INTELLIGENT
INFRASTRUCTURES

THINK PIECE | FEBRUARY 2022

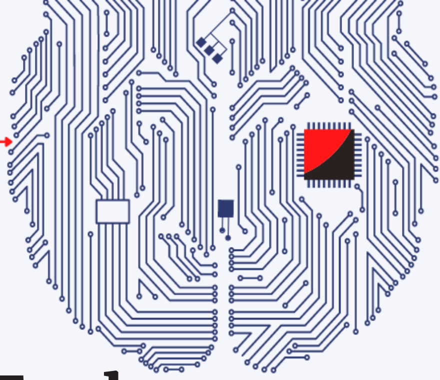

Credit Scoring Algorithms as Tools for Financial Inclusion
A Development Perspective

A comprehensive credit reporting system[1] is vital to a country's financial infrastructure and can contribute to increased financial inclusion, responsible finance, and financial stability. Credit reporting addresses a fundamental problem of credit markets: asymmetric distribution of information between borrowers and lenders which may lead to adverse selection, credit rationing, and moral hazard problems. Therefore, regulators and financial market participants increasingly recognize the importance of credit reporting systems in credit risk evaluation and overall credit portfolio management, financial supervision and financial sector stability, along with enhanced credit access.

The world of credit reporting is no stranger to disruption, having always leveraged new technologies and innovations aimed at improving efficiency, lowering cost of financing, and enhancing the speed of service. One of the earliest of these disruptions was credit scoring: a tool that allows financial institutions to evaluate individual consumers for their credit risk. Traditionally, credit scoring assessments relied on lending officials assigning a particular credit risk to a prospective customer based on criteria such as employment status, income, age, etc. This manual process not only resulted in administrative inefficiencies but also could be misused by lending officials who enjoyed wide discretionary powers in making these assessments.

The advent of new technologies facilitated the development of an automated process whereby a computer program takes information provided by the applicant as well as several outside sources and, using a complex set of weighted variables, produces a single number by which to rate the applicant's credit risk. This single number came to be known as a credit score – it is an objective score that could be justified, thereby overcoming the vagaries of subjective assessments. Simply put, a credit score is the "potential that a borrower or counterparty will fail to perform on an obligation". A low credit score implies a higher risk of default.

[1]A system that documents an individual's borrowings from different formal sector financial sources and credit history.

CONTRIBUTORS

IT for Change Team:

Author

Anuradha Ganapathy

Editorial guidance and review

Nandini Chami

Copyediting and proofreading

Sohel Sarkar and Sneha Bhagwat

Layout

Sreemoyee Mukherjee

This think piece is part of IT for Change's research project on Intelligent Infrastructures: For a People-Centered Digital Society, supported by IDRC (Canada) and the Fair, Green and Global Alliance (FGG).

Historically, credit scores were derived from historical financial data such as repayment history, defaults on prior credit transactions, and other forms of credit and payment-related behavior. Together, these are referred to as 'traditional' data. A FICO score,[2] for instance, principally looks at a consumer's payment history, credit history, the amount owed, and types of credit used. Similarly, a CIBIL score,[3] prepared by TransUnion CIBIL Ltd., considers payment history, extent of credit utilization, recently-sanctioned loans and credit cards, and a ratio of secured v/s unsecured loans. Over time, these automated scoring tools came to be regarded as better alternatives to the traditional subjective credit assessments that resulted in egregious forms of discrimination. Yet, these automated tools were not without exclusionary tendencies. Debates about financial inclusion pointed to their limitations in assessing the credit-worthiness of thin-file borrowers – individuals without sufficient financial payment history who remained invisible in formal lending systems. In such contexts, credit scoring algorithms that evaluate creditworthiness using 'non-traditional' or 'alternate' datasets have the potential to convert a large invisible, underserved population into formal consumers of credit. In the absence of a formally-accepted definition, 'alternate' data includes any or all data that falls outside the purview of 'traditional' data, with a precarious boundary between the two. Within this broad scope, alternate data may include mobile phone call records; periodic payments for non-loan products such as phone payments, rent, insurance, and utility bills; checking account transactions; data related to a consumer's educational and occupational history; consumer behavior; and data derived from a consumer's social media network. Lenders that use such forms of non-traditional data typically fall under the category of 'fintech' organizations.[4]

