes	No
Bootstrap Iterations	100	100	100	100
Baseline Controls	No	No	No	Yes
Y-mean	21.62	21.62	21.62	22.58

Appendix (For Online Publication)

Appendix 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

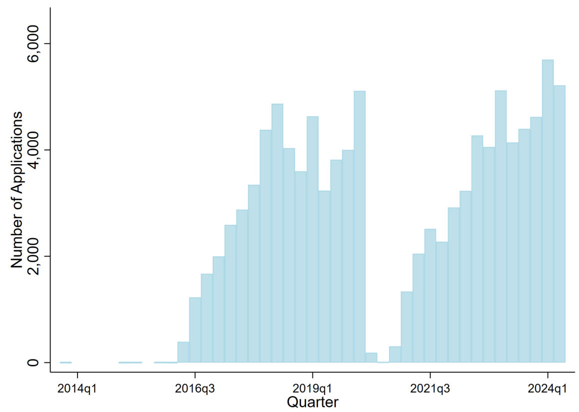

(B) Lender B

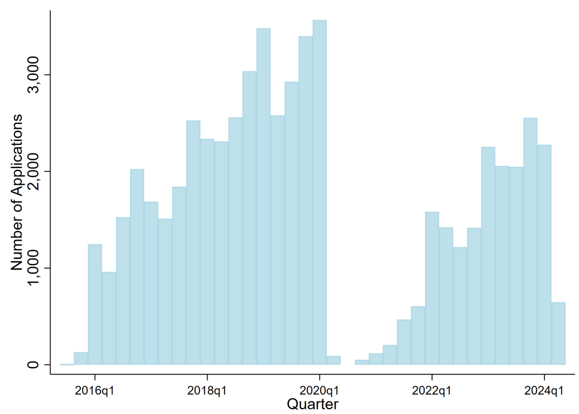

(C) Platform

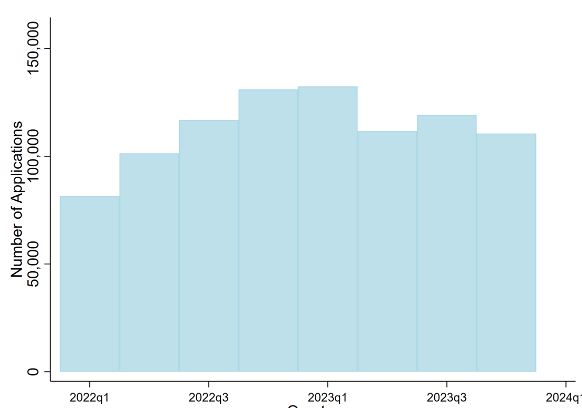

Figure A.3: Entrepreneur Age Distributions

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

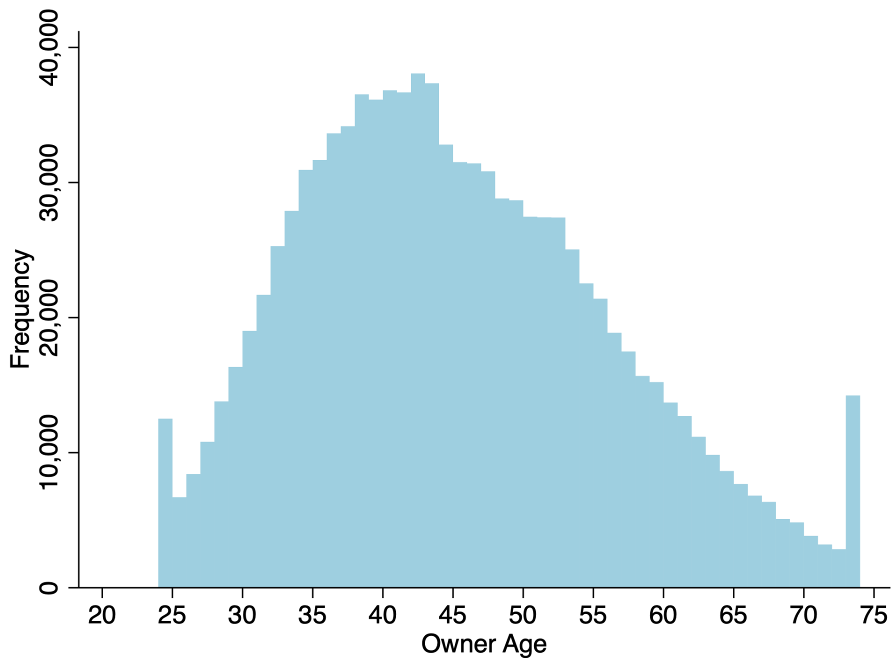

(B) Lender A and B Sample

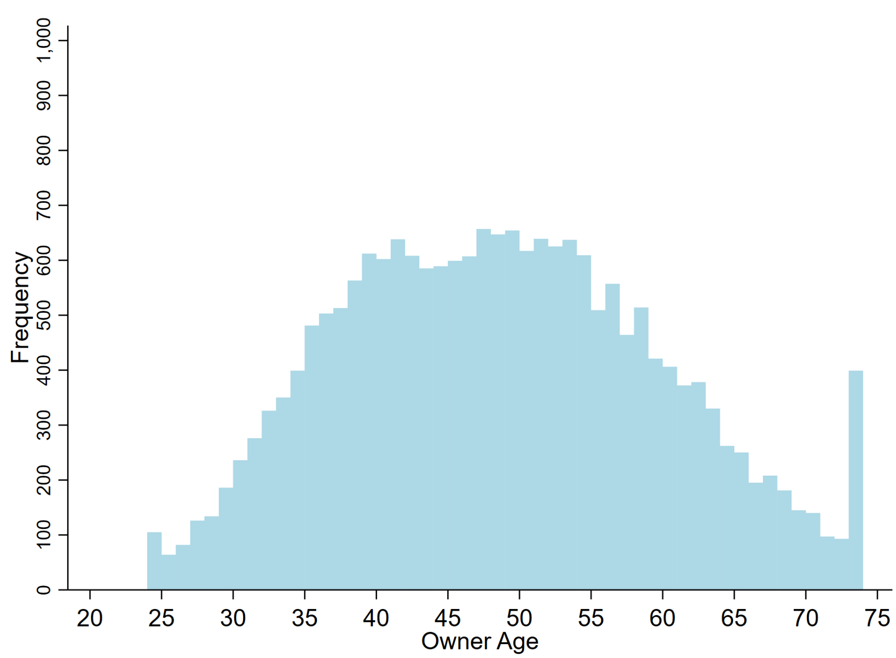

Figure A.4: Approval Model Performance

This figure uses data from Lender A, Lender B, and the Platform (N = 1,067,289) to report improvements in loan approval model performance from the baseline to the cash flow model. These statistics are reported in Table 3. Bars 1 to 2 show the improvement in R^2 using the OLS model. Bars 3 to 4 show the improvement in AUC using the machine learning model.

