

The Price of Asking: Social Frictions in Relational Lending and FinTech Adoption

Pedro Naso, Yuchen Sun, Mengye Wang, Yuchen Xu

1. Introduction

Access to credit is a central determinant of household welfare in low-income settings. Formal financial institutions have largely failed to serve ordinary households, leaving relational lending, i.e., borrowing from family members, neighbors, and peers, as the primary source of credit, sustained by social enforcement and reciprocity within kinship networks (Fafchamps and Lund 2003; Ferrara 2003; Chen, Ma, and Sinclair 2022). Yet borrowing from a social contact is not costless. It requires asking a friend, negotiating terms with someone one knows, and potentially damaging the relationship if repayment becomes difficult. These non-pecuniary costs are real, but standard models of household credit largely overlook them. While the existing literature has devoted considerable attention to supply-side frictions such as credit availability, collateral requirements, information asymmetries, and credit rationing, the demand-side frictions embedded in relational borrowing have received far less scrutiny.

Frictions, rather than preferences, can be the primary driver of apparently suboptimal financial decisions (Choukhmane and de Silva 2026). We apply this logic to the credit market: households embedded in relational networks may forgo borrowing, or choose not to approach a potential lender at all, not because they lack access or prefer to go without, but because the social cost of asking is too high. The challenge is that these frictions are difficult to separately identify from financial constraints or risk preferences in observational data.

We address this identification challenge using a lab-in-the-field experiment in rural Burundi.¹ We recruit real borrower-lender pairs drawn from existing social networks and introduce, alongside the relational credit option, an anonymous alternative with pre-set terms and no social obligation. By exogenously varying the cost of this anonymous option, we identify how borrowers' credit-channel choices respond to the price of anonymous borrowing and thereby recover the implied interest-rate premium associated with avoiding social frictions.

¹ Burundi offers a useful setting for this study because FinTech penetration remains minimal, so local credit markets closely resemble environments before digital lending became widespread. This has two advantages for our design. First, it allows us to study behavioral responses to the introduction of an arm's-length digital credit option in a setting with limited prior exposure to such products. Second, because participants are unlikely to have well-established habits of digital borrowing, the design reduces the scope for choices to be driven by prior experience with digital credit.

Our paper also speaks to the growing role of FinTech in promoting financial inclusion. The anonymous alternative in our design is implemented as a simulated FinTech application (FinApp) with fixed, non-negotiable terms and rule-based enforcement within the game. This setup captures the core features of many digital-credit products, e.g. speed, standardization, and the absence of interpersonal negotiation, while remaining transparent to participants.² The conventional view is that FinTech improves welfare primarily by expanding credit access to previously excluded borrowers. Our design highlights a complementary channel: FinTech may also improve outcomes for households already embedded in relational networks by offering a socially neutral exit from the frictions of peer borrowing. This distinction matters for how researchers and policymakers evaluate the welfare effects of financial innovation.

Our experiment seeks to answer the following research questions (RQ):

- **RQ1:** Are borrowers willing to choose the anonymous credit channel (FinApp) over peer lending even when FinApp charges a higher interest rate, in order to avoid the social frictions embedded in relational borrowing? If so, how large is the implied interest-rate premium?
- **RQ2:** Does the mental cost of asking (the emotional effort required to initiate and negotiate a loan request with a known person) suppress borrowing from peer lenders and drive substitution toward FinApp? Specifically in our design, are borrowers assigned to a more socially distant peer lender less likely to approach that lender and more likely to choose FinApp instead?
- **RQ3:** Does reputational concern (the fear that default will damage an ongoing relationship or generate future reciprocity obligations) suppress peer borrowing and drive substitution toward FinApp? Specifically in our design, are borrowers assigned to a project with default risk less likely to approach the peer lender and more likely to choose FinApp instead?
- **RQ4:** Does access to an anonymous credit alternative complement or crowd out relational lending?
- **RQ5:** Does access to an anonymous credit alternative improve household welfare by relaxing the social constraints on borrowing?

² The recent expansion of FinTech into low-income markets is widely seen as a solution to limited credit access, extending credit to previously excluded borrowers by replacing collateral requirements with algorithmic screening (Zins and Weill 2016; Dupas et al. 2022). The scale of this expansion is striking: the share of adults worldwide with a financial account rose from 51 percent in 2011 to 76 percent in 2021 (Demirgüç-Kunt et al. 2022).

This paper relates to three streams of literature.

First, we contribute to the literature on household borrowing decisions and the psychology of credit. A growing body of work documents that non-pecuniary factors shape credit demand in ways that standard models miss: psychological framing affects loan take-up (Bertrand et al. 2010), borrowers are more sensitive to loan maturity than to price (Karlán and Zinman 2008), and borrowing from friends activates distinct emotional responses, including shame, perceived obligation, and conflict over ownership of loaned funds, that are absent from anonymous transactions (Angulo, Goldstein, and Norton 2025). One strand of this literature shows that frictions, rather than preferences, can drive apparently suboptimal financial decisions: in the context of retirement savings, Choukhmane and de Silva (2026) show that low stock market participation reflects participation costs rather than a preference for safe assets, with most investors preferring stocks once frictions are removed. We draw on this logic in the credit market: borrowers who avoid relational credit and forgo investment may not lack the willingness to borrow, but face social frictions that make asking too costly. We build on this work by providing direct monetary estimates of the non-pecuniary costs of relational borrowing in a low-income setting, decomposing them into negotiation costs and reputational risks, and showing that these frictions are a quantitatively significant determinant of credit channel choice.

Second, we contribute to the literature on relational lending and informal credit markets. Prior work establishes the informational foundations of informal credit: kinship networks sustain lending through social enforcement and reciprocity (Ferrara 2003; Tomy and Wittenberg-Moerman 2024), and social ties substitute for formal contracts in allocating credit and sharing risk (Fafchamps and Lund 2003; Udry 1994). We complement this literature by documenting the other side of the ledger: the social frictions that relational lending imposes on borrowers. We further show how the introduction of a formal credit option reshapes both credit allocation and the social networks within which informal lending is embedded (Banerjee et al. 2024).

Third, we contribute to the literature on FinTech and financial inclusion in low-income settings. Existing work evaluates FinTech primarily through the lens of credit access, asking whether digital lending reaches previously unbanked households (Dupas et al. 2022; Suri, Bharadwaj, and Jack 2021; Mazer and Garz 2024; Allen, Gu, and Jagtiani 2021). We identify a distinct welfare channel that this literature has overlooked: FinTech improves outcomes not only by extending credit to the excluded, but by offering a socially costless alternative to borrowers who are already embedded in relational networks but find the social frictions of peer borrowing prohibitively costly. This has direct implications for how we measure the welfare effects of financial innovation in low-income settings.