The Rise of Fintech Lending

The last few years have seen the rapid emergence of fintech firms offering credit scoring services using alternate datasets. Singapore-based Lenddo, which calls itself a market leader in alternative credit scoring, has developed an algorithm that aggregates data from a multitude of sources, from social media accounts – friends, frequency of interaction, interests, etc. – to other smartphone applications – messaging and browser history, apps, WiFi network use, battery levels, etc. – to establish a rating that signals an individual's likelihood of repaying or defaulting on loans. UK-based lender Credit Kudos incorporates biometrics and behavioral analytics into its algorithms. In China, credit scores are calculated by Alipay & WeChat based on data points such as purchase history, the types of phones used, the augmented reality games played, and friends on social media. These scores determine much more than financial access. They can influence people's chances of finding employment or a partner on a dating site. CASHe, an online lending platform in India that provides short-term personal loans, uses a proprietary predictive algorithm called Social Loan Quotient to create credit profiles of users using alternative data such as mobile phone and social media footprint, education, monthly salary, and career experience.

[2] A widely used automated credit-scoring system in the USA developed by the Fair and Isaac Corporation (FICO).

[3] A credit score computed by Credit Information Companies in India.

[4] Technology-enabled financial innovation that could result in new business models, applications, processes or products.

Traditional Data Sets	Alternate / Non-Traditional Data Sets
Credit history (prior repayment behavior, value and volume of prior loan)	Mobile (type of phone, location, contacts, battery life)
Payments (mortgage, credit cards)	Social media / internet (social media footprint, size of network, relationships, email, messaging habits, search history, browser used)
Asset ownership (land, property)	Payments (rentals, utilities)
Official demographic data (age, address, occupation income, court / police records)	E-commerce (shopping patterns, websites frequented, purchases made)
	Employment / occupation (type of work, employer name, designation, income, employment history)
	Behavior / personality traits (psychometric test scores)

A report published by the credit reporting company Experien suggests that, as of May 2019, 65 per cent of lenders in the United States were using information beyond the traditional credit report to make a lending decision.

Why are fintech firms that use credit scoring algorithms gaining popularity as financial infrastructures? As a 2016 report by the Omidyar Network notes, the burgeoning field of digital credit, also referred to as Big Data Small Credit (BDSC), hinges at the intersection of three important global trends: (1) an unprecedented rise in the use of mobile phones in developing countries and the digital footprints they generate, (2) rapidly rising processing power that matches smart algorithms with massive data streams, and (3) billions of global consumers who remain invisible to formal financial services.

What is important to note therefore is that the availability of large amounts of non-traditional data, driven by increased mobile phone usage, by itself was not enough to differentiate the credit scoring models of fintech lenders from those used by traditional lenders. Indeed, the computation of the FICO score was an algorithmic process which used a set of instructions to transform inputs from multiple data sets such as a history of late payments, a person's debt-to-credit limit ratio, and other elements into a single numerical value. The difference is that the algorithms of traditional lenders were built by human programmers who took a call on which data points were relevant in making credit determinations, and how much weight to give to each element.

The bigger impetus for the growth in using algorithmic credit scoring models came from the exponential growth of computing power and data technology that allowed the mining of information from a diverse set of structured variables, giving rise to 'learning algorithms', that is, algorithms that could virtually use any data about a person, analyze whether it corresponds to a characteristic of known-to-be-creditworthy people, and accordingly extrapolate a credit score. Learning algorithms (as opposed to the earlier set of hand-programmed algorithms, often referred to as 'dumb algorithms') catapulted the "all data is credit data"narrative into an intelligence layer in the lending process and expanded the world of credit data to include elements that were not clearly linked to creditworthiness. In other words, the transition from dumb algorithms to learning algorithms normalized the idea that it is possible to find proxies for creditworthiness in datasets that may not be directly linked to an individual's repayment capacity; data a priori was not even required to be a linked to creditworthiness any more, because the algorithm could 'make' it perform that function.

