

Racial and Neighborhood Disparities in Legal Financial Obligations in Jefferson County, Alabama

Óscar Lara Yejas¹, Aakanksha Joshi², Andrew Martinez³, Leah Nelson⁴,
Skyler Speakman⁵, Krysten Thompson¹, Yuki Nishimura⁶,
Jordan Bond², and Kush R. Varshney⁷

¹IBM Silicon Valley Lab, San Jose, CA 95141, USA

²IBM Innovation Studio, Astor Place, New York, NY 10003, USA

³Center for Justice Innovation, New York, NY 10018, USA

⁴Alabama Appleseed, Montgomery, AL 36104, USA

⁵IBM Research – Africa, Nairobi, Kenya

⁶IBM Chief Analytics Office, Armonk, NY 10504, USA

⁷IBM Research – T. J. Watson Research Center, Yorktown Heights, NY 10598, USA

Abstract

Legal financial obligations (LFOs) such as court fees and fines are commonly levied on individuals who are convicted of crimes. It is expected that LFO amounts should be similar across social, racial, and geographic subpopulations convicted of the same crime. This work analyzes the distribution of LFOs in Jefferson County, Alabama and highlights disparities across different individual and neighborhood demographic characteristics. Data-driven discovery methods are used to detect subpopulations that experience higher LFOs than the overall population of offenders. Critically, these discovery methods do not rely on pre-specified groups and can assist scientists and researchers to investigate socially-sensitive hypotheses in a disciplined way. Some findings, such as individuals who are Black, live in Black-majority neighborhoods, or live in low-income neighborhoods tending to experience higher LFOs, are commensurate with prior expectation. However others, such as high LFO amounts in worthless instrument (bad check) cases experienced disproportionately by individuals living in affluent majority-white neighborhoods, are more surprising. More broadly than the specific findings, the methodology is shown to identify structural weaknesses that undermine the goal of equal justice under law that can be addressed through policy interventions.

1 Introduction

In various forms, excessive legal financial obligations (LFOs) menace all corners of the United States of America. Over the past several years, researchers and advocates from across the country have made a compelling case that fines, fees, and forfeitures related to court cases constitute a system of criminalized poverty, keeping low-wealth offenders and their families under surveillance and engaged with the justice system long after their court cases are otherwise resolved. For a variety of historical and political reasons, this system is deeply entrenched in Alabama, where every year, courts assess millions of dollars in court costs, fines, fees, and restitution (collectively, legal financial obligations or

LFOs) (Flynt 2004). Individuals against whom these LFOs are assessed face severe collateral consequences. Of nearly 1,000 Alabamians surveyed in 2018, 83% had forgone necessities like rent, food, medical bills, car payments, and child support; 44% used payday or title loans; 38% committed an additional crime, all to service their LFO debt (Nelson 2018). Even so, almost half were jailed at some point because of unpaid LFOs (Nelson 2018).

Particularly, Black people are overpoliced in Alabama compared to white people. For example, even though both Black and white people use marijuana at about the same rates, Black people in Alabama are more than four times as likely as white people to be arrested for marijuana possession (Nelson et al. 2018). And overall, about 55% of Alabama’s jail and prison population is Black, despite the state’s population being only about 27% Black (Nelson et al. 2018). Since every criminal conviction in Alabama results in some type of LFO, these disparities drive a form of racialized wealth extraction that depletes Black wealth and contributes to Alabama’s racial wealth gap.

A general understanding of the facts that Black people disproportionately bear the debt burden and that many people owe debt they will likely never have the ability to pay is not the same as knowing in detail what drives these disparities and injustices. While the human consequences of excessive LFO debt are well documented, relatively little is known about patterns in the distribution of such debt or precisely which administrative and criminal legal mechanisms drive it. The purpose of this study is to fill this gap in understanding by quantitatively analyzing the characteristics of debt for different social groups and measure existing disparities in LFOs.

Jefferson County, home to about 665,400 people, is Alabama’s most populous county. Its county seat, Birmingham, played a storied role in the civil rights movement of the 20th century as Black Americans fought for equal rights and treatment under law. In the decades since, the county—and the laws that govern it—have evolved, but racial inequities driven by history and habit persist. The population of the county is 53% white, 42% Black, with others making up the

balance.

In 2021, the Jefferson County judiciary agreed to participate in a project developed by research and advocacy organizations aimed at eliminating these inequities and improving the fair administration of justice. The findings presented here are essential to developing usable, evidence-based interventions. Toward this end, we acquired sensitive data from five and a half years over 2014–2019 created by the Jefferson County court and maintained by Alabama’s Administrative Office of Courts. With support from the judiciary, whose wish to gain clarity about the nature and scope of fine and fee inequities was a major driver of this research, the data was obtained and placed in a secure location for analysis. An important contribution of this work, setting aside the analysis conducted, is the collection and preparation of this dataset.

A significant challenge when analyzing bias is accounting for multiple potentially correlated factors (dimensions) such as income, education level, and neighborhood racial distribution, among others. As more dimensions are taken into account, the number of subpopulations given by each combination of factors grows exponentially. This motivates the use of intelligent algorithms to efficiently find the most vulnerable subpopulations, understand their characteristics, and measure the disparities they experience.

In this work, two methodologies are applied to characterize and analyze the fairness of LFOs in Jefferson County, Alabama: (1) descriptive statistics and (2) subpopulation discovery. The former focuses on comparing central tendency statistics (such as mean, median, and mode) across different pre-defined subpopulations. The latter intends to automatically find which subpopulations are experiencing significantly larger LFOs compared to the background population.

In addition to being a novel study design for the questions being explored, the analytic approach used here generated insights of a scale and scope that is significantly different from anything previously available to advocates and practitioners. By creating coherent picture of how LFOs operate across all available charges and demographics, this approach made it possible for the first time observe outcomes at scale and laid the groundwork for exploring options for addressing it.

