


Finance

The Past, Present, and Future of Credit Scores

Credit scores are changing. Where did the three-digit number come from, and what does its next chapter mean for consumer lending?



Emmanuel Awosika ■ 24 June 2026 ■ 15 min read

Credit before credit scores

Credit has a long and illustrious history, dating back thousands of years. The ancient Babylonians recorded loans of silver and grain on clay tablets, with maximum interest rates set by the Code of Hammurabi (<https://history.hanover.edu/courses/excerpts/165hammurabi.html>): 33 percent per year on grain loans and 20 percent for silver.¹ In Greece, maritime loans were used to finance trading voyages, where borrowers paid 20 to 30 percent in interest but could have the loan forgiven if the ship sank.² The Roman Empire relied on credit to fuel commercial activity: bankers accepted collateral deposits and provided loans at a maximum interest rate of 12 percent set by the Twelve Tables, though the cap was often circumvented in practice.

Today, the United States is the biggest consumer credit market in the world. No other country, except China, comes close to the US in the number of people using credit to finance purchases, or the amount those people are spending. Americans now owe more than \$5 trillion (<https://www.federalreserve.gov/releases/g19/current/>) in consumer credit, a figure that excludes mortgages, with auto loans, student

loans, and personal loans among the biggest categories. The system is not perfect, but it has helped drive economic growth, increase purchasing power, accelerate mobility, and provide financial security for many.

Although credit has evolved since the Babylonians, the core parts have not changed much. A borrower requests a loan. The lender assesses their ability and willingness to repay. The lender accepts or rejects the application. If accepted, the borrower is presented with a set of terms: repayment period, interest rate, collateral, fees, penalties. If the borrower accepts, the lender provides the principal, and the borrower repays the principal plus interest over the agreed period until the debt is settled and the transaction concluded.

The second step, assessing repayment capacity, is arguably the most important in the whole process. It determines whether the borrower can borrow at all and on what terms. If a lender believes a borrower is willing and able to repay, it accepts the application; if repayment capacity is suspect, it may refuse, or set more aggressive terms: higher collateral requirements, shorter repayment periods, higher interest rates, lower loan amounts, all to account for the risk.

The second step is also where things have been difficult for a very long time. The question is simple: how do you actually determine whether someone is willing and able to repay a loan? Ask the average person and the answer will be credit scores, and it is a good answer. Credit scores are used by lenders in the US to estimate the probability of a borrower defaulting, and they have been the dominant measure of creditworthiness for decades.

For most Americans, the credit score sits alongside the Social Security Number as one of the most consequential numbers in their lives. It determines whether you can get a loan and whether the terms will be favorable. A low score predicts reduced access to credit; a high score predicts greater access. The system has been criticized for perceived unfairness and its effect on inequality, but lenders argue that the score remains the most useful concrete measure of a person's financial history and their relationship with money and debt.

The origin of credit scores

Credit scores first appeared in the late 1950s, though the system we know today took years to develop. The first major event in that timeline was the founding of FICO, Fair, Isaac and Company (<https://www.myfico.com/credit-education/blog/history-of-the-fico-score>), by William Fair and Earl Isaac in 1956.³ FICO was the first company to produce a credit score that lenders could use off the shelf, and it became synonymous with credit scores themselves. Today FICO holds roughly a 90 percent share of the consumer lending category⁴, and despite decades of efforts to dethrone it, it remains the clear leader.

Fair, an engineer with degrees from Caltech, Stanford, and UC Berkeley, and Isaac, a mathematician with a master's degree from UCLA, were colleagues at the Stanford Research Institute (<https://www.sri.com/about/>) in Menlo Park, California when FICO was founded. Both men had a background in operations research (https://en.wikipedia.org/wiki/Operations_research), the discipline of using statistical modeling to solve complex optimization problems.⁵ In the Second World War, the Allies had enlisted

mathematicians and scientists to solve logistical problems that human judgment was unequipped to grasp: the best way to route convoys to minimize submarine losses, the optimal allocation of bombing runs, the positioning of radar stations for maximum coverage.

The core insight underpinning operations research was simple: many optimization problems could be modeled mathematically and solved reliably with data instead of relying on human judgment and experience. Like other innovations birthed by the military, it found civilian applications after the war. Retailers used it to optimize supply chains, insurers used it to price risk, oil companies used it to plan drilling. Decisions that once required experienced human judgment were decomposed into clusters of variables and made more reliably by algorithms that replaced intuition with probabilistic reasoning.