Crucially, as Matthew Bruckner points out, self-learning algorithms also effected a more fundamental shift in the objective of the credit scoring process. The first-generation algorithmic lenders, who created the dumb algorithms, were interested in designing a set of instructions their algorithm could execute to make better credit determinations than traditional credit scores. Designers of the new self-learning algorithms, on the other hand, were taking advantage of the exponential increases in processing power to design a set of instructions that would enable an algorithm to learn how to make creditworthiness determinations "on its own" by deciding which features were relevant and how to weigh them. The latter are often referred to as black box algorithms, where programmers can see what went in (vast amounts of data) and what came out (a credit determination) but NOT how or why the algorithm made a particular determination. The key difference therefore between the first and the second-generation algorithmic lenders is that in the latter, the normative lens that is applied in making the decision on who is "more" or "less" creditworthy was subsumed under a datafication process which allowed algorithms to make their own decisions based on the linkages in the data sets.

Fastidious supporters of learning algorithmic credit scoring models are invested in the promises of Big Data and its qualities, arguing that multiplying the number of variables will expand access to borrowers with thin credit files. A study of loan applications made by a fintech lender in India showed that evaluating creditworthiness based on social and mobile footprints could potentially extend financial access to those who have no formal credit scores, without severely impacting loan outcomes . The same study found that a model that includes mobile footprint, deep social footprint, and traditional credit score is better at predicting defaults than a model with only deep financial information and CIBIL score. Other reports found that as many as 1-in-4 of all minority applicants could transition from unscorable to scorable and can be eligible for reasonably priced credit when alternative data is used to make credit scoring assessments. Reports also suggest that algorithmic lenders could benefit from operational efficiencies such as reducing costs through enhanced under-writing capabilities, and enhanced speed in decision making through automation. These could translate into improved efficiencies for the credit market as a whole, where for example, more accurate credit underwriting can increase approval rates and decrease default rates.

More cautious observers regard algorithmic credit scoring as a double-edged sword, arguing that it is nothing more or less than a set of tools that can be applied to creating, refining, and scaling financial solutions for consumers. Its value to credit providers and consumers depends on how this tool is applied. Studies have advocated for digital footprints to complement rather than substitute credit bureau information, implying that lenders that use information from both sources – traditional credit bureaus and digital footprints – can make superior lending decisions compared to those that only access one of the two sources of information. A study on the statistical and economic significance of the use of call data records (CDRs) for credit scoring applications points to their enhanced contribution towards financial inclusion, especially when borrowers do not have any past behavior information that allows institutions to make a credit decision. Ethical use of such data, the same study recommends, would entail utilizing it in strictly positive terms, that is, not to deny credit to anyone, but to improve an individual's existing score. In some cases, the terms offered by algorithmic lending models were found to be an improvement over their traditional payday[5] lending counterparts, although the differences were not significant enough to consider them as safe or genuine alternatives to payday loans. There is also the concurrent recognition that the absence of adequate legal and regulatory frameworks for credit reporting poses a great challenge to the use of such models, particularly in the context of privacy violations, lack of transparency on data collection and usage practices, and the threat of predatory lending directed at vulnerable populations.

That said, this essay aims to go beyond the 'pros and cons' debate. We suggest that the starting point of regulating the use of machine learning algorithms in credit scoring is to see it not merely as a technological transformation underpinned by Big Data that brings in (or not) greater efficiency or objectivity, but as a fundamental political shift reflected in the willingness to dispense the task of shaping the normative framework of creditworthiness to a process of algorithmic datafication. In terms of the financial inclusion narrative that undergirds credit scoring algorithms, this means that the decision of who counts as creditworthy is based on what data points the algorithm counts as important. In the next section, we unpack three key issues that allow us to understand the political implications of this process. The first pertains to the shifting and unstable configurations of who counts as creditworthy. The second relates to the opacity of the algorithm as a feature of the model rather than a bug, meaning that what would be considered errors are likely built into its very design. The third relates to the logics of financial extraction that could underpin such models, and their implications for inclusion.