2 Previous Studies

Racial and ethnic disparities in the criminal legal system have been widely documented. Black Americans (20%) are more likely to report that they are unfairly stopped by the police compared to whites (3%) (Pew Research Center 2016). Moreover, Black Americans experience worse outcomes than whites in prosecutorial charging (Berdejó 2017), sentencing (Rehavi and Starr 2014), and pretrial detention (Sutton 2013). Ultimately, people of color are more likely to be incarcerated than white people (Bronson and Carson 2017).

However, one practice that has received less attention within the broader scope of research focusing on criminal legal system disparities is the levying of fines and fees, more formally referred to as legal financial obligations (LFOs). LFOs are often used by court systems across the United States to sanction people convicted of offenses—and now

constitute the most common sanction within the criminal legal system, especially for low-level misdemeanor charges (Bing, Pettit, and Slavinski 2022). Indeed, in the past several decades, states have added new monetary sanctions to their penal codes or have increased fine amounts (Harris 2016).

Studies have identified racial and ethnic disparities in the levying of legal financial obligations. However, such analyses are often methodologically limited or are not amenable to actionability. First, some studies have only examined disparities across a narrow subset of infractions, which does not offer a comprehensive account of how such disparities may vary across different infraction severity levels (Bing, Pettit, and Slavinski 2022). Others have analyzed a broader continuum of severity levels ranging from petty misdemeanors to felonies, but further studies are needed that can yield additional nuance, by examining disparities across specific charges and across severity levels (Stewart et al. 2022).

Second, research examining disparities in LFOs has largely focused on individual-level attributes, which often include race and ethnicity, and income (Bing, Pettit, and Slavinski 2022; Stewart et al. 2022; Meredith and Morse 2017). However, these studies do not always account for race/ethnicity and income simultaneously, nor do they focus on other key variables, such as age, which have been linked to likelihood of arrest.

Taken together, research examining disparities in legal financial obligations has been fragmented. Studies often examine a narrow set of severity levels or do not examine disparities across specific charges. Studies also focus on a narrow set of individual characteristics and often overlook the role of local geographic factors (e.g., neighborhood characteristics). Thus, robust analytical methods are needed that can comprehensively address these methodological limitations and offer a unified, detailed, and actionable assessment of disparities in LFOs.

3 Materials and Methods

As we summarize in this section, four datasets were gathered: *criminal master*, *LFOs*, *census tract*, and *charges*. As data preparation, sensitive fields were anonymized and curated, addresses were mapped to census tracts to obtain additional socioeconomic data, and all datasets were joined for further analysis and modeling. Two fairness analyses were carried out, namely (1) descriptive statistics and (2) subpopulation discovery.

Data

With support from the Jefferson County judiciary, four data tables created by the Jefferson County court and maintained by the Alabama Administrative Office of Courts were obtained and stored securely for preparation and analysis. They represent approximately 27,000 cases (approximately 13,000 are complete cases which means no missing data) constituting $5\frac{1}{2}$ years of data over the years 2014 to 2019 from 163 census tracts. The *criminal master* table contains demographic information of the offender, charge information, and case information. The *LFOs* table lists the total LFO amount, amount paid, and balance for individual offenders. The *charges* table is a taxonomy of possible

Factor	Values
Race	Black, White
Age	Young (< 30), Middle-Aged (30 to 54), Old (\geq 55)
Gender	Male, Female

Table 1: Offender features and possible values

Factor	Values
Median Income	Low (<\$30,000), High (\geq \$30,000)
Racial Distribution	Black-majority, White-majority
Perc. of High School Grads	Low (\leq 85%), High (>85%)
Perc. of College Grads	Low (\leq 30%), High (>30%)

Table 2: Census tract features and possible values

charges. The *census tract* table contains data on percentage of high school graduates, percentage of people graduating college, percentage of population by race, and median income for individual census tracts. The criminal master table was joined with the census tract table by using the Census Geocoder address look-up tool. The charge details were also appended to it. Further details about the raw data tables are provided in Appendix A.

The data was cleansed to remove missing values and outliers. Additionally, census tract demographics such as median income, high school-educated proportion, and college-educated proportion were binned for a more interpretable analysis. Tables 1 and 2 show the resulting feature sets and their values. *Race*, *Age* and *Gender* are characteristics of the offender. Given the racial distribution of Jefferson County, Alabama, only *Black* and *white* races are considered, which account for over 95% of the total population. *Median Income*, *Racial Distribution*, *Percentage of High School Graduates*, and *Percentage of College Graduates* are features of the census tract where the offender’s primary address is located. The Percentage of High School and College Graduates is calculated on the 25 years old or above population. Further details about the data preparation procedures are provided in Appendix B.

A *subpopulation* is defined as a subset of offenders determined by a combination of values of the three offender features and four census tract features, resulting in an exponential number of possible subpopulations. Given the cardinalities of the features here, there are 5103 different subpopulations.

Fairness Analysis

Two types of fairness analysis are conducted, namely (1) descriptive statistics and (2) subpopulation discovery. Descriptive statistics analysis compares central tendency measures (e.g., mean and median) of LFOs between manually hypothesized subpopulations, for instance, comparing LFOs between Black-majority and white-majority census tracts. In subpopulation discovery, we use a data-driven search method named *Multi-Dimensional Subset Scan* (MDScan) to find subpopulations that have higher-than- or lower-than-expected LFOs without relying on pre-specified subpopulations (Speakman et al. 2023). This discovery method can identify a subpopulation spanning multiple features. The

goal of both fairness analyses is to better understand how LFOs are distributed across subpopulations of offenders. Traditional statistical analysis requires the subpopulation to be described first (e.g., Black offenders) and then the analysis of the fines follow. In contrast, subpopulation discovery has subpopulations as the result of the analysis and not the prerequisite.

The primary target variable in the analysis is the *total amount of legal financial obligations being charged* to the offender. There are other relevant consequences an offender may experience, such as imprisonment time and loss of their driving privileges, but they are beyond the scope of this paper.