2. Experimental Design

The experiment consists of a five-round loan game with real borrower-lender pairs and real-value incentives. The borrower plays the role of a farmer, reflecting the economic environment of a typical rural participant in Burundi. The game has four investment rounds followed by a final settlement round. In each of the first four rounds, the borrower faces a risky agricultural investment opportunity: they may invest in improved seeds in order to increase output in the following round. The return to this investment is realized in the next round and depends on weather conditions. In all five rounds, the borrower also receives a fixed cash flow. The borrower can finance the investment by borrowing either from the peer lender or from an anonymous simulated FinTech application (FinApp). Any loan repayment is made out of the borrower's available cash on hand, including both the fixed income and any realized return from the investment.

2.1. Core concepts

(1) Relational Credit

Relational credit refers to borrowing from a known person within one's existing social network, where loan terms are determined through interpersonal negotiation rather than through standardized contracts. In our setting, this is the dominant form of borrowing available to ordinary households: borrowers approach family members, friends, or neighbors, discuss the loan terms directly, and may later renegotiate if repayment becomes difficult. The same social embeddedness that makes such credit flexible also makes it socially costly. Borrowing from a known person may require revealing financial need, negotiating with someone one expects to meet again, and managing the relational consequences of delayed repayment or default.

In the experiment, to replicate the social dynamics of real-world peer lending as closely as possible, we recruit borrower-lender pairs from participants' actual borrowing networks. The peer lender is invited from the borrower's real borrowing network identified in the baseline survey, subject to the borrower's randomly assigned tie-strength category. During the game, borrowers and lenders communicate via WhatsApp to negotiate the loan amount, the interest rate, and any repayment adjustments, thereby reproducing the interpersonal and renegotiable nature of real-world peer lending. Borrowers and lenders are seated in separate rooms and are not informed of each other's payoff structure or token balances, preserving the informational asymmetry typical of informal credit relationships.

(2) FinApp

In our experiment, the anonymous credit option is operationalized as a simulated FinTech application (FinApp). Real-world FinTech lenders often use large-scale behavioral and digital data to assess borrower creditworthiness and extend small loans to retail borrowers. Relative to other formal credit sources (e.g. traditional banks), FinTech lenders rely less on costly manual screening or conventional collateral requirements, making them a useful benchmark for an anonymous and standardized credit channel in our setting.

In the experiment, FinApp is designed to capture the contractual features that distinguish digital-style credit from relational borrowing: loans are available on demand, the interest rate is fixed and non-negotiable, repayment rules are standardized, and no interpersonal interaction is required. Unlike relational credit, borrowing from FinApp does not require the borrower to approach a known person, justify the request, negotiate terms, or manage relationship strain if repayment becomes difficult. FinApp therefore provides a socially neutral outside option against which demand for relational credit can be evaluated.

We make two design assumptions about FinApp. First, to approximate the logic of digital lending, FinApp is assumed to observe the borrower's repayment-relevant cash flows. Because borrowers in the experiment receive a stable per-round land income that continues over time, debt is, in principle, always recoverable given sufficient time. From FinApp's perspective, the borrower is therefore effectively risk-free within the design. As a result, the FinApp interest rate should not be interpreted as a market-clearing price that compensates for borrower default risk. Instead, it is best understood as the profit margin charged by FinApp, which we vary exogenously across treatment arms as a clean price shifter for anonymous borrowing. This feature allows us to treat variation in the FinApp rate as an exogenous change in the price of anonymous credit and to recover the implied interest-rate premium associated with avoiding the social frictions of relational lending. Second, one design simplification concerns repayment enforcement. FinApp is designated as the priority lender: in the event of default, the borrower's future cash flows are automatically garnished until the debt is fully repaid. This assumption is consistent with the practice of real-world FinTech applications, which often rely on strong penalty rules or automatic deductions for delinquent borrowers, thereby creating strong incentives to prioritize repayment when cash flow is available.

In our setting, this feature is not meant to reproduce every institutional detail of real-world digital credit, but rather to capture its impersonal character. FinApp transactions are anonymous and

socially neutral: they impose no interpersonal obligation, require no face-to-face negotiation, and generate no direct reputational exposure within the community.

(3) Interest Rate Premium for Anonymous Credit (FinApp)

Our main object of interest is the interest rate premium borrowers are prepared to incur in order to use FinApp rather than borrow from a peer lender. Conceptually, this premium reflects the monetary value borrowers place on avoiding the social frictions embedded in relational credit, including the discomfort of asking, the burden of interpersonal negotiation, and the fear of socially costly repayment difficulties.

In the experiment, this premium is identified primarily through revealed borrowing decisions under exogenously assigned FinApp interest rates. Within the design, because borrowers receive a stable per-round land income and FinApp has priority repayment, the borrower is effectively risk-free from FinApp's perspective (if the game is extendable to further interactions). The FinApp rate should therefore not be interpreted as a market-clearing price for default risk. Instead, it is an exogenously varied price of anonymous, standardized credit. Variation in FinApp take-up across rate arms therefore provides a revealed-preference measure of how borrowers respond to the price of avoiding the social frictions of borrowing from a known person. As the FinApp rate rises, take-up should decline, and the pattern of decline allows us to recover the monetary premium associated with anonymous credit.

To cross-check, we include a post-survey module in which the same borrower chooses between borrowing from a friend and borrowing from FinApp under different interest-rate comparisons. These choices provide a complementary measure of the borrower's switching point between relational borrowing and FinApp.

(4) Mental Cost of Asking

The first friction in relational borrowing is the mental cost of asking. This refers to the psychological and interpersonal burden involved in initiating a loan request with a known person, even before any repayment difficulty arises. Asking a friend, relative, or neighbor for money may require revealing financial need, risking embarrassment, anticipating refusal, or entering an uncomfortable negotiation over the loan amount and interest rate. In this sense, the cost of asking includes not only the emotional hurdle of making the request, but also the effort and discomfort involved in opening and conducting an interpersonal negotiation.

In our experimental design, this friction is identified through the borrower's decision of whether to approach the peer lender at all when an anonymous alternative is available. In principle, a borrower with access to both FinApp and a peer lender should have an incentive to approach the peer

first and then borrow from whichever source offers better terms. A borrower who bypasses peer negotiation altogether and goes directly to FinApp therefore reveals that initiating interpersonal borrowing carries a non-trivial cost. We study the mental cost of asking through two related behavioral margins: whether the borrower initiates a WhatsApp conversation with the peer lender before making a borrowing decision, and whether the borrower ultimately substitutes toward FinApp rather than relational credit.

The main design lever for this friction is network tie strength. Borrowers are assigned a partner drawn either from their top two nominees in the baseline borrowing network or from their third and fourth nominees. Because these nominees are ranked by the borrower's own stated borrowing preference, this treatment shifts the social ease of approaching the lender while preserving the basic structure of the borrowing decision.