What Or Who Counts as Creditworthy?

In Technologies of Speculation, Sun-ha Hong writes, "The moral and political question is not simply whether datafication delivers better knowledge but how it transforms what counts in our society: what counts for one's guilt and innocence, as grounds for suspicion and surveillance, as standards for health and happiness." In the case of credit scoring algorithms, the single-minded focus on their sophistication and effectiveness in producing a reliable score with high predictive capabilities prevents us from asking the more fundamental question of what is being counted in the definition of good, and what is being made visible as bad or risky. Why is this important? Let us take the situation of Kevin Johnson who returned home one evening to a letter from American Express informing him that his credit score had been downgraded from $10,800 to $3,800 and that he was now classified as a risky customer because "other customers who had used their card at establishments where [Kevin] recently shopped have a poor repayment history with American Express". While this determination of credit scoring by association and not by an individual's own merit, is now made possible by technology, a key question to consider

[5]A *payday loan* is a type of short-term borrowing where a lender will extend high-interest credit based on one's income. (Wikipedia definition)

is, in the calculation of his credit score, why are the establishments where Kevin shops more important than his impeccable credit record and repayment history? In other words, why is a speculative insight sparked by an algorithm more privileged as a data point than factual realities?

In another instance, fintech start up EarlySalary rejected an online loan application made by a young woman despite her eligibility, because its machine learning algorithm used GPS, social media data, and bank account data to detect that she was taking the loan out for her live-in boyfriend who was unemployed, and who had earlier applied for the loan himself and been rejected. As Tyler Reigeluth points out, the onset of Big Data has not only compressed qualitative and epistemological differences between data, information, and knowledge, it has also made invasive data collection and monetization feel like an infrastructural norm.

Equally, as Barocas and Selbst[6] note "[T]he notion of creditworthiness is itself a function of the way the credit industry has constructed the credit issuing and repayment system – one in which an individual's capacity to repay a minimum amount of an outstanding debt on a monthly basis is taken to be a non-arbitrary standard by which to determine in advance and all at once whether the applicant is worthy of credit". However, the proliferation of digital technologies and the rise of Big Data encourages the naturalization of an algorithmically-derived credit score as endowed with inherently 'objective' qualities, one that is capable of 'speaking the truth', creating a duality that we must deal with. On the one hand, it completely erases the arbitrariness associated with the term creditworthiness. On the other, by allowing numerous data points unrelated to financial habits to determine creditworthiness, it fosters new forms of arbitrariness, but this time the veneer of quantification and objectivity precludes us from calling out its shifting configurations.

Let us take the example of Wonga, a UK payday lender that claims to fully automate credit decisions using 8,000 different data points to sort customers, deliver credit decisions within six minutes, and wire money directly to their accounts in fifteen minutes. The emphasis on speed makes Wonga dependent on data that is available more or less instantly, which means that the availability of data at a particular point of time is most crucial to the decision. Indeed, users studying Wonga have found that its algorithm often shows them different credit amounts depending on the Internet Browsers used for login. So, the data "leaked out" by the individual at a particular point of time was being used to build an identity which would then be linked to a creditworthiness score. As Hong points out, the claim to better, more objective knowledge through data depends on shifting expectations around what looks and sounds like reliable truth. Arguably, these shifting configurations undergirding the scoring process are subsumed under Wonga's claims of superior algorithmic capabilities. In 2018, a little less than a decade after its launch, the payday lender shut shop, admitting that its algorithm had been prompting it to lend money to people who could not pay it back, forcing Wonga to eventually write off loans of 330,000 customers, and waive off interest and fees for an additional 45,000.

Thus, what counts as creditworthy becomes a function of the ways in which the unstructured problem of 'defining' creditworthiness is parsed. Creditworthiness becomes a subjective, more often than not arbitrary, judgement that systematically disadvantages protected classes in the name of objectivity and neutrality.