It is worth highlighting that the LFO amount varies depending on the *charge*. For example, *breaking/entering a vehicle* would generally result in smaller LFOs than *possession of controlled substances*. Also, charges can be either felonies or misdemeanors and they have different classes A, B, C, or D, depending on their severity. Therefore, all analyses presented in this study are conducted separately for each charge. Since meaningful analyses require a sufficiently large sample size, we filtered the dataset to the top 30 charges by volume of LFOs.

MDScan is an exploratory analysis method comprising an iterative ascent procedure that maximizes a *scoring function* (typically a likelihood ratio or goodness-of-fit measure) over the exponentially-many subpopulations spanning multiple features. A brute-force search would be computationally infeasible for even moderate sized datasets and manual attempts to increase the score may be viewed as data-dredging or ‘p-hacking’. MDScan is guaranteed to converge to a local optimum such that *any single change* to the identified subpopulation would decrease the score. In this work, the Berk-Jones non-parametric scoring function was used. Further details about MDScan are provided in Appendix C.

4 Results

In this section, we present the results of the fairness analysis, both descriptive statistics and subpopulation discovery.

Descriptive Statistics

First, let us present some key facts about the debt landscape drawn from the data that lead to hypotheses of LFO disparities. Although only 42% of the population is Black in Jefferson County, 58% of the cases correspond to Black offenders, and 63% of the cases are from Black-majority census tracts. Additionally, 76% of cases have LFOs that are not fully paid or collected, hinting that offenders may be facing financial difficulties. Also, 75% of cases correspond to census tracts characterized by low income (less than USD \$35K). We hypothesize that Black individuals, and people living in majority-Black neighborhoods and neighborhoods with a high proportion of low-income households experience disproportionately high debt burdens.

To investigate this hypothesis and quantify potential disparities in LFOs in Jefferson County, we compare central tendency measures between Black and white populations. Albeit unidimensional (as it only considers *race* as a single feature), this approach provides a good initial high-level

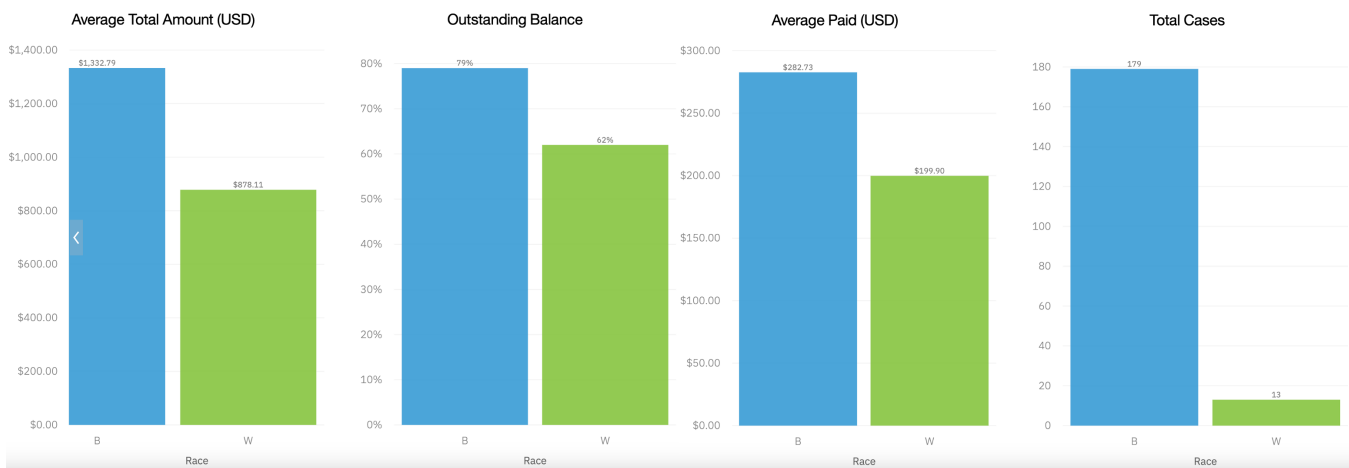


Figure 1: Descriptive statistics comparisons for possession of marijuana 1st Degree (Class C Felony). The blue bars represent Black offenders and the green bars represent white offenders.

understanding of the variability in LFOs, which is later complemented by subpopulation discovery.

For a given charge, the *average total LFOs*, *average outstanding balance*, *average amount paid*, and *total number of cases* were computed across the two populations (Black and white). As an example, Figure 1 shows this comparison for the specific charge of Possession of Marijuana 1st degree (class C felony). Three main results are:

1. On average, Black offenders are charged with 52% higher LFOs than white offenders for Possession of Marijuana 1st degree.
2. On average, Black offenders pay 42% more in LFOs than whites for Possession of Marijuana 1st degree.
3. Black offenders are over 13 times more likely than whites to get charged with a Class C felony for Possession of Marijuana 1st degree.

The results for this and other charges were placed in a dashboard that allows policymakers and scientists to easily interact with the statistics.

In Alabama, first-time possession of any amount of marijuana can be charged as either a Class A misdemeanor or a Class C felony, depending on whether the charging entity believes the marijuana is for personal use. Second-time possession for personal use is always charged as a felony. Fines and fees can run into the thousands of dollars, with felonies carrying steeper financial penalties than misdemeanors. Given longstanding research indicating that Black and white people use (and therefore possess) marijuana at close to the same rate (Burlew, McCuistian, and Szapocznik 2021), arrest rates and LFOs for the offense of possession should be roughly similar across the two races. However, the data tell a different story: for Possession of Marijuana 1st degree, Black offenders are charged 52% higher LFOs, pay 42% more, and are over 13 times more likely to be charged, compared to whites.

As a second insightful example, the map in Figure 2 shows the characteristics of the 320 total cases in Census

Tract 118.02. 29% more Black people (4.0K) live in the tract than white people (3.1K), but there are 2.85 times more cases involving a Black offender than white (237 vs. 83). Thus, on average, Black individuals are 2.2 times more likely to be charged in the tract (across all charges).