(5) Reputational Concern

The second friction in relational borrowing is reputational concern. This refers to the fear that borrowing from a known person may generate socially costly consequences if repayment becomes difficult. In informal credit relationships, repayment problems are not purely financial events: they may trigger awkward renegotiation, disappointment on the lender's side, requests for forgiveness or extensions, and concern on the borrower's side about damaging an ongoing relationship or creating future reciprocity obligations. Reputational concern is therefore distinct from the mental cost of asking. The latter arises at the moment of initiating the request; the former arises from anticipating the relational consequences of possible non-repayment.

In the experiment, we operationalize reputational concern by varying whether the assigned investment project places the borrower in a setting with or without default risk. In the project with default risk, a bad weather outcome may leave the borrower temporarily unable to repay the peer lender in full, making the relational consequences of borrowing salient and generating reputational concern. In the project without default risk, the return structure is calibrated so that repayment remains feasible even under bad weather, thereby removing reputational concern linked to possible non-repayment. This design allows us to examine whether borrowers become less willing to rely on peer lending, and more likely to substitute toward FinApp, when borrowing takes place under conditions with reputational concern rather than without it. As with the mental cost of asking, the main behavioral margins are negotiation initiation and credit channel choice.

2.2 Treatment arms

The experimental design varies three main dimensions of the borrowing environment: the FinApp interest rate, the strength of the borrower's tie to the peer lender, and whether the investment project places the borrower in a setting with or without reputational concern. The FinApp interest rate is the primary price treatment and provides the main revealed-preference source of identification for the interest-rate premium associated with avoiding the social frictions of relational borrowing, while the other two dimensions identify the mental cost of asking and reputational concern, respectively.

2.2.1. Main Arms

(1) FinApp Interest Rate (F3/5/8/10/15)

We vary the price of anonymous borrowing through the FinApp interest rate in order to identify how borrowers' credit-channel choices respond to the price of anonymous credit. Across the experiment, the FinApp rate takes five possible values: F3 (3%), F5 (5%), F8 (8%), F10 (10%), and F15 (15%). This variation provides the primary revealed-preference evidence on the interest-rate premium borrowers are prepared to incur in order to avoid the social frictions of relational borrowing.

(2) Borrowing Network Tie (TC vs. TL)

We vary the strength of the borrower's social relationship with the assigned peer lender to capture mental cost of asking. In the baseline survey, each participant is asked to nominate, in order of borrowing preference, up to four people they would think of approaching if they needed to borrow money. Because this ranking is based on the participant's own stated borrowing preference, it provides a natural measure of the social ease of approaching a potential lender and, by implication, of the mental cost of asking.

We define the first and second nominees as close ties (TC) and the third and fourth nominees as less-close ties (TL). Participants are randomly assigned to either the TC or TL condition. Within the assigned condition, they are informed of the two eligible nominees and are asked to invite one of the two to participate as their peer lender in the game. If they arrive with someone outside the assigned pair of eligible nominees, the pair is not allowed to participate.³ Our primary estimate is therefore an intention-to-treat effect based on the assigned tie-strength category, rather than on ex post realized closeness alone.

³ This procedure is designed to reduce selection bias arising from differential attrition. If we were to assign a single specific peer lender ex ante, participation could depend not only on whether that exact pair is jointly available, but also on whether the borrower is willing to invite that particular person to the game. In that case, treatment assignment could directly affect attendance, creating a selection problem. By allowing the borrower to invite either of the two nominees within the assigned tie-strength category, we preserve the experimental variation in closeness while reducing attrition driven by reluctance to invite a specific person or by logistical constraints.

(3) Investment Project Default Risk (PD vs. PN)

In the experiment, the borrower faces a repeated risky agricultural investment opportunity in the first four rounds. The return is realized in the following round and depends on a weather draw made by the borrower: good weather yields a gain, while bad weather results in a loss.

We operationalize reputational concern by varying whether the assigned investment project places the borrower in a setting with or without default risk. We denote the project with default risk as PD and the project without default risk as PN.

In PD, weather follows a 70:30 good-versus-bad distribution. The investment returns 1.5 times the amount invested under good weather and 0.5 times the amount invested under bad weather, implying an expected gross return of approximately 120%. Because a bad-weather realization can substantially reduce the borrower's available cash, the borrower may be temporarily unable to repay the peer lender in full. PD therefore creates default risk in the relational lending environment and, through that channel, generates reputational concern.

In PN, the weather distribution remains the same at 70:30, but the payoff structure is adjusted to 1.4 times under good weather and 0.7 times under bad weather, implying an expected gross return of approximately 119%, close to that of PD. The defining feature of PN is that, although the investment still performs poorly under bad weather, repayment remains feasible because borrowers in the experiment receive an unconditional per-round cash flow from land income. As a result, even in the bad-weather state, the borrower is not expected to experience repayment difficulty vis-a-vis the peer lender within the design. PN therefore removes default risk while preserving a similar expected return and the same intuitive perception of weather risk.

Compared with PD, PN has lower realized return risk, in the sense of a lower return variance (or standard deviation). We nevertheless choose this as the baseline comparison because PD and PN share the same perceived weather-risk distribution, with a 70:30 split between good and bad weather, and have very similar expected returns. This choice comes at the cost of allowing realized project risk to differ across the two projects. We make this tradeoff because, for rural Burundian farmers with limited numeracy, the perceived distribution of good versus bad weather is likely to be more salient than the exact variance of returns, and we do not expect participants to calculate project risk instantaneously from the payoff structure. Among the extension arms, we therefore include PN1, a no-default-risk project designed to more closely match the return volatility of PD as a robustness check.

2.2.2. Extension Arms

In addition to the main arms, we include several extension arms designed to test robustness and isolate additional mechanisms.

(1) Relational Credit Only

In this arm, borrowers do not have access to FinApp and may borrow only from the peer lender. This arm allows us to examine borrowing and investment behavior when relational credit is the only available external funding source, and serves as a benchmark for evaluating whether access to anonymous credit crowds out or complements peer borrowing.

(2) FinApp Only

In this arm, borrowers do not have access to the peer lender and may borrow only from FinApp. This arm allows us to examine borrowing and investment behavior when anonymous credit is the only available external funding source, and serves as a benchmark for assessing whether some borrowers who avoid peer lending would still choose to borrow and invest when the social frictions of relational credit are removed.

(3) Risk-matched Investment Project (PN1)

The baseline contrast between PD and PN keeps the probability of good versus bad weather fixed and varies only the payoff structure, thereby preserving participants' intuitive perception of weather risk. As discussed above, this comes at the cost of allowing realized return volatility to differ across the two projects. As a robustness check, we therefore include PN1, a no-default-risk project designed to better match the return variance of PD while still eliminating default risk.

