ZEST AI

The Path To A Fairer Credit Economy

Special Report: Three Ways AI/ML Can Increase Economic Inclusion In America

ZEST AI

Table of Contents

info@zest.ai

A Costly Flaw In The Credit System

Income and social inequality have been endemic in the U.S. from its origins. This spring and summer sparked historic awareness of racial injustice and disparities in the United States. From citizens to corporate and political leaders, there is a nationwide rallying cry to address the structures that perpetuate these racial and economic disparities.

Banks and credit unions want to do more to reach underserved communities, but they are wary of lending to borrowers whose risk they struggle to assess, often people of color, people without a college degree, recent immigrants, and many of the 40 million "credit invisible" Americans.

Research by the Consumer Financial Protection Bureau shows that as many as 15% of African Americans and Hispanics are credit invisible and another 13% of Blacks and 12% of Hispanics have records that are so thin they're treated as unscorable by commercially available credit scoring models. Compared with white credit applicants, people of color are almost twice as likely to be unscorable by traditional methods.

The credit invisible find themselves stuck in a vicious cycle through no fault of their own. If you can't get a credit card, it's hard to get a car, which makes it hard to buy a house or get the capital you need to start a business.

Generations of discrimination and rejection -- and dealing with exorbitantly priced "alternative" credit sources -- have caused many minorities simply to avoid seeking credit at all.

> "Systemic racism is a tragic part of America's history... We can do more and do better to break down systems that have propagated racism and widespread economic inequality."
>
> **Jamie Dimon**
> JPMorgan Chase CEO

Towards A More Holistic View Of Borrower Risk

This doesn't have to be. A lack of credit history doesn't make someone riskier than someone with a robust file. It just makes them harder to score using the traditional credit scoring system, which has been limited to a couple of dozen factors such as credit score, income and current debt outstanding. Overweighting a small number of factors ignores a good deal of information that can greatly impact a lender's decision to approve a loan -- and unfairly penalizes millions of Americans.

This is not the only reason for the yawning wealth gap in America, but it certainly doesn't help. One hundred and fifty years on from the end of the Civil War, the Black community still only owns less than 1% of the total wealth in America. The black homeownership rate, at 60% that of whites, is roughly unchanged since the 1960s.

~2X

Black households are nearly twice as likely to **lack access to credit** as white households.

About 15% of Black households were deemed "credit invisible" or unscorable, compared to about 16% of white households, the Consumer Finance Protection Bureau reported.

This prevents the socioeconomic mobility achieved through mortgages and business loans.

2X

Black-owned businesses get **turned down for bank financing** twice as often as white businesses do.

They also seek it out at rates twice as high, according to a Federal Reserve report.

These factors starve Black-owned businesses of capital to grow or achieve the financial resiliency to recover from major setbacks.

Source: Federal Reserve, Consumer Finance Protection Bureau

The Future Of Fair Lending Is Here

The good news is that change is at hand. Artificial Intelligence and Machine Learning (ML), when used properly, hold the key to ending racial disparity in credit underwriting while at the same time improving the safety and soundness of American financial institutions.

Consumers are demanding change amid the broader national call for equity. According to a recent Harris Poll survey, 7 out of 10 Americans would switch lenders for one with more inclusive practices. Six out of ten would switch to a lender that increased access to credit for people of color. The survey reflected a strong preference across the board for more racial and gender equity in loan underwriting.

Financial institutions are uniquely positioned to act. They can implement innovative technologies and techniques proven to break the cycle of racial and economic inequality and transform daily life for millions of people of color.

Our nation's banks and credit unions already have the data -- such as credit card transactions and cash flows from checking and savings accounts -- that can help solve this problem

Now they need to adopt the new modeling and fair lending approaches that can unlock the marginalized and "credit invisible," generating growth for themselves and opportunity for consumers, and ultimately a more inclusive credit economy.

The path to financial inclusion begins with re-evaluating the status quo and understanding these new technologies and methods. This white paper outlines a new approach to help banks move beyond the status quo to lend more inclusively without taking on added risk.

7/10

Americans would switch to a financial institution that has more inclusive lending practices.

Using ML To Fix What’s Broken

Challenges with Status Quo Methods

Let’s talk about the elephant in the room: Legacy methods for underwriting and fair lending analysis. Traditional credit models typically use regression math to make inferences based on linear relationships between a few variables (i.e., FICO score). They have practical limits on the amount of data that can be included and are difficult to update. Lenders using these legacy techniques find themselves in a fair-lending catch-22: They can either have more inclusive models that underperform or profitable models that continue to ignore minorities and the credit invisible. It’s a false choice, because with new tools and techniques you can have both fairer models and profitable growth.