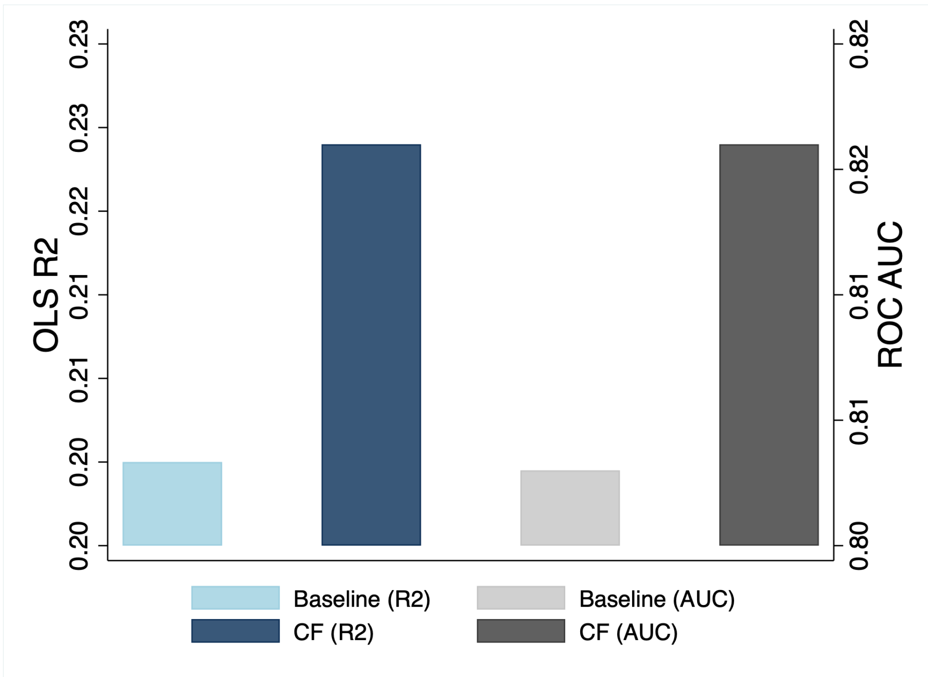

Figure A.5: AUC Improvement Distribution

This figure shows the distribution of AUC improvement from the FICO to the Cash Flow model for the 76 lenders in the Platform data. Lenders described as cash flow intensive have an improvement greater than 0 and are to the right of the red line.

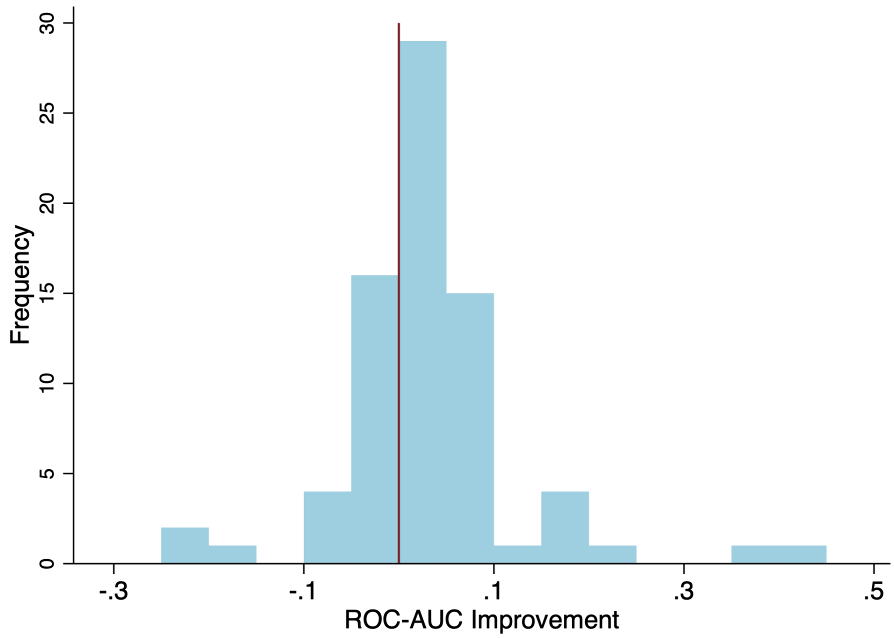

Figure A.6: Performance of Loan Default Prediction Models

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 9. In Figure A, we plot the receiver operating characteristic (ROC) curve for the full sample. In Figure B and Figure C, we plot the ROC curve for Young (< 40) and Mature ($\geq$ 40) Owners individually.

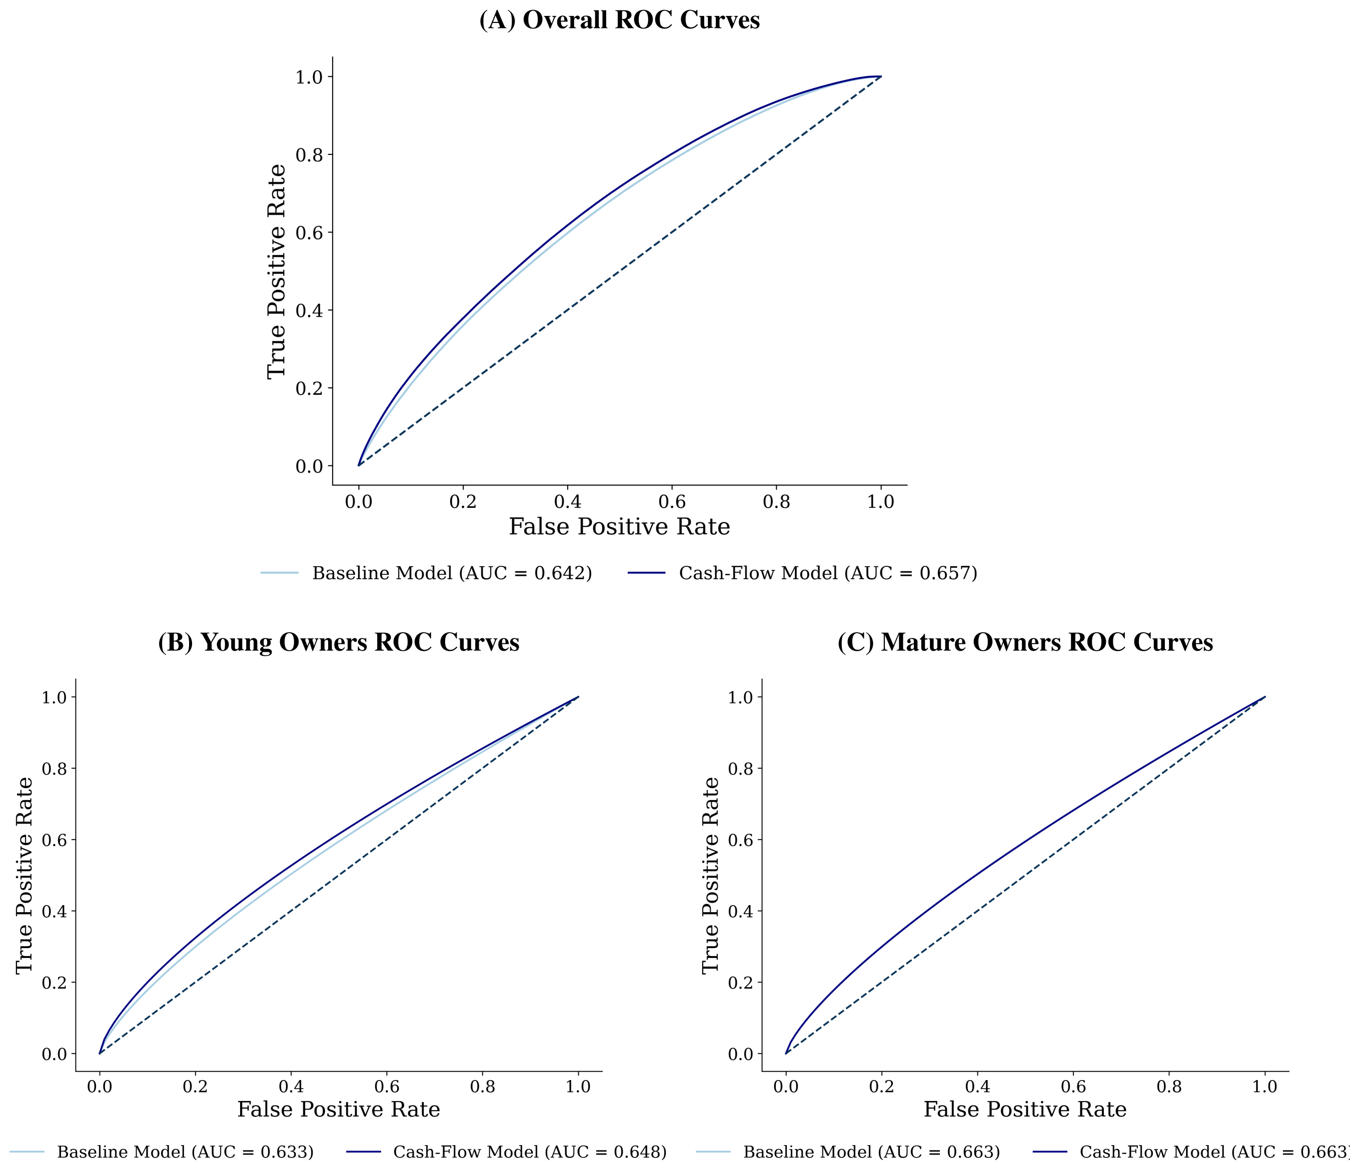

Figure A.7: Histogram of Baseline Model and CF Model Prediction Probabilities

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.