In PN1, good weather occurs with probability 55% and bad weather with probability 45%. The investment returns 160% under good weather and 70% under bad weather, implying an expected gross return of approximately 119.5%, again close to PD. In terms of return volatility, the standard deviation of gross returns is approximately 0.458 in PD, 0.321 in PN, and 0.448 in PN1. As in PN, the bad-weather payoff in PN1 remains high enough, together with the borrower's unconditional land income, to make repayment feasible. PN1 therefore does not generate default risk despite adverse weather outcomes, and helps separate the effect of reputational concern from any residual differences in return dispersion across projects.

(4) No-renegotiation arm

In the baseline design, debt renegotiation between the borrower and the peer lender is permitted. If the borrower cannot repay in full when repayment is due, the two parties may communicate and renegotiate the terms of repayment, including partial forgiveness, delayed

repayment, or a revision of the interest obligation. We allow this in the baseline because it makes relational lending resemble real-life borrowing from friends and acquaintances more closely: unlike formal lenders, peers often have discretion to be flexible when repayment difficulties arise. By contrast, formal lenders are more rigid and non-negotiable.

At the same time, this realistic flexibility may introduce a countervailing force for our research question on reputational concern. On the one hand, borrowers may avoid peer lending when default risk is present because they fear that non-repayment could damage the relationship, making FinApp more attractive. On the other hand, some borrowers may prefer peer lending precisely because it offers repayment flexibility that FinApp does not. In real life, this reflects the idea that, although defaulting on a friend may be socially costly, the financial consequences of missing repayment to a formal institution may be even more severe. Which of these two forces dominates on average is ultimately an empirical question. More broadly, it is also of interest to explore heterogeneity in channel preferences, that is, which types of borrowers are more likely to prefer relational credit versus FinApp when these forces are in tension.

To isolate reputational concern from this competing preference for flexibility, we include a no-renegotiation arm in which debt renegotiation between borrower and peer lender is not permitted. If repayment is due and cash flow is available, the borrower must repay according to the original terms. In this arm, peer lending becomes more similar to FinApp along the repayment-enforcement margin, allowing us to test whether the reluctance to borrow from peers under default risk is driven by reputational concern rather than by a preference for the flexibility embedded in relational credit.

2.3. Factorial Design, Power Analysis, and Sample Size

The primary experiment employs a $5 \times 2 \times 2$ factorial design crossing three treatment dimensions: FinApp interest rate (five levels: 3%, 5%, 8%, 10%, and 15%), network tie strength (close tie versus less-close tie), and investment project default risk (PD versus PN), yielding 20 cells.

Power calculations use the panel adjustment for repeated measures (Donner and Klar, 2000), under the following assumptions: $T = 4$ decision rounds, a within-person correlation of $\rho = 0.50$, and a significance level of $\alpha = 0.05$. These assumptions imply an effective sample size multiplier of $m = T/[1 + (T - 1)\rho] = 1.60$. For each adjacent pairwise comparison of FinApp adoption rates between neighboring interest-rate arms, we pool observations across the four cells within each rate arm (2 tie-strength conditions \times 2 project conditions), yielding an effective sample size of $N \times 4 \times 1.6multiplier$ per arm, where N denotes the number of participants per cell. Our power analysis is calibrated to an adoption profile in which FinApp take-up declines sharply at lower rates

and flattens at higher rates. Specifically, we assume FinApp take-up of 0.75, 0.50, 0.30, 0.15, and 0.05 at FinApp rates of 3%, 5%, 8%, 10%, and 15%, respectively, corresponding to adjacent declines of 25, 20, 15, and 10 percentage points.

Under this profile, a target of 22 borrower–lender pairs per cell, or 440 pairs in total, powers all four adjacent pairwise comparisons at or above 0.80 under the baseline assumption of $\rho = 0.5$. Specifically, power is 0.993 for the 25 percentage point comparison (3% versus 5%), 0.933 for the 20 percentage point comparison (5% versus 8%), 0.858 for the 15 percentage point comparison (8% versus 10%), and 0.803 for the 10 percentage point comparison (10% versus 15%).

In addition to the primary $5 \times 2 \times 2$ design, we include four extension arms for robustness and mechanism checks: a relational-credit-only arm, a FinApp-only arm, a risk-matched no-default project (PN1), and a no-renegotiation arm. To keep the total sample feasible, the FinApp-only, PN1, and no-renegotiation extension arms are implemented at the two benchmark FinApp rates of 5% and 10%; the relational-credit-only arm has no FinApp rate by construction. This yields 20 extension cells in total: 4 cells for the relational-credit-only arm, 4 cells for the FinApp-only arm, 4 cells for the PN1 arm, and 8 cells for the no-renegotiation arm. We target 10 borrower–lender pairs per extension cell, adding 200 pairs to the sample. Combined with the 440 pairs in the primary factorial design, this implies a total target sample of 640 borrower–lender pairs. Because these extension arms are designed for robustness and secondary mechanism checks rather than for the main treatment effects, we do not present separate formal power calculations for them here.

From the list of potential participants obtained from the baseline survey, we randomly select participants who satisfy the following conditions: they must have nominated at least four potential lenders, be reachable, and be willing to participate; their nominated partners must also be willing to participate. Eligible pairs are then randomly assigned to treatment arms, with randomization stratified by gender and age to ensure balance across treatment dimensions. To account for attrition and no-shows, we over-invite by approximately 20% for each session.

Borrower-round observations missing due to session incompleteness, no-show, or participant drop-out will be excluded from the primary analysis. We will report attrition rates by treatment arm and compare baseline characteristics between completed and dropped pairs to assess selection.

2.4. Payoffs

The following describes the payoff structure for the baseline treatment arm, in which participants have access to both relational credit and FinApp at a 5% interest rate and are assigned to Project PD, the investment project with default risk. Both borrowers and lenders are explicitly

instructed to maximize their final token holdings at the end of round five, which determine their real compensation.

Borrowers. Borrowers take the role of a farmer who owns a small plot of land. They start each round with a fixed land income of 30,000 tokens.⁴ From rounds 1 to 4, they may invest in an agricultural asset by purchasing high-quality seeds at a minimum cost of 150,000 tokens; they may invest more if their token balance permits. Since the minimum investment exceeds the per-round land income, borrowing is typically required unless the borrower has accumulated sufficient tokens from prior rounds. Investment returns depend on the realization of weather. In Project PD, good weather with probability 0.7 yields 1.5 times the amount invested, while bad weather with probability 0.3 yields 0.5 times the amount invested. Borrowers may borrow up to 120,000 tokens⁵ per round from either FinApp or their relational partner, but not both simultaneously within the same round. At the end of five rounds, the land asset is redeemed at a fixed value of 300,000 tokens and added to the borrower's final balance.⁶

Lenders. Lenders take the role of a wealthy landowner and receive a fixed income of 120,000 tokens per round in rounds 1 to 4.⁷ Each round, lenders decide whether to lend to their game partner and on what terms, negotiated via WhatsApp. The maximum loan amount per round is 120,000 tokens; any tokens not lent at the end of each round are automatically added to the lender's balance account and will be available at the end of the game.