[6] Cited in https://yjolt.org/credit-scoring-era-big-data

Essentially, the moral and the political aspects of what counts as better knowledge in a datafied society is perhaps more disruptive than the process of algorithmic calculation itself. We cannot ignore the fact that the claims to a higher-order intelligence, made possible by connecting vast amounts of data, go hand in hand with the increasing normalization of the idea that creditworthiness no longer needs to be datafied in terms of what is most meaningful or even linked to financial habits. Rather it appears to be a function of what aspects of our behaviors can be rationalized, combined, and organized to produce a new standard of truth that can be subsequently monetized for credit. This raises the important question of who benefits from what is counted as knowledge? Are credit scoring algorithms facilitating financial inclusion or reorganizing populations in terms of technological priorities with the objective of financial extraction?

The Algorithm Did It! Opacity as a Feature, Not Bug

We now move on from the issue of 'what' counts as knowledge to understanding the process of 'how' this decision is made. Studies on black box algorithms have shown that they can develop biases in ways that are not fully understood. As noted in the earlier section, in case of an error within a learning algorithm, a programmer usually cannot review the underlying instructions to find out why the error occurred and correct it. A 2013 study on Google Ads conducted by Professor Latanya Sweeney found that people with black-identifying names were "25 per cent more likely" to be shown ads suggestive of an arrest record as compared to people with white-identifying names, but it was unable to explain how or why Google's algorithm discriminated on race. In a datafied society, the connections between data, machines, and 'better knowledge' remain obscure for most of us, most of the time. As Andrew Tutt points out, "Our inability to understand, explain, or predict algorithmic errors is not only unsurprising, but destined to become commonplace. When and why machine-learning algorithms fail is difficult to predict and explain because what they do is probabilistic and emergent by design. What makes them valuable is what makes them uniquely hazardous."

This means that the way the algorithm learns to make decisions on creditworthiness is situated within a process that has numerous opportunities to pick up biases. Simply put, algorithms need training data to make decisions. But legacies of discrimination that are all-pervasive in society can "infect" such training data. Based on the available training data, the algorithm may draw a negative correlation between creditworthiness and certain pin codes or locations, for instance. This can be objective in that the algorithm may conclude that certain pin codes are highly prone to defaults, but also discriminatory because pin codes are linked to social geographies of race, caste, and class. In this way, the algorithm ends up perpetuating or even exacerbating existing biases. Essentially, all the proxies for financial footprints that are present in the training data such as social media presence, online browsing habits, shopping histories, number of connections, school or college attended, etc., are linked to key socio economic indicators. Even if the algorithm is specifically constructed so as to not consider protected categories such as race or gender, these will likely (and indirectly) become key considerations that impact the final score. For example, in response to David Heinemeier Hansson's now famous tweet asking Apple, and its underwriting partner Goldman Sachs, to explain why he and his wife were granted different credit limits despite their similar financial abilities, Goldman Sachs insisted that there is no gender bias in the algorithm because it does not use gender as an input. What this insistence fails to consider is that it is entirely possible for algorithms to discriminate based on gender, even when they are programmed to be 'blind' to that variable, because it is highly likely that they are drawing on inputs in the data set that correlate to gender. More crucially, imposing willful blindness to something as critical as gender only makes it harder for a company to detect, prevent, and reverse biases stemming from that variable.

