The Equitable Financial Inclusion Score (EFIS): A Microeconomic Framework for Credit Assessment in Underserved Populations

Kofi Hair-Ralston* Sarthak Pal Mahajan* Kshiteej Abhay Herode

This version: October 2026

Abstract

Many people with steady but informal livelihoods cannot borrow from formal lenders because they have no credit file. The difficulty is less that they are poor risks than that lenders cannot tell them apart from poor risks. We propose the Equitable Financial Inclusion Score (EFIS), a credit score built from four kinds of evidence that such borrowers often do have: a record of income, relationships with people who will vouch for them, a history of repaying small loans, and observable saving and insurance behavior. For each component we write down a simple model, state the assumptions under which it measures what it is meant to measure, and derive its main properties. We show how the component weights can be chosen to trade predictive accuracy against a fairness penalty, and we give conditions under which that choice is unique. We also identify where the components can be gamed and propose corrections. The framework has not yet been tested on loan data, and we describe the evidence that would be needed.

JEL Codes: D82, G21, O16, C58.

Keywords: financial inclusion, credit scoring, asymmetric information, signal extraction, social networks.

*Equal contribution. The order of the first two authors was determined by a coin flip.

Hair-Ralston: kofibhairralston@gmail.com. Mahajan: sarthak@klinn.works. Herode: herodekshiteej@gmail.com. An earlier version of this paper was posted on SSRN (abstract 5519538). We welcome comments. All errors are our own.

1 Introduction

Consider a woman who has run the same fruit stall in Lagos for seven years. She knows her customers, buys from the same suppliers, and puts money aside most weeks. When she asks a bank for a loan to expand, she is turned down, not because the bank has evidence that she is a poor risk but because it has no evidence at all. She has no credit file, no payslips, and no collateral the bank recognizes. About 1.4 billion adults worldwide had no account at a bank or mobile money provider in 2021 (Demirgüç-Kunt et al., 2022), and many more have accounts but no borrowing history that a lender can score.

The economics of her situation are old. When a lender cannot distinguish safe borrowers from risky ones, it must offer everyone the same terms, and those terms reflect the average risk of whoever applies (Akerlof, 1970; Stiglitz and Weiss, 1981). Safe borrowers then face interest rates that are too high for them, some of them stop applying, and the average quality of the pool falls further. A lender in this position may ration credit outright rather than raise rates. The borrowers hurt most are the creditworthy ones the lender cannot see.

The usual remedies each cover part of the problem. Collateral works for people who own assets. Group lending with joint liability uses the fact that neighbors know more about each other than a bank does (Besley and Coate, 1995; Ghatak and Guinnane, 1999), but it depends on particular social settings and places real costs on group members. Microcredit has reached many millions of people, although randomized evaluations find modest average effects on household income (Banerjee et al., 2015). More recently, lenders have begun to score applicants from mobile phone records and other digital traces (Björkegren and Grissen, 2020). These scores can predict repayment, but they are hard to explain to the borrower, and they may reproduce existing disparities if the inputs are correlated with group membership (Fuster et al., 2022).

This paper takes a different approach. Rather than search for a single new data source, we ask what a lender would want to know about a borrower and look for evidence of each part. We separate creditworthiness into four components. The first is earning capacity: how much the borrower can reliably earn, as opposed to what she happened to earn last month. The second is social capital: whether people whose own reliability is known are willing to vouch for her. The third is a repayment record built up through small loans. The fourth is financial behavior, measured by how she saves and whether she insures against losses. Each component is computed from data that informal borrowers often do have, and the overall score is a weighted combination of the four.

For each component we write down a small model and ask what it implies. Two of the results are corrections to the earlier version of this paper. In the earning-capacity model, we show that the score responds more slowly to new income when income reports are noisy and more quickly when earning capacity genuinely changes, which is the opposite of what our earlier draft stated (Section 3). In the repayment-record component, we show that the formula in our earlier draft rewarded borrowers for splitting one loan into many small ones, and we propose a version that removes most of that incentive (Section 5). We also show when the weights on the four components are uniquely determined, which requires replacing the fairness penalty we used before with one

that is convex (Section 7).

We want to be clear about what the paper does not do. We have not estimated the model on loan data, and nothing here shows that EFIS predicts repayment better than existing methods. Several claims in our earlier version, including that EFIS improves welfare relative to every alternative and that it makes all market participants better off, do not follow from the model and in at least one case are false in general; we discuss these in Section 8. The contribution is a framework that can be taken to data, with each piece specified clearly enough that it can be tested and, if it fails, replaced.

The rest of the paper is organized as follows. Section 2 sets out how the four components combine into one score. Sections 3 through 6 develop each component. Section 7 treats the choice of weights. Section 8 discusses what the framework does and does not guarantee. Section 9 covers rollout, data verification, privacy, and the use of the score in lending. Section 10 concludes. Proofs are in Appendix A.

2 The Combined Score

Let i index borrowers and t index periods. For each borrower we compute four component scores, $H_i(t)$ for earning capacity, $S_i(t)$ for social capital, $R_i(t)$ for the repayment record, and $B_i(t)$ for financial behavior. Collect them in the vector $C_i(t) = (H_i(t), S_i(t), R_i(t), B_i(t))^\top$. The EFIS score is

$$U_i(t) = \Lambda\big(w_0 + w^\top C_i(t)\big), \qquad \Lambda(x) := \frac{1}{1+e^{-x}}, \tag{2.1}$$

where $w = (w_H, w_S, w_R, w_B)^\top$ are weights and w_0 is an intercept. Because Λ is the logistic function, $U_i(t)$ lies strictly between zero and one and can be read as a predicted probability of repayment once the weights have been estimated on repayment outcomes (Section 7).[1]

The case for using four separate components rests on an assumption that we state openly rather than try to prove.

Assumption 2.1 (Non-redundancy)**.** For each component k, the component carries information about repayment that the other three do not. Formally, if Y_i indicates repayment, then $\mathbb{P}(Y_i = 1 \mid C_i) \neq \mathbb{P}(Y_i = 1 \mid C_{-k,i})$ on a set of borrowers with positive probability, where $C_{-k,i}$ is C_i with component k removed.