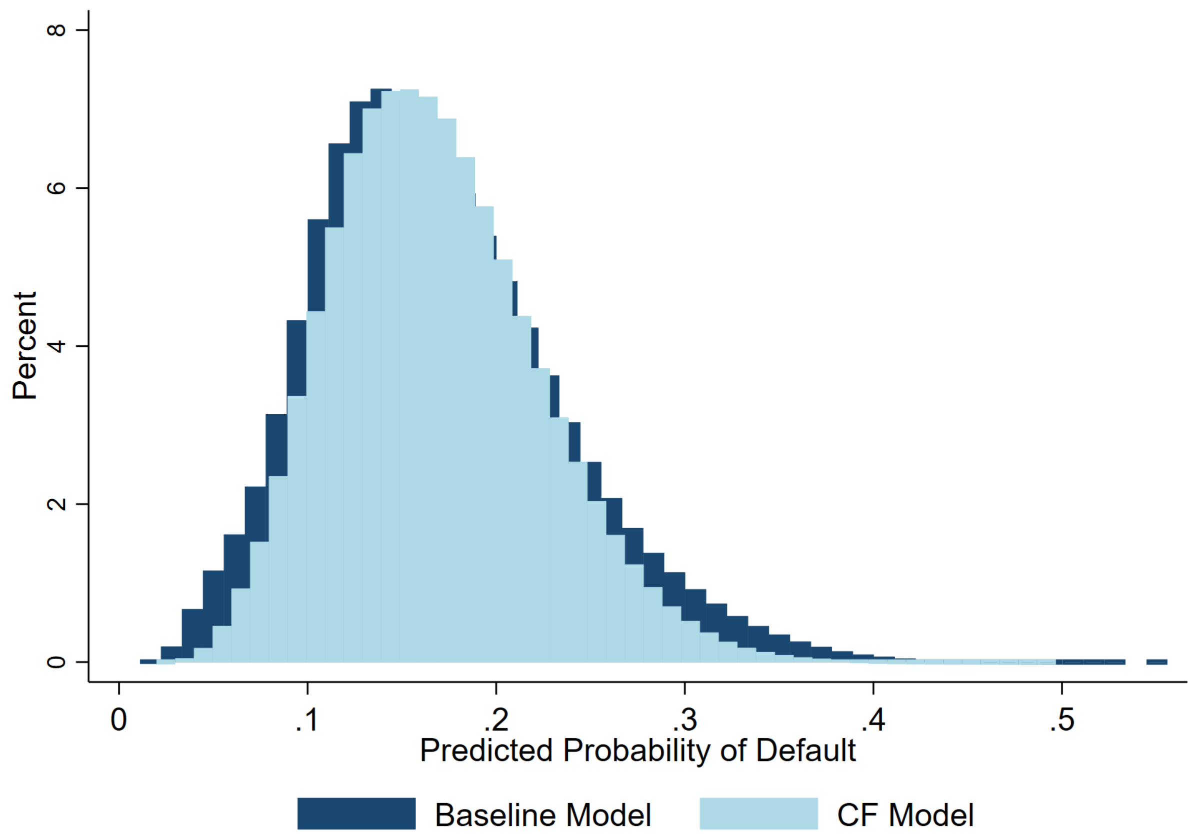

Figure A.8: Default Difference Distributions

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 6. 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.

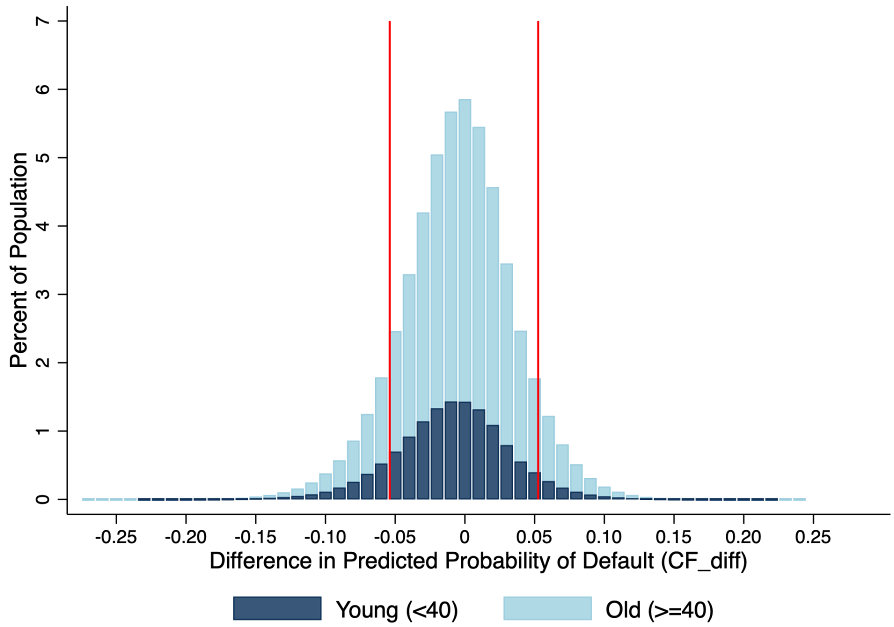

Figure A.9: Placebo Test of TACO Methodology by Age Group

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 6 for details). This is similar to Figure 4 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 1,000 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.

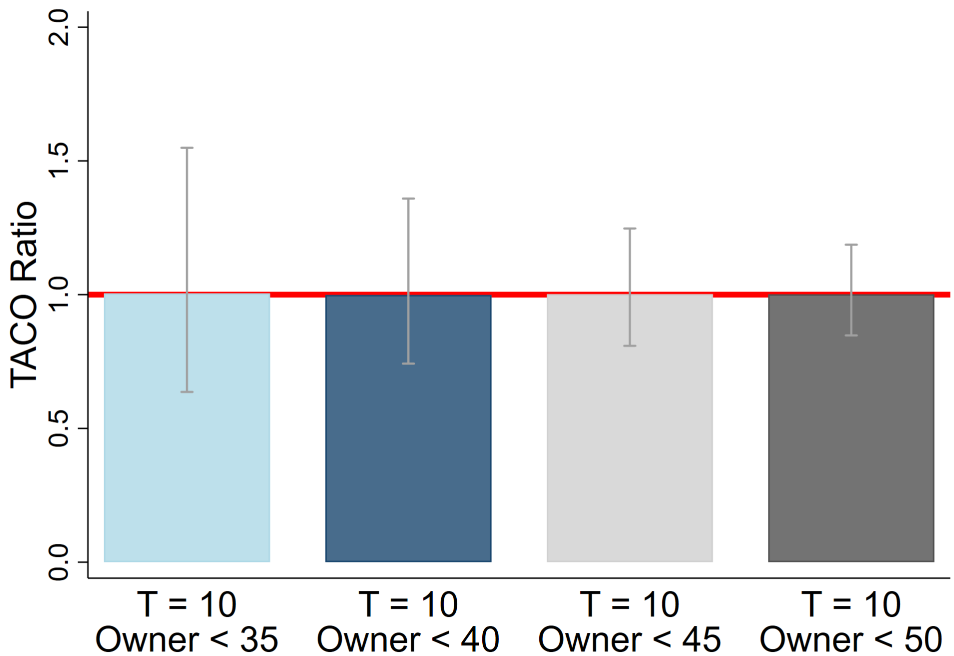

Figure A.10: Comparison of Benefit from Cash Flow Model vs. Baseline Model for Other Characteristics