Incentive structure. Soap bars are used as the experimental reward,⁸ with token-to-soap conversion thresholds calibrated independently for each role. Both borrowers and lenders earn between 1 and 4 bars, with each bar worth approximately 1,000 FBU. The conversion scales are:

- **Borrowers:** $\leq 320,000$ tokens = 1 bar; 320,001–470,000 = 2 bars; 470,001–620,000 = 3 bars;

⁴ 30,000 tokens corresponds to approximately 15 USD, calibrated to reflect the average monthly income of a typical rural household in the study area.

⁵ The maximum borrowing amount of 120,000 tokens was calibrated to reflect the average loan amount observed among participants in the baseline survey, ensuring that the stakes of the experiment are ecologically valid and meaningful relative to participants' real-world borrowing experience.

⁶ The land asset is included for borrower to balance total token earnings between lenders and borrowers, although participants are not told this in advance, in order to avoid ex post perceptions of unfairness.

⁷ The payoff structures of borrowers and lenders are calibrated to yield approximately equal expected total cash flows: borrowers receive 30,000 tokens per round plus a 300,000 token land redemption ($30,000 \times 5 + 300,000 = 450,000$ in the no-investment baseline), while lenders receive 120,000 tokens per round over four rounds (480,000).

⁸ Soap bars are used as the experimental reward for two reasons. First, soap is a staple commodity in rural agricultural households, a daily necessity that participants consistently value, making it a more meaningful and stable incentive than a small cash amount. Second, and more practically, distributing cash to male participants in particular risks the reward being spent immediately on alcohol, leaving the household with no tangible benefit. By distributing soap, the research team ensures that the reward benefits the entire household, including the participant's spouse and children. In addition to this incentive payment, we reimburse all participants for their transportation costs to attend the experimental session.

- **Lenders:** $\leq 475,000$ tokens = 1 bar; 475,001–495,000 = 2 bars; 495,001–515,000 = 3 bars; $\geq 515,001$ = 4 bars.¹⁰

2.5. Surveys

We conduct two surveys: a pre-experiment survey and a post-experiment survey.

The pre-experiment survey provides the sampling frame for the game. We conduct the pre-survey using a random-walk sampling procedure. The target population consists of adults aged 18 to 60 who have decision-making authority over their own or their household's finances and who reside in or near Muramvya, a rural area of Burundi with no current FinTech activity. In total, we interviewed 1,270 participants from 36 collines, the village-level administrative units in the study area.

The baseline survey collects demographic information, including gender, age, literacy, marital status, number of children, and place of residence, as well as detailed financial behavior, including past borrowing and lending activity, loan amounts, purposes, repayment status, interest rates, and non-monetary transfers. The survey also records the actual amounts borrowed and lent by each participant in the past 12 months, which we use to calibrate the loan amounts in the experiment so that the stakes remain ecologically valid and meaningful relative to participants' real-world borrowing experience. For each past borrowing and lending relationship, participants report the relationship type, interaction frequency, and whether the relationship was affected by lending experience. The survey also elicits risk preferences through incentivized lottery choices, financial literacy through percentage and compound-interest questions with self-reported confidence levels, and trust and social norms through questions about reciprocity and willingness to help others. Importantly, each participant is asked to nominate up to four people they would turn to if they needed to borrow money immediately, ranked

⁹ We calibrate the borrower's token-to-soap conversion thresholds separately for each project variant using all $2^4 = 16$ possible weather sequences under a fully leveraged strategy, defined as borrowing from FinApp and investing the maximum amount in each round. For each sequence, we compute final token holdings by tracking land income, investment returns, and debt repayment over time, then combine these outcomes using the corresponding weather probabilities for each project. Thresholds are chosen to satisfy several goals: a borrower who never invests always receives 2 pieces of soap; most active-investment outcomes fall in the 2-to-3-soap range, with non-trivial chances of both 1 and 4 soaps; the 4-soap threshold remains attainable enough to preserve incentives; and payoff distributions are broadly comparable across projects. Because variation in the FinApp interest rate has only a limited effect on final payoffs relative to project-level payoff differences, we do not recalibrate thresholds by interest rate. The resulting thresholds for the other two projects are: for Project PN, $\leq 390,000$, 390,001–540,000, 540,001–610,000, and $\geq 610,001$; for Project PN1, $\leq 340,000$, 340,001–540,000, 540,001–740,000, and $\geq 740,001$. Under these thresholds, the fully leveraged payoff distribution across the three projects is roughly 17–25% for 1 soap, 28–36% for 2 soaps, 21–31% for 3 soaps, and about 24% for 4 soaps.

¹⁰ The conversion thresholds for lenders are calibrated as follows. A lender who never lends, or who lends at a low rate of around 3%, receives 2 bars of soap. A lender who lends at the FinApp benchmark rate of 5% with full repayment across all rounds receives 3 bars. Reaching 4 bars requires negotiating a rate above approximately 8% with full repayment. A single borrower default is sufficient to reduce the lender to 1 bar, as the lost principal far exceeds accumulated interest income. Because the lender's income and decision environment are identical across all treatment arms, the same thresholds apply to all cases.

in order of preference. For each nominee, participants report the relationship type, interaction frequency, and an open-ended reason for the ranking. Finally, participants report awareness and prior use of digital financial services.

The post-experiment survey is administered separately to borrowers and lenders immediately after the game. It begins with a set of comprehension and implementation checks to verify that participants understood the game rules, including asset values, borrowing terms, repayment consequences, and key procedural details.

For borrowers, the survey then records summary choices made during the game, including whether they invested, whether they borrowed from FinApp, whether they borrowed from the peer lender, and how weather outcomes affected later decisions. It asks in detail which features of FinApp were attractive or unattractive, which features of peer borrowing were attractive or unattractive, and whether the possibility of renegotiation was an important reason for borrowing from the peer lender. It also includes revealed barriers to asking, such as whether there were rounds in which the borrower considered asking the partner for a loan but decided not to do so, as well as a counterfactual question on whether the decision would have differed if the partner had been a stranger. To capture the key mechanisms of the study, the borrower post-survey includes a switching module in which participants choose between borrowing from a friend at a fixed low interest rate and borrowing from FinApp at progressively higher rates, yielding a direct measure of the borrower's switching threshold between relational and anonymous credit. In addition, the borrower survey contains batteries on the mental cost of asking, reputational concern, social exposure, and external validity, including questions on real-life experiences of avoiding requests for loans, difficulties with repayment, and willingness to use a real mobile loan service. It also records relationship changes after the game and broader perceptions about when services like FinApp may be most useful.