Take a store owner making a decision about inventory. They have to decide what to buy, when, and how much. The early approach was to rely on a store buyer who decided based on experience and feel, but that approach had predictable limits. The operations research approach treated stocking as an optimization problem composed of variables: seasonality, since demand for some products follows seasonal cycles; supplier lead times, since how long delivery takes determines the best window to reorder; and sales velocity, since the rate at which customers buy determines the ideal amount of stock to hold. Feed in historical sales, lead times, and seasonal patterns, set objectives like maximizing revenue and minimizing carrying costs, and the algorithm produces a specific reorder quantity and reorder point for each product, with a speed and reliability no human could match.

Fair and Isaac judged that lending could be improved by the same techniques. Lending decisions were optimization problems: a lender had to figure out how to lend money without becoming insolvent while still making enough profit to make the investment rational. Satisfying that constraint required many things, none more important than choosing who to lend to.

The ideal borrower was someone who repaid on time and did not default. But it was hard to determine who fit that description. The loan officers tasked with evaluating applications relied on character judgments and intuition, which sounded reasonable in theory but created problems in practice, because the assessment of borrowers was highly subjective and opaque. Officers denied worthy applicants for reasons that had little to do with repayment capacity, and approved loans to people who seemed creditworthy on the surface but ended up defaulting.

From credit men to credit bureaus

The Manchester Guardian Society

Credit bureaus emerged directly out of the need to make lending less discretionary and more objective. They were preceded by credit men who traveled from town to town, comparing notes on debtors and publishing information about individuals who refused to settle debts. Some towns posted public bulletins listing newly identified defaulters to encourage prompt repayment and warn unsuspecting lenders. But the system was limited: a borrower could simply move to a new town the credit men had not reached, and many did.

Credit bureaus were more organized. They grew out of early efforts by local merchants, tired of being swindled and racking up bad debt, to build a more reliable way to assess creditworthiness. The first formal institution of this kind was the Society of Guardians for the Protection of Tradesmen against Swindlers, Sharpers and other Fraudulent Persons, later known as the Manchester Guardian Society, formed by a group of Manchester merchants in 1826. Members reported borrowers who defaulted, and the information was compiled into a circular published every month.

This approach was superior to the credit man in three ways. The information was collected centrally, rather than living in one man's head and notebook and disappearing when he died or moved. It propagated faster, since a monthly circular traveled far greater distances than a single man visiting one town at a time. And it was arguably more reliable: the guardians initially accepted member reports, but hired a dedicated accuracy officer in 1857 to verify claims and keep the data trustworthy.

The Manchester Guardian Society is notable because it grew into one of the biggest collectors of credit information in the UK, and later became part of Experian (<https://www.experianplc.com/about-us/our-history>), one of the largest credit reporting agencies in the world. The practice crossed the Atlantic. In New York in 1864, the Mercantile Agency (https://en.wikipedia.org/wiki/Dun_%26_Bradstreet), a commercial credit reporting firm, created the first formal ranking of creditworthy commercial borrowers. The Retail Credit Company (<https://en.wikipedia.org/wiki/Equifax>)⁶ began collecting consumer loan data in 1899 as revolving credit and bank cards started to proliferate, and by 1929 had offices across the US and Canada, with clerks working continuously to add millions of people to individual files.

What the bureaus got wrong

Although bureaus were an improvement over credit men, they still had limitations. Agents collected useful financial information such as total debt, new originations, repayment history, and defaults, which was easy because lenders were expected to report it. But they also collected non-financial information: marital status, address changes, community reputation, employment history, political affiliations, and more. A file might mix a few factors that genuinely predicted default with many that had little or no bearing on it.

This was both problematic and ineffective. Ineffective because much of the information could not tell a lender whether a borrower was creditworthy, and problematic because it invaded privacy by collecting sensitive personal data. Field agents often interpreted objective facts about borrowers subjectively, making assessments of future behavior that reflected their own biases and preconceptions.

It also did not help that bureaus were mostly local. A bureau in Kansas held information only on borrowers in Kansas; if a Kansas bank wanted to lend to someone in New York, it had to coordinate with a New York bureau. If a borrower moved from New York to Kansas, they might be unable to access loans locally because the Kansas bureaus had not collected enough on them. The fragmentation was worse still by segment: a bureau sourcing data from retailers offering revolving credit was useful only to a lender serving that same population, not to one serving, say, bank cardholders. This is why early credit scoring took the form of a bureau collecting data on a single lender's customers and building a bespoke model for that lender alone.