Our earlier draft presented this condition, and a related claim that the estimation errors of the four components are asymptotically independent, as theorems. Neither follows from the model. The components are built from different data, but they describe the same person, and they are linked mechanically: the social capital score uses other borrowers' EFIS scores, which in turn contain their own social capital scores. Whether each component adds information is an empirical

[1] Our earlier draft wrote the score both as $\sum_k w_k C_k$ and as $\Lambda(\sum_k w_k C_k)$ and omitted the intercept. We use the second form throughout, with an intercept, so that the score is a probability. Note that the logistic form does not by itself allow the components to interact; each component shifts the log-odds additively. Interactions would require adding product terms.

question. It can be checked directly by testing whether the estimated weight on each component is different from zero, or by comparing out-of-sample predictive accuracy with and without it.

3 Earning Capacity

Informal income is volatile. A lender looking at a month of low sales cannot tell whether the business is failing or the month was slow. Standard scoring treats that volatility as risk. The earning-capacity component instead treats each month's income as a noisy reading of an underlying capacity to earn, and estimates that capacity.

The model. Let $Y_{i,t}$ be borrower i's observed income in period t and $\theta_{i,t}$ her unobserved earning capacity. We assume

$$\theta_{i,t} = \phi_i \, \theta_{i,t-1} + \mu_i + \sigma_\theta \, \varepsilon_{i,t}, \tag{3.1}$$

$$Y_{i,t} = \theta_{i,t} + \sigma_y \, \nu_{i,t}, \tag{3.2}$$

where $\varepsilon_{i,t}$ and $\nu_{i,t}$ are independent standard normal shocks. The parameter ϕ_i measures how persistent earning capacity is. When $|\phi_i| < 1$, capacity fluctuates around a long-run level of $\mu_i/(1-\phi_i)$; when $\phi_i = 1$, capacity follows a random walk that drifts by μ_i per period.[2] The term $\sigma_\theta \varepsilon_{i,t}$ is a genuine change in capacity, such as losing a supplier, while $\sigma_y \nu_{i,t}$ is noise in the income record, such as a sale recorded in the wrong week.

The filter. Given the parameters, the best estimate of $\theta_{i,t}$ from incomes observed up to t is given by the Kalman filter (Kalman, 1960; Hamilton, 1994). Write $\hat{\theta}_{i,t|t}$ for this estimate and $P_{i,t|t}$ for its variance. Each period, the filter first predicts and then updates:

$$\hat{\theta}_{i,t|t-1} = \phi_i \, \hat{\theta}_{i,t-1|t-1} + \mu_i, \qquad P_{i,t|t-1} = \phi_i^2 \, P_{i,t-1|t-1} + \sigma_\theta^2, \tag{3.3}$$

$$K_{i,t} = \frac{P_{i,t|t-1}}{P_{i,t|t-1} + \sigma_y^2}, \qquad \hat{\theta}_{i,t|t} = \hat{\theta}_{i,t|t-1} + K_{i,t}\big(Y_{i,t} - \hat{\theta}_{i,t|t-1}\big), \tag{3.4}$$

with $P_{i,t|t} = (1 - K_{i,t}) \, P_{i,t|t-1}$. In words, the estimate moves toward each new income observation by a fraction $K_{i,t}$ of the surprise. Under normality, $\hat{\theta}_{i,t|t}$ is the conditional mean of capacity given the income history; without normality it is still the best estimate among those that are linear in past income.

How the filter weighs new income. The gain $K_{i,t}$ settles down quickly to a constant that depends on only two things, the persistence ϕ_i and the ratio of genuine volatility to measurement noise, $q := \sigma_\theta^2/\sigma_y^2$.

Proposition 3.1 (Steady-state gain)**.** *Suppose $\sigma_y > 0$, $\sigma_\theta > 0$, and $P_{i,0|0} \geq 0$. Then $P_{i,t|t}$ converges to a unique positive limit P^*, and the gain converges to $K^* = P^*/\sigma_y^2$. The limiting gain depends on the parameters only through ϕ_i and q, and it is strictly increasing in q.*

[2] Our earlier draft called μ_i "individual-specific growth." That reading is correct only when $\phi_i = 1$. When $|\phi_i| < 1$, μ_i sets the long-run level, and the model has no trend.

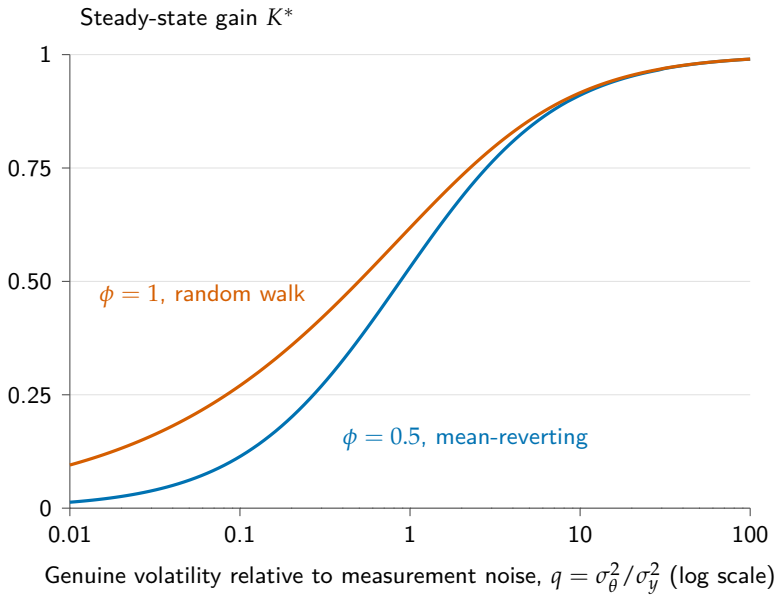


FIGURE 1. How much the earning-capacity estimate reacts to a new income observation. The vertical axis is the steady-state Kalman gain K^* from Proposition 3.1, the fraction of an income surprise that is passed through to estimated capacity. Each line fixes the persistence of earning capacity, ϕ; values of ϕ between the two lie between them. Computed in closed form from equation (A.2) in Appendix A; no data are used.