For lenders, the post-experiment survey records whether and why they lent, whether they ever forgave debt or accepted less favorable repayment terms, and how they felt about those decisions. It includes structured modules on lending motivations, lender-side social pressure, and supply-side distortions in credit provision, such as lending because it felt awkward to refuse, or refusing because of anticipated relationship tension if repayment failed. The lender survey also captures relationship changes after the game, real-life parallels in lending behavior, distance to formal financial institutions, harvest outcomes, risk preferences, and financial literacy. Together, the borrower and lender post-surveys allow us to link observed in-game behavior to the psychological and relational mechanisms

emphasized in the design, while also providing complementary measures of external validity and heterogeneity.

2.6. Experiment procedures and Data Recording

Experimental sessions proceed as follows. Borrowers and their invited partners arrive together at the session site and are first checked against the recruitment list to verify their identities. Only invited pairs are admitted to the session. Once admitted, borrowers and lenders are separated into different rooms. Enumerators then guide each participant individually through the experiment following a standardized protocol, including the practice game, comprehension checks, round-by-round decisions, and the recording of all borrowing, lending, repayment, and renegotiation choices. At the end of the game, participants complete the post-experiment survey individually. Participants are also reimbursed for transportation costs, with reimbursement varying by the distance between their residence and the session site.

All decisions and outcomes are recorded in real time by enumerators using KoboToolbox. Each round generates a structured data entry covering the weather draw, investment decision, borrowing channel, loan terms, repayment amounts, and renegotiation outcomes.

All peer-to-peer loan negotiations conducted via WhatsApp are collected for analysis. Messages are coded along two dimensions. First, we code the content of negotiations, including discussion of loan terms, offers of partial forgiveness, requests for extensions, and expressions of social pressure. Second, we code the tone of negotiations, classifying exchanges as adversarial or cooperative in order to construct the renegotiation tone measure. To protect participant privacy, no names or personally identifying information are recorded; all messages are linked only to anonymous participant ID codes consistent with the baseline survey.

2.7. Experiment schedule

The baseline survey was conducted in February 2026, and loan amounts used in the experiment were calibrated based on average loan amounts observed in the baseline survey. Enumerator training took place in early March 2026. A pilot phase was conducted from March 27 to April 04, 2026, to verify the experimental design and enumerator procedures. The formal experiment is scheduled to begin on April 13, 2026, and data collection is expected to conclude by July 10, 2026. Data come from pre-experiment and post-experiment surveys, the experimental sessions, and additional modules on risk aversion, confidence, the costs of relational lending (both monetary and psychological), and other relevant traits.

3. Variables

3.1. Primary outcomes (borrower-round level)

(1) Negotiation initiation. A binary indicator equal to one if the borrower initiates a WhatsApp loan request to the peer lender in a given round.

(2) FinApp take-up. A binary indicator equal to one if the borrower borrows from FinApp in a given round.

(3) Final token balance. The borrower's total token balance at the end of the five-round game. This serves as our main welfare outcome.

3.2. Secondary outcomes (measuring mechanisms)

We also collect a broad set of secondary outcomes to shed light on mechanisms and behavioral responses. These include additional game-based measures of borrowing, lending, investment, repayment, and renegotiation behavior; text-based measures derived from WhatsApp negotiations, including negotiation content and tone; borrower post-survey measures on the perceived attractiveness and drawbacks of FinApp and peer lending, the mental cost of asking, reputational concern, and relationship effects; and lender post-survey measures on lending motivations, social pressure, forgiveness, and relationship consequences. These outcomes are used primarily to interpret mechanisms, assess robustness, and study heterogeneity, rather than as core endpoints of the design.

3.3. Treatment variables (assigned at pair level)

(1) FinApp interest rate. A categorical treatment indicating the exogenously assigned FinApp rate. In the main design, the rate takes five values: 3%, 5%, 8%, 10%, and 15%.

(2) Network tie strength. A binary treatment indicating whether the borrower is assigned to the close-tie condition, in which the eligible peer lender is invited from the borrower's first two nominees in the baseline borrowing network, or to the less-close-tie condition, in which the eligible peer lender is invited from the borrower's third and fourth nominees.

(3) Investment project default risk. A binary treatment indicating whether the borrower is assigned to the project with default risk (PD) or the project without default risk (PN).

(4) Available borrowing option. A categorical treatment indicating which borrowing channels are available in a given arm, including relational credit only, FinApp only, or both channels. Additional extension arms modify renegotiation rules and, in separate specifications, introduce further borrowing options.

3.4. Controls and Heterogeneity Variables

We use a standard set of pre-specified controls and heterogeneity variables drawn from the baseline survey. These include demographic characteristics, prior borrowing and lending behavior, financial literacy, risk preferences, confidence, FinTech awareness, and distance to formal financial institutions. In supplementary dynamic specifications, we also consider lagged game outcomes such as the previous round's weather realization, token balance, repayment status, and negotiation outcome, together with round fixed effects. Because these lagged variables are themselves post-treatment, they are not part of the baseline causal specifications. They are used only for precision and heterogeneity analyses in clearly labeled auxiliary regressions.

4. Hypotheses and Empirical Strategy

The experiment generates panel data across four decision rounds for each participating pair. For the round-level analyses, we pool observations across rounds and use the borrower-round as the unit of observation. Standard errors are clustered at the borrower-lender pair level.

4.1. Hypothesis 1 (RQ1) : FinApp interest rate and channel choice

H1: As the FinApp interest rate increases, borrowers are more likely to initiate negotiations with peer lenders and less likely to borrow from FinApp.

To test this hypothesis, we exploit the exogenous variation in the FinApp interest rate across treatment cells. The main outcomes are, first, a binary indicator for whether the borrower initiates a WhatsApp loan request to the peer lender in a given round and, second, a binary indicator for whether the borrower borrows from FinApp in that round. The analysis is conducted at the borrower-round level, pooling the four investment rounds and including round fixed effects. Standard errors are clustered at the borrower-lender pair level.

Our baseline specification is:

$$Y_{it} = \sum_{k \in \{5,8,10,15\}} \beta_k \mathbf{1}(FinAppRate_i = k\%) + X_i' \gamma + \delta_t + \varepsilon_{it}$$

where Y_{it} is either negotiation initiation or FinApp take-up for borrower i in round t ; $\mathbf{1}(FinAppRate_i = k\%)$ is an indicator for the assigned FinApp rate $k \in \{5, 8, 10, 15\}$, with 3% as the omitted category; X_i is a vector of pre-specified baseline controls; and δ_t denotes round fixed effects. The coefficients of interest are the β_k , which trace out how channel choice responds to the price of anonymous credit.