The descriptive statistics analysis presented thus far relies on manually specifying the subpopulations and target variables. Through the dashboard, additional analysis could be conducted to assess different charges and census tracts. Though this approach allows scientists and policymakers the flexibility to freely explore different subpopulations and discover racial disparities, it is not scalable when considering multiple features (e.g., income, education level, age, etc.), due to the large number of combinations to explore.

Subpopulation Discovery

Upon running the automated MDScan algorithm on the most common 30 charges in Jefferson County, it was found that subsets of the Black population were consistently identified as having most anomalously high LFOs. Example results are presented for *1st* and *2nd Degree Possession of Marijuana*, as well as for *Negotiating a Worthless Negotiable Instrument* (i.e., giving a bad check). Among all charges analyzed, *1st Degree Possession of Marijuana* exhibits the highest level of disparity. As mentioned previously, there is longstanding research indicating that Black and white people use (and therefore possess) marijuana at close to the same rate (Burlew, McCuistian, and Szapocznik 2021), so arrest rates and LFOs for the offense of possession should be roughly similar across the two races. Given this established baseline, outcomes of MDScan for 1st and 2nd Degree Possession of Marijuana were of particular interest.

The most anomalous subpopulation found by MDScan for the *1st Degree Possession of Marijuana* charge is illustrated in Figure 3. Among the 194 people charged with 1st Degree Possession of Marijuana, the algorithm found that middle-aged, Black males are 9.8 times more likely to receive a high LFO. A high LFO was defined by the algorithm as being

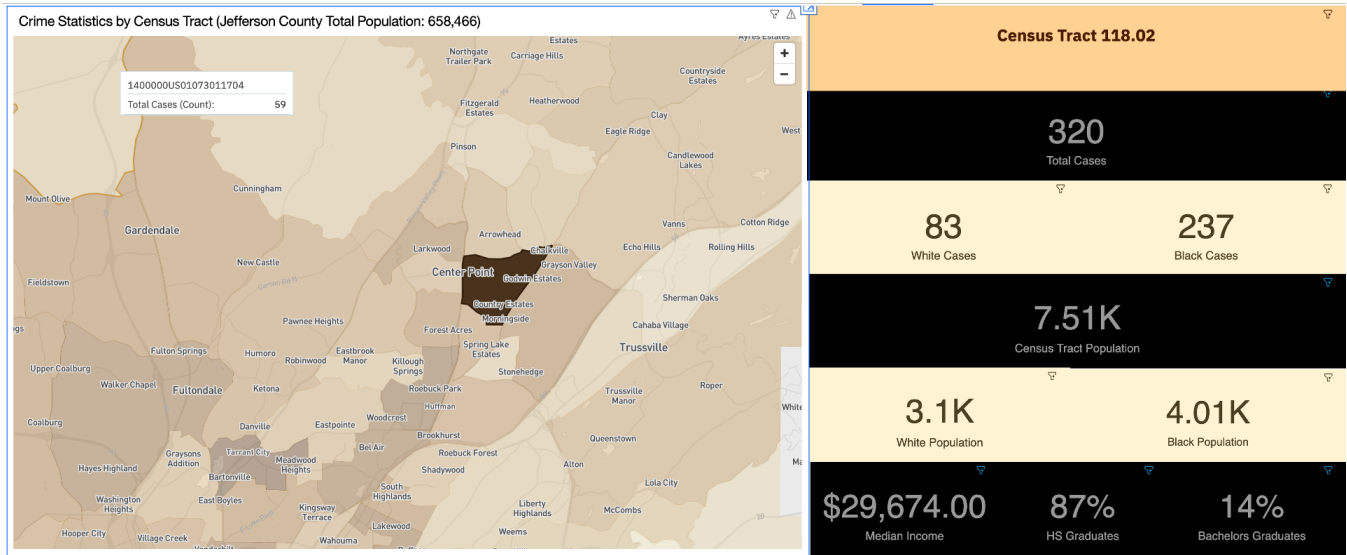
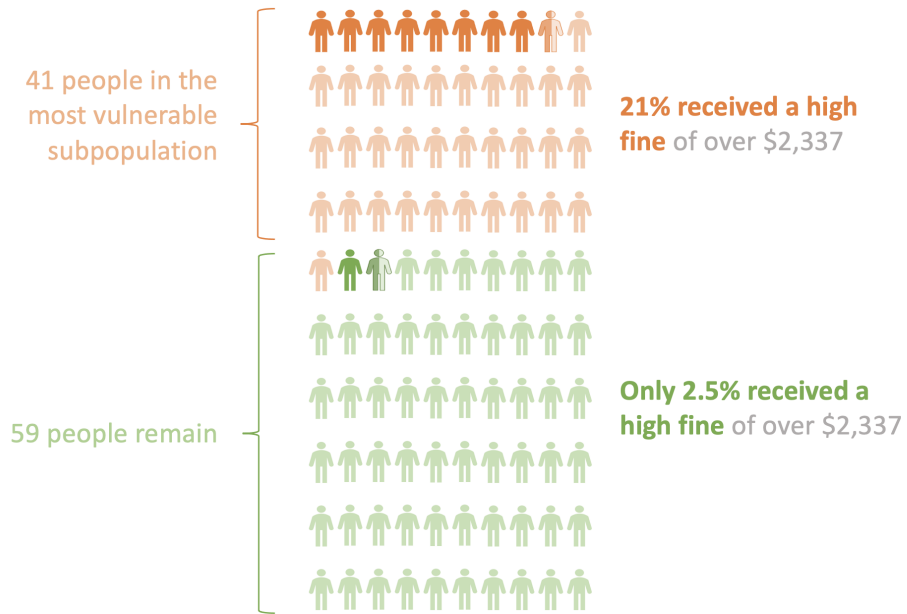


Figure 2: Census Tract Crime Statistics Map (all charges)

For every 100 people who were fined with **1st degree possession of marijuana**, **41 were in the most vulnerable subpopulation** identified by the algorithm.



The most vulnerable subpopulation consists of **middle-aged, Black males** and is **9.8 times more likely to receive a high fine** as compared to the others.