The proof is in Appendix A. Figure 1 plots K^* against q. When income reports are noisy relative to genuine changes in capacity (small q), the gain is small and the estimate barely reacts to any single month; it relies mainly on the borrower's track record. When genuine changes dominate (large q), the gain approaches one and the estimate follows recent income closely. This is the property that makes the filter useful for informal earners. A borrower whose income swings from month to month because of noisy records, and not because her business is changing, will have a stable estimated capacity. Our earlier draft stated the reverse, that the filter puts more weight on the long-run average for "genuinely volatile earners"; Proposition 3.1 shows the opposite is true.

From capacity to a score. The earning-capacity component combines estimated capacity with verified credentials:

$$H_i(t) = a_1 \frac{\hat{\theta}_{i,t|t}}{\bar{\theta}} + a_2 \sum_{j=1}^{n} \delta_j \mathbf{1}\{i \text{ holds credential } j\}\,, \tag{3.5}$$

where $\bar{\theta}$ is average estimated capacity in the population, the credentials are items such as trade licenses or completed training, and the weights δ_j are estimated from data. In practice the parameters ϕ_i, μ_i, σ_θ, and σ_y are not known. With short income histories, it is more realistic to estimate them for groups of similar borrowers (for example, by trade and region) by maximum likelihood and to treat them as known for each individual.

4 Social Capital

Neighbors, suppliers, and customers often know a great deal about whether someone is reliable. Joint-liability lending puts that knowledge to use by making neighbors responsible for one another's loans (Besley and Coate, 1995); there is also direct evidence that ties in a social network can serve as a kind of collateral (Karlan et al., 2009). The social capital component uses the same knowledge in a lighter form: a borrower earns credit when people with good records are willing to vouch for her.

Endorsements. Let $G = (V, E)$ be a network of borrowers in which each relationship (i, j) has a strength $w_{ij} \in [0, 1]$. A borrower j may endorse i. Let $\mathcal{E}_i(t)$ be the set of borrowers who endorse i at time t and who are within D steps of i in the network. The base social capital score is

$$\tilde{S}_i(t) = \sum_{j \in \mathcal{E}_i(t)} \frac{\omega_{ij}\, U_j(t-1)}{(1 + d_{ij})^{\gamma}}, \tag{4.1}$$

where d_{ij} is the number of steps between i and j, $\gamma > 0$ controls how fast endorsements lose value with distance, $U_j(t-1)$ is the endorser's own EFIS score last period, and ω_{ij} is the strength of the strongest path between them, defined as the largest product of relationship strengths along any path from i to j. The formula gives more weight to endorsers who have good records themselves, to endorsers who are closely tied to i, and to endorsers who are near i in the network.[3]

Using last period's scores of endorsers avoids a circular definition, since $U_j(t-1)$ is already known when $\tilde{S}_i(t)$ is computed. Distance decay alone does not prevent fake relationships, because a fake relationship is usually a direct one. That protection has to come from verifying the relationships themselves. We measure the strength of each relationship as the average of three separately verified measures,

$$w_{ij} = \tfrac{1}{3}\big(w_{ij}^{(1)} + w_{ij}^{(2)} + w_{ij}^{(3)}\big), \tag{4.2}$$

taken from communication records, shared membership in community organizations, and attestations from local leaders, with each measure in $[0, 1]$. A relationship that is confirmed by only one source therefore receives at most a third of full strength.[4]

Concentration. A borrower whose endorsements all come from one relationship is more exposed to that relationship going wrong. We therefore discount the score by the concentration of her direct ties,

$$S_i(t) = \tilde{S}_i(t)\,\big(1 - \zeta\ \mathrm{HHI}_i\big), \qquad \mathrm{HHI}_i = \sum_{j \in N(i)} \left(\frac{w_{ij}}{\sum_{k \in N(i)} w_{ik}} \right)^2, \tag{4.3}$$

where $N(i)$ is the set of i's direct neighbors and $\zeta \in [0, 1)$. The Herfindahl index HHI_i lies between $1/|N(i)|$ and one, so the restriction $\zeta < 1$ keeps the score from turning negative. One cost of this

[3] In our earlier draft the sum in (4.1) ran only over i's direct neighbors. Every direct neighbor is one step away, so the distance term was the same constant for everyone and had no effect. Allowing endorsements from up to D steps away gives the distance term a role.

[4] The earlier draft wrapped this average in $\min\{\cdot, 1\}$. Since each of the three measures is at most one, so is their average, and the cap never binds.

adjustment is that it falls hardest on borrowers with very few ties, who are often the newest users; in a rollout, ζ should be kept small until networks have had time to form.

5 The Repayment Record

A new borrower has no repayment history, and without one she cannot get the loans that would create a history. The repayment-record component breaks this circle by issuing a record of each successfully repaid loan, starting with small loans, so that a borrower can build a history from nothing.