Machine Learning Breaks the Cycle

Machine learning is a computing technique that makes predictions based on patterns observed in data. Lending offers a wealth of rich data to train ML models to predict delinquencies and defaults with higher precision. Why? ML models can ingest 10 to 100 times more data than logistic regression models, expanding to use a trended and credit-adjacent data such as cash flow, rent, and utility bills that greatly supplement borrower profiles. The real power of ML models comes from their ability to draw insight from millions of correlations among all these data variables. The emergence of ML in lending will entirely reshape the banking and credit industry in the next ten to fifteen years, touching customer acquisition, credit decisioning, anti-fraud, verification, servicing, and cross-selling.

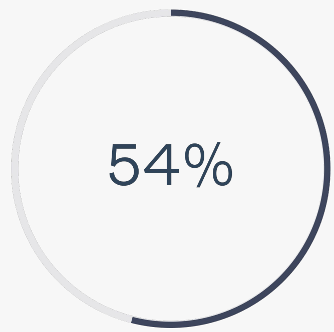

54% of financial services organizations with 5,000+ employees have adopted AI.
Source: Economist Intelligence Group.

The increased predictive power yields real economic gains. Lenders we’ve worked with that have switched to ML underwriting typically see 15% to 20% higher approvals and with that comes a jump in inclusion: More thin-file, no-file, and protected-status applicants get approved.

ML technology can safely widen credit access for low-income borrowers who have been left out of mainstream lending. The use of ML gives financial institutions a new set of fair lending tools and practices that promise fairer models and more accurate denial reasons to consumers on their credit-building journey. Let’s go into three ways Zest is applying innovation to make fair lending easier than it’s ever been.

Ever After: Automating The Search For The Fairest Models

Despite best efforts, many lenders still struggle with the balancing act of making underwriting models accurate and fair. Traditionally, there has been a trade-off between accuracy in risk prediction and minimizing disparate impact. Finding the source of bias in a legacy lending model is not terribly hard. The hard part is getting rid of that bias without wrecking the accuracy of the model. Why? The legacy techniques used to produce less discriminatory alternative models forces lenders to drop crucial variables to improve fairness. Usually those crucial variables, under business justification reasons, end up back in the model and we've failed to move the needle on inclusion. With the shift to ML underwriting, lenders have an arsenal of new techniques to mitigate protected-class bias in any lending model with no hit to accuracy.

Zest AI has patented a technique that automatically optimizes models for accuracy and fairness. If you've ever played with the bass and treble settings on your stereo, you can understand the process.

If you turn the bass and treble all the way down, you get only the bass; you have optimized for bass. But, if you turn both bass and treble all the way up, they end up canceling each other out; you've achieved the same thing as turning both knobs up part-way. Double optimization works the same way in underwriting models, except fairness and accuracy are the knobs instead of bass and treble. Zest's LDA search uses an ML technique called adversarial de-biasing that pairs a credit risk model tuned for maximum accuracy with a companion model that tries to guess the race of the borrower being scored.

The companion model looks at score distributions by race (or any other protected class status) and then instructs the risk model how to calibrate itself until it gets really good at producing evenly distributed, "color-blind" results to fool the companion model. The two models play off against each other repeatedly to create the efficient frontier of models you see on the next page.

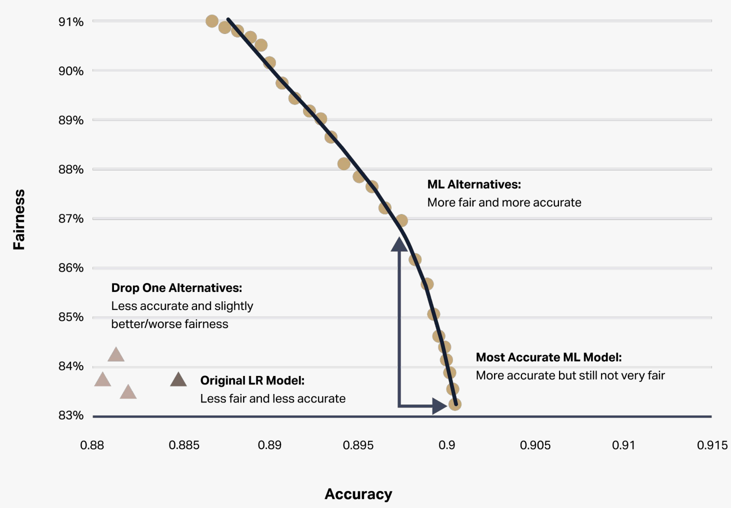

Machine learning and Zest LDA Search provides lenders with fairer models that sacrifice minimal incremental profit.