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 6 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) FICO Score

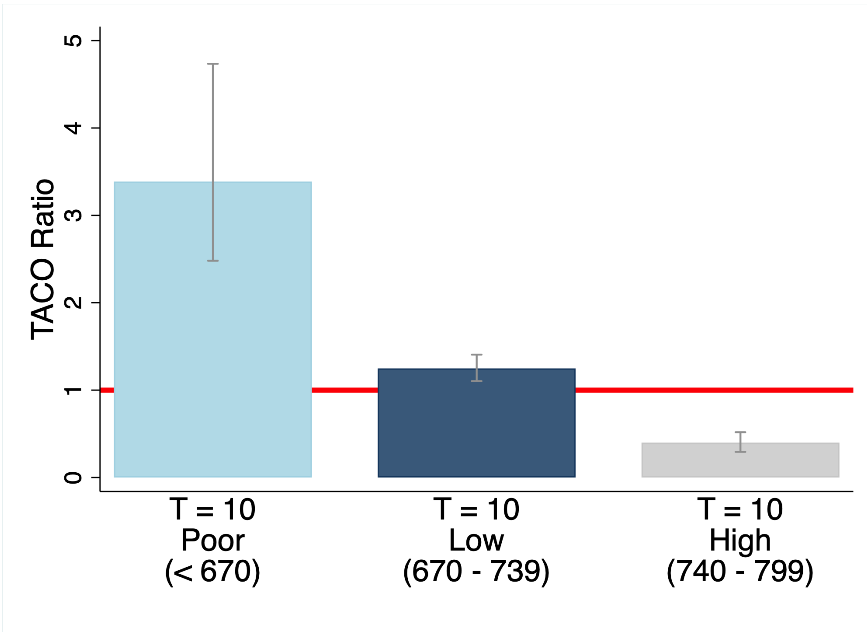

(B) Other Characteristics

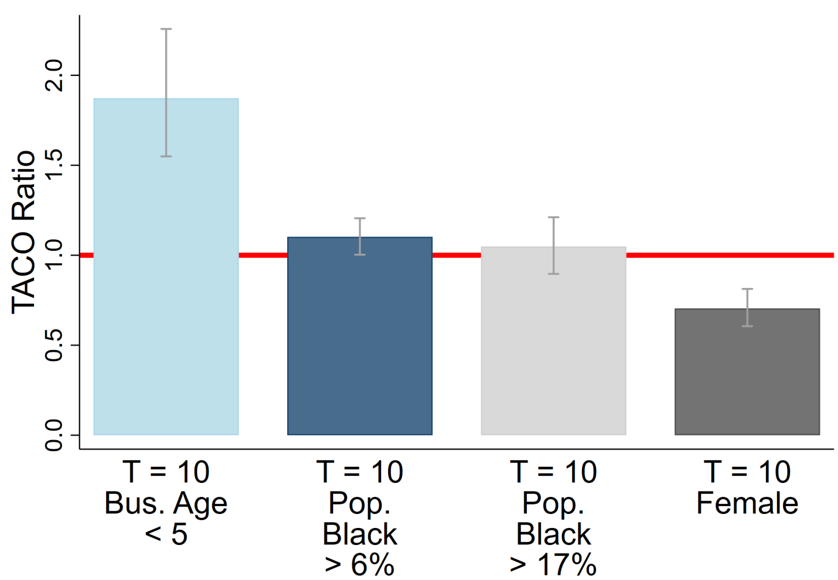

Figure A.11: Bootstrap Estimates of Differences in Predicting Default

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.17. 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 Section 8; 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

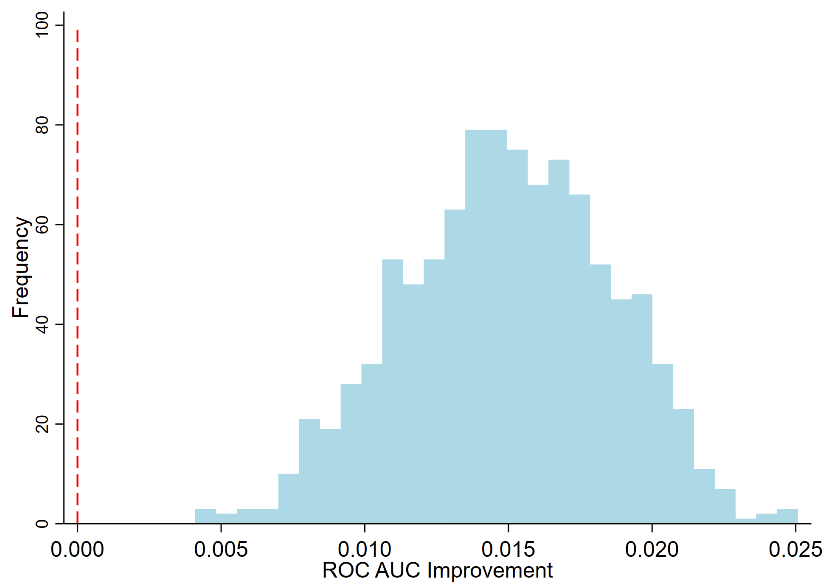

(B) Differences in H-Measure

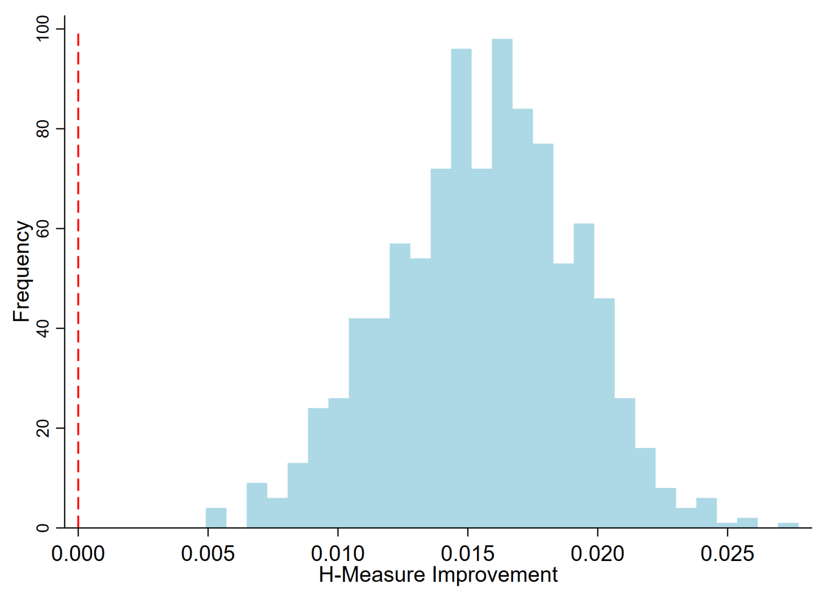

Figure A.12: APR Distribution

This figure shows distribution of APR 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