When does a record attract the right borrowers? Suppose a borrower who repays with probability π is offered a loan of size L at interest rate r. Let $V(L)$ be the value she gets from using the loan, and let $\rho(L)$ be the value of the record she receives if she repays, equal to the discounted expected increase in the loans she can obtain later. Following our earlier draft, suppose that the value of the loan and the record are realized only if her project succeeds, but that the obligation to repay is owed in any case, for example because it can be collected from other income. Her expected payoff from taking the loan is then

$$\pi \left[V(L) + \rho(L) \right] - (1+r)L. \tag{5.1}$$

Proposition 5.1 (Who takes small loans)**.** *Under payoff* (5.1)*, a borrower takes the loan if and only if* $\pi \geq \bar{\pi} := (1+r)L/[V(L)+\rho(L)]$*. Increasing the value of the record lowers the threshold* $\bar{\pi}$*.*

The proof is immediate, since (5.1) is increasing in π. The proposition says two things. First, small loans do sort borrowers: those who expect to repay take them, and those who do not, stay away. Second, making the record more valuable widens the pool rather than narrowing it, because it raises the payoff from borrowing for everyone. A more valuable record is therefore a way of drawing in more good borrowers, not a way of screening more strictly. The sorting also depends on the repayment obligation being owed whatever happens to the project. If borrowers repay only when the project succeeds, as under limited liability, the payoff becomes $\pi[V(L) - (1+r)L + \rho(L)]$, the decision no longer depends on π, and the sorting disappears. This is the standard adverse-selection problem of Stiglitz and Weiss (1981). In that case the lender has to rely on the other three components to separate borrowers.[5]

Turning repayments into a score. Our earlier draft proposed

$$R_i^{\text{old}}(t) = \sum_{k=1}^{m} \lambda^{s_k} \tanh\!\Big(\frac{L_k}{c}\Big) \mathbf{1}\{\text{loan } k \text{ repaid}\}\,, \tag{5.2}$$

where L_k is the size of loan k, s_k is the length of the run of consecutive repayments ending with loan k, $\lambda > 1$ rewards long runs, and c scales loan sizes. The intent was that tanh would keep very large loans from dominating. That is true loan by loan, but it has two consequences that the draft

[5]The separation condition in our earlier draft, $\pi_{\text{high}}\rho(L) > (1+r)L - \pi_{\text{high}}V(L) > \pi_{\text{low}}\rho(L)$, does not follow from (5.1). The correct condition is $\pi_{\text{high}}[V(L)+\rho(L)] \geq (1+r)L > \pi_{\text{low}}[V(L)+\rho(L)]$.

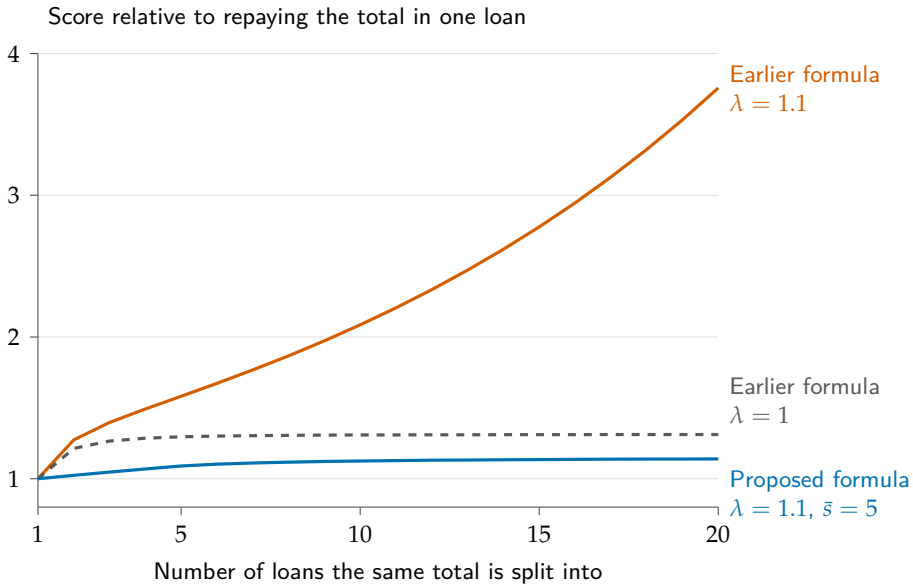


FIGURE 2. The earlier repayment score rewarded splitting one loan into many. A borrower repays a fixed total amount, equal to c, in n equal loans with no missed payments, so the kth loan ends a run of length k. The vertical axis is her repayment score divided by the score she would get from repaying the same total in a single loan. The earlier formula is equation (5.2); the proposed formula is equation (5.3). The dashed line isolates the effect of applying tanh loan by loan, with no run bonus.

overlooked. Because λ^{s_k} grows without limit, the score is not bounded: with $\lambda = 1.1$, the thirtieth consecutive repayment counts about 16 times as much as the first. And because tanh is concave, splitting a loan into smaller pieces always raises the score, since $n \tanh(x/n) > \tanh(x)$ for $n > 1$ and $x > 0$. Splitting also lengthens the run of repayments, which raises the score further. The formula therefore rewards exactly the gaming it was meant to prevent.

We propose instead

$$R_i(t) = \tanh\left(\frac{1}{c}\sum_{k=1}^{m} \lambda^{\min\{s_k,\bar{s}\}}\, L_k\, \mathbf{1}\{\text{loan } k \text{ repaid}\}\right), \tag{5.3}$$

which applies the concave transformation to the total amount repaid rather than to each loan, and caps the run bonus at $\bar{s}$ consecutive repayments. Under (5.3), repaying a total amount in more pieces can raise the score only through the run bonus, and by at most a factor of $\lambda^{\bar{s}-1}$ inside the tanh. The score still increases with every repayment, still rewards consistency, and now lies between zero and one.

Figure 2 shows the difference. Under the earlier formula, a borrower who repays a given total in 20 small loans instead of one ends up with a score almost four times higher, and the advantage keeps growing with the number of loans. Even with no run bonus at all, applying tanh loan by loan gives splitting an advantage of 31 percent in this example. Under the proposed formula, the advantage never exceeds 16 percent.

6 Financial Behavior

The last component tries to measure two traits that bear on whether a borrower will repay: patience and caution. A patient borrower values future access to credit and so has more reason to protect it. A cautious borrower is more likely to protect herself against shocks that would make repayment impossible. We infer both traits from what borrowers actually do with their money.