We expect higher FinApp rates to reduce FinApp take-up and increase negotiation initiation with the peer lender. In addition to this flexible specification, we also estimate a parsimonious version that treats the FinApp rate as a continuous variable, in order to summarize the average slope of channel substitution with respect to the price of anonymous credit. Together, these specifications allow us to test whether borrowers substitute away from FinApp as its price rises and to quantify the implied interest-rate premium associated with avoiding the social frictions of relational borrowing. Evidence on the range over which borrowers switch between peer lending and FinApp is further supplemented by the post-experiment switching module.

4.2. Hypothesis 2 (RQ2) : Mental cost of asking

H2: Borrowers assigned to a less-close peer lender are less likely to approach the peer lender and more likely to borrow from FinApp.

To test this hypothesis, we exploit the exogenous variation in network tie strength. Borrowers are randomly assigned to either the close-tie condition or the less-close-tie condition. Within the assigned condition, they are informed of the two eligible nominees from their baseline borrowing network and asked to invite one of the two to participate as their peer lender. Our main estimate therefore captures the intention-to-treat effect of assignment to tie strength. The main outcomes are identical as in Hypothesis 1. The analysis is conducted at the borrower-round level, pooling the four investment rounds and including round fixed effects. Standard errors are clustered at the borrower-lender pair level.

Our baseline specification is:

$$Y_{it} = \beta \text{LessCloseTie}_i + X_i' \gamma + \delta_t + \varphi_r + \varepsilon_{it}$$

where Y_{it} is either negotiation initiation or FinApp take-up for borrower i in round t ; LessCloseTie_i is an indicator equal to one if the borrower is assigned to a less-close tie and zero if assigned to a close tie; X_i is a vector of pre-specified baseline controls; δ_t denotes round fixed effects; and φ_r denotes FinApp rate fixed effects. The coefficient of interest is β . We expect $\beta < 0$ when the outcome is negotiation initiation and $\beta > 0$ when the outcome is FinApp take-up. These patterns would indicate that greater social distance raises the mental cost of asking and shifts borrowers away from relational credit and toward anonymous credit. As complementary evidence, we examine post-experiment survey measures on discomfort in initiating requests, awkwardness in discussing loan terms, and reasons for not approaching the peer lender.

As an extension, we also estimate specifications that interact $LessCloseTie_i$ with FinApp rate dummies to examine whether borrowers assigned to less-close peer lenders differ in their sensitivity to the price of FinApp credit.

4.3. Hypothesis 3 (RQ3): Reputational Concern

H3: Borrowers assigned to a project with default risk are less likely to approach the peer lender and more likely to borrow from FinApp.

To test this hypothesis, we exploit exogenous variation in project risk. Borrowers are randomly assigned to either Project PD, which involves default risk, or a project variant without default risk (PN). The main outcomes are identical as in Hypothesis 1. The analysis is conducted at the borrower-round level, pooling the four investment rounds and including round fixed effects. Standard errors are clustered at the borrower-lender pair level.

Our baseline specification is:

$$Y_{it} = \alpha + \beta PD_i + X_i' \gamma + \delta_t + \varphi_r + \varepsilon_{it}$$

where Y_{it} is either negotiation initiation or FinApp take-up for borrower i in round t ; PD_i is an indicator equal to one if the borrower is assigned to the project with default risk and zero otherwise; X_i is a vector of pre-specified baseline controls; δ_t denotes round fixed effects; and φ_r denotes FinApp rate fixed effects. The coefficient of interest is β . We expect $\beta < 0$ when the outcome is negotiation initiation and $\beta > 0$ when the outcome is FinApp take-up. These patterns would indicate that the presence of default risk makes borrowers less willing to rely on relational credit and more likely to use anonymous credit instead. As complementary evidence, we examine post-experiment survey measures on discomfort in initiating requests, awkwardness in discussing loan terms, and reasons for not approaching the peer lender.

As an extension, we also estimate specifications that interact PD_i with $LessCloseTie_i$ and FinApp rate dummies to test whether the effect of default risk depends on tie strength and on the price of FinApp credit.

Robustness and Additional Validation.

A potential concern is that the baseline comparison between PD and PN changes not only default risk but also the realized volatility of project returns. To address this, we use the PN1 extension arm as a robustness check. PN1 is constructed to eliminate default risk while more closely matching the return volatility of PD. We expect the PD-versus-PN1 comparison to yield results similar to the baseline PD-versus-PN comparison, but remain less assertive about this prediction because PN1 also

lowers the probability of good weather (from 70% to 55%), which may alter participants' risk perceptions and independently affect willingness to invest and borrow.

A second potential confound is the flexibility of relational credit. In the baseline design, peer loans can be renegotiated after repayment difficulty, whereas FinApp is rigid and non-negotiable. This flexibility may work against the effect of reputational concern, because some borrowers may still prefer peer lending precisely because friends can extend, reduce, or forgive debt. To assess this possibility, we use the no-renegotiation extension arm, in which peer loans become rigid along the repayment dimension. If the main PD effect is genuinely driven by reputational concern, it should remain present, and potentially become sharper, when this competing preference for flexibility is removed.

4.4. Hypothesis 4 (RQ 4): FinApp Access, Complementarity, and Crowding Out

H4: Access to FinApp may both substitute for and complement relational credit. It may crowd out peer lending by offering an alternative funding source, but it may also expand effective credit access for borrowers who would otherwise avoid approaching peer lenders because of social frictions. Whether one effect dominates or both coexist is an empirical question.

To test this hypothesis, we compare borrowing behavior across three borrowing-option environments: a relational-credit-only arm, a FinApp-only arm, and a baseline arm in which both channels are available. The main outcomes are peer borrowing, FinApp borrowing, and total borrowing across the game. We also examine investment participation and investment amounts as secondary outcomes, since complementarity may operate through increased access to productive investment rather than through simple substitution across borrowing channels.

Our baseline specification is:

$$Y_{it} = \alpha + \beta_1 \text{BothChannels}_i + \beta_2 \text{FinAppOnly}_i + X_i' \gamma + \delta_t + \varphi_r + \varepsilon_{it}$$

where Y_{it} is one of the borrowing outcomes for borrower i in round t ; BothChannels_i is an indicator equal to one if both relational credit and FinApp are available; FinAppOnly_i is an indicator equal to one if only FinApp is available; and the omitted category is the relational-credit-only arm. X_i is a vector of pre-specified baseline controls; δ_t denotes round fixed effects; and φ_r denotes FinApp rate fixed effects. Standard errors are clustered at the borrower-lender pair level. The coefficients of interest are β_1 and β_2 .

