

Case Study: Fair Credit Reporting Act

Steven M. Bellovin

<https://www.cs.columbia.edu/~smb>



What is the Fair Credit Reporting Act (FCRA)?

- The FCRA (15 U.S.C. §1681), from 1970, is one of the oldest privacy laws in the world
- Primarily, it governs relationships between consumers and credit bureaus
 - A credit bureau (technically, a “consumer reporting agency”) is an entity that reports consumer financial behavior and/or compiles reports on people
- The FCRA restricts who can receive credit reports and when
- It also gives consumers certain rights

Why Was the FCRA Passed?

- Credit bureaus date to the mid-19th century—but back then, they mostly collected gossip about the well-to-do, often from servants
 - Are they talking to multiple lenders?
 - Have they cut back spending on the household—dinners, staff, etc.?
 - Do they pay their employees promptly?
 - Are they gamblers?
- The spread of large computers in the 1960s made collecting large amounts of data on individuals feasible
- Academics and legislators had been worrying about giant computer databases

Consumer Rights Under the FCRA

- The right to free annual credit reports from each major credit bureau
- The right to receive a free credit report if it is used against you, e.g., to deny you a loan
- The right to dispute inaccurate information
- Under certain circumstances, especially connected with mortgage applications, the right to learn your *credit score* (15 U.S.C. §1681g) and “all of the key factors that adversely affected the credit score of the consumer in the model used” (15 U.S.C. §1681g(f)(1)(C))

What's a *Credit Score*?

- Legal definition: “a numerical value or a categorization derived from a *statistical tool or modeling system* used by a *person* who makes or arranges a loan to predict the likelihood of certain credit behaviors, including default” [emphasis added] (15 U.S.C. § 1681g(f)(2)(A)(i))
- By far the most common credit score is the FICO (Fair Isaac Corporation) score; the basic algorithm dates to 1958
- What if the “statistical tool or modeling system” is an AI algorithm?
- Is a “person” really involved?

Let's do a deep dive...

FICO Factors

The details are proprietary but some items have been disclosed

- Payment history, including late payments, foreclosures and repossessions, bankruptcies, etc.
- Debt burden, across different types of debts, including balance to limit ratio, number of accounts with balances, etc.
- Length of credit history (by law, data older than 7–10 years generally can't be used)
- Types of credit used
- Recent searches for credit

First Question...

First Question: How was the Model Trained?

First Question: How was the Model Trained?

- Where did the training data come from?
- How was it labeled?
- *Was the training data unbiased?*
 - Does it properly reflect *all* communities?
 - Minorities? Immigrants? Urban, suburban, rural? Young, middle-aged, retired?
 - Does it reflect different usage patterns and repayment probability among all of these groups?

Labeling the Training Data

Labeling the Training Data

- Training data may have to be *labeled*
- That is, humans may have to decide what different data points mean
- Does this labeling process reflect preexisting biases of the labelers?
- Does the training data reflect preexisting societal biases, e.g., historically segregated housing patterns?

Proxy Variables

Proxy Variables

- By law, certain things *cannot* be used in credit decisions, e.g., race, gender, religion, national origin, etc.
- But these categories often correlate with others
 - Names can denote gender or ethnic origin
 - Zip codes correlate with race (even though “redlining” was outlawed more than 50 years ago)
 - Hobbies can correlate with gender
- Are proxy variables to barred categories present in the training data?

Relevant and Irrelevant Variables

Relevant and Irrelevant Variables

- Why should name or hobbies be evaluated for a credit score?
 - But sometimes, there are real but surprising correlations—people with poorer credit scores tend to have worse driving records and hence higher car insurance premiums
- There may be unreported positive data, e.g., a running tab at a local store that you pay off promptly every month
- There may be inapplicable data: in Manhattan, many fewer people have cars and hence do not have car loans

*Problems with training data can be used
to contest almost any AI-based system.*

Building a Contestable FCRA-Compliant AI System

First Steps

- We've built a system and trained it
- We've verified that the training data is ok
- What's next?

First Steps

- We've built a system and trained it
- We've verified that the training data is ok
- What's next?
- *How do you verify the model?*

Verifying the Model

- You can't verify it against the training data—you know it already matches that
- You need new data
- What else?

Verifying the Model

- You can't verify it against the training data—you know it already matches that
- You need new data
- What else?
- What are the statistical patterns of grants and denials?
- Do they show bias?
- Do they show impermissible bias, bias not accounted for by actual economic differences?
- Would you recognize the bias, or is it due to subtle factors, e.g., different rates of certain diseases for different races?

Applicant Data

- Can you verify that the applicant's data was fed in correctly?
- What happens if you re-enter it?
- Is the system being continuously retrained, from new data? Can you rerun the application with data from a previous date?

Who Might Challenge a Determination?

Who Might Challenge a Determination?

- Someone who wants credit?
- A journalist doing an investigation?
- A regulator looking for bias?

What Do Challengers Already Know?

What Do Challengers Already Know?

- Applicants: their own personal financial situation
 - What loans and credit cards they have
 - What their payment history looks like
 - What their credit history is like
 - *Do these match what the credit bureau thinks?* Recall the right to correct erroneous information.
- Journalists and regulators
 - Pre-planned scenarios—but do these fake people have credit histories?

What Might Challengers Need to Learn?

What Might Challengers Need to Learn?

- What were the primary reasons for rejection?
 - Remember the right to learn “all of the key factors that adversely affected the credit score of the consumer in the model used” (15 U.S.C. §1681g(f)(1)(C))
 - Can an AI specify that?
- Other data in their credit reports—they have a right to obtain that
- How a FICO score is calculated? That’s proprietary and closely held—but should it be, if it’s part of the AI model?
 - What has previously been published or disclosed about the challenger’s FICO score?

And a Legal Question...

- A credit score must be “used by a person” to make a credit determination (15 U.S.C. § 1681g(f)(2)(A)(i))
- A pure AI system does not qualify
- But does that person simply blindly accept what the AI says?
 - What is the history of overrides?
 - Can an applicant contest a decision with a person?

Anything Else?

Questions?



Common raven, New York City, March 8, 2024