The model. Suppose borrower i chooses consumption $c_{i,t}$ to maximize

$$\mathbb{E}\sum_{t=0}^{\infty}\beta_i^t\left[\frac{c_{i,t}^{1-\sigma}}{1-\sigma}-\eta_i\,\frac{\ell_{i,t}^{1+\alpha}}{1+\alpha}\right], \tag{6.1}$$

where β_i is her discount factor (higher means more patient), $\sigma>0$ governs how much she dislikes swings in consumption, $\ell_{i,t}\geq 0$ is a loss she suffers, such as damaged stock, and η_i measures how much she dislikes such losses.[6]

Patience from consumption growth. If she can save and borrow at rate r, her choices satisfy the Euler equation $c_{i,t}^{-\sigma}=\beta_i(1+r)\,\mathbb{E}_t[c_{i,t+1}^{-\sigma}]$. Suppose that consumption growth $g_{i,t+1}\coloneqq\log(c_{i,t+1}/c_{i,t})$ is normal with mean $\bar{g}_i$ and variance s_i^2. Then $\mathbb{E}_t[(c_{i,t+1}/c_{i,t})^{-\sigma}]=\exp(-\sigma\bar{g}_i+\sigma^2 s_i^2/2)$, and taking logs of the Euler equation gives

$$\bar{g}_i=\frac{1}{\sigma}\big[\log\beta_i+\log(1+r)\big]+\frac{\sigma}{2}\,s_i^2,\qquad\text{so}\qquad\hat{\beta}_i=\exp\Big(\sigma\,\hat{g}_i-\log(1+r)-\tfrac{\sigma^2}{2}\,\hat{s}_i^2\Big), \tag{6.2}$$

where $\hat{g}_i$ and $\hat{s}_i^2$ are the sample mean and variance of her consumption growth (Hansen and Singleton, 1983). Faster planned consumption growth, which requires saving now, reveals greater patience.[7] Two cautions apply. The formula needs a value of σ, which cannot be learned from a single person's data when everyone faces the same interest rate, so in practice σ has to be taken from the literature or estimated across borrowers. And many informal borrowers cannot freely save and borrow at a single rate, which is the premise of the paper; where the Euler equation fails, $\hat{\beta}_i$ should be read as a summary of saving behavior rather than a deep preference parameter.

Caution from insurance purchases. Suppose borrower i, with resources c_i, can buy insurance against a loss ℓ for premium p. If she buys it, she consumes c_i-p and bears no loss; if not, she consumes c_i and faces the loss. Under (6.1), buying the insurance reveals that

$$\eta_i\;\geq\;(1+\alpha)\,\frac{u(c_i)-u(c_i-p)}{\mathbb{E}\big[\ell^{1+\alpha}\big]}\;\approx\;(1+\alpha)\,\frac{p\,c_i^{-\sigma}}{\mathbb{E}\big[\ell^{1+\alpha}\big]}, \tag{6.3}$$

where $u(c)=c^{1-\sigma}/(1-\sigma)$ and the approximation holds when the premium is small relative to

[6] Our earlier draft called η_i "loss aversion" and used γ for it. In behavioral economics, loss aversion means that losses weigh more than equal-sized gains relative to a reference point (Kahneman and Tversky, 1979). The parameter here is simply the weight on a separate cost of losses, so we call it loss sensitivity. We also renamed the curvature parameter from ρ to σ to avoid a clash with the record value $\rho(L)$.

[7] Our earlier draft had $\frac{1}{2}s_i^2$ where $\frac{\sigma}{2}s_i^2$ belongs in the first expression, and the corresponding error in the second. The draft also described $\hat{\beta}_i$ as coming from the savings rate; the formula uses consumption growth, which can be built from savings records together with income.

resources. Declining the insurance gives the reverse inequality. Each purchase decision therefore bounds η_i from one side, and several decisions narrow it down.[8]

The component. We combine these estimates as

$$B_i(t) = b_1\,\Lambda\big(\kappa(\hat{\beta}_i - \bar{\beta})\big) + b_2\tanh\Big(\frac{\hat{\eta}_i}{\bar{\eta}}\Big) + b_3\,F_i, \qquad F_i = \sum_{j=1}^{J} v_j\,\mathbf{1}\{\text{module } j \text{ completed}\}\ \text{score}_{ij}, \tag{6.4}$$

where $\bar{\beta}$ and $\bar{\eta}$ are population averages, $\kappa > 0$ sets how sharply the score responds to patience, and F_i records participation and performance in financial education modules with weights v_j. All three terms increase in the corresponding trait. The second term rewards more loss sensitivity, not an intermediate "appropriate" level as our earlier draft suggested; rewarding an intermediate level would require a function that peaks at some target value.[9]

7 Choosing the Weights

The weights in (2.1) are estimated from repayment outcomes, subject to two penalties: one that limits how much the score tracks protected characteristics, and one that keeps any single component from dominating.

Let $y_i \in \{0,1\}$ indicate whether borrower i repaid, for $i = 1,\dots,n$, and let $x_i = (1, C_i^\top)^\top$ include a constant. Write $\boldsymbol{\vartheta} = (w_0, w^\top)^\top$ and $p_i(\boldsymbol{\vartheta}) = \Lambda(x_i^\top \boldsymbol{\vartheta})$. The log-likelihood is

$$\mathcal{L}(\boldsymbol{\vartheta}) = \sum_{i=1}^{n}\big[y_i \log p_i(\boldsymbol{\vartheta}) + (1-y_i)\log\big(1 - p_i(\boldsymbol{\vartheta})\big)\big]. \tag{7.1}$$

Fairness penalty. Let A_{ij} indicate whether borrower i belongs to protected group j, for $j = 1,\dots,m$. Our earlier draft penalized the squared correlation between the score index $C_i^\top w$ and each $A_{\cdot j}$. A correlation does not change when w is multiplied by a positive number, so that penalty is constant along every ray from the origin and is not convex, and the problem can have several solutions. We penalize the sample covariance instead:

$$\Omega(w) = \sum_{j=1}^{m} \widehat{\mathrm{Cov}}\big(C_i^\top w,\, A_{ij}\big)^2 = \sum_{j=1}^{m}\big(w^\top \kappa_j\big)^2, \qquad \kappa_j \coloneqq \frac{1}{n}\sum_{i=1}^{n}(C_i - \bar{C})(A_{ij} - \bar{A}_j), \tag{7.2}$$

which is a convex quadratic in w.