The Apple card example points to another important consideration – the opacity of the algorithm does not merely affect access to credit, but also the terms on which such access is granted. For example, a lending platform that assesses creditworthiness based on the number of contacts stored on a person's smartphone would likely determine men to be more creditworthy in countries like India where men have greater social mobility (and likely more phone contacts) than women for socio-cultural reasons. Consequently, women (determined to be higher-risk individuals) would face higher interest rates on the same loan amount. Interestingly, the New York State Department of Financial Services' investigation into the Apple credit card matter did not find unlawful discrimination based on gender. Instead, it simply recommended that using credit scoring in its current form, and laws and regulations aimed at preventing lending discrimination are in need of strengthening and modernization to improve access to credit. By taking into account proxies for protected categories – for instance, the number of phone contacts as a proxy for gender – credit scoring algorithms, in effect, facilitate the same discriminations based on race, color, religion, national origin, sex, marital status, etc., that are prohibited by statutes such as the Equal Credit Opportunity Act in the US. Furthermore, the opacity of the algorithm gets a legal sanction on the premise that it did not 'intend' to discriminate on the basis of any protected characteristic. This despite the consensus that the way the algorithm is designed to operate makes it highly prone to (re)producing discriminatory outcomes and sustaining existing power asymmetries. In other words, existing legal approaches lack a robust epistemic framework to deal with issues of algorithmic discrimination, soft-balling the issue as unintentional rather than illegal.

It is important to note that there is no neat technical solution that can 'correct' these so-called bugs such that the algorithm can perform better in the next round. Framing the misrepresentations or exclusions produced by the algorithm as errors or unintentional consequences is thus misplaced. Apart from legitimizing harms as regrettable accidents that could not have been predicted and therefore cannot be anyone's specific responsibility, such framings also place technology in a realm beyond ethical reasoning.

What we need, therefore, is not a technical fix, but a more critical engagement with the institutional, historical, and societal legacies that pervade the deployment of scoring systems, along with scrutinizing their need in society. For example, how important is credit scoring to achieving financial inclusion if its assessments enable credit access to men who were hitherto credit invisible, but continue to discriminate against women? Furthermore, how do we account for the fact that merely becoming credit visible does not guarantee fair and equal treatment since the terms of visibility are made non-negotiable by the opacity of the algorithm? How important are the principles of efficiency and cost optimization if they override transparency and further exclusions? These are issues of power and politics, not of data fixes and technical solutions.

Inclusion through Extraction?

Referring to credit scoring algorithms as calculative infrastructures, Rob Aitkens writes that the process by which the credit invisibles are made visible as financial subjects capable of being inserted into financial value chains is by no means a form of seamless inclusion. Rather, these algorithms serve to segment and sort credit seekers in ways that allow credit providers to expand their loan pool while minimizing their risk exposure and lowering transaction costs. Arguably, there is no conclusive evidence that alternative credit scoring companies are using machine learning tools to maximize lender profitability at the expense of consumers rather than scoring for creditworthiness, but equally, there is no reason to believe that these companies have borrowers' interests at heart. A study on online payday lending notes that "lenders [using] sophisticated technology and advanced algorithms to predict which applicants are most likely to repay loans continue to charge interest rates ranging from 300 to 700 percent APR (annual percentage rate)". The data surplus generated by credit providers can also be used for price discrimination. Once their dominant position in data is established, digital monopolies may use the data not only to assess a potential borrower's creditworthiness or riskiness but also to identify the highest rate a given borrower would be willing to pay for a loan, resulting in a form of personalized pricing.

Ultimately, a combination of careful segmentation of the pool in ways that avoid risk, and algorithmic price discrimination that parses the credit visible into smaller categories becomes a method to deepen the financial value of loan portfolios held by lenders. Experts have also detailed how online 'lead generators' are using sophisticated algorithmic scoring techniques to zero in on consumers at precise moments when they are likely to be especially vulnerable to low-value, short-term credit products with usurious interest rates and highly unfavorable terms. In the language of economics, when consumer preferences are thus manipulated, the overestimation of benefits from the product or service can cause some consumers to purchase the product, even though its actual value to them is lower than the price, leading to a greater welfare loss than under price discrimination. Julia Angwin has written about financial manipulation practices that involve credit rating agencies abusing personal data by selling lists of people who were late in paying most of their mortgage bills to unscrupulous marketers who pitch fraudulent products. This underscores the possibility that certain alternative credit scorers may not be interested in predicting consumer creditworthiness, but rather in finding vulnerable, high-value targets for high-cost loan products. According to data cited by the Reserve Bank of India (RBI) in its annual Financial Stability Report 2021, fintech consumer credit has seen the highest increases in late repayments or defaults in the one year period from September 2020 to September 2021, raising concerns about predatory lending practices in an unregulated market that offers little protection to consumers.