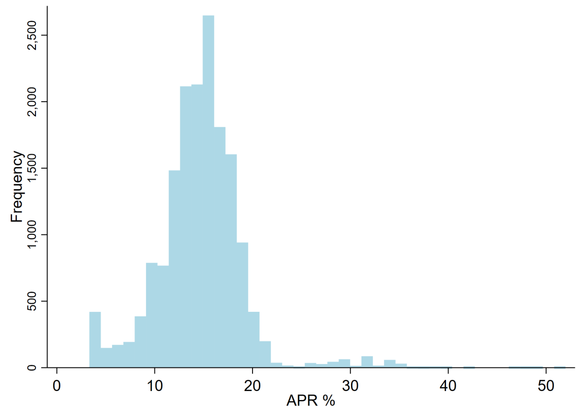

(B) Lender B

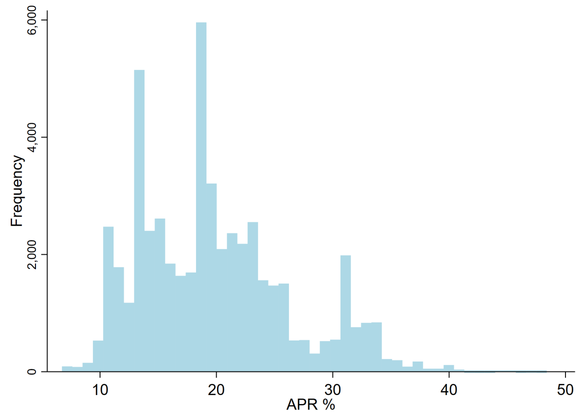

(C) Platform

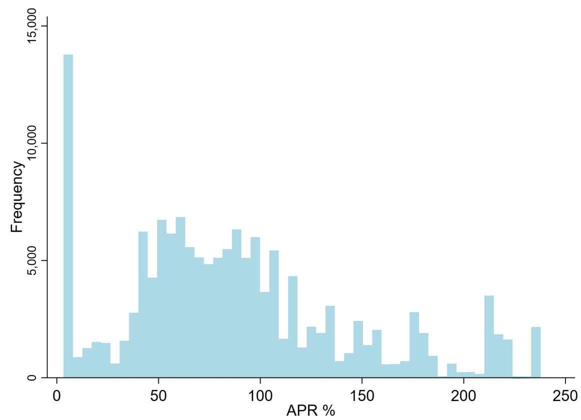

Table A.1: Industry Distribution

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 40).

		Loan Applicants		**Originated Loans**	
Industry	% National	% Applications	% < 40	% Loans	% < 40
Accomodation and Food Services	7.8	10.7	0.5	8.5	0.4
Agriculture, Forestry, Fishing and Hunting	0.4	1.0	0.6	0.7	0.3
Arts, Entertainment, and Recreation	2.2	3.3	0.6	2.2	0.5
Construction	12.3	20.9	0.6	16.2	0.4
Educational Services	1.4	1.9	0.5	1.2	0.4
Finance and Insurance	4.1	2.1	0.5	1.4	0.4
Health Care and Social Assistance	10.5	8.9	0.5	13.9	0.4
Information	1.3	4.0	0.6	1.3	0.4
Manufacturing	3.3	5.0	0.5	5.0	0.3
Mining, Quarrying, and Oil and Gas Extraction	0.3	0.2	0.5	0.1	0.4
Other Services (except Public Administration)	12.1	0.2	0.4	0.8	0.4
Professional, Scientific, and Technical Services	20.0	3.9	0.4	22.5	0.4
Real Estate and Rental and Leasing	5.8	2.9	0.5	2.3	0.3
Retail Trade	10.6	16.1	0.6	13.3	0.4
Transportation and Warehousing	3.2	14.8	0.6	7.2	0.4
Utilities	0.1	0.6	0.6	0.1	0.2
Wholesale Trade	4.5	3.4	0.5	3.5	0.3

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 (<40 years old) with those who are more mature (≥ 40).

	Young Owner (≤ 40)		Mature Owner (> 40)		
	N	Mean	N	Mean	Difference
Loan Variables:					
Approved (%)	34,609	35	89,039	43	-7.994***
Requested Loan Amount (Th$)	34,609	80	89,037	89	-8.107***
APR (%)	8,156	19	27,605	18	0.809***
Originated (%)	34,609	17	89,039	22	-4.962***
Non-Performing Loan (%)	5,009	15	15,181	15	-0.014
Originated Loan Amount (Th$)	5,900	80	19,597	89	-8.372***
Loan Maturity (Years)	5,900	2.71	19,597	2.94	-0.225***
Credit Score & Cash Flow (Bank Statement) Variables:					
FICO	34,609	714	89,039	731	-17.070***
Credits (Th$)	34,609	79	89,039	97	-17.973***
Balance (Th$)	34,609	26	89,039	32	-6.033***
(#) Insuff. Funds	34,609	0.38	89,039	0.30	0.083***
(#) Low or Neg. Bal.	34,609	2.70	89,039	1.95	0.749***
Withdrawals (Th$)	8,168	90	23,219	118	-27.659***
Credits (less new debt) (Th$)	4,062	140	15,217	145	-4.541
$\mathbb{1}$(Daily Pay Loan)	33,960	0.14	86,722	0.14	-0.004*
S.D. Credits (Th$)	34,609	10	89,039	12	-1.536***
S.D. Balance (Th$)	34,609	4.26	89,039	4.94	-0.672***
Borrower Characteristics:					
Owner Age	34,609	34	89,039	53	-19.183***
Female	33,919	0.24	87,996	0.24	0.006**
Business Age (Years)	34,609	5.68	89,039	11	-5.652***
Young Firm (< 5)	34,609	0.56	89,039	0.26	0.304***
Number of Employees	34,609	7.18	89,039	8.70	-1.521***
Pct Black Pop (%)	34,609	14	89,039	13	1.188***
High Pct Black Pop (> 6%)	34,609	0.52	89,039	0.48	0.037***

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 (< 40 years old) with those who are more mature (≥ 40) limited to owners with a FICO score below 740 (N = 5,220).

	All				Young Owner (≤ 40)		Mature Owner (> 40)		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	10,962	699	702	26	2,772	697	8,190	700	-2.24***
Credits (Th$)	10,962	102	61	128	2,772	94	8,190	104	-9.69***
Balance (Th$)	10,962	30	14	50	2,772	29	8,190	30	-1.66
(#) Insuff. Funds	10,962	0.09	0.00	0.34	2,772	0.08	8,190	0.09	-0.01
(#) Low or Neg. Bal.	10,962	0.55	0.00	1.23	2,772	0.58	8,190	0.54	0.04
Withdrawals (Th$)	1,354	115	57	170	342	98	1,012	121	-22.56**
Credits (less new debt) (Th$)	856	150	83	177	180	141	676	153	-12.02
$\mathbb{1}$(Daily Pay Loan)	10,770	0.05	0.00	0.22	2,730	0.04	8,040	0.05	-0.01**
S.D. Credits (Th$)	10,962	13	7.29	13	2,772	12	8,190	13	-0.77***
S.D. Balance (Th$)	10,962	5.00	2.79	5.65	2,772	4.79	8,190	5.08	-0.28**

Table A.4: Summary Statistics on Machine Learning Variables by Cash Flow-Intensive Lender Status

This table reports summary statistics on the Random Forest Model Performance definitions of (Platform) cash flow-intensity (CFI) (N = 879,889). Application Level statistics are limited to applications sent to both cash flow-intensive and not cash flow-intensive lenders. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

	All				CF Intensive Lender		Not CF Intensive Lender		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Lender Level:									
AUC Improvement	76	0.04	0.01	0.15	53	0.08	23	-0.05	0.13***
AUC FICO Model	76	0.68	0.70	0.11	53	0.68	23	0.68	0.00
AUC Cash Flow Model	76	0.70	0.72	0.10	53	0.73	23	0.64	0.08***
AUC P-Value (Difference)	73	0.04	0.00	0.15	53	0.03	20	0.07	-0.03
AUC-PR FICO Model	76	0.36	0.36	0.18	53	0.37	23	0.33	0.03
AUC-PR Cash Flow Model	76	0.38	0.38	0.19	53	0.41	23	0.31	0.10**
AUC-PR P-Value (Difference)	73	0.06	0.00	0.17	53	0.04	20	0.10	-0.06
H-Measure FICO Model	76	0.20	0.16	0.14	53	0.18	23	0.23	-0.05
H-Measure Cash Flow Model	76	0.22	0.22	0.11	53	0.23	23	0.18	0.05*
H-Measure P-Value (Difference)	73	0.04	0.00	0.13	53	0.04	20	0.04	0.00
Application-Forward Level:									
AUC Improvement	563,453	0.04	0.03	0.04	459,667	0.05	103,786	-0.01	0.06***
AUC FICO Model	563,453	0.73	0.73	0.05	459,667	0.72	103,786	0.76	-0.04***
AUC Cash Flow Model	563,453	0.75	0.76	0.04	459,667	0.75	103,786	0.75	0.01***
AUC P-Value (Difference)	563,354	0.00	0.00	0.06	459,667	0.00	103,687	0.02	-0.02***
AUC-PR FICO Model	563,453	0.36	0.41	0.14	459,667	0.38	103,786	0.26	0.12***
AUC-PR Cash Flow Model	563,453	0.39	0.42	0.16	459,667	0.42	103,786	0.25	0.17***
AUC-PR P-Value (Difference)	563,354	0.01	0.00	0.04	459,667	0.01	103,687	0.00	0.00***
H-Measure FICO Model	563,453	0.19	0.17	0.08	459,667	0.18	103,786	0.23	-0.05***
H-Measure Cash Flow Model	563,453	0.22	0.22	0.07	459,667	0.22	103,786	0.21	0.01***
H-Measure P-Value (Difference)	563,354	0.01	0.00	0.04	459,667	0.01	103,687	0.00	0.00***