Balance penalty. To keep the score explainable, we penalize departures from equal weights, $\Phi(w) = \sum_k (w_k - \frac{1}{4})^2$. This penalty is only meaningful if the four components are measured on

[8] The bound in our earlier draft had $\mathbb{E}[\ell] - p$ in a denominator. Insurance premiums usually exceed expected losses, so that expression is typically negative and the bound has no content. Inequality (6.3) follows directly from comparing utility with and without the policy.

[9] The first term in the earlier draft was garbled in typesetting and read as $1/\log(1 + e^{-\kappa(\hat{\beta}_i - \bar{\beta})})$, which grows without limit as patience rises. We use the logistic function, which we believe was intended.

comparable scales, so each should be standardized before estimation. The weights then solve

$$\hat{\boldsymbol{\vartheta}} = \arg\max_{\boldsymbol{\vartheta}} \; J(\boldsymbol{\vartheta}) := \mathcal{L}(\boldsymbol{\vartheta}) - \lambda_1 \, \Omega(w) - \lambda_2 \, \Phi(w), \tag{7.3}$$

with $\lambda_1, \lambda_2 \geq 0$ chosen by the analyst. The first-order conditions are $\sum_i (y_i - p_i) x_{ik} = \lambda_1 \, \partial\Omega/\partial\vartheta_k + \lambda_2 \, \partial\Phi/\partial\vartheta_k$ for each coefficient.

Proposition 7.1 (Existence and uniqueness of the weights)**.** *Suppose the $n \times 5$ matrix with rows $x_i^\top$ has full column rank, that both outcomes occur in the data ($0 < \sum_i y_i < n$), and that $\lambda_2 > 0$. Then problem* (7.3) *has exactly one solution.*

The proof is in Appendix A. The conditions matter. Without $\lambda_2 > 0$, the logistic likelihood has no maximum when some combination of components perfectly predicts repayment in the sample, a situation that is likely in the small samples of an early rollout. Our earlier draft claimed uniqueness "under standard regularity conditions" on the grounds that the objective is strictly concave; with the correlation penalty it is not, and even with the covariance penalty strict concavity alone does not guarantee that a maximum exists.

8 What the Framework Does Not Guarantee

Our earlier draft stated several further results as theorems: that scores converge to stable values, that borrowers have an incentive to behave truthfully, that EFIS raises welfare relative to any alternative, that it makes all market participants better off, and that the largest gains go to the currently excluded. None of these was proved, and some are not true in the generality stated. We discuss them here as hypotheses.

Convergence. Proposition 3.1 shows that the precision of the capacity estimate settles down. The estimate itself does not converge to a fixed number, because earning capacity keeps changing; at best it tracks capacity with a stable error. Under the earlier repayment formula (5.2), the score grew without limit for a borrower who kept repaying, so it could not converge at all. Under (5.3) it is bounded and increasing for such a borrower, and so it does converge. Whether the full score $U_i(t)$ settles down also depends on the social capital term, which feeds each borrower's score into her neighbors' scores; this feedback is stable if the endorsement weights are small enough, but we have not characterized the condition.

Incentives. Each component gives the borrower something to work toward: repaying, keeping real relationships, saving, and completing training. But that is not the same as incentive compatibility, which would require that no borrower can gain by misrepresenting herself. Section 5 shows one way the earlier formula could be gamed. Endorsement networks can also be gamed by groups of borrowers who vouch for each other, and borrowers may take up insurance or education modules mainly to raise their scores, which would weaken what those actions reveal. These are design problems to be studied, not properties already secured.

Welfare and who gains. Better information does not make everyone better off. When a lender cannot tell borrowers apart, safe borrowers subsidize risky ones through a common interest rate. Better scoring removes that subsidy: safe borrowers gain, but risky borrowers who were previously pooled with them can face higher rates or lose access, so the change is not a Pareto improvement in general. Total lending and average rates can move in either direction depending on the distribution of borrower types. We expect the largest gains to go to currently excluded borrowers who are in fact creditworthy, since they face the largest information gap, but this is a prediction to be tested, not a result.

9 Using the Score in Practice

Starting with no data. At launch, no borrower has an EFIS history and no lender has evidence that the score works. We propose a gradual rollout. In the first stage, the weights are estimated on existing loan records using whatever stand-ins for the four components are available. In the second stage, lenders use a blend of EFIS and their existing score,

$$U_i^{\text{blend}}(t) = h(t)\, U_i^{\text{EFIS}}(t) + \big(1 - h(t)\big)\, U_i^{\text{old}}(t), \qquad h(t) = \frac{\tau_{\text{EFIS}}(t)}{\tau_{\text{EFIS}}(t) + \tau_{\text{old}}}, \tag{9.1}$$

where τ denotes the precision (inverse error variance) of each score as a predictor of repayment. This choice of $h(t)$ gives the most accurate blend when the two scores have independent, unbiased errors.[10] As evidence accumulates, $\tau_{\text{EFIS}}(t)$ rises and the blend shifts toward EFIS. In the third stage, lenders use EFIS alone if it has proven more accurate.

Checking the data. Informal income records are easy to inflate. Income can be cross-checked against bank deposits, mobile money transactions, and estimates of business revenue from inventory and sales, and income reports that agree across sources can be trusted more. Relationships can be checked against communication records, the plausibility of the claimed relationship given where the two people live and work, and statements from community organizations, as in equation (4.2). Repayment records must be impossible to forge or transfer; any secure database maintained by the lender or a trusted third party can provide this.[11]

Privacy. EFIS asks for sensitive information, especially about relationships. Borrowers should hold their own records and decide which lenders see them. Statistics released from the data, such as the estimated patience parameters, can be protected with differential privacy, which limits how much any released figure can change when one person's data are added or removed (Dwork et al., 2006). Where scores are computed by a third party, standard cryptographic methods allow the computation to be carried out without that party seeing the underlying data. These protections

[10] The logic is straightforward: for two unbiased estimates with independent errors, the weighted average with the smallest error variance puts weight on each in proportion to its precision. If the errors are correlated, as is likely when both scores use some of the same information, the best weight also depends on the correlation.