Figure 3: Distribution (in percentage) of the most vulnerable sub-population vs. the rest for 1st degree Possession of Marijuana. Orange section represents the most anomalous subpopulation, while green section represents the rest. Sections in darker colors (both orange and green) represent individuals who received a high fine.

over USD \$2,337. The anomalous subpopulation constitutes 41% of the total population charged with 1st Degree Possession of Marijuana. However, it represents 85% of the total population with an LFO over USD \$2,337. Appendix D provides a worked example of obtaining the 9.8 odds ratio from other relevant quantities.

Among the 443 people charged with 2nd Degree Possession of Marijuana, MDScan found that young or middle-aged black people living in tracts with low college education are 60% more likely to receive a high LFO. A high LFO was defined by the algorithm as being over USD \$716. This anomalous subpopulation constitutes 67% of the total population charged with 2nd Degree Possession of Marijuana.

In Alabama, Negotiating a Worthless Negotiable Instrument (NWN) is the legal term for making a payment with a check that later bounces. Based on the analysis of the 800 people with this charge, MDScan found that middle-aged males living in tracts with high income, high college educations and white majority are 4.67 times more likely to be assessed a high LFO (over USD \$803). This subpopulation comprises 5% of the total population charged with this offense. This subpopulation (living in high-income, high-college-education, majority-white neighborhoods) is unusual compared to the overall trends identified in the analysis for other charges. The debt load in NWN cases partly increases with the amount of the check that triggered the charge. For any number of reasons, individuals who live in majority-white, affluent neighborhoods convicted of NWN charges seem to be writing bad checks in higher amounts than other individuals convicted of the same crime in Jefferson County.

5 Limitations

As noted previously, this analysis was performed on cases where we had complete information on the 7 factors being accounted for about the offenders. Missing information about offenders could be due to challenges with data collection and storage or could reflect more deeper social challenges in getting access to information such as zip codes for certain offenders. The study can be continued to understand those data collection challenges and assess if the disparity widens or narrows down when that population is added into the mix.

The results section of the paper focuses on the most vulnerable subpopulation identified by the algorithm. This does not mean there are not other vulnerable subpopulations in the dataset. The study can be continued to identify other vulnerable subpopulations in the dataset for the charges of interest.

6 Actionability of Results

The legitimacy of the American legal system is premised on the notion that all who are subject to it are entitled to, and receive, “equal justice under law.” But for many reasons, equal justice under law is difficult to achieve. Impartial consideration of what causes the legal system to fall short of its aspiration to deliver equal justice to all is crucial in all reform efforts. It is even more critical in a place where, in

living memory, Black residents’ peaceful demands for equal treatment were met with state violence. Jefferson County, Alabama has repeatedly been at the epicenter of battles over how courts at all levels will interpret their duty to ensure equal justice. It has been the crucible of reform and the crucible of reaction. Given the county’s history, concerns about bias in any criminal legal reform effort are warranted.

The idea for this research was generated by a multidisciplinary team that included data scientists, judges, and social policy advocates who work closely with people impacted by LFO debt they struggle to pay. The fundamental purpose of that collaborative, called the Jefferson County Equitable Fines and Fees (JEFF) Project, is to explore and map how LFOs operate in Jefferson County, to the end of identifying ways to end harmful inequities and improve the fair administration of justice. LFOs implicate two thorny areas of social policy: criminal justice and revenue. For practical, political, and cultural reasons, evidence-based change in these two policy areas is difficult to accomplish.

Because they do not rely on human decision-making about the sequence of queries input to identify anomalies or disparities, the discovery methods used to identify subpopulations in this analysis are more insulated from the usual forms of bias that may infect and compromise efforts to advance the cause of equal justice. For instance, the finding that LFOs disproportionately impact Black residents and people who live in low-wealth and majority-Black neighborhoods shocks the conscience, but it does not challenge lived experience and the bias it can generate. Jefferson County’s history of racism is well documented and known to its residents, and would reasonably be expected to inform any reform effort (ami 2022). But the finding that individuals living in affluent majority-white neighborhoods are disproportionately likely to experience high LFOs in worthless check cases is genuinely surprising and points to the distinctive promise of this methodology: by revealing results we would not have thought to seek, it hones our ability to identify the structural weaknesses that undermine the goal of equal justice under law.

In 2022, the results of this analysis were presented to members of the judiciary of Alabama’s 10th Judicial Circuit, which has jurisdiction over the majority of Jefferson County. The findings prompted concern and additional funding was obtained to support research into drivers of disparities. Over the next few years, it is expected that the Circuit will develop, pilot, and track the consequences of targeted interventions aimed at eliminating the disparities first identified through this analysis. If successful, LFO debt disparities in Jefferson County should shrink; the overall burden of LFO debt should no longer fall most heavily on low-income neighborhoods with disparately high Black populations. The evidence generated by this analysis will not serve as a silver bullet, but by thoroughly and for the first time illuminating the system as it exists, it created space for new thinking about what can and should be done.

Since this analysis was completed, the JEFF collaborative has conducted additional qualitative and quantitative research and analysis within Jefferson County, which in turn served as the basis of practice and policy recommendations

presented to high-level stakeholders within all branches of Alabama’s state government. The JEFF Project also serves as the framework for a national initiative that will ask similar questions about the fairness and efficiency of LFO schemes in jurisdictions in five different U.S. states.

7 Ethical Considerations

Our work sought to take the measure of demographic and neighborhood disparities and create a map of how these disparities (which reflect the racist and classist values of the men who framed Alabama’s constitution) as they play out in Jefferson County, Alabama in the form of legal financial obligations. We did this in the hope that illustrating the depth of disparities, and teasing out details about what drives them, would support the creation a roadmap for a more equitable alternative system.

Two years on, our findings are being leveraged by system actors to do just that. But misuse potential exists because the same information that illuminates opportunities to reduce demographic and neighborhood disparities could also be used to identify ways to increase them. If state actors wished to increase the debt assessed against any given vulnerable subpopulation, they could use our methods to identify opportunities to do so. To mitigate this risk and ensure that we understood this community’s concerns, our team partnered with a community-based organization whose staff includes individuals whose lives and neighborhoods are negatively impacted by legal financial obligations.