Table A.5: Summary Statistics by Cash Flow-Intensive Lender Status

This table reports summary statistics by Lender (Platform) cash flow-intensity (CFI) (N = 879,889). This table is limited to applications sent to both cash flow-intensive and not cash flow-intensive lenders. We compare application-forward level data by CFI Lender using a two-sample t-test. Survival is measured by whether the business is open according to Google as of data collection in September, 2024. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

	All				CF Intensive Lender		Not CF Intensive Lender		
	N	Mean	Median	SD	N	Mean	N	Mean	Difference
Loan Variables:									
Approved (%)	563,453	21.0	0.00	41	459,667	23.6	103,786	9.18	14***
# Forwards	563,453	10.8	10.00	4.97	459,667	11.0	103,786	9.65	1.37***
Requested Loan Amount (Th$)	563,453	166	100	203	459,667	168	103,786	158	9.76***
APR (%)	99,280	78	71	49	91,267	80	8,013	54.9	25***
Originated (%)	563,453	3.48	0.00	18	459,667	3.55	103,786	3.19	0.37***
Originated Loan Amount (Th$)	19,630	57.4	34	68	16,324	55.7	3,306	65.7	-10.05***
Loan Maturity (Years)	19,630	1.18	1.00	1.16	16,324	1.04	3,306	1.87	-0.83***
Credit Score & Cash Flow (Bank Statement) Variables:									
FICO	563,453	696	690	57	459,667	696	103,786	697	-1.08***
Credits (Th$)	563,453	113	56	164	459,667	115	103,786	104	11***
Balance (Th$)	563,453	38.5	14	69	459,667	39.0	103,786	36.3	2.69***
(#) Insuff. Funds	563,453	0.38	0.00	1.13	459,667	0.37	103,786	0.41	-0.05***
(#) Low or Neg. Bal.	563,453	0.37	0.00	1.03	459,667	0.36	103,786	0.42	-0.06***
Withdrawals (Th$)	563,453	114	57	166	459,667	116	103,786	105	11***
Credits (less new debt) (Th$)	563,453	95	47	139	459,667	97	103,786	87	9.44***
$\mathbb{1}$(Daily Pay Loan)	563,453	0.35	0.00	0.48	459,667	0.35	103,786	0.35	-0.00
Borrower Age:									
Owner Age	563,453	45.2	44	11	459,667	45.2	103,786	45.3	-0.14***
Young Owner (< 35)	563,453	0.20	0.00	0.40	459,667	0.20	103,786	0.20	0.00***
Young Owner (< 40)	563,453	0.37	0.00	0.48	459,667	0.37	103,786	0.37	0.01***
Young Owner (< 45)	563,453	0.55	1.00	0.50	459,667	0.55	103,786	0.54	0.01***
Young Owner (< 50)	563,453	0.70	1.00	0.46	459,667	0.70	103,786	0.69	0.00***
Other Borrower Characteristics:									
Female	550,832	0.28	0.00	0.45	449,393	0.28	101,439	0.29	-0.01***
Business Age (Years)	563,453	7.87	5.33	7.26	459,667	7.84	103,786	7.98	-0.13***
Young Firm (< 5)	563,453	0.45	0.00	0.50	459,667	0.45	103,786	0.43	0.01***
Number of Employees	563,453	7.82	4.00	11	459,667	7.92	103,786	7.38	0.55***
Business Survival	32,048	0.42	0.00	0.49	28,808	0.42	3,240	0.42	-0.01
Pct Black Pop (%)	563,453	12.9	6.00	17	459,667	12.9	103,786	13.1	-0.25***
High Pct Black Pop (> 6%)	563,453	0.50	0.00	0.50	459,667	0.50	103,786	0.50	-0.01***

Table A.6: Effect of Assignment to Cash Flow-Intensive Lender on Approval and APR (Within-Application), All Applicant Ages

This table uses data from the Platform to test the effect of assignment to a cash flow-intensive lender on approval and APR for 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. Panel A shows if cash flow-intensive lenders are more likely to approve young entrepreneurs and Panel B shows the impact on APR. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood or APR for young applicants forwarded to cash flow-intensive lenders relative to older 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.

Panel A: Effect on Loan Approval

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	1.31***	1.40***	1.65***	2.08***
	(0.25)	(0.22)	(0.22)	(0.25)
Observations	829,167	829,167	829,167	829,167
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.315	0.315	0.315	0.315
Y-mean	19.83	19.83	19.83	19.83

Panel B: Effect on the Interest Rate

Dependent Variable:	APR (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	-0.94	-1.55***	-1.70***	-1.44***
	(0.80)	(0.54)	(0.46)	(0.46)
Observations	105,129	105,129	105,129	105,129
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.878	0.878	0.878	0.878
Y-mean	81.78	81.78	81.78	81.78

Table A.7: Effect of Assignment to Cash Flow-Intensive Lender on Approval and APR, Without Application Fixed Effects

This table uses data from the Platform to test the effect of assignment to a cash flow-intensive lender on approval and APR for 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. This is the same specification as Table 5 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. Panel A shows if cash flow-intensive lenders are more likely to approve young entrepreneurs and Panel B shows the impact on APR. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood or APR 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: Effect on Loan Approval

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	1.94***	1.85***	1.68***	1.60***
	(0.26)	(0.23)	(0.22)	(0.21)
Observations	444,774	600,472	754,764	880,614
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.091	0.091	0.091	0.091
Y-mean	19.38	19.44	19.55	19.66

Panel B: Effect on the Interest Rate

Dependent Variable:	APR (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	2.34***	1.40***	1.20***	1.22***
	(0.56)	(0.43)	(0.37)	(0.34)
Observations	72,926	98,810	124,941	146,545
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.800	0.797	0.797	0.797
Y-mean	87.86	88.92	88.73	88.22

Table A.8: Effect of Assignment to Cash Flow-Intensive Lender on Approval and APR (Within-Application), Limited to Applications Sent to Both Types of Lender

This table uses data from the Platform to test the effect of assignment to a cash flow-intensive lender on approval and APR for 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. This is the same specification as Table 5 but limited to applications sent to both cash flow-intensive and not cash flow-intensive lenders (N = 571,628). 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 shows if cash flow-intensive lenders are more likely to approve young entrepreneurs and Panel B shows the impact on APR. The interaction term "Young=1 × CFI=1" represents the difference in approval likelihood or APR 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: Effect on Loan Approval

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	2.56***	2.38***	2.21***	2.09***
	(0.31)	(0.27)	(0.25)	(0.25)
Observations	283,792	380,107	479,642	563,453
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.292	0.294	0.295	0.294
Y-mean	20.73	20.74	20.86	20.97

Panel B: Effect on the Interest Rate

Dependent Variable:	APR (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	-1.82**	-2.03***	-1.83***	-1.44***
	(0.84)	(0.59)	(0.50)	(0.46)
Observations	40,228	53,909	68,579	81,207
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.876	0.876	0.876	0.876
Y-mean	73.95	74.60	74.51	74.23