Experts also caution against overestimating the potential efficiency gains due to algorithmic credit scoring. While there may be empirical data to corroborate the hypothesis that algorithmic credit scoring could improve access to credit for thin-file and no-file borrowers, they suggest that these have to be weighed against the larger exploitative practices generated by such algorithms. These practices could give rise to new inefficiencies such as the use of data-driven insights by lenders to "skim the most creditworthy segment of the market for themselves", or over-investment in gathering private information in a way that is socially wasteful (rent seeking), or using borrower insights to exploit cognitive and behavioral biases rather than for improving design and marketing of credit products.

Additionally, there is some evidence to suggest that credit access may not have expanded to include new and underserved borrower groups. Rather, alternate credit scorers may be employing target variables that work to the detriment of historically-disadvantaged groups. In the aftermath of Wonga's collapse, an anonymous person "with direct knowledge of the figures" was found saying that the well-publicized low default rates were calculated on a per-loan basis, but "significantly more than half of customers eventually failed to repay, spiraling into debt as they took out new loans to pay off the earlier ones". This seems to support Aitkens' claim that placing credit invisible populations in ways that allow them to be easily coded in relation to their networks, cultures, locations, habits, etc., may actually facilitate a form of adverse incorporation. At least in some contexts, the credit invisible may simply be converted into the credit trapped.

Towards Alternate Imaginaries of Algorithmic Credit Scoring

Against this backdrop, to start imagining alternate forms of algorithmic credit scoring is to understand that the process of determining creditworthiness cannot be reduced to a technical exercise performed by a code. Rather, it is a normative process driven by values about what constitutes creditworthiness, who sets these terms, and who benefits from being made visible on these terms. The use of credit scoring infrastructures needs to foreground concerns of structural inequalities and social injustices and ask a more fundamental question – how fair is it that certain sections of society who have access to other sources of credit do not have to subject themselves to such forms of datafication, leaving resource poor groups and marginalized communities to disproportionately bear this burden? Importantly, the ethics of legitimizing the role of technology in defining creditworthiness and arbitrating key development goals such as financial inclusion and economic empowerment, when it is not explicitly driven by such objectives, need to be critically examined. The normative task of constructing creditworthiness must center the voices of people it is meant to serve and prioritize their protection.

Secondly, we need to decenter the notion of objectivity from this discussion and, instead, privilege values of explainability and accountability. If anything, credit scoring algorithms tell us that perfectly objective processes can also be perfectly biased and discriminatory, and a willful blindness or neutrality to key variables in the decision-making process may actually have detrimental effects on society. We know now that if human scoring models have historically been discriminatory, black box algorithmic models reproduce such forms of discrimination, albeit in more 'sophisticated' ways that are not unconstitutional or illegal. In other words, if humans discriminate, algorithms discriminate better. Technical fixes such as perfecting the algorithm by adding more data or tools for debiasing data sets are, therefore, inadequate responses because they do not solve the problem of opacity. We need to equip our legal institutional mechanisms with vocabulary that moves beyond benign labels of 'unintentional bias' or 'unintended consequences' and develop standards that hold such black box systems accountable to higher levels of transparency and explainability, using human rights and social justice principles.

Finally, the principles of financialization and profit maximization that are at the center of decisions on creditworthiness often cause more harm than good. They can likely encourage a form of adverse incorporation, or even foster new forms of marginalization. In contrast, an intelligent credit scoring system is one that pays adequate attention to the ways in which power asymmetry between the lender and the borrower can be reduced, such that the terms of credit do not leave the borrower worse off. Regulation can enforce this by protecting consumers from the individual and collective harms of algorithmic profiling in predatory credit markets, strengthening recourse options, and introducing more robust measures for governing data practices of firms deploying credit scoring infrastructures.

IT for Change
Fair.Green.Global. Dialogue and dissent
IDRC · CRDI
Canada