Lessons from the Atlantic

The convoy problem


Both Fair and Isaac had served in the war, and both had been influenced by Allied efforts, particularly the British military's, to apply mathematical modeling to complex decisions under uncertainty. One example was Germany's effort to force a surrender without invading by targeting the convoys carrying food, ammunition, fuel, and equipment across the Atlantic from the US to the UK. After various defensive tactics failed repeatedly, the Admiralty brought in a team of scientists led by the physicist Patrick Blackett (<https://www.nobelprize.org/prizes/physics/1948/blackett/facts/>)² to analyze the problem from first principles.

Blackett's team studied convoy sizes, attack patterns, and escort strategy, and produced recommendations that contradicted the common wisdom of the military but dramatically improved outcomes once implemented.

Bigger convoys, fewer losses

The Admiralty had long believed large convoys were risky and capped them at 40 vessels. Blackett's team showed the opposite was true, and recommended increasing convoy sizes, which lowered loss rates overall. The belief rested on two intuitive but wrong assumptions.

The first was that larger convoys are easier to attack. An 80-vessel convoy covers more surface area than a 40-vessel one, which officers assumed gave the U-boats more targets. But the Germans could only attack with as many U-boats as they had already deployed in the area; a commander could not conjure more submarines because a convoy was larger. If a wolf pack of ten U-boats carried ten torpedoes, it could sink at most ten vessels whether the convoy held 40 or 80. Ten losses out of 40 is a 25 percent loss rate; ten out of 80 is 12.5 percent. The same attack produced a proportionally lower loss rate. The Admiralty had been counting ships sunk rather than the share of ships that survived.



Blackett's point in one picture. The two piles of losses are identical; only the denominator changes.

The second assumption was that larger convoys are harder to defend because they are more spread out. Blackett disproved this with basic geometry. The perimeter of a shape grows in proportion to its width, but the area grows with the square of the width, so the interior grows much faster than the exterior. Double the radius of a circle and its perimeter doubles while its area quadruples. Ships sailed in a compact cluster, not a straight line, and only the ships on the outer edge were exposed; those in the interior were shielded. Adding ships to a compact formation mostly added protected interior, not exposed edge.

The commanders had been thinking about a convoy the way you think about a line of bottles at a shooting range: more bottles, more targets. Blackett's team showed that a convoy was more like a pack of people huddled together with their backs to each other, where those on the inside are protected by those on the outside, and adding people increases the number who can be protected rather than the number who can be hit. Applying the insight, the British moved more supplies across the Atlantic while losing fewer ships.

Building the FICO score

Fair and Isaac set out to apply the same modeling the British used to optimize convoys to consumer lending. At Stanford Research Institute both men had worked on operations research for the US military, and in each case the work revolved around one question: given a set of variables and a

defined objective, what decision maximizes the probability of the desired outcome? The military asked it when making life-and-death decisions. Fair and Isaac believed the same framework could serve many areas of civilian life, and consumer credit sat at the top of their list.

Consumer credit was only beginning to take root in the United States, but the signs were clear that lending would become a major part of the modern economy, and that its failure would carry a significant cost. Despite the scale at which credit was growing, the system still had no reliable way to answer who should be allowed to borrow and on what terms. The bureaus collected data, but lenders found it hard to use, because it was statistically noisy, poorly organized, and limited in its applicability, a mix of financial facts reported by lenders and contextual information collected by field agents.

Multivariate analysis

Multivariate analysis, which the pair developed and first wrote about in a research paper in the mid-1950s, sought to build a better picture of a borrower's creditworthiness by modeling a loan as a decision problem that could be solved reliably: analyze the variables, evaluate the potential outcomes, account for the objective, and produce the decision that maximizes the probability of the desired result. The lender's objective was simple: make a profit by lending to people willing and able to repay principal plus interest, and avoid a loss. Whether that outcome occurred depended on several variables that determined whether a borrower would repay.

Unlike the early bureaus, FICO refined its data collection to focus on the data points that actually predicted repayment. Because it relied on more predictive factors, such as total debt on record and repayment history, and dropped extraneous data with little predictive power, such as marital status and employment history, the score solved two problems for lenders at once: it considered more variables than any human officer could weigh, and it weighted each according to its predictive power, distilling creditworthiness into a single metric.