Table A.9: Effect of Assignment to Cash Flow-Intensive Lender on Approval and APR by FICO Range, Supplementary Analysis (Within-Application)

This table uses data from the Platform to test the effect of assignment to a cash flow-intensive lender on approval (Panel A) and APR (Panel B) for young entrepreneurs. We split the sample into two FICO score groups, one below the threshold for "Very Good" of 739, and one above. Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Approval with Full Sample Split Across Age Groups

Dependent Variable:	Approved (%)							
Young Owner (Young) Defined As:	< 35		< 40		< 45		< 50	
FICO Range:	< 740	≥ 740	< 740	≥ 740	< 740	≥ 740	< 740	≥ 740
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1	3.14***	2.23***	3.02***	1.46**	2.62***	2.03***	2.43***	2.05***
	(0.33)	(0.74)	(0.30)	(0.62)	(0.28)	(0.57)	(0.27)	(0.54)
Observations	337,903	80,549	460,142	104,792	579,505	130,852	675,804	153,363
Application FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.311	0.335	0.313	0.337	0.312	0.340	0.312	0.339
Y-mean	18.13	25.53	18.23	25.68	18.31	25.99	18.39	26.18

Panel B: APR with Full Sample Split Across Age Groups

Dependent Variable:	APR (%)							
Young Owner (Young) Defined As:	< 35		< 40		< 45		< 50	
FICO Range:	< 740	≥ 740	< 740	≥ 740	< 740	≥ 740	< 740	≥ 740
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Young=1 × CFI=1	-1.85*	-1.82	-2.70***	-0.99	-2.75***	-0.44	-2.08***	-0.54
	(1.07)	(1.31)	(0.77)	(0.88)	(0.66)	(0.72)	(0.62)	(0.66)
Observations	37,981	13,945	52,166	18,259	66,192	23,083	77,762	27,358
Application FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.870	0.871	0.870	0.874	0.870	0.876	0.869	0.875
Y-mean	89.22	59.23	90.24	59.69	90.17	59.37	89.75	59.16

Table A.10: Effect of Assignment to Cash Flow-Intensive Lender on Approval for Low-Risk Lenders (Within-Application)

This table uses data from the Platform to test the effect of assignment to a cash flow-intensive lender on approval. In Panel A, we restrict the sample to lenders who approve less than 10% of loan applications. In Panel B, we restrict the sample to lenders for whom the 25th percentile FICO score for approved applications is above 669 (this score is close to the average 25th percentile in the overall sample, and is also the industry threshold separating Poor/Fair from Good). Standard errors are clustered by applicant. ***, **, * correspond to statistical significance at the 1, 5, and 10 percent levels, respectively.

Panel A: Within Lenders with Low Approval Rates

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	5.13***	4.63***	4.09***	3.51***
	(0.32)	(0.29)	(0.27)	(0.27)
Observations	92,853	124,266	156,169	182,651
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.391	0.393	0.390	0.386
Y-mean	4.45	4.42	4.41	4.41

Panel B: Within Lenders with High Average FICO Scores Among Approved Appplications

Dependent Variable:	Approved (%)			
	Young (< 35)	Young (< 40)	Young (< 45)	Young (< 50)
	(1)	(2)	(3)	(4)
Young=1 × CFI=1	3.23***	2.50***	2.15***	2.04***
	(0.66)	(0.56)	(0.51)	(0.49)
Observations	132,874	177,196	223,164	261,941
Application FE	Yes	Yes	Yes	Yes
Lender FE	Yes	Yes	Yes	Yes
R-squared	0.422	0.426	0.427	0.426
Y-mean	20.49	20.47	20.66	20.79

Table A.11: 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
Other Borrower Characteristics:									
Female	11,408	0.23	0.00	0.42	5,710	0.22	5,698	0.23	-0.01
Business Age (Years)	11,535	11.4	9.29	7.77	5,776	11.3	5,759	11.4	-0.12
Young Firm (< 5)	11,535	0.18	0.00	0.39	5,776	0.19	5,759	0.18	0.01
Number of Employees	11,535	9.84	5.00	13	5,776	9.78	5,759	9.89	-0.11
Pct Black Pop (%)	11,535	11.0	5.00	16	5,776	11.0	5,759	11.0	-0.04
High Pct Black Pop (> 6%)	11,535	0.44	0.00	0.50	5,776	0.44	5,759	0.44	-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
Other Borrower Characteristics:									
Female	15	0.23	0.23	0.03	5	0.23	10	0.22	0.01
Business Age (Years)	15	11.0	11	1.00	5	11.1	10	11.0	0.06
Young Firm (< 5)	15	0.20	0.20	0.05	5	0.20	10	0.19	0.00
Number of Employees	15	10.8	11	2.08	5	10.3	10	11.1	-0.75
Pct Black Pop (%)	15	11.3	11	2.45	5	11.2	10	11.3	-0.04
High Pct Black Pop (> 6%)	15	0.45	0.45	0.06	5	0.46	10	0.45	0.01

Table A.12: 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 (1)	Low (2)	High (3)	Low (4)	High (5)	Low (6)	High (7)	Low (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

Table A.13: 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.17. 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 1,000 bootstrap iterations. Definitions of the performance metrics *ROC AUC*, *AUC-PR* and *H-measure* are provided in Section 8; 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	0.617	0.240	0.051
Cash-Flow Model	0.634	0.251	0.065
Difference	0.017***	0.011***	0.014***
Preferred Specification (Borrower Default):			
FICO Model	0.642	0.255	0.071
Cash-Flow Model	0.657	0.272	0.086
Difference	0.015***	0.017***	0.016***
Full Specification with APR (Borrower Default):			
FICO Model	0.691	0.322	0.127
Cash-Flow Model	0.700	0.321	0.135
Difference	0.009***	-0.001***	0.008***

Table A.14: 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 6). 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 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 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 (< 740):	1,522,000	0.686	761,000	0.532	761,000	0.840	1.578***
High FICO (≥ 740)	1,522,000	0.314	761,000	0.468	761,000	0.160	0.342***
Full Sample:							
Young Firm (< 5)	1,522,000	0.291	761,000	0.203	761,000	0.380	1.873***
High Pct Black Pop (> 6%)	1,522,000	0.464	761,000	0.441	761,000	0.486	1.103**
High Pct Black Pop (> 17%)	1,522,000	0.211	761,000	0.206	761,000	0.216	1.049
Female	1,514,978	0.218	756,157	0.256	758,821	0.180	0.704***
Low FICO (< 740):							
Young Firm (< 5)	1,522,000	0.214	761,000	0.107	761,000	0.321	2.989***
High Pct Black Pop (> 6%)	1,522,000	0.323	761,000	0.233	761,000	0.414	1.778***
High Pct Black Pop (> 17%)	1,522,000	0.146	761,000	0.107	761,000	0.184	1.713***
Female	1,514,978	0.150	756,157	0.142	758,821	0.158	1.108***
High FICO (≥ 740):							
Young Firm (< 5)	1,522,000	0.077	761,000	0.095	761,000	0.059	0.616
High Pct Black Pop (> 6%)	1,522,000	0.141	761,000	0.208	761,000	0.073	0.349***
High Pct Black Pop (> 17%)	1,522,000	0.065	761,000	0.099	761,000	0.032	0.324***
Female	1,514,978	0.068	756,157	0.114	758,821	0.023	0.199***