The graph shows what happens when you double-optimize lending models using standard logistic regression approaches versus ML techniques. The y-axis represents fairness, measured by the Adverse Impact Ratio. The x-axis represents the accuracy of the underwriting decision, measured by the area under the ROC curve, a standard measure of predictive power in data science. Each diamond is a model. Zest's LDA Search produced a series of models along an efficient frontier of accuracy and fairness in different proportions. These models are all very close to one another, meaning that the trade-off between accuracy and fairness is quite small. Contrast the line on the right to the diamonds on the left, which represent a lender's attempts to use "drop one" methods (explained more fully below) to create fairer alternatives to the original LR model. Lenders using LR have to give up a lot of accuracy for only marginal fairness gains. This forces them to "justify" the use of models that risk perpetuating historical inequities.

ML offers lenders real choices with only minor sacrifices of *incremental* profit gain. One auto lender model achieved a 4% increase in approvals for African-Americans for a mere 0.2% drop in performance (loosely translated: about two bucks in incremental profit per loan). A bank that offers personal loans achieved a 6% increase in approvals for borrowers of color for a mere 0.1% drop in performance.

Because of the companion model approach to de-biasing, these gains are achieved without considering protected status as an input in the original risk model. Fairness is considered after the initial model is created. This has the effect of directing the modeling process toward a fairer model without violating regulatory prohibitions on disparate treatment. Lenders can choose the extent to which fairness enters into the process, develop many alternative models, and select one for use in production.

What these results show is that the double-optimization of ML models gives lenders better alternative models to choose from.

Providing Consumers With More Accurate Denial Reasons

Of course, even when models have been optimized for fairness and accuracy, some consumers will be denied credit. Providing these consumers with accurate information about why they may have been denied is critical; it is the only way for them to know what behavior to change or what questions to ask in correcting errors in their credit profile. What is less obvious, however, is that identifying principal denial reasons requires advanced math whether a lender uses an ML model or a LR underwriting model, though, as we will explain, the need is far greater for ML models.

Some lenders use one of two seemingly reasonable methods to identify the principal reasons that their ML and LR models denied an application for credit: “drop one” and its cousin “impute median.” With drop one, lenders test which model variables contribute most to the model score by removing one variable at a time and measuring the change in score as a means of quantifying the importance or influence of that lost variable.

With impute median, instead of dropping a variable, they replace each variable, one at a time, with the median value of that variable in the dataset and measure the result the same way.

Those methods sound reasonable, essentially saying let’s see whether so and so would have been denied if they didn’t have X variable in their credit file or if their X variable were the same as everyone else’s. But, in practice, those methods are often inaccurate for at least two reasons. First, once you change the data that the model considers, you have moved from the real world into a hypothetical one. You end up trying to explain situations that would never happen in the real world, such as where the income variable is missing (because it was kicked-out during the drop one analysis) but the debt-to-income ratio is available. This situation is impossible in reality, and so is the resulting explanation.

Second, those methods produce inaccurate principal reasons when used to explain ML and even LR underwriting models because those methods don't account for the fact that variables interact, that variables are not always independent, and that, in ML models, variables may point in different directions. Of course, LR models are largely blind to variable interactions, making the drop one and impute median methods more accurate when applied to LR models (at least, in the experiments we have conducted). But ML models rely, in part, on modeling interactions among lots of variables to gain predictive power, meaning that drop one and impute median identify the wrong denial reasons almost every time when used on ML models.

We ran an experiment on an auto lending ML model and dataset from one Florida lender and showed that the drop-one method identified the correct principal reason code only 11% of the time, while the impute median method was almost always wrong (see table at right). We've seen similar results in our work for dozens of financial institutions.

Drop One and Impute Median are almost always wrong when run on ML models

Percentage of time the reasons given by each method correctly matched any of the actual top 3.