The first customers

Their first product was a custom credit decisioning engine tailored to each lender's population. To win customers, they sent letters to 50 of the biggest banks and financial institutions in the US. Only one, the Kansas National Bank, through its Public Finance Corporation subsidiary, responded and agreed to a pilot. Later FICO would be contracted by Carte Blanche and a number of other institutions, all heavily involved in issuing credit and all wanting a more consistent, less troublesome way to assess borrowers and underwrite loans.

To start, Fair and Isaac reviewed the files of delinquent borrowers and traced them back to the original loan origination to identify which variables had predicted default. Doing this across thousands of files let them find statistical patterns and stress-test the model. If many delinquent borrowers in a program shared a quality such as high debt loads or weak employment prospects, that quality became a variable the algorithm analyzed and a data point that factored into future decisions.

From custom models to a universal score

Around this time the credit industry was under increasing scrutiny for its perceived role in reinforcing inequality through opaque and subjective decisions that excluded segments of the population from access to credit. The Fair Credit Reporting Act (<https://www.consumerfinance.gov/compliance/compliance-resources/other-applicable-requirements/fair-credit-reporting-act/>) of 1970⁸ limited what data bureaus could collect and imposed penalties on lenders found to use personal data such as race, gender, and marital status to approve or deny loans. The old approach of looking someone in the eye to judge their character, and lending on a handshake, was no longer defensible. By providing an objective statistical basis for approving or rejecting an application, Fair and Isaac helped solve a thorny problem for lenders.

Today a FICO score is used by lenders of every size, but the general-purpose FICO score did not arrive until 1989², three decades after the company was founded. In the years before, FICO built a separate model trained on each lender's own data. It took that long to reach a universal score for several reasons.

Why standardization took decades

1. **Data fragmentation.** The US had thousands of credit bureaus operating across all 50 states with little coordination. For a general-purpose score to be predictive, the model needed to be trained on a broad cross-section of the borrower population, and no single bureau had access to that much data.
2. **Poor standardization.** Bureaus did not all store the same information, so their records were hard to reconcile. A scoring engine has to be trained on data with a consistent structure to be applicable across contexts, and with no major lender or consortium enforcing it, the industry could not standardize around a single reported dataset.
3. **Low incentives.** Before the Fair Credit Reporting Act, lenders faced few consequences for opacity. The Act not only restricted the data they could collect but gave borrowers the right to ask why a request was declined. Because a decision could be denied for many reasons, explaining it clearly was hard and invited lawsuits. A credit score offered an easy alternative: if a lender denied an application, it could simply cite the score.

That, in the end, was the score's quiet triumph. It did not just predict repayment more reliably than a loan officer. It gave every lender, regulator, and securitization market a single number they could all agree to treat the same way, and it is that shared agreement, more than the underlying mathematics, that has kept the score at the center of American lending for more than half a century. ■

■ Footnotes

1. The Code of Hammurabi, from Babylon around 1750 BC, is among the oldest written legal codes and set some of the first legal ceilings on interest, which merchants routinely tried to evade. ←

2. A loan repaid only if the ship arrives safely is called a bottomry loan. Because the lender carries the voyage's risk, it doubles as one of the earliest forms of marine insurance. ↩
3. Fair and Isaac met at the Stanford Research Institute and started the company in San Rafael, California, reportedly with a few hundred dollars each. The name now covers both the firm and its score. ↩
4. FICO holds no data of its own. It scores files kept by the three main bureaus, Equifax, Experian, and TransUnion, which in 2006 launched a joint rival called [VantageScore](#) to loosen FICO's grip. ↩
5. The field was named operational research in Britain in the late 1930s, when scientists were first asked to study how to use radar and other new systems in combat. The American term became operations research. ↩
6. The Retail Credit Company, founded in Atlanta in 1899, built enormous files on ordinary Americans. After public alarm over its practices helped prompt new privacy law, it renamed itself Equifax in 1975. ↩
7. Patrick Blackett won the 1948 Nobel Prize in Physics. The teams that ran this wartime analysis were sometimes called Blackett's circus, and their methods became the discipline Fair and Isaac later borrowed. ↩
8. The Act gave people the right to see their own credit files and dispute mistakes in them, and it barred lenders from basing decisions on factors such as race and gender. ↩
9. The first general purpose FICO score, released in 1989, was built with the bureau Equifax. It set the familiar scale that runs from 300 to 850. ↩