Table A.15: 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 1,000 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)	0.137	709,342	0.131	346,996	0.110	0.123	362,346	0.150	0.177**
Young Owner (< 40)	0.342	709,342	0.265	346,996	0.226	0.291*	362,346	0.303	0.435***
Young Owner (< 45)	0.696	709,342	0.417	346,996	0.377	0.605*	362,346	0.456	0.837**
Young Owner (< 50)	1.291	709,342	0.567	346,996	0.535	1.152	362,346	0.598	1.487*
Low FICO (< 740):									
Young Owner (< 35)	0.036	709,342	0.098	346,996	0.070	0.076**	362,346	0.124	0.142***
Young Owner (< 40)	0.079	709,342	0.179	346,996	0.116	0.131***	362,346	0.240	0.317***
Young Owner (< 45)	0.130	709,342	0.282	346,996	0.191	0.237***	362,346	0.368	0.582***
Young Owner (< 50)	0.184	709,342	0.382	346,996	0.273	0.376***	362,346	0.487	0.948***
High FICO (≥ 740):									
Young Owner (< 35)	0.093	709,342	0.033	346,996	0.039	0.041	362,346	0.026	0.027**
Young Owner (< 40)	0.221	709,342	0.086	346,996	0.110	0.123	362,346	0.063	0.067***
Young Owner (< 45)	0.420	709,342	0.136	346,996	0.186	0.228	362,346	0.088	0.096***
Young Owner (< 50)	0.691	709,342	0.185	346,996	0.262	0.356	362,346	0.111	0.125***

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 (< 740):	0.440	1,522,000	0.686	761,000	0.532	1.138***	761,000	0.840	5.246***
High FICO (≥ 740)	2.271	1,522,000	0.314	761,000	0.468	0.879***	761,000	0.160	0.191***
Full Sample:									
High Pct Black Pop (> 6%)	0.854	1,522,000	0.464	761,000	0.441	0.789	761,000	0.486	0.947
High Pct Black Pop (> 17%)	0.255	1,522,000	0.211	761,000	0.206	0.259	761,000	0.216	0.275
Young Firm (< 5)	0.362	1,522,000	0.291	761,000	0.203	0.254***	761,000	0.380	0.612***
Female	0.298	1,514,978	0.218	756,157	0.256	0.344***	758,821	0.180	0.220***
Low FICO (< 740):									
High Pct Black Pop (> 6%)	0.173	1,522,000	0.323	761,000	0.233	0.303***	761,000	0.414	0.705***
High Pct Black Pop (> 17%)	0.073	1,522,000	0.146	761,000	0.107	0.120***	761,000	0.184	0.225***
Young Firm (< 5)	0.097	1,522,000	0.214	761,000	0.107	0.120***	761,000	0.321	0.473***
Female	0.080	1,514,978	0.150	756,157	0.142	0.166	758,821	0.158	0.187***
High FICO (≥ 740):									
High Pct Black Pop (> 6%)	0.456	1,522,000	0.141	761,000	0.208	0.263	761,000	0.073	0.079***
High Pct Black Pop (> 17%)	0.156	1,522,000	0.065	761,000	0.099	0.109***	761,000	0.032	0.033***
Young Firm (< 5)	0.215	1,522,000	0.077	761,000	0.095	0.105	761,000	0.059	0.062*
Female	0.184	1,514,978	0.068	756,157	0.114	0.128***	758,821	0.023	0.023***

Table A.16: Applying TACO to Approval Decisions: Triple Interaction

This table examines the effects of the predicted default probabilities generated by the Cash Flow (CF) and Baseline models on applicant approval. The CF and Baseline models are trained using data from Lender A and Lender B, following the methodology in Section 6.3. The trained models are used to make out-of-sample default probability predictions on all applicants to the Platform forwarded to at least one lender. The difference $h_i = g(X_i) - f(X_i)$ is then used in an OLS regression to predict applicant approval, where g and f represent the CF-based underwriting model and the Baseline model, respectively. Here, "Most Hurt by CF" ("Most Helped by CF") indicates the top (bottom) tail of h_i by percentile within bootstrap iteration. Interaction rows report the difference for CFI lenders using the CFI classification from Section 5, as well as triple interactions for "Young Owners", which is defined here as owners under 40 years of age. Columns 1–3 include lender and application fixed effects. Column 4 reports a linear specification in h_i without application fixed effects but with controls equivalent to the Baseline model. Standard errors are based on the bootstrap distribution across 100 iterations; ***, **, * denote p-values below 0.01, 0.05, 0.10, respectively.

Dependent Variable:	Approved (%)			
	10% Tails (1)	20% Tails (2)	30% Tails (3)	Continuous h_i (4)
Most Hurt by CF × (CFI = 1)	-1.03	-0.83	-0.83	
	(1.22)	(0.70)	(0.51)	
Most Helped by CF × (CFI = 1)	1.52***	1.20***	0.90***	
	(0.59)	(0.49)	(0.39)	
Most Hurt by CF × (CFI = 1) × Young Owner (< 40)	-0.39	-0.18	0.22	
	(0.63)	(0.42)	(0.29)	
Most Helped by CF × (CFI = 1) × Young Owner (< 40)	0.57	0.63	0.58*	
	(0.48)	(0.39)	(0.30)	
Continuous h_i				-25.93***
				(4.34)
Continuous h_i × (CFI = 1)				-39.36***
				(4.77)
Continuous h_i × (CFI = 1) × Young Owner (< 40)				0.10
				(1.10)
Observations	880,462	880,462	880,462	983,920
Bootstrap Iterations	100	100	100	100
Application FE	Yes	Yes	Yes	—
Lender FE	Yes	Yes	Yes	Yes
Baseline Controls	—	—	—	Yes
Y-mean	21.62	21.62	21.62	22.36

Table A.17: Approval and Default Machine Learning Features

This table lists the features included in the ML analysis (Table A.13 and 11). 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

Appendix 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 primary applicant (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.
- Number of Employees: The number of employees in a firm.
- Business Survival: Survival is measured by whether the business is open according to Google as of data collection in September, 2024.
- 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.

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.
- Number of Employees: The number of employees in a firm.
- Business Survival: Survival is measured by whether the business is open according to Google as of data collection in September, 2024.
- 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.

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.
- Number of Employees: The number of employees in a firm.
- 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.

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 90%. This leaves us with a sample of 76 lenders.

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 $\times$ Credits: An interaction term between variables *Balance (Th$)* and *Credits (Th$)*.
- (#) Insuff. Funds $\times$ (#) 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$)*.
- Never Low or Neg. Balance: An indicator variable that takes the value of 1 if *(#) Low or Neg. Balance* is equal to zero, and zero otherwise.
- Never Insuff. Funds: An indicator variable that takes the value of 1 if *(#) Insuff. Funds* is equal to zero, and zero otherwise.
- (#) Insuff. Funds > 5: An indicator variable that takes the value of 1 if *(#) Insuff. Funds* is greater than 5, and zero otherwise.
- Balance Volatility Ratio: The ratio of the log of variables *S.D. Balance (Th$)* 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.

Appendix C. Machine Learning Methods

A random forest classifier is a ML 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.[45] 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.

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.[46] 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

[45] 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.

[46] 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.

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 sampled with replacement and then used to train a model for each bootstrap iteration.[47] This mirrors the common bootstrapping approach used by, e.g., Fuster et al. (2022). Each model's performance 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.

Performance Metrics We evaluate the performance of the models using three measures, which are all reported in Table 9. The first, known as the receiver-operating characteristic (ROC) area under the curve (AUC), is widely used to evaluate binary classification models.[48] The ROC curve

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

[48] 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.

describes the true positive rate (TPR) and the false positive rate (FPR) at various thresholds.[49] 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.

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. Next, we compute the H-Measure for each model, an alternative to AUC proposed by Hand (2009).[50] 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 1,000 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.11.

Lender Cash Flow Intensity Measure To construct a measure of cash-flow intensity (mentioned in Section 5.2) for each lender in our Platform data we first restrict the sample to those with at least 50 application forwards and approval rates below 95 percent. For each lender, we then perform 100 stratified random splits of its application-level data (80 percent training, 20 percent testing). On each split we run the Baseline and Cash Flow model in a Random Forest classifiers using identical tuning parameters (Estimators: 100, Maximum Tree Depth: 10, Minimum Samples to Split: 15, Minimum Child Weight: 20). We create the average ROC AUC for all lenders for whom the model was able to run and calculate the percent improvement from the baseline to the cash flow model. We then define a lender as "Cash Flow Intensive" if their improvement is positive. Further details about lender level performance metrics are contained in Table A.4.

[49] 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.

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