A Data Collection Details

Figure 4 shows a relational data model of the four data tables used in this research. They are detailed as follows.

1. *Criminal Master*: contains demographic information of the offender, including race, sex, and age; charge information, such as type (*felony* or *misdemeanor*) and class (A, B, C, or D); and case information, such as offense, arrest, filing, and indictment dates.
2. *LFO Amounts*: for a given offender, the data includes total LFOs, amount paid, and outstanding balance.
3. *Charges*: provides a dictionary of all possible charges in the system, including types (felony or misdemeanor) and class (A, B, C, or D).
4. *Census Tract*: a census tract is a small geographic territory within a county used for statistical purposes. Census tracts are designed to be relatively homogeneous in terms of demographics, economic status, and living conditions (Domínguez-Berjón et al. 2005). On average, a census tract population is around 4,000 inhabitants. Since no granular socio-economic data are available for offenders, census tract data can be used as a proxy to understand the impact of education, income, housing, etc. Examples of census tract data are percentage of high school graduates, percentage of people graduating college, percentage of population by race, and median income.

B Data Preparation Details

Data Encoding

Many government systems still use legacy data management tools such as COBOL, which rely on older data encoding formats. Most of the datasets we analyzed were provided in fixed-width format, where fields have a constant number of characters regardless of the value length, and unused characters are filled with spaces. This added an additional layer of complexity, especially around formatting numeric fields with decimals. Using data dictionaries, all datasets were parsed into CSV format via Python scripts.

Mapping Addresses to Census Tracts

The U.S. government provides an online Census Geocoder address look-up tool (geo 2022) for public use. The tool allows a user to submit an individual address or upload batches of up to 9,999 addresses. It also provides various geographic data for an address, including latitude, longitude, and corresponding census tract. The tool is not perfect and the FAQs provide helpful information when inconsistencies are encountered. For example, if the same group of addresses is run through the tool twice, the number of positive matches may vary. This is a known issue and related to the tool’s processing load capabilities. There are also specific situations where a match will not be found because the address or address range does not exist in the database. An address not being in the database can be because the address is non-residential or non-commercial, because it was recently constructed or demolished, the address range data was missing, or for other reasons. If an address is located in a sparsely populated area, tract information may not be provided because it could identify individuals or companies which is a violation of U.S. Code Title 13.

From the original Criminal Master dataset, addresses were extracted, cleaned, and de-duplicated. A total of 19,414 addresses were run through the Census Geocoder tool. Of the 19,414 addresses, 12,577 (65%) had a match while 3,598 (19%) did not. Another set of 3,028 (16%) addresses returned empty content. This can happen when processing addresses in bulk and requires multiple runs of the data through the tool. A small fraction of 211 (1%) addresses were returned with match-type designation “tie”. This means there were multiple possible results and further investigation was necessary.

Addresses in the criminal master dataset for prisons, county jails, or other county organizations (where people charged with a crime might reside), returned “no match” since commercial addresses are excluded from the Census Geocoder database. Also, in some cases, the address was marked as “TRANSIENT” which meant the person was homeless. This would also return “no match”.

Data Merge

Once addresses were mapped to census tracts, it was possible to join the Criminal Master dataset with the Census Tract dataset, with the goal of expanding the demographics data available (as only race, sex, and date of birth of the offender were available). Since the offender’s income and education

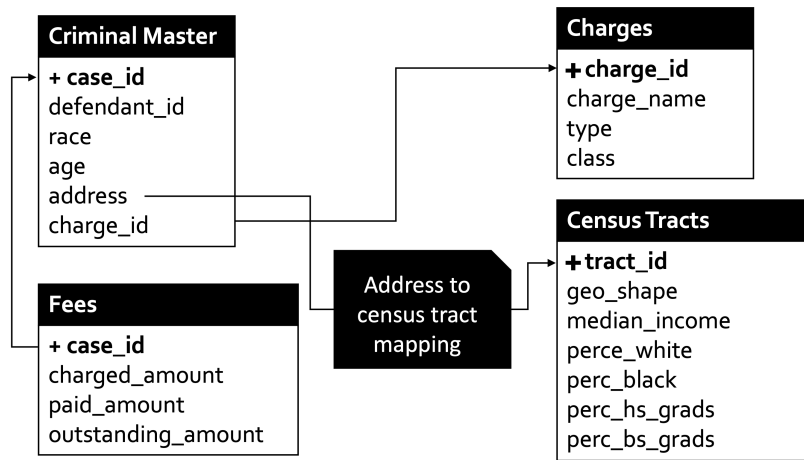


Figure 4: Relational data model. This diagram includes only a subset of the most relevant columns for each of the datasets.

level were unavailable in the data obtained from the court, neighborhood (census tract) demographics were used as a proxy. Racial distribution (percentage of Black and white populations) are also part of the Census Tract dataset. The birth date was used to compute the age of the offender at the time of the offense.

Finally, our dataset was merged with the Charges dataset, which is a dictionary of all charges in the state of Alabama, with their corresponding type (felony or misdemeanor), and class (A, B, C, or D, depending on the severity).

Data Cleaning

Case-level and census tract-level data was cleansed to remove missing values and outliers. Specifically, cases in the data with no matching census tracts, no education level, census tracts with less than 1000 in total population, non-felony or misdemeanor charge type cases, and cases with no charge type were removed from the data. Additionally, census tract demographics such as median income, high school-educated proportion, college-educated proportion were binned (i.e., categorized) for a more interpretable analysis. Details of the binning are described in Tables 1 and 2. For the census tract-level data, extra features were generated to conduct analysis, including *total number of cases*, *number of cases per 1000 people*, *median total amount*, *median amount paid*, *median outstanding balance*, *percentage with outstanding balance*, and *white-to-black population ratio*.

C MDSCan Details

Multidimensional subset scanning (MDSCan) is a general bias scan method. It detects and identifies which subgroups of features have statistically significant predictive bias for a binary, multi-class or numeric outcome.