Technique	1st Reason	2nd Reason	3rd Reason
Drop One	11%	11%	13%
Impute Median	-0%	-0%	1%

- We generated denial reasons for a machine learning model built to approve loans for a mid-sized auto lender using various methods.
- We compared Drop One and Impute Median to the Shapley game-theoretic baseline to assess their accuracy.
- Drop One was wrong ~90% of the time; Impute Median was almost always correct.

These results beg the question: How can any math equation accurately capture and explain the interaction of so many variables? And how can we know with certainty that such an equation is accurately identifying the most important factors influencing the model's decision?

The answer is simple: the mathematics of games. In the 1960's and 70's, certain mathematicians, sociologists, and economists became interested in what has come to be called game theory. Those scholars developed ways to quantify how each player on a sports team contributed to the final score of the game, taking into account baskets, touchdowns, or goals the player scored, as well as the player's assists, passes, and blocks. Game theory pioneer Lloyd Shapley eventually won the Nobel Prize in economics because of this work. Shapley methods offer some of the best ways to explain how ML models make decisions. In the case of an ML model, the "players" are the model variables, the "game" is the model, and the "score" is the model's output (in credit, this usually represents the probability of defaulting on a loan).

To comply with federal fair lending laws, creditors need to know the precise importance of each variable in their models to determine whether one reason is more or less important than the others. Shapley's method precisely quantifies the importance of every variable in generating a given score for a given applicant, and it takes into account complex variable interactions of the kind you see in ML models. Zest, along with other well-respected academics, havew proven this out -- using Shapley's method (and variations on its formal extensions) to explain millions of ML lending decisions for many different kinds of credit products. These methods accurately identify principal denial reasons every time.

Game-theoretic approaches to explainability are consistent with federal regulations and guidance, and also the most effective way to ensure compliance. Our customers use them and we're helping inform bank regulators about the benefits of these approaches so they can issue guidance that speeds their adoption so that all consumers can get the information they deserve.

Building a Better Yardstick: Using ML To Improve Race Prediction

Underpinning all these fair lending innovations is the need to get one data point right: the protected status of a borrower. An accurate estimate of an applicant's race, gender, or age is crucial because, outside of the mortgage business where race information is collected by law, lenders have to guess the race of the borrower to do their fair lending analysis. The tools the industry uses to do so are of limited use. A healthy application of machine learning in one area -- race estimation -- can go a long way toward creating a fairer, more inclusive system.

Today, banks and credit unions estimate the race of borrowers using an algorithm called BISG, or Bayesian Improved Surname Geocoding, which constructs a probability of assignment to race or ethnic group based on surname and the demographic characteristics associated with place of residence. BISG was developed by RAND in 2008 for use in health care, not financial services. According to a paper published by RAND in 2016, BISG was 90% to 96% accurate for the four largest racial/ethnic groups.

In 2014 the CFPB introduced BISG as the new and improved way to proxy race for fair lending purposes. The problem with BISG is that it's often wrong. A 2014 Charles River Associates auto lending study, sponsored in part by some lending institutions, found that BISG correctly identified African-American borrowers a mere 24% of the time at an 80% confidence threshold. Hispanic and Asian borrowers were correctly identified 77% and 60% of the time, respectively.

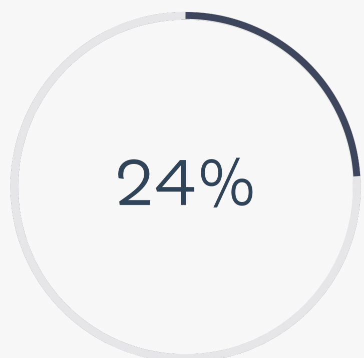

According to a Charles River Associated auto lending study, BISG correctly identified African-Americans a mere 24% of the time.

Why is this happening? BISG is an overly simple model that relies on outdated variables such as the 2010 Census. Using only surnames is also a hindrance as surnames become less predictive of race over time with interracial marriage increasing 3 percentage points every 10 years among newlyweds [according to a Pew Study](). BISG also assumes we live in segregated neighborhoods, yet 80% of the U.S. lives in diverse cities with many urban ZIP codes having more than 100,000 residents -- making it clear why BISG struggles with people of color.

We picked three Zest employees at random and plugged their names and zip codes into an open-source BISG package. All three were mislabeled (see below).

Lenders using a flawed yardstick to assess the racial disparity of loan approvals and denials can end up falsely confident that their credit models are as fair as they ought to be when thousands of denials counted as a white are actually Black declines. Not understanding where real disparity occurs also makes it impossible for the lender to identify problematic lending policies and for regulators to assess harm. Consumers deserve better, as do the lenders and regulators who make the decisions that affect the lives of millions of borrowers. We're now going to walk you through a new ML-based approach that delivers far more accurate race identification.