[11] Our earlier draft proposed blockchain tokens and included contract code. A distributed ledger is one way to meet this requirement, but nothing in the economics depends on it, and we no longer treat it as part of the framework.

have costs in accuracy and computation that would need to be measured in practice.

Fairness checks. The fairness penalty in (7.2) limits the score's correlation with group membership but does not guarantee equal outcomes. Lenders should also compare approval rates across groups. A common rule of thumb in United States employment law, the four-fifths rule, flags a problem when $\mathbb{P}(\text{approved} \mid \text{group } g) / \mathbb{P}(\text{approved} \mid \text{group } g') < 0.8$ for any pair of groups.

Matching borrowers and lenders. A score also makes it easier to run a market in which several lenders compete for each borrower. If each lender states the value it places on lending to each scored borrower, the matching of borrowers to lenders that maximizes total value can be found by linear programming, and with payments set by the Vickrey-Clarke-Groves rule, stating true values is each lender's best strategy regardless of what others do (Vickrey, 1961; Clarke, 1971; Groves, 1973; Leonard, 1983).[12] A repaid-loan record can also act as a form of collateral: a borrower who can lose a valuable record by defaulting has a reason to repay even when the lender cannot seize assets, as in models of reputation in credit markets (Diamond, 1989).

10 Conclusion

Lenders turn away many creditworthy people because they cannot see what makes them creditworthy. EFIS proposes to fill that gap with four kinds of evidence that informal borrowers often have, each with a model that says what the evidence means. Working through those models carefully changed several of our earlier conclusions. The earning-capacity filter discounts noisy income rather than genuinely volatile income. The repayment score we first proposed rewarded borrowers for splitting loans and grew without limit; the revised score does neither. The weights are unique only with a convex fairness penalty and a penalty that keeps them bounded. And better scoring is not a Pareto improvement, because it ends the cross-subsidy from safe to risky borrowers.

The main open question is empirical. EFIS needs to be tested on real lending, ideally in a field experiment in which some applicants are scored with it and some are not, so that repayment and access to credit can be compared. Further questions include how scores should work for migrants whose economic lives span several countries, how often scores can be updated before borrowers start to target the score rather than the behavior it measures, and whether better credit at the household level adds up to measurable gains in local economies.

References

George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

[12] Our earlier draft claimed this property for a mechanism that did not specify how lenders would be paid or charged. Without a payment rule of this kind, truthful bidding is not generally optimal; in a first-price auction, for example, bidders shade their bids. The earlier sections on reserve prices and reputation bonds also contained derivations that did not follow from their stated objectives, and we have removed them from this version.

Abhijit Banerjee, Dean Karlan, and Jonathan Zinman. Six randomized evaluations of microcredit: Introduction and further steps. *American Economic Journal: Applied Economics*, 7(1):1–21, 2015.

Timothy Besley and Stephen Coate. Group lending, repayment incentives and social collateral. *Journal of Development Economics*, 46(1):1–18, 1995.

Daniel Björkegren and Darrell Grissen. Behavior revealed in mobile phone usage predicts credit repayment. *World Bank Economic Review*, 34(3):618–634, 2020.

Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

Asli Demirgüç-Kunt, Leora Klapper, Dorothe Singer, and Saniya Ansar. *The Global Findex Database 2021: Financial Inclusion, Digital Payments, and Resilience in the Age of COVID-19*. World Bank, Washington, DC, 2022.

Douglas W. Diamond. Reputation acquisition in debt markets. *Journal of Political Economy*, 97(4): 828–862, 1989.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284, 2006.

Andreas Fuster, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. Predictably unequal? The effects of machine learning on credit markets. *Journal of Finance*, 77(1):5–47, 2022.

Maitreesh Ghatak and Timothy W. Guinnane. The economics of lending with joint liability: Theory and practice. *Journal of Development Economics*, 60(1):195–228, 1999.

Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

James D. Hamilton. *Time Series Analysis*. Princeton University Press, Princeton, NJ, 1994.

Lars Peter Hansen and Kenneth J. Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91(2):249–265, 1983.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

Rudolph E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

Dean Karlan, Markus Mobius, Tanya Rosenblat, and Adam Szeidl. Trust and social collateral. *Quarterly Journal of Economics*, 124(3):1307–1361, 2009.

Herman B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.

Joseph E. Stiglitz and Andrew Weiss. Credit rationing in markets with imperfect information. *American Economic Review*, 71(3):393–410, 1981.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

A Proofs

Proof of Proposition 3.1

Drop the subscript i and write $r := \sigma_y^2$ and $Q := \sigma_\theta^2$, so that $q = Q/r$. We proceed in four steps.

Step 1: a recursion for the variance. Substituting the prediction step (3.3) into the update step (3.4), with $M_t := P_{t|t-1} = \phi^2 P_{t-1|t-1} + Q$, gives

$$P_{t|t} = (1 - K_t)\, M_t = \Big(1 - \frac{M_t}{M_t + r}\Big) M_t = \frac{r M_t}{M_t + r}.$$

So $P_{t|t} = f(P_{t-1|t-1})$ with

$$f(P) := \frac{r(\phi^2 P + Q)}{\phi^2 P + Q + r}, \qquad P \geq 0.$$

Step 2: properties of f. The function f is continuous on $[0, \infty)$. Its derivative is $f'(P) = r^2\phi^2/(\phi^2 P + Q + r)^2 \geq 0$, so f is nondecreasing. Since $Q > 0$, the numerator of f is positive, so $f(P) > 0$; and since $\phi^2 P + Q < \phi^2 P + Q + r$, we have $f(P) < r$. Hence f maps $[0, \infty)$ into $(0, r)$.