Traditional statistical analysis requires the subpopulation to be described first (e.g., Black offenders) and *then* the analysis of the LFOs follow. In contrast, subpopulation discovery has subpopulations as the result of the analysis and not the prerequisite. This contrast persists with regression techniques as researchers must specify which feature(s) will be

used in the model before it is fitted/trained on the outcome of interest. Regression techniques and tests of significance are examples of *confirmation analysis*. While these methods are ubiquitous and useful, John Tukey reminds researchers that *exploratory analysis* is also critical: “Finding the question is often more important than finding the answer” (Tukey 1980).

MDSCan is an exploratory analysis method originally developed as a tool for bio-surveillance (Neill and Kumar 2013). It has since been extended to detect predictive bias (Zhang and Neill 2016), heterogeneous treatment effects (McFowland, Somanchi, and Neill 2023), and systematic deviations in data and models, more generally (Speakman et al. 2023). MDSCan is an iterative ascent procedure that maximizes a scoring function (typically a likelihood ratio or goodness-of-fit measure) over the exponentially-many subpopulations spanning multiple features/dimensions. A brute-force search would be computationally infeasible for even moderate-sized datasets and manual attempts to increase the score may be viewed as data-dredging or ‘p-hacking’. MDSCan is guaranteed to converge to a local optimum such that *any single change* to the identified subpopulation would decrease the score. These guarantees and efficiencies are due to the linear-time subset scanning (LTSS) property of the scoring function (Neill 2012; Speakman et al. 2016).

In this work, a non-parametric scoring function was maximized over all possible subpopulations. Non-parametric scoring functions make fewer assumptions on the underlying distribution of the outcome of interest than their parametric counterparts and this is critical if the outcome distribution is bi-modal and/or heavily skewed. Non-parametric measures do not rely on means or standard deviations but rather on ranked order statistics. For example, parametric assumptions may report the mean LFO amount for a crime, whereas a non-parametric would report that 30% of LFOs exceeded a given threshold. Common examples of non-parametric goodness-of-fit measures include Kolmogorov-Smirnov, Higher Criticism, Anderson-Darling, and Berk-

	FINE < 2337	FINE >= 2337	TOTAL	P (High Fine)	20/194	0.10
Vulnerable	64	17	81	P (High Fine Vulnerable)	17/81	0.21
Other	110	3	113	P (High Fine Other)	3/113	0.026
TOTAL	174	20	194			

The odds of getting a high fine for everyone = $P(\text{High Fine}) / 1 - P(\text{High Fine})$ 0.11

The odds of getting a high fine for Vulnerable group = $P(\text{High Fine} | \text{Vulnerable}) / 1 - P(\text{High Fine} | \text{Vulnerable})$ 0.265

The odds of getting a high fine for Other group = $P(\text{High Fine} | \text{Other}) / 1 - P(\text{High Fine} | \text{Other})$ 0.027

The **ODDS RATIO** for **Vulnerable** vs **Other** 0.265/0.027 **9.80**

The **Vulnerable** group is **9.80 times more likely to receive a high fine** as compared to the **Other** group

Figure 5: Analysis table for 1st Degree Possession of Marijuana

Charge	Subpopulation	Total cases	Total cases for sub-population	Threshold (US dollar)	Odds ratio
UPCS POSS CONTR. SUBS.	OLD, YOUNG, TRACTS WITH HIGH INCOME, MALE, TRACTS WITH LOW HS EDUCATION, TRACTS WITH BLACK MAJORITY, BLACK	1,729	18	\$1,598	7.66
NWNI NEGOTIATING WORTHLES	TRACTS WITH HIGH COLLEGE EDUCATION, TRACTS WITH WHITE MAJORITY, MIDDLE AGED, MALE, TRACTS WITH HIGH INCOME	800	41	\$803	4.67
TET3 THEFT OF PROPERTY 3R	MALE, TRACTS WITH HIGH HS EDUCATION, TRACTS WITH LOW INCOME	415	26	\$987	8.98
TOP2 THEFT OF PROPERTY 2N	TRACTS WITH LOW COLLEGE EDUCATION, MALE, TRACTS WITH HIGH INCOME, MIDDLE AGED	414	40	\$2,828	8.76
VAPM POSS MARIJUANA 2ND	MIDDLE AGED, YOUNG, BLACK, TRACTS WITH LOW COLLEGE EDUCATION	443	298	\$716	1.6
VDR1 USE/POSSESS DRUG PAR	TRACTS WITH HIGH COLLEGE EDUCATION, TRACTS WITH HIGH INCOME	306	34	\$837	7.26
PINT PUBLIC INTOXICATION	TRACTS WITH LOW COLLEGE EDUCATION, MIDDLE AGED, TRACTS WITH HIGH INCOME, FEMALE, TRACTS WITH HIGH HS EDUCATION	290	19	\$462	6.31
BEMV BREAK/ENTER VEHICLE	MIDDLE AGED, TRACTS WITH LOW INCOME, BLACK	279	59	\$597	3.34
BUR3 BURGLARY 3RD DEGREE	TRACTS WITH LOW COLLEGE EDUCATION, MALE, TRACTS WITH LOW INCOME, TRACTS WITH HIGH HS EDUCATION	295	27	\$1,359	3.76
SEX3 SORNA VIOLATION	MIDDLE AGED, TRACTS WITH HIGH INCOME, BLACK, TRACTS WITH LOW HS EDUCATION, TRACTS WITH BLACK MAJORITY	198	3	\$1,504	37
FLEE ATTEMPT TO ELUDE	YOUNG, TRACTS WITH HIGH INCOME, TRACTS WITH HIGH HS EDUCATION, TRACTS WITH WHITE MAJORITY	185	19	\$824	7.49
VAPF POSS MARIJUANA 1ST	MALE, MIDDLE AGED, BLACK	194	81	\$2,337	9.74
FRCC FRAUD USE CREDIT/DEB	TRACTS WITH LOW HS EDUCATION, MIDDLE AGED, TRACTS WITH HIGH INCOME	174	6	\$868	12.87
CPF2 POSS FORGED INSTR 2N	MIDDLE AGED, TRACTS WITH HIGH INCOME, MALE, BLACK	148	9	\$1,902	26.91
ASS2 ASSAULT 2ND DEGREE	TRACTS WITH WHITE MAJORITY, TRACTS WITH LOW COLLEGE EDUCATION, YOUNG, WHITE	160	9	\$653	17.17
REST RESISTING ARREST	MALE, MIDDLE AGED, YOUNG, TRACTS WITH LOW HS EDUCATION	160	65	\$530	4.3