	CATEGORY	RESULT
	**Surname**	Matthews
	Zip code	91207
	BISG Race/Ethnicity	White
	Actual Race/Ethnicity	African American
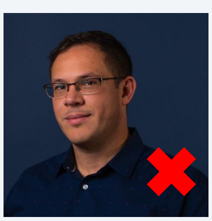	**Surname**	Flo
	Zip code	91362
	BISG Race/Ethnicity	White
	Actual Race/Ethnicity	Hispanic
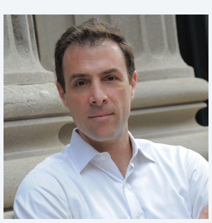	**Surname**	Upbin
	Zip code	90068
	BISG Race/Ethnicity	???
	Actual Race/Ethnicity	White

The Zest Approach To Race Prediction

The Zest data science team has built a simple but powerful ML model called Zest Race Predictor that uses 11 variables including first and last name, and a wider variety of geographic factors. In a test on roughly one million records from the Florida voter database, ZRP correctly identified 30% more African Americans than BISG (420,000 people) and correctly identified 70% fewer whites as non-white. In a test with one auto lender, ZRP offered a 7% improvement in AIR over the baseline ML model. Though early results are promising, there is plenty of work to do to make this better and generalize it to a national population.

The Impact Of Switching From BISG To Zest Race Predictor

400K

Fewer African Americans Counted As Whites

1M

Fewer Whites Counted As People Of Color

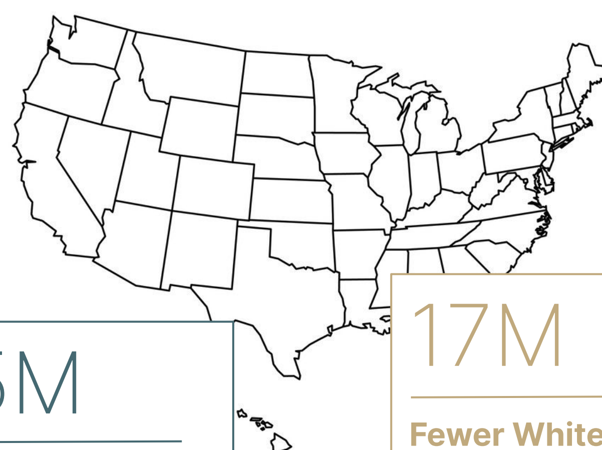

5M

Fewer African Americans Counted As Whites

17M

Fewer Whites Counted As People Of Color

Conclusion

America is at an inflection point. Business leaders in every industry are questioning their practices around diversity and inclusion. They're stepping up and offering plans and investment to address racial inequities. Financial institutions can answer the call for more inclusive lending and help break the cycle of racial inequality. Now is the time to think big and act.

What's exciting is that lenders now have the power to use technology to generate this equity directly through fairer lending decisions. The U.S. credit reporting system is a marvel of data warehousing and analytics, but the data itself is encoded with the effects of generations of discrimination and economic oppression. We can directly mitigate this systemic bias using advanced technology and shrink the disparity in approval rates between whites and people of color by substantial margins.

The path to fairer lending is realistic and no longer a catch-22. As former Citigroup executive Ray McGuire said, "Corporate America needs to have courage in the months and years ahead." The health of our society and economy depends on what we all do next.

ZEST AI

About Zest AI

Zest AI makes the power of machine learning safe to use in credit underwriting. Lenders using Zest AI software make better decisions and better loans—increasing revenue, reducing risk, and automating compliance. Zest AI was founded in 2009 with the mission of making fair and transparent credit available to everyone and is now one of the fastest-growing fintech software companies. The company is headquartered in Los Angeles, California. Learn more at www.zest.ai and connect with us on Twitter and LinkedIn.

With Zest, you can:

Grow

Identify the best borrowers and safely increase approvals

Protect

Mitigate risk, limit losses, maintain performance

Optimize

Exchange pricing, segmentation, and recoveries

Automate

Gain efficiencies and return on high-value activities

Better lending for you and your customers with more powerful insights and more accurate risk assessments, Zest gives you the ability to approve more credit-worthy borrowers, reduce existing losses, offer better rates, and develop more customized policies and rules engines for your business.

LEARN MORE

info@zest.ai