Step 3: the fixed point is unique. A fixed point $P = f(P)$ with $P \geq 0$ satisfies $P(\phi^2 P + Q + r) = r(\phi^2 P + Q)$, that is,

$$\phi^2 P^2 + (Q + r - r\phi^2)\, P - rQ = 0. \tag{A.1}$$

If $\phi = 0$, this is linear and has the single root $P = rQ/(Q + r) > 0$. If $\phi \neq 0$, it is a quadratic with leading coefficient $\phi^2 > 0$ and constant term $-rQ < 0$. The product of its two roots is $-rQ/\phi^2 < 0$, so one root is positive and the other negative. In both cases there is exactly one nonnegative fixed point, $P^* > 0$. Dividing (A.1) by r^2 and writing $k := P/r$ gives $\phi^2 k^2 + (q + 1 - \phi^2)k - q = 0$, whose positive root is

$$\frac{P^*}{r} = \frac{-(q + 1 - \phi^2) + \sqrt{(q + 1 - \phi^2)^2 + 4\phi^2 q}}{2\phi^2} \tag{A.2}$$

for $\phi \neq 0$, and $q/(q + 1)$ for $\phi = 0$. Thus P^*/r depends only on ϕ and q.

Step 4: convergence and the gain. Let $P_t := P_{t|t}$. Since $P_1 = f(P_0) \in (0, r)$, all iterates from $t = 1$ on lie in $[0, r]$. Because f is nondecreasing, if $P_2 \geq P_1$ then $P_3 = f(P_2) \geq f(P_1) = P_2$, and by induction the sequence is nondecreasing; similarly, if $P_2 \leq P_1$ it is nonincreasing. A monotone sequence in $[0, r]$ converges, say to $\bar{P}$. By continuity of f, $\bar{P} = \lim P_{t+1} = \lim f(P_t) = f(\bar{P})$, so $\bar{P}$ is a nonnegative fixed point, and by Step 3, $\bar{P} = P^*$. For the gain, $K_t = M_t/(M_t + r)$ and $P_t = rM_t/(M_t + r) = rK_t$, so $K_t = P_t/r \to P^*/r = K^*$.

Step 5: K^* increases with q. For $\phi = 0$, $K^* = q/(q+1)$ is strictly increasing. For $\phi \neq 0$, K^* is the positive root of $g(K, q) := \phi^2 K^2 + (q + 1 - \phi^2)K - q = 0$. We have $\partial g/\partial q = K - 1$ and $\partial g/\partial K = 2\phi^2 K + q + 1 - \phi^2$. At the positive root, the quadratic formula gives $2\phi^2 K^* + (q + 1 - \phi^2) = \sqrt{(q + 1 - \phi^2)^2 + 4\phi^2 q} > 0$. By the implicit function theorem,

$$\frac{dK^*}{dq} = -\frac{\partial g/\partial q}{\partial g/\partial K} = \frac{1 - K^*}{2\phi^2 K^* + q + 1 - \phi^2}.$$

The denominator is positive, and $K^* = P^*/r < 1$ by Step 2, so $dK^*/dq > 0$. □

Proof of Proposition 7.1

Let X be the $n \times 5$ matrix with rows $x_i^\top$. We show that J is strictly concave and that its upper level sets are bounded; together these imply a unique maximizer.

Step 1: $\mathcal{L}$ is strictly concave. For a single observation, $y \log \Lambda(z) + (1 - y) \log(1 - \Lambda(z))$ has second derivative $-\Lambda(z)(1 - \Lambda(z)) < 0$ in z. Hence the Hessian of $\mathcal{L}$ is $-X^\top D X$, where D is diagonal with entries $p_i(1 - p_i) > 0$. For any $v \neq 0$, full column rank gives $Xv \neq 0$, so $v^\top X^\top D X v = \sum_i p_i(1 - p_i)(x_i^\top v)^2 > 0$. Thus $\mathcal{L}$ is strictly concave.

Step 2: the penalties are concave after the sign change. $\Omega(w) = \sum_j (w^\top \kappa_j)^2$ is a sum of squares of linear functions, hence convex, and Φ is a convex quadratic. So $-\lambda_1 \Omega - \lambda_2 \Phi$ is concave in $\boldsymbol{\vartheta}$, and J, the sum of a strictly concave and a concave function, is strictly concave.

Step 3: $J(\boldsymbol{\vartheta}) \to -\infty$ as $\|\boldsymbol{\vartheta}\| \to \infty$. Note that $\mathcal{L} \leq 0$ and $\Omega \geq 0$, so $J \leq -\lambda_2 \Phi(w)$. Take any sequence $\boldsymbol{\vartheta}^{(m)} = (w_0^{(m)}, w^{(m)})$ with norm tending to infinity, and suppose for contradiction that $J(\boldsymbol{\vartheta}^{(m)})$ stays above some constant $-b$ along a subsequence. On that subsequence $\lambda_2 \Phi(w^{(m)}) \leq b$, and since $\lambda_2 > 0$, the vectors $w^{(m)}$ are bounded. Then $|w_0^{(m)}| \to \infty$. Passing to a further subsequence, either $w_0^{(m)} \to +\infty$ or $w_0^{(m)} \to -\infty$. In the first case, pick an observation with $y_i = 0$, which exists because $\sum_i y_i < n$; since $C_i^\top w^{(m)}$ is bounded, $p_i \to 1$ and the term $\log(1 - p_i) \to -\infty$. In the second case, pick an observation with $y_i = 1$; then $p_i \to 0$ and $\log p_i \to -\infty$. Since all other terms of $\mathcal{L}$ are nonpositive, $\mathcal{L} \to -\infty$ and hence $J \to -\infty$, a contradiction.

Step 4: conclusion. Fix any $\boldsymbol{\vartheta}^{(0)}$. By Step 3, the set $\{\boldsymbol{\vartheta} : J(\boldsymbol{\vartheta}) \geq J(\boldsymbol{\vartheta}^{(0)})\}$ is bounded; it is closed because J is continuous, and so it is compact. J therefore attains a maximum on it, which is a global maximum. If two distinct points both attained the maximum, strict concavity would make J strictly larger at their midpoint, a contradiction. So the maximizer is unique. □