Figure 6: Additional charges and their MDScan analysis outcomes

Jones. All of these measures satisfy the LTSS property and may be optimized within MDScan (McFowland, Somanchi, and Neill 2023; McFowland, Speakman, and Neill 2013). The Berk-Jones statistic was used in this work; it maximizes the divergence between the expected and observed number of LFOs that exceeded a threshold dollar amount.

D Additional Results

Worked Example of Odds Ratio Computation

Figure 5 illustrates how to obtain the odds likelihood value from the subpopulation and the high dollar threshold value for 1st Degree Possession of Marijuana. The same can be repeated for other charges as well, once the subpopulation and the high dollar threshold value have been identified by MDScan.

Results from Other Charges

Figure 6 illustrates a table with other additional charges and their outcomes from the MDScan analysis.

Acknowledgements

We would like to acknowledge the support and contributions from Dr. John Speir, Dr. Sarah Picard, Dr. Julian Adler, Hon. Stephen Wallace, Hon. Michael Streety, Rebecca Duane, Joppe Geluykens, and Michelle Mullins.

References

2022. Brief of the Central Alabama Fair Housing Center, the Fair Housing Center of Northern Alabama, and the Center for Fair Housing as Amici Curiae in Support of Appellees and Respondents. Supreme Court of the United States.

2022. US Government Census Geocoder tool. <https://www.census.gov/programs-surveys/geography/technical-documentation/complete-technical-documentation/census-geocoder.html>.

Berdej , C. 2017. Criminalizing Race: Racial Disparities in Plea Bargaining. *Boston College Law Review*, 59: 1187.

Bing, L.; Pettit, B.; and Slavinski, I. 2022. Incomparable Punishments: How Economic Inequality Contributes to the Disparate Impact of Legal Fines and Fees. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 8(2): 118–136.

Bronson, J.; and Carson, A. 2017. Prisoners in 2017. *U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics*, 137–156.

Burlew, K.; McCuistian, C.; and Szapocznik, J. 2021. Racial/Ethnic Equity in Substance Use Treatment Research: The Way Forward. *Addiction Science Clinical Practice*, 16(1): 50.

Dom nguez-Berj n, M.; Borrell, C.; L pez-Gimenez, R.; and Pastor, V. 2005. Mortality and Socioeconomic Deprivation in Census Tracts of an Urban Setting in Southern Europe. *Journal of Urban Health*, 82: 225–236.

Flynt, W. 2004. *Alabama in the Twentieth Century*. Tuscaloosa, AL: University of Alabama Press.

Harris, A. 2016. *A Pound of Flesh: Monetary Sanctions as Punishment for the Poor*. Russell Sage Foundation.

McFowland, E., III; Somanchi, S.; and Neill, D. B. 2023. Efficient Discovery of Heterogeneous Quantile Treatment Effects in Randomized Experiments via Anomalous Pattern Detection. arXiv:1803.09159.

McFowland, E., III; Speakman, S.; and Neill, D. B. 2013. Fast Generalized Subset Scan for Anomalous Pattern Detection. *Journal of Machine Learning Research*, 14(12): 1533–1561.

Meredith, M.; and Morse, M. 2017. Discretionary Disenfranchisement: The Case of Legal Financial Obligations. *Journal of Legal Studies*, 46(2): 309–338.

Neill, D.; and Kumar, T. 2013. Fast Multidimensional Subset Scan for Outbreak Detection and Characterization. *Online Journal of Public Health Informatics*, 5.

Neill, D. B. 2012. Fast Subset Scan for Spatial Pattern Detection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(2): 337–360.

Nelson, L. 2018. Under Pressure: How Fines and Fees Hurt People, Undermine Public Safety, and Drive Alabama’s Racial Wealth Divide. <https://alabamaappleseed.org/underpressure/>.

Nelson, L.; Tucker, W.; Dills, K. A.; and Redford, A. 2018. Alabama’s War on Marijuana: Assessing the Human and Fiscal Toll of Criminalization. <https://alabamaappleseed.org/waronmarijuana/>.

Pew Research Center. 2016. On Views of Race and Inequality, Blacks and Whites Are Worlds Apart. <https://www.pewresearch.org/social-trends/2016/06/27/on-views-of-race-and-inequality-blacks-and-whites-are-worlds-apart/>.

Rehavi, M. M.; and Starr, S. B. 2014. Racial Disparity in Federal Criminal Sentences. *Journal of Political Economy*, 122(6): 1320–1354.

Speakman, S.; Somanchi, S.; III, E. M.; and Neill, D. B. 2016. Penalized Fast Subset Scanning. *Journal of Computational and Graphical Statistics*, 25(2): 382–404.

Speakman, S.; Tadesse, G.; Cintas, C.; Ogallo, W.; Akumu, T.; and Oshingbesan, A. 2023. Detecting Systematic Deviations in Data and Models. *Computer*, 56(2): 82–92.

Stewart, R.; Watters, B.; Horowitz, V.; Larson, R.; Sargent, B.; and Uggen, C. 2022. Native Americans and Monetary Sanctions. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 8: 137–156.

Sutton, J. 2013. Structural Bias in the Sentencing of Felony Defendants. *Social Science Research*, 42: 1207–21.

Tukey, J. W. 1980. We Need Both Exploratory and Confirmatory. *The American Statistician*, 34(1): 23–25.

Zhang, Z.; and Neill, D. B. 2016. Identifying Significant Predictive Bias in Classifiers. arXiv:1611.08292.