If FinApp crowds out relational lending, we expect peer borrowing to be lower in the both-channels arm than in the relational-credit-only arm. If FinApp complements existing borrowing opportunities by relaxing socially costly constraints, we expect total borrowing and investment to be higher when FinApp is available, especially in the both-channels arm. Comparing the both-channels arm with the FinApp-only arm also helps assess whether relational credit and anonymous credit act as substitutes or jointly expand borrowing capacity.

4.5. Hypothesis 5 (RQ 5): FinApp Access and Welfare

H5: Access to FinApp improves borrower welfare by relaxing the social constraints embedded in relational borrowing. Relative to settings in which only peer lending is available, access to FinApp should improve borrowers' final token balances. These gains may be particularly important for borrowers who face greater social frictions in relational credit.

To test this hypothesis, we compare borrower welfare across borrowing-option environments, focusing on the relational-credit-only arm, the FinApp-only arm, and the baseline arm in which both channels are available. Our main outcome is the borrower's final token balance at the end of the game. Our baseline specification is:

$$FinalBalance_i = \alpha + \beta_1 BothChannels_i + \beta_2 FinAppOnly_i + X_i' \gamma + \varphi_r + \varepsilon_{ir}$$

where $FinalBalance_i$ is the borrower's final token balance for borrower i ; $BothChannels_i$ is an indicator equal to one if the borrower has access to both relational credit and FinApp; $FinAppOnly_i$ is an indicator equal to one if the borrower has access only to FinApp; and the omitted category is the relational-credit-only arm. X_i denotes a vector of pre-specified baseline controls; and φ_r denotes FinApp rate fixed effects. Standard errors are clustered at the borrower-lender pair level. The coefficients of interest are β_1 and β_2 . If FinApp relaxes socially costly borrowing constraints, we expect final token balances in the both-channels arm, and potentially also in the FinApp-only arm, to be higher than in the relational-credit-only benchmark.

We also examine heterogeneity in these welfare effects by network tie strength and project default risk. If FinApp primarily improves welfare by allowing borrowers to avoid social frictions rather than only by adding another source of funds, then the welfare gains from FinApp access should be larger for borrowers assigned to less-close ties or to projects with default risk.

5. Ethical Considerations

We obtained ethics clearance from the Research Committee at the Geneva Graduate Institute in Switzerland in December 2025. We also received approval from the University of Burundi and research permits for the areas in Muramvya where the experiment will take place. In all surveys and during the experiment, local enumerators obtain informed verbal consent from all participants in Kirundi, the participants' native language, before registration. Participation is fully voluntary, and participants may withdraw at any time without penalty. Participant identities are recorded only for matching and payment purposes. All analysis uses anonymized data identified only by participant codes. To preserve the integrity of the experiment, participants are asked not to disclose the content of the game to other potential participants. We do not use deception and make clear that the game is a simulation. The simulated FinTech application does not correspond to any real financial product accessible to participants.

6. Conclusion

This paper studies the social frictions embedded in relational lending and their implications for the demand for anonymous credit. Using a lab-in-the-field experiment in rural Burundi, we examine how borrowers substitute between peer lending and a simulated FinTech alternative when the price of anonymous credit changes. We further isolate two mechanisms—the mental cost of asking and reputational concern—by varying network tie strength and project default risk. More broadly, the value of FinTech may extend beyond expanding access to credit: it may also offer a socially cleaner alternative to households already embedded in informal financial networks.

References

- Allen, Franklin, Xian Gu, and Julapa Jagtiani. 2021. "A Survey of Fintech Research and Policy Discussion." *Review of Corporate Finance* 1 (3–4): 259–339.
- Angulo, Ariella N., Noah J. Goldstein, and Michael I. Norton. 2025. "Friendship Fallout and Bailout Backlash: The Psychology of Borrowing and Lending." *Journal of Consumer Psychology* 35 (1): 3–21.
- Banerjee, Abhijit, Emily Breza, Arun G. Chandrasekhar, Esther Duflo, Matthew O. Jackson, and Cynthia Kinnan. 2024. "Changes in Social Network Structure in Response to Exposure to Formal Credit Markets." *Review of Economic Studies* 91 (3): 1331–1372.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman. 2010. "What's Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment." *Quarterly Journal of Economics* 125 (1): 263–306.
- Chen, Zhiwu, Chicheng Ma, and Andrew J. Sinclair. 2022. "Banking on the Confucian Clan: Why China

- Developed Financial Markets So Late.” *Economic Journal* 132 (644): 1378–1413.
- Choukhmane, Taha, and Tim de Silva. 2026. “What Drives Investors’ Portfolio Choices? Separating Risk Preferences from Frictions.” *Journal of Finance* 81 (1): 5–48.
- Demirgüç-Kunt, Asli, Leora Klapper, Dorothe Singer, and Saniya Ansar. 2022. “*The Global Findex Database 2021: Financial Inclusion, Digital Payments, and Resilience in the Age of COVID-19.*” Washington, DC: World Bank.
- Donner, Allan, and Neil Klar. 2000. *Design and Analysis of Cluster Randomization Trials in Health Research*. London: Arnold.
- Dupas, Pascaline, Jonathan Robinson, and Valentina Brailovskaya. 2022. “The Impact of Digital Credit in Low-Income Countries.” Center for Effective Global Action (CEGA), March 31. <https://cega.berkeley.edu/article/the-impact-of-digital-credit-in-low-income-countries/>
- Fafchamps, Marcel, and Susan Lund. 2003. “Risk-Sharing Networks in Rural Philippines.” *Journal of Development Economics* 71 (2): 261–287.
- Ferrara, Eliana La. 2003. “Kin Groups and Reciprocity: A Model of Credit Transactions in Ghana.” *American Economic Review* 93 (5): 1730–1751.
- Karlan, Dean S., and Jonathan Zinman. 2008. “Credit Elasticities in Less-Developed Economies: Implications for Microfinance.” *American Economic Review* 98 (3): 1040–1068.
- Mazer, Rafe, and Seth Garz. 2024. “Fast Growth and Slow Policy: A Decade of Digital Credit in Kenya.” *Oxford Review of Economic Policy* 40 (1): 82–103.
- Suri, Tavneet, Prashant Bharadwaj, and William Jack. 2021. “Fintech and Household Resilience to Shocks: Evidence from Digital Loans in Kenya.” *Journal of Development Economics* 153: 102697.
- Tomy, Rimmy E., and Regina Wittenberg-Moerman. 2024. “Community Membership and Reciprocity in Lending: Evidence from Informal Markets.” *Journal of Accounting and Economics* 78 (1): 101697.
- Udry, Christopher. 1994. “Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria.” *Review of Economic Studies* 61 (3): 495–526.
- Zins, Alexandra, and Laurent Weill. 2016. “The Determinants of Financial Inclusion in Africa.” *Review of Development Finance* 6 (1): 46–57.